



HAL
open science

Synthesis in presence of dynamic links

Béatrice Bérard, Benedikt Bollig, Patricia Bouyer, Matthias Függer, Nathalie Sznajder

► **To cite this version:**

Béatrice Bérard, Benedikt Bollig, Patricia Bouyer, Matthias Függer, Nathalie Sznajder. Synthesis in presence of dynamic links. *Information and Computation*, 2022, 289 (Part B), pp.104856. 10.1016/j.ic.2021.104856 . hal-03518879

HAL Id: hal-03518879

<https://inria.hal.science/hal-03518879>

Submitted on 10 Jan 2022

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.

Synthesis in Presence of Dynamic Links

Béatrice Bérard^a, Benedikt Bollig^b, Patricia Bouyer^b, Matthias Függer^{b,c}, Nathalie Sznajder^a

^a*Sorbonne Université, CNRS, LIP6, 75005 Paris, France*

^b*Université Paris-Saclay, CNRS, ENS Paris-Saclay, Laboratoire Méthodes Formelles, 91190 Gif-sur-Yvette, France*

^c*Inria, France*

Abstract

The problem of distributed synthesis is to automatically generate a distributed algorithm, given a target communication network and a specification of the algorithm’s correct behavior.

Previous work has focused on static networks with an a priori fixed message size. This approach has two shortcomings: Recent work in distributed computing is shifting towards dynamically changing communication networks rather than static ones, and an important class of distributed algorithms are so-called full-information protocols, where nodes piggy-pack previously received messages onto current messages.

In this work, we consider the synthesis problem for a system of two nodes communicating in rounds over a dynamic link whose message size is not bounded. Given a network model, i.e., a set of link directions, in each round of the execution, the adversary choses an arbitrary link from the network model, restricted only by the specification, and delivers messages according to the current link’s directions. Motivated by communication buses with direct acknowledge mechanisms, we further assume that nodes are aware of which messages have been delivered.

We show that the synthesis problem is decidable for a network model if and only if the network model does not contain the empty link that dismisses both nodes’ messages. We then extend the characterization to sequences of communication links that may contain empty links. We show that the synthesis problem is decidable in this case if and only if the number of consecutive empty links in all possible sequences is uniformly bounded from above.

Keywords: synthesis, linear-time temporal logic, distributed computing, dynamic links, game theory

1. Introduction

Starting from Church’s work [1] on synthesizing circuits from arithmetic specifications in the 1960s, automatic synthesis of programs or circuits has been widely studied.

In the case of a reactive system, given a specification, the goal is to find an implementation for a system that repeatedly receives inputs from the environment and generates outputs such that the system’s behavior adheres to the specification. Early work [2, 3, 4] was synthesizing algorithms that require knowledge of the complete system state, inherently yielding single-process solutions.

Single-process synthesis is related to finding a strategy for a player representing the process that has to win against the adversarial environment, and has been studied in the context of games [5, 6, 7] and with automata techniques [4, 8].

For systems with more than one process, different models for how communication and computation is organized have been studied. Their two extremes are message-triggered asynchronous computation [9, 10] and round-wise synchronous computation.

An example for the latter is the work by Pnueli and Rosner [11], who considered synchronous distributed systems with an a priori fixed communication network. In their model, the network is given by a directed communication graph, whose nodes are the processes and with a link from process p to q if p can send messages to q (or write to and read from a shared variable). Messages are from a fixed, finite alphabet per link. A solution to the synthesis problem is a distributed algorithm that operates in rounds, repeatedly reading inputs, exchanging messages, and setting outputs. Already the case of two processes with separate inputs and outputs, and without a communication link to each other, was shown to be undecidable for linear temporal logic (LTL) specifications [12] on the inputs and outputs. As a positive result, the paper presents a solution for unidirectional process chains.

Still in the case of static architectures and bounded messages, Kupferman and Vardi [13, 14] extended decidability results to branching time specifications and proved sufficient conditions on communication networks for decidability, while Finkbeiner and Schewe [15] presented a characterization of networks where synthesis is decidable. Since specifications are allowed to talk about message contents, however, they are powerful enough to break existing communication links between processes, leading to undecidability like in the two-process system without communication [11]. Gastin *et al.* [16] proved a necessary and sufficient condition for decidability on a class of communication networks if specifications are only on the inputs and outputs of processes used by these to communicate with the environment and not on messages sent between processes. Like [16], our work only allows “input-output” specifications that talk about environmental inputs in_p and outputs out_p of processes p and do not constrain message contents exchanged between processes. In addition, we allow specifications to talk about communication links, e.g., constraining the network dynamics under which an algorithm has to work correctly as in $\text{GF}(\text{link} = \leftarrow) \implies \text{GF}((in_2 = 1) \implies (out_1 = 1))$. For such specifications, we obtain decidability in several cases where the framework of [15] does not provide decidability.

Like in the single-process scenario, synthesis in distributed systems can be modeled as a game, which, in this context, are partial information games played between a cooperating set of processes against the environment [17, 18, 19, 20]. With the exception of Berwanger *et al.* [20], all the above approaches assume static, reliable networks. In [20], Berwanger *et al.* study games in which information that players have about histories is hierarchically ordered, and this order may change dynamically during a play. The main difference to our work is that we consider a memory model where messages carry the complete *causal* history allowing for unbounded communication messages, while [20] is based on local observations so that, at every round, a bounded amount of information is transmitted between players. Further, while asynchronous solutions to the synthesis problem considered potentially unbounded messages [10, 9], previous synchronous solutions assume an a priori fixed message size. Also Madhusudan *et al.* [10] assume that processes that communicate infinitely often encounter each other within a bounded number of steps.

The above assumptions have two shortcomings:

Modeling unreliability. Distributed computing has a long history of studying algorithms that provide services in presence of unstable or unreliable components [21]. Indeed, classical process and link failures can be treated as particular dynamic network behavior [22]. Early work by Akkoyunlu *et al.* [23] considered the problem of two groups of gangsters coordinating a coup despite an unreliable channel between both parties; later on generalized to the Byzantine generals problem [24]. Protocols like the Alternating Bit Protocol [25] aim at tolerating message loss between a sender and receiver node, and Aho *et al.* [26] studies optimal transmission rates over unreliable links. Afek *et al.* [27] discuss protocols that implement reliable links on top of unreliable links. Further, for algorithms that have to operate in dynamic networks, see, e.g., [28, 29, 30], network changes are the normal case rather than the exception.

Synthesis with unstable or faulty components has been studied by Vélner and Rabinovich [31] for two player games in presence of information loss between the environment and the inputs of a process. The approach is restricted to a single process, however. Dimitrova and Finkbeiner [32] study synthesis of fault-tolerant distributed algorithms in synchronous, fully connected networks. Processes are partitioned into correct and faulty. It is assumed that at every round at least one process is correct and the output of a correct process must not depend on the local inputs of faulty processes. While unreliable links can be mapped to process failures, the above assumptions are a priori too restrictive to cover dynamic networks.

Modeling full-information protocols. An important class of distributed algorithms are full-information protocols, where nodes piggy-pack previously received messages onto current messages [21, 33]. By construction, such algorithms do not have bounded message size. This kind of causal memory has been considered in [9, 10, 34, 35] for synthesis and control of Zielonka automata over Mazurkiewicz traces with various objectives, ranging from local-state reachability to ω -branching behaviors. Zielonka automata usually model *asynchronous processes* (there is no global clock so that processes evolve at their own speed until they synchronize) and *symmetric communication* (whenever processes synchronize, they mutually exchange their complete history).

In this work we consider the synthesis problem for a system of two nodes communicating in synchronous rounds, where specifications are given as LTL formulas or, more generally, ω -regular languages. The nodes are connected via a dynamic link. As in Coulouma *et al.* [29] and Charron-Bost *et al.* [30], a network is specified by a set of communication graphs, the *network model*, also referred to as oblivious message adversary in literature [29]. A distributed algorithm operates in rounds as in the model by Pnueli and Rosner [11], with the difference that the

communication graph is chosen by an adversary per round. Motivated by communication buses, like the industry standard I²C bus [36] and CAN bus [37], with direct acknowledge mechanisms after message transfers, we assume that nodes are aware if messages have been delivered successfully. In contrast to the Pnueli-Rosner setting, we suppose full-information protocols where processes have access to their causal history. That is, the dynamic links have unbounded message size. Unlike in Zielonka automata over traces, however, we consider *synchronous processes* and potentially *asymmetric communication*. In particular, the latter implies that a process may learn all about the other's history without revealing its own. Observe that, when restricting to Zielonka automata, synthesis of asynchronous distributed systems is *not* a generalization of the synchronous case.

We show that the synthesis problem for two processes is decidable for a network model if and only if it does not contain the empty link that dismisses both processes' messages. As we assume that LTL specifications can not only reason about inputs and outputs, but also about the communication graph, our result covers synthesis for dynamic systems where links change in more restricted ways. In particular, this includes processes that do not send further messages after their message has been lost, bounded interval omission faults, etc.

We then extend this characterization to sets of network sequences. We show that in case the adversarial environment can choose from an ω -regular language of network sequences, the synthesis problem is decidable if and only if the number of consecutive empty links is uniformly bounded from above.

The work is an extended version of [38].

Outline. We define the synthesis problem for the dynamic two-process model in Section 2. In Section 3, we discuss the asymmetric model where communication to process 1 never fails. Central to the analysis is to show that, despite the availability of unbounded communication links, finite-memory distributed algorithms actually suffice. We then prove that the synthesis problem is decidable (Theorem 2). In Section 5 we reduce the general case of dynamic communication to the asymmetric case, obtaining our main result of decidability in network models that do not contain the empty link (Theorem 1). In Section 6, we extend our results to the case of network sequences where the number of successive empty links is uniformly bounded. We conclude in Section 7.

2. The Synthesis Problem

We start with a few preliminaries. Let $\mathbb{N} = \{0, 1, 2, \dots\}$. For a (possibly infinite) alphabet A , the set of finite words over A is denoted by A^* , the set of nonempty finite words by A^+ , and the set of countably infinite words by A^ω . We let ε be the empty word and denote the concatenation of $w_1 \in A^*$ and $w_2 \in A^* \cup A^\omega$ by $w_1 \cdot w_2$ or simply $w_1 w_2$.

Fix the set of processes $P = \{1, 2\}$. Every process $p \in P$ comes with fixed finite sets X_p and Y_p of possible *inputs* and *outputs*, respectively. We assume that there are at least two possible inputs and outputs per process, i.e., $|X_p| \geq 2$ and $|Y_p| \geq 2$.

We consider systems where computation and communication proceed in rounds. In round $r = 0, 1, 2, \dots$, process $p \in P$ receives an input $x_p^r \in X_p$ and it produces an output $y_p^r \in Y_p$. The decision on y_p^r depends on the knowledge that process p has about the execution up to round r . In addition to all local inputs x_p^0, \dots, x_p^r , this knowledge can also include inputs of the other process, which may be communicated through communication links.

Following Charron-Bost *et al.* [30], we consider a dynamic communication topology in terms of a *network model*, i.e., a fixed nonempty set $\mathcal{N} \subseteq \{\times, \leftarrow, \rightarrow, \leftrightarrow\}$ of potentially occurring communication graphs. In round $r \geq 0$, a graph $\rightleftharpoons^r \in \mathcal{N}$ is chosen non-deterministically with the following intuitive meaning:

- \times No communication takes place. The knowledge of process p that determines y_p^r only includes the knowledge at round $r - 1$ as well as the new input x_p^r .
- \leftarrow Process 1 becomes aware of the whole input sequence $x_2^0 \dots x_2^r$ that process 2 has received so far. This includes x_2^r , which is transmitted without delay. The case \rightarrow is analogous.
- \leftrightarrow Both processes become aware of the whole input sequence of the other process.

As discussed in the introduction, the knowledge of process p at round r also includes the communication link \rightleftharpoons^r at round r , which is therefore common knowledge.

2.1. Histories and Views

Let us be more formal. Recall that we fixed the sets P , X_p , Y_p , and \mathcal{N} . We let $\Sigma = X_1 \times \mathcal{N} \times X_2$ be the set of *input signals*. For ease of notation, we write $\langle x_1 \rightrightarrows x_2 \rangle$ instead of $(x_1, \rightrightarrows, x_2) \in \Sigma$. Moreover, for $\rightrightarrows \in \mathcal{N}$, we let $\Sigma_{\rightrightarrows} = X_1 \times \{\rightrightarrows\} \times X_2$. A word $w \in \Sigma^*$ represents a possible *history*, a sequence of signals to which the system has been exposed so far. For a process p , we inductively define the *view* $\llbracket w \rrbracket_p$ of p on w by replacing inputs that are invisible to p by the symbol \perp (we suppose $\perp \notin X_1 \cup X_2$). First of all, let $\llbracket \varepsilon \rrbracket_1 = \llbracket \varepsilon \rrbracket_2 = \varepsilon$. Moreover, for $u \in \Sigma^*$:

$$\begin{aligned} \llbracket u \langle x_1 \leftrightarrow x_2 \rangle \rrbracket_1 &= u \langle x_1 \leftrightarrow x_2 \rangle & \llbracket u \langle x_1 \leftrightarrow x_2 \rangle \rrbracket_2 &= u \langle x_1 \leftrightarrow x_2 \rangle \\ \llbracket u \langle x_1 \leftarrow x_2 \rangle \rrbracket_1 &= u \langle x_1 \leftarrow x_2 \rangle & \llbracket u \langle x_1 \rightarrow x_2 \rangle \rrbracket_2 &= u \langle x_1 \rightarrow x_2 \rangle \\ \llbracket u \langle x_1 \rightarrow x_2 \rangle \rrbracket_1 &= \llbracket u \rrbracket_1 \langle x_1 \rightarrow \perp \rangle & \llbracket u \langle x_1 \leftarrow x_2 \rangle \rrbracket_2 &= \llbracket u \rrbracket_2 \langle \perp \leftarrow x_2 \rangle \\ \llbracket u \langle x_1 \times x_2 \rangle \rrbracket_1 &= \llbracket u \rrbracket_1 \langle x_1 \times \perp \rangle & \llbracket u \langle x_1 \times x_2 \rangle \rrbracket_2 &= \llbracket u \rrbracket_2 \langle \perp \times x_2 \rangle \end{aligned}$$

With this, we let $Views_1 = \{\llbracket w \rrbracket_1 \mid w \in \Sigma^+\}$ and $Views_2 = \{\llbracket w \rrbracket_2 \mid w \in \Sigma^+\}$ be the sets of possible *views* of processes 1 and 2.

Example 1. The view $\llbracket w \rrbracket_1$ is illustrated in Figure 1 for four different words w . For the history in Figure 1(c), for example, we have $\llbracket \langle x_1^0 \leftarrow x_2^0 \rangle \langle x_1^1 \leftrightarrow x_2^1 \rangle \langle x_1^2 \rightarrow x_2^2 \rangle \langle x_1^3 \rightarrow x_2^3 \rangle \rrbracket_1 = \langle x_1^0 \leftarrow x_2^0 \rangle \langle x_1^1 \leftrightarrow x_2^1 \rangle \langle x_1^2 \rightarrow \perp \rangle \langle x_1^3 \rightarrow \perp \rangle$. Note that Figures 1(c) and (d) depict different continuations of the word from Figure 1(b). In case (c), x_2^2 remains unknown to process 1, whereas in (d), \leftarrow will reveal both x_2^2 and x_2^3 to process 1.




Figure 1: View $\llbracket w \rrbracket_1$ of process 1 for some histories w ; the white part is unknown in the view, and replaced by \perp .

2.2. Linear-Time Temporal Logic

Let $\Omega = Y_1 \times Y_2$ be the set of *output signals*. An *execution* is a word from $(\Sigma \times \Omega)^\omega$, which records, apart from the input signals, the outputs at every round. A convenient specification language to define the *valid* system executions is *linear-time temporal logic* (LTL) interpreted over words from $(\Sigma \times \Omega)^\omega$. The logic can, therefore, talk about inputs, outputs, and communication links at a given position. Moreover, it has the usual temporal modalities. Mind, however, that our specifications cannot speak about *received* values, thus differing from Finkbeiner and Schewe [15] where the specification may force a constant reception of 0 from process 1 to process 2, i.e., disable communication between processes. Formally, the set $LTL(\mathcal{N})$ of LTL formulas is given by the grammar

$$\begin{aligned} \varphi ::= & (in_p = x) \mid (out_p = y) \mid (link = \rightrightarrows) \mid && \text{atomic formulas} \\ & X\varphi \mid F\varphi \mid G\varphi \mid \varphi U \varphi \mid && \text{temporal modalities} \\ & \neg\varphi \mid \varphi \vee \psi \mid \varphi \wedge \psi \mid \varphi \implies \psi \mid \varphi \iff \psi && \text{Boolean connectives} \end{aligned}$$

where $p \in P$, $x \in X_p$, $y \in Y_p$, and $\rightrightarrows \in \mathcal{N}$. Let $e = \alpha_0 \alpha_1 \alpha_2 \dots$ be an execution with $\alpha_i \in \Sigma \times \Omega$ for all $i \in \mathbb{N}$ and $\alpha_0 = (\langle x_1^0 \rightrightarrows x_2^0 \rangle, (y_1^0, y_2^0))$. For $r \in \mathbb{N}$, let $e_{\geq r}$ denote its suffix $\alpha_r \alpha_{r+1} \alpha_{r+2} \dots$, i.e., $e = e_{\geq 0}$. Boolean connectives are interpreted as usual. Moreover:

$$\begin{aligned} e \models (in_p = x) & \quad \text{if } x_p^0 = x & e \models X\varphi & \quad \text{if } e_{\geq 1} \models \varphi \\ e \models (out_p = y) & \quad \text{if } y_p^0 = y & e \models F\varphi & \quad \text{if } \exists r \geq 0 : e_{\geq r} \models \varphi \\ e \models (link = \rightrightarrows) & \quad \text{if } \rightrightarrows^0 = \rightrightarrows & e \models G\varphi & \quad \text{if } \forall r \geq 0 : e_{\geq r} \models \varphi \\ e \models \varphi U \psi & \quad \text{if } \exists r \geq 0 : (e_{\geq r} \models \psi \wedge \forall 0 \leq r' < r : e_{\geq r'} \models \varphi) \end{aligned}$$

Finally, we let $L(\varphi) = \{e \in (\Sigma \times \Omega)^\omega \mid e \models \varphi\}$ be the set of executions that satisfy φ .

Remark 1. In general, the sequence of communication graphs in an execution is arbitrary from \mathcal{N}^ω , modeling a highly dynamic network without any restrictions on stability, eventual convergence, etc. Note that the specification is allowed to speak about the communication links along a history, however, with the possibility to restrict the behavior of the dynamic network and impose process behavior to depend on the network dynamics.

Example 2. Suppose $X_1 = X_2 = Y_1 = Y_2 = \{0, 1\}$ and $\mathcal{N} = \{\leftarrow, \rightarrow\}$. Consider

$$\begin{aligned}\varphi_1 &= \mathbf{G}((out_1 = 1) \iff (out_2 = 1)) \\ \varphi_2 &= \mathbf{GF}((in_1 = 1) \wedge (in_2 = 1)) \iff \mathbf{GF}((out_1 = 1) \wedge (out_2 = 1)) \\ \psi &= (\mathbf{GF}(link = \leftarrow) \wedge \mathbf{GF}(link = \rightarrow)) \implies \varphi_1 \wedge \varphi_2.\end{aligned}$$

Formula φ_1 says that, in each round, both processes agree on their output. Formula φ_2 postulates that both processes simultaneously output 1 infinitely often if, and only if, both inputs are simultaneously 1 infinitely often. Finally, ψ requires φ_1 and φ_2 to hold if both communication links occur infinitely often. We will come back to these formulas later to illustrate the synthesis problem. \triangleleft

2.3. Synthesis Problem

A *distributed algorithm* is a pair $f = (f_1, f_2)$ of functions $f_1 : Views_1 \rightarrow Y_1$ and $f_2 : Views_2 \rightarrow Y_2$ that associate with each view an output. Given $w = \sigma_0\sigma_1\sigma_2 \dots \in \Sigma^\omega$, we define the execution $f(\llbracket w \rrbracket) = (\sigma_0, (y_1^0, y_2^0))(\sigma_1, (y_1^1, y_2^1)) \dots \in (\Sigma \times \Omega)^\omega$ where $y_p^r = f_p(\llbracket \sigma_0 \dots \sigma_r \rrbracket_p)$. For a finite word $w \in \Sigma^*$, we define $f(\llbracket w \rrbracket) \in (\Sigma \times \Omega)^*$ similarly (in particular, $f(\llbracket \varepsilon \rrbracket) = \varepsilon$).

Let $L \subseteq (\Sigma \times \Omega)^\omega$ and $\varphi \in \text{LTL}(\mathcal{N})$. We say that f *fulfills* L (respectively φ) if, for all $w \in \Sigma^\omega$, we have $f(\llbracket w \rrbracket) \in L$ (respectively $f(\llbracket w \rrbracket) \in L(\varphi)$). Moreover, we say that L (respectively φ) is *realizable* if there is some distributed algorithm that fulfills L (respectively φ).

We are now ready to define our main decision problem:

Definition 1. For a fixed network model \mathcal{N} (recall that we also fixed P, X_p, Y_p), the *synthesis problem* is defined as:

SYNTHESIS(\mathcal{N}) **Input:** $\varphi \in \text{LTL}(\mathcal{N})$
 Question: Is φ realizable?

Example 3. Consider the formulas $\varphi_1, \varphi_2, \psi$ from Example 2 over $\mathcal{N} = \{\leftarrow, \rightarrow\}$. We easily see that φ_1 is realizable by the distributed algorithm where both processes always output 1. However, $\varphi_1 \wedge \varphi_2$ is *not* realizable: if the communication link is always \leftarrow (an analogous argument holds for \rightarrow), process 2 has no information about any of the inputs of process 1. Thus, it is impossible for the processes to agree on their outputs in every round while respecting φ_2 .

Finally, formula ψ is realizable. We can now assume that both \leftarrow and \rightarrow occur infinitely often. A sequence of signals can be divided into maximal finite blocks with identical communication links as illustrated in Figure 2 for the prefix of an execution. The distributed algorithm proceeds as follows. By default, both processes output 0, with the following exception: at the first position of each block, a process outputs 1 if, and only if, the preceding block contains a round where both processes simultaneously received 1. Note that this preceding block is entirely contained in the view of both processes. The algorithm's outputs are illustrated in Figure 2. At rounds 4 and 6, they are 1 because the corresponding preceding blocks contain an input pair of 1's. As every block has finite size, satisfaction of φ_2 is guaranteed. \triangleleft

It is well known that the synthesis problem is undecidable if processes are not connected:

| round | | signal | | | |
|-------|---|--------|---------------|---|---|
| 0 | 0 | 0 | \leftarrow | 1 | 0 |
| 1 | 0 | 1 | \rightarrow | 0 | 0 |
| 2 | 0 | 1 | \rightarrow | 1 | 0 |
| 3 | 0 | 0 | \rightarrow | 0 | 0 |
| 4 | 1 | 1 | \leftarrow | 0 | 1 |
| 5 | 0 | 1 | \leftarrow | 1 | 0 |
| 6 | 1 | 0 | \rightarrow | 0 | 1 |
| 7 | 0 | 0 | \leftarrow | 1 | 0 |
| 8 | 0 | 1 | \leftarrow | 0 | 0 |
| 9 | 0 | 1 | \rightarrow | 1 | 0 |

Figure 2: Fulfilling ψ

Fact 1 (Pnueli-Rosner). *The problem $\text{SYNTHESIS}(\{\times\})$ is undecidable.*

One also observes that undecidability of the synthesis problem is upward-closed:

Lemma 1. *Let $\mathcal{N}_1 \subseteq \mathcal{N}_2$. If $\text{SYNTHESIS}(\mathcal{N}_1)$ is undecidable, then so is $\text{SYNTHESIS}(\mathcal{N}_2)$.*

Indeed, formula $\varphi_1 \in \text{LTL}(\mathcal{N}_1)$ is realizable iff the formula $\varphi_2 \in \text{LTL}(\mathcal{N}_2)$ is realizable where we let $\varphi_2 = (\mathbf{G} \bigvee_{\text{link} \in \mathcal{N}_1} (\text{link} = \text{link})) \implies \varphi_1$.

Therefore, we will now focus on network models that do not contain \times . Our main result is the following:

Theorem 1. *For a network model \mathcal{N} , $\text{SYNTHESIS}(\mathcal{N})$ is decidable if and only if $\times \notin \mathcal{N}$.*

The “only if” direction follows from Fact 1 and Lemma 1.

Let us give a short high-level proof plan for the “if” direction. We will first consider $\mathcal{N} = \{\leftrightarrow, \leftarrow\}$:

Theorem 2. *The problem $\text{SYNTHESIS}(\{\leftrightarrow, \leftarrow\})$ is decidable (in 3-fold exponential time).*

The proof of Theorem 2 is spread over Sections 3 and 4. As our setting features a dynamic architecture and unbounded message size in terms of causal histories, the proof of the theorem requires some new techniques. In particular, we cannot apply the information-fork criterion from [15], since our specifications can only describe the link between the processes, and cannot constrain the contents of the messages.

Theorem 2 crucially relies on the fact that, for every realizable specification φ , there is a distributed algorithm with a sort of *finite memory* fulfilling it (Lemma 2 in Section 3). This allows us to reduce the problem of finding a distributed algorithm to finding a winning strategy in a decidable game, which we will call a (2, 1)-player game thereafter (Lemma 5 in Section 4). The game involves two cooperating players, where one player has imperfect information, and an antagonistic environment.

Once we have Theorem 2, we reduce the remaining cases needed for Theorem 1 to the particular network model $\{\leftrightarrow, \leftarrow\}$. This is done in Section 5. By Lemma 1, it is enough to do the reduction for $\{\leftrightarrow, \leftarrow, \rightarrow\}$.

3. Finite-Memory Distributed Algorithms for $\mathcal{N} = \{\leftrightarrow, \leftarrow\}$

In this section, we suppose $\mathcal{N} = \{\leftrightarrow, \leftarrow\}$. We will provide the first ingredient of the proof of Theorem 2: Every realizable specification φ comes with a distributed algorithm with finite memory fulfilling it (Lemma 2).

Remark 2. For the sake of technical simplification, we assume in Sections 3 and 4, without loss of generality, that input sequences start with a symbol from $\Sigma_{\leftrightarrow} = X_1 \times \{\leftrightarrow\} \times X_2$. Instead of the original formula $\hat{\varphi}$, we then simply take $\varphi = X\hat{\varphi}$. That is, we can henceforth consider that $\text{Views}_1 = \{\llbracket w \rrbracket_1 \mid w \in \Sigma_{\leftrightarrow}\Sigma^*\}$ and $\text{Views}_2 = \{\llbracket w \rrbracket_2 \mid w \in \Sigma_{\leftrightarrow}\Sigma^*\}$, and that a distributed algorithm f fulfills $\varphi \in \text{LTL}(\mathcal{N})$ if, for all $w \in \Sigma_{\leftrightarrow}\Sigma^\omega$, we have $f(\llbracket w \rrbracket) \in L(\varphi)$.

3.1. Finite-Memory Distributed Algorithms

Deterministic Parity Word Automata. The definition of a finite-memory distributed algorithm relies on deterministic parity word automata; see [39]:

Definition 2. A *deterministic parity word automaton (DPWA)* over a finite alphabet A is a tuple $\mathcal{A} = (S, \iota, \delta, \mathcal{F})$, where S is a finite set of states, $\iota \in S$ is the *initial state*, $\delta : S \times A \rightarrow S$ is the transition function, and $\mathcal{F} : S \rightarrow \mathbb{N}$ is the acceptance condition, mapping each state to a *priority*.

The DPWA \mathcal{A} defines a language of infinite words $L(\mathcal{A}) \subseteq A^\omega$ as follows. We extend δ to a function $\delta : S \times A^* \rightarrow S$ letting $\delta(s, \varepsilon) = s$ and $\delta(s, aw) = \delta(\delta(s, a), w)$. Let $w = a_0a_1a_2\dots \in A^\omega$. We define $\text{Visit}_{\mathcal{A}}^\infty(w) = \{s \in S \mid s = \delta(\iota, a_0\dots a_i) \text{ for infinitely many } i \in \mathbb{N}\}$. We say that w is *accepted* by \mathcal{A} if $\min\{\mathcal{F}(s) \mid s \in \text{Visit}_{\mathcal{A}}^\infty(w)\}$ is even, i.e., the least priority seen infinitely often is even. We let $L(\mathcal{A}) = \{w \in A^\omega \mid w \text{ is accepted by } \mathcal{A}\}$.

Existence of Finite-Memory Distributed Algorithms. We are now ready to state that, if there is a distributed algorithm that fulfills a specification $\varphi \in \text{LTL}(\mathcal{N})$, then there is also a distributed algorithm f with finite “synchronization memory” in the following sense: There is a DPWA \mathcal{A} over $\Sigma \times \Omega$ such that the output of a process for a history wu with $u \in \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^*$ only depends on u and the state that \mathcal{A} reaches after reading $f(\llbracket w \rrbracket)$. Let $\Sigma_{\perp \leftarrow} = \{\perp\} \times \{\leftarrow\} \times X_2$.

Lemma 2. *Let $\varphi \in \text{LTL}(\mathcal{N})$. There is a DPWA $\mathcal{A} = (S, \iota, \delta, \mathcal{F})$, with $\delta : S \times (\Sigma \times \Omega) \rightarrow S$, such that the following are equivalent:*

- (1) *There is a distributed algorithm $f = (f_1, f_2)$ that fulfills φ .*
- (2) *There is a distributed algorithm $f = (f_1, f_2)$ that fulfills φ and such that, for all words $w, w' \in \{\varepsilon\} \cup \Sigma_{\leftrightarrow} \Sigma^*$ satisfying $\delta(\iota, f(\llbracket w \rrbracket)) = \delta(\iota, f(\llbracket w' \rrbracket))$, the following hold:*
 - $f_1(wu) = f_1(w'u)$ for all $u \in \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^*$
 - $f_2(wu) = f_2(w'u)$ for all $u \in \Sigma_{\leftrightarrow} \Sigma_{\perp \leftarrow}^*$

Note that the acceptance condition and the language of \mathcal{A} are not important in the lemma.

3.2. Distributed Algorithms as Strategy Trees

Section 3.2 is devoted to the proof of Lemma 2. The first step is to represent a distributed algorithm as a *strategy tree*, whose branching structure reflects the algorithm’s choices depending on the various inputs. We then build a tree automaton that accepts a strategy tree iff it represents a distributed algorithm fulfilling the given formula φ . The challenge is to define the tree automaton in such a way that its strategies can be cast into hierarchical multiplayer games with *finite sets of observations*, and that winning strategies within these games are equivalent to distributed algorithms. We show in this section that this is possible by collapsing potentially unboundedly long input sequences into an unbounded branching structure. With this construction, we can show that, if the tree automaton recognizes *some* strategy tree, then it also accepts one that represents a finite-memory distributed algorithm.

Trees and Parity Tree Automata. Let A be a nonempty (possibly infinite) alphabet and D be a nonempty (possibly infinite) set of *directions*. An A -labeled D -tree is a mapping $t : D^* \rightarrow A$. In particular, ε is the root with label $t(\varepsilon)$, and ud , with $d \in D$, is the d -successor of node $u \in D^*$, with label $t(ud)$.

Definition 3. A (nondeterministic) *parity tree automaton (PTA)* over A -labeled D -trees is a tuple $\mathcal{T} = (S, \iota, \Delta, \mathcal{F})$ with finite set of states S , initial state $\iota \in S$, acceptance condition $\mathcal{F} : S \rightarrow \mathbb{N}$, and (possibly infinite) set of transitions $\Delta \subseteq S \times A \times S^D$.

A *run* of \mathcal{T} on an A -labeled D -tree t is an S -labeled D -tree $\rho : D^* \rightarrow S$ where $\rho(\varepsilon) = \iota$ (the root is assigned the initial state) and, for all $u \in D^*$, $(\rho(u), t(u), d \in D \mapsto \rho(ud)) \in \Delta$. The latter is the transition *applied* at u , and we denote it by $\text{trans}_\rho(u)$.

A path of run ρ is a word $\xi = d_0 d_1 d_2 \dots \in D^\omega$, inducing the sequence $\varepsilon, d_0, d_0 d_1, d_0 d_1 d_2, \dots$ of nodes visited along ξ . We let $\text{Inf}(\xi)$ be the set of states that occur infinitely often as the labels of these nodes. Path ξ is *accepting* if $\min\{\mathcal{F}(s) \mid s \in \text{Inf}(\xi)\}$ is even. Run ρ is *accepting* if all its paths are accepting. Finally, \mathcal{T} defines the language of A -labeled D -trees $L(\mathcal{T}) = \{t : D^* \rightarrow A \mid \text{there is an accepting run of } \mathcal{T} \text{ on } t\}$.

Lemma 3. *Let A be a singleton alphabet, D a nonempty (possibly infinite) set of directions, and \mathcal{T} a PTA over A -labeled D -trees (as A is a singleton, we say that \mathcal{T} is input-free). Call a run ρ of \mathcal{T} on the unique A -labeled D -tree rational if, for all $w, w' \in D^*$ with $\rho(w) = \rho(w')$, we have $\text{trans}_\rho(w) = \text{trans}_\rho(w')$. If $L(\mathcal{T}) \neq \emptyset$, then there is a rational accepting run of \mathcal{T} .*

Proof. The lemma essentially follows from the fact that parity games are positionally determined [40, 41].

Let $\mathcal{T} = (S, \iota, \Delta, \mathcal{F})$ be the given PTA. Without loss of generality, we assume that, for every $s \in S$, there exists an outgoing transition in Δ , i.e., a transition of the form $(s, (s_d)_{d \in D})$.

We transform \mathcal{T} into a turn-based parity game with a possibly infinite arena. We assume some familiarity with parity games; for an introduction see, e.g., [41]. The set of vertices of the arena is $S \cup \Delta$, with initial vertex ι . A vertex from S belongs to player 0 (whose aim is to show that the tree automaton has an accepting run on the unique A -labeled D -tree), and a vertex from Δ belongs to player 1 (whose aim is to show that there is no such run).

Accordingly, for every $t = (s, (s_d)_{d \in D}) \in \Delta$, player 0 has a transition from s to t in the parity game. Upon such a move of player 0, player 1 has to show that one of the branches will not be accepting and, to do so, picks one of the states s_d . Thus, for every $t = (s, (s_d)_{d \in D}) \in \Delta$ and $d \in D$, player 1 has a transition from t to s_d . The priority assigned to $s \in S$ is simply $\mathcal{F}(s)$. On the other hand, every vertex from Δ obtains priority $\max\{\mathcal{F}(s) \mid s \in S\} + 1$.

Now, assume that \mathcal{T} has an accepting run ρ on the unique A -labeled D -tree. This run can readily be transformed into a winning strategy $g : S(\Delta S)^* \rightarrow \Delta$ of player 0 in the parity game, which maps a possible history of a play to the next vertex chosen by player 0. The first move chosen by player 0 is from ι to $\text{trans}_\rho(\varepsilon)$, i.e., $g(\iota) = (\iota, (\rho(d))_{d \in D})$. Player 1 will then choose some $s \in \{\rho(d) \mid d \in D\}$, whereupon player 0 chooses $\text{trans}_\rho(d)$ for some $d \in D$ such that $s = \rho(d)$, and so on. Let us define g more formally. Fix an arbitrary well-founded total order on D . For $\pi \in S(\Delta S)^*$, $t = (\hat{s}, (s_d)_{d \in D}) \in \Delta$, and $s \in S$ such that πts is the prefix of a valid play in the parity game, we define a node $w_{\pi ts} \in D^*$ inductively as follows: $w_\iota = \varepsilon$ and $w_{\pi ts} = w_\pi \cdot d$ where d is the smallest direction with $s_d = s$. With this, we let $g(\pi) = \text{trans}_\rho(w_\pi)$. As ρ is accepting, any play compatible with g is winning for player 0.

As parity games are positionally determined, the existence of the winning strategy g of player 0 implies the existence of a *positional* winning strategy $g' : S \rightarrow \Delta$ of player 0, whose choice only depends on the current state $s \in S$. We transform g' into a rational accepting run ρ' of \mathcal{T} inductively as follows. We first set $\rho'(\varepsilon) = \iota$. Let $s \in S$ and $g'(s) = (s, (s_d)_{d \in D})$. For $d \in D$, we let $g'_d(s)$ refer to s_d . With this, given $w \in D^*$ and $d \in D$, we set $\rho'(wd) = g'_d(\rho'(w))$.

That is, ρ' is a run and, for all $w, w' \in D^*$ with $\rho'(w) = \rho'(w')$, we have $\text{trans}_{\rho'}(w) = \text{trans}_{\rho'}(w')$, meaning that ρ' is rational. It remains to show that ρ' is accepting. Pick any path $\xi = d_0 d_1 d_2 \dots \in D^\omega$. The sequence $\iota g'(\iota) g'_{d_0}(\iota) g'(g'_{d_0}(\iota)) g'_{d_1}(g'_{d_0}(\iota)) \dots$ is compatible with the strategy g' . As g' is a winning strategy, the subsequence $\iota g'_{d_0}(\iota) g'_{d_1}(g'_{d_0}(\iota)) \dots$, which equals $\rho'(\varepsilon) \rho'(d_0) \rho'(d_0 d_1) \dots$, satisfies the parity condition imposed by \mathcal{F} . Thus, ρ' is accepting. \square

Strategy Trees. Recall that our goal is to show Lemma 2 using strategy trees as a representation of distributed algorithms. Strategy trees are trees over the (infinite) set of directions $D = \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^*$, with the aim to isolate the positions where a resynchronization occurs, via a letter from Σ_{\leftrightarrow} . By Remark 2, we only have to consider $\Sigma_{\leftrightarrow} \Sigma^* = (\Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^*)^+ = D^+$. Hence, to avoid additional notation, we can identify nonempty words in D^* with words in $\Sigma_{\leftrightarrow} \Sigma^*$. It will always be clear from the context whether the underlying alphabet is D or Σ .

Intuitively, a node $u \in D^*$ represents a given history, and the label of u represents the outputs for possible continuations from $\Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^*$. More precisely, the set Λ of labels is the set of pairs $\lambda = (\lambda_1, \lambda_2)$ where $\lambda_1 : \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^* \rightarrow Y_1$ and $\lambda_2 : \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^* \rightarrow Y_2$. For $w \in \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^*$, we define $\lambda(w) \in (\Sigma_{\leftrightarrow} \times \Omega)(\Sigma_{\leftarrow} \times \Omega)^*$ as expected (cf. the definition of $f(w)$ for a distributed algorithm f). Similarly, for $w \in \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^\omega$, we obtain a word $\lambda(w) \in (\Sigma_{\leftrightarrow} \times \Omega)(\Sigma_{\leftarrow} \times \Omega)^\omega$.

A *strategy tree* is a Λ -labeled D -tree $t : D^* \rightarrow \Lambda$. For $u \in D^*$, let $(\lambda_1^u, \lambda_2^u)$ refer to $t(u)$. The distributed algorithm associated with t is denoted by f_t and is defined as $f_t = (f_1, f_2)$ as follows (recall that $\Sigma_{\leftarrow} = \{\perp\} \times \{\leftarrow\} \times X_2$):

- $f_1(uu') = \lambda_1^u(u')$ for all $u \in \{\varepsilon\} \cup \Sigma_{\leftrightarrow} \Sigma^*$ and $u' \in \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^*$
- $f_2(uu') = \lambda_2^u(u')$ for all $u \in \{\varepsilon\} \cup \Sigma_{\leftrightarrow} \Sigma^*$, and $u' \in \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^*$

In $\lambda_1^u(u')$ and $\lambda_2^u(u')$, we consider the unique decomposition of u over D so that f_1 and f_2 are well-defined.

Remark 3. The mapping $t \mapsto f_t$ is a bijection. In particular, for every distributed algorithm f , there is a strategy tree t such that $f_t = f$.

Example 4. Suppose $X_1 = X_2 = Y_1 = Y_2 = \{0, 1\}$. Figure 3 depicts a part of a strategy tree t . Its nodes are gray-shaded. The labels of nodes of t are themselves represented as (infinite) trees. Consider the input sequence $w = \langle 1 \leftrightarrow 1 \rangle \langle 0 \leftarrow 0 \rangle \langle 1 \leftrightarrow 1 \rangle \langle 0 \leftarrow 0 \rangle \in \Sigma_{\leftrightarrow} \Sigma^*$. To know what f_t outputs for the first two signals, we look at the blue-colored nodes of the trees associated with the root of t . To determine the outputs for the two remaining signals, we look at the red-colored nodes of the trees associated with node d . We thus get $f_t(w) = (\langle 1 \leftrightarrow 1 \rangle, (0, 0)) (\langle 0 \leftarrow 0 \rangle, (0, 1)) (\langle 1 \leftrightarrow 1 \rangle, (1, 0)) (\langle 0 \leftarrow 0 \rangle, (1, 1))$ for the whole word w . \triangleleft

Now, Lemma 2 is a consequence of the following lemma:

Lemma 4. Let $\varphi \in \text{LTL}(\mathcal{N})$. There is a DPWA $\mathcal{A} = (S, \iota, \delta, \mathcal{F})$, with $\delta : S \times (\Sigma \times \Omega) \rightarrow S$, such that the following are equivalent:




Figure 3: A strategy tree t .

- (1) There is a strategy tree t such that f_i fulfills φ .
- (2) There is a strategy tree t such that (a) f_i fulfills φ , and (b) for all $w, w' \in D^*$ with $\delta(t, f_i(w)) = \delta(t, f_i(w'))$, we have $t(w) = t(w')$.

Proof. Let $\varphi \in \text{LTL}(\mathcal{N})$ be the given formula. We first define \mathcal{A} and then prove its correctness in terms of the statement of Lemma 4 using a PTA \mathcal{T}_φ over strategy trees.

The DPWA \mathcal{A} . It is well known that there is a DPWA $\mathcal{A}_\varphi = (S_\varphi, \iota_\varphi, \delta_\varphi, \mathcal{F}_\varphi)$ over $\Sigma \times \Omega$, with doubly exponentially many states and exponentially many priorities, such that $L(\mathcal{A}_\varphi) = L(\varphi)$; see [42, 43]. We refer to states of \mathcal{A}_φ as $\mathcal{s} \in S_\varphi$. Let $\mathbb{P} = \{\mathcal{F}_\varphi(\mathcal{s}) \mid \mathcal{s} \in S_\varphi\}$, whose size is thus exponential.

Starting from \mathcal{A}_φ , we now define the DPWA $\mathcal{A} = (S, \iota, \delta, \mathcal{F})$ such that, for words that contain infinitely many \leftrightarrow , it is enough to look at the sequence of states reached by \mathcal{A} right before the \leftrightarrow -positions to determine whether the word is in $L(\mathcal{A}_\varphi)$ or not. The idea is to keep track of the minimal priority seen between two \leftrightarrow -positions. Accordingly, the set of states is $S = S_\varphi \times \mathbb{P}$, with initial state $\iota = (\iota_\varphi, \max \mathbb{P})$. Concerning the transitions, for $(\mathcal{s}, \theta) \in S$ and $\alpha = (\langle x_1 \rightrightarrows x_2 \rangle, (y_1, y_2)) \in \Sigma \times \Omega$, we define, letting $\theta' = \mathcal{F}_\varphi(\delta_\varphi(\mathcal{s}, \alpha))$.

$$\delta((\mathcal{s}, \theta), \alpha) = \begin{cases} (\delta_\varphi(\mathcal{s}, \alpha), \min\{\theta, \theta'\}) & \text{if } \rightrightarrows = \leftarrow \\ (\delta_\varphi(\mathcal{s}, \alpha), \theta') & \text{if } \rightrightarrows = \leftrightarrow. \end{cases}$$

Finally, the acceptance condition is given by $\mathcal{F}((\mathcal{s}, \theta)) = \theta$ for all $(\mathcal{s}, \theta) \in S$.

The following claim states that \mathcal{A} is correct wrt. executions with infinitely many synchronization points, while the acceptance condition is looking only at states reached right before these synchronizing points:

Claim 1. Let $w_0, w_1, w_2, \dots \in (\Sigma_{\leftrightarrow} \times \Omega)(\Sigma_{\leftarrow} \times \Omega)^*$. Moreover, let $w = w_0 w_1 w_2 \dots$ be the concatenation of all w_i . Set $s_0 = \iota$ and, for $i \in \mathbb{N}$, $s_{i+1} = (\mathcal{s}_{i+1}, \theta_{i+1}) = \delta(\iota, w_0 \dots w_i)$. Then,

$$w \in L(\mathcal{A}_\varphi) \iff \text{the sequence } s_0, s_1, s_2, \dots \text{ satisfies } \mathcal{F} \iff w \in L(\mathcal{A}).$$

Proof of Claim 1. For state $\mathcal{s} \in S_\varphi$ and $w = \alpha_0 \dots \alpha_{n-1} \in (\Sigma \times \Omega)^*$, let

$$\text{Visit}_{\mathcal{A}_\varphi}(\mathcal{s}, w) = \{\delta(\mathcal{s}, \alpha_0), \dots, \delta(\mathcal{s}, \alpha_1 \dots \alpha_{n-1})\}$$

be the set of states that are traversed by \mathcal{A}_φ when reading w from state β . Note that $Visit_{\mathcal{A}_\varphi}(\beta, w)$ does not necessarily contain β . For all $i \in \mathbb{N}$, we have $\theta_{i+1} = \min\{\mathcal{F}_\varphi(\beta) \mid \beta \in Visit_{\mathcal{A}_\varphi}(\beta_i, w_i)\}$. With this, we get:

$$\begin{aligned}
& w \in L(\mathcal{A}_\varphi) \\
\iff & \min\{\mathcal{F}_\varphi(\beta) \mid \beta \in Visit_{\mathcal{A}_\varphi}^\infty(w)\} \text{ is even} \\
\iff & \min\{\theta \mid \theta = \theta_{i+1} \text{ for infinitely many } i\} \text{ is even} \\
\iff & \text{the sequence } s_0, s_1, s_2, \dots \text{ satisfies the acceptance condition } \mathcal{F} \\
\iff & w \in L(\mathcal{A})
\end{aligned}$$

The last equivalence is due to the fact that the priority component is monotonically decreasing when \mathcal{A} is reading a word $\alpha_0 \dots \alpha_{n-1} \in (\Sigma_{\leftrightarrow} \times \Omega)(\Sigma_{\leftarrow} \times \Omega)^*$. In fact, for $(\beta'_0, \theta'_0) \in S$ and $(\beta'_{i+1}, \theta'_{i+1}) = \delta((\beta'_i, \theta'_i), \alpha_i)$, we have $\theta'_1 \geq \theta'_2 \geq \dots \geq \theta'_n$. (Claim 1) ■

The PTA \mathcal{T}_φ . To get finite-memory algorithms, we will rely on Lemma 3, which is based on tree automata. In fact, a crucial ingredient of the proof is a PTA \mathcal{T}_φ over Λ -labeled D -trees such that

$$L(\mathcal{T}_\varphi) = \{t \mid t \text{ is a strategy tree such that } f_t \text{ fulfills } \varphi\}.$$

It is defined by $\mathcal{T}_\varphi = (S, \iota, \Delta, \mathcal{F})$ where S, ι , and \mathcal{F} are taken from \mathcal{A} , and Δ is given by

$$\Delta = \left\{ (s = (\beta, \theta), \lambda, (s_d)_{d \in D}) \mid \begin{array}{l} s_d = \delta(s, \lambda(\langle d \rangle)) \text{ for all } d \in \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^* \quad \text{(T1)} \\ \lambda(\langle w \rangle) \in L(\mathcal{A}_\varphi[\beta]) \text{ for all } w \in \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^\omega \quad \text{(T2)} \end{array} \right\}.$$

Here, $\mathcal{A}_\varphi[\beta] = (S_\varphi, \beta, \delta_\varphi, \mathcal{F}_\varphi)$ is the automaton \mathcal{A}_φ but where ι_φ has been replaced by β as the initial state. While condition (T1) ‘‘unfolds’’ \mathcal{A} into the tree structure taking care of input sequences with infinitely many synchronization points, condition (T2) guarantees that the distributed algorithm behaves correctly should there be no more synchronization.

Correctness of \mathcal{T}_φ is shown below (Claim 2). Let us first finish the proof of Lemma 4:

Putting It Together. We obtain Lemma 4 as a corollary from Lemma 3 using \mathcal{T}_φ .

Direction (2) \implies (1) is trivial. Let us show (1) \implies (2) and suppose $L(\mathcal{T}_\varphi) \neq \emptyset$. Consider the input-free PTA $\mathcal{T}'_\varphi = (S, \iota, \Delta', \mathcal{F})$ obtained from \mathcal{T}_φ by replacing the transition relation with $\Delta' = \{(s, (s_d)_{d \in D}) \mid (s, \lambda, (s_d)_{d \in D}) \in \Delta\}$. Note that $L(\mathcal{T}'_\varphi) \neq \emptyset$. By Lemma 3, there is an accepting run ρ of \mathcal{T}'_φ such that, for all $w, w' \in D^*$ with $\rho(w) = \rho(w')$, we have $trans_\rho(w) = trans_\rho(w')$. For all transitions $\theta = (s, (s_d)_{d \in D}) \in \Delta'$, fix $\lambda^\theta \in \Lambda$ such that $(s, \lambda^\theta, (s_d)_{d \in D}) \in \Delta$. Let $t : D^* \rightarrow \Lambda$ be the strategy tree defined by $t(w) = \lambda^{trans_\rho(w)}$.

We have $t \in L(\mathcal{T}_\varphi)$. Therefore, f_t fulfills φ , i.e., (2a) holds. It remains to show (2b). Let $w, w' \in D^*$ with $\delta(\iota, f_i(\langle w \rangle)) = \delta(\iota, f_i(\langle w' \rangle))$. By induction, we can show that $\rho(w) = \delta(\iota, f_i(\langle w \rangle)) = \delta(\iota, f_i(\langle w' \rangle)) = \rho(w')$, i.e., $t(w) = t(w')$, which proves (2b). Indeed, $\delta(\iota, f_i(\langle \varepsilon \rangle)) = \iota = \rho(\varepsilon)$ and, for $u \in D^*$ and $d \in D$, we have $\delta(\iota, f_i(\langle ud \rangle)) = \delta(\iota, f_i(\langle u \rangle) \cdot \lambda^u(\langle d \rangle)) = \delta(\delta(\iota, f_i(\langle u \rangle)), \lambda^u(\langle d \rangle)) = \delta(\rho(u), \lambda^u(\langle d \rangle)) = \rho(ud)$. The last equation is by (T1) in the definition of the transition relation Δ of \mathcal{T}_φ .

This concludes the proof of Lemma 4. □

It remains to show correctness of \mathcal{T}_φ , which relies on Claim 1.

Claim 2. We have $L(\mathcal{T}_\varphi) = \{t \mid t \text{ is a strategy tree such that } f_t \text{ fulfills } \varphi\}$.

Proof of Claim 2. We have to consider two inclusions:

Inclusion \subseteq : Suppose $t \in L(\mathcal{T}_\varphi)$. For $u \in D^*$, let $\lambda^u = (\lambda_1^u, \lambda_2^u)$ refer to $t(u)$. There is an accepting run $\rho : D^* \rightarrow S$ of \mathcal{T}_φ on t . Let $w \in \Sigma_{\leftrightarrow} \Sigma^\omega$. We will show, using Claim 1, that $f_i(\downarrow w) \in L(\mathcal{A}_\varphi)$.

- Suppose $w = d_0 d_1 \dots d_{n-1} u$ where $d_0, \dots, d_{n-1} \in \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^*$, with $n \in \mathbb{N}$, and $u \in \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^\omega$. In particular, seen as a word over Σ , w contains only finitely many letters from Σ_{\leftrightarrow} . We have

$$f_i(\downarrow w) = \lambda^\varepsilon(\downarrow d_0) \cdot \lambda^{d_0}(\downarrow d_1) \cdot \lambda^{d_0 d_1}(\downarrow d_2) \cdot \dots \cdot \lambda^{d_0 d_1 \dots d_{n-2}}(\downarrow d_{n-1}) \cdot \lambda^{d_0 d_1 \dots d_{n-1}}(\downarrow u).$$

By the definition of Δ , we have

$$\begin{aligned} (\vartheta_1, \theta_1) &:= \rho(d_0) = \delta(\iota, \lambda^\varepsilon(\downarrow d_0)) \\ (\vartheta_2, \theta_2) &:= \rho(d_0 d_1) = \delta(\rho(d_0), \lambda^{d_0}(\downarrow d_1)) \\ &\vdots \\ (\vartheta_n, \theta_n) &:= \rho(d_0 d_1 d_2 \dots d_{n-1}) = \delta(\rho(d_0 d_1 \dots d_{n-2}), \lambda^{d_0 d_1 \dots d_{n-2}}(\downarrow d_{n-1})) \end{aligned}$$

and $\lambda^{d_0 d_1 \dots d_{n-1}}(\downarrow u) \in L(\mathcal{A}_\varphi[\vartheta_n])$. This implies

$$\begin{aligned} \vartheta_1 &= \delta_\varphi(\iota_\varphi, \lambda^\varepsilon(\downarrow d_0)) \\ \vartheta_2 &= \delta_\varphi(\vartheta_1, \lambda^{d_0}(\downarrow d_1)) \\ &\vdots \\ \vartheta_n &= \delta_\varphi(\vartheta_{n-1}, \lambda^{d_0 d_1 \dots d_{n-2}}(\downarrow d_{n-1})). \end{aligned}$$

Therefore, together with $\lambda^{d_0 d_1 \dots d_{n-1}}(\downarrow u) \in L(\mathcal{A}_\varphi[\vartheta_n])$, we obtain $f_i(\downarrow w) \in L(\mathcal{A}_\varphi)$.

- Suppose $w = d_0 d_1 d_2 \dots$ where $d_0, d_1, d_2, \dots \in \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^*$ for all $n \in \mathbb{N}$. Thus, w contains infinitely many letters from Σ_{\leftrightarrow} . We have

$$f_i(\downarrow w) = \lambda^\varepsilon(\downarrow d_0) \cdot \lambda^{d_0}(\downarrow d_1) \cdot \lambda^{d_0 d_1}(\downarrow d_2) \cdot \lambda^{d_0 d_1 d_2}(\downarrow d_3) \cdot \dots$$

Moreover, we have

$$\begin{aligned} s_1 &:= \rho(d_0) = \delta(\iota, \lambda^\varepsilon(\downarrow d_0)) \\ s_2 &:= \rho(d_0 d_1) = \delta(\rho(d_0), \lambda^{d_0}(\downarrow d_1)) \\ s_3 &:= \rho(d_0 d_1 d_2) = \delta(\rho(d_0 d_1), \lambda^{d_0 d_1}(\downarrow d_2)) \\ &\vdots \end{aligned}$$

As ρ is an accepting run on t , the sequence ι, s_1, s_2, \dots satisfies \mathcal{F} . By Claim 1, we obtain $f_i(\downarrow w) \in L(\mathcal{A}_\varphi)$.

Inclusion \supseteq : Suppose t is a strategy tree such that f_i fulfills φ . Again, for $u \in D^*$, let $\lambda^u = (\lambda_1^u, \lambda_2^u)$ refer to $t(u)$. We will construct an accepting run $\rho : D^* \rightarrow S$ of \mathcal{T}_φ on t . First of all, we let $\rho(\varepsilon) = \iota$.

Suppose that we defined $\rho(u)$ for $u = d_0 d_1 \dots d_{n-1} \in D^*$ where $d_0, \dots, d_{n-1} \in D = \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^*$, with $n \in \mathbb{N}$. For $d \in D$, we let

$$\rho(ud) = \delta(\rho(u), \lambda^u(\downarrow d)).$$

Claim 3. For all $u \in D^*$ and $u' \in \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^* \cup \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^\omega$, the following hold:

$$f_i(\downarrow uu') = f_i(\downarrow u) \cdot \lambda^u(\downarrow u') \tag{1}$$

$$\rho(u) = \delta(\iota, f_i(\downarrow u)) \tag{2}$$

Proof of Claim 3. The first statement is due to the definition of f_i . The second statement follows from an easy induction on u (see also end of Section 3):

$$\begin{aligned}
\rho(\varepsilon) &= \delta(\iota, \varepsilon) \\
\rho(ud) &= \delta(\rho(u), \lambda^u(\langle d \rangle)) \\
&= \delta(\delta(\iota, f_i(\langle u \rangle)), \lambda^u(\langle d \rangle)) \\
&= \delta(\iota, f_i(\langle u \rangle) \cdot \lambda^u(\langle d \rangle)) \\
&= \delta(\iota, f_i(\langle ud \rangle))
\end{aligned}$$

Note that the last equality is due to (1). (Claim 3) ■

Let $u \in D^*$ and $(\mathcal{J}, R) = \rho(u)$. Let us establish that $(\rho(u), \lambda^u, (\rho(ud))_{d \in D})$ is a transition of \mathcal{T}_φ :

(T1) We have $\rho(ud) = \delta(\rho(u), \lambda^u(\langle d \rangle))$ by the definition of ρ .

(T2) Let $u' \in \Sigma_{\rightarrow} \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^\omega$. As f_i fulfills φ , we have $f_i(\langle uu' \rangle) \in L(\mathcal{A}_\varphi)$. By Claim 3(1), $\lambda^u(\langle u' \rangle) \in L(\mathcal{A}_\varphi[\delta_\varphi(\iota_\varphi, f_i(\langle u \rangle)])$. By means of Claim 3(2) and the definition of \mathcal{A} wrt. to \mathcal{A}_φ , we can deduce $\lambda^u(\langle u' \rangle) \in L(\mathcal{A}_\varphi[\mathcal{J}])$.

Finally, we show that ρ is accepting. Let $d_0, d_1, d_2, \dots \in D$ and consider the path $\xi = d_0 d_1 d_2 \dots$ along with the induced infinite sequence

$$\rho(\varepsilon), \rho(d_0), \rho(d_0 d_1), \rho(d_0 d_1 d_2), \rho(d_0 d_1 d_2 d_3), \dots$$

Recall that we have

$$f_i(\langle w \rangle) = \lambda^\varepsilon(\langle d_0 \rangle) \cdot \lambda^{d_0}(\langle d_1 \rangle) \cdot \lambda^{d_0 d_1}(\langle d_2 \rangle) \cdot \lambda^{d_0 d_1 d_2}(\langle d_3 \rangle) \cdot \dots$$

as well as

$$\begin{aligned}
\rho(d_0) &= \delta(\iota, \lambda^\varepsilon(\langle d_0 \rangle)) \\
\rho(d_0 d_1) &= \delta(\rho(d_0), \lambda^{d_0}(\langle d_1 \rangle)) \\
\rho(d_0 d_1 d_2) &= \delta(\rho(d_0 d_1), \lambda^{d_0 d_1}(\langle d_2 \rangle)) \\
&\vdots
\end{aligned}$$

As $f_i(\langle w \rangle) \in L(\mathcal{A}_\varphi)$, by Claim 1, we have that ξ is accepting. (Claim 2) ■

4. From Finite-Memory Distributed Algorithms to Games

4.1. Games with Imperfect Information

The existence of finite-memory distributed algorithms shown in Section 3 paves the way for a reduction of the synthesis problem to (2, 1)-player games with imperfect information, where two players form a coalition against an environment in order to fulfill some objective. The main differences between games and the synthesis problem are twofold: Games are played in an arena, on a finite set of nodes (or states), while the input of the synthesis problem is a logical specification. More importantly, in a game, communication between players occurs implicitly, by observing the nodes that are visited. Hence, communication between players is bounded by the finite nature of the arena, whereas in the synthesis problem, processes can send an unbounded amount of information at each communication point. Recall that $P = \{1, 2\}$ is the set of processes. In the context of games, however, its elements are referred to as *players*.

Definition 4. A (2, 1)-player game is a tuple $\mathcal{G} = (V, v_0, W, \Gamma, (A_p, O_p, obs_p)_{p \in P}, \tau)$. Here, V is the finite set of nodes containing the initial node $v_0 \in V$. We assume a parity winning condition $W : V \rightarrow \mathbb{N}$. Moreover, Γ is the finite set of actions of the environment, A_p is the finite set of actions of player p , O_p is the finite set of observations of p , and $obs_p : V \times \Gamma \rightarrow O_p$ determines what p actually observes for a given node and environment action. Finally, $\tau : V \times \Gamma \times (A_1 \times A_2) \rightarrow V$ is the transition function.

The game proceeds in rounds $r \in \mathbb{N}$, the first round starting in v_0 . When a round starts in $v \in V$, the environment first chooses an action $\gamma \in \Gamma$. Players 1 and 2 do not see γ , but only $obs_1(v, \gamma)$ and $obs_2(v, \gamma)$, respectively. Once the players receive these observations, they simultaneously choose actions $a_1 \in A_1$ and $a_2 \in A_2$. The next state is $\tau(v, \gamma, (a_1, a_2))$, etc.

Accordingly, a *play* (starting from v_0) is a sequence $\pi = (v_0, \gamma_0)(v_1, \gamma_1) \dots \in (V \times \Gamma)^\omega$ such that, for all $r \in \mathbb{N}$, there is $(a_1, a_2) \in A_1 \times A_2$ such that $v_{r+1} = \tau(v_r, \gamma_r, (a_1, a_2))$. The observation that a player p collects in play π until round r is defined as $\llbracket (v_0, \gamma_0) \dots (v_r, \gamma_r) \rrbracket_p^{\text{game}} = obs_p(v_0, \gamma_0) \dots obs_p(v_r, \gamma_r) \in \mathcal{O}_p^*$. The play is *winning* (for the coalition of players 1 and 2) if $v_0 v_1 v_2 \dots$ satisfies the parity winning condition in the expected manner.

A *strategy* for player p is a mapping $g_p : \mathcal{O}_p^* \rightarrow A_p$. A *strategy profile* is a pair $g = (g_1, g_2)$ of strategies. We say that play $\pi = (v_0, \gamma_0)(v_1, \gamma_1) \dots$ is *compatible* with g if, for all $r \in \mathbb{N}$, we have $v_{r+1} = \tau(v_r, \gamma_r, (a_1^r, a_2^r))$ where $a_p^r = g_p(\llbracket (v_0, \gamma_0) \dots (v_r, \gamma_r) \rrbracket_p^{\text{game}})$. Strategy profile g is *winning* if all plays that are compatible with g are winning.

The following fact has been shown by Peterson and Reif [17] for games and corresponds to the undecidability result of Pnueli and Rosner [11] for two processes without communication.

Fact 2 (Peterson-Reif). *The following problem is undecidable: Given a (2, 1)-player game \mathcal{G} , is there a winning strategy profile?*

Therefore, we have to impose a restriction. It turns out that, when we translate the synthesis problem for $\mathcal{N} = \{\leftrightarrow, \leftarrow\}$ to games in Section 4.2, player 1 (who corresponds to process 1) will have perfect information. We say that player p has *perfect information* in \mathcal{G} if $\mathcal{O}_p = V \times \Gamma$ and obs_p is the identity function.

The following result is by van der Meyden and Wilke [19, Theorem 6] with a proof in [44, Theorem 1].

Fact 3 (van der Meyden-Wilke). *The following problem is decidable: Given a (2, 1)-player game \mathcal{G} such that player 1 has perfect information, is there a winning strategy profile?*

Note that the transition function of our game is deterministic so that we actually obtain decidability in exponential time exploiting a standard technique: We use a small tree automaton to represent the *global* (full information) winning strategies and another small alternating tree automaton for the local ones of player 2 that conform with some global strategy. The alternating automaton can be checked for nonemptiness in exponential time [45].

4.2. Reduction to Games

The analogies between synthesis and games suggest a natural translation of the former into the latter. However, the crucial difference being the access to histories, we rely on the fact that certain histories in distributed algorithms enjoy a finite abstraction. In fact, it is enough to reveal a bounded amount of information to player 2 at every environment action from Σ_{\leftrightarrow} .

Lemma 5. *Let $\varphi \in \text{LTL}(\mathcal{N})$ with $\mathcal{N} = \{\leftrightarrow, \leftarrow\}$. We can effectively construct a (2, 1)-player game \mathcal{G}_φ such that player 1 has perfect information and the following holds: There is a distributed algorithm that fulfills φ iff there is a winning strategy profile in \mathcal{G}_φ .*

Proof. By Remark 2, input sequences that do not start with a symbol from Σ_{\leftrightarrow} are discarded. Hence, we assume that those sequences are all trivially “winning”, i.e., $(\Sigma_{\leftarrow} \times \Omega)(\Sigma \times \Omega)^\omega \subseteq L(\varphi)$. Let $\mathcal{A} = (S, \iota, \delta, \mathcal{F})$ be the DPWA according to Lemma 2. Recall that $S = S_\varphi \times \mathbb{P}$, where S_φ is taken from \mathcal{A}_φ and $\mathbb{P} = \{\mathcal{F}_\varphi(\beta) \mid \beta \in S_\varphi\}$, and that the transition function is of the form $\delta : S \times (\Sigma \times \Omega) \rightarrow S$.

We construct the game $\mathcal{G}_\varphi = (V, v_0, W, \Gamma, (A_p, \mathcal{O}_p, obs_p)_{p \in P}, \tau)$ as follows. Obviously, player 1 corresponds to process 1 and player 2 to process 2. We simply set $V = S$ and $v_0 = \iota = (\iota_\varphi, \emptyset)$, and, for $(\beta, \theta) \in S$, $W((\beta, \theta)) = \mathcal{F}_\varphi(\beta)$.

Moreover, $\Gamma = \Sigma$, the idea being that the environment chooses the inputs and the network graph. Accordingly, processes 1 and 2 choose their outputs so that $A_1 = Y_1$ and $A_2 = Y_2$.

Player 1’s observations are $\mathcal{O}_1 = V \times \Sigma$ and we set $obs_1(s, \langle x_1 \rightleftharpoons x_2 \rangle) = (s, \langle x_1 \rightleftharpoons x_2 \rangle)$. Thus, Player 1 has full information. Player 2’s observations are $\mathcal{O}_2 = (S \times \Sigma_{\leftrightarrow}) \cup \Sigma_{\leftarrow}$ and we set

$$obs_2(s, \langle x_1 \rightleftharpoons x_2 \rangle) = \begin{cases} (s, \langle x_1 \leftrightarrow x_2 \rangle) & \text{if } \rightleftharpoons = \leftrightarrow \\ (\perp \leftarrow x_2) & \text{if } \rightleftharpoons = \leftarrow. \end{cases}$$

That is, when the environment chooses a synchronizing input signal, the current state of \mathcal{A} is revealed to player 2, which corresponds to passing the (abstracted) history to process 2. Finally, the transitions are given by $\tau(s, \langle x_1 \rightleftharpoons x_2 \rangle, (y_1, y_2)) = \delta(s, (\langle x_1 \rightleftharpoons x_2 \rangle, (y_1, y_2)))$.

Correctness of the reduction. There are two directions to show.

Claim 4. *If there is a winning strategy profile in \mathcal{G}_φ , then there is a distributed algorithm that fulfills φ .*

Proof of Claim 4. Let $g = (g_1, g_2)$ be a winning strategy profile in \mathcal{G}_φ , with $g_p : \mathcal{O}_p^+ \rightarrow Y_p$.

We define $\nu : \Sigma^* \rightarrow V$ and $\eta : \Sigma^* \rightarrow (V \times \Sigma)^*$ inductively by

$$\begin{aligned} \nu(\varepsilon) &= \iota \\ \eta(\varepsilon) &= \varepsilon \\ \nu(w \langle x_1 \rightleftharpoons x_2 \rangle) &= \tau(\nu(w), \langle x_1 \rightleftharpoons x_2 \rangle, (y_1, y_2)) \\ \eta(w \langle x_1 \rightleftharpoons x_2 \rangle) &= \eta(w) \cdot (\nu(w), \langle x_1 \rightleftharpoons x_2 \rangle) \end{aligned}$$

where $y_p = g_p(\llbracket \eta(w \langle x_1 \rightleftharpoons x_2 \rangle) \rrbracket_p^{\text{game}})$. That is, $\nu(w)$ is the node which is visited after input word w under strategy profile g , and $\eta(w)$ is the path corresponding to w in the game, starting at ι and applying g .

For every $p \in \{1, 2\}$ and $w \in \Sigma_{\leftrightarrow} \Sigma^*$, we define

$$f_p(\llbracket w \rrbracket_p) = g_p(\llbracket \eta(w) \rrbracket_p^{\text{game}}).$$

This is well-defined by construction of the game, using the fact that g is known by both players. Indeed, $\llbracket w \rrbracket_1 = w$, then it is possible to compute $\llbracket \eta(w) \rrbracket_1^{\text{game}}$ from $\llbracket w \rrbracket_1$. For player 2, one can show inductively that for all $w, w' \in \Sigma_{\leftrightarrow} \Sigma^*$ such that $\llbracket w \rrbracket_2 = \llbracket w' \rrbracket_2$, $\llbracket \eta(w) \rrbracket_2^{\text{game}} = \llbracket \eta(w') \rrbracket_2^{\text{game}}$.

Let $w = \sigma_0 \sigma_1 \sigma_2 \dots \in \Sigma_{\leftrightarrow} \Sigma^\omega$. We have to show that $f(\llbracket w \rrbracket) \in L(\varphi) = L(\mathcal{A}_\varphi)$. Let us determine the sequence s_0, s_1, s_2, \dots of states of \mathcal{A} visited while reading $f(\llbracket w \rrbracket)$. Set $s_0 = \iota$ and, for every $r \in \mathbb{N}$,

$$s_{r+1} = \delta(s_r, [\sigma_r, (f_1(\sigma_0 \dots \sigma_r), f_2(\llbracket \sigma_0 \dots \sigma_r \rrbracket_2))]).$$

For all $r \in \mathbb{N}$, we have

$$\begin{aligned} s_{r+1} &= \delta(s_r, [\sigma_r, (f_1(\sigma_0 \dots \sigma_r), f_2(\llbracket \sigma_0 \dots \sigma_r \rrbracket_2))]) \\ &= \delta(s_r, [\sigma_r, (g_1(\eta(\sigma_0 \dots \sigma_r)), g_2(\llbracket \eta(\sigma_0 \dots \sigma_r) \rrbracket_2^{\text{game}}))]) \\ &= \tau(s_r, \sigma_r, (g_1(\eta(\sigma_0 \dots \sigma_r)), g_2(\llbracket \eta(\sigma_0 \dots \sigma_r) \rrbracket_2^{\text{game}}))). \end{aligned}$$

Since g is winning and by the winning condition W of the game, we obtain $f(\llbracket w \rrbracket) \in L(\mathcal{A}_\varphi)$. (Claim 4) ■

Claim 5. *If there is a distributed algorithm that fulfills φ , then there is a winning strategy profile in \mathcal{G}_φ .*

Proof of Claim 5. Let $f = (f_1, f_2)$ be a distributed algorithm that fulfills φ . Due to Lemma 2, we can assume that for all words $w, w' \in \{\varepsilon\} \cup \Sigma_{\leftrightarrow} \Sigma^*$ satisfying $\delta(\iota, f(\llbracket w \rrbracket)) = \delta(\iota, f(\llbracket w' \rrbracket))$, we have $f_2(wu) = f_2(w'u)$ for all $u \in \Sigma_{\leftrightarrow} \Sigma_{\leftarrow}^*$.

We have to define a strategy profile $g = (g_1, g_2)$ for the game. Recall that $g_p : \mathcal{O}_p^+ \rightarrow A_p$. For every $s \in S$, we will define an ‘‘access string’’ $w_s \in \Sigma^*$ as follows: Set $w_\iota = \varepsilon$. Moreover, for $s \in S \setminus \{\iota\}$, fix any word $w_s \in \Sigma_{\leftrightarrow} \Sigma^*$ such that $\delta(\iota, f(\llbracket w_s \rrbracket)) = s$. If no such word exists, we let $w_s = \varepsilon$.

Note that, if the first environment action is not from Σ_{\leftrightarrow} , then we can output anything. So fix an arbitrary pair $(y_1, y_2) \in Y_1 \times Y_2$. Now, g is given as follows:

$$\begin{aligned} g_1 : & \begin{cases} (V \times \Sigma)^+ \rightarrow Y_1 \\ (v_0, \sigma_0) \dots (v_n, \sigma_n) \mapsto \begin{cases} f_1(\sigma_0 \dots \sigma_n) & \text{if } \sigma_0 \in \Sigma_{\leftrightarrow} \\ y_1 & \text{otherwise} \end{cases} \end{cases} \\ g_2 : & \begin{cases} \mathcal{O}_2^+ \rightarrow Y_2 \\ o \mapsto y_2 & \text{for } o \in (\Sigma_{\leftarrow}) \mathcal{O}_2^* \\ o \cdot (s, \langle x_1 \leftrightarrow x_2 \rangle) \cdot u \mapsto f_2(w_s \cdot \langle x_1 \leftrightarrow x_2 \rangle \cdot u) & \text{for } o \in \{\varepsilon\} \cup (S \times \Sigma_{\leftrightarrow}) \mathcal{O}_2^* \text{ and } u \in \Sigma_{\leftarrow}^* \end{cases} \end{aligned}$$

It remains to show that g is winning. So let $\pi = (s_0, \sigma_0)(s_1, \sigma_1)(s_2, \sigma_2) \dots$ be a play that is compatible with g , with $s_r = (\mathcal{J}_r, \theta_r)$. By our assumption that $(\Sigma_{\leftarrow} \times \Omega)(\Sigma \times \Omega)^\omega \subseteq L(\varphi)$, we only need to consider the case $\sigma_0 \in \Sigma_{\leftrightarrow}$. For all $r \in \mathbb{N}$, we have

$$(\mathcal{J}_{r+1}, \theta_{r+1}) = \tau((\mathcal{J}_r, \theta_r), \sigma_r, (a_1^r, a_2^r)) = \delta((\mathcal{J}_r, \theta_r), (\sigma_r, (a_1^r, a_2^r)))$$

where $a_p^r = g_p(\llbracket \pi_{\leq r} \rrbracket_p^{\text{game}})$ with $\pi_{\leq r} = (s_0, \sigma_0) \dots (s_r, \sigma_r)$.

It is enough to show that, for all $r \in \mathbb{N}$, we have

$$(\mathcal{J}_{r+1}, \theta_{r+1}) = \delta((\mathcal{J}_r, \theta_r), [\sigma_r, (f_1(\sigma_0 \dots \sigma_r), f_2(\llbracket \sigma_0 \dots \sigma_r \rrbracket_2))]).$$

We proceed by induction on the number k of letters from Σ_{\leftrightarrow} in $\pi_{\leq r}$. So suppose

$$\pi_{\leq r} = \underbrace{(s_0, \sigma_0) \dots (s_{m-1}, \sigma_{m-1})}_{=: w} (s_m, \langle x_1^m \leftrightarrow x_2^m \rangle) (s_{m+1}, \langle x_1^{m+1} \leftarrow x_2^{m+1} \rangle) \dots (s_r, \langle x_1^r \leftarrow x_2^r \rangle).$$

Set $\llbracket \varepsilon \rrbracket_2^{\text{game}} = \varepsilon$ and let $u = \langle \perp \leftarrow x_2^{m+1} \rangle \dots \langle \perp \leftarrow x_2^r \rangle$. Then,

$$\begin{aligned} s_{r+1} &= \delta(s_r, [\sigma_r, (g_1(\llbracket \pi_{\leq r} \rrbracket_1^{\text{game}}), g_2(\llbracket \pi_{\leq r} \rrbracket_2^{\text{game}}))]) \\ &= \delta(s_r, [\sigma_r, (g_1(\pi_{\leq r}), g_2(\llbracket w \rrbracket_2^{\text{game}} \cdot (s_m, \langle x_1^m \leftrightarrow x_2^m \rangle) \cdot u))]) \\ &= \delta(s_r, [\sigma_r, (f_1(\sigma_0 \dots \sigma_r), f_2(w_{s_m} \cdot \langle x_1^m \leftrightarrow x_2^m \rangle \cdot u))]) \\ &\stackrel{(*)}{=} \delta(s_r, [\sigma_r, (f_1(\sigma_0 \dots \sigma_r), f_2(\sigma_0 \dots \sigma_{m-1} \cdot \langle x_1^m \leftrightarrow x_2^m \rangle \cdot u))]) \\ &= \delta(s_r, [\sigma_r, (f_1(\sigma_0 \dots \sigma_r), f_2(\llbracket \sigma_0 \dots \sigma_r \rrbracket_2))]) \end{aligned}$$

Equation (*) is trivial for $m = 0$ (i.e., $k = 1$), as then $w_{s_m} = \varepsilon$ (by definition). Otherwise, it follows from the induction hypothesis: The word $\sigma_0 \dots \sigma_{m-1}$ contains $k - 1$ signals from Σ_{\leftrightarrow} . That is, if $m \geq 1$, then we have $s_m = \delta(\iota, f(\sigma_0 \dots \sigma_{m-1}))$. (Claim 5) ■

This concludes the proof of Lemma 5. □

Now we are ready to prove Theorem 2, stating that the problem $\text{SYNTHESIS}(\{\leftrightarrow, \leftarrow\})$ is decidable.

Proof of Theorem 2. Decidability follows by combining Lemma 5 and Fact 3.

Let us comment on the complexity. The size of \mathcal{A}_φ is doubly exponential in the length of the formula, and so are the size of \mathcal{A} and, therefore, the size of \mathcal{G}_φ . Deciding the winner of our (2, 1)-player game where one player has perfect information can be done in exponential time so that the overall decision procedure runs in 3-fold exponential time. □

Note that $\text{SYNTHESIS}(\{\leftrightarrow\})$, which is equivalent to centralized synthesis in presence of one single process, is 2EXPTIME-complete [3], from which we inherit the best known lower bound for our problem. Moreover, hierarchical information further increases the complexity: for static pipelines with variable number of processes, the problem is no longer elementary [11].

As, in the proof, the given LTL formula is translated into a DPWA, synthesis is decidable even when the specification is given by any common finite automaton over ω -words (starting with a nondeterministic Büchi automaton, we actually save one exponential wrt. LTL):

Corollary 1. *Over $\mathcal{N} = \{\leftrightarrow, \leftarrow\}$, the following problem is decidable: Given an ω -regular language $L \subseteq (\Sigma \times \Omega)^\omega$, is L realizable?*




Figure 4: Illustration of $\langle \cdot \rangle : \Sigma^* \rightarrow (\Sigma')^*$




Figure 5: The mappings $\langle \cdot \rangle : \Sigma^* \rightarrow (\Sigma')^*$ and $\langle \cdot \rangle : (\Sigma \times \Omega)^* \rightarrow (\Sigma' \times \Omega')^*$

5. Reduction from $\{\leftrightarrow, \leftarrow, \rightarrow\}$ to $\{\leftrightarrow, \leftarrow\}$

In this section, we show decidability for the network model $\mathcal{N} = \{\leftrightarrow, \leftarrow, \rightarrow\}$, with input alphabet $\Sigma = X_1 \times \mathcal{N} \times X_2$ and output alphabet $\Omega = Y_1 \times Y_2$. Recall that this also implies decidability for the network model $\{\leftarrow, \rightarrow\}$.

The idea is to reduce the problem to the case of the network model $\mathcal{N}' = \{\leftrightarrow, \leftarrow\}$ that we considered in Sections 3 and 4, choosing as input alphabet $\Sigma' = X'_1 \times \mathcal{N}' \times X'_2$ where $X'_1 = X'_2 = (X_1 \cup X_2) \uplus \{\#\}$, and as output alphabet $\Omega' = Y'_1 \times Y'_2$ where $Y'_1 = Y'_2 = (Y_1 \cup Y_2) \uplus \{\#\}$. To do so, we will rewrite the given specification $\varphi \in \text{LTL}(\mathcal{N})$ towards an (automata-based) specification over \mathcal{N}' in such a way that process 1 can always simulate the “more informed” process and process 2 simulates the other process. Roughly speaking, what we are looking for is a translation $\langle \cdot \rangle : \Sigma^* \rightarrow (\Sigma')^*$ of histories w over \mathcal{N} to histories $\langle w \rangle$ over \mathcal{N}' such that the view of process 1 in $\langle w \rangle$ is “congruent” to the view of the more informed process in w , and the view of process 2 in $\langle w \rangle$ is “congruent” to the view of the less informed process in w . Note that [20] also uses a simulation technique to cope with dynamically changing hierarchies.

Example 5. Before defining $\langle \cdot \rangle$ formally, we illustrate it in Figure 4 for a history w . Round 0 uses \leftarrow so that there is nothing to change. Round 1 employs \rightarrow so that process 1 henceforth simulates process 2 and vice versa. To make sure that the corresponding views in $\langle w \rangle$ are still “congruent”, we insert the dummy signal $\langle \# \leftrightarrow \# \rangle$. Actually, the gray-shaded view of process 1 in w after round 2 contains the same information as the gray-shaded view of process 2 in $\langle w \rangle$ after round 3. Though w encounters \leftrightarrow in round 3, we decide not to change roles again; we will only do so when facing another \leftarrow (like in round 5). \triangleleft

Formally, $\langle \cdot \rangle : \Sigma^* \rightarrow (\Sigma')^*$ is given by the sequential *transducer* shown in Figure 5. For the moment, we ignore the red part. A transition with label $\alpha \mid \beta$ reads α and transforms it into β . As the transducer is deterministic, it actually defines a function. When we include the red part, i.e., the symbols from Ω and Ω' , we obtain an extension to $\langle \cdot \rangle : (\Sigma \times \Omega)^* \rightarrow (\Sigma' \times \Omega')^*$. Finally, these mappings are extended to infinite words as expected.

Observe that the state of the transducer reached after reading $w \in \Sigma^*$ (or $w \in (\Sigma \times \Omega)^*$) reveals the process that process 1 is currently simulating. We denote this process by $\text{sim}_1(w)$. Accordingly, $\text{sim}_2(w) = 3 - \text{sim}_1(w)$ is the process that process 2 simulates after input sequence w . For the example word w in Figure 4, we get $\text{sim}_1(w) = 1$ and $\text{sim}_2(w) = 2$.

Note that, for all $w, w' \in \Sigma^*$ and $p \in \{1, 2\}$, such that $\llbracket w \rrbracket_p = \llbracket w' \rrbracket_p$, we have $\text{sim}_p(w) = \text{sim}_p(w')$. This is because the simulated process only depends on the sequence of links.

Note that the mappings $\langle \cdot \rangle$ are all injective. Indeed, at the first position that distinguishes w and w' , the transducer produces letters that distinguish $\langle w \rangle$ and $\langle w' \rangle$. There is an analogous statement for views.

Lemma 6. *For all $w, w' \in \Sigma^*$ and $p \in \{1, 2\}$, the following hold:*

- (a) $\llbracket \langle w \rangle \rrbracket_p = \llbracket \langle w' \rangle \rrbracket_p \implies \llbracket w \rrbracket_{\text{sim}_p(w)} = \llbracket w' \rrbracket_{\text{sim}_p(w')}$
- (b) $\llbracket w \rrbracket_p = \llbracket w' \rrbracket_p \implies \llbracket \langle w \rangle \rrbracket_{\text{sim}_p(w)} = \llbracket \langle w' \rangle \rrbracket_{\text{sim}_p(w')}$

Proof. We prove (a) and (b) separately.

Part (a): First, assume $p = 1$. Observe that for $w \in \Sigma^*$, $\langle\langle w \rangle\rangle \in (\Sigma')^*$. Further, within the domain $(\Sigma')^*$, $\llbracket \cdot \rrbracket_1$ is the identity. Together with the injectivity of $\langle\langle \cdot \rangle\rangle$, $\llbracket \langle\langle w \rangle\rangle \rrbracket_p = \llbracket \langle\langle w' \rangle\rangle \rrbracket_p$ implies $w = w'$; the lemma's statement follows for $p = 1$.

Second, assume $p = 2$ and that $\llbracket \langle\langle w \rangle\rangle \rrbracket_2 = \llbracket \langle\langle w' \rangle\rangle \rrbracket_2$. For $p' \in P$, let $\langle\langle \cdot \rangle\rangle_{p'}$ denote the transduction defined by the same transducer but with initial state p' . In particular, we have $\langle\langle \cdot \rangle\rangle = \langle\langle \cdot \rangle\rangle_1$. We start by observing that the function $\llbracket \cdot \rrbracket_2$ is length preserving and the projection onto a sequence of communication graphs is the same in $\langle\langle w \rangle\rangle$ and $\langle\langle w' \rangle\rangle$. Moreover, the latter are of the form

$$\begin{aligned}\langle\langle w \rangle\rangle &= \hat{u} \sigma \overbrace{\langle z_1 \leftarrow x_1 \rangle \dots \langle z_n \leftarrow x_n \rangle}^{\hat{v}} \\ \langle\langle w' \rangle\rangle &= \hat{u} \sigma' \overbrace{\langle z'_1 \leftarrow x_1 \rangle \dots \langle z'_n \leftarrow x_n \rangle}^{\hat{v}'}\end{aligned}$$

for some $\hat{u} \in (\Sigma')^*$ and $\sigma \in \{\varepsilon\} \cup \Sigma'_{\leftrightarrow}$ such that $\sigma \neq \varepsilon$ or $\hat{u} = \sigma = \varepsilon$.

- Suppose $\hat{u} = \sigma = \varepsilon$. Then, by the definition of $\langle\langle \cdot \rangle\rangle$, we have $w = \hat{v}$ and $w' = \hat{v}'$. We deduce $\llbracket w \rrbracket_{\text{sim}_2(w)} = \llbracket w' \rrbracket_{\text{sim}_2(w')}$ with $\text{sim}_2(w) = 2$.
- Suppose that $\varepsilon \neq \sigma = \langle \chi_1 \leftrightarrow \chi_2 \rangle \neq \langle \# \leftrightarrow \# \rangle$. Then, $w = uv$ and $w' = u'v'$ for some u, v, u', v' such that $\langle\langle u \rangle\rangle = \hat{u}$ and $\langle\langle u' \rangle\rangle = \hat{u}$ and $\langle\langle v \rangle\rangle_{\text{sim}_1(u)} = \sigma \hat{v}$ and $\langle\langle v' \rangle\rangle_{\text{sim}_1(u')} = \sigma \hat{v}'$. By injectivity of $\langle\langle \cdot \rangle\rangle$, we have $u = u'$.
 - Suppose $\text{sim}_1(u) = 1$, i.e., $\text{sim}_2(u) = 2$. By the definition of $\langle\langle \cdot \rangle\rangle$, we obtain $v = \sigma \hat{v}$ and $v' = \sigma \hat{v}'$. Therefore, $\llbracket w \rrbracket_2 = \llbracket w' \rrbracket_2 = u \sigma \langle \perp \leftarrow x_1 \rangle \dots \langle \perp \leftarrow x_n \rangle$.
 - Suppose $\text{sim}_1(u) = 2$, i.e., $\text{sim}_2(u) = 1$. By the definition of $\langle\langle \cdot \rangle\rangle$, we can deduce that $v = \langle \chi_2 \leftrightarrow \chi_1 \rangle \langle x_1 \rightarrow z_1 \rangle \dots \langle x_n \rightarrow z_n \rangle$ and $v' = \langle \chi_2 \leftrightarrow \chi_1 \rangle \langle x_1 \rightarrow z'_1 \rangle \dots \langle x_n \rightarrow z'_n \rangle$. We conclude $\llbracket w \rrbracket_1 = \llbracket w' \rrbracket_1 = u \langle \chi_2 \leftrightarrow \chi_1 \rangle \langle x_1 \rightarrow \perp \rangle \dots \langle x_n \rightarrow \perp \rangle$.
- Suppose that $\sigma = \langle \# \leftrightarrow \# \rangle$. Then, $n \geq 1$. Moreover, $w = uv$ and $w' = u'v'$ for some u, v, u', v' such that $\langle\langle u \rangle\rangle = \hat{u}$ and $\langle\langle u' \rangle\rangle = \hat{u}$ and $\langle\langle v \rangle\rangle_{\text{sim}_1(u)} = \sigma \hat{v}$ and $\langle\langle v' \rangle\rangle_{\text{sim}_1(u')} = \sigma \hat{v}'$. By injectivity of $\langle\langle \cdot \rangle\rangle$, we have $u = u'$.
 - Suppose $\text{sim}_1(u) = 1$. By the definition of $\langle\langle \cdot \rangle\rangle$, we obtain $v = \langle x_1 \rightarrow z_1 \rangle \dots \langle x_n \rightarrow z_n \rangle$ and $v' = \langle x_1 \rightarrow z'_1 \rangle \dots \langle x_n \rightarrow z'_n \rangle$. Therefore, $\text{sim}_2(w) = \text{sim}_2(w') = 1$. We have that $\llbracket w \rrbracket_1 = \llbracket uv \rrbracket_1 = \llbracket u \rrbracket_1 \langle x_1 \rightarrow \perp \rangle \dots \langle x_1 \rightarrow \perp \rangle = \llbracket uv' \rrbracket_1 = \llbracket w' \rrbracket_1$.
 - Suppose $\text{sim}_1(u) = 2$. Now, by the definition of $\langle\langle \cdot \rangle\rangle$, we obtain $v = \langle x_1 \leftarrow z_1 \rangle \dots \langle x_n \leftarrow z_n \rangle$ and $v' = \langle x_1 \leftarrow z'_1 \rangle \dots \langle x_n \leftarrow z'_n \rangle$. Therefore, $\text{sim}_2(w) = \text{sim}_2(w') = 2$. We have that $\llbracket w \rrbracket_2 = \llbracket uv \rrbracket_2 = \llbracket u \rrbracket_2 \langle \perp \leftarrow x_1 \rangle \dots \langle \perp \leftarrow x_n \rangle = \llbracket uv' \rrbracket_2 = \llbracket w' \rrbracket_2$.

Part (b): Suppose $\llbracket w \rrbracket_p = \llbracket w' \rrbracket_p$. Note that this implies $\text{sim}_p(w) = \text{sim}_p(w') =: p_w$.

First, assume that one of the following holds:

- $w = u \langle x_1 \leftrightarrow x_2 \rangle$, or
- $w = u \langle x_1 \leftarrow x_2 \rangle$ and $p = 1$, or
- $w = u \langle x_1 \rightarrow x_2 \rangle$ and $p = 2$.

Then, $w = w'$ and we are done.

For the remaining cases, we proceed by induction. The statement is obvious for $w = \varepsilon$.

Now, assume $w = u \langle x_1 \leftarrow x_2 \rangle$ and $p = 2$. Then, $p_w = 2$ and $\llbracket w \rrbracket_2 = \llbracket u \rrbracket_2 \langle \perp \leftarrow x_2 \rangle = \llbracket w' \rrbracket_2$. Thus, we have $w' = u' \langle x'_1 \leftarrow x_2 \rangle$ for some u' and x'_1 such that $\llbracket u \rrbracket_2 = \llbracket u' \rrbracket_2$. The latter implies $\text{sim}_2(u) = \text{sim}_2(u') =: p_u$. By induction hypothesis, we get $\llbracket \langle\langle u \rangle\rangle \rrbracket_{p_u} = \llbracket \langle\langle u' \rangle\rangle \rrbracket_{p_u}$.

Suppose $p_u = 1$. Then,

$$\begin{aligned} \llbracket \langle w \rangle \rrbracket_{p_w} &= \llbracket \langle u \langle x_1 \leftarrow x_2 \rangle \rangle \rrbracket_2 \\ &= \llbracket \langle u \rangle \langle \# \leftrightarrow \# \rangle \langle x_1 \leftarrow x_2 \rangle \rrbracket_2 \\ &= \langle u \rangle \langle \# \leftrightarrow \# \rangle \langle \perp \leftarrow x_2 \rangle \end{aligned}$$

and from $p_u = 1$ and $\llbracket u \rrbracket_2 = \llbracket u' \rrbracket_2$

$$\begin{aligned} &= \langle u' \rangle \langle \# \leftrightarrow \# \rangle \langle \perp \leftarrow x_2 \rangle \\ &= \llbracket \langle u' \rangle \langle \# \leftrightarrow \# \rangle \langle x'_1 \leftarrow x_2 \rangle \rrbracket_2 \\ &= \llbracket \langle u' \langle x'_1 \leftarrow x_2 \rangle \rangle \rrbracket_2 = \llbracket \langle w' \rangle \rrbracket_{p_w} \end{aligned}$$

Suppose $p_u = 2$. Then,

$$\begin{aligned} \llbracket \langle w \rangle \rrbracket_{p_w} &= \llbracket \langle u \langle x_1 \leftarrow x_2 \rangle \rangle \rrbracket_2 \\ &= \llbracket \langle u \rangle \langle x_1 \leftarrow x_2 \rangle \rrbracket_2 \\ &= \llbracket \langle u \rangle \rrbracket_2 \langle \perp \leftarrow x_2 \rangle \end{aligned}$$

and by induction hypothesis

$$\begin{aligned} &= \llbracket \langle u' \rangle \rrbracket_2 \langle \perp \leftarrow x_2 \rangle \\ &= \llbracket \langle u' \rangle \langle x'_1 \leftarrow x_2 \rangle \rrbracket_2 \\ &= \llbracket \langle u' \langle x'_1 \leftarrow x_2 \rangle \rangle \rrbracket_2 = \llbracket \langle w' \rangle \rrbracket_{p_w} \end{aligned}$$

Now, assume $w = u \langle x_1 \rightarrow x_2 \rangle$ and $p = 1$. Then, $p_w = 2$ and $\llbracket w \rrbracket_1 = \llbracket u \rrbracket_1 \langle x_1 \rightarrow \perp \rangle = \llbracket w' \rrbracket_1$. Thus, we have $w' = u' \langle x_1 \rightarrow x'_2 \rangle$ for some u' and x'_2 such that $\llbracket u \rrbracket_1 = \llbracket u' \rrbracket_1$. The latter implies $\text{sim}_1(u) = \text{sim}_1(u') =: p_u$. By induction hypothesis, we get $\llbracket \langle u \rangle \rrbracket_{p_u} = \llbracket \langle u' \rangle \rrbracket_{p_u}$.

Suppose $p_u = 1$. Then,

$$\begin{aligned} \llbracket \langle w \rangle \rrbracket_{p_w} &= \llbracket \langle u \langle x_1 \rightarrow x_2 \rangle \rangle \rrbracket_2 \\ &= \llbracket \langle u \rangle \langle \# \leftrightarrow \# \rangle \langle x_2 \leftarrow x_1 \rangle \rrbracket_2 \\ &= \langle u \rangle \langle \# \leftrightarrow \# \rangle \langle \perp \leftarrow x_1 \rangle \end{aligned}$$

and from $p_u = 1$ and $\llbracket u \rrbracket_1 = \llbracket u' \rrbracket_1$

$$\begin{aligned} &= \langle u' \rangle \langle \# \leftrightarrow \# \rangle \langle \perp \leftarrow x_1 \rangle \\ &= \llbracket \langle u' \rangle \langle \# \leftrightarrow \# \rangle \langle x'_2 \leftarrow x_1 \rangle \rrbracket_2 \\ &= \llbracket \langle u' \langle x_1 \rightarrow x'_2 \rangle \rangle \rrbracket_2 = \llbracket \langle w' \rangle \rrbracket_{p_w} \end{aligned}$$

Suppose $p_u = 2$. Then,

$$\begin{aligned} \llbracket \langle w \rangle \rrbracket_{p_w} &= \llbracket \langle u \langle x_1 \rightarrow x_2 \rangle \rangle \rrbracket_2 \\ &= \llbracket \langle u \rangle \langle x_2 \leftarrow x_1 \rangle \rrbracket_2 \\ &= \llbracket \langle u \rangle \rrbracket_2 \langle \perp \leftarrow x_1 \rangle \end{aligned}$$

and by induction hypothesis

$$\begin{aligned} &= \llbracket \langle u' \rangle \rrbracket_2 \langle \perp \leftarrow x_1 \rangle \\ &= \llbracket \langle u' \rangle \langle x'_2 \leftarrow x_1 \rangle \rrbracket_2 \\ &= \llbracket \langle u' \langle x_1 \rightarrow x'_2 \rangle \rangle \rrbracket_2 = \llbracket \langle w' \rangle \rrbracket_{p_w} \end{aligned}$$

This concludes the proof of Lemma 6. □

Moreover, the transducer can be applied to ω -regular languages in the following sense:

Lemma 7. *Given a DPWA \mathcal{A} over the alphabet $\Sigma \times \Omega$, there is a DPWA \mathcal{A}' over $\Sigma' \times \Omega'$ of linear size such that $L(\mathcal{A}') = \langle\langle L(\mathcal{A}) \rangle\rangle := \{\langle\langle w \rangle\rangle \mid w \in L(\mathcal{A})\}$.*

The proof is by a classical construction of a synchronized product of the transducer and the DPWA \mathcal{A} . Now, decidability for \mathcal{N} is due to Lemma 7 and the following result, whose proof crucially relies on injectivity of $\langle\langle \cdot \rangle\rangle$ and Lemma 6:

Lemma 8. *Let $\varphi \in \text{LTL}(\mathcal{N})$. The following statements are equivalent:*

- (i) *There is a distributed algorithm f (over \mathcal{N}) such that, for all $w \in \Sigma^\omega$, $f(w) \in L(\varphi)$.*
- (ii) *There is a distributed algorithm f' (over \mathcal{N}') such that, for all $w \in \Sigma^\omega$, $f'(\langle\langle w \rangle\rangle) \in \langle\langle L(\varphi) \rangle\rangle$.*

Proof. We start by showing (i) \rightarrow (ii). Let $f = (f_1, f_2)$ be a distributed algorithm over \mathcal{N} that fulfills $L(\varphi)$. Let $f' = (f'_1, f'_2)$ be a distributed algorithm over \mathcal{N}' such that, for all $w \in \Sigma^+$, $w' \in (\Sigma')^+$, and $p \in \{1, 2\}$,

$$\begin{aligned} f'_p(\llbracket \langle\langle w \rangle\rangle \rrbracket_p) &:= f_{\text{sim}_p(w)}(\llbracket w \rrbracket_{\text{sim}_p(w)}) \\ f'_p(w' \langle\# \leftrightarrow \# \rangle) &:= \#. \end{aligned} \tag{3}$$

Note that this is well-defined due to Lemma 6(a). We have to show that, for all $w \in \Sigma^\omega$, we get $f'(\langle\langle w \rangle\rangle) \in \langle\langle L(\varphi) \rangle\rangle$. This follows from the fact that, for all $w \in \Sigma^*$, we have $f'(\langle\langle w \rangle\rangle) = \langle\langle f(w) \rangle\rangle$ which we show by induction (in the following, let $f'(u)$ stand for $(f'_1(\llbracket u \rrbracket_1), f'_2(\llbracket u \rrbracket_2))$):

- From the definitions, we obtain $f'(\langle\langle \varepsilon \rangle\rangle) = \langle\langle f(\varepsilon) \rangle\rangle$.
- For $w \in \Sigma^*$ with $\text{sim}_1(w) = 1$ and $\hat{w} = w \langle x_1 \leftarrow x_2 \rangle$ (therefore, $\text{sim}_1(\hat{w}) = 1$), we have

$$\begin{aligned} f'(\langle\langle w \langle x_1 \leftarrow x_2 \rangle \rangle\rangle) &= f'(\langle\langle w \rangle\rangle \langle x_1 \leftarrow x_2 \rangle) \\ &= f'(\langle\langle w \rangle\rangle) \cdot (\langle x_1 \leftarrow x_2 \rangle, f'(\langle\langle w \rangle\rangle \langle x_1 \leftarrow x_2 \rangle)) \\ &= f'(\langle\langle w \rangle\rangle) \cdot (\langle x_1 \leftarrow x_2 \rangle, f'(\langle\langle \hat{w} \rangle\rangle)) \\ &\stackrel{(3)}{=} f'(\langle\langle w \rangle\rangle) \cdot (\langle x_1 \leftarrow x_2 \rangle, (f_1(\llbracket \hat{w} \rrbracket_1), f_2(\llbracket \hat{w} \rrbracket_2))) \\ &\stackrel{\text{IH}}{=} \langle\langle f(w) \rangle\rangle \cdot (\langle x_1 \leftarrow x_2 \rangle, (f_1(\llbracket \hat{w} \rrbracket_1), f_2(\llbracket \hat{w} \rrbracket_2))) \end{aligned}$$

and from $\text{sim}_1(f(w)) = \text{sim}_1(w) = 1$ (because the projection of $f(w)$ to Σ equals w)

$$\begin{aligned} &= \langle\langle f(w) \rangle\rangle \cdot (\langle x_1 \leftarrow x_2 \rangle, (f_1(\llbracket \hat{w} \rrbracket_1), f_2(\llbracket \hat{w} \rrbracket_2))) \\ &= \langle\langle f(w \langle x_1 \leftarrow x_2 \rangle) \rangle\rangle = \langle\langle f(\hat{w}) \rangle\rangle \end{aligned}$$

- For $w \in \Sigma^*$ with $\text{sim}_1(w) = 1$ and $\hat{w} = w \langle x_1 \leftrightarrow x_2 \rangle$ (therefore, $\text{sim}_1(\hat{w}) = 1$), we have

$$\begin{aligned} f'(\langle\langle w \langle x_1 \leftrightarrow x_2 \rangle \rangle\rangle) &= f'(\langle\langle w \rangle\rangle \langle x_1 \leftrightarrow x_2 \rangle) \\ &= f'(\langle\langle w \rangle\rangle) \cdot (\langle x_1 \leftrightarrow x_2 \rangle, f'(\langle\langle w \rangle\rangle \langle x_1 \leftrightarrow x_2 \rangle)) \\ &= f'(\langle\langle w \rangle\rangle) \cdot (\langle x_1 \leftrightarrow x_2 \rangle, f'(\langle\langle \hat{w} \rangle\rangle)) \\ &\stackrel{(3)}{=} f'(\langle\langle w \rangle\rangle) \cdot (\langle x_1 \leftrightarrow x_2 \rangle, (f_1(\llbracket \hat{w} \rrbracket_1), f_2(\llbracket \hat{w} \rrbracket_2))) \\ &\stackrel{\text{IH}}{=} \langle\langle f(w) \rangle\rangle \cdot (\langle x_1 \leftrightarrow x_2 \rangle, (f_1(\llbracket \hat{w} \rrbracket_1), f_2(\llbracket \hat{w} \rrbracket_2))) \end{aligned}$$

and from $\text{sim}_1(f(w)) = \text{sim}_1(w) = 1$

$$\begin{aligned} &= \langle\langle f(w) \rangle\rangle \cdot (\langle x_1 \leftrightarrow x_2 \rangle, (f_1(\llbracket \hat{w} \rrbracket_1), f_2(\llbracket \hat{w} \rrbracket_2))) \\ &= \langle\langle f(w \langle x_1 \leftrightarrow x_2 \rangle) \rangle\rangle = \langle\langle f(\hat{w}) \rangle\rangle \end{aligned}$$

We next show (ii) \rightarrow (i). Let $f' = (f'_1, f'_2)$ be a distributed algorithm over the network model \mathcal{N}' such that, for all $w \in \Sigma^\omega$, $f'(\llbracket w \rrbracket) \in \langle L(\varphi) \rangle$. We can assume that, for all $p \in P$ and $u \in (\Sigma')^*$, we have $f'_p(u \langle \# \leftrightarrow \# \rangle) = \#$. Let $f = (f_1, f_2)$ be the distributed algorithm over \mathcal{N} defined, for all $w \in \Sigma^+$ and $p \in \{1, 2\}$, by

$$f_p(\llbracket w \rrbracket_p) := f'_{\text{sim}_p(w)}(\llbracket \llbracket w \rrbracket \rrbracket_{\text{sim}_p(w)}). \quad (4)$$

This is well-defined due to Lemma 6(b). We have to show that, for all $w \in \Sigma^\omega$, $f(\llbracket w \rrbracket) \in L(\varphi)$. By injectivity of $\langle \cdot \rangle$, this follows from the fact that, for all $w \in \Sigma^*$, we have

$$\langle f(\llbracket w \rrbracket) \rangle = f'(\llbracket \llbracket w \rrbracket \rrbracket).$$

To show the latter, we again proceed by induction:

- From the definitions, we obtain $\langle f(\llbracket \varepsilon \rrbracket) \rangle = f'(\llbracket \llbracket \varepsilon \rrbracket \rrbracket)$.
- For $w \in \Sigma^*$ with $\text{sim}_1(w) = 1$ and $\hat{w} = w \langle x_1 \leftarrow x_2 \rangle$, we have

$$\begin{aligned} \langle f(\llbracket w \langle x_1 \leftarrow x_2 \rangle \rrbracket) \rangle &= \langle f(\llbracket w \rrbracket) \langle \langle x_1 \leftarrow x_2 \rangle, (f_1(\llbracket \hat{w} \rrbracket_1), f_2(\llbracket \hat{w} \rrbracket_2)) \rangle \rangle \\ &\stackrel{(4)}{=} \langle f(\llbracket w \rrbracket) \langle \langle x_1 \leftarrow x_2 \rangle, (f'_1(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_1), f'_2(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_2)) \rangle \rangle \end{aligned}$$

and from $\text{sim}_1(f(\llbracket w \rrbracket)) = \text{sim}_1(w) = 1$

$$\begin{aligned} &= \langle f(\llbracket w \rrbracket) \rangle \langle \langle x_1 \leftarrow x_2 \rangle, (f'_1(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_1), f'_2(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_2)) \rangle \rangle \\ &\stackrel{\text{IH}}{=} f'(\llbracket \llbracket w \rrbracket \rrbracket) \langle \langle x_1 \leftarrow x_2 \rangle, (f'_1(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_1), f'_2(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_2)) \rangle \rangle \\ &= f'(\llbracket \llbracket w \langle x_1 \leftarrow x_2 \rangle \rrbracket \rrbracket) = f'(\llbracket \llbracket \hat{w} \rrbracket \rrbracket) \end{aligned}$$

- For $w \in \Sigma^*$ with $\text{sim}_1(w) = 1$ and $\hat{w} = w \langle x_1 \leftrightarrow x_2 \rangle$, we have

$$\begin{aligned} \langle f(\llbracket w \langle x_1 \leftrightarrow x_2 \rangle \rrbracket) \rangle &= \langle f(\llbracket w \rrbracket) \langle \langle x_1 \leftrightarrow x_2 \rangle, (f_1(\llbracket \hat{w} \rrbracket_1), f_2(\llbracket \hat{w} \rrbracket_2)) \rangle \rangle \\ &\stackrel{(4)}{=} \langle f(\llbracket w \rrbracket) \langle \langle x_1 \leftrightarrow x_2 \rangle, (f'_1(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_1), f'_2(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_2)) \rangle \rangle \end{aligned}$$

and from $\text{sim}_1(f(\llbracket w \rrbracket)) = \text{sim}_1(w) = 1$

$$\begin{aligned} &= \langle f(\llbracket w \rrbracket) \rangle \langle \langle x_1 \leftrightarrow x_2 \rangle, (f'_1(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_1), f'_2(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_2)) \rangle \rangle \\ &\stackrel{\text{IH}}{=} f'(\llbracket \llbracket w \rrbracket \rrbracket) \langle \langle x_1 \leftrightarrow x_2 \rangle, (f'_1(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_1), f'_2(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_2)) \rangle \rangle \\ &= f'(\llbracket \llbracket w \langle x_1 \leftrightarrow x_2 \rangle \rrbracket \rrbracket) = f'(\llbracket \llbracket \hat{w} \rrbracket \rrbracket) \end{aligned}$$

- For $w \in \Sigma^*$ with $\text{sim}_1(w) = 1$ and $\hat{w} = w \langle x_1 \rightarrow x_2 \rangle$, we have

$$\begin{aligned} \langle f(\llbracket w \langle x_1 \rightarrow x_2 \rangle \rrbracket) \rangle &= \langle f(\llbracket w \rrbracket) \langle \langle x_1 \rightarrow x_2 \rangle, (f_1(\llbracket \hat{w} \rrbracket_1), f_2(\llbracket \hat{w} \rrbracket_2)) \rangle \rangle \\ &\stackrel{(4)}{=} \langle f(\llbracket w \rrbracket) \langle \langle x_1 \rightarrow x_2 \rangle, (f'_2(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_2), f'_1(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_1)) \rangle \rangle \end{aligned}$$

and from $\text{sim}_1(f(\llbracket w \rrbracket)) = \text{sim}_1(w) = 1$

$$\begin{aligned} &= \langle f(\llbracket w \rrbracket) \rangle \langle \langle \# \leftrightarrow \# \rangle, (\#, \#) \langle \langle x_2 \leftarrow x_1 \rangle, (f'_1(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_1), f'_2(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_2)) \rangle \rangle \rangle \\ &\stackrel{\text{IH}}{=} f'(\llbracket \llbracket w \rrbracket \rrbracket) \langle \langle \# \leftrightarrow \# \rangle, (\#, \#) \langle \langle x_2 \leftarrow x_1 \rangle, (f'_1(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_1), f'_2(\llbracket \llbracket \hat{w} \rrbracket \rrbracket_2)) \rangle \rangle \rangle \\ &= f'(\llbracket \llbracket w \rrbracket \rrbracket) \langle \langle \# \leftrightarrow \# \rangle, (\#, \#) \langle \langle x_2 \leftarrow x_1 \rangle, \\ &\quad (f'_1(\llbracket \llbracket w \rrbracket \langle \# \leftrightarrow \# \rangle \langle x_2 \leftarrow x_1 \rangle \rrbracket_1), f'_2(\llbracket \llbracket w \rrbracket \langle \# \leftrightarrow \# \rangle \langle x_2 \leftarrow x_1 \rangle \rrbracket_2)) \rangle \rangle \rangle \\ &= f'(\llbracket \llbracket w \rrbracket \langle \# \leftrightarrow \# \rangle \langle x_2 \leftarrow x_1 \rangle \rrbracket) \\ &= f'(\llbracket \llbracket w \langle x_1 \rightarrow x_2 \rangle \rrbracket \rrbracket) = f'(\llbracket \llbracket \hat{w} \rrbracket \rrbracket) \end{aligned}$$

- For $w \in \Sigma^*$ with $\text{sim}_1(w) = 2$ and $\hat{w} = w\langle x_1 \rightarrow x_2 \rangle$, we have

$$\begin{aligned} \langle\langle f(w\langle x_1 \rightarrow x_2 \rangle) \rangle\rangle &= \langle\langle f(w)(\langle x_1 \rightarrow x_2 \rangle), (f_1(\llbracket \hat{w} \rrbracket_1), f_2(\llbracket \hat{w} \rrbracket_2)) \rangle\rangle \\ &\stackrel{(4)}{=} \langle\langle f(w)(\langle x_1 \rightarrow x_2 \rangle), (f'_2(\llbracket \langle \hat{w} \rangle \rrbracket_2), f'_1(\llbracket \langle \hat{w} \rangle \rrbracket_1)) \rangle\rangle \end{aligned}$$

and from $\text{sim}_1(f(w)) = \text{sim}_1(w) = 2$

$$\begin{aligned} &= \langle\langle f(w)(\langle x_2 \leftarrow x_1 \rangle), (f'_1(\llbracket \langle \hat{w} \rangle \rrbracket_1), f'_2(\llbracket \langle \hat{w} \rangle \rrbracket_2)) \rangle\rangle \\ &\stackrel{\text{IH}}{=} f'(\langle\langle w \rangle\rangle)(\langle x_2 \leftarrow x_1 \rangle, (f'_1(\llbracket \langle \hat{w} \rangle \rrbracket_1), f'_2(\llbracket \langle \hat{w} \rangle \rrbracket_2))) \\ &= f'(\langle\langle w \rangle\rangle)(\langle x_2 \leftarrow x_1 \rangle, (f'_1(\llbracket \langle w \rangle \langle x_2 \leftarrow x_1 \rangle \rrbracket_1), f'_2(\llbracket \langle w \rangle \langle x_2 \leftarrow x_1 \rangle \rrbracket_2))) \\ &= f'(\langle\langle w \rangle \langle x_2 \leftarrow x_1 \rangle \rangle) \\ &= f'(\langle\langle w\langle x_1 \rightarrow x_2 \rangle \rangle) = f'(\langle\langle \hat{w} \rangle \rangle) \end{aligned}$$

- For $w \in \Sigma^*$ with $\text{sim}_1(w) = 2$ and $\hat{w} = w\langle x_1 \leftrightarrow x_2 \rangle$, we have

$$\begin{aligned} \langle\langle f(w\langle x_1 \leftrightarrow x_2 \rangle) \rangle\rangle &= \langle\langle f(w)(\langle x_1 \leftrightarrow x_2 \rangle), (f_1(\llbracket \hat{w} \rrbracket_1), f_2(\llbracket \hat{w} \rrbracket_2)) \rangle\rangle \\ &\stackrel{(4)}{=} \langle\langle f(w)(\langle x_1 \leftrightarrow x_2 \rangle), (f'_2(\llbracket \langle \hat{w} \rangle \rrbracket_2), f'_1(\llbracket \langle \hat{w} \rangle \rrbracket_1)) \rangle\rangle \end{aligned}$$

and from $\text{sim}_1(f(w)) = \text{sim}_1(w) = 2$

$$\begin{aligned} &= \langle\langle f(w)(\langle x_2 \leftrightarrow x_1 \rangle), (f'_1(\llbracket \langle \hat{w} \rangle \rrbracket_1), f'_2(\llbracket \langle \hat{w} \rangle \rrbracket_2)) \rangle\rangle \\ &\stackrel{\text{IH}}{=} f'(\langle\langle w \rangle\rangle)(\langle x_2 \leftrightarrow x_1 \rangle, (f'_1(\llbracket \langle \hat{w} \rangle \rrbracket_1), f'_2(\llbracket \langle \hat{w} \rangle \rrbracket_2))) \\ &= f'(\langle\langle w \rangle\rangle)(\langle x_2 \leftrightarrow x_1 \rangle, (f'_1(\llbracket \langle w \rangle \langle x_2 \leftrightarrow x_1 \rangle \rrbracket_1), f'_2(\llbracket \langle w \rangle \langle x_2 \leftrightarrow x_1 \rangle \rrbracket_2))) \\ &= f'(\langle\langle w \rangle \langle x_2 \leftrightarrow x_1 \rangle \rangle) \\ &= f'(\langle\langle w\langle x_1 \leftrightarrow x_2 \rangle \rangle) = f'(\langle\langle \hat{w} \rangle \rangle) \end{aligned}$$

- For $w \in \Sigma^*$ with $\text{sim}_1(w) = 2$ and $\hat{w} = w\langle x_1 \leftarrow x_2 \rangle$, we have

$$\begin{aligned} \langle\langle f(w\langle x_1 \leftarrow x_2 \rangle) \rangle\rangle &= \langle\langle f(w)(\langle x_1 \leftarrow x_2 \rangle), (f_1(\llbracket \hat{w} \rrbracket_1), f_2(\llbracket \hat{w} \rrbracket_2)) \rangle\rangle \\ &\stackrel{(4)}{=} \langle\langle f(w)(\langle x_1 \leftarrow x_2 \rangle), (f'_1(\llbracket \langle \hat{w} \rangle \rrbracket_1), f'_2(\llbracket \langle \hat{w} \rangle \rrbracket_2)) \rangle\rangle \end{aligned}$$

and from $\text{sim}_1(f(w)) = \text{sim}_1(w) = 2$

$$\begin{aligned} &= \langle\langle f(w)(\langle \# \leftrightarrow \# \rangle), (\#, \#)(\langle x_1 \leftarrow x_2 \rangle), (f'_1(\llbracket \langle \hat{w} \rangle \rrbracket_1), f'_2(\llbracket \langle \hat{w} \rangle \rrbracket_2)) \rangle\rangle \\ &\stackrel{\text{IH}}{=} f'(\langle\langle w \rangle\rangle)(\langle \# \leftrightarrow \# \rangle, (\#, \#)(\langle x_1 \leftarrow x_2 \rangle), (f'_1(\llbracket \langle \hat{w} \rangle \rrbracket_1), f'_2(\llbracket \langle \hat{w} \rangle \rrbracket_2))) \\ &= f'(\langle\langle w \rangle\rangle)(\langle \# \leftrightarrow \# \rangle, (\#, \#)(\langle x_1 \leftarrow x_2 \rangle, \\ &\quad (f'_1(\llbracket \langle w \rangle \langle \# \leftrightarrow \# \rangle \langle x_1 \leftarrow x_2 \rangle \rrbracket_1), f'_2(\llbracket \langle w \rangle \langle \# \leftrightarrow \# \rangle \langle x_1 \leftarrow x_2 \rangle \rrbracket_2))) \\ &= f'(\langle\langle w \rangle \langle \# \leftrightarrow \# \rangle \langle x_1 \leftarrow x_2 \rangle \rangle) \\ &= f'(\langle\langle w\langle x_1 \leftarrow x_2 \rangle \rangle) = f'(\langle\langle \hat{w} \rangle \rangle) \end{aligned}$$

This concludes the proof of Lemma 8. \square

In other words, an instance $\varphi \in \text{LTL}(\mathcal{N})$ of the synthesis problem can be reduced to the existence of a distributed algorithm f' over \mathcal{N}' , Σ' , and Ω' that fulfills $L = M \cup \langle\langle L(\varphi) \rangle\rangle$ where $M \subseteq (\Sigma' \times \Omega')^\omega$ is the set of words whose projection to Σ' is *not* contained in $\langle\langle \Sigma^\omega \rangle\rangle$. Using Lemma 7, we obtain a DPWA for L (of doubly exponential size) so that, by Corollary 1, the problem is decidable. Again, the overall procedure runs in 3-fold exponential time.

This concludes the proof of our main result, Theorem 1.

6. \times -block-boundedness

In this section, we propose an extension of the decidability result for subsets of $\{\times, \leftarrow, \rightarrow, \leftrightarrow\}^\omega$ where the number of consecutive empty links in an input sequence is uniformly bounded. We reduce the synthesis problem over network model $\mathcal{N} = \{\times, \leftarrow, \leftrightarrow\}$ to the synthesis problem over $\mathcal{N}' = \{\leftarrow, \leftrightarrow\}$. The general case can be obtained by combining this reduction with the previous one. We also show that relaxing the hypothesis over the number of consecutive empty links leads to undecidability, so the boundedness condition is necessary and sufficient for the synthesis to be decidable.

6.1. Decidability of $\mathcal{N} = \{\times, \leftarrow, \leftrightarrow\}$ for bounded sequences of \times

As previously, without loss of generality, we assume that input sequences start with a symbol from Σ_{\leftrightarrow} .

Definition 5. Let $w \in \mathcal{N}^\omega$ be a sequence of communication graphs. A *maximal \times -block* in w is a sequence of communication links $v \in \times^* \cup \{\times^\omega\}$ such that $uvu' = w$, u ends in a graph different from \times , and u' is either empty or starts with a graph different from \times . Given $K \in \mathbb{N}$, we say sequence w is *K - \times -block-bounded* (respectively *K -exact- \times -block-bounded*) if all maximal \times -blocks in w have size at most (resp. exactly) K .

Let $\varphi \in \text{LTL}(\mathcal{N})$ be a specification formula, and K a natural number. We say that a distributed algorithm f *K - \times -block-bounded-fulfills* (respectively *K -exact- \times -block-bounded-fulfills*) φ if, for all $w \in \Sigma^\omega$ such that $w|_{\mathcal{N}}$ (the projection of w on \mathcal{N}) is K - \times -block-bounded (respectively K -exact- \times -block-bounded), we have $f(w) \in L(\varphi)$. Moreover, we say that φ is *K - \times -block-bounded-realizable* (respectively *K -exact- \times -block-bounded-realizable*) if there is some distributed algorithm that K - \times -block-bounded-fulfills (respectively K -exact- \times -block-bounded-fulfills) φ . By slight abuse of notation we will also use these definitions with a language $L \subseteq (\Sigma \times \Omega)^\omega$ instead of a formula φ .

The problem we aim to solve in this section is hence $\text{SYNTHESIS}(\mathcal{N}, K)$, defined as follows:

SYNTHESIS(\mathcal{N}, K) **Input:** $\varphi \in \text{LTL}(\mathcal{N})$
 Question: Is φ K - \times -block-bounded-realizable?

To solve it, we use the exact variant of the \times -block-bounded synthesis, namely $\text{SYNTHESIS}_=(\mathcal{N}, K)$ defined by

SYNTHESIS₌(\mathcal{N}, K) **Input:** $\varphi \in \text{LTL}(\mathcal{N})$
 Question: Is φ K -exact- \times -block-bounded-realizable?

Both problems can be defined with an ω -regular language L as input instead of φ .

We now proceed in two steps. The first one is a reduction of $\text{SYNTHESIS}(\mathcal{N}, K)$ to $\text{SYNTHESIS}_=(\mathcal{N}, K)$, obtained by padding any block of less than K empty links to exactly K empty links. The second one reduces $\text{SYNTHESIS}_=(\mathcal{N}, K)$ to $\text{SYNTHESIS}(\{\leftrightarrow, \leftarrow\})$ by expanding these blocks of exactly K empty links using only links in $\mathcal{N}' = \{\leftrightarrow, \leftarrow\}$.

6.1.1. Reduction of $\text{SYNTHESIS}(\mathcal{N}, K)$ to $\text{SYNTHESIS}_=(\mathcal{N}, K)$

As mentioned above, the translation to K -exact- \times -block-bounded executions consists in adding the right number of dummy symbols ($\langle \# \times \# \rangle, \langle \# \# \rangle$) to obtain exactly K consecutive empty links. The mapping transforms some $w \in (\Sigma \times \Omega)^\omega$ into $\text{PAD}_K(w)$ in $A^*[(B^K A^+)]^\omega$, where $A = (X_1 \times \mathcal{N}' \times X_2) \times \Omega$ and $B = ((X_1 \times \{\times\} \times X_2) \times \Omega) \cup \{(\langle \# \times \# \rangle, \langle \# \# \rangle)\}$. This can be done by a transducer with $K + 1$ states to accommodate the possible sizes of empty link blocks less than or equal to K . For a history $w \in \Sigma^\omega$, $\text{PAD}_K(w)$ is defined in the obvious way and both versions are injective mappings.

Let $X_p^\# = X_p \cup \{\#\}$ and $Y_p^\# = Y_p \cup \{\#\}$ for $p = 1, 2$. We define $\Sigma^\# = X_1^\# \times \mathcal{N} \times X_2^\#$ and $\Omega^\# = Y_1^\# \times Y_2^\#$. As before, we can apply the transducer to ω -regular languages as follows:

Lemma 9. *Given a DPWA \mathcal{A} over the alphabet $\Sigma \times \Omega$, there is a DPWA \mathcal{A}' over $\Sigma^\# \times \Omega^\#$ of linear size such that $L(\mathcal{A}') = \text{PAD}_K(L(\mathcal{A})) := \{\text{PAD}_K(e) \mid e \in L(\mathcal{A})\}$.*

Given $\varphi \in \text{LTL}(\mathcal{N})$, we define the ω -regular language $L'_\varphi = \text{PAD}_K(L(\varphi)) \cup \overline{L_P}$ with $\overline{L_P}$ being the set of words from $(\Sigma^\# \times \Omega^\#)^\omega$ whose projection on $\Sigma^\#$ is not in $\text{PAD}_K(\Sigma^\omega)$.

We now show the reduction using the following lemma:

Lemma 10. Let $K \in \mathbb{N}$ and $\varphi \in \text{LTL}(N)$. The following statements are equivalent:

(i) There is a distributed algorithm f over N such that, for all K - \times -block-bounded $w \in \Sigma^\omega$, $f(\llbracket w \rrbracket) \in L(\varphi)$.

(ii) There is a distributed algorithm f' over N such that, for all K -exact- \times -block-bounded $w \in (\Sigma^\#)^\omega$, $f'(\llbracket w \rrbracket) \in L'_\varphi$.

Proof. We first show that (i) implies (ii). Let $f = (f_1, f_2)$ be a distributed algorithm fulfilling φ , for all K - \times -block-bounded inputs. For $p \in \{1, 2\}$ and for all $w \in \Sigma^+$, $w' \in (\Sigma^\#)^+$,

$$\begin{aligned} f'_p(\llbracket \text{PAD}_K(w) \rrbracket_p) &:= f_p(\llbracket w \rrbracket_p) \\ f'_p(\llbracket w' \langle \# \times \# \rangle \rrbracket_p) &:= \#. \end{aligned} \quad (5)$$

This is well defined because $\text{PAD}_K(w)$ does not modify the views of the processes, since it does not modify the pattern of the sequence of communication networks, and because the mapping is injective.

For all other inputs w , $f'_p(w) = y$ for some $y \in Y_p^\#$.

Let $w' \in (\Sigma^\#)^\omega$ and $f'(\llbracket w' \rrbracket)$ be the corresponding K -exact- \times -block-bounded execution. If $f'(\llbracket w' \rrbracket) \notin \overline{L}_P$, let $w \in \Sigma^\omega$ be the K - \times -block-bounded input sequence such that $\text{PAD}_K(w) = w'$. We show by induction that $f'(\llbracket w' \rrbracket) = \text{PAD}_K(f(\llbracket w \rrbracket))$.

The base case follows immediately from the definitions.

- For $w \in \Sigma^*$, $\langle x_1 \Rightarrow x_2 \rangle$, such that $\text{PAD}_K(w.\langle x_1 \Rightarrow x_2 \rangle) = \text{PAD}_K(w).\langle x_1 \Rightarrow x_2 \rangle$, we have

$$\begin{aligned} f'(\llbracket \text{PAD}_K(w.\langle x_1 \Rightarrow x_2 \rangle) \rrbracket) &= f'(\llbracket \text{PAD}_K(w) \rrbracket).\langle x_1 \Rightarrow x_2 \rangle, f'(\llbracket \text{PAD}_K(w.\langle x_1 \Rightarrow x_2 \rangle) \rrbracket) \\ &= f'(\llbracket \text{PAD}_K(w) \rrbracket).\langle x_1 \Rightarrow x_2 \rangle, f(w.\langle x_1 \Rightarrow x_2 \rangle) \end{aligned}$$

- For $w \in \Sigma^*$, $\langle x_1 \Rightarrow x_2 \rangle$, such that $\text{PAD}_K(w.\langle x_1 \Rightarrow x_2 \rangle) = \text{PAD}_K(w).\langle \# \times \# \rangle^k.\langle x_1 \Rightarrow x_2 \rangle$, for $k < K$, we have

$$\begin{aligned} f'(\llbracket \text{PAD}_K(w.\langle x_1 \Rightarrow x_2 \rangle) \rrbracket) &= f'(\llbracket \text{PAD}_K(w) \rrbracket).\langle \# \times \# \rangle, (\#, \#)^k.\langle x_1 \Rightarrow x_2 \rangle, f'(\llbracket \text{PAD}_K(w).\langle \# \times \# \rangle^k.\langle x_1 \Rightarrow x_2 \rangle \rrbracket) \\ &= f'(\llbracket \text{PAD}_K(w) \rrbracket).\langle \# \times \# \rangle, (\#, \#)^k.\langle x_1 \Rightarrow x_2 \rangle, f(w.\langle x_1 \Rightarrow x_2 \rangle) \\ &= \text{PAD}_K(f(\llbracket w \rrbracket)).\langle \# \times \# \rangle, (\#, \#)^k.\langle x_1 \Rightarrow x_2 \rangle, f(w.\langle x_1 \Rightarrow x_2 \rangle) \\ &= \text{PAD}_K(f(\llbracket w.\langle x_1 \Rightarrow x_2 \rangle \rrbracket)) \end{aligned}$$

Hence, since f K - \times -block-bounded-fulfills $L(\varphi)$, we have $f'(\llbracket w' \rrbracket) \in \text{PAD}_K(L(\varphi))$, and f' K -exact- \times -block-bounded-fulfills L'_φ .

We now show that (ii) implies (i). Let $f' = (f'_1, f'_2)$ be a distributed algorithm fulfilling L'_φ for all K -exact- \times -block-bounded inputs. For $p \in \{1, 2\}$ and all $w \in \Sigma^*$, we let

$$f_p(\llbracket w \rrbracket_p) := f'_p(\llbracket \text{PAD}_K(w) \rrbracket_p).$$

This is well-defined since for all $w, w' \in \Sigma^*$, if $\llbracket w \rrbracket_p = \llbracket w' \rrbracket_p$, then $\llbracket \text{PAD}_K(w) \rrbracket_p = \llbracket \text{PAD}_K(w') \rrbracket_p$.

Let $w \in \Sigma^\omega$ and $f(\llbracket w \rrbracket)$ be the corresponding K - \times -block-bounded execution. First observe that, since f' K -exact- \times -block-bounded-fulfills L'_φ , on an input prefix that might be continued into a word in $\text{PAD}_K(\Sigma^\omega)$, if f' does not output $(\#, \#)$, it cannot be winning because it cannot be in $\text{PAD}_K(L(\varphi))$ and it cannot ensure to be in \overline{L}_P . Hence, for any $w \in \Sigma^*$, $\langle x_1 \Rightarrow x_2 \rangle$, such that $\text{PAD}_K(w.\langle x_1 \Rightarrow x_2 \rangle) = \text{PAD}_K(w).\langle \# \times \# \rangle^k.\langle x_1 \Rightarrow x_2 \rangle$, for $k < K$, we have $f'(\text{PAD}_K(w).\langle \# \times \# \rangle^\ell) = (\#, \#)$, for $\ell \leq k$. We show that $\text{PAD}_K(f(\llbracket w \rrbracket)) = f'(\llbracket \text{PAD}_K(w) \rrbracket)$. Again, the base case follows from the definitions. Then,

- For $w \in \Sigma^*$, $\langle x_1 \Rightarrow x_2 \rangle$, such that $\text{PAD}_K(w.\langle x_1 \Rightarrow x_2 \rangle) = \text{PAD}_K(w).\langle x_1 \Rightarrow x_2 \rangle$, we have

$$\begin{aligned} \text{PAD}_K(f(\llbracket w.\langle x_1 \Rightarrow x_2 \rangle \rrbracket)) &= \text{PAD}_K(f(\llbracket w \rrbracket).\langle x_1 \Rightarrow x_2 \rangle, f(w.\langle x_1 \Rightarrow x_2 \rangle)) \\ &= \text{PAD}_K(f(\llbracket w \rrbracket).\langle x_1 \Rightarrow x_2 \rangle, f'(\llbracket \text{PAD}_K(w.\langle x_1 \Rightarrow x_2 \rangle) \rrbracket)) \\ &= \text{PAD}_K(f(\llbracket w \rrbracket).\langle x_1 \Rightarrow x_2 \rangle, f'(\text{PAD}_K(w).\langle x_1 \Rightarrow x_2 \rangle)) \\ &\stackrel{\text{IH}}{=} f'(\llbracket \text{PAD}_K(w) \rrbracket).\langle x_1 \Rightarrow x_2 \rangle, f'(\text{PAD}_K(w).\langle x_1 \Rightarrow x_2 \rangle)) \\ &= f'(\llbracket \text{PAD}_K(w.\langle x_1 \Rightarrow x_2 \rangle) \rrbracket) \end{aligned}$$

- For $w \in \Sigma^*$, $\langle x_1 \Rightarrow x_2 \rangle$, such that $\text{PAD}_K(w.\langle x_1 \Rightarrow x_2 \rangle) = \text{PAD}_K(w).(\# \times \#)^k.\langle x_1 \Rightarrow x_2 \rangle$, for $k < K$, we have

$$\begin{aligned}
\text{PAD}_K(f(\!|w.\langle x_1 \Rightarrow x_2 \rangle\!|)) &= \text{PAD}_K(f(\!|w|)|.\langle x_1 \Rightarrow x_2 \rangle, f(w.\langle x_1 \Rightarrow x_2 \rangle))) \\
&= \text{PAD}_K(f(\!|w|)|.\langle x_1 \Rightarrow x_2 \rangle, f'(\text{PAD}_K(w.\langle x_1 \Rightarrow x_2 \rangle))) \\
&= \text{PAD}_K(f(\!|w|)|).\langle \# \times \# \rangle, (\#, \#)^k.\langle x_1 \Rightarrow x_2 \rangle, f'(\text{PAD}_K(w.\langle x_1 \Rightarrow x_2 \rangle))) \\
&\stackrel{\text{IH}}{=} f'(\!|\text{PAD}_K(w)|).\langle \# \times \# \rangle, (\#, \#)^k.\langle x_1 \Rightarrow x_2 \rangle, f'(\text{PAD}_K(w.\langle x_1 \Rightarrow x_2 \rangle)) \\
&= f'(\!|\text{PAD}_K(w.\langle x_1 \Rightarrow x_2 \rangle)\!|)
\end{aligned}$$

Since f' K -exact- \times -block-bounded-fulfills L'_φ , we know that $f'(\!|\text{PAD}_K(w)|) \in \text{PAD}_K(L(\varphi))$. Hence $\text{PAD}_K(f(\!|w|)) \in \text{PAD}_K(L(\varphi))$ and $f(\!|w|) \in L(\varphi)$. \square

6.1.2. Reduction of $\text{SYNTHESIS}_=(\mathcal{N}, K)$ to $\text{SYNTHESIS}_=(\{\leftrightarrow, \leftarrow\})$

We next show how to transform any execution with sequences of exactly K consecutive empty links into an execution where all links are from $\mathcal{N}' = \{\leftrightarrow, \leftarrow\}$.

Given an input alphabet $\Sigma = X_1 \times \mathcal{N} \times X_2$ and output alphabet $\Omega = Y_1 \times Y_2$ for $\text{SYNTHESIS}_=(\mathcal{N}, K)$, we introduce a new symbol $\$$ and choose as new input alphabet $\Sigma' = X'_1 \times \mathcal{N}' \times X'_2$ where $X'_1 = X_1 \uplus \{\$\}$, $X'_2 = X_2 \uplus \{\$\}$, and as output alphabet $\Omega' = Y'_1 \times Y'_2$ where $Y'_1 = Y_1 \uplus \{\$\}$ and $Y'_2 = Y_2 \uplus \{\$\}$. We let $A = X_1 \times \mathcal{N}' \times X_2$ and $B = X_1 \times \{\times\} \times X_2$. Then, in an execution, any sequence

$$w = (\langle x_1 \times x'_1 \rangle, (y_1, y'_1)) \dots (\langle x_K \times x'_K \rangle, (y_K, y'_K))$$

in $(B \times \Omega)^K$ is transformed into

$$(\langle x_1 \leftarrow \$ \rangle, (y_1, \$)) \dots (\langle x_K \leftarrow \$ \rangle, (y_K, \$)) (\langle \$ \leftarrow x'_1 \rangle, (\$, y'_1)) \dots (\langle \$ \leftarrow x'_K \rangle, (\$, y'_K)),$$

the rest of the execution being unchanged. Hence, even if process 1 potentially receives information from process 2, this information is useless because it is always equal to $\$$. Moreover, being in \mathcal{N}' , it is always totally informed about the history of process 2. Thus, transforming an \times into a \leftarrow , potentially giving away all the history of process 2, does not convey any new information. If process 1 receives the inputs of process 2, because of the reduction, before it is supposed to know them, process 1 cannot use this information because the specification will ensure that it outputs only $\$$ during that time. The corresponding mapping, denoted by EXPAND_K , can be formally defined as follows: for every $x_1, \dots, x_K \in X_1$, $x'_1, \dots, x'_K \in X_2$, for every $y_1, \dots, y_K \in Y_1$ and $y'_1, \dots, y'_K \in Y_2$, for every $u \in \Sigma^*$, for every $w \in \Sigma^*.A^+B^k$ with $k < K - 1$, for every $w' \in \Sigma^*.A^+$,

$$\begin{aligned}
\text{EXPAND}_K(\langle x_1 \leftrightarrow x'_1 \rangle, (y_1, y'_1)) &= \langle x_1 \leftrightarrow x'_1 \rangle, (y_1, y'_1) \\
\text{EXPAND}_K(u.\langle x_1 \Rightarrow x'_1 \rangle, (y_1, y'_1)) &= \text{EXPAND}_K(u).\langle x_1 \Rightarrow x'_1 \rangle, (y_1, y'_1) \text{ if } \Rightarrow \in \{\leftrightarrow, \leftarrow\} \\
\text{EXPAND}_K(w.\langle x_1 \times x'_1 \rangle, (y_1, y'_1)) &= \text{EXPAND}_K(w).\langle x_1 \leftarrow \$ \rangle, (y_1, \$) \tag{6}
\end{aligned}$$

$$\begin{aligned}
\text{EXPAND}_K(w'.\langle x_1 \times x'_1 \rangle, (y_1, y'_1)) \dots (\langle x_K \times x'_K \rangle, (y_K, y'_K)) &= \text{EXPAND}_K(w'.\langle x_1 \times x'_1 \rangle, (y_1, y'_1)) \dots (\langle x_{K-1} \times x'_{K-1} \rangle, (y_{K-1}, y'_{K-1})) \\
&\quad (\langle x_K \leftarrow \$ \rangle, (y_K, \$)).(\langle \$ \leftarrow x'_1 \rangle, (\$, y'_1)) \dots (\langle \$ \leftarrow x'_K \rangle, (\$, y'_K)) \tag{7}
\end{aligned}$$

Observe that $\text{EXPAND}_K(w'.\langle x_1 \times x'_1 \rangle, (y_1, y'_1)) \dots (\langle x_K \times x'_K \rangle, (y_K, y'_K)) = \text{EXPAND}_K(w').\langle x_1 \leftarrow \$ \rangle, (y_1, \$) \dots (\langle x_K \leftarrow \$ \rangle, (y_K, \$)) (\langle \$ \leftarrow x'_1 \rangle, (\$, y'_1)) \dots (\langle \$ \leftarrow x'_K \rangle, (\$, y'_K))$, by applying $K - 1$ times equation (6) and then equation (7), which makes the catching up of the input/output of process 2.

This mapping can also be realized by a transducer, where remembering the delayed inputs and outputs of process 2 will produce a number of states in $O(|X_2||Y_2|^{(K+1)})$. One obtains:

Lemma 11. *Given a DPWA \mathcal{A} over the alphabet $\Sigma \times \Omega$, there is a DPWA \mathcal{A}' over $\Sigma' \times \Omega'$ of polynomial size in $\Sigma \times \Omega$ and exponential in K such that $L(\mathcal{A}') = \text{EXPAND}_K(L(\mathcal{A})) := \{\text{EXPAND}_K(e) \mid e \in L(\mathcal{A})\}$.*

Given a history $w \in \Sigma^\omega$, we also define $\text{EXPAND}_K(w)$ in the natural way. Further, one observes the following basic properties:

Lemma 12. For all $w, w' \in \Sigma^+$, for $p \in \{1, 2\}$, if $\llbracket w \rrbracket_p = \llbracket w' \rrbracket_p$, then $\llbracket \text{EXPAND}_K(w) \rrbracket_p = \llbracket \text{EXPAND}_K(w') \rrbracket_p$, and if $\llbracket \text{EXPAND}_K(w) \rrbracket_1 = \llbracket \text{EXPAND}_K(w') \rrbracket_1$ then $\llbracket w \rrbracket_1 = \llbracket w' \rrbracket_1$. Moreover, if $\text{EXPAND}_K(w) \notin \Sigma'^*(X_1 \times \{\leftarrow\} \times \{\$\})$, then if $\llbracket \text{EXPAND}_K(w) \rrbracket_2 = \llbracket \text{EXPAND}_K(w') \rrbracket_2$ then $\llbracket w \rrbracket_2 = \llbracket w' \rrbracket_2$.

We let $L_K = \{w' \in (\Sigma' \times \Omega)^\omega \mid \text{there exists } w \in \Sigma^\omega \text{ such that } \text{EXPAND}_K(w) = \pi_{\Sigma'}(w')\}$ with $\pi_{\Sigma'}$ be the projection on Σ' , and \overline{L}_K its complement. We then show the reduction using the following lemma:

Lemma 13. Let $K \in \mathbb{N}$ and $L \subseteq \Sigma \times \Omega$ be an ω -regular language. The following statements are equivalent:

(i) There is a distributed algorithm f over \mathcal{N} such that, for all K -exact- \times -block-bounded $w \in \Sigma^\omega$, $f(\llbracket w \rrbracket) \in L$.

(ii) There is a distributed algorithm f' over \mathcal{N}' such that, for all $w \in (\Sigma')^\omega$, $f'(\llbracket w \rrbracket) \in \text{EXPAND}_K(L) \cup \overline{L}_K$.

Proof. First let $f = (f_1, f_2)$ be a distributed algorithm K -exact- \times -block-bounded-fulfilling the specification L . We define $f' = (f'_1, f'_2)$ as follows: For all $w \in \Sigma^+$,

$$f'_1(\llbracket \text{EXPAND}_K(w) \rrbracket_1) = \begin{cases} \$ & \text{if } \text{EXPAND}_K(w) \in \Sigma'^*(\{\$\} \times \{\leftarrow\} \times X'_2) \\ f_1(\llbracket w \rrbracket_1) & \text{otherwise.} \end{cases}$$

Now let $w' \in \Sigma'^+$ for which there is no $w \in \Sigma^+$ such that $\text{EXPAND}_K(w) = w'$. If there exists some $u' \in \Sigma'^+$ such that $\text{EXPAND}_K(w) = w'u'$, then taking u' minimal,

$$f'_1(\llbracket w' \rrbracket_1) = \begin{cases} f_1(\llbracket w \rrbracket_1) & \text{if } w' \text{ ends with an element from } X'_1 \times \{\leftarrow\} \times \{\$\} \\ \$ & \text{otherwise.} \end{cases} \quad (8)$$

In all other cases, $f'_1(\llbracket w' \rrbracket_1) = \$$.

The first case of equation (8) corresponds to an expansion of an input sequence by rule (7), when the input sequence given to f'_1 is of the form $\text{EXPAND}_K(u.\langle x_1 \times x'_1 \rangle \dots \langle x_{K-1} \times x'_{K-1} \rangle . \langle x_K \leftarrow \$ \rangle)$. It is a truncated sequence of $\text{EXPAND}_K(w) = \text{EXPAND}_K(u.\langle x_1 \times x'_1 \rangle \dots \langle x_{K-1} \times x'_{K-1} \rangle . \langle x_K \times x'_K \rangle)$ for some x'_K . It is well defined: let $u' \in \Sigma'^+$ and $w \in \Sigma^+$ be such that $w'u' = \text{EXPAND}_K(w)$. Then for all $u'' \in \Sigma'^+$, for all $w'' \in \Sigma^+$ such that $w'u'' = \text{EXPAND}_K(w'')$, $\llbracket w \rrbracket_1 = \llbracket w'' \rrbracket_1$. This comes from the form of the common prefix w' ending with $\langle x_K \leftarrow \$ \rangle$, so that $u' = \langle \$ \leftarrow x'_1 \rangle \dots \langle \$ \leftarrow x'_K \rangle$ and $u'' = \langle \$ \leftarrow x'_1 \rangle \dots \langle \$ \leftarrow x'_K \rangle$ for some x'_1, \dots, x'_K and x''_1, \dots, x''_K , hence not impacting the view of Process 1.

For all $w \in \Sigma^+$,

$$f'_2(\llbracket \text{EXPAND}_K(w) \rrbracket_2) = \begin{cases} \$ & \text{if } \text{EXPAND}_K(w) \in \Sigma'^*(X_1 \times \{\leftarrow\} \times \{\$\}) \\ f_2(\llbracket w \rrbracket_2) & \text{otherwise.} \end{cases}$$

For any history $w = \alpha_1 \dots \alpha_k$, for any $0 \leq r < k$, we let $w_{-r} = \alpha_1 \dots \alpha_{k-r}$. Let now $w' \in \Sigma'^+$ for which there is no $w \in \Sigma^+$ such that $\text{EXPAND}_K(w) = w'$. If there exists some $u' \in \Sigma'^+$ such that $\text{EXPAND}_K(w) = w'u'$, then taking u' minimal,

$$f'_2(\llbracket w' \rrbracket_2) = \begin{cases} f_2(\llbracket w_{-|u'|} \rrbracket_2) & \text{if } w' \text{ ends with an element from } \{\$\} \times \{\leftarrow\} \times X_2 \\ \$ & \text{otherwise.} \end{cases} \quad (9)$$

Again, the first case of equation (9) handles the case where the input sequence given to f'_2 is of the form $w' = \text{EXPAND}_K(u.\langle x_1 \leftarrow \$ \rangle \dots \langle x_{K-1} \leftarrow \$ \rangle . \langle x_K \leftarrow \$ \rangle . \langle \$ \leftarrow x'_1 \rangle \dots \langle \$ \leftarrow x'_k \rangle)$, with $k < K$. Fix some $x_0 \in X_2$, and let $u' = (\langle \$ \leftarrow x_0 \rangle)^{K-k}$. Then

$$\begin{aligned} w'u' &= \text{EXPAND}_K(u.\langle x_1 \leftarrow \$ \rangle \dots \langle x_{K-1} \leftarrow \$ \rangle . \langle x_K \leftarrow \$ \rangle . \langle \$ \leftarrow x'_1 \rangle \dots \langle \$ \leftarrow x'_k \rangle . (\langle \$ \leftarrow x_0 \rangle)^{K-k}) \\ &= \text{EXPAND}_K(u.\langle x_1 \times x'_1 \rangle \dots \langle x_k \times x'_k \rangle \langle x_{k+1} \times x_0 \rangle \dots \langle x_K \times x_0 \rangle). \end{aligned}$$

Further, $w = u.\langle x_1 \times x'_1 \rangle \dots \langle x_k \times x'_k \rangle \langle x_{k+1} \times x_0 \rangle \dots \langle x_K \times x_0 \rangle$ and $w_{-|u'|} = u.\langle x_1 \times x'_1 \rangle \dots \langle x_k \times x'_k \rangle$.

In all other cases, $f'_2(\llbracket w' \rrbracket_2) = \$$.

By Lemma 12, (f'_1, f'_2) are well-defined.

To show that for all $w' \in (\Sigma')^\omega$, $f'(\llbracket w' \rrbracket) \in \text{EXPAND}_K(L) \cup \overline{L}_K$, it is enough to show that for all $w \in \Sigma^\omega$ K -exact- \times -block-bounded sequences, $f'(\llbracket \text{EXPAND}_K(w) \rrbracket) \in \text{EXPAND}_K(L)$. Let then $w \in \Sigma^\omega$ K -exact- \times -block-bounded sequence. We obtain that $f'(\llbracket \text{EXPAND}_K(w) \rrbracket) \in \text{EXPAND}_K(L)$ from the fact that $f(\llbracket w \rrbracket) \in L$ and $f'(\llbracket \text{EXPAND}_K(w) \rrbracket) = \text{EXPAND}_K(f(\llbracket w \rrbracket))$. Indeed, this can be shown by induction. We omit the straightforward base case.

- If $w = u\langle x_1 \Rightarrow x_2 \rangle$ with $u \in \Sigma_{\leftrightarrow} \Sigma^*$, $\Rightarrow \in \{\leftarrow, \leftrightarrow\}$, $x_1 \in X_1$, $x_2 \in X_2$, then

$$\begin{aligned}
f'(\text{EXPAND}_K(w)) &= f'(\text{EXPAND}_K(u).\langle x_1 \Rightarrow x_2 \rangle, (f'_1(\llbracket \text{EXPAND}_K(u.\langle x_1 \Rightarrow x_2 \rangle) \rrbracket_1), f'_2(\llbracket \text{EXPAND}_K(u.\langle x_1 \Rightarrow x_2 \rangle) \rrbracket_2))) \\
&= \text{EXPAND}_K(f(u)).\langle x_1 \Rightarrow x_2 \rangle, (f'_1(\llbracket \text{EXPAND}_K(u.\langle x_1 \Rightarrow x_2 \rangle) \rrbracket_1), f'_2(\llbracket \text{EXPAND}_K(u.\langle x_1 \Rightarrow x_2 \rangle) \rrbracket_2))) \\
&= \text{EXPAND}_K(f(u)).\langle x_1 \Rightarrow x_2 \rangle, (f_1(\llbracket u.\langle x_1 \Rightarrow x_2 \rangle \rrbracket_1), f_2(\llbracket u.\langle x_1 \Rightarrow x_2 \rangle \rrbracket_2))) \\
&= \text{EXPAND}_K(f(w)).
\end{aligned}$$

- If $w = u\langle x_1 \times x_2 \rangle$ with $u \in \Sigma_{\leftrightarrow} \Sigma^*.A.B^k$, $k < K - 1$, $x_1 \in X_1$, $x_2 \in X_2$, then

$$\begin{aligned}
f'(\text{EXPAND}_K(w)) &= f'(\text{EXPAND}_K(u).\langle x_1 \leftarrow \$ \rangle) \\
&= f'(\text{EXPAND}_K(u).\langle x_1 \leftarrow \$ \rangle, (f'_1(\llbracket \text{EXPAND}_K(u.\langle x_1 \times x_2 \rangle) \rrbracket_1), f'_2(\llbracket \text{EXPAND}_K(u.\langle x_1 \times x_2 \rangle) \rrbracket_2))) \\
&= f'(\text{EXPAND}_K(u).\langle x_1 \leftarrow \$ \rangle, (f_1(\llbracket u.\langle x_1 \times x_2 \rangle \rrbracket_1), \$)) \\
&= \text{EXPAND}_K(f(u)).\langle x_1 \leftarrow \$ \rangle, (f_1(\llbracket u.\langle x_1 \times x_2 \rangle \rrbracket_1), \$)) \\
&= \text{EXPAND}_K(f(u).\langle x_1 \times x_2 \rangle, (f_1(\llbracket u.\langle x_1 \times x_2 \rangle \rrbracket_1), f_2(\llbracket u.\langle x_1 \times x_2 \rangle \rrbracket_2))) \\
&= \text{EXPAND}_K(f(u.\langle x_1 \times x_2 \rangle)).
\end{aligned}$$

- If $w = u\alpha_1 \dots \alpha_K$ with $\alpha_i = \langle x_i \times x'_i \rangle$ and $u \in \Sigma_{\leftrightarrow} \Sigma^*.A^+$ then, writing

$$\begin{aligned}
y &= f'_1(\llbracket \text{EXPAND}_K(u.\alpha_1 \dots \alpha_{K-1}).\langle x_K \leftarrow \$ \rangle \rrbracket_1), \\
y' &= f'_2(\llbracket \text{EXPAND}_K(u.\alpha_1 \dots \alpha_{K-1}).\langle x_K \leftarrow \$ \rangle \rrbracket_2), \\
y_i &= f'_1(\llbracket \text{EXPAND}_K(u\alpha_1 \dots \alpha_{K-1}).\langle x_K \leftarrow \$ \rangle \dots \langle \$ \leftarrow x'_i \rangle \rrbracket_1) \\
y'_i &= f'_2(\llbracket \text{EXPAND}_K(u\alpha_1 \dots \alpha_{K-1}).\langle x_K \leftarrow \$ \rangle \dots \langle \$ \leftarrow x'_i \rangle \rrbracket_2),
\end{aligned}$$

we have:

$$\begin{aligned}
f'(\text{EXPAND}_K(u.\alpha_1 \dots \alpha_K)) &= f'(\text{EXPAND}_K(u.\alpha_1 \dots \alpha_{K-1}).\langle x_K \leftarrow \$ \rangle.\langle \$ \leftarrow x'_1 \rangle \dots \langle \$ \leftarrow x'_K \rangle) \\
&= f'(\text{EXPAND}_K(u.\alpha_1 \dots \alpha_{K-1}).\langle x_K \leftarrow \$ \rangle, (y, y')).\langle \$ \leftarrow x'_1 \rangle, (y_1, y'_1)) \dots \langle \$ \leftarrow x'_K \rangle, (y_K, y'_K)) \\
&= \text{EXPAND}_K(f(u.\alpha_1 \dots \alpha_{K-1}).\langle x_K \leftarrow \$ \rangle, (y, y')).\langle \$ \leftarrow x'_1 \rangle, (y_1, y'_1)) \dots \langle \$ \leftarrow x'_K \rangle, (y_K, y'_K))
\end{aligned}$$

Observe now that, by (8) and (9), $y = f_1(\llbracket u.\alpha_1 \dots \alpha_K \rrbracket_1)$, $y' = \$$, and for all $1 \leq i \leq K$, $y_i = \$$, and $y'_i = f_2(\llbracket u.\alpha_1 \dots \alpha_i \rrbracket_1)$. We obtain then that

$$f'(\text{EXPAND}_K(u.\alpha_1 \dots \alpha_K)) = \text{EXPAND}_K(f(u.\alpha_1 \dots \alpha_{K-1}\alpha_K)).$$

We show now that (ii) implies (i). Let $f' = (f'_1, f'_2)$ be a distributed algorithm such that for all $w \in \Sigma'^\omega$, $f'(w) \in \text{EXPAND}_K(L) \cup \overline{L_K}$.

Then we define $f = (f_1, f_2)$ as follows: For all $w \in \Sigma^*$,

$$f_1(\llbracket w \rrbracket_1) = \begin{cases} f'_1(\llbracket w'_{-K} \rrbracket_1) \text{ where } w' = \text{EXPAND}_K(w) & \text{if } w \in \Sigma^*.B^K \\ f'_1(\llbracket \text{EXPAND}_K(w) \rrbracket_1) & \text{otherwise.} \end{cases}$$

Let $w \in \Sigma^*.A.B^k$ for $k < K$. We let $x_0^1 \in X_1$ and $x_0^2 \in X_2$ two fixed input values and $\text{FILL}(w) = w.\langle x_0^1 \times x_0^2 \rangle^{K-k} \in \Sigma^*.A.B^K$. Then, we let

$$f_2(\llbracket w \rrbracket_2) = \begin{cases} f'_2(\llbracket w'_{-(K-k)} \rrbracket_2) \text{ where } w' = \text{EXPAND}_K(\text{FILL}(w)) & \text{if } w \in \Sigma^*.A.B^k, \text{ for some } k < K \\ f'_2(\llbracket \text{EXPAND}_K(w) \rrbracket_2) & \text{otherwise.} \end{cases}$$

It is easy to see that if $\llbracket w \rrbracket_2 = \llbracket w' \rrbracket_2$ then $\llbracket \text{FILL}(w) \rrbracket_2 = \llbracket \text{FILL}(w') \rrbracket_2$. Then, by Lemma 12 we establish that these strategies are well defined.

Let $w \in \Sigma^\omega$ be a K -exact- \times -block-bounded input sequence. We show that $\text{EXPAND}_K(f(w)) = f'(\text{EXPAND}_K(w))$, again omitting the base case of the induction.

- If $w = u\langle x_1 \Rightarrow x_2 \rangle$ with $u \in \Sigma^*$, $\Rightarrow \in \{\leftarrow, \leftrightarrow\}$, $x_1 \in X_1$, $x_2 \in X_2$, then

$$\begin{aligned}
\text{EXPAND}_K(f(\langle u\langle x_1 \Rightarrow x_2 \rangle)) &= \text{EXPAND}_K(f(\langle u \rangle, \langle x_1 \Rightarrow x_2 \rangle, f(u.\langle x_1 \Rightarrow x_2 \rangle))) \\
&= \text{EXPAND}_K(f(\langle u \rangle, \langle x_1 \Rightarrow x_2 \rangle, f(u.\langle x_1 \Rightarrow x_2 \rangle))) \\
&= \text{EXPAND}_K(f(\langle u \rangle, \langle x_1 \Rightarrow x_2 \rangle, f'(\text{EXPAND}_K(u.\langle x_1 \Rightarrow x_2 \rangle))) \\
&= f'(\langle \text{EXPAND}_K(u) \rangle, \langle x_1 \Rightarrow x_2 \rangle, f'(\text{EXPAND}_K(u.\langle x_1 \Rightarrow x_2 \rangle))) \\
&= f'(\langle \text{EXPAND}_K(u.\langle x_1 \Rightarrow x_2 \rangle) \rangle) = f'(\langle \text{EXPAND}_K(u.\langle x_1 \Rightarrow x_2 \rangle) \rangle)
\end{aligned}$$

- If $w = u\langle x_1 \times x_2 \rangle$ with $u \in \Sigma^*.A^+B^k$, $k < K - 1$, $x_1 \in X_1$, $x_2 \in X_2$, then

$$\begin{aligned}
\text{EXPAND}_K(f(\langle u\langle x_1 \times x_2 \rangle)) &= \text{EXPAND}_K(f(\langle u \rangle, \langle x_1 \times x_2 \rangle, f(u.\langle x_1 \times x_2 \rangle))) \\
&= \text{EXPAND}_K(f(\langle u \rangle, \langle x_1 \leftarrow \$ \rangle, (f_1(\llbracket u.\langle x_1 \times x_2 \rangle \rrbracket_1), \$))) \\
&= f'(\langle \text{EXPAND}_K(u) \rangle, \langle x_1 \leftarrow \$ \rangle, (f'_1(\llbracket \text{EXPAND}_K(u.\langle x_1 \times x_2 \rangle) \rrbracket_1), \$)) \\
&= f'(\langle \text{EXPAND}_K(u.\langle x_1 \times x_2 \rangle) \rangle) \text{ by (6)}
\end{aligned}$$

- If $w = u\alpha_1 \dots \alpha_K$ with $\alpha_i = \langle x_i \times x'_i \rangle$ and $u \in \Sigma^*.A^+$, we let $y_K = f_1(\llbracket u\alpha_1 \dots \alpha_K \rrbracket_1)$ and for $1 \leq i \leq K$, $y'_i = f_2(\llbracket u\alpha_1 \dots \alpha_i \rrbracket_2)$. Then,

$$\begin{aligned}
\text{EXPAND}_K(f(\langle u\alpha_1 \dots \alpha_K \rangle)) &= \text{EXPAND}_K(f(\langle u\alpha_1 \dots \alpha_{K-1} \rangle, \langle x_K \leftarrow \$ \rangle, (y_K, \$)).(\langle \$ \leftarrow x'_1 \rangle, (\$, y'_1)) \dots (\langle \$ \leftarrow x'_K \rangle, (\$, y'_K))) \\
&= f'(\langle \text{EXPAND}_K(u\alpha_1 \dots \alpha_{K-1}) \rangle, \langle x_K \leftarrow \$ \rangle, (y_K, \$)).(\langle \$ \leftarrow x'_1 \rangle, (\$, y'_1)) \dots (\langle \$ \leftarrow x'_K \rangle, (\$, y'_K)).
\end{aligned}$$

Observe that, by definition of f_1 , $y_K = f'_1(\llbracket \text{EXPAND}_K(u\alpha_1 \dots \alpha_{K-1}).\langle x_K \leftarrow \$ \rangle \rrbracket_1)$. Moreover, for all i , let $w'_i = \text{EXPAND}_K(u\alpha_1 \dots \alpha_i.\langle x_0^1 \times x_0^2 \rangle^{K-i})$, then $(w'_i)_{-(K-i)} = \text{EXPAND}_K(u.\langle x_1 \leftarrow \$ \rangle \dots \langle x_i \leftarrow \$ \rangle \langle x_0^1 \leftarrow \$ \rangle \dots \langle x_0^1 \leftarrow \$ \rangle \langle \$ \leftarrow x'_1 \rangle \dots \langle \$ \leftarrow x'_i \rangle)$. By definition of f_2 , $y'_i = f'_2(\llbracket (w'_i)_{-(K-i)} \rrbracket_2) = f'_2(\llbracket \text{EXPAND}_K(u\alpha_1 \dots \alpha_K)_{-(K-i)} \rrbracket_2)$. Then,

$$\begin{aligned}
\text{EXPAND}_K(f(\langle u\alpha_1 \dots \alpha_K \rangle)) &= f'(\langle \text{EXPAND}_K(u\alpha_1 \dots \alpha_{K-1}) \rangle, \langle x_K \leftarrow \$ \rangle, (y_K, \$)).(\langle \$ \leftarrow x'_1 \rangle, (\$, y'_1)) \dots (\langle \$ \leftarrow x'_K \rangle, (\$, y'_K)) \\
&= f'(\langle \text{EXPAND}_K(u\alpha_1 \dots \alpha_K) \rangle).
\end{aligned}$$

Then, since $f'(\langle \text{EXPAND}_K(w) \rangle) \in \text{EXPAND}_K(L) \cup \overline{L_K}$ by hypothesis, and $f'(\langle \text{EXPAND}_K(w) \rangle) \in L_K$ by definition, we deduce that $f'(\langle \text{EXPAND}_K(w) \rangle) \in \text{EXPAND}_K(L)$. Hence $\text{EXPAND}_K(f(\langle w \rangle)) \in \text{EXPAND}_K(L)$ and, by injectivity of EXPAND_K , $f(\langle w \rangle) \in L$. \square

From Lemma 10, Lemma 13, and Corollary 1, we obtain:

Theorem 3. $\text{SYNTHESIS}_{\{\leftrightarrow, \leftarrow, \times\}}(K)$ and $\text{SYNTHESIS}_{\{\leftrightarrow, \leftarrow, \times\}}(K)$ are decidable.

6.2. Undecidability of $\mathcal{N} = \{\leftrightarrow, \leftarrow, \times\}$ for finite sequences of \times

In this section, we show that the synthesis problem when considering only \times -block-finite executions is undecidable. We say that a sequence $w \in \mathcal{N}^\omega$ is \times -block-finite if any maximal \times -block in w is finite, and we adapt the definitions related to \times -block-bounded to that variant. We then consider the synthesis problem

SYNTHESIS $_{\times}(\mathcal{N})$ Input $\varphi \in \text{LTL}(\mathcal{N})$
Question: Is φ \times -block-finite realizable?

Theorem 4. *The problem $\text{SYNTHESIS}_{\times}(\{\leftrightarrow, \leftarrow, \times\})$ is undecidable.*

Proof. Let M be a deterministic Turing Machine with state set Q , input alphabet Σ and tape alphabet Γ . Using the same encoding as in [16], we can build an $\text{LTL}(\mathcal{N})$ specification φ such that φ is \times -block-finite realizable if and only if M does not halt.

To that end, suppose that $X_1 = X_2 = \{0, 1\}$ and that $Y_1 = Y_2 = \Gamma \cup Q \cup \{\#\}$, with $\#$ being a symbol not appearing in $Q \cup \Gamma$. The idea is to encode a natural number n by a sequence of inputs in $0^*1^n0\{0, 1\}^\omega$, and then to enforce that, when given this input, process 2 outputs the n -th configuration reached by M starting on the empty tape, denoted by C_n . Indeed, it is possible to write an $\text{LTL}(\mathcal{N})$ specification (see [16] for details) $\varphi = \varphi_1 \wedge \varphi_2 \wedge \varphi_3 \wedge \varphi_4$ that requires:

- (φ_1) If the input sequence for process i is in $0^q 1^n 0\{0, 1\}^\omega$ then the output of process i is in $\#^{q+n} \Gamma^* Q \Gamma^+ \#\omega$.
- (φ_2) If the input sequence for process i is in $0^q 10\{0, 1\}^\omega$, then the output of process i is $\#^{q+1} C_1 \#\omega$, with C_1 being the encoding of the initial configuration of M , starting on the empty tape.
- (φ_3) If the input sequence is in $(\langle 0 \times 0 \rangle)^q (\langle 1 \times 1 \rangle)^p \langle 0 \times 0 \rangle . \Sigma^\omega$ (i.e., if the input sequences define the same number, and are “synchronized”), then the process outputs must be equal.
- (φ_4) If the input sequence is in $(\langle 0 \times 0 \rangle)^q . \langle 1 \times 0 \rangle . (\langle 1 \times 1 \rangle)^n \langle 0 \times 0 \rangle . \Sigma^\omega$ (i.e., the input sequence on process 2 encode a number n and the input sequence on process 1 encodes the number $n + 1$, and they are “synchronized”), then the output of process 2 must encode a configuration C and the output of process 1 must encode a configuration C' such that C' is the successor configuration of C .

We can also define an LTL(\mathcal{N}) formula *halt* which is true iff the output of process 2 encodes the halting configuration of M . One then observes that the distributed algorithm where each process writes $\#^{q+n} C_n \#\omega$ on any input of the form $0^q 1^n 0\{0, 1\}^\omega$ fulfills φ .

Assume that there is a distributed algorithm f fulfilling φ . One can show, by induction on p (as in [11, 16]):

Lemma 14. *For all $p > 0$, for all $q \geq 0$, for all $w \in \Sigma^\omega$,*

$$\llbracket w \rrbracket_2 \in (\langle \perp \times 0 \rangle)^q (\langle \perp \times 1 \rangle)^p \langle \perp \times 0 \rangle . (\{\perp\} \times \{*\} \times X_2)^{|C_p|} . \Sigma^\omega \implies \pi_{Y_2}(f(\llbracket w \rrbracket)) = \#^{q+p+1} C_p \#\omega.$$

where $|C_p|$ is the size of C_p and $\pi_{Y_2}(f(\llbracket w \rrbracket))$ is the projection on the outputs of process 2 of $f(\llbracket w \rrbracket)$.

This lemma ensures that, if process 2 is given the encoding of a natural number p , and if the communication link is empty for sufficiently long, then any winning strategy will output the encoding of the p -th configuration of M . We then reduce the non-halting problem for M on the empty tape, to $\text{SYNTHESIS}_*(\mathcal{N})$ with specification $\varphi \wedge \text{G-halt}$. \square

Like in the previous sections, the results also hold for a slightly adapted problem: Given an ω -regular language L of allowed communication sequences and an LTL formula φ , is φ realizable over the allowed sequences? Noting that for any such language L one can decide if there exists a K such that L is K - \times -block-bounded, and if so, compute this K , this synthesis problem is decidable if and only if the number of consecutive empty links is uniformly bounded from above in L .

7. Conclusion

Highly dynamic networks can be modeled by a set \mathcal{N} of communication graphs from which the adversarial environment may choose arbitrary sequences. We showed that synthesis in such synchronous two-node system is decidable for LTL specifications if and only if the network model does not contain the empty network. Our model covers full-information protocols where nodes communicate their complete unbounded causal history.

While network models account for highly unstable networks, minimal stability guarantees can often be made in practical settings. In particular, arbitrarily long partitions of the two nodes by empty links may be overly pessimistic for a network model that contains the empty link. We showed that the synthesis problem is decidable in this context if and only if the number of consecutive empty links in an input sequence is uniformly bounded.

Future work is concerned with establishing the precise complexity of our problem: our work shows a 3-fold exponential time upper bound, and a 2-fold exponential time lower bound follows from the work by Pnueli and Rosner [3]. Finally, we plan to extend our model to distributed systems of arbitrary size. We conjecture that synthesis is solvable over a given network model if and only if, in each communication graph, any two nodes are connected via a directed path. This would yield an analogue of the information-fork criterion by Finkbeiner and Schewe [15], which applies to static architectures. It remains to be seen whether the ideas presented in [15] can be lifted to dynamic architectures with causal memory.

Acknowledgment. We thank Dietmar Berwanger for valuable feedback. We are grateful to the reviewers for their careful reading and their comments, which led to an improved presentation. We also thank an anonymous reviewer of a previous version for pointing out that using parity automata instead of Rabin automata saves one exponential. This work was partly supported by ANR FREDDA (ANR-17-CE40-0013).

References

- [1] A. Church, Applications of recursive arithmetic to the problem of circuit synthesis – summaries of talks, Institute for Symbolic Logic, Cornell University (1957).
- [2] M. O. Rabin, Automata on infinite objects and Church’s problem, Vol. 13, American Mathematical Soc., 1972. doi:10.1090/cbms/013.
- [3] A. Pnueli, R. Rosner, A framework for the synthesis of reactive modules, in: International Conference on Concurrency (Concurrency 88), Springer, 1988, pp. 4–17. doi:10.1007/3-540-50403-6_28.
- [4] A. Pnueli, R. Rosner, On the synthesis of a reactive module, in: 16th ACM SIGPLAN-SIGACT symposium on Principles of Programming Languages (POPL’89), 1989, pp. 179–190. doi:10.1145/75277.75293.
- [5] J. R. Büchi, L. H. Landweber, Solving sequential conditions by finite-state strategies, in: The Collected Works of J. Richard Büchi, Springer, 1990, pp. 525–541. doi:10.1007/978-1-4613-8928-6_29.
- [6] M. Abadi, L. Lamport, P. Wolper, Realizable and unrealizable specifications of reactive systems, in: International Colloquium on Automata, Languages, and Programming (ICALP’89), Springer, 1989, pp. 1–17. doi:10.1007/BFb0035748.
- [7] W. Thomas, On the synthesis of strategies in infinite games, in: Annual Symposium on Theoretical Aspects of Computer Science (STACS’95), Springer, 1995, pp. 1–13. doi:10.1007/3-540-59042-0_57.
- [8] O. Kupferman, M. Y. Vardi, Church’s problem revisited, *Bulletin of Symbolic Logic* 5 (2) (1999) 245–263. doi:10.2307/421091.
- [9] P. Gastin, B. Lerman, M. Zeitoun, Distributed games with causal memory are decidable for series-parallel systems, in: International Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS’04), Springer, 2004, pp. 275–286. doi:10.1007/978-3-540-30538-5_23.
- [10] P. Madhusudan, P. Thiagarajan, S. Yang, The MSO theory of connectedly communicating processes, in: International Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS’05), Springer, 2005, pp. 201–212. doi:10.1007/11590156_16.
- [11] A. Pnueli, R. Rosner, Distributed reactive systems are hard to synthesize, in: 31st Annual Symposium on Foundations of Computer Science (FoCS’90), IEEE, 1990, pp. 746–757. doi:10.1109/FSCS.1990.89597.
- [12] A. Pnueli, The temporal semantics of concurrent programs, *Theoretical Computer Science* 13 (1) (1981) 45–60. doi:10.1016/0304-3975(81)90110-9.
- [13] O. Kupferman, M. Y. Vardi, Synthesis with incomplete information, in: *Advances in Temporal Logic*, Springer, 2000, pp. 109–127. doi:10.1007/978-94-015-9586-5_6.
- [14] O. Kupferman, M. Y. Vardi, Synthesizing distributed systems, in: 16th Annual IEEE Symposium on Logic in Computer Science (LICS’01), IEEE, 2001, pp. 389–398. doi:10.1109/LICS.2001.932514.
- [15] B. Finkbeiner, S. Schewe, Uniform distributed synthesis, in: 20th Annual IEEE Symposium on Logic in Computer Science (LICS’05), IEEE, 2005, pp. 321–330. doi:10.1109/LICS.2005.53.
- [16] P. Gastin, N. Sznajder, M. Zeitoun, Distributed synthesis for well-connected architectures, *Formal Methods in System Design* 34 (3) (2009) 215–237. doi:10.1007/s10703-008-0064-7.
- [17] G. L. Peterson, J. H. Reif, Multiple-person alternation, in: 20th Annual Symposium on Foundations of Computer Science (FOCS’79), IEEE, 1979, pp. 348–363. doi:10.1109/SFCS.1979.25.
- [18] S. Mohalik, I. Walukiewicz, Distributed games, in: International Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS’03), Springer, 2003, pp. 338–351. doi:10.1007/978-3-540-24597-1_29.
- [19] R. Van der Meyden, T. Wilke, Synthesis of distributed systems from knowledge-based specifications, in: International Conference on Concurrency Theory (CONCUR’05), Springer, 2005, pp. 562–576. doi:10.1007/11539452_42.
- [20] D. Berwanger, A. B. Mathew, M. van den Bogaard, Hierarchical information and the synthesis of distributed strategies, *Acta Informatica* 55 (8) (2018) 669–701. doi:10.1007/s00236-017-0306-5.
- [21] N. A. Lynch, *Distributed algorithms*, Elsevier, 1996.
- [22] B. Charron-Bost, A. Schiper, The heard-of model: computing in distributed systems with benign faults, *Distributed Computing* 22 (1) (2009) 49–71. doi:10.1007/s00446-009-0084-6.
- [23] E. A. Akkoyunlu, K. Ekanadham, R. V. Huber, Some constraints and tradeoffs in the design of network communications, in: 5th ACM symposium on Operating Systems Principles, 1975, pp. 67–74. doi:10.1145/800213.806523.
- [24] L. Lamport, R. Shostak, M. Pease, The byzantine generals problem, *ACM Transactions on Programming Languages and Systems* 4 (3) (1982) 382–401. doi:10.1145/357172.357176.
- [25] K. A. Bartlett, R. A. Scantlebury, P. T. Wilkinson, A note on reliable full-duplex transmission over half-duplex links, *Communications of the ACM (CACM)* 12 (5) (1969) 260–261. doi:10.1145/362946.362970.
- [26] A. V. Aho, A. D. Wyner, M. Yannakakis, J. D. Ullman, Bounds on the size and transmission rate of communications protocols, *Computers & Mathematics with Applications* 8 (3) (1982) 205–214. doi:10.1016/0898-1221(82)90043-8.
- [27] Y. Afek, H. Attiya, A. Fekete, M. Fischer, N. Lynch, Y. Mansour, D.-W. Wang, L. Zuck, Reliable communication over unreliable channels, *Journal of the ACM (JACM)* 41 (6) (1994) 1267–1297. doi:10.1145/195613.195651.
- [28] F. Kuhn, N. Lynch, R. Oshman, Distributed computation in dynamic networks, in: 42nd ACM Symposium on Theory of Computing (STOC’10), 2010, pp. 513–522. doi:10.1145/1806689.1806760.
- [29] É. Coudouma, E. Godard, J. Peters, A characterization of oblivious message adversaries for which consensus is solvable, *Theoretical Computer Science* 584 (2015) 80–90.
- [30] B. Charron-Bost, M. Függer, T. Nowak, Approximate consensus in highly dynamic networks: The role of averaging algorithms, in: 42nd International Colloquium on Automata, Languages, and Programming (ICALP’15), 2015, pp. 528–539. doi:10.1007/978-3-662-47666-6_42.
- [31] Y. Velner, A. Rabinovich, Church synthesis problem for noisy input, in: International Conference on Foundations of Software Science and Computational Structures (FoSSaCS’11), Springer, 2011, pp. 275–289. doi:10.1007/978-3-642-19805-2_19.
- [32] R. Dimitrova, B. Finkbeiner, Synthesis of fault-tolerant distributed systems, in: International Symposium on Automated Technology for Verification and Analysis (ATVA’09), Springer, 2009, pp. 321–336. doi:10.1007/978-3-642-04761-9_24.

- [33] R. Fagin, Y. Moses, J. Y. Halpern, M. Y. Vardi, Reasoning about knowledge, MIT press, 2003.
- [34] B. Genest, H. Gimbert, A. Muscholl, 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, Vol. 7966 of Lecture Notes in Computer Science, Springer, 2013, pp. 275–286. doi:10.1007/978-3-642-39212-2_26.
- [35] H. Gimbert, On the control of asynchronous automata, in: 37th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'17), Vol. 93 of LIPIcs, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2018, pp. 30:1–30:15. doi:10.4230/LIPIcs.FSTTCS.2017.30.
- [36] I2C-bus specification and user manual, <https://www.nxp.com/docs/en/user-guide/UM10204.pdf> (2014).
- [37] ISO 11898-1:2015. Road vehicles — Controller area network (can) — part 1: Data link layer and physical signalling, <https://www.iso.org/standard/63648.html> (2014).
- [38] B. Bérard, B. Bollig, P. Bouyer, M. Függer, N. Sznajder, Synthesis in presence of dynamic links, in: 11th International Symposium on Games, Automata, Logics, and Formal Verification (GandALF'20), 2020. doi:10.4204/EPTCS.326.3.
- [39] W. Thomas, Automata on infinite objects, in: J. van Leeuwen (Ed.), Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics, Elsevier and MIT Press, 1990, pp. 133–191. doi:10.1016/B978-0-444-88074-1.50009-3.
- [40] E. A. Emerson, C. S. Jutla, Tree automata, mu-calculus and determinacy, in: Proceedings of FOCS'91, IEEE Computer Society, 1991, pp. 368–377.
- [41] W. Zielonka, Infinite games on finitely coloured graphs with applications to automata on infinite trees, TCS 200 (1-2) (1998) 135–183.
- [42] M. Y. Vardi, P. Wolper, Reasoning about infinite computations, Information and Computation 115 (1) (1994) 1–37. doi:10.1006/inco.1994.1092.
- [43] N. Piterman, From nondeterministic buchi and streett automata to deterministic parity automata, in: 21st Annual IEEE Symposium on Logic in Computer Science (LICS'06), IEEE, 2006, pp. 255–264.
- [44] R. Van der Meyden, T. Wilke, Synthesis of distributed systems from knowledge-based specifications. unsw-cse-tr-0504, Tech. rep., UNSW Sydney (2 2005).
- [45] C. Löding, Automata on infinite trees, in: J.-E. Pin (Ed.), Handbook of Automata Theory, Volume I. Theoretical Foundations, European Mathematical Society, 2021, p. 265–302. doi:10.4171/Automata.