Consensus Game Acceptors

Dietmar Berwanger and Marie van den Bogaard

LSV, CNRS & ENS Cachan, Université Paris-Saclay, France

Abstract We study a game for recognising formal languages, in which two players with imperfect information need to coordinate on a common decision, given private input strings correlated by a finite graph. The players have a joint objective to avoid an inadmissible decision, in spite of the uncertainty induced by the input.
We show that the acceptor model based on consensus games characterises context-sensitive languages, and conversely, that winning strategies in such games can be described by context-sensitive languages. We also discuss consensus game acceptors with a restricted observation pattern that describe nondeterministic linear-time languages.

1 Introduction

The idea of viewing computation as an interactive process has been at the origin of many enlightening developments over the past three decades. With the concept of alternation, introduced around 1980 by Chandra and Stockmeyer, and independently by Kozen [6], computation steps are attributed to conflicting players seeking to reach or avoid certain outcome states. This approach relies on determined games with perfect information, and it lead to important and elegant results, particularly in automata theory. Around the same time, Peterson and Reif [18] initiated a study on computation via games with imperfect information, also involving teams of players. This setting turned out to be even more expressive, but also overwhelmingly difficult to comprehend. (See [3, 10], for more recent accounts.)

In this paper, we propose a game model of a language acceptor based on coordination games between two players with imperfect information. Compared to the model of Reif and Peterson, our setting is extremely simple: the games are played on a finite graph, plays are of finite duration, they involve only one yes/no decision, and the players have no means to communicate. Moreover, they are bound to take their decisions in consensus. Given an input word that may yield different observations to each of the players, they have to settle simultaneously and independently on a common decision, otherwise they lose.

We model such systems as consensus game acceptors, a particular case of coordination games with perfect recall, also described as multiplayer concurrent games or synchronous distributed games with incomplete information in the computer-science literature. Our motivation for studying the acceptor model is to obtain lower bounds on the complexity of basic computational problems regarding these more general games, specifically (1) solvability: whether a winning strategy exists, for a given game, and (2) implementability: which computational resources are needed to implement a winning strategy, if any exists.

Without the restrictions to consensus and to a single decision per play, the solvability problem for coordination games with safety winning conditions is known to be undecidable for two or more players [18, 19]. Furthermore, Janin [11] points out that there exist two-player safety games that admit a winning strategy but none that can be implemented by a Turing machine.

Our first result establishes a correspondence between context-sensitive languages and winning strategies in consensus games: We prove that every context-sensitive language L corresponds to a consensus game in which the characteristic function of L describes a winning strategy, and conversely, every consensus game that admits a joint winning strategy also admits one characterised by a context-sensitive language. Games with imperfect information for one player against the environment (which, here we call Input) display a similar correspondence with regular languages; they correspond to consensus games where the two player receive identical (or functionally dependent) observations. In extension, we consider consensus games where the observations of the two players can be ordered, and we show that the resulting acceptor model subsumes context-free languages and moreover allows to describe languages decidable in nondeterministic time.

The correspondence has several consequences in terms of game complexity. On the one hand, it reveals that consensus games preserve surprisingly much of the computational complexity found in games with imperfect information, in spite of the restriction to a single decision and to consensus. Consensus games are relevant because they represent coordination games in the limiting case where signalling is impossible. The classical constructions for proving undecidability of synchronous distributed games typically simulate a communication channel that may lose one message and involve an unbounded number of non-trivial decisions by which the players describe configurations of a Turing machine [19, 2, 22]. In contrast, our undecidability argument for acceptor games relies on the impossibility to attain common knowledge when knowledge hierarchies can grow unboudedly, and this can be witnessed by a single decision. Apart of this, we obtain a simple game family in which winning strategies are PSpace-hard to run, in the length of the play, or, in a positive perspective, where winning strategies can be implemented by a linear-bounded automata whenever they exist.

2 Preliminaries

For classical notions of formal languages, in particular context-sensitive languages, we refer, e.g., to the textbook of Salomaa [21]. We use the characterisation of context-sensitive languages in terms of nondeterministic linear-bounded automata, given by Kuroda [12] and the following well-known results from the same article: (1) For a fixed context-sensitive language L over an alphabet Σ, the problem whether a given word $w \in \Sigma^*$ belongs to L is PSpace-hard. (2) The problem of determining whether a given context-sensitive language represented by a linear-bounded automaton contains any non-empty word is undecidable.

2.1 Consensus game acceptors

Consensus acceptors are games between two cooperating players 1 and 2, and an additional agent called Input. Given a finite *observation alphabet* Γ common to both players, a *consensus game acceptor* $G = (V, E, (\beta^1, \beta^2), v_0, \Omega)$ is described by a finite set V of *states*, a *transition* relation $E \subseteq V \times V$, and a pair of *observation* functions $\beta^i : V \to \Gamma$ that label every state with an observation, for each player $i = 1, 2$. There is a distinguished initial state $v_0 \in V$ with no incoming transition. States with no outgoing transitions are called final states; the admissibility condition $\Omega : V \to \mathcal{P}(\{0, 1\})$ maps every final state $v \in V$ to a nonempty subset of admissible decisions $\Omega(v) \subseteq \{0, 1\}$. The observations at the initial and the final states do not matter, we may assume that they are labelled with the same observation # for both players.

The game is played as follows: Nature chooses a finite path $\pi = v_0, v_1, \ldots, v_{n+1}$ in G from the initial state v_0, following transitions $(v_\ell, v_{\ell+1}) \in E$, for all $\ell \leq n$, to a final state v_{n+1}. Then, each player i receives a private sequence of observations $\beta^i(\pi) := \beta^i(v_1), \beta^i(v_1), \ldots, \beta^i(v_n)$ and is asked to take a *decision* $a^i \in \{0, 1\}$, independently and simultaneously. The players win if they agree on an admissible decision, that is, $a^1 = a^2 \in \Omega(v_{n+1})$; otherwise they lose. Without risk of confusion we sometimes write $\Omega(\pi)$ for $\Omega(v_{n+1})$.

We say that two plays π, π' are *indistinguishable* to a player i, and write $\pi \sim^i \pi'$, if $\beta^i(\pi) = \beta^i(\pi')$. This is an equivalence relation, and its classes, called the *information sets* of Player i, correspond to observation sequences $\beta^i(\pi)$. A *strategy* for Player i is a mapping $s^i : V^* \to \{0, 1\}$ from plays π to decisions $s^i(\pi) \in \{0, 1\}$ such that $s^i(\pi) = s^i(\pi')$, for any pair $\pi \sim^i \pi'$ of indistinguishable plays. A joint strategy is a pair $s = (s^1, s^2)$; it is *winning*, if $s^1(\pi) = s^2(\pi) \in \Omega(\pi)$, for all plays π. In this case, the components s^1 and s^2 are equal, and we use the term *winning strategy* to refer to the joint strategy or either of its components. Finally, a game is *solvable*, if there exists a (joint) winning strategy.

In the terminology of distributed systems, consensus game acceptors correspond to *synchronous* systems with *perfect recall* and *known initial state*. They are a particular case of distributed games with safety objectives [16], coordination games with imperfect information [4], or multi-player concurrent games [1].

Strategies and knowledge. We say that two plays π and π' are *connected*, and write $\pi \sim^* \pi'$, if there exists a sequence of plays $\pi_1, \ldots, \pi_k$ such that $\pi \sim^1 \pi_1 \sim^2 \cdots \sim^1 \pi_k \sim^2 \pi'$. Then, a mapping $f : V^* \to \{0, 1\}$ from plays to decisions is a strategy that satisfies the consensus condition if, and only if, $f(\pi) = f(\pi')$, for all $\pi \sim^* \pi'$. In terms of distributed knowledge, this means that, for every play π, the events $\{\pi \in V^* \mid f(\pi) = 1\}$ and $\{\pi \in V^* \mid f(\pi) = 0\}$ are common knowledge among the players. (For an introduction to knowledge in distributed systems, see the book of Fagin, Halpern, Moses, and Vardi [9, Ch. 10, 11].) Such a consensus strategy — or, more precisely, the pair (f, f)— may still fail, due to prescribing inadmissible decisions. We say that a decision $a \in \{0, 1\}$ is *safe* at a play π if $a \in \Omega(\pi')$, for all $\pi' \sim^* \pi$. Then, a consensus strategy f is winning, if and only if, it prescribes a safe decision $f(\pi)$, for every play π.

It is sometimes convenient to represent a strategy for a player i as a function $f^i : \Gamma^* \to \{0, 1\}$. Every such function describes a valid strategy, because observation sequences identify information sets; we refer to an *observation-based* strategy in contrast to the *state-based* representation $s^i : V^* \to \{0, 1\}$. Note that the components of a joint winning strategy need no longer be equal in the observation-based representation. However, once the strategy for one player is fixed, the strategy of the other player is determined by the consensus condition, so there is no risk of confusion in speaking of a winning strategy rather than a joint strategy pair.

As an example, consider the game depicted in Figure 1, with observation alphabet $\Gamma = \{a, b, \triangleleft, \triangleright, \square\}$. States v at which the two players receive different observations are split, with $\beta^1(v)$ written in the upper part and $\beta^2(v)$ in the lower part; states at which the players receive the same observation carry only one symbol. The admissible decisions at final states are indicated on the outgoing arrow. Notice that upon receiving the observation sequence a^2b^2, for instance, the first player is constrained to choose decision 1, due to the following sequence of indistinguishable plays that leads to a play where deciding 0 is not admissible.

$$\begin{pmatrix} a, a \\ a, \triangleleft \\ b, \triangleright \\ b, b \end{pmatrix} \sim^2 \begin{pmatrix} a, a \\ \triangleleft, \triangleleft \\ \triangleright, \triangleright \\ b, b \end{pmatrix} \sim^1 \begin{pmatrix} a, \triangleleft \\ \triangleleft, \triangleright \\ \triangleright, \triangleleft \\ b, \triangleright \end{pmatrix} \sim^2 \begin{pmatrix} \triangleleft, \triangleleft \\ \triangleright, \triangleright \\ \triangleleft, \triangleleft \\ \triangleright, \triangleright \end{pmatrix} \sim^1 \begin{pmatrix} \triangleleft, \square \\ \triangleright, \square \\ \triangleleft, \square \\ \triangleright, \square \end{pmatrix} \sim^2 \begin{pmatrix} \square, \square \\ \square, \square \\ \square, \square \\ \square, \square \end{pmatrix}$$

In contrast, decision 0 may be safe when Player 1 receives input a^3b^2, for instance. Actually, the strategy $s^1(w)$ that prescribes 1 if, and only if, $w \in \{a^n b^n \ : \ n \in \mathbb{N}\}$ determines a joint winning strategy. Next, we shall make the relation between games and languages more precise.

3 Describing languages by games

We consider languages L over a letter alphabet Σ. The empty word ε is excluded from the language, and also from its complement $\bar{L} := (\Sigma^* \setminus \{\varepsilon\}) \setminus L$. As acceptors for such languages, we consider games over an observation alphabet $\Gamma \supseteq \Sigma$, and we assume that no observation sequence in Σ^+ is omitted: for every word $w \in \Sigma^+$, and each player i, there exists a play π that yields the observation sequence $\beta^i(\pi) = w$. Every consensus game acceptor can be modified to satisfy this condition without changing the winning strategies.

Given an acceptor game G, we associate to every observation-based strategy $s \in S^1$ of the first player, the language $L(s) := \{w \in \Sigma^* \ : \ s(w) = 1\}$. We say that the game G *covers* a language $L \subseteq \Sigma^*$, if G is solvable and

- $L = L(s)$, for *some* winning strategy $s \in S^1$, and
- $L \subseteq L(s)$, for *every* winning strategy $s \in S^1$.

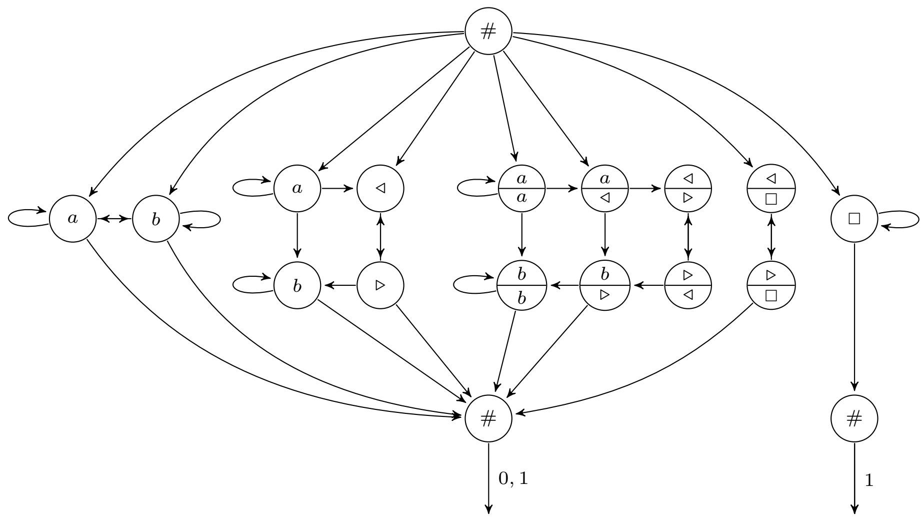

Figure 1. A consensus game acceptor

If, moreover, $L = L(s)$ for *every* winning strategy in G, we say that G *characterises* L. In this case, all winning strategies map L to 1 and $\bar{L}$ to 0.

As suggested above, the consensus game acceptor represented in Figure 1 covers the language $\{a^n b^n \ : \ n \in \mathbb{N}\}$. To characterise a language rather than covering it, we need to add constraints that require to reject inputs.

Given two games G, G', we define the *conjunction* $G \wedge G'$ as the acceptor game obtained by taking the disjoint union of G and G' and identifying the initial states. Then, winning strategies of the component games can be turned into winning strategies of the composite game, if they agree on the observation sequences over the common alphabet.

Lemma 1. *Let G, G' be two acceptor games over observation alphabets Γ, Γ'. Then, an observation-based strategy r is winning in $G \wedge G'$ if, and only if, there exist observation-based winning strategies s, s' in G, G' that agree with r on Γ^* and on Γ'^*, respectively.*

Whenever a language and its complement are covered by two acceptor games, we can construct a new game that characterises the language. The construction involves *inverting* the decisions in a game, that is, replacing the admissible decisions for every final state $v \in V$ with $\Omega(v) = \{0\}$ by $\Omega(v) := \{1\}$ and vice versa; final states v with $\Omega(v) = \{0, 1\}$ remain unchanged.

Lemma 2. *Suppose two acceptor games G, G' cover a language $L \subseteq \Sigma^*$ and its complement $\bar{L}$, respectively. Let G'' be the game obtained from G' by inverting the admissible decisions. Then, the game $G \wedge G''$ characterises L.*

3.1 Domino frontier languages

We use domino systems as an alternative to encoding machine models and formal grammars (See [23] for a survey.). A *domino system* $\mathcal{D} = (D, E_h, E_v)$ is described by a finite set of *dominoes* together with a horizontal and a vertical compatibility relation $E_h, E_v \subseteq D \times D$. The generic domino tiling problem is to determine, for a given system $\mathcal{D}$, whether copies of the dominoes can be arranged to tile a given space in the discrete grid $\mathbb{Z} \times \mathbb{Z}$, such that any two vertically or horizontally adjacent dominoes are compatible. Here, we consider finite rectangular grids $Z(\ell, m) := \{0, \dots, \ell + 1\} \times \{0, \dots, m\}$, where the first and last column, and the bottom row are distinguished as border areas. Then, the question is whether there exists a *tiling* $\tau : Z(\ell, m) \to D$ that assigns to every point $(x, y) \in Z(\ell, m)$ a domino $\tau(x, y) \in D$ such that:

- if $\tau(x, y) = d$ and $\tau(x + 1, y) = d'$ then $(d, d') \in E_h$, and
- if $\tau(x, y) = d$ and $\tau(x, y + 1) = d'$ then $(d, d') \in E_v$.

The *Border-Constrained Corridor* tiling problem takes as input a domino system $\mathcal{D}$ with two distinguished border dominoes $\#$ and $\Box$, together with a sequence $w = w_1, \dots, w_\ell$ of dominoes $w_i \in D$, and asks whether there exists a height m such that the rectangle $Z(\ell, m)$ allows a tiling τ with w in the top row, $\#$ in the first and last column, and $\Box$ in the bottom row:

- $\tau(i, 0) = w_i$, for all $i = 1, \dots, \ell$;
- $\tau(0, y) = \tau(\ell + 1, y) = \#$, for all $y = 0, \dots, m - 1$;
- $\tau(x, m) = \Box$, for all $x = 1, \dots, \ell$.

Domino systems can be used to recognise formal languages. For a domino system $\mathcal{D}$ with side and bottom border dominoes as above, the *frontier language* $L(\mathcal{D})$ is the set of words $w \in D^*$ that yield positive instances of the border-constrained corridor tiling problem. We use the following correspondence between context-sensitive languages and domino systems established by Latteux and Simplot.

Theorem 3 ([13, 14]). *For every context-sensitive language $L \subseteq \Sigma^*$, we can effectively construct a domino system $\mathcal{D}$ over a set of dominoes $D \supseteq \Sigma$ with frontier language $L(\mathcal{D}) = L$.*

Figure 2 describes a domino system for recognising the language $a^n b^n$ also covered by the game in Figure 1. In the following, we show that domino systems can generally be described in terms of consensus game acceptors.

3.2 Uniform encoding of domino problems in games

Game formulations of domino tiling problems are standard in complexity theory, going back to the early work of Chlebus [7]. However, these reductions are typically non-uniform: they construct, for every input instance consisting of a domino system together with a border constraint, a different game which depends, in particular, on the size of the constraint. Here, we use imperfect information to construct a *uniform* reduction that associates to a fixed domino system $\mathcal{D}$ a game $G(\mathcal{D})$, such that for every border constraint w, the question whether $\mathcal{D}, w$ allows a correct tiling is reduced to the question of whether decision 1 is safe in a certain play associated to w in $G(\mathcal{D})$.

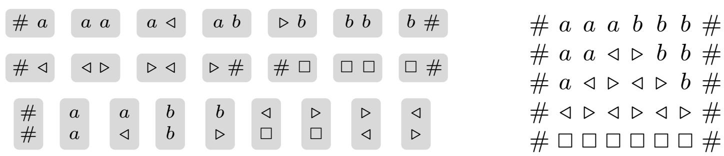

(a) domino system for $a^n b^n$ (b) tiling $a^3 b^3$

Figure 2. Characterising a language with dominoes

Proposition 4. *For every domino system D, we can construct, in polynomial time, a consensus game acceptor that covers the frontier language of D.*

Proof. Let us fix a domino system $\mathcal{D} = (D, E_h, E_v)$ with a left border domino $\#$ and a bottom domino $\Box$. We construct an acceptor game G for the alphabet $\Sigma := D \setminus \{\#, \Box\}$ to cover the frontier language $L(\mathcal{D})$.

The game is built as follows. There are domino states of two types: singleton states d for each $d \in D \setminus \{\#\}$ and pair states (d, b) for each $(d, b) \in E_v$. At singleton states d, the two players receive the same observation d.

At states (d, b), the first player observes d and the second player b. The domino states are connected by moves $d \to d'$ for every $(d, d') \in E_h$, and $(d, b) \to (d', b')$ whenever (d, d') and (b, b') are in E_h. There is an initial state v_0 and two final states $\widehat{z}$ and z, all associated to the the observation $\#$ for the border domino. From v_0 there are moves to all compatible domino states d with $(\#, d) \in E_h$, and all pair states (d, b) with $(\#, d)$ and $(\#, b) \in E_h$. Conversely, the final state z is reachable from all domino states d with $(d, \#) \in E_h$, and all pair states (d, b) with $(d, \#)$ and $(b, \#) \in E_h$; the final $\widehat{z}$ is reachable only from the singleton bottom domino state $\Box$. Finally, admissible decisions are $\Omega(z) = \{0, 1\}$ and $\Omega(\hat{z}) = \{1\}$. Clearly, G is an acceptor game, and the construction can be done in polynomial time.

Note that any sequence $x = d_1, d_2, \dots, d_\ell \in D^\ell$ that forms a horizontally consistent row in a tiling by $\mathcal{D}$ corresponds in the game to a play $\pi_x = v_0, d_1, d_2, \dots, d_\ell, z$ or $\pi_x = v_0, \Box^\ell, \widehat{z}$. Conversely, every play in G corresponds either to one possible row, in case Nature chooses a single domino in the first move, or to two rows, in case it chooses a pair. Moreover, a row x can appear on top of a row $y = b_1, b_2, \dots, b_\ell \in D^\ell$ in a tiling if, and only if, there exists a play ρ in G such that $\pi_x \sim^1 \rho \sim^2 \pi_y$, namely $\rho = v_0, (d_1, b_1), (d_2, b_2), \dots (d_\ell, b_\ell), z$.

Now, we claim that, at an observation sequence $\pi = w$ for $w \in \Sigma^\ell$ the decision 0 is safe if, and only if, there exists no correct corridor tiling by $\mathcal{D}$ with w in the top row. According to our remark, there exists a correct tiling of the corridor with top row w, if and only if, there exists a sequence of rows corresponding to plays $\pi_1, \ldots, \pi_m$, and a sequence of witnessing plays $\rho_1, \ldots, \rho_{m-1}$ such that $w = \pi_1 \sim^1 \rho_1 \sim^2 \pi_2 \cdots \sim^1 \rho_{m-1} \sim^2 \pi_m = \Box^\ell$. However, the decision 0 is unsafe in the play $\Box^\ell$ and therefore at w as well. Hence, every winning strategy s for G must prescribe $s(w) = 1$, for every word w in the frontier language of $\mathcal{D}$, meaning that $L(s) \subseteq L(\mathcal{D})$.

Finally, consider the mapping $s : D^* \to A$ that prescribes $s(w) = 1$ if, and only, if $w \in L(\mathcal{D})$. The observation-based strategy s in the acceptor game G is winning since $s(\Box^*) = 1$, and it witnesses the condition $L(s) = L(\mathcal{D})$. This concludes the proof that the constructed acceptor game G covers the frontier language of $\mathcal{D}$. □

4 Characterising context-sensitive languages

Theorem 5. *For every context-sensitive language $L \subseteq \Sigma^*$, we can construct effectively a consensus game acceptor that characterises L.*

Proof. Let $L \subseteq \Sigma^*$ be an arbitrary context-sensitive language, represented, e.g., by a linear-bounded automaton. By Theorem 3, it is possible to construct a domino system $\mathcal{D}$ with frontier language L. Then, by Proposition 4, we can construct an acceptor game G that covers $L(\mathcal{D}) = L$. Due to the Immerman-Szelepcsényi Theorem, context-sensitive languages are effectively closed under complement, so we construct an acceptor game G' that covers $\bar{L}$ following the same procedure. Finally, we combine the games G and G' as described in Lemma 2 to obtain an acceptor game that characterises L. □

One interpretation of the characterisation is that, for every context-sensitive language, there exists a consensus game that is as hard to play as it is to decide membership in the language. On the one hand, this implies that winning strategies for consensus games are in general PSpace-hard. Indeed, there are instances of acceptor games that admit winning strategies, however, any machine that computes the decision to take in a play requires space polynomial in the length of the play.

Theorem 6. *There exists a solvable consensus game acceptor for which every winning strategy is* PSpace*-hard.*

Proof. There exist context-sensitive languages with a PSpace-hard word problem [12]. Let us fix such a language $L \subseteq \Sigma^*$ together with a consensus game G that characterises it, according to Theorem 5. This is a solvable game, and every winning strategy can be represented as an observation-based strategy s for the first player. Then, the membership problem in L reduces (in linear time) to the problem of deciding the value of s in a play in G: For any input word $w \in \Sigma^*$, we have $w \in L$ if, and only if, $s(w) = 1$. In conclusion, it is PSpace-hard to decide whether $s(\pi) = 1$, for every winning strategy s in G. □

On the other hand, it follows that determining whether a consensus game admits a winning strategy is no easier than solving the emptiness problem of context-sensitive languages, which is well known to be undecidable.

Theorem 7. *The question whether an acceptor game admits a winning strategy is undecidable.*

Proof. We reduce the emptiness problem for a context-sensitive grammar to the solvability problem for a acceptor game.

For an arbitrary context-sensitive language $L \in \Sigma^*$ given as a linear bounded automaton, we construct an acceptor game G that characterises L, in polynomial time, according to Theorem 5. Additionally, we construct an acceptor game G' that characterises the empty language over Σ^*: this can be done, for instance, by connecting a clique over letters in Σ observable for both players to a final state at which only the decision 0 is admissible. Now, for any word $w \in \Sigma^*$, the game G' requires decision 0 at every observation sequences $w \in \Sigma^*$, whereas G requires decision 1 whenever $w \in L$. Accordingly, the acceptor game $G \wedge G'$ is solvable if, and only if, L is empty. As the emptiness problem for context-sensitive languages is undecidable [12], it follows that the solvability problem is undecidable for consensus game acceptors. □

We have seen that every context-sensitive language corresponds to a consensus game acceptor such that language membership tests reduce to winning strategy decisions in a play. Conversely, every solvable game admits a winning strategy that is the characteristic function of some context-sensitive language. Intuitively, a strategy should prescribe 0 at a play π whenever there exists a connected play π' at which 0 is the only admissible decision. Whether this is the case can be verified by a nondeterministic machine using space linear in the size of π.

Theorem 8. *Every solvable acceptor game admits a winning strategy that is implementable by a nondeterministic linear bounded automaton.*

5 Games for weaker language classes

The relation between the observation sequences received by the players in a synchronous game on a finite graph can also be explained in terms of letter-to-letter transducers, that is, finite-state automata where the transitions are labelled with input and output letters (See, e.g., [20, Ch. IV]). For a game G, the relation $\{(\beta^1(\pi), \beta^2(\pi)) \in \Gamma^* \times \Gamma^* \ : \ \pi \text{ a play in } G\}$ and its inverse are recognised by letter-to-letter transducers with the same transition structure as G. Conversely, every transducer τ can be turned into a game by letting one player observe the input and the other player the output letter of every transition. The consensus condition requires decisions to be invariant under the transitive closure τ^* of the described relation over Γ^*, which corresponds to iterating letter-to-letter transductions. Denoting by $L_{\text{acc}} \subseteq \Gamma^*$ the language of observation sequences for plays in which only the decision 1 is admissible, the Σ-language covered by G is $L := \Sigma^*\tau^*L_{\text{acc}}$. To characterise L, we additionally need to ensure $\bar{L} = \Sigma^*\tau^*L_{\text{rej}}$, for the language L_{rej} of observation sequences for plays in which only decision 0 is admissible. Thus, every consensus game acceptor can be described as a collection of three finite-state devices: two automata recognising the accepting and rejecting *seed* languages L_{acc} and L_{rej}, and a (nondeterministic) letter-to-letter transducer τ relating the observation sequences of the players.

Properties of iterated letter-to-letter transductions, or equivalently, length-preserving transductions, have been investigated in [15], also leading to the conclusion that iterated transducers capture context-sensitive languages. In the following, we investigate restrictions of consensus game acceptors towards capturing weaker language classes.

Firstly, we remark that regular languages correspond to games where the two players receive the same observation at every node.

Proposition 9. *Every regular language $L \subseteq \Sigma^*$ is characterised by a consensus game acceptor with identical observations for the players.*

Here, the consensus condition is ineffective, the model reduces to one-player games with imperfect information. To characterise a regular language L, we can build a game from a deterministic automaton for L, by moving symbols from transitions into the target states and allowing Nature to go from every accepting state in the automaton to a final game state v_{acc} with $\Omega(v_{\text{acc}}) = \{1\}$, and from every rejecting state to a final state v_{rej} with $\Omega(v_{\text{rej}}) = \{0\}$. Conversely, given a consensus game acceptor G with identical observations, the accepting seed language L_{acc} mentioned above yields the language characterised by G. Clearly, winning strategies in such games are regular.

We say that a consensus game acceptor has *ordered observations* if its alphabet Γ can be ordered so that $\beta^1(v) \geq \beta^2(v)$, for every state $v \in V$. One consequence of this restriction is that the implementation complexity of winning strategy drops from PSpace to NP.

Proposition 10. *Every solvable acceptor game with ordered observations admits a winning strategy that is characterised by a language in* NP*.*

Without loss of generality we can assume that the symbols occurring in L_{rej} or L_{acc} are disjoint from the input alphabet Σ and order them below. Then, given a sequence of observations $\pi \in \Sigma^*$, any sequence of indistinguishable plays that starts with observation π and leads to L_{rej} or L_{acc} is of length at most $|\Gamma| \times |\pi|$. To decide whether to prescribe 0 or 1 at π, a nondeterministic machine can guess and verify such a sequence in at most cubic time.

Despite this drop of complexity, games with ordered observations are sufficiently expressive to cover context-free languages.

Lemma 11. *Every context-free language is covered by a consensus game acceptor with ordered observations.*

Firstly, any Dyck language over a finite alphabet Σ of parentheses, possibly with extra neutral symbols, can be covered by a consensus game acceptor over

Σ extended with one additional $\square$ symbol ordered below Σ. The accepting seed language is $L_{\mathrm{acc}} = \square^*$, and the plays allow to project either a neutral symbol or an innermost pair of parentheses from the observation sequence of the first player by replacing them with $\square$ observations for the second player. Next, we observe that the class of languages covered by consensus game acceptors with ordered observations is effectively closed under intersection and letter-to-letter transductions, and thus particularly under letter-to-letter homomorphisms. The statement then follows by the Chomsky-Schützenberger [8] representation theorem, in the non-erasing variant proved by Okhotin [17]: every context-free game is the letter-to-letter homomorphic image of a Dyck language with neutral symbols with a regular language.

Since it is undecidable whether two context-free languages have non-empty intersection [21], the above lemma also imples that the solvability problem for consensus game acceptors is undecidable, already when observations are ordered. Concretely, we can represent the standard formulation of Posts Correspondence Problem as a solvability problem for such restricted consensus games.

Corollary 12. *The question whether a consensus game acceptor with ordered observations admits a winning strategy is undecidable.*

Due to the characterisation of nondeterministic linear-time languages as homomorphic images of intersections of three context-free languages due to Book and Greibach [5], we can draw the following conclusion.

Theorem 13. *For every language L decidable in nondeterministic linear time, we can effectively construct a consensus game acceptor with ordered observations that covers L.*

Acknowledgement. This work was supported by the European Union Seventh Framework Programme under Grant Agreement 601148 (CASSTING).

References

[1] R. Alur, T. A. Henzinger, and O. Kupferman, *Alternating-time temporal logic*, J. ACM, 49 (2002), pp. 672–713.

[2] A. Arnold and I. Walukiewicz, *Nondeterministic controllers of nondeterministic processes*, in Logic and Automata, vol. 2, Amsterdam University Press, 2007.

[3] S. Azhar, G. Peterson, and J. Reif, *Lower bounds for multiplayer noncooperative games of incomplete information*, Journal of Computers and Mathematics with Applications, 41 (2001), pp. 957–992.

[4] D. Berwanger and Ł. Kaiser, *Information tracking in games on graphs*, Journal of Logic, Language and Information, 19 (2010), pp. 395–412.

[5] R. V. Book and S. A. Greibach, *Quasi-realtime languages*, Mathematical Systems Theory, 4 (1970), pp. 97–111.

[6] A. K. Chandra, D. Kozen, and L. J. Stockmeyer, *Alternation*, Journal of the ACM, 28 (1981), pp. 114–133.

[7] B. S. Chlebus, *Domino-tiling games*, Journal of Computer and System Sciences, 32 (1986), pp. 374 – 392.

[8] N. Chomsky and M. Schützenberger, *The algebraic theory of context-free languages*, in Computer Programming and Formal Systems, vol. 35 of Studies in Logic and the Foundations of Mathematics, Elsevier, 1963, pp. 118 – 161.

[9] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi, *Reasoning about Knowledge*, MIT Press, 1995.

[10] R. A. Hearn and E. D. Demaine, *Games, Puzzles, and Computation*, A. K. Peters, Ltd., Natick, MA, USA, 2009.

[11] D. Janin, *On the (high) undecidability of distributed synthesis problems*, in Proc. of Theory and Practice of Computer Science (SOFSEM 2007), vol. 4362 of Lecture Notes in Computer Science, Springer, 2007, pp. 320–329.

[12] S.-Y. Kuroda, *Classes of languages and linear-bounded automata*, Information and Control, 7 (1964), pp. 207–223.

[13] M. Latteux and D. Simplot, *Context-sensitive string languages and recognizable picture languages*, Information and Computation, 138 (1997), pp. 160 – 169.

[14] M. Latteux and D. Simplot, *Recognizable picture languages and domino tiling*, Theoretical Computer Science, 178 (1997), pp. 275 – 283.

[15] M. Latteux, D. Simplot, and A. Terlutte, *Iterated length-preserving rational transductions*, in Mathematical Foundations of Computer Science 1998, Springer, 1998, pp. 286–295.

[16] S. Mohalik and I. Walukiewicz, *Distributed games*, in FSTTCS 2003: Foundations of Software Technology and Theoretical Computer Science, Springer, 2003, pp. 338–351.

[17] A. Okhotin, *Non-erasing variants of the Chomsky-Schützenberger theorem*, in Developments in Language Theory, H.-C. Yen and O. Ibarra, eds., vol. 7410 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2012, pp. 121–129.

[18] G. L. Peterson and J. H. Reif, *Multiple-Person Alternation*, in Proc 20th Annual Symposium on Foundations of Computer Science, (FOCS 1979), IEEE, 1979, pp. 348–363.

[19] A. Pnueli and R. Rosner, *Distributed reactive systems are hard to synthesize*, in Proceedings of the 31st Annual Symposium on Foundations of Computer Science, FoCS '90, IEEE Computer Society Press, 1990, pp. 746–757.

[20] J. Sakarovitch, *Elements of automata theory*, Cambridge University Press, 2009.

[21] A. Salomaa, *Formal Languages*, Academic Press, New York, NY , USA, 1973.

[22] S. Schewe, *Distributed synthesis is simply undecidable*, Inf. Process. Lett., 114 (2014), pp. 203–207.

[23] P. van Emde Boas, *The convenience of tilings*, in Complexity, Logic, and Recursion Theory, vol. 18 of Lecture Notes in Pure and Applied Mathematics, Marcel Dekker Inc, 1997, pp. 331–363.