

1 Controller Synthesis for Broadcast Networks with 2 Data

3 Corto Mascle   

4 LaBRI, Université de Bordeaux, France

5 MPI-SWS, Kaiserslautern, Germany

6 — Abstract —

7 We study the distributed controller synthesis problem in a parameterised setting. We search for
8 strategies that guarantee that no error state is reached, no matter the number of processes involved.
9 In the model at hand, processes communicate through unreliable broadcasts. Additionally, messages
10 are signed with data from an infinite alphabet, representing identifiers. Processes have local registers,
11 with which they can compare and store those data. Our main result is that controller synthesis is
12 decidable for this model. We also characterise the complexity of the problem for each number of
13 registers per process.

14 **2012 ACM Subject Classification** Theory of computation → Verification by model checking

15 **Keywords and phrases** Parameterised verification, distributed synthesis, broadcast networks

16 **Digital Object Identifier** 10.4230/LIPIcs.CVIT.2016.23

17 **1** Introduction

18 Distributed synthesis is a famously difficult problem, with many undecidability results on a
19 variety of models [19, 21, 12, 6]. The difficulty can be attributed to the presence of multiple
20 players with partial information. In order to find relevant models where the problem is
21 decidable, we take inspiration from the parameterised approach to distributed verification:
22 instead of verifying systems with large numbers of components, we abstract away that
23 number, and ask whether a property holds for any number of components. This is also useful
24 to verify protocols which are supposed to work on arbitrarily large networks.

25 In this paper we present a parameterised approach to distributed controller synthesis.
26 We first study one of the most popular parameterised models, reconfigurable broadcast
27 networks [9]. Then we go further and study the extension of those systems with data, as
28 introduced in [8]. In this extension, agents have unique identifiers which they can use to sign
29 messages and local registers which let them store and compare the signatures of received
30 messages. This considerably extends the expressivity of the model. Each agent possesses two
31 primary operations: broadcasting a letter from a finite alphabet along with the content of
32 one of its registers or receiving a message, and comparing its datum with its registers and/or
33 storing it in them. Broadcasts are lossy: when an agent sends a message, each other agent
34 may or may not receive it; the set of agents receiving the broadcast is non-deterministic. A
35 fundamental problem on such models is the *coverability problem*, which asks if a system has
36 a run from an initial configuration to one with at least one agent in a given state.

37 We formalise the controller synthesis problem as follows: processes have controllable
38 states, from which they can choose the next action, and uncontrollable ones, from which an
39 adversary may decide the next step. The question is whether there exists a local strategy
40 that chooses actions from controllable states so that for all N , a system made of N processes
41 applying this strategy cannot reach an error state. We establish the decidability of this
42 problem, and show tight complexity bounds on the problem, depending on the number of
43 registers that each agent has access to.



© Corto Mascle;
licensed under Creative Commons License CC-BY 4.0
42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:45



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

44 **Related work** This paper extends the work on verification of broadcast networks with
 45 data done in [8] and [13]. This model follows a recent effort to enrich parameterised models
 46 with data, usually to represent identifiers. For instance, the classic framework of population
 47 protocols has recently been extended with data [7], and some first decidability results have
 48 been established [26]. Let us also mention Petri nets with data [18]. The decidability of
 49 reachability in this model is an important open question, which saw some recent progress [16].
 50 The relation between those models is still blurry, but we can draw hope from the known
 51 relations between these models without data: population protocols are tightly linked to Petri
 52 nets [10] and a common restriction of population protocols, called immediate observation,
 53 can be encoded in reconfigurable broadcast networks [2].

54 The problem we consider here fits in a family of what could be called parameterised games,
 55 which involve one player against an arbitrary number of opponents. Some other instances
 56 are concurrent parameterised games [3], where players choose their actions in parallel, and
 57 population control models, a full-information turn-based formalism [4]. In this work, we
 58 focus on distributed strategies: we want each process to act based only on its local history.
 59 Similar ideas have been explored in [5] and [24, Chapter 11].

60 **Structure** Section 2 describes the model and the main problem. The central result of this
 61 work is the decidability of controller synthesis for broadcast networks of register automata.
 62 We present it incrementally: In Section 3 we use the much simpler case of broadcast networks
 63 without data to illustrate our approach. In Section 4 we extend this proof to the subclass
 64 of systems called **signature BGR** where processes can only send messages with their initial
 65 identifier. Finally, we present the main decidability result in Section 5. We highlight the
 66 common structure between these sections by using similar sequences of definitions and
 67 lemmas. For instance, Definition 11, Lemma 12, and Theorem 13 have counterparts in
 68 Section 4: Definition 17, Lemma 18, and Theorem 24.

69 In Section 6 we discuss a different choice of definition of local strategies, and argue that
 70 the problem is unaffected by this choice. To complete the picture, in Section 7 we show tight
 71 complexity bounds for the **SAFESTRAT** problems when each process has a single register.

Nb of registers \ Problem	$r = 0$	$r = 1$	$r \geq 2$
Coverability	P [9]	NP [13]	$\mathbf{F}_{\omega^\omega}$ [13]
Safe strategy synthesis	NP (Thm 13)	NEXPTIME (Thm 33)	$\mathbf{F}_{\omega^\omega}$ (Thm 16)

■ **Table 1** Complexity of **COVER** and **SAFESTRAT**, depending on the number of registers. All problems are complete for the indicated class. The results of the last column hold for any fixed $r \geq 2$ and when r is part of the input.

72 *This paper uses hyperlinks. Occurrences of a term are linked to its **definition**. The reader can*
 73 *click on words and symbols (or just hover over them on some PDF readers) to see the definition.*

74 **2 Preliminaries**

75 **2.1 Register transducers**

76 We start by describing the transition system of individual processes, which are register
 77 transducers. Then, in Section 2.2 we will define our broadcast network model as the
 78 composition of a finite but arbitrary number of those processes.

79 We fix an infinite set of *data* \mathbb{D} . We define a notion of **register transducer** that is well-suited
80 for the definition of our distributed systems. They receive and send *messages*, which are
81 pairs (m, d) made of a *letter* m from a finite alphabet and a *datum* d from \mathbb{D} .

82 ► **Definition 1** (Register transducer). A **register transducer** with r registers over domain \mathbb{D} is
83 given by a tuple $\mathcal{R} = (Q, \mathcal{M}, q_{init}, \Delta)$ with Q a finite set of states, with q_{init} the initial state,
84 \mathcal{M} a finite alphabet, and Δ a set of transitions which are of three kinds:

- 85 ■ $q \xrightarrow{\text{br}(m,i)} q'$ **broadcast transitions** that broadcast a message (m, d) with d the content of
86 register i ,
- 87 ■ $q \xrightarrow{\text{rec}(m,=i)} q'$ **equality transitions** that read a message (m, d) and check that d is in
88 register i ,
- 89 ■ $q \xrightarrow{\text{rec}(m,\downarrow i)} q'$ **record transitions** that read a message (m, d) where d is not in any register
90 and put d in register i .

91 Transitions of the two last kinds are called **reception transitions**. The **size** of \mathcal{R} , written $|\mathcal{R}|$,
92 is $|Q| + |\Delta| + r$.

93 A **local configuration** of \mathcal{R} is an element of $Q \times \mathbb{D}^r$, describing the current state and the
94 content of each register. A **local configuration** (q, c) is **initial** if $q = q_{init}$ and all registers
95 have the same content, i.e., there exists $d \in \mathbb{D}$ such that $c(i) = d$ for all $i \in [1, r]$.

96 Given a record transition $q \xrightarrow{\text{rec}(m,\downarrow i)} q'$ and a datum d , we can apply δ to go from (q, c)
97 to (q', c') by reading (m, d) if for all $j \in [1, r]$, $c(j) \neq d$ and $c'(i) = d$, and for all $j \neq i$,
98 $c'(j) = c(j)$.

99 Given an equality transition $q \xrightarrow{\text{rec}(m,=i)} q'$ and a datum d , we can apply δ to go from
100 (q, c) to (q', c') by reading (m, d) if $c(i) = d$ and $c' = c$.

101 If one of those cases applies, we write $(q, c) \xrightarrow{\text{rec}(m,d)}_{\delta} (q', c')$ and call it a **reception step**.

102 Given a broadcast transition $\delta = q \xrightarrow{\text{br}(m,i)} q'$ and a datum d , we can apply δ to go from
103 (q, c) to (q', c') by broadcasting (m, d) if $c(i) = d$ and $c' = c$. If those conditions are met we
104 write $(q, c) \xrightarrow{\text{br}(m,d)}_{\delta} (q', c')$ and call it a **broadcast step**.

105 A **local step** $(q, c) \xrightarrow{\text{op}(m,d)}_{\delta} (q', c')$ between two local configurations is either a re-
106 ception step or a broadcast step. A **local run** u of \mathcal{R} is a sequence of local steps $u =$
107 $(q_0, c_0) \xrightarrow{\text{op}_1(m_1,d_1)}_{\delta_1} (q_1, c_1) \xrightarrow{\text{op}_2(m_2,d_2)}_{\delta_2} \dots \xrightarrow{\text{op}_n(m_n,d_n)}_{\delta_n} (q_n, c_n)$. It is **initial** if (q_0, c_0) is
108 an initial configuration. In that case the common datum d of registers in c_0 is called is called
109 the **initial datum** of u . Its **input** $\mathbf{In}(u) \in (\mathcal{M} \times \mathbb{D})^*$ is the sequence of messages received by in-
110 put transitions $(q_{i-1}, c_{i-1}) \xrightarrow{\text{rec}(m_i,d_i)}_{\delta_i} (q_i, c_i)$ in u . Similarly, its **output** $\mathbf{Out}(u) \in (\mathcal{M} \times \mathbb{D})^*$
111 is the sequence of messages sent by output transitions $(q_{i-1}, c_{i-1}) \xrightarrow{\text{br}(m_i,d_i)}_{\delta_i} (q_i, c_i)$ in u .

112 Its **d -input** $\mathbf{In}_d(u) \in \mathcal{M}^*$ is the sequence of letters associated to datum d in $\mathbf{In}(u)$, and
113 its **d -output** $\mathbf{Out}_d(u) \in \mathcal{M}^*$ is the sequence of letters associated to datum d in $\mathbf{Out}(u)$.

114 ► **Remark 2.** Record transitions can only be taken if the received value is not already in the
115 registers. This is not a restriction on the expressivity: instead of storing the same datum in
116 several registers, the system use its registers to store each datum once, and use a function
117 $[1, r] \rightarrow [1, r]$, stored in the states, to assign registers to their content.

118 2.2 Broadcast Networks and Games with Registers

119 Let $r \in \mathbb{N}$ and let $\mathcal{R} = (Q, \mathcal{M}, q_{init}, \Delta)$ be a register transducer with r registers. The
120 **broadcast network of register automata** (BNRA for short) described by \mathcal{R} is the infinite

23:4 Controller Synthesis for Broadcast Networks with Data


121 transition system described below. We call register transducer *protocols* when we use them
 122 to define BNRA.

123 A *configuration* is a function $\gamma : \mathbb{A} \rightarrow Q \times \mathbb{D}^F$ with \mathbb{A} a finite set of *agents*. It maps each
 124 agent to a local configuration.

125 We write $\text{st}(\gamma)$ for the state component of γ and $\text{data}(\gamma)$ for its register component. A
 126 configuration γ is *initial* if for all $a \in \mathbb{A}$, $\text{st}(\gamma)(a) = q_{\text{init}}$, $\text{data}(\gamma)(a, i) = \text{data}(\gamma)(a, i')$ for all
 127 i, i' and $\text{data}(\gamma)(a, i) \neq \text{data}(\gamma)(a', i')$ for all $a \neq a'$ and i, i' . Intuitively, each agents starts
 128 with a unique identifier that is contained in all of its registers.

129 Given two configurations γ, γ' over \mathbb{A} , a *step* $\gamma \rightarrow \gamma'$ is defined when there exist $a_0 \in \mathbb{A}$,
 130 $m \in \mathcal{M}$, $d \in \mathbb{D}$ and a transition $\delta_0 \in \Delta_O$ such that $\gamma(a_0) \xrightarrow{\text{br}(m,d)}_{\delta_0} \gamma'(a_0)$, and for all $a \neq a_0$,
 131 either $\gamma'(a) = \gamma(a)$, or there is a transition $\delta \in \Delta_I$ such that $\gamma(a) \xrightarrow{\text{rec}(m,d)}_{\delta} \gamma'(a)$.

132 A (global) *run* ρ is a sequence of steps $\gamma_0 \rightarrow \gamma_1 \rightarrow \gamma_2 \cdots \gamma_{n-1} \rightarrow \gamma_n$. It is *initial* if γ_0 is
 133 an initial configuration. The *projection* of ρ on an agent a is the local run $\pi_a(\rho)$ made of all
 134 transitions taken by a in ρ . We write $\rho : \gamma \xrightarrow{*} \gamma'$ when ρ is a run from γ to γ' .



■ **Figure 1** A protocol which can do two things: On the left, it receives **start** and a sequence of a and b while checking that they carry the same datum (i.e. come from the same sender). At any point it may stop by sending **stop** with the received datum, to confirm that the communication has been received. In the right part, it broadcasts a **start** message with its initial identifier, then a sequence of messages a and b with that same identifier. When it receives it back with the letter **stop**, it terminates.

135 ► **Definition 3** (Coverability problem). The *coverability problem COVER* asks, given a protocol
 136 \mathcal{R} and a message m_{err} , whether there is an initial run in which m_{err} is broadcast.

137 If there is an initial run ρ in which an agent broadcasts m_{err} , then we say that ρ *covers*
 138 m_{err} , and that m_{err} is *coverable*.

139 ► **Remark 4.** The coverability problem is usually defined with an error state q_{err} : is there a
 140 an initial run where an agent reaches q_{err} ? We define it with a message as it will be more
 141 convenient for some definitions, and the two versions are easily inter-reducible.

142 We will examine the controller synthesis problem on this model. The COVER problem
 143 defined earlier is the particular case in which no state is controllable.

144 ► **Definition 5** (Broadcast Game with Registers). A *Broadcast Game with Registers* with r re-
 145 gisters $\mathcal{G} = (\mathcal{R}, Q_{\text{ctrl}}, Q_{\text{env}}, m_{\text{err}})$ is defined by a protocol with r registers $\mathcal{R} = (Q, \mathcal{M}, q_{\text{init}}, \Delta)$,
 146 a partition of its states $Q = Q_{\text{ctrl}} \sqcup Q_{\text{env}}$, and an error letter m_{err} .


147 A *control strategy* for \mathcal{G} is a function $\sigma : \Delta^* \rightarrow \Delta$ observing a sequence of transitions and
 148 choosing the next one ¹.

¹ We choose to not give access to the data to Controller, as we want to be able to rename data at will. We will discuss the version of the game where Controller can see the data in Section 6.

149 A σ -local run is an initial local run $u = (q_0, c_0) \xrightarrow{\text{op}_1(m_1, d_1)}_{\delta_1} \dots \xrightarrow{\text{op}_n(m_n, d_n)}_{\delta_n} (q_n, c_n)$
 150 such that for all $i \in [1, n]$, if $q_{i-1} \in Q_{\text{ctrl}}$ then $\sigma(\delta_1 \dots \delta_{i-1}) = \delta_i$. A σ -run is an initial run
 151 whose projection on every agent a is a σ -local run.

152 A control strategy is *winning* if no σ -run covers m_{err} . In this game and all games we
 153 construct from it the player trying to construct a winning control strategy will be called
 154 Controller and her opponent Environment.

155 ► **Definition 6** (Controller synthesis problem). *The safe strategy problem SAFESTRAT takes*
 156 *as input a BGR \mathcal{G} , and asks whether there is a winning control strategy for \mathcal{G} .*



157 ■ **Figure 2** A BGR. The round state C belongs to Controller, square states to Environment.

157 ► **Example 7.** In the BGR displayed in Figure 2, Controller has a single winning control
 158 strategy, which is to always choose a different letter from the one chosen by Environment
 159 from E . Indeed, doing otherwise would let an agent broadcast either aa or bb with its initial
 160 datum. In the first case, Environment could send an agent in the second row to receive aa
 161 and then broadcast m_{err} . In the second case, Environment could send two agents to the
 162 third row, who receive bb and broadcast a with the same datum. An agent in the second row
 163 could then receive both a broadcasts and then broadcast m_{err} . By contrast, it is easy to
 164 check that if Controller always picks a different letter from the one chosen by Environment
 165 in E , there cannot be two broadcasts of a or of b with the same datum. Hence agents sent
 166 to the second and third row will be unable to broadcast anything.

167 2.3 Subword order toolbox

168 A *well quasi-order* is a set equipped with a preorder relation (S, \preceq) such that in every infinite
 169 sequence s_0, s_1, \dots there exist $i < j$ such that $s_i \preceq s_j$.

170 Given two words $v = a_1 \dots a_m$ and $w = b_1 \dots b_n$ in Σ^* , we say that v is a *subword* of
 171 w and write $v \sqsubseteq w$ if v can be obtained from w by removing letters, i.e., there are indices
 172 $i_1 < \dots < i_m$ such that $v = b_{i_1} \dots b_{i_m}$.

173 Given a set of words W , we define its *upward-closure* $W \uparrow = \{u \in \Sigma^* \mid \exists w \in W, w \sqsubseteq u\}$
 174 and its *downward-closure* similarly $W \downarrow = \{u \in \Sigma^* \mid \exists w \in W, u \sqsubseteq w\}$. We say that W is
 175 *downward-closed* if $W = W \downarrow$, and *upward-closed* if $W = W \uparrow$. The set of minimal elements
 176 of an upward-closed set I is called its *basis*. Given a finite basis B , we define its *norm* as the
 177 maximum length of its words: $\|B\| = \max\{|w| \mid w \in B\}$.

178 A seminal result in the study of well quasi-orders is *Higman's lemma* [15, 14], which states
 179 that (Σ^*, \sqsubseteq) is a well quasi-order for all finite alphabet Σ . As a corollary, we obtain that
 180 every upward-closed set of words $I \subseteq \Sigma^*$ has a finite basis B such that $I = B \uparrow$. Another

181 corollary is that the upward-closure of any language over a finite alphabet is regular. This is
 182 a consequence of the previous property, along with the following fact.

183 ► **Lemma 8** (Folklore). *Given a finite set of words B over a finite alphabet Σ , one can*
 184 *construct a deterministic automaton $\mathcal{A}_{B\uparrow}$ recognising $B\uparrow$ with at most $(\|B\| + 1)^{|B|}$ states.*

185 In this work we will show that our main problem is $\mathbf{F}_{\omega^\omega}$ -complete, meaning that it is
 186 decidable but with a very high complexity, much higher than the Ackermann function for
 187 instance. For a formal definition of this complexity class $\mathbf{F}_{\omega^\omega}$ (and the related class of
 188 functions $\mathcal{F}_{\omega^\omega}$), see [22]. For our purpose, we will only use the Length function theorem,
 189 stated below.

190 A finite or infinite sequence of words w_0, w_1, \dots is *good* if there exist $i < j$ such that
 191 $w_i \sqsubseteq w_j$, and *bad* otherwise. Higman's lemma states that every bad sequence of words over
 192 a finite alphabet is finite, but we do not have a bound on their size. However, if we add a
 193 constraint so that each word can only have finitely many successors, then a uniform bound
 194 exists. Given a function $g : \mathbb{N} \rightarrow \mathbb{N}$ and an integer $n \in \mathbb{N}$, we say that a sequence of words
 195 w_1, \dots is *(g, n) -controlled* if $|w_i| \leq g^{(i)}(n)$ for all $i \geq 1$ (where $g^{(i)}$ denotes g applied i times).

196 ► **Theorem 9** (*Length Function Theorem* [23]). *Let Σ be a finite alphabet and $g : \mathbb{N} \rightarrow \mathbb{N}$ a*
 197 *primitive recursive function. There exists a function $f \in \mathcal{F}_{\omega^{|\Sigma|-1}}$ such that, for all $n \in \mathbb{N}$,*
 198 *every (g, n) -controlled bad sequence w_1, w_2, \dots has at most $f(n)$ terms.*

199 2.4 Games toolbox

200 We assume familiarity with automata and regular games, and simply fix some terms notations
 201 (see [11, Chapter 2] for an in-depth presentation). A *two-player game* \mathcal{G} is given by a directed
 202 graph $G = (V, E)$ called the *arena*, along with a partition of V in two, $V = V_0 \sqcup V_1$, an initial
 203 vertex v_{init} , a colouring function $c : V \rightarrow C$ mapping vertices to a finite set of colours C ,
 204 and a language $\mathcal{L} \subseteq C^\omega$ of infinite sequences of colours, called the *objective*. There are two
 205 players, called P_0 and P_1

206 A (finite or infinite path) in G starting in v_{init} is called a *play*. A play $v_0 \rightarrow v_1 \rightarrow \dots$ is
 207 *winning* for P_0 if $c(v_0)c(v_1)\dots \in \mathcal{L}$, and *losing* for P_0 otherwise.

208 A *strategy* for player P_i is a function $\sigma_{\mathcal{G}} : V^*V_i \rightarrow V$. A *$\sigma_{\mathcal{G}}$ -play* is a path $v_0 \rightarrow v_1 \rightarrow \dots$
 209 in G such that for all $j \geq 1$, if $v_{j-1} \in V_i$ then $v_j = \sigma_{\mathcal{G}}(v_0 \dots v_{j-1})$. A *strategy* for P_0 (resp.
 210 P_1) is *winning* if all infinite $\sigma_{\mathcal{G}}$ -plays are winning (resp. losing) for P_0 .

211 We call \mathcal{G} a *reachability game* (resp. *safety game*), when the objective is of the form LV^ω
 212 (resp. $V^\omega \setminus LV^\omega$) with L a regular language of finite words (represented by a deterministic
 213 finite automaton). It is well-known that those games are *determined*, i.e., in every game one
 214 of the two players has a *winning strategy*.

215 ► **Proposition 10** (Folklore). *One can compute the winner of a finite reachability game in*
 216 *polynomial time. Furthermore if P_0 has a winning strategy, then she has one that guarantees*
 217 *that she wins in at most $|V| \cdot |\mathcal{A}|$ steps.*

218 3 An introductory case: Broadcast networks without data

219 We start by showing the proof principles in an easy case, when processes do not have registers.
 220 In that case communication is made only through letters of \mathcal{M} . In this section we will forget
 221 the data in messages, and only consider letters. We simplify notations: we write $\mathbf{br}(m)$ for a
 222 broadcast of letter m and $\mathbf{rec}(m)$ for a reception of m . We obtain *Reconfigurable Broadcast*
 223 *Networks*, as introduced in [9]. From now on we will use the term RBN for this model.

224 COVER has been shown decidable and P-complete for those systems [9]. We could not
 225 find a result in the literature stating that SAFESTRAT is NP-complete, but closely-related
 226 results were proven in [5] and [24]. We prove it to illustrate our method.

227 To begin with, we show that we can characterise winning control strategies as the ones
 228 which force the set of letters sent to stay within some set $I \subseteq \mathcal{M} \setminus \{m_{err}\}$, called an *invariant*.

229 ► **Definition 11** (Invariants for RBN). *An invariant for an RBN over alphabet \mathcal{M} is a set of*
 230 *letters $I \subseteq \mathcal{M}$. We say that it is sufficient for a control strategy σ if:*

- 231 ■ $m_{err} \notin I$, and
- 232 ■ If a σ -local run receives only messages of I then it broadcasts only messages of I .

233 ► **Lemma 12** (Invariants characterise winning control strategies). *A control strategy σ is*
 234 *winning if and only if there exists a sufficient invariant $I \subseteq \mathcal{M}$ for it.*

235 **Proof sketch.** Suppose σ is winning. Let I be the set of messages such that there exists a
 236 σ -run in which they are broadcast. As σ is winning, $m_{err} \notin I$. If a σ -local run u receives
 237 only messages of I , then we can build a σ -run where an agent follows u : for each letter m
 238 received in u , we make other agents execute a σ -run where m is broadcast: this is possible
 239 as $m \in I$. We match this broadcast with the reception in u . We obtain a σ -run where all
 240 letters broadcast in u are broadcast.

241 For the other direction, suppose by contradiction that σ has a sufficient invariant I and
 242 that there is a σ -run ρ in which m_{err} is broadcast. Let m be the first message broadcast in
 243 ρ that is not in I , and a the agent broadcasting it. Those are well-defined as $m_{err} \notin I$. Let
 244 ρ' be the prefix of ρ stopping right after that broadcast. The projection $\pi_a(\rho')$ of ρ' on agent
 245 a contains a broadcast of m but no reception of any $m' \notin I$, a contradiction. ◀

246 This lets us turn the distributed game into a sequential one: If we are given an invariant
 247 I , checking whether there is a strategy that maintains it comes down to a two-player safety
 248 game. We obtain an algorithm for strategy synthesis: guess an invariant, and then solve the
 249 resulting safety game, which can be done in polynomial time.

250 ► **Theorem 13.** *Deciding the winner of a BGR without registers is NP-complete.*

251 **Proof.** For the upper bound, by Lemma 12, it suffices to guess a set $I \subseteq \Sigma$ such that $m_{err} \notin I$
 252 and then check if there is a strategy that guarantees that we can only broadcast a message
 253 outside of I if we received one beforehand. This is easily encoded into a safety game: Take
 254 the states and transitions of the BGR, without the operations, add a sink state with no
 255 outgoing transitions, and redirect every reception of a message $m \notin I$ to it. The objective of
 256 the first player is to avoid transitions broadcasting letters of $\Sigma \setminus I$.

257 Clearly there is a winning control strategy for the BGR if and only if there is an invariant
 258 I and a strategy avoiding transitions broadcasting messages outside of I in this safety game.
 259 This can be checked in polynomial time, by Proposition 10.

260 The lower bound is shown in Appendix A ◀

261 To conclude this section, we present an argument in favour of *parameterized* distributed
 262 synthesis. The fact that we have an arbitrary amount of agents makes the existence of a
 263 winning control strategy less likely. One might wonder what happens if we simply want
 264 a strategy that works for a bounded, or even fixed amount of agents. We show that the
 265 problem becomes undecidable in this case, even for 3 agents. Hence considering arbitrary
 266 numbers of agents can be a good approximation as it spectacularly reduces the difficulty of
 267 the problem.

268 ▶ **Theorem 14.** *Given a BGR $\mathcal{G} = (\mathcal{R}, Q_{ctrl}, Q_{env}, m_{err})$, it is undecidable whether there is a*
 269 *control strategy such that no σ -run with 3 agents covers m_{err} .*

270 4 Signature BGR

271 In this section we establish decidability of SAFESTRAT in a subcase of interest, that illustrates
 272 well the decidability proof for the general case, while requiring less technical complications.

273 ▶ **Definition 15.** *A signature protocol is one where every broadcast is made with the value*
 274 *of register 1, and all receptions are made on other registers. The associated BGR are called*
 275 *signature BGR.*

276 In other words, there are no transitions of the form $\xrightarrow{\text{br}(m,i)}$ with $i \geq 2$ or $\xrightarrow{\text{rec}(m,\downarrow 1)}$ or
 277 $\xrightarrow{\text{rec}(m,=1)}$. Such a protocol keeps its initial datum in register 1 and uses it for broadcasts,
 278 while the other registers are used to store and compare received values. An interesting
 279 property of those systems is that the datum of a message identifies its sender: Each agent
 280 only sends messages with its initial datum, and since those are unique, messages containing
 281 the same datum necessarily come from the same agent. In this section we will call *output* the
 282 d -output of a local run u with d its initial datum, and write it $\text{Out}_{sign}(u)$.

283 ▶ **Theorem 16.** *The SAFESTRAT problem is decidable $\mathbf{F}_{\omega\omega}$ -complete for signature BGR.*

284 The lower bound is provided by [13], as they show $\mathbf{F}_{\omega\omega}$ -hardness already for COVER
 285 with $r = 2$ registers. Fix $\mathcal{G} = (\mathcal{R}, Q_{ctrl}, Q_{env}, m_{err})$ a BGR with r registers. To prove the
 286 theorem, we once again use a characterisation of winning strategies in terms of invariants.
 287 Here an invariant is a downward-closed set of words of \mathcal{M}^* . A witness for non-coverability
 288 of a message m_{err} is a downward-closed set I of words that contains ε and not m_{err} and
 289 such that an agent whose d -inputs are all in I has an output in I^2 . Intuitively, if all agents
 290 respect that condition, then we can only obtain runs where all agents output words in the
 291 invariant, and thus no-one broadcasts m_{err} .

292 The downward-closed property comes from the fact that if an agent outputs a word w ,
 293 then other agents can receive any subsequence of letters of that word, as broadcasts can be
 294 lost. It is crucial as it gives us a finite representation of invariants, their basis.

295 ▶ **Definition 17** (Invariants for signature BGR). *An invariant for a signature BGR over an*
 296 *alphabet \mathcal{M} is a downward-closed set $I \subseteq \mathcal{M}^*$. We say that it is sufficient for a control*
 297 *strategy σ if it satisfies the following conditions:*

- 298 1. $\varepsilon \in I$ and $m_{err} \notin I$
- 299 2. For all σ -local run u , if $\text{In}_d(u) \in I$ for all $d \in \mathbb{D}$ then $\text{Out}_{sign}(u) \in I$.

300 The next step is to show that a winning strategy always comes with a sufficient invariant.

301 ▶ **Lemma 18** (Invariants characterise winning strategies). *A control strategy σ is winning if*
 302 *and only if there exists a sufficient invariant $I \subseteq \mathcal{M}^*$ for it.*

303 Recall that, as a corollary of Higman's lemma, upward-closed sets of words can be finitely
 304 described by their finite basis.

305 To solve SAFESTRAT, we cannot enumerate potential strategies as there are uncountably
 306 many. Instead, our algorithm enumerates invariants (represented by the basis of their

² As I is downward-closed, $\varepsilon \in I$ is synonymous with I being non-empty.

307 complement) and checks for each one whether there is a strategy such that the conditions
 308 listed in Lemma 18 are satisfied. While the first item is straightforward to check, the second
 309 is not. To verify it, we design a game in which the two players construct a local run, and the
 310 received data are chosen by Environment.

311 4.1 Invariant game for signature BGR

312 Intuitively, the two players construct a local run by picking the transitions from their
 313 respective states, and Environment picks the data received at each step, when they are not
 314 already determined by the chosen transition. If at some point the d -input gets out of I for
 315 some d then the game stops and Controller wins. If the output gets out of I then the game
 316 stops and Environment wins. If none of the two happen and the game goes on forever then
 317 Controller wins. This characterises the capacity of Controller to keep outputs within a given
 318 invariant, but if we made the choice of data explicit this game would be infinite.

319 We reduce it to a finite reachability game, called the invariant game which we can solve
 320 by a simple fix-point computation.

321 A first observation is that it is always in Environment's best interest to choose fresh data
 322 that were never seen before, as they come with the smallest d -input. Thus whenever we
 323 receive a datum that is not in the registers we can assume that the associated d -input is
 324 empty. This means that we do not need to remember the d -inputs associated to every datum
 325 of the local run, but only those that are currently in the registers.

326 To formalise this, we need to define the inputs and output of sequences of transitions.
 327 The idea is that we can assume that every datum that disappears from the registers will
 328 never appear again. In order to check whether some d -input gets out of I , we only need to
 329 keep track of the sequences of letters received with the data currently in the registers. We
 330 call those the recent inputs. Furthermore, in our model of register transducers, a received
 331 datum always appears in at most one register at a time, and while it is not forgotten, it
 332 stays in that one register. This will let us read the recent inputs directly from the sequence
 333 of transitions.

334 Given a sequence of transitions $\delta_1 \cdots \delta_k$ of \mathcal{R} , we define its *output* as the sequence of
 335 letters sent by broadcasts. For all registers $i \in [1, r]$, we also define its *recent input on i* as
 336 the sequence of letters received with an equality transition with register i since it was last
 337 updated. Formally, the output of $\delta_1 \cdots \delta_k$ is defined inductively as $\text{Out}(\varepsilon) = \varepsilon$ and

$$338 \quad \text{Out}(\delta_1 \cdots \delta_{k+1}) = \begin{cases} \text{Out}(\delta_1 \cdots \delta_k) & \text{if } \delta_{k+1} \text{ is a reception,} \\ \text{Out}(\delta_1 \cdots \delta_k)m & \text{if } \delta_{k+1} = \xrightarrow{\text{br}(m,1)} \text{ for some } m. \end{cases}$$

339 The recent input on i is defined as $\text{recentIn}_i(\varepsilon) = \varepsilon$ and:

$$340 \quad \text{recentIn}_i(\delta_1 \cdots \delta_{k+1}) = \begin{cases} m & \text{if } \delta_{k+1} = \xrightarrow{\text{rec}(m, \downarrow i)} \text{ for some } m \text{ and } i, \\ \text{recentIn}_i(\delta_1 \cdots \delta_k)m & \text{if } \delta_{k+1} = \xrightarrow{\text{rec}(m, =i)} \text{ for some } m \text{ and } i, \\ \text{recentIn}_i(\delta_1 \cdots \delta_k) & \text{otherwise.} \end{cases}$$

341 Note that we always have $\text{recentIn}_1(\delta_1 \cdots \delta_k) = \varepsilon$, as we assumed that no reception is
 342 made using register 1.

343 The *invariant game* $\mathcal{IG}(\mathcal{G}, I)$ goes as follows. The set of vertices is simply $Q_{\mathcal{R}}$. From
 344 each vertex $q \in Q_{\mathcal{R}}$, players choose a transition from q in $\Delta_{\mathcal{R}}$. Controller chooses the next
 345 transition when the current vertex is in Q_{ctrl} , Environment when it is in Q_{env} .

23:10 Controller Synthesis for Broadcast Networks with Data

- 346 ■ If at some point the play $\pi = \delta_1 \cdots \delta_k$ is such that $\text{recentIn}_i(\pi) \notin I$ for some $i \geq 2$ then
 347 Controller wins.
- 348 ■ If at some point the play $\pi = \delta_1 \cdots \delta_k$ is such that $\text{Out}(\pi) \notin I$ then Environment wins.
- 349 ■ If the play goes on forever without any of those things happening then Controller wins.

350 We start by showing that we can solve this game by considering it as a regular safety
 351 game. We obtain as a corollary that if Environment wins then he can win in a bounded
 352 number of steps. Define $\varphi(\mathcal{R}, B) := |\mathcal{R}|(|B| + 1)^{|\mathcal{R}|(|B|+1)}$

353 ► **Lemma 19** (Decidability of the invariant game). *Given a BGR over protocol \mathcal{R} and a finite
 354 set of words B , we can decide in exponential time whether Controller has a winning strategy
 355 in $\mathcal{IG}(\mathcal{G}, (B^\uparrow)^\complement)$. Furthermore, if Environment has a winning strategy then he has a strategy
 356 to win in at most $\varphi(\mathcal{R}, B)$ steps.*

357 **Proof.** By Lemma 8, B^\uparrow is a regular language, recognised by a deterministic finite automaton
 358 $\mathcal{A}_{B^\uparrow} = (Q_B, \mathcal{M}, \Delta_B, q_0^B, F_B)$ with $(|B| + 1)^{|B|+1}$ states.

359 We can construct a deterministic automaton \mathcal{B} over the alphabet $\Delta_{\mathcal{R}}$ that reads plays
 360 $\delta_1 \cdots \delta_k$ of $\mathcal{IG}(\mathcal{G}, (B^\uparrow)^\complement)$ and accepts exactly the winning plays for Environment. Its set of
 361 states is $(Q_B)^r$, plus a rejecting sink state \perp and an accepting sink state \top , which is the
 362 only accepting state. The first component keeps track of the state reached in \mathcal{A}_{B^\uparrow} by the
 363 output of the sequence of transitions. The others keep track, for each register $i \geq 2$, of the
 364 state reached by the recent input on i in \mathcal{A}_{B^\uparrow} . Transitions of that automaton are easy to
 365 infer from the definition of output and recent input on i . The automaton goes to \perp if the
 366 recent input on i is in B^\uparrow for some i , or if it sees a reception transition of a message $m \in B^\uparrow$.
 367 It goes to \top if the output is in B^\uparrow .

368 By Proposition 10, we can solve this game in polynomial time in the size of the automaton
 369 \mathcal{B} and the size of the arena of $\mathcal{IG}(\mathcal{G}, (B^\uparrow)^\complement)$ (i.e., $|\mathcal{R}|$), that is, in exponential time in
 370 $|B| + |B| + |\mathcal{R}|$. Furthermore, if Environment has a winning strategy then he has one that
 371 guarantees that he wins in at most $\varphi(\mathcal{R}, B) = |\mathcal{A}_{B^\uparrow}|^r |\mathcal{R}|$ steps. ◀

372 We have two things to prove: First that a winning strategy for Controller in $\mathcal{IG}(\mathcal{G}, I)$
 373 yields a control strategy σ for which I is a sufficient invariant. Then, that a winning strategy
 374 for Environment in $\mathcal{IG}(\mathcal{G}, I)$ implies that there is no control strategy for which I is a sufficient
 375 invariant.

376 ► **Lemma 20.** *Let $I \subseteq \mathcal{M}^*$ be a downward-closed set of words containing ε and not m_{err} . If
 377 Controller wins the invariant game $\mathcal{IG}(\mathcal{G}, I)$ then there is a control strategy σ such that I is
 378 a sufficient invariant for σ .*

379 ► **Lemma 21.** *Let σ be a control strategy. Let $I \subseteq \mathcal{M}^*$ be a downward-closed set of words
 380 containing ε and not m_{err} , and let B be the basis of I^\complement .
 381 If Environment wins the invariant game $\mathcal{IG}(\mathcal{G}, I)$ then there is a σ -local run of length at
 382 most $\varphi(\mathcal{R}, B)$ with an output not in I and all d -inputs in I .*

383 **Proof.** By Proposition 10 there exists $\tau_{\mathcal{IG}}$ a winning strategy $\tau_{\mathcal{IG}}$ for Environment in the
 384 invariant game $\mathcal{IG}(\mathcal{G}, I)$ such that Environment always wins in at most $\varphi(\mathcal{R}, B)$ steps.

385 We construct a σ -local run of length at most $\varphi(\mathcal{R}, B)$ with an output not in I and all
 386 d -inputs in I . To do so, we apply $\tau_{\mathcal{IG}}$ to choose transitions and we choose data by always
 387 picking a datum never seen before in the run, when the datum is not determined by the
 388 transition.

389 Let (s_0, c_0) be an initial configuration of \mathcal{R} . We define iteratively a sequence of steps
 390 $(s_{k-1}, c_{k-1}) \xrightarrow{\text{op}_k(m_k, d_k)}_{\delta_k} (s_k, c_k)$ as follows. Suppose we defined them up to (s_{k-1}, c_{k-1}) ,

391 and let u_{k-1} be the local run defined so far. We first choose δ_k : If $s_{k-1} \in Q_{\text{ctrl}}$ then
 392 $\delta_k = \sigma(u_{k-1})$, otherwise $\delta_k = \tau_{\mathcal{IG}}(\delta_1 \cdots \delta_{k-1})$.

393 We then choose d_k :

- 394 ■ If δ_k is a broadcast transition of letter m , we set $d_k = c_k(1)$ (the initial datum of the
 395 local run).
- 396 ■ If δ_k is a record transition, we pick a datum d_k that does not appear in u_{k-1} before.
- 397 ■ If $\delta_k = s_{k-1} \xrightarrow{\text{rec}(m,=i)} s_k$ is an equality transition of letter m , we set $d_k = c_{k-1}(i)$.

398 Clearly we maintain the fact that u_k is a σ -local run and $\delta_1 \cdots \delta_k$ is a $\tau_{\mathcal{IG}}$ -play in $\mathcal{IG}(\mathcal{G}, I)$.
 399 We stop when $\delta_1 \cdots \delta_k$ is winning for Environment in $\mathcal{IG}(\mathcal{G}, I)$, which happens for some
 400 $K \leq \varphi(\mathcal{R}, B)$. Let $u = u_K$ be the local run obtained at the end.

401 It remains to show that the output of u is not in I while all its d -inputs are in I . To do
 402 so, we rely on the following claim:

403 ▷ **Claim 22.** For all register i and index k , $\text{recentIn}_i(\delta_1 \cdots \delta_k) = \mathbf{In}_{u_k}(c_k(i))$. Furthermore,
 404 $\text{Out}(\delta_1 \cdots \delta_k) = \mathbf{Out}_{\text{sign}}(u_k)$

405 **Proof.** By a straightforward induction on k . ◀

406 By definition $\delta_1 \cdots \delta_K$ is a winning $\tau_{\mathcal{IG}}$ -play for Environment, hence its output is not
 407 in I , thus $\mathbf{Out}_{\text{sign}}(u) = \mathbf{Out}_{\text{sign}}(u_K)$ is not in I either. Let $d \in \mathbb{D}$ a datum appearing in
 408 u , and let k be such that (s_k, c_k) is the last configuration in which d appears. Let i be the
 409 register such that $c_k(i) = d$. Then we have $\mathbf{In}_u(d) = \mathbf{In}_{u_k}(c_k(i)) = \text{recentIn}_i(\delta_1 \cdots \delta_k)$. As
 410 $\tau_{\mathcal{IG}}$ is winning for Environment, $\text{recentIn}_i(\delta_1 \cdots \delta_k) \in I$, and thus $\mathbf{In}_u(d) \in I$.

411 We have found a σ -local run of length at most $\varphi(\mathcal{R}, B)$ whose output is not in I while all
 412 its d -inputs are. ◀

413 Our next step is to bound the minimal size of a sufficient invariant for some winning
 414 control strategy σ when there is one. The idea is as follows: Take an invariant I such
 415 that the basis $\{w_1, \dots, w_k\}$ of I^c has as few elements as possible. We can assume that
 416 $|w_1| \leq \dots \leq |w_k|$. Then we know that, for all i , $\{w_1, \dots, w_i\}$ is not a sufficient invariant
 417 for σ . Hence by Lemma 21 we get a σ -local run of bounded size breaking the invariant
 418 $\{w_1, \dots, w_i\}$, which forces $\{w_{i+1}, \dots, w_k\}$ to contain a word of bounded size. This bounds
 419 the size of w_{i+1} with respect to w_1, \dots, w_i , as stated in the lemma below.

420 Define $\psi(n) = |\mathcal{R}|(n+1)^{|\mathcal{M}|^{n+1}+1}$

421 ► **Lemma 23** (Bounding the size of the invariant). *Let \mathcal{G} a signature BGR. There is a winning
 422 control strategy for \mathcal{G} if and only if there is a sequence of words $w_0, \dots, w_k \in \mathcal{M}^*$ such that*

- 423 ■ *Controller wins $\mathcal{IG}(\mathcal{G}, \{w_1, \dots, w_k\}^{\uparrow c})$,*
- 424 ■ *and for all $i \in [1, k]$, $|w_i| \leq \psi(|w_{i-1}|)$.*

425 We will now leverage the Length Function Theorem to bound the size of the basis of
 426 I^c in Lemma 23.

427 ► **Theorem 24.** *SAFESTRAT is decidable and in \mathbf{F}_{ω} for signature BGR.*

428 **Proof.** Let \mathcal{G} a BGR. We apply the Length Function Theorem with $\Sigma = \mathcal{M}$ and $g(n) =$
 429 $n(n+1)^{n^{n+1}+1}$. We obtain a function $f \in \mathcal{F}_{\omega, |\mathcal{M}|-1}$ such that every (g, n) -controlled bad
 430 sequence of words w_0, w_1, \dots, w_k has at most $f(n)$ terms.

431 We use a non-deterministic algorithm that guesses a sequence of words w_1, \dots, w_k such
 432 that $w_1 = m_{\text{err}}$ and $|w_i| \leq |w_{i+1}| \leq \psi(|w_i|)$ for all i . One can straightforwardly check that
 433 then we have $|w_i| \leq g^{(i)}(|\mathcal{R}| + |\mathcal{M}| + 1)$ for all i .

23:12 Controller Synthesis for Broadcast Networks with Data

434 Let $B = \{w_0, w_1, \dots, w_k\}$. The algorithm checks that there exists a strategy σ such that
 435 the complement of $\{w_0, w_1, \dots, w_k\}^\uparrow$ is a sufficient invariant for σ , by solving the invariant
 436 game $\mathcal{IG}(\mathcal{G}, (\{w_0, w_1, \dots, w_k\}^\uparrow)^\complement)$. This can be done in exponential time in $|\mathcal{R}| + k + |w_k|$, by
 437 Lemma 19. We accept if there is such a strategy and reject otherwise.

438 By Lemma 23, this algorithm is correct. We can make it deterministic with an exponential
 439 blow-up in the time complexity. The time required by this algorithm is therefore $h(f(|\mathcal{R}| +$
 440 $|\mathcal{M}| + 1))$ with h a primitive recursive function. As $\mathcal{F}_{\omega, |\mathcal{M}|-1}$ is closed under composition with
 441 primitive recursive functions, the algorithm takes a time bounded by a function of $\mathcal{F}_{\omega, |\mathcal{M}|-1}$.

442 As a consequence, the problem is in $\mathbf{F}_{\omega^\omega}$ (see [22] for details). ◀

5 General case

444 In this section we generalise the previous result to all BGR.

445 ▶ **Theorem 25.** *The SAFESTRAT problem is decidable in $\mathbf{F}_{\omega^\omega}$ for general BGR.*

446 We fix a BGR $\mathcal{G} = (\mathcal{R}, Q_{ctrl}, Q_{env}, m_{err})$ for the rest of this section.

447 The general structure of the proof is the same as before, but the removal of the signature
 448 hypothesis makes it significantly more technical. The main difference between the signature
 449 and general models is that in the latter a process can send acknowledgements to a process it
 450 received messages from, as in the right protocol in Figure 1.

451 We make the following observation: Say an agent receives a message (m, d) with d its
 452 initial datum; this is possible in general BGR but not in signature ones. Then this means
 453 that other agents, which did not have this datum initially, received enough messages with
 454 datum d to be able to broadcast (m, d) . Intuitively, we can copy these agents many times,
 455 which allows us to assume that we have an unlimited supply of messages (m, d) . In sum, we
 456 will show that if an agent sends a message (m, d) with d that is not its initial datum, then
 457 from this point on we can assume that messages (m, d) are for free. This intuition justifies
 458 the definition of decomposition, which summarises the sequence of letters sent with a given
 459 datum during a run. It details the sequence of letters sent by the agent with that datum
 460 initially, and the points at which each letter is first broadcast with that datum by another
 461 agent. These decompositions were already used for the verification of those systems [13].

462 ▶ **Definition 26.** *A decomposition is a tuple $\text{dec} = (v_0, m_1, \dots, v_{k-1}, m_k, v_k)$ with m_0, \dots, m_k
 463 distinct letters of \mathcal{M} and $v_i \in \mathcal{M}^*$ for all i .*

464 *A word $w \in \mathcal{M}^*$ matches dec if $w = w_0 \dots w_k$ where each w_i can be obtained by inserting
 465 letters from $\{m_1, \dots, m_i\}$ in v_i .*

466 ▶ **Example 27.** Let $\mathcal{M} = \{a, b, c\}$. Then $\text{dec} = (\text{abba}, a, \text{cbc}, b, \text{cc})$ is a decomposition. The
 467 word abbacabaacbabcbca matches dec as we can cut in in three parts abbacabaacbabcbca , and
 468 cabaac can be obtained by adding some a to cbc and babcbca can be obtained by adding
 469 some a and b to cc .

470 We write \mathcal{L}_{dec} for the language of words that match dec. Given a family of upward-closed
 471 sets of words $(J_m)_{m \in \mathcal{M}}$, we define $\mathcal{D}((J_m)_{m \in \mathcal{M}})$ as the set of decompositions

$$472 \quad \mathcal{D}((J_m)_{m \in \mathcal{M}}) = \{(v_0, m_1, \dots, v_{k-1}, m_k, v_k) \mid \forall i, \mathcal{L}_{(v_0, m_1, \dots, v_{i-1})} \cap J_{m_i} \neq \emptyset\}.$$

473 With an additional downward-closed set I , we also define

$$474 \quad \mathcal{D}(I, (J_m)_{m \in \mathcal{M}}) = \{(v_0, m_1, \dots, v_{k-1}, m_k, v_k) \mid v_0 \dots v_k \in I, \forall i, \mathcal{L}_{(v_0, m_1, \dots, v_{i-1})} \cap J_{m_i} \neq \emptyset\}.$$

475 Finally, the set of words producible by $I, (J_m)_{m \in \mathcal{M}}$ is

$$476 \quad \mathcal{L}(I, (J_m)_{m \in \mathcal{M}}) = \bigcup_{\text{dec} \in \mathcal{D}(I, (J_m)_{m \in \mathcal{M}})} \mathcal{L}_{\text{dec}}.$$

477 We say that a local run u with initial datum d is *compatible* with a decomposition $\text{dec} =$
 478 $(v_0, m_1, \dots, v_{k-1}, m_k, v_k)$ if $u = u_0 \cdots u_k$ where $v_i \sqsubseteq \text{Out}_d(u_i)$ and $\text{In}_d(u_i) \in \{m_1, \dots, m_i\}^*$
 479 for all i .

480 Here I should be thought of as the set of words over \mathcal{M} that can be broadcast by an
 481 agent with its initial datum. Meanwhile, J_m represents the set of words w over \mathcal{M} such that
 482 an agent can broadcast (m, d) with d not its initial datum while having received before only
 483 (a subword of) w with that datum. It can be read as the “cost” of a message m : in order to
 484 receive a message (m, d) you should first broadcast a sequence of letters of J_m with datum d .

485 A decomposition (v_0, m_1, \dots, v_k) is a scenario of the sequence of letters broadcast over
 486 a datum d during a run: The agent who has d as initial datum broadcasts $v_0 \cdots v_k$ with
 487 it, while m_1, \dots, m_k mark the points at which each of those letters is first broadcast with
 488 datum d by another agent.

489 Then, we can see $\mathcal{D}(I, (J_m)_{m \in \mathcal{M}})$ as the set of decompositions (v_0, m_1, \dots, v_k) that are
 490 compatible with the invariant $I, (J_m)_{m \in \mathcal{M}}$. The condition $v_0 \cdots v_k \in I$ means that an agent
 491 with d as initial datum should be able to broadcast $v_0 \cdots v_k$ with it. The other condition
 492 says that for all i there is a word $w \in \mathcal{L}_{(v_0, m_1, \dots, v_{i-1})} \cap J_{m_i}$. This should be read as follows:

- 493 ■ $w \in J_{m_i}$ means that if we can broadcast the sequence w with datum d , we can make an
 494 agent broadcast (m_i, d)
- 495 ■ $w \in \mathcal{L}_{(v_0, m_1, \dots, v_{i-1})}$ means that we can broadcast the sequence w with datum d , as we
 496 can obtain it from $v_0 \cdots v_{i-1}$ by adding enough m_1, \dots, m_{i-1} .

497 5.1 Characterisation of winning strategies with invariants

498 An *invariant* for general BGR is made of a downward-closed set of words $I \subseteq \mathcal{M}^*$ (the
 499 sequences of letters that may be produced over some datum) and an upward-closed set of
 500 words $J_m \subseteq \mathcal{M}^*$ for each letter m (the sequences of letters that allow an agent to send a
 501 message (m, d) with d that is not its initial datum).

502 ► **Definition 28** (Invariants for BGR). *An invariant for general BGR is a pair $(I, (J_m)_{m \in \mathcal{M}})$*
 503 *with $I \subseteq \mathcal{M}^*$ a downward-closed set of words and, for all m , $J_m \subseteq \mathcal{M}^*$ an upward-closed set*
 504 *of words. We say that it is sufficient for a control strategy σ if the following conditions hold.*

- 505 1. $\varepsilon \in I$, $m_{\text{err}} \notin I$ and $J_{m_{\text{err}}} \cap I = \emptyset$
- 506 2. $\mathcal{L}(I, (J_m)_{m \in \mathcal{M}}) \subseteq I$
- 507 3. For all initial σ -local run u with initial datum d , if:
 - 508 (i) u is compatible with a decomposition $\text{dec} \in \mathcal{D}((J_m)_{m \in \mathcal{M}})$, and
 - 509 (ii) for all $d' \neq d$, $\text{In}_{d'}(u) \in I$,
 then we have that
 - 510 (a) $\text{Out}_d(u) \in I$
 - 511 (b) for all $m \in \mathcal{M}$ and $d' \neq d$, if u contains a broadcast of (m, d') then $\text{In}_{d'}(u) \in J_m$.

513 We once again prove that every winning control strategy has a sufficient invariant. The
 514 proof is presented in Appendix C

515 ► **Lemma 29** (Invariants characterise winning strategies). *A control strategy σ is winning if*
 516 *and only if there exists a sufficient invariant $(I, (J_m)_{m \in \mathcal{M}})$ for it.*

517 **5.2 The invariant game**

518 We have characterised winning control strategies using invariants. The next step is to consider
 519 an invariant $(I, (J_m)_{m \in \mathcal{M}})$ and show that we can construct a game in which the two players
 520 determine whether there is a control strategy for which this invariant is sufficient.

521 We proceed as in Section 4. First we consider a game played on \mathcal{R} where players pick a
 522 sequence of transitions (the next transition is chosen by the player owning the current state),
 523 and Environment picks the data when needed. The goal of Environment is to eventually
 524 obtain a local run u that satisfies either 3a or 3b but neither 3i nor 3ii, i.e., contradicting the
 525 invariant. The goal of Controller is to avoid this forever.

526 We define formally the invariant game $\mathcal{IG}(\mathcal{G}, I, (J_m)_{m \in \mathcal{M}})$ for general BGR in Appendix C.
 527 We show that Controller wins that game if and only if she has a control strategy in \mathcal{G} for
 528 which $(I, (J_m)_{m \in \mathcal{M}})$ is sufficient. Furthermore, we show that when Environment wins we
 529 can obtain a local run contradicting the invariant of bounded size with respect to \mathcal{G} , I and
 530 $(J_m)_{m \in \mathcal{M}}$. This lets us bound the size of a sufficient invariant when it exists, using the
 531 Length Function Theorem.

532 ► **Theorem 30** (Main theorem). *SAFESTRAT is decidable and $\mathbf{F}_{\omega\omega}$ -complete.*

533 **6 Allowing agents to see data**

534 So far we only considered control strategies that chose transitions based on the previous
 535 sequence of transitions, and not the sequence of data received. It is natural to wonder what
 536 happens if we use strategies of the form $\sigma : (\Delta \mathbb{D}^r)^* \rightarrow \Delta$. For this section we will only
 537 consider the signature case to make things easier. We conjecture that the following proof
 538 can be adapted to the general case.

539 A central ingredient in this proof is Ramsey’s theorem on infinite hypergraphs, which
 540 extends naturally Ramsey’s theorem on graphs [20]. It states that if we colour every subset
 541 of size k of an infinite set while using finitely many colours, then there is an infinite subset
 542 in which every k -subset has the same colour.

543 We now define *data-aware control strategies*. They are functions $\sigma : \mathbb{D}(\Delta \times \mathbb{D})^* \rightarrow \Delta$.
 544 The next transition is chosen based on the local run taken so far, including the initial datum
 545 and the data received. Notions of σ -local runs and σ -runs are extended naturally.

546 ► **Theorem 31.** *There is a winning data-aware control strategy for \mathcal{G} if and only if there is
 547 a winning control strategy for \mathcal{G} .*

548 **Proof sketch.** We show that there is a function h such that whenever a strategy is losing
 549 there is a losing σ -run of a certain shape where each agent has a local run of length at most
 550 $h(|\mathcal{G}|)$. Assume we have a winning data-aware control strategy. We then colour every set
 551 of $h(|\mathcal{G}|)$ data according to the behaviour of that strategy on local runs of length $\leq h(|\mathcal{G}|)$
 552 where those data appear. We apply Ramsey’s theorem to obtain an infinite set of data on
 553 which the strategy behaves the same on “short” local runs. This defines a control strategy
 554 which does not fail on runs where local runs are of length $\leq h(|\mathcal{G}|)$. By definition of h , the
 555 resulting strategy is winning. ◀

556 By combining this with Theorem 16, we obtain the following result.

557 ► **Corollary 32.** *The existence of a winning data-aware control strategy for a BGR is decidable
 558 and $\mathbf{F}_{\omega\omega}$ -complete.*

7 The case of one register

The construction of [13] for the \mathbf{F}_{ω} lower bound only requires two registers. We studied in Section 3 the complexity of those problems when protocols do not have registers. The remaining gap is for protocols with one register. We call them *1BGR*. They offer an intermediate step in terms of tractability and expressivity between the protocols without registers and the general case. Those protocols can sign messages with their initial identifier, and check that several messages have the same datum, but not simultaneously. They relate to Petri nets and population protocols with data, as those only allow each process to store one datum. In particular, the subclass of *IO population protocols with data* can be seen as a particular case of BGR with one register.

We investigate the complexity of SAFESTRAT for 1BGR. In this case, a record transition essentially resets the memory of the process. This lets us split the invariant game used in the general case into simpler games: the output game and the echo games.

► **Theorem 33.** *SAFESTRAT is NEXPTIME-complete for 1BGR.*

A strategy $\sigma : V^* \rightarrow V$ for a two-player game is *positional* if its output only depends on the current state, that is, for all $w, w' \in V^*$ and $v \in V$ we have $\sigma(wv) = \sigma(w'v)$. We rely on the following criterion, which can be used to show that a player can win with a positional strategy. A language \mathcal{L} is *submixing* (or *concave*) if whenever we have words u_0, u_1, \dots and v_0, v_1, \dots such that $u_0u_1 \dots \notin \mathcal{L}$ and $v_0v_1 \dots \notin \mathcal{L}$ then $u_0v_0u_1v_1 \dots \notin \mathcal{L}$. It was shown in [17] that if an objective is submixing then player P_0 has a positional optimal strategy in all games with this objective.

We rely on the characterisation of winning control strategies by the invariant game, as stated in Lemma 45 and 48. It turns out that for 1BGR, the invariant game can be split into several simpler games. Essentially, we consider the recording of a new value in the register as a reset of the game. We define two different games: in the output game the players build the part of the local run before the first record transition. In the echo game the players build an interval of the local run between two record transitions.

We show that in the first game Controller can always use a positional strategy (Lemma 57) while in the second one it is Environment who can stick to positional strategies (Lemma 58). In both cases we use the submixing property of their objectives to prove it.

We also prove that the winners of those games determine the winner of the 1BGR (Lemma 59). The positionality of Environment's strategy in the echo game then lets us bound the size of the invariants necessary to witness the existence of a winning control strategy for Controller (Lemma 62). We exhibit an NEXPTIME algorithm, in which the non-deterministic guess is the invariant and a positional strategy for Controller in the output game. The lower bound follows from a reduction from the exponential grid tiling problem.

8 Conclusion

We showed decidability of SAFESTRAT for a powerful parameterised distributed model. We showcased a method for distributed controller synthesis through invariants by using it for increasingly complex versions of the model. We also match every resulting complexity class with a lower bound, which tends to show that this method makes sense for this model. The most promising future direction is to develop invariants for other models of distributed systems in order to obtain more decidability results. We can also investigate the relation between other distributed models with data and BGR, especially 1BGR.

603 — References

- 604 1 Georg Bachmeier, Michael Luttenberger, and Maximilian Schlund. Finite automata for the sub-
605 and superword closure of cfls: Descriptive and computational complexity. In Adrian-Horia
606 Dediu, Enrico Formenti, Carlos Martín-Vide, and Bianca Truthe, editors, *Language and
607 Automata Theory and Applications - 9th International Conference, LATA 2015, Nice, France,
608 March 2-6, 2015, Proceedings*, volume 8977 of *Lecture Notes in Computer Science*, pages
609 473–485. Springer, 2015. doi:10.1007/978-3-319-15579-1_37.
- 610 2 A. R. Balasubramanian and Chana Weil-Kennedy. Reconfigurable broadcast networks and
611 asynchronous shared-memory systems are equivalent. In Pierre Ganty and Davide Bresolin,
612 editors, *Proceedings 12th International Symposium on Games, Automata, Logics, and Formal
613 Verification, GandALF 2021, Padua, Italy, 20-22 September 2021*, volume 346 of *EPTCS*,
614 pages 18–34, 2021. doi:10.4204/EPTCS.346.2.
- 615 3 Nathalie Bertrand, Patricia Bouyer, and Anirban Majumdar. Concurrent parameterized
616 games. In Arkadev Chattopadhyay and Paul Gastin, editors, *39th IARCS Annual Conference
617 on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2019,
618 December 11-13, 2019, Bombay, India*, volume 150 of *LIPICs*, pages 31:1–31:15. Schloss
619 Dagstuhl - Leibniz-Zentrum für Informatik, 2019. URL: <https://doi.org/10.4230/LIPICs.FSTTCS.2019.31>, doi:10.4230/LIPICs.FSTTCS.2019.31.
- 620 4 Nathalie Bertrand, Miheer Dewaskar, Blaise Genest, Hugo Gimbert, and Adwait Amit Godbole.
621 Controlling a population. *Log. Methods Comput. Sci.*, 15(3), 2019. doi:10.23638/LMCS-15(3:
622 6)2019.
- 623 5 Nathalie Bertrand, Paulin Fournier, and Arnaud Sangnier. Distributed local strategies in broad-
624 cast networks. In Luca Aceto and David de Frutos-Escrig, editors, *26th International Conference
625 on Concurrency Theory, CONCUR 2015, Madrid, Spain, September 1-4, 2015*, volume 42 of
626 *LIPICs*, pages 44–57. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015. URL: <https://doi.org/10.4230/LIPICs.CONCUR.2015.44>, doi:10.4230/LIPICs.CONCUR.2015.44.
- 627 6 Raven Beutner, Bernd Finkbeiner, and Jesko Hecking-Harbusch. Translating asynchronous
628 games for distributed synthesis. In Wan J. Fokkink and Rob van Glabbeek, editors, *30th
629 International Conference on Concurrency Theory, CONCUR 2019, August 27-30, 2019,
630 Amsterdam, the Netherlands*, volume 140 of *LIPICs*, pages 26:1–26:16. Schloss Dagstuhl
631 - Leibniz-Zentrum für Informatik, 2019. URL: <https://doi.org/10.4230/LIPICs.CONCUR.2019.26>, doi:10.4230/LIPICs.CONCUR.2019.26.
- 632 7 Michael Blondin and François Ladouceur. Population protocols with unordered data. In
633 Kousha Etessami, Uriel Feige, and Gabriele Puppis, editors, *50th International Colloquium
634 on Automata, Languages, and Programming, ICALP 2023, July 10-14, 2023, Paderborn,
635 Germany*, volume 261 of *LIPICs*, pages 115:1–115:20. Schloss Dagstuhl - Leibniz-Zentrum für
636 Informatik, 2023. URL: <https://doi.org/10.4230/LIPICs.ICALP.2023.115>, doi:10.4230/
637 LIPICs.ICALP.2023.115.
- 638 8 Giorgio Delzanno, Arnaud Sangnier, and Riccardo Traverso. Parameterized verification
639 of broadcast networks of register automata. In *Reachability Problems, RP 2013*, volume
640 8169 of *Lecture Notes in Computer Science*, pages 109–121. Springer, 2013. doi:10.1007/
641 978-3-642-41036-9_11.
- 642 9 Giorgio Delzanno, Arnaud Sangnier, and Gianluigi Zavattaro. Parameterized verification of ad
643 hoc networks. In *CONCUR 2010*, volume 6269 of *Lecture Notes in Computer Science*, pages
644 313–327. Springer, 2010. doi:10.1007/978-3-642-15375-4_22.
- 645 10 Javier Esparza, Pierre Ganty, Jérôme Leroux, and Rupak Majumdar. Verification of popu-
646 lation protocols. *Acta Informatica*, 54(2):191–215, 2017. URL: <https://doi.org/10.1007/s00236-016-0272-3>, doi:10.1007/S00236-016-0272-3.
- 647 11 Nathanaël Fijalkow, Nathalie Bertrand, Patricia Bouyer-Decitre, Romain Brenguier, Arnaud
648 Carayol, John Fearnley, Hugo Gimbert, Florian Horn, Rasmus Ibsen-Jensen, Nicolas Mar-
649 key, Benjamin Monmege, Petr Novotný, Mickael Randour, Ocan Sankur, Sylvain Schmitz,
650 Olivier Serre, and Mateusz Skomra. Games on graphs. *CoRR*, abs/2305.10546, 2023.

- 655 URL: <https://doi.org/10.48550/arXiv.2305.10546>, arXiv:2305.10546, doi:10.48550/
656 ARXIV.2305.10546.
- 657 12 Hugo Gimbert. Distributed asynchronous games with causal memory are undecidable. *Log.*
658 *Methods Comput. Sci.*, 18(3), 2022. URL: [https://doi.org/10.46298/lmcs-18\(3:30\)2022](https://doi.org/10.46298/lmcs-18(3:30)2022),
659 doi:10.46298/LMCS-18(3:30)2022.
- 660 13 Lucie Guillou, Corto Mascle, and Nicolas Waldburger. Parameterized broadcast networks with
661 registers: from NP to the frontiers of decidability. In Naoki Kobayashi and James Worrell,
662 editors, *Foundations of Software Science and Computation Structures - 27th International*
663 *Conference, FoSSaCS 2024, Held as Part of the European Joint Conferences on Theory*
664 *and Practice of Software, ETAPS 2024, Luxembourg City, Luxembourg, April 6-11, 2024,*
665 *Proceedings, Part II*, volume 14575 of *Lecture Notes in Computer Science*, pages 250–270.
666 Springer, 2024. doi:10.1007/978-3-031-57231-9_12.
- 667 14 Leonard H Haines. On free monoids partially ordered by embedding. *Journal of Combinatorial*
668 *Theory*, 6(1):94–98, 1969.
- 669 15 Graham Higman. Ordering by divisibility in abstract algebras. *Proceedings of the London*
670 *Mathematical Society*, s3-2(1):326–336, 1952. doi:10.1112/plms/s3-2.1.326.
- 671 16 Lukasz Kaminski and Slawomir Lasota. Bi-reachability in petri nets with data. In Rupak
672 Majumdar and Alexandra Silva, editors, *35th International Conference on Concurrency*
673 *Theory, CONCUR 2024, September 9-13, 2024, Calgary, Canada*, volume 311 of *LIPICs*,
674 pages 31:1–31:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. URL: <https://doi.org/10.4230/LIPICs.CONCUR.2024.31>, doi:10.4230/LIPICs.CONCUR.2024.31.
- 675 17 Eryk Kopczynski. Half-positional determinacy of infinite games. In Michele Bugliesi, Bart Pren-
676 eel, Vladimiro Sassone, and Ingo Wegener, editors, *Automata, Languages and Programming,*
677 *33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings,*
678 *Part II*, volume 4052 of *Lecture Notes in Computer Science*, pages 336–347. Springer, 2006.
679 doi:10.1007/11787006_29.
- 680 18 Ranko Lazic, Thomas Christopher Newcomb, Joël Ouaknine, A. W. Roscoe, and James Worrell.
681 Nets with tokens which carry data. *Fundamenta Informaticae*, 88(3):251–274, 2008.
- 682 19 Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *Proc. ACM POPL*,
683 pages 179–190, 1989.
- 684 20 Frank P Ramsey. On a problem of formal logic. In *Classic Papers in Combinatorics*, pages
685 1–24. Springer, 1987.
- 686 21 Sven Schewe and Bernd Finkbeiner. Synthesis of asynchronous systems. In Germán Puebla,
687 editor, *Logic-Based Program Synthesis and Transformation, 16th International Symposium,*
688 *LOPSTR 2006, Venice, Italy, July 12-14, 2006, Revised Selected Papers*, volume 4407 of *Lecture*
689 *Notes in Computer Science*, pages 127–142. Springer, 2006. doi:10.1007/978-3-540-71410-1_10.
- 690 22 Sylvain Schmitz. Complexity hierarchies beyond elementary. *ACM Transactions on Computa-*
691 *tion Theory*, 8(1):3:1–3:36, 2016. doi:10.1145/2858784.
- 692 23 Sylvain Schmitz and Philippe Schnoebelen. Multiply-recursive upper bounds with higman’s
693 lemma. In *International Colloquium on Automata, Languages and Programming, ICALP*
694 *2011*, volume 6756 of *Lecture Notes in Computer Science*, pages 441–452. Springer, 2011.
695 doi:10.1007/978-3-642-22012-8_35.
- 696 24 Daniel Stan. *Randomized strategies in concurrent games. (Stratégies randomisées dans les*
697 *jeux concurrents)*. PhD thesis, University of Paris-Saclay, France, 2017. URL: [https://tel.](https://tel.archives-ouvertes.fr/tel-01519354)
698 [archives-ouvertes.fr/tel-01519354](https://tel.archives-ouvertes.fr/tel-01519354).
- 699 25 Larry J. Stockmeyer. Planar 3-colorability is polynomial complete. *SIGACT News*, 5(3):19–25,
700 1973. doi:10.1145/1008293.1008294.
- 701 26 Steffen van Bergerem, Roland Guttenberg, Sandra Kiefer, Corto Mascle, Nicolas Waldburger,
702 and Chana Weil-Kennedy. Verification of population protocols with unordered data. In Karl
703 Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson, editors, *51st International*
704 *Colloquium on Automata, Languages, and Programming, ICALP 2024, July 8-12, 2024,*
705
706

23:18 Controller Synthesis for Broadcast Networks with Data

- 707 *Tallinn, Estonia*, volume 297 of *LIPICs*, pages 156:1–156:20. Schloss Dagstuhl - Leibniz-
708 Zentrum für Informatik, 2024. URL: <https://doi.org/10.4230/LIPICs.ICALP.2024.156>,
709 doi:10.4230/LIPICs.ICALP.2024.156.
- 710 **27** Peter van Emde Boas. The convenience of tilings. In *Complexity, Logic, and Recursion Theory*,
711 pages 331–363. CRC Press, 2019.

A

 Missing proof from Section 3

► **Lemma 12** (Invariants characterise winning control strategies). *A control strategy σ is winning if and only if there exists a sufficient invariant $I \subseteq \mathcal{M}$ for it.*

Proof. \Rightarrow Suppose σ is winning. Let I be the set of messages such that there exists a σ -run in which they are broadcast. We show that I is a sufficient invariant for σ .

As σ is winning, m_{err} can never be broadcast, thus $m_{err} \notin I$. Suppose by contradiction that we have a σ -local run $s_0 \xrightarrow{\text{op}_1(m_1)}_{\delta_1} s_1 \xrightarrow{\text{op}_2(m_2)}_{\delta_2} \dots \xrightarrow{\text{op}_k(m_k)}_{\delta_k} s_k$ with $s_0 = s_{init}$ broadcasting some $m_{out} \notin I$ and only receiving messages of I . Then we can construct a σ -run in which m_{out} is broadcast.

We proceed by induction: for all $i \in [0, k]$, we show that there is a run ϱ_i whose projection on some agent a is $s_0 \xrightarrow{\text{op}_1(m_1)}_{\delta_1} \dots \xrightarrow{\text{op}_i(m_i)}_{\delta_i} s_i$. For $i = 0$ this is immediate. Let $i > 0$, suppose we have constructed ϱ_{i-1} , and let us construct ϱ_i . Let \mathbb{A}_{i-1} be the set of agents of ϱ_{i-1} .

■ If $s_{i-1} \xrightarrow{\text{op}_i(m_i)}_{\delta_i} s_i$ is a broadcast step, then we simply execute ϱ_{i-1} and then make a apply that broadcast, which no other agent receives.

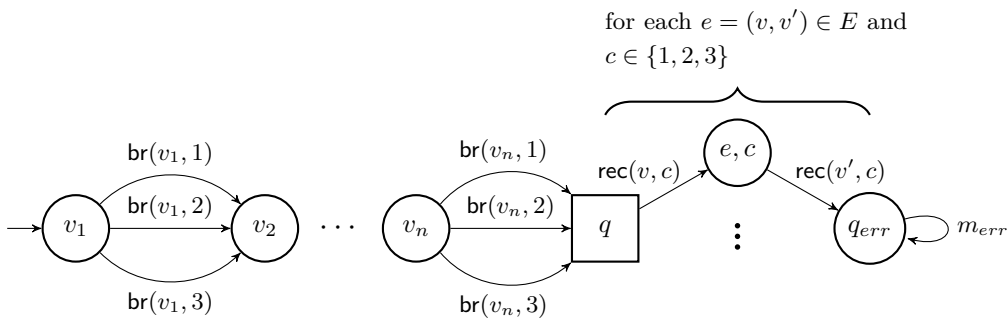
■ If $s_{i-1} \xrightarrow{\text{op}_i(m_i)}_{\delta_i} s_i$ is a reception step, in which a message type m is received, then we have $m \in I$, by construction of the σ -local run. Hence there exists a σ -run ϱ_m over a set of agents \mathbb{A}_m in which m is broadcast. Up to renaming agents, we can assume that \mathbb{A}_{i-1} and \mathbb{A}_m are disjoint. We then define ϱ_i over $\mathbb{A}_{i-1} \sqcup \mathbb{A}_m$ by executing ϱ_{i-1} over \mathbb{A}_{i-1} , then executing ϱ_m over \mathbb{A}_m up to the point before an agent a_m broadcasts m . Finally, we make a_m broadcast m and a receive it.

In both cases we obtain a σ -run in which the local run of a is $s_0 \xrightarrow{\text{op}_1(m_1)}_{\delta_1} \dots \xrightarrow{\text{op}_i(m_i)}_{\delta_i} s_i$. In particular, for $i = k$, we get a σ -run in which m_{out} is broadcast. As $m_{out} \notin I$, this contradicts the definition of I . Hence I satisfies both items of the lemma.

\Leftarrow Suppose there exists $I \subseteq \mathcal{M}$ satisfying the conditions of the lemma. Suppose by contradiction that there is a σ -run ϱ in which m_{err} is broadcast. Let m be the first message broadcast in ϱ that is not in I , and a the agent broadcasting it. Those are well-defined as $m_{err} \notin I$. Let ϱ' be the prefix of ϱ stopping right after that broadcast. The projection $\pi_a(\varrho')$ of ϱ' on a contains a broadcast of m but no reception of any $m' \notin I$, a contradiction. ◀

► **Lemma 34.** *The SAFESTRAT problem is NP-hard for BGR without registers.*

We reduce from the graph 3-colouring problem [25].



■ **Figure 3** Illustration of the lower bound proof from Theorem 13.

742 Consider an undirected graph $G = (V, E)$ with $V = \{v_1, \dots, v_n\}$. We build a BGR with
 743 no registers as in Figure 3. From the initial state Controller chooses one of $(v_i, 1), (v_i, 2), (v_i, 3)$
 744 for each $i \in [1, n]$ and broadcasts it. Then Environment picks an edge $e = (v, v') \in E$ and
 745 $c \in \{1, 2, 3\}$ and tries to reach q_{err} by receiving (v, c) and (v', c) .

746 A strategy for Controller comes down to a colouring of V . It is winning if Environment
 747 can find an edge e and c such that both ends of e are coloured with c . In other words,
 748 Controller wins if and only if the selected colouring of V is a valid 3-colouring of G .

749 **B** Missing proofs from Section 4

750 We show that the invariants defined for signature BGR are accurate witnesses for winning
 751 control strategies. The idea behind this construction already existed in [13].

752 ► **Lemma 18** (Invariants characterise winning strategies). *A control strategy σ is winning if
 753 and only if there exists a sufficient invariant $I \subseteq \mathcal{M}^*$ for it.*

754 **Proof.** \Rightarrow Suppose σ is winning. Let I be the set of words such that there exists a σ -run, an
 755 agent a and a datum d such that w is a subword of the d -output of the projection of that run
 756 on a . The empty word is in I as it is the output of a local run of length 0, which is a σ -run. As
 757 σ is winning, m_{err} can never be broadcast, thus $m_{err} \notin I$. For the other condition, consider
 758 a σ -local run $u = (s_0, c_0) \xrightarrow{\text{op}_1(m_1, d_1)}_{\delta_1} (s_1, c_1) \xrightarrow{\text{op}_2(m_2, d_2)}_{\delta_2} \dots \xrightarrow{\text{op}_k(m_k, d_k)}_{\delta_k} (s_k, c_k)$ whose
 759 d -input is in I for every datum d .

760 Then we can construct a σ -run in which some agent has output $\text{Out}_{sign}(u)$, thus proving
 761 that $\text{Out}_{sign}(u) \in I$. This construction is illustrated in Figure 4.

762 Let D be the set of data appearing in u . For each datum $d \in D$, let w_d be the d -input of
 763 u . As $w_d \in I$, there exists a σ -run ϱ_d such that w_d is a subword of the output of an agent a_d .
 764 Let \mathbb{A}_d be the set of agents of that σ -run.

765 Up to renaming data and agents, we can assume that the initial datum of a_d in ϱ_d is d ,
 766 and that the σ -runs $(\varrho_d)_{d \in D}$ operate over disjoint sets of data and agents.

767 We take a fresh agent a . We construct a σ -run ϱ over $\{a\} \sqcup \bigsqcup_{d \in D} \mathbb{A}_d$ as follows. We make
 768 a follow the local run u . Whenever a needs to receive a message (m, d) , we run ϱ_d over \mathbb{A}_d
 769 until a message (m, d) is broadcast by a_d , and make a receive it. Then we continue running
 770 u . As $\text{In}_d(u)$ is a subword of the output of $\pi_{a_d}(\varrho_d)$ for all $d \in D$, we eventually run u in full.

771 This yields a valid σ -run in which u is fully executed by a . Hence we have a σ -run in
 772 which agent a outputs $\text{Out}_{sign}(u)$. By definition of I , we thus have $\text{Out}_{sign}(u) \in I$.

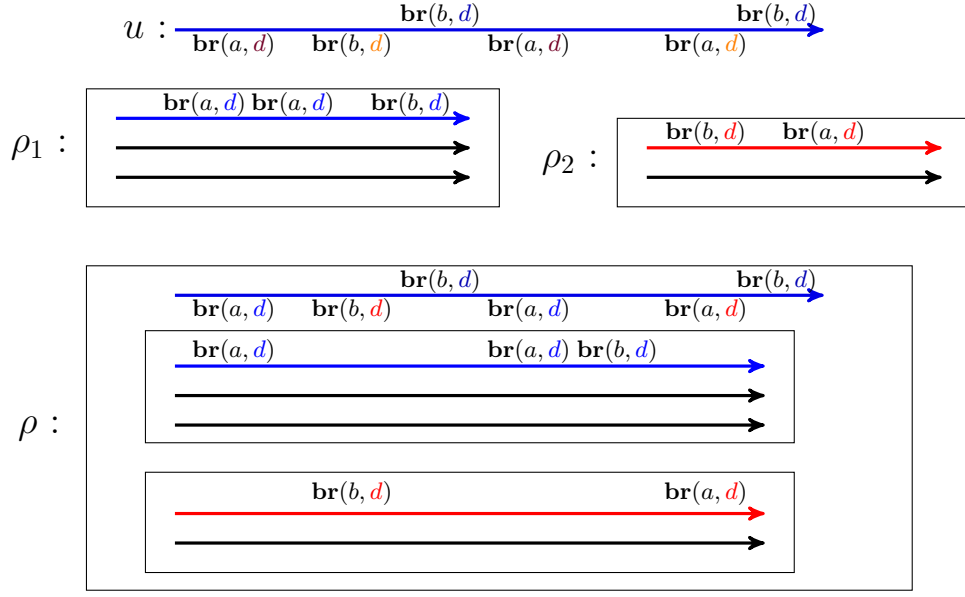
773 \Leftarrow Suppose there exists $I \subseteq \mathcal{M}^*$ satisfying the conditions of the lemma. Suppose by
 774 contradiction that there is a σ -run ϱ in which m_{err} is broadcast.

775 Let ϱ_- be the maximal prefix of ϱ such that the output of each agent is in I . It is
 776 well-defined as $\varepsilon \in I$, thus the prefix of ϱ with no step satisfies that condition. As $m_{err} \notin I$
 777 and I is downward-closed, I does not contain any word containing m_{err} . Hence the output
 778 of ϱ is not in I , and thus ϱ_- is a strict prefix of ϱ .

779 Let a be the agent making the broadcast of the step right after ϱ_- in ϱ , and let m be the
 780 message it broadcasts. Let ϱ_+ be the prefix of ϱ made of ϱ_- and that extra step.

781 Let w_- be the output of a in ϱ_- . For all $d \in \mathbb{D}$, the d -input of a in ϱ_- must be a subword
 782 of the output of another agent. By definition of ϱ_- , the d -input of a in ϱ_- is thus in I for all
 783 d . As the d -input of a in ϱ_- and ϱ_+ is the same, the d -input of a in ϱ_+ is in I for all d . By
 784 maximality of ϱ_- , the output of a in ϱ_+ is not in I .

785 This contradicts the second condition on I given by the lemma. \blacktriangleleft



■ **Figure 4** Illustration of the proof of Lemma 18. Most information is omitted, we only represent schematically the relevant broadcasts and receptions. Data are represented by colours. If we have a local run u outputting bb and for each d a run in which an agent outputs the d -input of u , then we can rename some data and compose those runs to form a run in which an agent outputs bb . Local runs of relevant agents are coloured with their initial datum.

786 ▶ **Lemma 20.** *Let $I \subseteq \mathcal{M}^*$ be a downward-closed set of words containing ε and not m_{err} . If*
 787 *Controller wins the invariant game $\mathcal{IG}(\mathcal{G}, I)$ then there is a control strategy σ such that I is*
 788 *a sufficient invariant for σ .*

789 **Proof.** Let $\sigma_{\mathcal{IG}}$ be a winning strategy for Controller in $\mathcal{IG}(\mathcal{G}, I)$. We define σ as the control
 790 strategy in which Controller follows $\sigma_{\mathcal{IG}}$. That is, given a local run $u = (s_0, c_0) \xrightarrow{\text{op}_1(m_1, d_1)}_{\delta_1}$
 791 $\dots \xrightarrow{\text{op}_k(m_k, d_k)}_{\delta_k} (s_k, c_k)$, we set $\sigma(u) = \sigma_{\mathcal{IG}}(\delta_1 \dots \delta_k)$.

792 We show that I is a sufficient invariant for σ . Assume by contradiction that we have
 793 a σ -local run $u = (s_0, c_0) \xrightarrow{\text{op}_1(m_1, d_1)}_{\delta_1} \dots \xrightarrow{\text{op}_k(m_k, d_k)}_{\delta_k} (s_k, c_k)$ such that u has an output
 794 outside of I , and its d -input is in I for all $d \in \mathbb{D}$. We then show that $\pi = \delta_1 \dots \delta_k$ is a losing
 795 $\sigma_{\mathcal{IG}}$ -play for Controller in $\mathcal{IG}(\mathcal{G}, I)$. As u is a σ -local run, by definition of σ , $\delta_1 \dots \delta_k$ is a
 796 $\sigma_{\mathcal{IG}}$ -play.

797 For all $j \in [0, k]$ let u_j be the prefix of u up to (s_j, c_j) and $\pi_j = \delta_1 \dots \delta_j$.

798 ▷ **Claim 35.** For all $j \in [0, k]$ and $i \in [2, r]$ we have $\text{recentIn}_i(\pi_j) \sqsubseteq \text{In}_{u_j}(c_j(i))$ and
 799 $\text{Out}(\pi_j) = \text{Out}_{\text{sign}}(u_j)$.

800 **Proof.** By a straightforward induction on j . ◁

801 We can instantiate the previous claim with $j = k$ to obtain $\text{Out}(\pi) = \text{Out}_{\text{sign}}(u)$. As we
 802 assumed that $\text{Out}_{\text{sign}}(u) \notin I$, we have $\text{Out}(\pi) \notin I$. As the d -input of u is in I for all $d \in \mathbb{D}$,
 803 and I is downward-closed, the letters of all messages received in u are in I . Moreover, by the
 804 previous claim, for all $j \in [0, k]$, we have $\text{recentIn}_i(\pi_j) \sqsubseteq \text{In}_{u_j}(c_j(i)) \sqsubseteq \text{In}_u(c_j(i)) \in I$. As I is
 805 downward-closed, we have $\text{recentIn}_i(\pi_j) \in I$ for all i, j .

23:22 Controller Synthesis for Broadcast Networks with Data

806 As a result, π is a losing $\sigma_{\mathcal{IG}}$ -play for Controller. This contradicts the assumption that
 807 $\sigma_{\mathcal{IG}}$ is a winning strategy for $\mathcal{IG}(\mathcal{G}, I)$. In consequence, there is no σ -local run u whose
 808 output is outside of I , and whose d -input is in I for all $d \in \mathbb{D}$.

809 This means that I is a sufficient invariant for σ . ◀

810 ► **Lemma 23** (Bounding the size of the invariant). *Let \mathcal{G} a signature BGR. There is a winning
 811 control strategy for \mathcal{G} if and only if there is a sequence of words $w_0, \dots, w_k \in \mathcal{M}^*$ such that*

- 812 ■ *Controller wins $\mathcal{IG}(\mathcal{G}, \{w_1, \dots, w_k\}^{\uparrow^c})$,*
- 813 ■ *and for all $i \in [1, k]$, $|w_i| \leq \psi(|w_{i-1}|)$.*

814 **Proof.** Suppose there is a winning control strategy σ . By Lemma 18 there is a downward-
 815 closed sufficient invariant $I \subseteq \mathcal{M}^*$ for σ . By Lemma 21, Controller wins $\mathcal{IG}(\mathcal{G}, I)$, so the
 816 first condition is satisfied.

817 For the second condition, as I^c is upward-closed it has a finite basis B . Let w_0, w_1, \dots, w_k
 818 be the elements of B sorted by length. We can assume that we took I so that k is minimal. For
 819 all $j \in [1, k]$, we define $B_j = \{w_i \mid i < j\}$ and $I_j = B_j^{\uparrow^c}$. Note that we have $I \subseteq I_k \subseteq \dots \subseteq I_0$.
 820 As I contains ε and not m_{err} , we can assume $w_0 = m_{err}$. By minimality of k , for all $j \in [1, k]$
 821 the set I_j is not a sufficient invariant for σ .

822 By Lemma 21, there is a σ -local run of length at most $\varphi(\mathcal{R}, B_j)$ whose output is not in
 823 I_j and whose d -inputs are all in I_j . As I is a sufficient invariant for σ , one of those d -inputs
 824 must not be in I . We choose one of those and call it w . As a consequence, there exists w_ℓ
 825 with $\ell \geq j$ such that $w_\ell \sqsubseteq w$, and thus $|w_\ell| \leq |w| \leq \varphi(\mathcal{R}, B_j)$. As $|w_i| \leq |w_{i+1}|$ for all i , this
 826 implies $|w_j| \leq \varphi(\mathcal{R}, B_j) = |\mathcal{R}|(|B_j| + 1)^{(|B_j|+1)}$. As w_{j-1} is of maximal length among words
 827 of B_j , we have $|B_j| = |w_{j-1}|$ and $|B_j| \leq |\mathcal{M}|^{|w_{j-1}|+1}$.

828 As a result, $|w_j| \leq |\mathcal{R}|(|w_{j-1}| + 1)^{|\mathcal{M}|^{|w_{j-1}|+1} + 1} = \psi(|w_{j-1}|)$. Thus the second condition
 829 of the lemma is also satisfied.

830 The other direction follows by Lemma 20 and Lemma 18. ◀

831 **C** Missing proofs from Section 5

832 **C.1** Characterisation of winning strategies (Section 5)

833 ► **Lemma 29** (Invariants characterise winning strategies). *A control strategy σ is winning if
 834 and only if there exists a sufficient invariant $(I, (J_m)_{m \in \mathcal{M}})$ for it.*

835 This section is dedicated to the proof of this lemma. To do so, we need an argument
 836 that resembles the construction illustrated in Figure 4. However, the construction gets more
 837 involved in this case.

838 We can start by proving the easier direction of the equivalence, given by the following
 839 lemma. The general structure of this construction already existed in [13]. However, the
 840 different nature of the objects used here and there make it difficult to use their proof as a
 841 black box. We have to go through all the steps here.

842 ► **Lemma 36.** *If there exists a sufficient invariant $(I, (J_m)_{m \in \mathcal{M}})$ for a control strategy σ
 843 then σ is winning.*

844 **Proof.** Suppose σ has a sufficient invariant $(I, (J_m)_{m \in \mathcal{M}})$. Suppose by contradiction that
 845 there is a σ -run ρ in which m_{err} is broadcast.

846 Let a be an agent broadcasting m_{err} in ρ , let u be its local run.

847 We first show that the local run of a in ρ does not satisfy 3a and 3b.

848 If m_{err} is broadcast in u with its initial datum then, as $m_{err} \notin I$ and I is downward-closed,
 849 we cannot have $\mathbf{Out}_d(u) \in I$. On the other hand, if m_{err} is broadcast in u with another
 850 datum d' then as $J_{m_{err}} = \emptyset$, $\mathbf{In}_{d'}(u) \notin J_{m_{err}}$. Hence u does not satisfy 3a and 3b.

851 Let ϱ_- be the maximal prefix of ϱ such that the local runs of all agents satisfy 3a and 3b.
 852 It is well-defined: we saw that the full run ϱ does not satisfy this requirement, and as $\varepsilon \in I$,
 853 the prefix of ϱ with no step satisfies it. Furthermore ϱ_- must be a strict prefix of ϱ .

854 Let a be the agent making the broadcast of the step right after ϱ_- in ϱ , and let m be the
 855 message it broadcasts. Let ϱ_+ be the prefix of ϱ made of ϱ_- and that extra step.

856 By maximality of ϱ_- , there must be an agent whose local run in ϱ_+ does not satisfy 3a
 857 and 3b. This agent can only be a : all agents satisfied both conditions in ϱ_- , an agent cannot
 858 switch from satisfying to not satisfying those conditions without making a broadcast (for
 859 3b, this is due to the fact that all J_m are upward-closed), and a is the only one who made a
 860 broadcast in the last step.

861 As a consequence, the local run u_+ of a in ϱ_+ must dissatisfy either 3a or 3b. It remains
 862 to show that u_+ satisfies both 3i and 3ii to obtain a contradiction.

863 We start by showing that the local run u_- of a in ϱ_- satisfies 3i and 3ii.

864 ■ Let d be the initial datum of u_- . Let m_1, \dots, m_k be the letters such that (m_i, d) is
 865 broadcast by an agent that is not a during ϱ_- . Let us cut ϱ_- into sections $\varrho_0 \cdots \varrho_k$ such
 866 that $\varrho_0 \cdots \varrho_i$ is the maximal prefix of ϱ_- in which (m_i, d) has not been broadcast by
 867 any agent apart from a . For each i let u_i be the projection of ϱ_i on a . We thus have
 868 $u_- = u_0 \cdots u_k$. Let a_{m_i} be the first agent different from a who broadcasts (m_i, d) and let
 869 u_{m_i} be the projection of ϱ_- on a_{m_i} . Let w_i be the sequence of letters broadcast in ϱ_i
 870 with datum d .

871 Consider the decomposition $\mathbf{dec} = (v_0, m_1, \dots, v_{k-1}, m_k, v_k)$ where $v_i = \mathbf{Out}_d(u_i)$. By
 872 definition u_- must be compatible with it. Let $i \in [1, k]$. As u_{m_i} satisfies 3b, we
 873 have $\mathbf{In}_d(u_{m_i}) \in J_{m_i}$. By definition, we must have $\mathbf{In}_d(u_{m_i}) \sqsubseteq w_0 \cdots w_{i-1}$ and thus
 874 $w_0 \cdots w_{i-1} \in J_{m_i}$ as J_{m_i} is upward-closed. Furthermore, each w_j (the letters sent
 875 in ϱ_j with datum d) can be obtained from v_j (the ones sent by a) by adding letters
 876 of $\{m_1, \dots, m_j\}$ (the broadcasts of other agents). As a result, we have $w_0 \cdots w_{i-1} \in$
 877 $\mathcal{L}_{(v_0, m_1, \dots, v_{i-1})}$. We obtain that $w_0 \cdots w_{i-1} \in J_{m_i} \cap \mathcal{L}_{(v_0, m_1, \dots, v_{i-1})}$, thus $J_{m_i} \cap \mathcal{L}_{(v_0, m_1, \dots, v_{i-1})}$
 878 is not empty. In conclusion, $\mathbf{dec} \in \mathcal{D}((J_m)_{m \in \mathcal{M}})$.

879 ■ We now show that u_- satisfies 3ii.

880 Let $d' \neq d$. If d' does not appear in ϱ_- then $\mathbf{In}_{d'}(u_-) = \varepsilon \in \mathcal{L}(I, (J_m)_{m \in \mathcal{M}})$. Otherwise,
 881 let a' be the agent whose initial datum in ϱ is d' . We set w' the sequence of letters
 882 broadcast with datum d' in ϱ_- . Clearly $\mathbf{In}_{d'}(u_-) \sqsubseteq w'$. In order to show that $\mathbf{In}_{d'}(u_-)$,
 883 it suffices to show that $w' \in \mathcal{L}(I, (J_m)_{m \in \mathcal{M}})$.

884 We use the same arguments as for the previous item: we cut ϱ_- into sections $\varrho'_0 \cdots \varrho'_k$
 885 according to the times at which new letters are broadcast with d' by agents other than
 886 a . We then construct a decomposition $\mathbf{dec}' = (v'_0, m'_1, \dots, v'_{k-1}, m'_k, v'_k)$ where m'_i is the
 887 message broadcast with d' at the start of ϱ'_i and v'_i is the sequence of letters broadcast by
 888 a' in ϱ'_i .

889 We argue as before that $w' \in \mathcal{L}_{\mathbf{dec}'}$ and $\mathbf{dec}' \in \mathcal{D}(I, (J_m)_{m \in \mathcal{M}})$.

890 We have shown that u_- satisfied 3i and 3ii. To obtain a contradiction, we must show that
 891 u_+ satisfies them as well. By definition, u_+ is u_- with an additional broadcast at the end.

892 ■ Let $\mathbf{dec} = (v_0, m_0, \dots, v_k) \in \mathcal{D}((J_m)_{m \in \mathcal{M}})$ be a decomposition such that u_- is compatible
 893 with \mathbf{dec} . We have $u_- = u_0 \cdots u_k$ where $v_i \sqsubseteq \mathbf{Out}_d(u_i)$ and $\mathbf{In}_d(u_i) \in \{m_1, \dots, m_i\}^*$
 894 for all i . Let u_k^+ be u_k to which we append the last broadcast in u_+ . We obtain

895 $u_- = u_0 \cdots u_{k-1} u_k^+$. Since $v_k \sqsubseteq \mathbf{Out}_d(u_k) \sqsubseteq \mathbf{Out}_d(u_k^+)$ and $\mathbf{In}_d(u_k) = \mathbf{In}_d(u_k^+) \in$
 896 $\{m_1, \dots, m_k\}^*$, we conclude that u_+ is compatible with dec. Hence u_+ satisfies 3i.
 897 ■ As $\mathbf{In}_{d'}(u_-) = \mathbf{In}_{d'}(u_+)$ for all $d' \in \mathbb{D}$, u_+ satisfies 3ii.

898 In conclusion, we have constructed a σ -local run u_+ such that u_+ satisfies 3i and 3ii but
 899 not 3a and 3b, yielding a contradiction.

900 ◀

901 We must now prove the other implication of Lemma 29. Intuitively, the argument goes
 902 as follows.

903 We define a notion of **partial run**. This is a run of a set of agents, but some messages can
 904 be received without being broadcast. They are called **unmatched receptions**. A **local run** is a
 905 particular case of **partial run**, with a single agent.

906 We assume that σ is winning. We take I as the downward-closure of the set of words
 907 $w \in \mathcal{M}^*$ such that there is a σ -run ϱ in which the sequence of messages w is broadcast,
 908 all with the same datum. For each m , we set J_m to be the upward-closure of the set of
 909 words $w = m_1 \cdots m_n$ such that there is a σ -partial run in which the sequence of **unmatched**
 910 **receptions** is of the form $(m_1, d) \cdots (m_n, d)$ for some $d \in \mathbb{D}$, and (m, d) is broadcast at some
 911 point. This should be understood as follows: if we have a run in which $(m_1, d) \cdots (m_n, d)$
 912 is broadcast, then we can compose it with the partial run above, match all the **unmatched**
 913 **receptions** and obtain an extra broadcast of m .

914 The difficulty is to show that those sets form a **sufficient invariant** for σ . In particular,
 915 we need to take a σ -local run u satisfying 3i and 3ii and show that it satisfies 3a and 3b.
 916 We do that by building σ -runs in which the local run of some agent is u .

917 We rely on several technical lemmas. Lemma 38, 39 and 40. Their statements are involved
 918 but they come with illustrations that should give helpful intuition. Before reading the details
 919 of those lemma we recommend that the reader reads the proof of Theorem 29 at the end of
 920 this section, to better understand how those lemmas are used.

921 C.1.1 Definitions for partial runs

922 For the following proof we need to introduce the notion of *partial run*, which describes the
 923 projection of a run on a subset of agents. We then show a key technical lemma that allows
 924 us to construct a run from a local run and a set of suitable partial runs.

925 We will use this lemma to prove a characterisation of **winning control strategies** using
 926 some invariants, like in the previous sections.

927 ► **Definition 37.** *Let γ, γ' two configurations.*

928 A **partial step** $\gamma \rightarrow_p \gamma'$ is defined if either $\gamma \rightarrow \gamma'$ (normal step) or there exist $m \in \mathcal{M}$,
 929 $d \in \mathbb{D}$ such that for all agent a either $\gamma(a) = \gamma'(a)$ or $\gamma(a) \xrightarrow{\text{rec}(m,d)}_\delta \gamma'(a)$ for some reception
 930 transition δ (**unmatched reception** of (m, d)).

931 A **partial run** ϱ is a sequence of partial steps. It is **initial** if it starts in an initial config-
 932 uration. Its **d -input** $\mathbf{In}_d(\varrho)$ is the sequence $m_0 \cdots m_k$ of letters corresponding to **unmatched**
 933 **receptions** with datum d in ϱ . Its **d -output** $\mathbf{Out}_d(\varrho)$ is the sequence of letters corresponding
 934 to broadcasts with datum d in ϱ .

935 Note that a local run can be seen as a partial run with a single agent. Given a control
 936 strategy σ , a **σ -partial run** is a partial run in which the local runs of all agents are σ -local
 937 runs.

A datum d is *initial* in ϱ if it appears in the first configuration. We extend the notion of compatible to partial runs: A partial run ϱ is *compatible* over d with a decomposition $\text{dec} = (v_0, m_1, \dots, v_k)$ if $\varrho = \varrho_0 \cdots \varrho_k$ and for all $i \in [0, k]$, $v_i \sqsubseteq \text{Out}_d(\varrho_i)$ and $\text{In}_d(\varrho_i) \in \{m_1, \dots, m_i\}^*$, with d an initial datum of some agent in ϱ .

The following lemmas give us ways to compose partial runs to obtain complete runs.

Suppose we have a partial run ϱ compatible with a decomposition $\text{dec} = (v_0, m_1, \dots, v_k)$ over an initial datum d .

Suppose that we have, for each non-initial datum d' , a run $\varrho_{d'}$ such that $\text{In}_{d'}(\varrho) \sqsubseteq \text{Out}_{d'}(\varrho_{d'})$.

Also suppose that for each $j \in [1, k]$ we have a partial run ϱ_j such that $\text{In}_d(\varrho'_j) \in \mathcal{L}_{(v_0, m_1, \dots, v_{j-1})}$ and which contains a broadcast of (m, d) , and no unmatched receptions on data other than d .

■ First, we show that given a word $w \in \mathcal{L}_{\text{dec}}$ we can use the ϱ_i to extend ϱ and obtain a σ -partial run which is still compatible with dec and whose d -output contains w as a subword. This is done by composing ϱ with many copies of each ϱ_i to fill in the missing broadcasts.

■ Then, we show that we can again use many copies of the ϱ_i to eliminate the unmatched receptions with datum d . We do this by carefully adding the necessary copies of ϱ_i , by decreasing i . Each time we fill in a missing broadcast of m_i while possibly adding new ones for some of the m_j with $j < i$. This terminates as the number of unmatched receptions of each letter m_i decreases with respect to the lexicographic ordering.

■ We show that for each non-initial d' we can eliminate the unmatched receptions with datum d' by composing that partial run with the σ -runs $\varrho_{d'}$. We use the broadcasts in $\varrho_{d'}$ to match the unmatched receptions in ϱ over d' .

■ Finally, we combine the two first steps to show that given a run compatible with a decomposition dec over some datum d and a word $w \in \mathcal{L}_{\text{dec}}$, we can extend this run to obtain another run whose d -output contains w .

C.1.2 Extending the output

► **Lemma 38.** *Let $\text{dec} = (v_0, m_1, \dots, v_k)$ be a decomposition, let $w \in \mathcal{L}_{\text{dec}}$.*

Let d a datum and ϱ an initial σ -partial run compatible with dec over d .

Suppose that for all $j \in [1, k]$ there exist an initial σ -partial run ϱ'_j such that $\text{In}_d(\varrho'_j) \in \mathcal{L}_{\text{dec}_j}$ where $\text{dec}_j = (v_0, m_1, \dots, v_{j-1})$, $\text{In}_{d'}(\varrho'_j) = \varepsilon$ for all $d' \neq d$ and $m_j \sqsubseteq \text{Out}_d(\varrho'_j)$.

Then, there is a partial run $\tilde{\varrho}$ such that

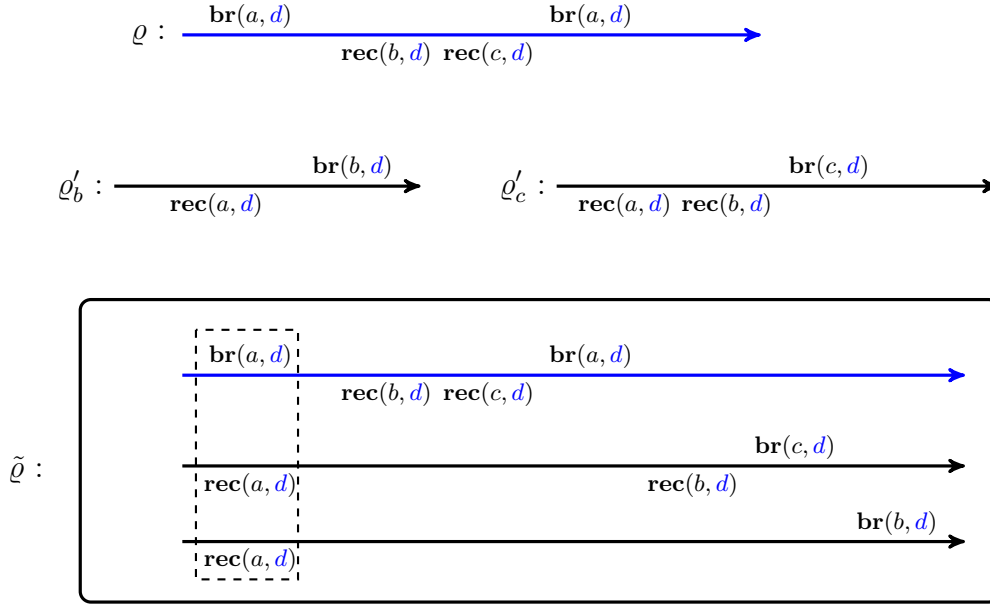
■ $\tilde{\varrho}$ is compatible with dec over d ,

■ $w \sqsubseteq \text{Out}_d(\tilde{\varrho})$

■ for all $d' \neq d$, either $\text{In}_{d'}(\tilde{\varrho}) = \varepsilon$ or $\text{In}_{d'}(\tilde{\varrho}) = \text{In}_{d'}(\varrho)$

Proof. As $w \in \mathcal{L}_{\text{dec}}$, we have $w = w_0 \cdots w_k$, where each w_i can be obtained by adding some letters of $\{m_1, \dots, m_i\}$ to v_i . As u is compatible with dec , $u = u_0 \cdots u_k$ with $v_i \sqsubseteq \text{Out}_d(u_i)$ for all i and $\text{In}_d(u_i) \in \{m_1, \dots, m_i\}^*$. As a consequence, to obtain a d -output that contains w , it suffices to show that we can add a letter from $\{m_1, \dots, m_i\}$ at any point of u_i . We do so using $\tilde{\varrho}_i$: Since $\text{In}_d(\tilde{\varrho}_i) \in \mathcal{L}_{(v_0, m_1, \dots, v_{i-1})} \downarrow$, we can split $\tilde{\varrho}_i$ into $\tilde{\varrho}_{i,0}, \dots, \tilde{\varrho}_{i,j-1}$ so that $\text{In}_d(\tilde{\varrho}_{j,i}) \sqsubseteq \tilde{w}_{j,i}$ where $\tilde{w}_{j,i}$ can be obtained by adding letters from $\{m_1, \dots, m_j\}$ to v_i .

We use the following composition operation: consider ϱ and one of the ϱ'_j . We can build a new run in which we execute both runs in parallel over disjoint sets of agents. We match each $\tilde{\varrho}_{i,j}$ with ϱ_j so that the broadcasts of ϱ_j with d forming v_i are received in $\tilde{\varrho}_{i,j}$ and the only remaining missing broadcasts in that section of the run are with letters m_1, \dots, m_i .



■ **Figure 5** An illustration of the proof of Lemma 38. The partial run ϱ is compatible with decomposition $(a, b, a, c, \varepsilon)$. We have $\mathbf{In}_d(\varrho_b) = a \in \mathcal{L}_{(a)}$ and $\mathbf{In}_d(\varrho_b) = ab \in \mathcal{L}_{(a,b,a)}$. We build a partial run $\tilde{\varrho}$ such that $aacb \sqsubseteq \mathbf{Out}_d(\tilde{\varrho})$. Note that $\tilde{\varrho}$ is also compatible with decomposition $(a, b, a, c, \varepsilon)$. We ignore data other than d in this picture.

983 We obtain a run section whose d -output still contains v_i and whose d -input only contains
 984 m_1, \dots, m_i . This lets us get to a point where the next step in $\tilde{\varrho}_j$ is a broadcast of (m_j, d)
 985 and ϱ has been executed up to the beginning of ϱ_j . We may then use the (m_j, d) broadcast
 986 at any moment in the rest of ϱ to extend the d -output. As a consequence, we can compose ϱ
 987 with the ϱ'_i as many times as necessary to obtain a run $\tilde{\varrho}$ whose d -output contains w .

988 Each composition maintains the fact that the run is compatible with dec . Further, for
 989 all $d' \neq d$, either d' does not appear in ϱ and $\mathbf{In}_d(\tilde{\varrho}) = \varepsilon$ or d' appears in ϱ and then
 990 $\mathbf{In}_{d'}(\tilde{\varrho}) = \mathbf{In}_{d'}(\varrho)$. ◀

991 C.1.3 Unmatched receptions with initial data

992 ▶ **Lemma 39.** Let $\text{dec} = (v_0, m_1, \dots, v_k)$ be a decomposition, d a datum, ϱ an initial σ -partial
 993 run compatible with dec over d .

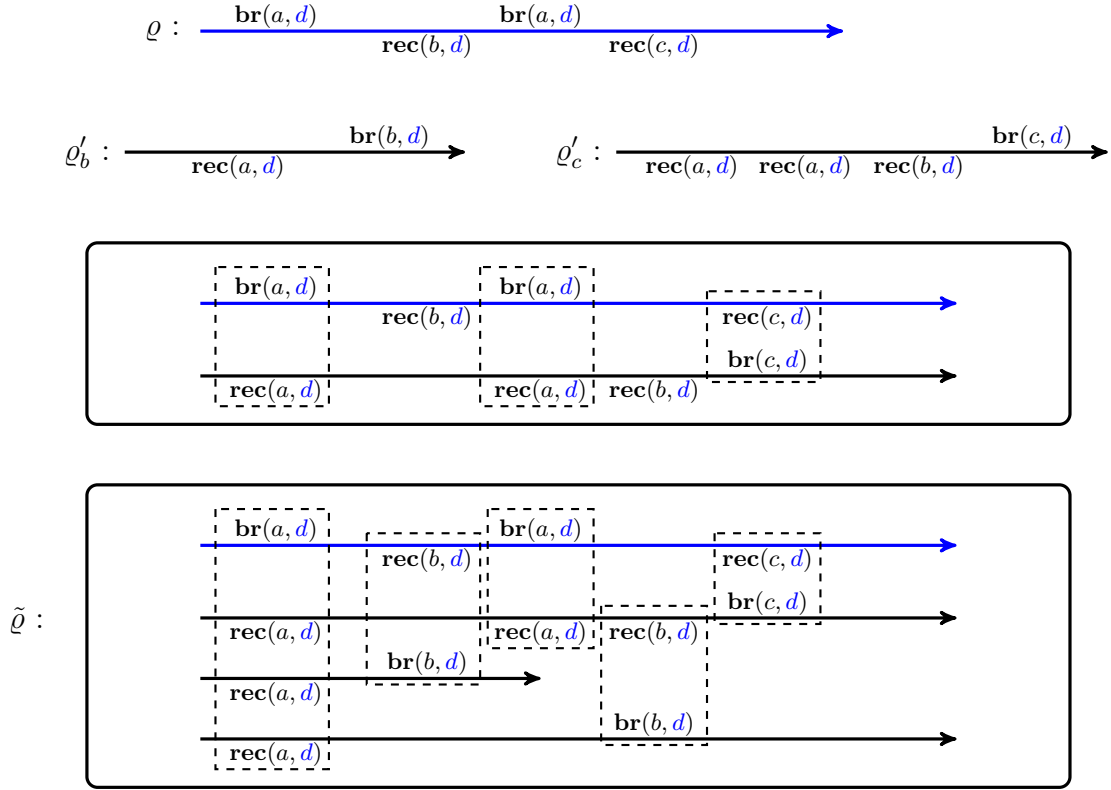
994 Suppose that for all $j \in [1, k]$ there exist an initial partial run ϱ'_j in which d is not initial
 995 such that $\mathbf{In}_{\varrho'_j}(d) \in \mathcal{L}_{\text{dec}_j}$ where $\text{dec}_j = (v_0, m_1, \dots, v_{j-1})$, $\mathbf{In}_{\varrho'_j}(d') = \varepsilon$ for all $d' \neq d$ and
 996 $m_j \sqsubseteq \mathbf{Out}_d(\varrho'_j)$.

997 Then, there exist a σ -partial run $\tilde{\varrho}$ such that

- 998 ■ $\mathbf{In}_d(\tilde{\varrho}) = \varepsilon$,
- 999 ■ $\mathbf{Out}_d(\varrho) \sqsubseteq \mathbf{Out}_d(\tilde{\varrho})$,
- 1000 ■ for all $d' \neq d$, $\mathbf{In}_{d'}(\tilde{\varrho}) = \mathbf{In}_{d'}(\varrho)$

1001 **Proof.** We proceed in the same way as in the previous part: the goal is now to use the
 1002 partial runs ϱ'_j to eliminate the d -input of ϱ .

1003 As ϱ is compatible with dec over d , we can split ϱ into $\varrho_0, \dots, \varrho_k$ with $w_i \sqsubseteq \mathbf{Out}_d(\varrho_i)$ and
 1004 $\mathbf{In}_d(\varrho_i) \in \{m_1, \dots, m_i\}^*$ for all i . Again, we rename agents and data so that the sets of



■ **Figure 6** An illustration of the proof of Lemma 39. The partial run ϱ is compatible with decomposition $(a, b, a, c, \varepsilon)$. We have $\mathbf{In}_d(\varrho_b) = a \in \mathcal{L}_{(a)}$ and $\mathbf{In}_d(\varrho_c) = aab \in \mathcal{L}_{(a,b,a)}$. We build $\tilde{\varrho}$ such that $\mathbf{In}_d(\tilde{\varrho}) = \varepsilon$. We start by using ϱ'_c to eliminate the unmatched receptions of c (while adding some unmatched receptions of b), then we use ϱ'_b to eliminate the unmatched receptions of b . We ignore data other than d in this picture.

1005 agents of ϱ and of every ϱ'_j are all disjoint and the only shared datum between any two of
 1006 these runs is d .

1007 We once again use the composition operation described in the proof of Lemma 38: consider
 1008 ϱ and one of the ϱ'_j . We execute both runs in parallel and match each $\tilde{\varrho}_{i,j}$ with ϱ_j so that
 1009 the broadcasts of ϱ_j with d forming v_i are received in $\tilde{\varrho}_{i,j}$, leaving only unmatched receptions
 1010 with letters m_1, \dots, m_i . We obtain a run section whose d -output still contains v_i and whose
 1011 d -input only contains m_1, \dots, m_i . We can do that until the next step in $\tilde{\varrho}_j$ is a broadcast
 1012 of (m_j, d) and ϱ has been executed up to the beginning of ϱ_j . We may then use the (m_j, d)
 1013 broadcast at any moment in the rest of ϱ to match an unmatched reception of ϱ . As a
 1014 consequence, we can compose ϱ with the ϱ'_i as many times as necessary to obtain a run $\tilde{\varrho}$
 1015 with no unmatched receptions on d .

1016 Each composition maintains the fact that the run is compatible with dec. When we
 1017 do a composition with ϱ'_j to match a reception of (m_j, d) , we may add some receptions of
 1018 m_1, \dots, m_{j-1} to the run (the ones of ϱ'_j). However, every composition decreases the number
 1019 of unmatched receptions of m_k, \dots, m_1 for the lexicographic ordering.

1020 As a result, in the end we obtain a run $\tilde{\varrho}$ without any unmatched reception on datum d .
 1021 As ϱ is fully contained in $\tilde{\varrho}$, $\mathbf{Out}_d(\varrho) \sqsubseteq \mathbf{Out}_d(\tilde{\varrho})$. Moreover, for all $d' \neq d$, either d' does not
 1022 appear in ϱ and then $\mathbf{In}_{d'}(\tilde{\varrho}) = \varepsilon$ or d' appears in ϱ and $\mathbf{In}_{d'}(\tilde{\varrho}) = \mathbf{In}_{d'}(\varrho)$ ◀

1023 **C.1.4 Unmatched receptions with non-initial data**

1024 ► **Lemma 40.** *Let ϱ be an initial σ -partial run, d' a datum, ϱ' an initial σ -run. If*
 1025 $\mathbf{In}_d(\varrho) \sqsubseteq \mathbf{Out}_{d'}(\varrho')$ *and d' an initial datum value in ϱ' but not in ϱ , then there exists*
 1026 *an initial σ -partial run $\tilde{\varrho}$ such that*

- 1027 ■ $\mathbf{In}_{d'}(\tilde{\varrho}) = \varepsilon$
- 1028 ■ for all $d'' \neq d'$, $\mathbf{In}_{d''}(\tilde{\varrho}) = \mathbf{In}_{d''}(\varrho)$
- 1029 ■ for all $d'' \neq d'$, $\mathbf{Out}_{d''}(\varrho) \sqsubseteq \mathbf{Out}_{d''}(\tilde{\varrho})$

1030 **Proof.** Up to renaming agents, assume that ϱ and ϱ' have disjoint agents. We rename data
 1031 in ϱ' so that ϱ' has no shared data with ϱ besides d' .

1032 We build $\tilde{\varrho}$ by running ϱ and ϱ' over their respective agents separately. We use the
 1033 broadcasts made by ϱ' with d' to match the unmatched receptions with datum d' in ϱ : this
 1034 gives us a new partial run ϱ with no unmatched reception with datum d' . Furthermore, for
 1035 every datum d'' , either the sequence of broadcasts and unmatched receptions is the same as
 1036 before, or $\mathbf{In}_d(\varrho) = \varepsilon$ (if d'' appears in $\varrho_{d'}$).

1037 The d'' -output can only increase as ϱ is fully executed within $\tilde{\varrho}$. ◀

1038 **C.1.5 How to obtain a word $w \in \mathcal{L}(I, (J_m)_{m \in \mathcal{M}})$**

1039 We now combine Lemmas 38 and 39 to obtain one last useful technical lemma for the proof
 1040 of Lemma 29. It will be used to prove the second condition of Definition 28 when showing
 1041 that an invariant $(I, (J_m)_{m \in \mathcal{M}})$ is sufficient for a strategy σ .

1042 ► **Lemma 41.** *Let σ be a control strategy, I a downward-closed set of words, and $(J_m)_{m \in \mathcal{M}}$*
 1043 *upward-closed ones.*

1044 *Suppose that for all $w \in I$ there is an initial σ -run and a datum d such that $w \sqsubseteq \mathbf{Out}_d(\varrho)$.*
 1045 *Suppose also that for all $m \in \mathcal{M}$ and $w \in J_m$ there is a σ -partial run ϱ and a datum d that*
 1046 *is not initial in ϱ such that $\mathbf{In}_d(\varrho) \sqsubseteq w$, $m \sqsubseteq \mathbf{Out}_d(\varrho)$ and $\mathbf{In}_{d'}(\varrho) = \varepsilon$ for all $d' \neq d$.*

1047 *Then for all $w \in \mathcal{L}(I, (J_m)_{m \in \mathcal{M}})$, there is a σ -run ϱ and a datum d such that $w \sqsubseteq \mathbf{Out}_d(\varrho)$.*

1048 **Proof.** Let $w \in \mathcal{L}(I, (J_m)_{m \in \mathcal{M}})$, w matches a decomposition $\mathbf{dec} = (v_0, m_1, \dots, v_k)$ such
 1049 that $v_0 \cdots v_k \in I$ and, for all j , $\mathcal{L}_{(v_0, m_1, \dots, v_{j-1})} \cap J_{m_j} \neq \emptyset$. Hence we have a σ -run ϱ and
 1050 a datum d such that $v_0 \cdots v_k \sqsubseteq \mathbf{Out}_d(\varrho)$. Note that as ϱ has no unmatched reception, in
 1051 particular, it is compatible with \mathbf{dec} . Also, for all j we have a σ -partial run ϱ_j and a datum
 1052 d_j not initial in ϱ_j such that $\mathbf{In}_{d_j}(\varrho_j) \in \mathcal{L}_{(v_0, m_1, \dots, v_{j-1})} \downarrow$, $m_j \sqsubseteq \mathbf{Out}_{d_j}(\varrho_j)$ and $\mathbf{In}_{d'}(\varrho_j) = \varepsilon$
 1053 for all $d' \neq d_j$.

1054 By Lemma 38, this means that we can obtain a σ -partial run whose d -output contains w ,
 1055 with no unmatched receptions on data other than d , and compatible with \mathbf{dec} .

1056 We can then use Lemma 39 to eliminate all unmatched receptions and obtain a σ -run
 1057 whose d -output contains w . ◀

1058 **C.1.6 Proof of the characterisation lemma**

1059 **Proof of Lemma 29.** \Rightarrow Suppose σ is winning. Consider R the set of σ -runs.

1060 Let $I = \{\mathbf{Out}_d(\varrho) \mid \varrho \in R, d \in \mathbb{D}\} \downarrow$ be the set of all outputs of all σ -runs.

1061 For all $m \in \mathcal{M}$, we set J_m as the upward-closure of the set of $\mathbf{In}_d(\varrho)$ with ϱ a σ -partial
 1062 run such that d is not an initial datum of ϱ , ϱ contains a broadcast of (m, d) and $\mathbf{In}_{d'}(\varrho) = \varepsilon$
 1063 for all $d' \neq d$.

1064 Let us now prove that $(I, (J_m)_{m \in \mathcal{M}})$ is sufficient for σ . As σ is winning, m_{err} is never
 1065 broadcast, and thus never received, in any σ -run. Hence $m_{err} \notin I$. Furthermore, if we had a

1066 word $w \in I \cap J_{m_{err}}$, then we would have a σ -run ϱ and a σ -partial run ϱ' such that m_{err}
 1067 is broadcast in ϱ' , $\mathbf{In}_{d'}(\varrho') \sqsubseteq w \sqsubseteq \mathbf{Out}_d(\varrho)$ and $\mathbf{In}_{d''}(\varrho') = \varepsilon$ for all $d'' \neq d'$. We can assume
 1068 $d = d'$, as we can rename data.

1069 As a result, we could form a σ -run by renaming data and agents such that their sets of
 1070 data and agents are disjoint except for d . We then execute the two runs in parallel, and
 1071 match the unmatched receptions of ϱ' with broadcasts in ϱ , to obtain a σ -run, with no
 1072 unmatched receptions. This contradicts the fact that σ is winning. Hence $I \cap J_{m_{err}} = \emptyset$.
 1073 Further, as an empty run is a σ -run, we have $\varepsilon \in I$.

1074 For the second point, we can simply apply Lemma 41.

1075 It remains to show that a σ -local run satisfying 3i and 3ii also satisfies 3a and 3b. Let u
 1076 be a σ -local run satisfying 3i and 3ii.

1077 ■ First, we construct a σ -run ϱ whose projection on some agent is u , which shows that
 1078 u satisfies 3a. Let d be the initial datum of u . As u satisfies 3i, there is some $\mathbf{dec} =$
 1079 $(v_0, m_1, \dots, v_k) \in \mathcal{D}((J_m)_{m \in \mathcal{M}})$ such that u is compatible with \mathbf{dec} .

1080 By definition of $(J_m)_{m \in \mathcal{M}}$, for each j we have a σ -partial run ϱ_j and a non-initial datum
 1081 d_j such that $\mathbf{In}_{d_j}(\varrho_j) \in \mathcal{L}_{(v_0, m_1, \dots, v_{j-1})} \downarrow$, there are no unmatched receptions with data
 1082 other than d_j , and $m_j \sqsubseteq \mathbf{Out}_{d_j}(\varrho_j)$.

1083 We can thus apply Lemma 39, to obtain a σ -partial run with no unmatched reception
 1084 over d such that $\mathbf{Out}_d(u) \sqsubseteq \mathbf{Out}_d(\varrho)$.

1085 Furthermore, as u satisfies 3ii, by definition of I , for all $d' \neq d$ there is a σ -run $\varrho_{d'}$
 1086 such that $\mathbf{In}_{d'}(u) \sqsubseteq \mathbf{Out}_{d'}(\varrho_{d'})$. We can then apply Lemma 40 on ϱ , with every $d' \neq d$
 1087 appearing in u to obtain a σ -run ϱ' such that $\mathbf{Out}_d(u) \sqsubseteq \mathbf{Out}_d(\varrho')$. This shows that
 1088 $\mathbf{Out}_d(u) \in I$, by definition.

1089 ■ Let $d' \neq d$ and $m \in \mathcal{M}$ be such that (m, d') is broadcast in u . We can apply Lemma 39
 1090 on u and then Lemma 40 on the resulting run, with every $d'' \notin \{d, d'\}$. We obtain a
 1091 σ -partial run in which (m, d') is broadcast and whose d' -input is the same as u . As a
 1092 consequence, u satisfies 3b by definition of $(J_m)_{m \in \mathcal{M}}$.

1093 This concludes the proof of that direction.

1094 \Leftarrow By Lemma 36.

1095 ◀

1096 C.2 The invariant game

1097 The *invariant game* associated with BGR \mathcal{G} and invariant $(I, (J_m)_{m \in \mathcal{M}})$, which we denote
 1098 by $\mathcal{IG}(\mathcal{G}, I, (J_m)_{m \in \mathcal{M}})$ is defined as follows: The set of vertices is $Q_{\mathcal{R}}$: the current state in
 1099 the protocol and a set of registers, which are the ones supposed to contain the initial datum.
 1100 The initial vertex is q_{init} . From each vertex $q \in Q_{\mathcal{R}}$, players choose a transition from q in
 1101 $\Delta_{\mathcal{R}}$. Controller chooses the next transition when q is in Q_{ctrl} , Environment when it is in
 1102 Q_{env} . The state is updated to the target of the transition.

1103 For all play π , we define $\mathbf{reg}(\pi)$ as the set of registers on which there were no record
 1104 transition in π . Intuitively, $\mathbf{reg}(\pi)$ is the set of registers that contain the initial datum of
 1105 the local run.

1106 Given a play π , we define its *initial input* $\mathbf{initIn}(\pi)$ as the sequence of letters received with
 1107 equality transitions with registers of \mathbf{reg} . This represents the sequence of letters received
 1108 with the initial datum. Formally, $\mathbf{initIn}(\varepsilon) = \varepsilon$, and

$$1109 \quad \text{initln}(\delta_1 \cdots \delta_{k+1}) = \begin{cases} \text{initln}(\delta_1 \cdots \delta_k)m \text{ if } \delta_{k+1} \text{ is an equality transition } \xrightarrow{\text{rec}(m,=i)} \\ \text{with } i \in \text{reg}(\delta_1 \cdots \delta_k), \\ \text{initln}(\delta_1 \cdots \delta_k) \text{ otherwise.} \end{cases}$$

1110 For all registers i , we define its recent input on i , written $\text{recentln}_i(\pi)$ like in the previous
 1111 section: it is the sequence of messages received with equality transitions over register i since
 1112 its last reset.

1113 We define the *output* $\text{Out}(\pi)$ of π in a different way as in the signature case. It is the
 1114 sequence of letters broadcast from registers that were in reg at the time of the broadcast.
 1115 Intuitively, this is the sequence of letters that are broadcast with the initial datum in the
 1116 local run. Formally,

$$1117 \quad \text{Out}(\delta_1 \cdots \delta_{k+1}) = \begin{cases} \text{Out}(\delta_1 \cdots \delta_k)m \text{ if } \delta_{k+1} \text{ is a broadcast transition } \xrightarrow{\text{br}(m,i)} \\ \text{with } i \in \text{reg}(\delta_1 \cdots \delta_k), \\ \text{Out}(\delta_1 \cdots \delta_k) \text{ otherwise.} \end{cases}$$

1118 Given a decomposition $\text{dec} = (v_0, m_1, \dots, v_k)$, we say that a play π is *compatible* with
 1119 dec if $\pi = \pi_0 \cdots \pi_k$ and for all j we have $\text{initln}(\pi_j) \in \{m_1, \dots, m_j\}^*$ and $v_j \sqsubseteq \text{Out}(\pi_j)$.

1120 The objective of the game is then described as follows.

- 1121 (A) If at some point the play $\pi = \delta_1 \cdots \delta_k$ is not compatible with any decomposition of
 1122 $\mathcal{D}((J_m)_{m \in \mathcal{M}})$ then Controller wins.
- 1123 (B) If at some point the play $\pi = \delta_1 \cdots \delta_k$ is such that $\text{recentln}_i(\pi) \notin I$ for some $i \notin \text{reg}$
 1124 then Controller wins.
- 1125 (C) If at some point the play $\pi = \delta_1 \cdots \delta_k$ is such that $\text{Out}(\pi) \notin I$ then Environment wins.
- 1126 (D) If at some point of the play a broadcast transition with $i \notin \text{reg}(\pi)$ and $\xrightarrow{\text{br}(m,i)}$ is taken
 1127 while $\text{recentln}_i(\pi) \notin J_m$ (with π the play formed so far) then Environment wins.
- 1128 (E) If the play goes on forever without any of those things happening then Controller wins.

1129 ► **Lemma 42** (Deciding the invariant game). *There is an elementary function $\varphi(N)$ such
 1130 that:*

1131 *Given a BGR \mathcal{G} over a protocol \mathcal{R} and finite sets of words B and $(B_m)_{m \in \mathcal{M}}$, we can
 1132 decide in time $\varphi(|\mathcal{R}| + \|B\| + |B| + \sum_{m \in \mathcal{M}} \|B_m\| + |B_m|)$ whether Controller has a winning
 1133 strategy in $\mathcal{IG}(\mathcal{G}, (B^\uparrow)^c, (B_m^\uparrow)_{m \in \mathcal{M}})$.*

1134 *Furthermore, if Environment has a winning strategy then he has a strategy to win in at
 1135 most $\varphi(|\mathcal{R}| + \|B\| + |B| + \sum_{m \in \mathcal{M}} \|B_m\| + |B_m|)$ steps.*

1136 **Proof.** By Lemma 8, B^\uparrow is a regular language, recognised by a deterministic finite automaton
 1137 $\mathcal{A}_{B^\uparrow} = (Q_B, \mathcal{M}, \Delta_B, q_0^B, F_B)$ with $(\|B\| + 1)^{|B|+1}$ states. Similarly, for each m we can
 1138 construct a deterministic automaton $\mathcal{A}_{B_m^\uparrow} = (Q_{B_m}, \mathcal{M}, \Delta_{B_m}, q_0^{B_m}, F_{B_m})$

1139 Let $I = B^{\uparrow c}$ and for each $m \in \mathcal{M}$, $J_m = B_m^\uparrow$.

1140 We define a deterministic automaton over the alphabet $\Delta_{\mathcal{R}}$ that reads plays $\delta_1 \cdots \delta_k$ of
 1141 $\mathcal{IG}(\mathcal{G}, I, (J_m)_{m \in \mathcal{M}})$ and accepts exactly the winning plays for Environment.

1142 Consider the alphabet $\mathcal{M} \sqcup \bar{\mathcal{M}}$, where $\bar{\mathcal{M}} = \{\bar{m} \mid m \in \mathcal{M}\}$ is a copy of \mathcal{M} .

1143 We define the useful automata in the following claims. Let us define $K = |\mathcal{R}| + |B| +$
 1144 $\|B\| + \sum_{m \in \mathcal{M}} \|B_m\| + |B_m|$

1145 ▷ **Claim 43.** We can construct an NFA of exponential size in K and recognising the language
 1146 $\{v_0\bar{m}_1 \cdots v_k \mid (v_0, m_1 \cdots, v_k) \in \mathcal{D}((J_m)_{m \in \mathcal{M}})\}$.

1147 **Proof.** Consider the language of decompositions defined as $\{v_0\bar{m}_1 \cdots v_{k-1}\bar{m}_k v_k \mid v_0, \dots, v_k \in$
 1148 $\mathcal{M}^*, \bar{m}_1, \dots, \bar{m}_k \in \bar{\mathcal{M}} \text{ distinct.}\}$

1149 This language is recognised by an automaton of exponential size which simply checks
 1150 that each letter of $\bar{\mathcal{M}}$ appears at most once.

1151 We can turn this automaton into a non-deterministic transducer \mathcal{T} that reads a decom-
 1152 position $v_0\bar{m}_1 \cdots v_{k-1}\bar{m}_k v_k$, outputs all the letters of \mathcal{M} that it reads, and can output letters
 1153 of $\bar{\mathcal{M}}$ arbitrarily as soon as it has read them before. If some letter of $\bar{\mathcal{M}}$ is repeated then the
 1154 run is rejected. The set of images of $v_0\bar{m}_1 \cdots v_{k-1}\bar{m}_k v_k$ is exactly $\mathcal{L}_{(v_0, m_1, \dots, v_k)}$.

1155 By composing this transducer with an automaton recognising J_m , we obtain an automaton
 1156 \mathcal{A}_m recognising decompositions that have an image in J_m by the transducer, i.e., the language
 1157 $\{v_0\bar{m}_1 \cdots v_{k-1}\bar{m}_k v_k \mid \mathcal{L}_{(v_0, m_1, \dots, v_k)} \cap J_m \neq \emptyset\}$.

1158 It is then easy to obtain an automaton recognising $\{v_0\bar{m}_1 \cdots v_k \mid (v_0, m_1 \cdots, v_k) \in$
 1159 $\mathcal{D}((J_m)_{m \in \mathcal{M}})\}$ using a product of the automata $(\mathcal{A}_m)_{m \in \mathcal{M}}$.

1160 The resulting automaton is of exponential size in K . ◁

1161 ▷ **Claim 44.** We can construct a deterministic automaton of double-exponential size in K
 1162 recognising plays compatible with a decomposition of $\mathcal{D}((J_m)_{m \in \mathcal{M}})$.

1163 **Proof.** We use the automaton recognising $\{v_0\bar{m}_1 \cdots v_k \mid (v_0, m_1 \cdots, v_k) \in \mathcal{D}((J_m)_{m \in \mathcal{M}})\}$
 1164 defined in the first claim.

1165 We can define a non-deterministic transducer that takes as input a sequence of transitions
 1166 $\pi = \delta_1 \cdots \delta_p$ and outputs some decomposition with which it is compatible. The transducer
 1167 keeps track of $\text{reg}(\pi)$ while reading the play.

1168 The transducer simply guesses a sequence $\bar{m}_1 \cdots \bar{m}_k$ of distinct letters of $\bar{\mathcal{M}}$. It outputs
 1169 them in that order at arbitrary moments while reading π .

1170 When it reads a broadcast transition $\xrightarrow{\text{br}(m,i)}$ over a register currently in $\text{reg}(\pi)$, it
 1171 non-deterministically outputs m or not.

1172 When it reads an equality transition $\xrightarrow{\text{rec}(m,=i)}$ over a register currently in $\text{reg}(\pi)$, if \bar{m}
 1173 has not been broadcast before it goes to a rejecting sink state.

1174 The set of images of a play π are the decompositions it is compatible with. We compose
 1175 this transducer with the automaton from the first claim to get an automaton recognising the
 1176 set of plays compatible with some decomposition of $\mathcal{D}((J_m)_{m \in \mathcal{M}})$. ◁

1177 We have automata for I and each J_m , as well as for plays compatible with a decomposition
 1178 of $\mathcal{D}((J_m)_{m \in \mathcal{M}})$. From those it is straightforward to define an automaton \mathcal{C} reading plays
 1179 and accepting the ones winning for Environment. We can then determinise it at the cost of an
 1180 exponential blow-up.

1181 By Proposition 10, we can solve this game in polynomial time in the size of the resulting
 1182 automaton \mathcal{C} and the size of the arena of $\mathcal{IG}(\mathcal{G}, I, (J_m)_{m \in \mathcal{M}})$ (i.e., $|\mathcal{R}|$), that is, in double-
 1183 exponential time in $\|B\| + |B| + |\mathcal{R}|$.

1184 Furthermore, if Environment has a winning strategy then he has one that guarantees
 1185 that he wins in at most double-exponentially many steps in K . ◀

1186 We have to show that Controller wins the invariant game if and only if she has a winning
 1187 control strategy.

23:32 Controller Synthesis for Broadcast Networks with Data

1188 ► **Lemma 45** (From the invariant game to control strategies). *Let $I \subseteq \Sigma^*$ be a downward-closed*
 1189 *set and $(J_m)_{m \in \mathcal{M}}$ upward-closed sets such that I contains ε and not m_{err} , $J_{m_{err}} \cap I = \emptyset$,*
 1190 *and $\mathcal{L}(I, (J_m)_{m \in \mathcal{M}}) \subseteq I$.*

1191 *If Controller wins the invariant game $\mathcal{IG}(\mathcal{G}, I, (J_m)_{m \in \mathcal{M}})$ then there is a control strategy*
 1192 *σ such that $(I, (J_m)_{m \in \mathcal{M}})$ is a sufficient invariant for σ .*

1193 **Proof.** Let $\sigma_{\mathcal{IG}}$ be a winning strategy for Controller in $\mathcal{IG}(\mathcal{G}, I, (J_m)_{m \in \mathcal{M}})$.

1194 We define σ as the control strategy in which Controller follows $\sigma_{\mathcal{IG}}$. That is, given a
 1195 local run $u = (s_0, c_0) \xrightarrow{\text{op}_1(m_1, d_1)}_{\delta_1} \cdots \xrightarrow{\text{op}_k(m_\ell, d_\ell)}_{\delta_\ell} (s_\ell, c_\ell)$, we set $\sigma(u) = \sigma_{\mathcal{IG}}(\delta_1 \cdots \delta_\ell)$. Let
 1196 d be the initial datum of u . We show that $(I, (J_m)_{m \in \mathcal{M}})$ is a sufficient invariant for σ . To
 1197 do so, we assume by contradiction that we have a σ -local run $u = (s_0, c_0) \xrightarrow{\text{op}_1(m_1, d_1)}_{\delta_1}$
 1198 $\cdots \xrightarrow{\text{op}_\ell(m_\ell, d_\ell)}_{\delta_\ell} (s_\ell, c_\ell)$ such that u satisfies 3i and 3ii but does not satisfy either 3a or 3b.
 1199 Let d be its initial datum.

1200 We then show that $\pi = \delta_1 \cdots \delta_\ell$ is a losing $\sigma_{\mathcal{IG}}$ -play for Controller in $\mathcal{IG}(\mathcal{G}, I)$. As u is a
 1201 σ -local run, by definition of σ , $\delta_1 \cdots \delta_\ell$ is a $\sigma_{\mathcal{IG}}$ -play.

1202 For all $j \in [0, \ell]$ let u_j be the prefix of u up to (s_j, c_j) and $\pi_j = \delta_1 \cdots \delta_j$.

1203 ▷ **Claim 46.** For all index j and $i \notin \text{reg}(\pi_j)$ we have $\text{recentIn}_i(\pi_j) \sqsubseteq \text{In}_{u_j}(c_j(i))$ and
 1204 $\text{initIn}(\pi_j) \sqsubseteq \text{In}_{u_j}(d)$. Furthermore, if $\text{reg}(\pi_j) \neq \emptyset$ then $\text{Out}(\pi_j) = \text{Out}_d(u_j)$.

1205 **Proof.** By a straightforward induction on j . ◁

1206 As u satisfies 3i, it is compatible with a decomposition $\text{dec} = (v_0, m_1, \dots, v_k)$ in
 1207 $\mathcal{D}((J_m)_{m \in \mathcal{M}})$. We thus have $u = u^0 \cdots u^k$ with $v_i \sqsubseteq \text{Out}_d(u^i)$ and $\text{In}_d(u^i) \in \{m_1, \dots, m_i\}^*$
 1208 for all i .

1209 Let j be the maximal index such that $\text{reg}(\pi_j) \neq \emptyset$, and i_0 the maximal index such that
 1210 $\pi^0 \cdots \pi^{i_0}$ is a prefix of π_j .

1211 Hence we can cut π in the same way: $\pi = \pi^0 \cdots \pi^{i_0}$ where π^i is the sequence of transitions
 1212 of u^i . We can infer using the previous claim that $v_j \sqsubseteq \text{Out}_d(u^i) = \text{Out}(\pi^i)$ for all $i \leq i_0$ and
 1213 $\text{initIn}(\pi^j) \sqsubseteq \text{In}_d(u^i) \in \{m_1, \dots, m_i\}^*$.

1214 As a consequence, π_j is compatible with $\text{dec}' = (v_0, m_1, \dots, m_{i_0}, \varepsilon)$. Furthermore, we
 1215 have $\text{initIn}(\pi_j) = \text{initIn}(\pi)$ and we can conclude that π is compatible with dec' , which is in
 1216 $\mathcal{D}((J_m)_{m \in \mathcal{M}})$.

1217 We can also infer that all its prefixes π' are compatible with a decomposition of
 1218 $\mathcal{D}((J_m)_{m \in \mathcal{M}})$: it suffices to consider the decomposition $(v_0, m_1, \dots, m_i, \varepsilon)$, with i the maximal
 1219 index such that m_i is appears in $\text{initIn}(\pi')$.

1220 Furthermore, as u satisfies 3ii, for all $d' \neq d$, we have $\text{In}_d(u_\ell) \in I$. For all j ,

$$1221 \quad \text{recentIn}_{\pi_j}(i) \sqsubseteq \text{In}_{c_j(i)}(u_j) \sqsubseteq \text{In}_{c_\ell(i)}(u_\ell)$$

1222 . As I is downward-closed, we have $\text{recentIn}_{\pi_j}(i) \in I$ for all $j \in [0, \ell]$.

1223 We know that either u does not satisfy 3a or does not satisfy 3b.

1224 Let us first assume that u does not satisfy 3b. Let $d' \neq d$ and $m \in \mathcal{M}$ be such that u
 1225 contains a broadcast of (m, d') while $\text{In}_{d'}(u) \notin J_m$. Let j be the index of the first broadcast
 1226 of (m, d') in u and i the register containing d' at that point. Then δ_j is a broadcast transition
 1227 $\xrightarrow{\text{br}(m, i)}$, while $\text{recentIn}_{\pi_j}(i) \sqsubseteq \text{In}_{d'}(u_j) \sqsubseteq \text{In}_{d'}(u)$. As J_m is upward-closed, $\text{recentIn}_{\pi_j}(i) \notin J_m$,
 1228 which means that π is losing for Controller.

1229 Now we assume that u does not satisfy 3a. Let $u = u_- u_+$ be such that u_- is the maximal
 1230 prefix of u in which d appears at all times. We can cut $\pi = \pi_- \pi_+$ the same way: π_- is the
 1231 sequence of transitions of u_- , and is also the maximal prefix of π such that $\text{reg}(\pi_-) \neq \emptyset$.

1232 ▷ **Claim 47.** Suppose that π is a winning play for Controller. Then there is a decomposition
 1233 $\text{dec} = (v_0, m_1, \dots, v_k) \in \mathcal{D}((J_m)_{m \in \mathcal{M}})$ such that $v_0 \cdots v_k \sqsubseteq \mathbf{Out}(\pi)$ and $\mathbf{Out}_d(u) \in \mathcal{L}_{\text{dec}}$.

1234 **Proof.** Let $\text{dec} = (v_0, m_1, \dots, v_k)$ be a decomposition of $\mathcal{D}((J_m)_{m \in \mathcal{M}})$ such that u is com-
 1235 patible with dec . It exists as u satisfies 3i.

1236 Furthermore, we choose it so that $|v_0 \cdots v_k|$ is minimal. Among the ones with minimal
 1237 $|v_0 \cdots v_k|$, we choose one with k maximal.

1238 Suppose $v_0 \cdots v_k$ is not a subword of $\mathbf{Out}(\pi)$. Then, as $\mathbf{Out}(\pi) = \mathbf{Out}(\pi_-) = \mathbf{Out}_d(u_-)$,
 1239 we get that $v_0 \cdots v_k \not\sqsubseteq \mathbf{Out}_d(u_-)$.

1240 Let i be the minimal index such that $v_0 \cdots v_i \not\sqsubseteq \mathbf{Out}_d(u_-)$, and let v_{i-} be the maximal
 1241 prefix of v_i such that $v_0 \cdots v_{i-} v_{i-} \sqsubseteq \mathbf{Out}_d(u_-)$, and m the letter right after v_{i-} in v_i .
 1242 Let v_{i+} such that $v_i = v_{i-} m v_{i+}$. The letter m must be broadcast with d in u_+ . The
 1243 same broadcast appears in π_+ , say at step j on register i_0 . As we assumed that π is
 1244 winning, we have $\text{recentIn}_{\pi_j}(i_0) \in J_m$. Hence $\mathbf{In}_d(u_j) \in J_m$, as J_m is upward-closed and
 1245 $\text{recentIn}_{\pi_j}(i_0) \sqsubseteq \mathbf{In}_d(u_j)$.

1246 We have three cases:

- 1247 ■ $m \in \{m_1, \dots, m_{i-1}\}$: then it is easily checked that we can remove m from v_i without
 1248 affecting the properties of dec , contradicting the minimality of $|v_0 \cdots v_k|$.
- 1249 ■ $m = m_\ell$ for some $\ell \in [i, k]$: then we can use the following decomposition:

$$1250 (v_0, m_1, \dots, v_{i-1}, m_i, v_{i-}, m, v_{i+}, \dots, m_{\ell-1}, v_{\ell-1} v_\ell, m_{\ell+1}, \dots, v_k)$$

1251 instead of dec , again contradicting the minimality of $|v_0 \cdots v_k|$.

- 1252 ■ $m \notin \{m_1, \dots, m_k\}$. Then we use the following decomposition instead of dec :

$$1253 (v_0, m_1, \dots, v_{i-1}, m_i, v_{i-}, m, v_{i+}, m_{i+1}, \dots, v_k). \text{ This contradicts the minimality of } |v_0 \cdots v_k|$$

1254 .

1255 As a consequence, we obtain that $v_0 \cdots v_k$ is a subword of $\mathbf{Out}(\pi)$. It remains to show
 1256 that $\mathbf{Out}_d(u) \in \mathcal{L}_{\text{dec}}$. To do that, let us assume that u_+ contains a broadcast with d of
 1257 a letter that is not in $\{m_1, \dots, m_k\}$. Let m be the letter in the first such broadcast of
 1258 u_+ , i the corresponding register, and j the index of the step. Since we assumed that π
 1259 is winning, we have $\text{recentIn}_{\pi_j}(i) \in J_m$. Hence $\mathbf{In}_d(u_j) \in J_m$, as J_m is upward-closed and
 1260 $\text{recentIn}_{\pi_j}(i) \sqsubseteq \mathbf{In}_d(u_j)$. Moreover, every letter in $\mathbf{In}_d(u_j)$ must be in $\{m_1, \dots, m_k\}$, as u is
 1261 compatible with dec .

1262 As a result, $\mathbf{In}_d(u_j) \in J_m \cap \mathcal{L}_{\text{dec}}$, hence $J_m \cap \mathcal{L}_{\text{dec}} \neq \emptyset$ and thus $(v_0, m_1, \dots, v_k, m, \varepsilon) \in$
 1263 $\mathcal{D}((J_m)_{m \in \mathcal{M}})$. Moreover, u is compatible with this decomposition. This contradicts the
 1264 maximality of k .

1265 In conclusion, we have shown that dec matches all the conditions of the claim. ◁

1266 Suppose π is winning, then by this claim we have a decomposition $\text{dec} = (v_0, m_1, \dots, v_k) \in$
 1267 $\mathcal{D}((J_m)_{m \in \mathcal{M}})$ such that $v_0 \cdots v_k \sqsubseteq \mathbf{Out}(\pi)$ and $\mathbf{Out}_d(u) \in \mathcal{L}_{\text{dec}}$.

1268 As π is winning, we have $\mathbf{Out}(\pi) \in I$, and thus $\mathbf{Out}_d(u) \in \mathcal{L}(I, (J_m)_{m \in \mathcal{M}})$. Since
 1269 $\mathcal{L}(I, (J_m)_{m \in \mathcal{M}}) \subseteq I$, we get $\mathbf{Out}_d(u) \in I$, and thus u satisfies 3a, a contradiction.

1270 In conclusion, we obtained that π is a losing $\sigma_{\mathcal{IG}}$ -play, which contradicts the assumption
 1271 that $\sigma_{\mathcal{IG}}$ is winning. As a consequence, $(I, (J_m)_{m \in \mathcal{M}})$ is a sufficient invariant for σ . ◀

1272 ▶ **Lemma 48** (From control strategies to the invariant game). *Let σ be a control strategy.*

1273 *Let $I \subseteq \Sigma^*$ be a downward-closed set and $(J_m)_{m \in \mathcal{M}}$ upward-closed sets such that I
 1274 contains ε and not m_{err} , $J_{m_{\text{err}}} \cap I = \emptyset$, and $\mathcal{L}(I, (J_m)_{m \in \mathcal{M}}) \subseteq I$.*

23:34 Controller Synthesis for Broadcast Networks with Data

1275 Let B be the basis of I^c and B_m the basis of J_m for all m .

1276 If Environment wins the invariant game $\mathcal{IG}(\mathcal{G}, I, (J_m)_{m \in \mathcal{M}})$ then there is a σ -local run of
 1277 length at most $\varphi(|\mathcal{R}| + |B| + \|B\| + \sum_{m \in \mathcal{M}} |B_m| + \|B_m\|)$ satisfying 3i and 3ii and dissatisfying
 1278 either 3a or 3b.

1279 **Proof.** Let $N = |\mathcal{R}| + |B| + \|B\| + \sum_{m \in \mathcal{M}} |B_m| + \|B_m\|$. By Lemma 42 and Proposi-
 1280 tion 10 there exists $\tau_{\mathcal{IG}}$ a winning strategy $\tau_{\mathcal{IG}}$ for Environment in the invariant game
 1281 $\mathcal{IG}(\mathcal{G}, I, (J_m)_{m \in \mathcal{M}})$ such that Environment always wins in at most $\varphi(N)$ steps. We construct
 1282 a σ -local run of length at most $\varphi(N)$ satisfying 3i and 3ii and dissatisfying either 3a or
 1283 3b. To do so, we apply $\tau_{\mathcal{IG}}$ to choose transitions and we choose data by always picking a
 1284 datum never seen before in the run, when the datum is not determined by the transition.
 1285 Let (s_0, c_0) be an initial configuration of \mathcal{R} . We define iteratively a sequence of steps
 1286 $(s_{\ell-1}, c_{\ell-1}) \xrightarrow{\text{op}_\ell(m_\ell, d_\ell)}_{\delta_\ell} (s_\ell, c_\ell)$ as follows. Suppose we defined them up to $(s_{\ell-1}, c_{\ell-1})$, and
 1287 let $u_{\ell-1}$ be the local run defined so far. We first choose δ_ℓ :

- 1288 ■ If $s_{\ell-1} \in Q_{\text{ctrl}}$ then $\delta_\ell = \sigma(\delta_1 \cdots \delta_{\ell-1})$,
- 1289 ■ otherwise $\delta_\ell = \tau_{\mathcal{IG}}(\delta_1 \cdots \delta_{\ell-1})$.

1290 We then choose d_ℓ :

- 1291 ■ If δ_ℓ is a broadcast transition of letter m , we set d_ℓ as the initial datum of the local run.
- 1292 ■ If δ_ℓ is a record transition, we pick a datum d_k that does not appear in $u_{\ell-1}$.
- 1293 ■ If $\delta_\ell = s_{\ell-1} \xrightarrow{\text{rec}(m, i)} s_\ell$ is an equality transition of letter m , we set $d_\ell = c_{\ell-1}(i)$.

1294 Clearly we maintain the fact that u_ℓ is a σ -local run and $\delta_1 \cdots \delta_\ell$ is a $\tau_{\mathcal{IG}}$ -play in
 1295 $\mathcal{IG}(\mathcal{G}, I, (J_m)_{m \in \mathcal{M}})$. We stop when $\delta_1 \cdots \delta_\ell$ is winning for Environment in $\mathcal{IG}(\mathcal{G}, I, (J_m)_{m \in \mathcal{M}})$,
 1296 which happens for some $\ell \leq \varphi(N)$. Let M be the final value of ℓ and $u = u_M$ be the local
 1297 run obtained at the end. Let d be its initial datum. It remains to show that u satisfies 3i
 1298 and 3ii and dissatisfies either 3a or 3b. To do so, we rely on the following claim:

1299 \triangleright **Claim 49.** For all register i and index ℓ such that $i \notin \text{reg}(\delta_1 \cdots \delta_\ell)$, $\text{recentln}_i(\delta_1 \cdots \delta_\ell) =$
 1300 $\mathbf{In}_{u_\ell}(c_\ell(i))$. Furthermore, $\text{Out}(\delta_1 \cdots \delta_\ell) = \mathbf{Out}_d(u_\ell)$ and $\text{initln}(\delta_1 \cdots \delta_\ell) = \mathbf{In}_d(u_\ell)$.

1301 **Proof.** By a straightforward induction on ℓ . ◁

1302 Let $\pi_\ell = \delta_1 \cdots \delta_\ell$ for all ℓ , and let $\pi = \pi_M$.

1303 First we show that u satisfies 3i: As π is winning for Environment, it is compatible
 1304 with some decomposition $\text{dec} = (v_0, m_1, \dots, v_k) \in \mathcal{D}((J_m)_{m \in \mathcal{M}})$. Thus $\pi = \pi^0 \cdots \pi^k$ with
 1305 $v_j \sqsubseteq \text{Out}(\pi^j)$ and $\text{initln}(\pi^j) \in \{m_1, \dots, m_j\}^*$, for all j .

1306 We divide u like π , $u = u^0 \cdots u^k$. As a consequence of the claim, we obtain $v_j \sqsubseteq \mathbf{Out}_d(u^j)$
 1307 and $\mathbf{In}_d(u^j) \in \{m_1, \dots, m_j\}^*$, for all j . Thus u is compatible with dec .

1308 Now, we show that u satisfies 3ii. Let $d' \neq d$. If d' is stored in a register at some point
 1309 in u , let ℓ be the maximal index such that $c_\ell(i) = d'$ for some i . There can be no step
 1310 involving d' after ℓ , as d' would need to be stored in a register, contradicting the maximality
 1311 of ℓ . As a consequence, $\mathbf{In}_{d'}(u) = \text{recentln}_{\pi_\ell}(i)$. As π is winning for Environment, we have
 1312 $\text{recentln}_{\pi_\ell}(i) \in I$. If $\mathbf{In}_{d'}(u) = \varepsilon$ then clearly $\mathbf{In}_{d'}(u) \in I$ by assumption on I . If d' is never
 1313 stored in a register then $\mathbf{In}_{d'}(u) = \varepsilon \in I$.

1314 We have shown that u satisfies 3i and 3ii.

1315 If $\text{Out}(\pi) \notin I$ then $\mathbf{Out}_d(u) \notin I$, by the claim, hence u does not satisfy 3a.

1316 If $\text{Out}(\pi) \in I$, since π is winning for Environment, there must be an index ℓ such that
 1317 the ℓ th transition of π is a broadcast transition $\xrightarrow{\text{br}(m, i)}$, but $\text{recentln}_\pi(i) \notin J_m$. In that case,
 1318 we have $\mathbf{In}_{c_{\ell+1}(i)}(u_{\ell+1}) \notin J_m$ and $u_{\ell+1}$ contains a broadcast of $(m, c_{\ell+1}(i))$.

1319 As a consequence, we have found a prefix $u_{\ell+1}$ of u which does not satisfy 3b. As 3i and
1320 3ii hold for u , it is easy to see that they must also hold for all its prefixes.

1321 In all cases we have found a σ -local run of length at most $\varphi(N)$ which satisfies 3i and 3ii
1322 but dissatisfies either 3a or 3b.

1323 This concludes our proof. ◀

1324 ▶ **Lemma 50** (Bounding invariants). *There is an elementary function $\psi(N)$ such that the*
1325 *following statement holds.*

1326 *Let \mathcal{G} a BGR. There is a winning control strategy for \mathcal{G} if and only if there is a sequence*
1327 *of words $w_1, \dots, w_k \in \mathcal{M}^*$ and subsets $B, (B_m)_{m \in \mathcal{M}}$ of $\{w_1, \dots, w_k\}$ such that*

1328 ■ *B contains m_{err} and not ε and $B_{m_{err}} \uparrow \cap B^{\uparrow c} = \emptyset$*

1329 ■ *$\mathcal{L}(B^{\uparrow c}, (B_m \uparrow)_{m \in \mathcal{M}}, \subseteq) B^{\uparrow c}$*

1330 ■ *Controller wins $\mathcal{IG}(\mathcal{G}, B^{\uparrow c}, (B_m \uparrow)_{m \in \mathcal{M}})$,*

1331 ■ *B and all B_m are antichains for the subword order \sqsubseteq ,*

1332 ■ *$B \cup \bigcup_{m \in \mathcal{M}} B_m = \{w_1, \dots, w_k\}$,*

1333 ■ *for all $i \in [1, k]$, $|w_i| \leq \psi(|w_{i-1}|)$.*

1334 **Proof.** By Lemma 45, if there are such sets of words B and $(B_m)_{m \in \mathcal{M}}$, then there is a control
1335 strategy such that $(B^{\uparrow c}, (B_m \uparrow)_{m \in \mathcal{M}})$ is a sufficient invariant for σ . Hence, by Lemma 29, σ
1336 is a winning control strategy.

1337 Conversely, suppose there is a winning control strategy σ . By Lemma 29 there is a
1338 sufficient invariant $(I, (J_m)_{m \in \mathcal{M}})$ for σ . As I^c is upward-closed it has a finite basis B .
1339 Similarly, each J_m has a finite basis B_m .

1340 The first two conditions hold by definition, as $(I, (J_m))$ is a sufficient invariant.

1341 By Lemma 48 Controller wins $\mathcal{IG}(\mathcal{G}, I, (J_m)_{m \in \mathcal{M}})$, so the third condition of the lemma
1342 is satisfied.

1343 For the third condition, by definition, all basis are antichains.

1344 Let w_0, w_1, \dots, w_k be the elements of $B \cup \bigcup_{m \in \mathcal{M}} B_m$ sorted by length, i.e., $|w_i| \leq |w_{i+1}|$
1345 for all i . We can assume that we chose I and $(J_m)_{m \in \mathcal{M}}$ so that k would be minimal.

1346 By minimality of k , for all $j \in [1, k]$, $(B', (B'_m)_{m \in \mathcal{M}})$ is not a sufficient invariant for σ ,
1347 with $B' = B \cap \{w_i \mid i < j\}$ and for all m , $B'_m = B_m \cap \{w_i \mid i < j\}$. Let $I' = B'^{\uparrow c}$ and
1348 $J'_m = B'_m \uparrow$ for all m . Note that $I \subseteq I'$ while $J'_m \subseteq J_m$ for all m .

1349 A possibility is that $m_{err} \in I'$. As $m_{err} \in B$, we then have $|w_j| \leq 1$.

1350 Another possibility is that $I' \cap J'_{m_{err}} \neq \emptyset$. As a consequence, there is a word $w \in B'_{m_{err}}$
1351 with no subword in B' . As this word is of length at most $|w_{j-1}|$, we conclude that there is a
1352 word of length at most $|w_{j-1}|$ in $B \setminus B'$, hence $|w_j| \leq |w_{j-1}|$.

1353 Thirdly, we may have $\mathcal{L}(I', (J'_m)_{m \in \mathcal{M}}) \not\subseteq I'$. Then there is a decomposition $\text{dec} =$
1354 $(v_0, m_1, \dots, v_k) \in \mathcal{D}(I', (J'_m)_{m \in \mathcal{M}})$ and $w \in \mathcal{L}_{\text{dec}}$ such that $w \notin I'$.

1355 It is easy to construct deterministic automata recognising $\mathcal{L}(I', (J'_m)_{m \in \mathcal{M}})$ and I' of
1356 double-exponential size in $|\mathcal{R}|, |\mathcal{M}|, B'$ and $(B'_m)_{m \in \mathcal{M}}$, by using Lemma 8 and Claim 43.

1357 Hence we can find such a w of at most double-exponential size, and thus the decomposition
1358 $\text{dec} = (v_0, m_1, \dots, v_k)$ also has at most double-exponential size. Now note that $v_0 \cdots v_k$ is
1359 in I' , but cannot be in I : otherwise, we would have $w \in \mathcal{L}(I, (J'_m)_{m \in \mathcal{M}}) \subseteq \mathcal{L}(I, (J_m)_{m \in \mathcal{M}})$,
1360 while $w \notin I' \supseteq I$, a contradiction. Hence there is a word of at most double-exponential size
1361 in $|\mathcal{G}|, B'$ and $(B'_m)_{m \in \mathcal{M}}$ (thus of at most triple-exponential size in $|w_{j-1}|$) that is in B' but
1362 not B . As a consequence, $|w_j|$ is at most triply-exponential in $|w_{j-1}| + |\mathcal{M}| + |\mathcal{R}|$.

1363 The last case is when there is a run σ -local run which satisfies 3i and 3ii but dissatisfies
1364 either 3a or 3b, with respect to $(I', (J'_m)_{m \in \mathcal{M}})$. By Lemma 48, there is such a σ -local run u
1365 of length at most $K = \varphi(|\mathcal{R}| + \|B'\| + \sum_{m \in \mathcal{M}} \|B'_m\|) \leq \varphi(|\mathcal{R}| + (|\mathcal{M}| + 1)|w_{j-1}|)$.

1366 As $J'_m \subseteq J_m$ for all m , we have $\mathcal{D}((J'_m)_{m \in \mathcal{M}}) \subseteq \mathcal{D}((J_m)_{m \in \mathcal{M}})$. As a consequence, u
 1367 satisfies 3i with respect to $(I, (J_m)_{m \in \mathcal{M}})$.

1368 As $I \subseteq I'$, if u satisfies 3a with respect to $(I, (J_m)_{m \in \mathcal{M}})$ then it also satisfies it with
 1369 respect to $(I', (J'_m)_{m \in \mathcal{M}})$.

1370 Two cases remain: either u satisfies 3ii with respect to $(I', (J'_m)_{m \in \mathcal{M}})$ and not $(I, (J_m)_{m \in \mathcal{M}})$,
 1371 or satisfies 3b with respect to $(I, (J_m)_{m \in \mathcal{M}})$ and not $(I', (J'_m)_{m \in \mathcal{M}})$.

1372 We examine the two cases:

1373 ■ Suppose u satisfies 3ii with respect to $(I', (J'_m)_{m \in \mathcal{M}})$ and not $(I, (J_m)_{m \in \mathcal{M}})$. Let d' be a
 1374 datum such that $\mathbf{In}_{d'}(u) \notin I$. As $\mathbf{In}_{d'}(u) \in I'$, we found a word of length at most K that
 1375 is in I' but not I .

1376 ■ Suppose u satisfies 3b with respect to $(I, (J_m)_{m \in \mathcal{M}})$ and not $(I', (J'_m)_{m \in \mathcal{M}})$. Then there
 1377 exist $m \in \mathcal{M}$ and $d' \neq d$ such that u contains a broadcast of (m, d') and $\mathbf{In}_{d'}(u) \notin J'_m$,
 1378 while $\mathbf{In}_{d'}(u) \in J_m$. Furthermore, we have $|\mathbf{In}_{d'}(u)| \leq |u| \leq K$

1379 In both cases, there exists w_ℓ with $\ell \geq j$ such that $w_\ell \sqsubseteq w$, and thus $|w_\ell| \leq |w| \leq K$.

1380 As $|w_j| \leq |w_\ell|$, we have $|w_j| \leq K$. As $\|B'\| \leq |w_{j-1}|$ and $\|B'_m\| \leq |w_{j-1}|$, we obtain
 1381 $|w_j| \leq \varphi(|\mathcal{R}| + (|\mathcal{M}| + 1)|w_{j-1}|)$.

1382 We can then simply take a suitable elementary function so that $|w_{j-1}| \leq \psi(|\mathcal{R}| + |\mathcal{M}| +$
 1383 $|w_j|)$ ◀

1384 C.3 Main theorem

1385 ► **Theorem 30 (Main theorem).** *SAFESTRAT is decidable and \mathbf{F}_{ω} -complete.*

1386 **Proof.** It was shown in [13] that the coverability problem is \mathbf{F}_{ω} -hard. As coverability is the
 1387 particular case of SAFESTRAT where there are no controller nodes, this immediately yields
 1388 the same lower bound for SAFESTRAT.

1389 Let us now show the upper bound. Let \mathcal{G} a BGR. We once again apply the Length
 1390 Function Theorem.

1391 Consider a sequence of words $w_1, \dots, w_k \in \mathcal{M}^*$ and subsets $B, (B_m)_{m \in \mathcal{M}}$ of $\{w_1, \dots, w_k\}$
 1392 satisfying the conditions of Lemma 50.

1393 We use a fresh letter $\# \notin \mathcal{M}$. For each w_i we define $w'_i = \#^{|\mathcal{R}|+|\mathcal{M}|}w_i\#\#$ if $w_i \in B$, and
 1394 $w'_i = \#^{|\mathcal{R}|+|\mathcal{M}|}w_i\#m$ with m such that $w_i \in B_m$ otherwise.

1395 By the second condition of Lemma 50, w'_i is well-defined for all i .

1396 Note that the sequence $w'_1 \cdots w'_k$ is an antichain: as $\#$ does not appear in any w_i , $w'_i \sqsubseteq w'_j$
 1397 implies that $w_i \sqsubseteq w_j$, and that they both belong to B or to some common B_m . This is
 1398 impossible as all those sets are antichains.

1399 Furthermore, for all i , we have $|w'_{i+1}| \leq \psi(|\mathcal{R}| + |\mathcal{M}| + |w_i|) + |\mathcal{R}| + |\mathcal{M}| + 2 \leq \psi(|w'_i|) + |w'_i|$.
 1400 As $g : n \mapsto \psi(n) + n$ is a primitive recursive function, by the Length function theorem we obtain
 1401 a function $f \in \mathcal{F}_{\omega, |\mathcal{M}|}$ such that every (g, n) -controlled bad sequence of words w_0, w_1, \dots, w_k
 1402 has at most $f(n)$ terms.

1403 As m_{err} is in B , $|w_0| \leq 1$, thus $|w'_0| \leq |\mathcal{R}| + |\mathcal{M}| + 3$. We therefore have $|w_i| \leq g^{(i)}(|\mathcal{R}| +$
 1404 $|\mathcal{M}| + 3)$ for all i . As a consequence, we have $k \leq f(|\mathcal{R}| + |\mathcal{M}| + 3)$.

1405 Our algorithm guesses a sequence of words of sorted by length w_1, \dots, w_k with $k \leq$
 1406 $f(|\mathcal{R}| + |\mathcal{M}| + 3)$ such that $|w_{i+1}| \leq \psi(|w_i|)$ for all i . The algorithm then guesses subsets B
 1407 and $(B_m)_{m \in \mathcal{M}}$ that cover $\{w_i \mid i \in [1, k]\}$.

1408 It checks that Controller wins $\mathcal{IG}(\mathcal{G}, B \uparrow^c, (B_m \uparrow)_{m \in \mathcal{M}})$. We accept if she does and reject
 1409 otherwise.

1410 This can be done in double-exponential time in $|\mathcal{R}| + k + |w_k|$, by Lemma 42. We can
 1411 make this algorithm deterministic with an exponential blow-up in the time complexity. By
 1412 Lemma 50, this algorithm is correct.

1413 The time required by this algorithm is therefore $h(f(|\mathcal{R}| + |\mathcal{M}| + 3))$ with h a primitive
 1414 recursive function. As $\mathcal{F}_{\omega|\mathcal{M}|}$ is closed under composition with primitive recursive functions,
 1415 the algorithm takes a time bounded by a function of $\mathcal{F}_{\omega|\mathcal{M}|}$. As a consequence, the problem
 1416 is in $\mathbf{F}_{\omega\omega}$. ◀

1417 **D** Missing proofs from Section 6

1418 We say that a local run has *organised data* if

- 1419 ■ if whenever a datum is received for the first time, it is greater than the initial datum and
- 1420 all data received previously.
- 1421 ■ Each datum is recorded at most once in the registers.

1422 ► **Proposition 51.** *There is a function $h : \mathbb{N} \rightarrow \mathbb{N}$ such that for all BGR \mathcal{G} , if a control*
 1423 *strategy is losing then there exists a σ -run ρ in which every local run has length at most*
 1424 *$h(|\mathcal{G}|)$ and has organised data in which m_{err} is broadcast.*

1425 **Proof.** Let us start by defining an *execution tree* as a tree of the following form:

- 1426 ■ There are two types of nodes, *word nodes* and *run nodes*
- 1427 ■ The children of a word node are run nodes, and the children of a run node are word
- 1428 nodes.
- 1429 ■ For all run node ν with a label u and all $d \in \mathbb{D}$ such that $\mathbf{In}_d(u) \neq \varepsilon$, ν has a child with
- 1430 a label w such that $\mathbf{In}_d(u) \sqsubseteq w$.
- 1431 ■ For all word node ν labelled w , for all child ν' of ν labelled u , $w \sqsubseteq \mathbf{Out}_{sign}(u)$

1432 Consider the following algorithm: We start with an *execution tree* made only of a root
 1433 labelled m_{err} . We maintain a set of word nodes O , initially containing only the root. The
 1434 word nodes in O are called *open*, others are called *closed*

1435 While O is not empty, we apply the following steps:

- 1436 ■ If there is a *run node* ν whose children are all *closed*, let ν' be its parent, labelled w . We
- 1437 remove every node that was added to the tree after ν' (in particular, we remove all of its
- 1438 descendants). Then, we remove ν from O .
- 1439 ■ Otherwise, let B be the set of labels of open nodes, we define $I = B^{\uparrow c}$. By Lemma 21, there
- 1440 exists a σ -local run u of length at most $\varphi(\mathcal{R}, B)$ such that $\mathbf{Out}_{sign}(u) \notin I$, $\mathbf{In}_d(u) \in I$
- 1441 for all d and no datum is recorded twice. Let ν be an open node with a label w such that
- 1442 $w \sqsubseteq \mathbf{Out}_{sign}(u)$. It exists by definition of B and I . We add a child ν' to ν , labelled by u .
- 1443 Then, consider the set $\{\mathbf{In}_d(u) \mid d \in \mathbb{D}\} \setminus \{\varepsilon\}$, let B_u be its set of minimal elements for
- 1444 \sqsubseteq . For each $v \in B_u$ we add a child labelled v to ν' , and we add all those children to O .

1445 Note that when we remove a node we remove all nodes added after that one. As a
 1446 consequence, at all times we can enumerate open nodes $O = \{\nu_0, \nu_1, \dots, \nu_k\}$ in their order
 1447 of appearance, and we obtain $|w_i| \leq |u_i| \leq \varphi(\mathcal{R}, \{w_1, \dots, w_{i-1}\})$ for all i , with w_1, \dots, w_k
 1448 the labels of ν_1, \dots, ν_k and u_1, \dots, u_k the labels of their respective parents. Additionally, we
 1449 maintain the fact that the sequence w_1, \dots, w_k is a *bad sequence*. We can then apply the
 1450 *Length Function Theorem* to bound k by $f(|\mathcal{R}|)$ with f a function of $\mathbf{F}_{\omega\omega}$.

1451 We also obtain a bound $h(|\mathcal{R}|)$ on the length of run node labels. As a consequence, the
 1452 number of data appearing in each local run is bounded by that same bound, and thus the
 1453 degree of the tree is at most $h(|\mathcal{R}|)$.

1454 As every node is a leaf or an open node or the child of an open node, we get a bound
 1455 $b(|\mathcal{R}|)$ on the size of the tree. As a consequence, the set of trees we see is finite. In order to
 1456 show that the algorithm terminates, we simply have to show that we cannot loop.

1457 Given the tree at some point of the algorithm, let ν_0, \dots, ν_k be the set of word nodes in
 1458 their order of creation. For each i , let x_i be 0 if ν_i is open and 1 if it is closed. It is easy to
 1459 check that the sequence $x_0 \cdots x_k$ increases at each step for the lexicographic ordering. As
 1460 a result, we never see the same tree twice. The algorithm therefore terminates in at most
 1461 $c(|\mathcal{R}|)$ steps.

1462 \triangleright **Claim 52.** After each iteration, for all closed node ν labelled w , ν is a leaf and there is
 1463 a σ -run in which every local run has organised data and has length at most $h(|\mathcal{R}|)$ and in
 1464 which the sequence w is broadcast by some agent.

1465 *Proof.* We proceed by induction on the number of iterations. This property is clearly true at
 1466 the beginning as there are no closed nodes.

1467 For the induction step, note that we never add children to closed nodes and only turn
 1468 leaves into closed nodes. Hence we maintain the fact that every closed node is a leaf.

1469 Furthermore, say we turn an open node labelled w into a closed one. We do so when
 1470 it has a child ν' whose children are all closed. Let u be the label of ν' and w_1, \dots, w_n the
 1471 labels of its children. By induction hypothesis, for each i we have a σ -run ϱ_i in which each
 1472 local run has organised data and length at most $h(|\mathcal{R}|)$ and in which an agent broadcasts w_i .

1473 For each d received in u , we know that there is an i such that $\mathbf{In}_d(u) \sqsubseteq w_i$. We define
 1474 ϱ_d as ϱ_i where data have been renamed so that w_i is broadcast with datum d and all other
 1475 data are fresh and do not appear in u .

1476 Let d_1, \dots, d_m be the data received in u , in order of appearance. Let d_0 be the initial
 1477 datum of u . For all $j \in [2, m]$, in increasing order, let A_{j-1} be such that all data appearing
 1478 in ϱ_{j-1} are below A_{j-1} . Let $A_0 = d_0 + 1$. Define ϱ'_j as ϱ_{d_j} where each datum d' has been
 1479 renamed into $d' + A_{j-1}$. The runs ϱ'_j use disjoint sets of data and in each ϱ'_j an agent
 1480 broadcasts $\mathbf{In}_{d'_j}(u)$ with datum $d'_j = d_j + A_{j-1}$. In particular we have $d_0 < d'_1 < \dots < d'_m$.
 1481 We have also maintained the fact that all local runs in those runs have organised data.

1482 We rename each datum d_j with in u to d'_j and obtain a local run u' with organised data
 1483 and $\mathbf{In}_{d'_j}(u) = \mathbf{In}_{d'_j}(u')$. We can execute all runs ϱ'_j and u over disjoint sets of agents, and
 1484 use the broadcasts in each ϱ'_j to match the receptions in u . This gives us a σ -run ϱ_i in
 1485 which each local run has increasing data and length at most $h(|\mathcal{R}|)$ and in which an agent
 1486 broadcasts w (by executing u). \triangleleft

1487 As the algorithm terminates, eventually the root is closed. By the claim above, we have a
 1488 σ -run ϱ in which each local run has length at most $h(|\mathcal{R}|)$ and in which an agent broadcasts
 1489 m_{err} . \blacktriangleleft

1490 We recall Ramsey's theorem on infinite hypergraphs. Given a set S and $k \in \mathbb{N}$, we use
 1491 the notation $\binom{S}{k}$ for the set of subsets of S of size k .

1492 \blacktriangleright **Theorem 53** (Ramsey's theorem on infinite hypergraphs). *Let V be an infinite set of vertices*
 1493 *and $k \in \mathbb{N}$. Let $\mathbf{col} : \binom{V}{k} \rightarrow C$ with C a finite set of colours. Then there exists an infinite*
 1494 *subset $V' \subseteq V$ and $c \in C$ such that $\mathbf{col}(\binom{V'}{k}) = \{c\}$.*

1495 \blacktriangleright **Theorem 31.** *There is a winning data-aware control strategy for \mathcal{G} if and only if there is*
 1496 *a winning control strategy for \mathcal{G} .*

1497 **Proof.** The right-to-left direction is clear.

1498 For the left-to-right direction, suppose there is a winning data-aware control strategy σ
 1499 for \mathcal{G} . Let $K = h(|\mathcal{R}|)$ with h as defined in Proposition 51. Let \mathbf{R}_K be the set of σ -local runs
 1500 with organised data of length at most K .

1501 We define a function $\mathbf{col} : \binom{\mathbb{D}}{K+1} \rightarrow 2^{\mathbf{R}_K}$ as follows. Let D be a set of $R + 1$ data. Let
 1502 d_0, \dots, d_R be the elements of \mathbb{D} in increasing order. Then $\mathbf{col}(D)$ is the set of σ -local runs
 1503 with organised data of length at most R such that the initial datum is d_0 the other data
 1504 appearing in the run are d_1, \dots, d_k for some k . With the organised data property and those
 1505 conditions, the local run is fully determined by its sequence of transitions. As a result,
 1506 $|\mathbf{col}(D)| \leq |\Delta|^K$.

1507 In a local run with at most B steps, at most $B + 1$ data appear. As a result, every
 1508 element of \mathbf{R}_K has an antecedent by \mathbf{col} . We can now apply Theorem 53. We obtain an
 1509 infinite set $\mathbb{D}' \subseteq \mathbb{D}$ of data and a set of local runs R such that $\mathbf{col}(\binom{\mathbb{D}'}{K+1}) = \{R\}$.

1510 Let $d_0, \dots, d_K \in \mathbb{D}'$ with $d_0 < \dots < d_K$. Define the strategy $\sigma' : \Delta^* \rightarrow \Delta$ which, given
 1511 a sequence of transitions, takes the same decision as σ over the unique local run with that
 1512 sequence of transitions, organised data, and using data $\{d_0, \dots, d_k\}$ for some k , with d_0 the
 1513 initial datum.

1514 If σ' was losing, we would have a run in which every local run has length at most K and
 1515 organised data in which m_{err} is broadcast. This run is, however, also a σ -run, which is a
 1516 contradiction. As a result, σ' is winning. ◀

1517 **E** Missing proofs from Section 7

1518 In this section we prove the following result.

1519 ▶ **Theorem 33.** *SAFESTRAT is NEXPTIME-complete for 1BGR.*

1520 We start with the upper bound.

1521 ▶ **Proposition 54.** *SAFESTRAT is in NEXPTIME on 1BGR.*

1522 For the rest of this section we fix a 1BGR $\mathcal{G} = (\mathcal{R}, Q_{ctrl}, Q_{env}, m_{err})$.

1523 We will use the following criterion for the existence of positional strategies.

1524 ▶ **Proposition 55** ([17]). *If an objective is submixing then player P_0 has a positional optimal*
 1525 *strategy in all games with this objective.*

1526 The *output game* is played on \mathcal{R} , with players picking transitions from their respective
 1527 states. It has two parameters: an invariant $(I, (J_m)_{m \in \mathcal{M}})$, and a set of record transitions
 1528 $T \subseteq \Delta$. We will use the term $(I, (J_m)_{m \in \mathcal{M}}, T)$ -*output game* for the output game with those
 1529 parameters.

1530 The winning condition is defined as follows:

- 1531 **(O1)** If at some point the play is not compatible with any decomposition of $\mathcal{D}((J_m)_{m \in \mathcal{M}})$,
 1532 then Controller wins.
- 1533 **(O2)** If the output of the play is not in I then Environment wins.
- 1534 **(O3)** If we reach a record transition then Controller wins if it is in T and Environment wins
 1535 otherwise.
- 1536 **(O4)** If the play goes on forever without any of the previous things happening then Controller
 1537 wins.

1538 ► **Lemma 56.** *If Controller wins an output game then she has a positional winning strategy.*

1539 The *echo game* is also played on \mathcal{R} , with players picking transitions from their respective
 1540 states. It has as parameters an invariant $(I, (J_m)_{m \in \mathcal{M}})$, a set of record transitions $T \subseteq \Delta$ and
 1541 a record transition $t = q \xrightarrow{\text{rec}(m, \downarrow 1)} q'$. The play starts by taking transition t , and continues
 1542 from q' .

- 1543 (E1) If at some point the recent input on 1 is not in I , Controller wins.
 1544 (E2) If at some point we make a broadcast with letter m while the recent input on 1 is not
 1545 in J_m , then Environment wins.
 1546 (E3) If we reach a record transition the game stops: If that transition is in T then Controller
 1547 wins, otherwise Environment does.
 1548 (E4) If the play goes on forever without any of those things happening then Controller wins.

1549 ► **Lemma 57.** *If Controller wins an output game then she has a positional winning strategy.*

1550 **Proof.** We show that Controller's objective in an output game is submixing. This proves the
 1551 lemma by applying [17, Theorem 4.5].

1552 Consider two losing plays for Controller π and π' , and a third play $\bar{\pi} = \pi_0 \pi'_0 \pi_1 \dots$
 1553 obtained by shuffling the two. We show that $\bar{\pi}$ is also losing for Controller. As π and π' are
 1554 losing for Controller, we can consider them as finite: the victory of Environment is witnessed
 1555 by a finite prefix. We can cut π and π' into $\pi = \pi_0 \pi_1 \dots \pi_m$ and $\pi' = \pi'_0 \pi'_1 \dots \pi'_m$ so that
 1556 $\bar{\pi} = \pi_0 \pi'_0 \pi_1 \dots \pi_m \pi'_m$ (note that π'_m can be empty).

1557 Clearly no transition of T is seen in π or π' , thus not in $\bar{\pi}$ as well. Let $\tilde{\pi}$ a prefix of $\bar{\pi}$, we
 1558 show that it is compatible with some decomposition of $\mathcal{D}((J_m)_{m \in \mathcal{M}})$. Let $\tilde{m}_1, \dots, \tilde{m}_k$ be
 1559 the set of letters received along $\tilde{\pi}$, in that order. Let $\tilde{\pi} = \tilde{\pi}_0 \dots \tilde{\pi}_k$ so that for each i the first
 1560 step of $\tilde{\pi}_i$ is the first reception of m_i . Let \tilde{v}_i be the sequence of letters broadcast in $\tilde{\pi}_i$, for
 1561 all i . Let $\tilde{\text{dec}} = (\tilde{v}_0, \tilde{m}_1, \dots, \tilde{v}_k)$. Clearly $\tilde{\pi}$ is compatible with $\tilde{\text{dec}}$.

1562 It remains to show that $\tilde{\text{dec}} \in \mathcal{D}((J_m)_{m \in \mathcal{M}})$. Let $i \in [1, k]$, we need to find a word in
 1563 $\mathcal{L}_{(\tilde{v}_0, \tilde{m}_1, \dots, \tilde{v}_{i-1})} \cap J_{m_i}$. For that, we observe that the reception of \tilde{m}_i happens in either a segment
 1564 from π or from π' . We assume that it is from π , the other case is symmetric. Since every
 1565 prefix of π is compatible with some decomposition of $\mathcal{D}((J_m)_{m \in \mathcal{M}})$, in particular the prefix of
 1566 π up to that reception of \tilde{m}_i is compatible with one. Thus there exists $\text{dec} = (v_0, m_1, \dots, v_\ell)$
 1567 such that $\mathcal{L}_{\text{dec}} \cap J_{\tilde{m}_i} \neq \emptyset$ and with which π is compatible. Let $w \in \mathcal{L}_{\text{dec}} \cap J_{\tilde{m}_i}$, w can be
 1568 obtained from $v_0 \dots v_\ell$ by adding letters from $\{m_1, \dots, m_j\}$ to each v_j .

1569 As this prefix of π is fully contained in $\tilde{\pi}$, we can find the same sequence of broadcast
 1570 $v_0 \dots v_\ell$ in $\tilde{\pi}$. Moreover, for each j , the first reception of m_j can only be earlier in $\tilde{\pi}$ than in
 1571 π , hence dec allows us to find $v_0 \dots v_\ell$ and to add the same letters at the same places. As a
 1572 consequence, $w \in \mathcal{L}_{(\tilde{v}_0, \tilde{m}_1, \dots, \tilde{v}_{i-1})}$.

1573 It follows that every prefix of $\bar{\pi}$ is compatible with some decomposition of $\mathcal{D}((J_m)_{m \in \mathcal{M}})$.
 1574 As a consequence, Controller does not win at any point in $\bar{\pi}$.

1575 If some record transition outside of T is seen in π or π' then in $\bar{\pi}$ as well. Otherwise, it
 1576 means the output of some prefix of π is not in I . As the output of that prefix must be a
 1577 subword of the output of $\bar{\pi}$, and I is downward-closed, we obtain that the output of $\bar{\pi}$ is not
 1578 in I .

1579 In conclusion, Controller does not win at any point in $\bar{\pi}$ while Environment does. As
 1580 a consequence, Controller's objective is submixing and thus if Controller wins she can win
 1581 with a positional strategy. ◀

1582 ► **Lemma 58.** *If Environment wins an echo game then he has a positional winning strategy.*

1583 **Proof.** We show that Environment's objective in an echo game has the submixing property,
1584 and again apply [17, Theorem 4.5].

1585 Consider two losing plays for Environment π and π' , and $\bar{\pi}$ a submixing of the two. We
1586 can cut π and π' into $\pi = \pi_0\pi_1 \cdots$ and $\pi' = \pi'_0\pi'_1 \cdots$ so that $\bar{\pi} = \pi_0\pi'_0\pi_1 \cdots$.

1587 At all times in $\bar{\pi}$ if we make a broadcast with letter m while the recent input is w ,
1588 then that broadcast was made in π or π' with a recent input that is a subword of w . As
1589 Environment loses in π and π' , and as J_m is upward-closed, $w \in J_m$.

1590 Every record transition seen in $\bar{\pi}$ must be seen in π or π' , hence must be in T .

1591 As a consequence, Environment cannot win $\bar{\pi}$, hence Controller wins. The objective
1592 of Environment is therefore submixing, and thus if Environment wins he can win with a
1593 positional strategy. ◀

1594 E.1 Characterisation of winning strategies

1595 ► **Lemma 59.** *Controller wins the $(I, (J_m)_{m \in \mathcal{M}})$ -invariant game if and only if there is a set
1596 of record transitions T such that she wins the $(I, (J_m), T)$ -output game and the $(I, (J_m), T, t)$ -
1597 echo game for all $t \in T$.*

1598 **Proof.** Suppose Controller wins the $(I, (J_m)_{m \in \mathcal{M}})$ -invariant game with a strategy σ . Let T
1599 be the set of record transitions taken in a σ -play in which no player has won yet.

1600 We start with the $(I, (J_m), T)$ -output game : let Controller apply the same strategy σ in
1601 that game.

1602 ▷ **Claim 60.** Let π be a play such that no record transition has been seen yet.

1603 Then π is winning for a player in the invariant game if and only if it is winning for that
1604 player in the output game.

1605 **Proof.** Take a look at the winning conditions in the invariant game. Condition A and C are
1606 the same as 1 and 2. Condition B and D cannot happen: as we have not seen any record
1607 transition, $\text{reg}(\pi') = \{1\}$ for all prefixes π' of π . ◀

1608 As a consequence, a σ -play can only be winning for Environment if we reach a record
1609 transition $t \notin T$ while Controller has not won. However, this means that the play obtained
1610 before reaching t is not winning for Controller in the invariant game either, by the previous
1611 claim. This contradicts the definition of T . Hence σ is winning for Controller in the
1612 $(I, (J_m), T)$ -output game.

1613 Let $t \in T$. We now show that we have a winning strategy for Controller in the
1614 $(I, (J_m), T, t)$ -echo game.

1615 ▷ **Claim 61.** Let $t\pi_+$ be a play in the $(I, (J_m), T, t)$ -echo game such that no record transition
1616 has been seen yet (apart from the first step). Let π_-t be a play ending with t in the
1617 $(I, (J_m))$ -invariant game such that no player wins in it.

1618 Then π_+ is winning for a player in the echo game if and only if $\pi_-t\pi_+$ is winning for that
1619 player in the output game.

1620 **Proof.** First of all note that $\text{reg}(\pi_-t) = \emptyset$ as t updates the only register. As π_-t is not
1621 winning for either player, no prefix of it fulfils either A or C. We can then conclude that
1622 there is no play starting with π_-t that fulfils either of those conditions.

1623 Furthermore, since no player wins in π_- and t updates the only register, conditions B, D
1624 are satisfied by $\pi_-t\pi_+$ if and only if they are satisfied by $t\pi_+$ if and only if $t\pi_+$ satisfies 1, 2
1625 respectively.

1626 This proves the claim. ◁

1627 By definition of t there exists a play reaching t in the invariant game in which no player
1628 has won yet. Let π_- be the prefix of that play before reaching t . We define the strategy σ_E
1629 as $\sigma_E(\pi) = \sigma(\pi_- \pi)$.

1630 Let us consider a σ_E -play $t\pi_+$ in the echo game and show that it cannot be winning for
1631 Environment.

1632 By the claim above, a σ_E -play can only be winning for Environment if we reach a record
1633 transition $t' \notin T$ while Controller has not won. However, this means that the play π obtained
1634 before reaching t' is such that $\pi_- \pi$ is not winning for Controller in the invariant game either,
1635 by the previous claim. This contradicts the definition of T .

1636 We have established that a winning strategy in the $(I, (J_m)_{m \in \mathcal{M}})$ -invariant game yields
1637 a set of record transitions T and winning strategies in the $(I, (J_m), T)$ -output game and the
1638 $(I, (J_m), T, t)$ -echo game for all $t \in T$.

1639 For the reverse direction, let us consider a set of record transitions T , σ_O a winning
1640 strategy in the $(I, (J_m), T)$ -output game and, for all $t \in T$, σ_t a winning strategy in the
1641 $(I, (J_m), T, t)$ -echo game.

1642 We define a strategy σ in the invariant game as follows: If π does not contain any record
1643 transition then $\sigma(\pi) = \sigma_O(\pi)$. Otherwise, let π' be the largest suffix of π with no record
1644 transition and t the record transition just before π' . We set $\sigma(\pi) = \sigma_t(t\pi')$.

1645 It remains to show that σ is a winning strategy in the invariant game. Suppose by
1646 contradiction that there exists a finite σ -play winning for Environment. Let π be such a
1647 σ -play of minimal size. If π contains no record transition then by the first claim it is also
1648 winning for Environment in the output game. As σ mimics σ_O while no record transition
1649 has been seen, this is a contradiction with the fact that σ_O is winning.

1650 On the other hand, if π contains a record transition, then we can decompose it as
1651 $\pi = \pi_- t\pi_+$ with t a record transition and π_+ the maximal suffix of π with no record
1652 transition.

1653 Then by minimality of π , no player wins in π_- . As a result, by the second claim, $t\pi_+$
1654 is winning for Environment in the $(I, (J_m), T, t)$ -echo game. This is a contradiction as by
1655 definition of π , $t\pi_+$ is a σ_t -play, and σ_t is winning for Controller.

1656 This concludes our proof. ◀

1657 **► Lemma 62.** *If there exists a winning control strategy for a BGR then there exist I such
1658 that every word in the basis of I^c is of length $\leq |\mathcal{R}|(|\mathcal{M}| + 1)$ and $(J_m)_{m \in \mathcal{M}}$ in which every
1659 word in the basis has length $\leq |\mathcal{R}|$ for all m and T a set of record transitions such that
1660 Controller wins the $I, (J_m), T$ -output game and the $I, (J_m), T, t$ -echo game for all $t \in T$.*

1661 **Proof.** Suppose there exists a winning control strategy σ , then we have some $I, (J_m)$ such
1662 that Controller wins $I, (J_m), T$ -output game and the $I, (J_m), T, t$ -echo game for all $t \in T$.
1663 We can assume that the sum of the lengths in the basis of the J_m is minimal.

1664 We remove a word w from the basis of J_m . By minimality of J_m the resulting invariant
1665 is not sufficient. Hence Environment wins one of the games. Since I has not changed but
1666 (J_m) has decreased, Controller still wins the output game. As a consequence, Environment
1667 wins the $I, (J_m), T, t$ -echo game for some $t \in T$. Let σ_{echo} be a positional winning strategy
1668 for Environment in the new instance of that game. There must be a σ_{echo} -play that is losing
1669 for him in the previous instance. As we have decreased $\mathcal{L}(I, (J_m)) \downarrow$, the only possibility is
1670 that there is a play in which we broadcast m while the recent input is not in J_m . As J_m
1671 is upward-closed and σ_{echo} is positional, we can cut all cycles from this play: We obtain

1672 a σ_{echo} -play whose recent input is not in J_m , of length at most $|\mathcal{R}|$. As a consequence,
 1673 $|w| \leq |\mathcal{R}|$.

1674 We have shown that all words in the basis of all J_m have length at most $|\mathcal{R}|$. Let us now
 1675 bound the words in the basis of I^c .

1676 Consider I with a basis of minimal size such that $I, (J_m)$ is a sufficient invariant for σ . We
 1677 remove a word w from the basis of I^c , thus increasing I . By minimality of I^c , Environment
 1678 wins one of the games. It cannot be the output game as I has increased and the J_m are the
 1679 same. If it is an echo game then let σ_{echo} be a positional winning strategy for Environment
 1680 in the new instance of that game. There must be a play whose recent input was previously
 1681 out of $\mathcal{L}(I, (J_m))\downarrow$ but is now in it. We can once again cut cycles on that play. Once we do
 1682 so, we obtain a play of length $\leq |\mathcal{R}|$ whose recent input w_{in} is in $\mathcal{L}(I, (J_m))\downarrow$ but was not
 1683 previously. As a consequence, there exists $\text{dec} = (v_0, m_1, \dots, v_k)$ such that $w_{in} \in \mathcal{L}_{\text{dec}}\downarrow$ and
 1684 for all i , $\mathcal{L}_{(v_0, m_1, \dots, v_{i-1})} \cap J_{m_i} \neq \emptyset$. As we have bounded the lengths of words in the basis
 1685 of each J_m by $|\mathcal{R}|$, there exist u_1, \dots, u_k , all of size at most $|\mathcal{R}|$, such that $u_i \in J_{m_i}$ and
 1686 $u_i \in \mathcal{L}_{(v_0, m_1, \dots, v_{i-1})}\downarrow$.

1687 For each u_i and for w_{in} at most $|\mathcal{R}|$ letters from v_0, \dots, v_k suffice to maintain these
 1688 properties. We define v'_0, \dots, v'_k as the words obtained by removing all other letters. Let
 1689 $\text{dec}' = (v'_0, m_1, \dots, v'_k)$. We therefore have $|v'_0 \cdots v'_k| \leq |\mathcal{R}|(|\mathcal{M}| + 1)$, and $(v'_0, m_1, \dots, v'_k) \in$
 1690 $\mathcal{D}(I, (J_m))$ and $w_{in} \in \mathcal{L}_{\text{dec}'}\downarrow$.

1691 As a consequence, we must have that $v'_0 \cdots v'_k$ is in I but was previously not. As a result,
 1692 $w \sqsubseteq v'_0 \cdots v'_k$ and thus $|w| \leq |\mathcal{R}|(|\mathcal{M}| + 1)$. ◀

1693 ▶ **Theorem 33.** *SAFESTRAT is NEXPTIME-complete for 1BGR.*

1694 **Proof.** We guess a positional strategy σ for Controller in the output game. We also guess a
 1695 set of record transitions T , a set of words B of length $\leq |\mathcal{R}|(|\mathcal{M}| + 1)$ and a family of sets of
 1696 words $(B_m)_{m \in \mathcal{M}}$, where all words have length at most $|\mathcal{R}|$.

1697 We then try to check if Environment has a winning strategy in one of the games. For the
 1698 output game, we enumerate all positional strategies for Controller. As the size of words in
 1699 the basis of I is bounded by $|\mathcal{R}|(|\mathcal{M}| + 1)$, if such a strategy allows a losing play, it allows
 1700 one of length at most $|\mathcal{R}|(|\mathcal{R}| + 1)(|\mathcal{M}| + 1)$. As a consequence, we can check in exponential
 1701 time whether one of those strategies is winning.

1702 For the echo games, we enumerate all positional strategies for Environment.

1703 ▷ **Claim 63.** We can check that a positional strategy σ_{echo} is not winning for Environment
 1704 in an echo game in non-deterministic exponential time.

1705 **Proof.** Let π be a play won by Controller, at all times if we make a broadcast with letter m ,
 1706 the recent input on 1 is in J_m . For each m broadcast in the play, we can select a sequence
 1707 of at most $|\mathcal{R}|$ preceding receptions forming an element of the basis of J_m . Those elements
 1708 witness the fact that the recent input is in J_m .

1709 We can thus easily construct an NFA of exponential size recognising finite plays in which
 1710 Environment does not win.

1711 Since σ_{echo} is positional, there is an automaton with $|\mathcal{R}|$ states recognising the set of
 1712 σ_{echo} -plays. We first check whether there is an infinite word whose prefixes are all accepted
 1713 by the NFA. If not, we check whether the NFA accepts a play ending with a transition of T .

1714 Finally, we project it to obtain an NFA \mathcal{A} recognising the recent inputs of σ_{echo} -plays not
 1715 won by Environment. We also build an exponential-size NFA \mathcal{B} recognising $\mathcal{L}(I, (J_m)_{m \in \mathcal{M}})$.
 1716 As shown in [1], if there is a word in $\mathcal{L}(\mathcal{A})\downarrow$ that is not in $\mathcal{L}(\mathcal{B})\downarrow$, then there is one of
 1717 polynomial size in $|\mathcal{A}|$ and $|\mathcal{B}|$.

1718 As a consequence, we can check that non-inclusion in non-deterministic exponential time.
 1719 In sum, we can check in non-deterministic exponential time that the given strategy is not
 1720 winning for Environment. \triangleleft

1721 This lets us decide in non-deterministic exponential time if there exist $I, (J_m), T$ such
 1722 that Controller wins the output game and all the echo games. As a result, SAFESTRAT
 1723 is in NEXPTIME. All that is left to do is show the matching lower bound, which is done
 1724 below. \blacktriangleleft

1725 E.2 NEXPTIME-hardness of SafeStrat for 1BGR

1726 The *exponential grid tiling problem* asks, given a set of colours C , a number N in unary and
 1727 a set of tiles $T \subseteq C^{\{\text{up}, \text{down}, \text{left}, \text{right}\}}$, whether there is a tiling of the $2^N \times 2^N$ grid, i.e., a
 1728 function $\tau : [0, 2^N - 1] \times [0, 2^N - 1] \rightarrow T$ such that for all $x, y, x', y' \in [0, 2^N - 1]$,

- 1729 ■ if $x = x'$ and $y = y' + 1$ then $\tau(x, y).down = \tau(x, y').up$
- 1730 ■ if $x = x' + 1$ and $y = y'$ then $\tau(x, y).left = \tau(x', y).right$
- 1731 ■ if $x = 0$ (resp. $x = 2^N - 1, y = 0, y = 2^N - 1$) then $\tau(x, y).left = c_{border}$ (resp.
 1732 $right, down, up$)

1733 This problem is NEXPTIME-complete [27].

1734 ► **Lemma 64.** *SAFESTRAT is NEXPTIME-hard on 1BGR.*

1735 **Proof.** We reduce from the exponential grid tiling problem.

1736 Let C be a set of colours containing a border colour B , let $T = \{t_1, \dots, t_k\}$ be a set of
 1737 tiles and N an integer in unary. We use the alphabet of letters $\mathcal{M} = \{0, 1, \bar{0}, \bar{1}\} \cup T \cup \bar{T}$,
 1738 where $\bar{T} = \{\bar{t}_1, \dots, \bar{t}_k\}$ is a copy of T .

1739 We design a 1BGR in which Controller wins if and only if there is a valid tiling of the
 1740 $2^N \times 2^N$ grid with those tiles.

1741 Essentially, Environment may use some agents to broadcast coordinates (x, y) and (\bar{x}, \bar{y})
 1742 in the grid, respectively using letters $\{0, 1\}$ and $\{\bar{0}, \bar{1}\}$. Environment can also make an agent
 1743 receive coordinates (x, y) (resp. (\bar{x}, \bar{y})), while checking that they all have the same datum. He
 1744 then makes Controller choose a tile t (resp. \bar{t}), which is broadcast with that same identifier.
 1745 A strategy for Controller amounts to two functions $\tau, \bar{\tau} : [0, 2^N - 1] \times [0, 2^N - 1] \rightarrow T$.

1746 The agents that broadcast coordinates (x, y) and (\bar{x}, \bar{y}) can then receive tiles t and \bar{t} with
 1747 their own identifier, and check that:

- 1748 ■ If $x = \bar{x}$ and $y = \bar{y}$ then $t = \bar{t}$
- 1749 ■ If $x + 1 = \bar{x}$ and $y = \bar{y}$ then $t.right = \bar{t}.left$
- 1750 ■ If $x = \bar{x}$ and $y + 1 = \bar{y}$ then $t.up = \bar{t}.down$
- 1751 ■ If $x = 0$ (resp. $y = 0, x = 2^N - 1, y = 2^N - 1$) then $t.left = B$ (resp. $down, left, right$).

1752 The first item forces Controller to choose $\tau = \bar{\tau}$. The other items make sure that she
 1753 picks a valid tiling of the grid.

1754 From the initial state Environment chooses between three modes:

- 1755 ■ He can receive a sequence of $2N$ bits in $\{0, 1\}$ with the same datum and then let Controller
 1756 broadcast a letter of T with that same identifier.
- 1757 ■ He can receive a sequence of $2N$ bits in $\{\bar{0}, \bar{1}\}$ with the same datum and then let Controller
 1758 broadcast a letter of \bar{T} with that same identifier.
- 1759 ■ He can broadcast a sequence of letters of the form $x_1\bar{x}_1 \cdots x_N\bar{x}_N y_1\bar{y}_1 \cdots y_N\bar{y}_N$ with
 1760 $x_1, y_1, \dots, x_N, y_N \in \{0, 1\}$, all with his initial datum. He then receives one letter t' of
 1761 T and one letter \bar{t}' of \bar{T} with his initial datum. If $t = t'$ then he stops, otherwise he
 1762 broadcasts m_{err} .

- 1763 ■ He can broadcast a sequence of letters of the form $x_1\bar{x}_1 \cdots x_N\bar{x}_N y'_1\bar{y}_1 \cdots y'_N\bar{y}_N$ with
 1764 $x_1, y_1, y'_1, \dots, x_N, y_N, y'_N \in \{0, 1\}$, all with his initial datum. He makes sure that
 1765 $\langle y_1 \cdots y_N \rangle_2 = \langle y'_1 \cdots y'_N \rangle_2 + 1$. He then receives one letter t' of T and one letter \bar{t}
 1766 of \bar{T} with his initial datum. If $up(t') \neq down(t)$ or $\langle y'_1 \cdots y'_N \rangle_2 = 0$ and $down(t') \neq B$ or
 1767 $\langle y_1 \cdots y_N \rangle_2 = 2^N - 1$ and $up(t) \neq B$, he broadcasts m_{err} . Otherwise he stops broadcasting
 1768 m_{err} .
- 1769 ■ Similarly, he can broadcast a sequence of letters of the form $x'_1\bar{x}_1 \cdots x'_N\bar{x}_N y_1\bar{y}_1 \cdots y_N\bar{y}_N$
 1770 with $x_1, x'_1, y_1, \dots, x_N, x'_N, y_N \in \{0, 1\}$, all with his initial datum. He makes sure that
 1771 $\langle x_1 \cdots x_N \rangle_2 = \langle x'_1 \cdots x'_N \rangle_2 + 1$. He then receives one letter t' of T and one letter \bar{t} of
 1772 \bar{T} with his initial datum. If $right(t') \neq left(t)$ or $\langle x'_1 \cdots x'_N \rangle_2 = 0$ and $left(t') \neq B$ or
 1773 $\langle x_1 \cdots x_N \rangle_2 = 2^N - 1$ and $right(t) \neq B$, he broadcasts m_{err} . Otherwise he stops without
 1774 broadcasting m_{err} .

1775 If there is a valid tiling, Controller can play the corresponding strategy. In order to
 1776 broadcast m_{err} , Environment must make an agent a broadcast coordinates with its initial
 1777 datum, and then receive two tiles that do not satisfy the conditions mentioned above. The
 1778 agents that send those tiles must receive exactly $2N$ letters from a , as they are signed by its
 1779 initial datum. Thus their broadcasts are the tiles of the valid tiling at those coordinates, and
 1780 the agent will not be able to broadcast m_{err} , as they match all the conditions.

1781 If there is no valid tiling, Controller's strategy will either induce two different tilings or
 1782 two identical invalid ones. In both cases Environment can detect the mistake by making an
 1783 agent a broadcast the coordinates corresponding to the mistake, making two agents answer
 1784 with the faulty tiles, and make a broadcast m_{err} by observing the mistake. ◀