

Gradual Guaranteed Coordination in Repeated Win-Lose Coordination Games

Valentin Goranko^{a,b} and Antti Kuusisto^{c,d} and Raine Rönholm^{d 1}

Abstract. We investigate repeated win-lose coordination games and analyse when and how rational players can guarantee eventual coordination in such games. Our study involves both the setting with a protocol shared in advance as well as the scenario without an agreed protocol. In both cases, we focus on the case without any communication amongst the players once the particular game to be played has been revealed to them. We identify classes of coordination games in which coordination cannot be guaranteed in a single round, but can eventually be achieved in several rounds by following suitable coordination protocols. In particular, we study coordination using protocols invariant under structural symmetries of games under some natural assumptions, such as: priority hierarchies amongst players, different patience thresholds, use of focal groups, and gradual coordination by contact.

1 INTRODUCTION

1.1 Repeated pure win-lose coordination games

Pure win-lose coordination games (WLC-games) were studied in our work [8] (and the extended version [9]) as strategic form games in which all players receive the same payoffs: 1 (win) or 0 (lose). All players have the same goal, to coordinate on any winning strategy profile. Here we extend that study to *repeated (pure) win-lose coordination games*, where the players make simultaneous choices in a discrete succession of rounds. The game continues until the players either succeed to coordinate or, alternatively, until they run out of time, i.e., fail to coordinate in a predetermined (possibly infinite) number of rounds. We note that these games are not of the type of repeated games usually studied in game theory, where not reachability objectives but accumulated payoffs are considered.

Scenarios modelled by repeated pure coordination games occur naturally in real life, for example when a group of people who cannot communicate try to get together in one of several possible meeting places and go around in search for each other. Another common scenario is the phenomenon called ‘pavement tango’², where two people try to pass each other but end up blocking each others’ way by repeatedly moving to the same direction before either bumping into each other or eventually succeeding to resolve the situation. These scenarios can be regarded as a canonical examples of the kinds of repeated coordination games that we consider here.

To give a simple example of a scenario where coordination can be *guaranteed* with certainty, but not in the first round, consider a

setting where three agents are to coordinate by meeting in one of two possible locations. We assume that the agents can move from one location to the other in discrete, synchronous rounds, on the tick of the clock. They cannot communicate with each other during the play, so clearly they cannot guarantee coordination right away in the first round. However, assuming that each agent can observe if there is someone else in the same location, a simple (predesigned) protocol will guarantee coordination within two rounds, as follows: Each agent, after moving to some location in the first round, observes if there is anyone else there. If so, she stays in that location; else she moves to the other one in the second round.

In this paper we study several natural variants of repeated coordination games and investigate conditions under which such games can be solved, i.e., coordination can be guaranteed. We adopt and extend the technical framework of [8] and [9], where WLC-games are represented by abstract winning relations over sets of players’ choices. We assume a common belief among the players that:

1. *all players know the structure of the game;*
2. *all players have the same goal, viz. to coordinate, by selecting together a winning profile, and*
3. *all players are rational, i.e., act towards achieving that goal.*

In this paper we assume that the players want to *guarantee* coordination (with certainty), in as few rounds as possible. That is, their top priority is to coordinate with certainty, and if that is possible, a further aim is to minimize the time for guaranteed coordination.

We focus on two communication scenarios between the players:

1. The case of *no communication at all*, neither before nor during the play of the game. Here the players do not share any prenegotiated strategies or conventions, but play independently of each other.
2. The case where the players may have agreed on a joint protocol (global non-deterministic strategy for all games), but only *before* the concrete game to play is presented to them. Here again the players cannot communicate once the game has been presented.

The players’ strategies, while possibly synchronised by a joint protocol, are assumed to be indifferent to the particularities of concrete games such as, e.g., ‘names of choices’. Technically this means invariance under ‘renamings’ of players and choices (see Section 3).

1.2 Structure and content of the paper

After providing preliminaries on WLC-games in Section 2, we introduce the basics of repeated WLC-games and give some examples in Section 3. In that section we also discuss possible applications of rationality principles for solving one-shot WLC-games in the context

¹ ^aStockholm University, Sweden; ^bUniversity of Johannesburg, South Africa (visiting professorship); ^cUniversity of Helsinki, Finland; ^dTampere University, Finland

² Douglas Adams and John Lloyd have jokingly coined the word ‘droitwich’ in [1] to name that phenomenon.

of repeated games. In particular, we consider *symmetry-based principles* and show that there are WLC-games that are unsolvable in the single-round scenario of [8], but become solvable as repeated games, simply due to reasoning based on symmetries.

The main new results are in Section 4, where we consider, e.g., different assumptions about visibility and distinctions in the behaviour and roles of different players. For most of the studied cases, we establish conditions under which protocols for guaranteed coordination exist and find bounds for the number of rounds needed to guarantee coordination. We consider three different cases of *visibility* between players: *minimal*, where the players can only observe their own choices but not those of the others; *local*, where they can only see those other players with whom they currently share a winning profile; and *complete visibility*, where the players always see the choices of all of the other players (like in the ‘pavement tango’ scenario).

First, in Section 4.1, we construct coordination protocols assuming a commonly known, naturally established or prenegotiated *strict priority hierarchy* of players. The hierarchy can be used, e.g., to determine an order according to which the players choose with which other players they attempt to coordinate in different, consecutive rounds. We construct such protocols under different assumptions of visibility and show that they have significantly different maximum times for coordination.

Then, in Section 4.2, we study coordination based on the notion of a *patience threshold*, which means the maximum number of rounds for which a player is willing to stay in the same location (i.e., repeat the same choice) without observing any change. We show that assuming different patience thresholds for all players leads to a natural way of coordinating in all repeated WLC-games. Next, in Section 4.3, we discuss coordination by joining so called *focal/anchor groups* which remain stationary on a winning profile while the remaining players try to coordinate with them.

Lastly, in Section 4.4, we explore *coordination by contact*, where players who get into contact (by making simultaneous choices in the same winning profile) can synchronise all their further actions. We show that under local visibility, supposing a first-contact has been made, the subsequent game can be solved in linear time in the number of choices. Furthermore, the contact assumption leads to very quick (logarithmic number of rounds in the number of players) coordination under complete visibility and using a hierarchy of players. Concluding remarks and directions for further work are given in Section 5.

1.3 Related work

There is a vast literature on coordination games in addition to [8]. However, most of these works only bear a somewhat superficial connection to the present study, as they consider coordination games with different (though possibly equal for all players) payoffs. This changes the focus to a more traditional, quantitative game-theoretic study where strategy dominance and equilibrium analysis are more relevant. Still, some essential concepts, such as common belief in rationality, focal points, conventions, and symmetries, are of common importance to both approaches.

We list some of the more relevant and notable references. First, the classical works of Schelling [13] and Lewis [11] lay much of the foundations of coordination games and demonstrate the importance of focal points, salience, and conventions. The follow-up study on coordination by Gauthier [5] is also relevant, and so are Sugden’s survey on rational choice [15] and his paper [16] on focal points. Essential general references include the book [10] on the theory of

rational choice in selecting equilibria and the books [2], [12] on repeated games. The work [4] on repeated coordination games is particularly relevant as, e.g., it emphasises the importance of symmetries. However, it indeed considers a more general set of payoffs than the current paper, and thereby the focus is different. Our first paper and main reference on the topic is [8], substantially extended in [9]. The work [7] extends the results in the latter to WLC games with extra structure including, e.g., priority hierarchies of players. In all these works we have considered one-shot games only. Other relevant references include [14], [6], [3].

Lastly, we note that the present work is relevant in a broader sense to distributed computing and algorithms, and some aspects of our work (e.g., communication by contacts) are more specifically related to gossip protocols.

2 PRELIMINARIES

To give precise definitions of the necessary technical notions, here we adopt the notation and terminology for (pure) win-lose coordination games from [8], to which we refer the reader for further background.

Definition 2.1. An n -player **win-lose coordination game** (WLC-game) is a relational structure $G = (A, C_1, \dots, C_n, W_G)$ where A is a finite and non-empty domain of **choices**, each C_i is a non-empty unary relation (representing the choices of player i) such that $C_1 \cup \dots \cup C_n = A$, and $W_G \subseteq C_1 \times \dots \times C_n$ is an n -ary **winning relation**. For technical reasons we assume that the players have pairwise disjoint choice sets, i.e., $C_i \cap C_j = \emptyset$ for all $i \neq j$. A tuple $\sigma \in C_1 \times \dots \times C_n$ is called a **choice profile** for G and the choice profiles in W_G are called **winning choice profiles**.

In this paper we make the following additional assumptions:

1. There is at least one winning choice profile, i.e. W_G is nonempty;
2. No player has a surely losing choice (i.e., a choice that does not belong to any winning choice profile).


The assumption 1 is reasonable, as games with an empty winning relation can never be won. The assumption 2 (which formally implies assumption 1, as A is non-empty) is also justified, assuming that no rational player would ever select a surely losing choice if there are other choices available (cf. the ‘Non-losing principle’ NL in [8]).

Consider a WLC-game $G = (A, C_1, \dots, C_n, W_G)$, and let $X = \{p_1, \dots, p_k\} \subseteq \{1, \dots, n\}$ be a nonempty set of players³ such that $p_1 < \dots < p_k$. Let $\bar{X} := \{1, \dots, n\} \setminus X$ denote the set of remaining players $p'_1 < \dots < p'_m$. For any choice profile $\sigma = (c_1, \dots, c_n) \in C_1 \times \dots \times C_n$, we define $\sigma_X := (c_{p_1}, \dots, c_{p_k})$. Let T be the set of tuples τ in W_G that contain σ_X as a subtuple. Then the **winning extension** of σ_X is the relation $W_G(\sigma_X) := \{\tau_{\bar{X}} \mid \tau \in T\}$, i.e., the set of winning tuples that contain σ_X but with the subtuple σ_X itself projected away. Let $G^* := (A, C_{p'_1}, \dots, C_{p'_m}, W_G(\sigma_X))$ and let $G(\sigma_X)$ to be the same as G^* but with all surely losing choices removed (if there are any). The game $G(\sigma_X)$ is called the **subgame (of G) induced by the winning extension of σ_X** . When σ_X is a singleton tuple, i.e., a single player’s choice c , we simply write $W_G(c)$ and $G(c)$ instead of $W_G(\sigma_X)$ and $G(\sigma_X)$.

For technical convenience, we will use the visual presentation of WLC-games as hypergraphs from [8]. The choices of each player are displayed as columns of nodes, starting from the choices of player 1 on the left and ending with the column with choices of player n . The winning relation consists of lines that represent the winning profiles.

³ We routinely identify the player indices $1, \dots, n$ with players.

Example 2.2. Here are two examples of WLC-games: a 2-player game G_1 with 3 choices for both players and a total of 5 winning profiles presented as edges; and a 3-player WLC-game G_2 with 2 choices for each player and 4 winning profiles, each represented as a triple of choices connected by (solid or dotted) lines.



We note that G_2 represents a coordination game where it suffices that players 1 and 3 make the ‘same’ choice (i.e., both choose from the top row or, alternatively, both choose from the bottom row) in order to guarantee coordination.

3 REPEATED WLC-GAMES AND SOLVABILITY BY RATIONAL PRINCIPLES

In this section we introduce repeated WLC-games and discuss solvability of such games by using rational principles of coordination studied in [8].

3.1 Basics of repeated WLC-games

Consider a scenario where the players are expected to play a WLC-game G repeatedly until they eventually coordinate (by selecting some winning choice profile). This leads to a **repeated play** of G which consists of consecutive one-step plays of G until (if ever) the players select a choice profile in W_G . Thus every WLC-game can be associated with the corresponding **repeated WLC-game**.


We assume that the players can remember the history of the repeated play and use this information when planning their next choice. The history of a play after k rounds is encoded in a sequence \mathcal{H}_k defined formally below.

Definition 3.1. Consider an n -player WLC-game G . Let \mathcal{H}_k be a k -sequence of choice profiles in G . We call \mathcal{H}_k a **history** of G , and the pair (G, \mathcal{H}_k) is referred to as a **stage k (or a k th stage) in a repeated play of G** . Formally, $\mathcal{H}_k = \{H_i\}_{i \in \{1, \dots, k\}}$ where each H_i is an n -ary relation $H_i = \{(c_1, \dots, c_n)\}$ for some $(c_1, \dots, c_n) \in C_1 \times \dots \times C_n$.

Note that a stage k has a history containing k earlier choice tuples in a repeated play, and thus we define the **initial stage** (the 0th stage (G, \mathcal{H}_0)) so that $\mathcal{H}_0 = \emptyset$. If coordination succeeds after k rounds, then the k th stage is called the **final stage** of the repeated play. Alternatively, the repeated play may also continue for an infinite number of rounds without coordination.

Consider a stage (G, \mathcal{H}_k) in a repeated play of G with $H_k = \{(c_1, \dots, c_n)\}$. We sometimes say that the choice c_i (which is the most recent choice) of player i is the **current location** of i . This terminology naturally comes from coordination games where choices are different (physical) locations and selecting a choice is interpreted as moving to the selected location.

When presenting stages in repeated WLC-games graphically, we label the previously chosen nodes in different ways. Below is a graphical representation of a second stage in a repeated play of the game $G_{2,2}$, which is a ‘coordination game version’ of the matching pennies game or the ‘pavement tango’ scenario from the introduction. Here the players have failed to coordinate in round 1 and then failed again by swapping their choices in round 2.



We now generalize the definitions of protocols and principles from [8] to repeated WLC-games. A **protocol** Σ is a function that outputs a set $C \subseteq C_i$ of choices with the input of a player i and a stage (G, \mathcal{H}_k) of any WLC-game G . A **principle** P is any nonempty set of protocols. As in [8], a protocol describes a (non-deterministic) strategy in any given WLC-game in the role of any player i . Here, a protocol also ‘sees’ all the history of the current stage and thus gives a (memory-based) strategy for any repeated WLC-game.

Principles are properties of protocols and thus they can be seen as ‘reasoning or behaviour styles’ which can be applied in repeated games. Principles can usually be defined by using simple descriptions in natural language. Here is an example of a principle for repeated games: “*Change your choice in every second round*”. Clearly there are several different protocols that belong to this principle.

We say that a player **follows** a principle P if he uses a protocol from P . A principle P **solves a WLC-game G in k rounds** if players are *guaranteed* to coordinate in at most k rounds when every player follows P , and coordination is *not guaranteed* in fewer than k rounds⁴. Similarly, we say that a **protocol Σ solves G in k rounds** if the singleton principle $\{\Sigma\}$ solves G in k rounds.

3.2 Structural equivalences and protocols

As argued in [8], it is natural to assume that rational principles only consider ‘structural properties’ of a game. In order to define this formally, a notion of *renaming* between WLC-games is introduced in [8]. We need to generalize this notion here to involve stages of repeated games, as the history of a repeated game creates additional structure that can be used by the players in later rounds. The formal definition below is easier to understand by also considering the related Example 3.3. The intuitive idea is to relax isomorphisms between game graphs (including histories) to allow permutations of the players $1, \dots, n$.


Definition 3.2 (cf. [8]). Consider stages (G, \mathcal{H}_k) and (G', \mathcal{H}'_k) in n -player WLC-games G and G' , where $k \in \mathbb{N}$. A pair (β, π) is a **renaming** between (G, \mathcal{H}_k) and (G', \mathcal{H}'_k) if there is an n -player WLC-game G'' and a history \mathcal{H}''_k of a repeated play of G'' such that the following conditions hold:

1. β is a permutation of $\{1, \dots, n\}$ such that (G'', \mathcal{H}''_k) is obtained from (G, \mathcal{H}_k) by ‘permuting the players’. That is, we use β to permute the choice sets C_i of players and the tuples both in the winning relation and in the history.
2. π is an isomorphism between (G'', \mathcal{H}''_k) and (G', \mathcal{H}'_k) . Here it is useful that the relations H_i in the history are part of the relational structures together with the WLC-games G'' and G' .

If (G, \mathcal{H}_k) and (G', \mathcal{H}'_k) have the same domain A , we say that (β, π) is a **renaming of (G, \mathcal{H}_k)** . We say that the choices $c \in C_i$ and $c' \in C_j$ are **structurally equivalent**, denoted by $c \sim c'$, if there is a renaming (β, π) of (G, \mathcal{H}_k) such that $\beta(i) = j$ and $\pi(c) = c'$. It is easy to see that \sim is an equivalence relation on the set A of all choices. We denote the equivalence class of a choice c by $[c]$.

⁴ Note that this means the time for coordination in the *worst case scenario*, as it is possible that players can coordinate by chance in fewer than k rounds when following P .

Example 3.3. There is a renaming between the stages (G, \mathcal{H}_2) and (G', \mathcal{H}'_2) below. Indeed, after swapping the players of the stage on the left, there is an isomorphism to the stage on the right. Also note that the choices c and d are structurally equivalent in the initial stage (G, \mathcal{H}_0) , but this equivalence is broken when player 2 selects c in the first round.



We say that a protocol Σ is **structural** if it is ‘indifferent’ with respect to renamings. This means that, given any stages (G, \mathcal{H}_k) , (G', \mathcal{H}'_k) with a renaming (β, π) between (G, \mathcal{H}_k) and (G', \mathcal{H}'_k) , for any i and any $c \in C_i$, we have $c \in \Sigma((G, \mathcal{H}_k), i)$ if and only if $\pi(c) \in \Sigma((G', \mathcal{H}'_k), \beta(i))$. Intuitively, this reflects the idea that when following a structural protocol, one acts independently of the names of choices and the names (or ordering) of players. However, a structural protocol may ‘see’ all the history of the game and remember in which order all choices have been played. We say that a **principle is structural** if it consists only of structural protocols.

3.3 Solvability of repeated WLC-games with rational principles for one-step games

All the principles P defined for one-step WLC-games (such as the ones defined in [8]) can be also used in repeated WLC-games. Perhaps the simplest way of doing this is when players simply apply P in each stage of the game the same way as they would in the one-step WLC-game, completely disregarding the history. It is clear that if a principle P solves the one-step WLC-game G , then it solves the repeated version of G in one round. However, assuming the history is ignored, if P does not solve G , then it alone does not guarantee coordination in any number of rounds in the repeated play of G .

It is thus natural to ask the question how the history could be used together with different rational principles. Let us first consider the well known principle from standard game theory which prescribes *iterated elimination of dominated choices*. (This can be interpreted to mean, e.g., that the players use the principle ‘Collective rational choices’, CRC, of [8].) Since repeated coordination attempts do not affect domination between choices, there is no obvious way to use the history to strengthen this principle for repeated games. For similar reasons, most of the rational principles for one-step WLC-games presented in [8] do not give anything more for repeated WLC-games. The only exception here are the *symmetry based principles* which we will discuss in the next subsection.

It is worth noting that even if a (rational) principle P does not solve a repeated game, it may still be useful when combined with other principles.

3.4 Symmetry principles in repeated games


Since we have generalized the definition of renamings to cover stages of repeated games, we can define the symmetry principle ES from [8] exactly as defined in that article. A justification for this principle and a related concrete example will be given after the definition.

Definition 3.4 ([8]). Consider a choice profile $\vec{c} = (c_1, \dots, c_n)$ in a WLC-game G and let $U_i := C_i \cap \bigcup_j [c_j]$ for each i (recall the notation $[c_i]$ from Def. 3.2). We say that (c_1, \dots, c_n) **exhibits a bad**

symmetry if $U_1 \times \dots \times U_n \not\subseteq W_G$. Moreover, a choice c **generates a bad symmetry** if every choice profile that contains c exhibits a bad symmetry. We can now define the principle of **Elimination of bad symmetries (ES)** as follows: “Never play choices that generate bad symmetries, if possible”. We also assume that all protocols in ES are structural.

The justification for ES is that—assuming no communication or conventions—rational players should be indifferent between structurally equivalent choices. Therefore players may assume that everyone follows a structural protocol. By this assumption, if a choice generates a bad symmetry, then coordination is not guaranteed by choosing it and thus it should be avoided, if possible.

Example 3.5. Consider the game $G_{2,3}$ (or $G(3(1 \times 1))$) in the notation established in [8] with 2 players and 3 winning profiles, pictured below. Intuitively, this can be seen as a game where two players try to meet at one of three locations—under the assumption that they can somehow always see (or infer) which location was chosen by the other player.



We assume that the players are not able to communicate before or during the play, but their choices are recorded in a history. The nodes l_1^1, l_2^1, l_3^1 are the choices of player 1 and l_1^2, l_2^2, l_3^2 are the choices of player 2; the circles around some of the choices represent a history, to be discussed below. Indeed, we will sketch an argument showing that—even without communication—the players are guaranteed to win this simple game in at most two rounds by using symmetry-based reasoning principles. We note that all choices here are initially structurally equivalent and generate a bad symmetry. Therefore, $G_{2,3}$ is *structurally unsolvable* (cf. [8]), i.e., unsolvable by any structural principle as a one-step coordination game.

Now, due to the structural equivalence of all choices, when following ES in the repeated coordination game, the players make their first choices randomly. Suppose that they fail to coordinate in the first round, e.g., due to picking the choices l_1^1 and l_2^2 (with dotted circles around them in the figure). When the history of the first round is added to the next stage of the game, we have $l_1^1 \sim l_2^2$, $l_2^1 \sim l_1^2$ and $l_3^1 \sim l_3^2$. Hence the choices l_1^1, l_2^1, l_1^2 and l_2^2 generate a bad symmetry, but the choices l_3^1 and l_3^2 do not. Thus, following ES, the players will choose (l_3^1, l_3^2) in the second round, thereby succeeding to coordinate.

For further examples and discussion on the rationality of the symmetry principle ES, see [8].

The reasoning in Example 3.5 can easily be generalized to those 2-player WLC-games where the winning relation forms three isomorphic components in the bipartite game graph. If players select their choices from different components in the first round, then, in the second round, they can both select a choice from the third remaining component. This alone does not typically guarantee coordination, but other principles in [8] can be additionally applied in order to try to coordinate in the remaining third component.

Another natural generalization of the game $G_{2,3}$ is $G_{2,m}$ with two players and m winning profiles that can all be drawn horizontally, i.e., each choice coordinates with precisely one choice of the other player. For any odd m , the protocol from Example 3.5 can now be generalized so that after every round where the players have failed to

coordinate, they always select a new (previously not selected) choice, until eventually a single unselected choice remains. Clearly coordination is then guaranteed whenever m is odd and it then takes at most $\lceil m/2 \rceil$ rounds. It is also easy to see that coordination cannot be guaranteed in $G_{2,m}$ by any structural protocol when m is even.

4 PROTOCOLS FOR GRADUAL GUARANTEED COORDINATION

In this section we consider several multi-player scenarios where players guarantee coordination gradually with the help of their observations, memory, and also possibly some additional features, such as players' *priority hierarchies* and *patience thresholds*, to be defined formally further. In most of our examples it is natural to assume that the players use a joint protocol which is agreed upon before they are presented with the actual game to play.

First, we need to distinguish three basic cases with respect to the visibility amongst players during the play.

Definition 4.1. We say that a repeated WLC-game has:

- **complete visibility** if each player can see (and remember) the choices of all other players after every round. Here the protocols used by a player may depend on the complete history \mathcal{H}_k .
- **local visibility** if each player can only see their own choices and the current choices of those other players who he shares a winning profile with (i.e., those other players who he is 'currently coordinating' with). In such games the protocol used by a player may only depend on those choices, in \mathcal{H}_k , which he has seen earlier in the game.
- **minimal visibility** if each player can only see the choices that he has made during the play. Here the protocols used by a player may only depend on their own choices in \mathcal{H}_k .

One could naturally consider a more general notion for visibility by using *visibility graphs*. With this approach, each player can observe (and remember) only choices made by the players she can 'see'. We leave this analysis for a further work.

4.1 Hierarchical coordination

If players have a commonly known strict hierarchy (technically, a linear priority order) amongst them, they can use that hierarchy in several natural scenarios involving gradual coordination. When a hierarchy is present, it naturally allows protocols which violate some symmetries between different players.

We consider the scenario where all players attempt, one-by-one, to coordinate with all those who are 'senior' to them in the hierarchy.

Theorem 4.2. Consider WLC-games with a commonly known total hierarchy amongst the players. Then there is a protocol which guarantees coordination in every repeated n -player WLC-game G within N rounds, where

1. $N = n$, if the players have complete visibility,
2. $N = 1 + (m - 1)(n - 1)$, where m is the maximal number of choices per player, if the players have local visibility,
3. $N = \left(\frac{1}{|C_\ell|} \prod_{i=1}^n |C_i| \right) - w_\ell + 1$, if the players have minimal visibility, where ℓ is the most senior player and w_ℓ is the largest number of winning choice profiles containing a fixed choice of player ℓ . (Recall that $|C_i|$ is the cardinality of C_i .)

Proof. Suppose, w.l.o.g., that the priority hierarchy arranges the players into the order $1 < 2 < \dots < n$, with n most senior.

1. Under complete visibility, the protocol is straightforward: in the first round all players choose randomly. If they coordinate, the game ends. Otherwise, the most senior player n remains stationary in his location, i.e., keeps repeating his first choice u , and, recursively, the play continues in the subgame $G(u)$, i.e., the game induced by the winning extension of u in G (see Sec 2). In $G(u)$, the remaining $n - 1$ players try to coordinate in exactly the same way, with the player $n - 1$ now being the most senior player. Clearly coordination is guaranteed to ultimately occur within n rounds.

2. Under local visibility, first the most senior player n fixes a choice c_n in the 1st round. Then the next most senior player, $n - 1$ 'finds' n on some winning profile containing c_n . This takes at most $m - 1$ more rounds and fixes the choice c_{n-1} of player $n - 1$. Meanwhile, all the other players wait exactly $m - 1$ additional rounds (after round 1), to make sure that players n and $n - 1$ have coordinated. Thereafter the procedure continues likewise. Thus, each player $1, \dots, n - 1$ will take at most $m - 1$ additional rounds to coordinate with the more senior ones.

3. Lastly, under minimal visibility, the protocol consists, intuitively, in traversing all choice profiles in a systematic (lexicographic) way until reaching a winning choice profile, as follows (recall the assumption that no choice is surely losing). First, the most senior player n fixes a 'best choice' u , which is in w_ℓ winning choice profiles. Note that such a best choice need not be unique, and therefore we cannot assume that the other players know it. Then each other player i , for $i < n$, fixes an ordering of the set C_i of his choices⁵. Now, in every round, each player i picks a choice, starting from the first one in the fixed order, so that each of his choices is repeated (if necessary) for exactly $r_i := |C_1 \times \dots \times C_{i-1}|$ consecutive rounds (where $r_1 := 1$) before moving to the next choice. This is done until a winning choice profile is reached, in a cyclic order, i.e., player i begins again from his first choice after his last choice gets repeated r_i times.

This protocol guarantees that all possible choice profiles containing player n 's choice u will be tried systematically by the remaining players $1, \dots, n - 1$, and that a winning choice profile will be reached with that procedure within at most N rounds, for N defined as in the theorem. Indeed, $\prod_{i=1}^{n-1} |C_i|$ is the total number of possible joint choices of players $1, \dots, n - 1$, and w_ℓ of them coordinate with the choice u , so in the worst case none of these w_ℓ joint choices will be explored until the end of the procedure, but the first one of them to be reached will result in coordination. \square

Clearly the upper bounds in Theorem 4.2 can be improved in many games G based on the particular structure of G . For example, in the extreme case where a player has a surely winning move, it obviously makes sense to choose it right away. Also, the upper bound in case 3 can be improved if not only the most senior player's choice is made optimally, but also the orderings of the choices of the other players are suitably fixed in some suitable order (say, the order of decreasing sizes of winning extensions), or if the most senior player has a unique best choice, that can be guessed by the other players, et cetera. Likewise, the upper bound in case 2 can be improved, but this requires more involved reasoning which we omit here.

Moreover, the number of rounds needed for coordination in the complete visibility case can be reduced essentially to a logarithm in

⁵ This ordering may be assumed to first list those choices that occur in the winning choice profiles of player n 's best choices. This assumption would generally improve the efficiency of the procedure, but it is not needed for the proof.

the number of agents if they can coordinate in gradually increasing groups where the groups coordinate with each other as if they were single agents; this idea is described in more detail in Section 4.4.

On the other hand, there are cases where the respective upper bounds for 1 and 3 are realized in the worst case. In the case of complete visibility, this is when the winning profiles are the branches of an unordered tree with $|C_n|$ successors of the root representing the choices of player n and each of these with $|C_{n-1}|$ successors representing the choices of player $n - 1$, et cetera. down to the leaves representing the choices of player 1. In the case of minimal visibility, the worst case is reached, e.g., in games where every choice belongs to a single winning choice profile.

4.2 Coordination based on patience thresholds

In scenarios with repeated games, it often makes sense for players to take time to ‘wait’ for their partners to move, for example to learn their strategic behaviour. However, the waiting must *stop at some stage*, after which the player moves to a new location in ‘search’ of the other players. This leads to the notion of a *patience threshold*.

Formally, the **patience threshold** of a player i is a positive integer $p(i)$ which gives an upper bound of the number of rounds which that player is willing to stay in the same location if no change occurs in his observable part of the game. Thus, a player with patience threshold k will stay in the same location while observing no change in the game for at most k consecutive rounds.

Assuming complete visibility, we will next design a protocol that guarantees coordination in all WLC-games without a hierarchy amongst the players, but where all players have fixed and *different* patience thresholds. We note that the players are not assumed to know the patience thresholds of each other. The assumption of different patience thresholds is reasonable in scenarios where the game moves are done in essentially continuous time. In such settings, the players could in general signal their intention to move in the next round once they have run out of patience, but since we assume no form of communication during play, such signalling is impossible. Nevertheless, this can be compensated by the different patience thresholds assumption. Alternatively, the assumption can be satisfied by design, when a team of artificial agents is designed with the intention to be able to coordinate without communication. While the assumption may seem ad hoc at first, it is in fact quite natural and common. Indeed, this is how many coordination problems, both in real life (e.g., the ‘pavement tango’) and in artificial scenarios (asynchronous distributed computing) are solved in practice.

Proposition 4.3. *Assuming complete visibility and different patience thresholds for each player, there is a protocol based on these patience thresholds which solves every n -player WLC-game in $N + 1$ rounds, where N is the sum of the patience thresholds of all the $n - 1$ players excluding the most patient one.*

Proof. The coordination protocol is similar to the one in the case of total hierarchy and complete visibility in Theorem 4.2. Suppose, w.l.o.g., that the patience thresholds of the players decrease with their natural order, i.e., $p(1) > \dots > p(n)$. The protocol prescribes that in the first round, all players make random choices. If they coordinate, the game ends. Otherwise, they all stay in their first locations for $p(n)$ rounds, after which player n runs out of patience and moves to another choice (location) u and remains there until the end of the play. Thereafter the play continues recursively in the subgame $G(u)$ where the remaining $n - 1$ players try to coordinate. When only the most patient player is left to move, she does not wait anymore but

coordinates with the rest in the next (and last) round. Clearly, the coordination is guaranteed to occur within $N + 1$ rounds. \square


Note that the assumption that all players have different patience thresholds is essential, as otherwise two players with the same patience threshold could start moving in a lock-step fashion and possibly never coordinate.

The visibility assumption is essential in Proposition 4.3, too, as the next simple example shows. Consider the game $G_{2,3}$ from Example 3.5 where both players have minimal visibility (i.e., they see and remember only their own choices). Suppose they start with non-matching choices and, when each of them runs out of patience, they keep moving in any order that preserves the mismatch forever. This endless discoordination scenario can occur regardless of the players’ patience thresholds, different or equal. Thus, coordination in this case is not guaranteed in any number of rounds.

4.3 Coordination by joining anchor groups

We call a group of players who have selected their choices from the same winning choice profile a **coordinating group**. Under the assumption of complete visibility, after every round the players can observe which coordinating groups have been formed by the previous choices. One natural mode of behaviour (also in real life scenarios) is that players try to coordinate with a coordinating group of *maximal size*, that is, with one of the **majority groups**; the players in majority groups remain stationary and wait for the others to coordinate with them, as long as they are still in a majority group. However, after each round the set of majority groups may change, so this behaviour is revised accordingly.

This leads to the **Majority group principle (MG)** which dictates that a player should pick a choice from a winning extension of one of the majority groups—unless he is part of a majority group, in which case he should simply repeat his previous choice.⁶ The 3-player game $G_{3,2}$, displayed below, is solvable by MG in 2 rounds as a majority group of size 2 (at least) is formed in the first round. (Cf. the example given in the introduction.)



The 5-player game $G_{5,3}$ (where there are 3 disjoint winning profiles) is likewise solvable by MG in 3 rounds: after the first round there could be two majority groups of size 2, but then the player outside these groups will join one of them and thus there will be a majority group of size 3 after round 2.

Many natural variations and extensions can be defined for the MG. For instance, if there is exactly one group of players whose winning extension is a Cartesian product of all possible choices (so, coordination with that group is guaranteed to succeed), then the **Unique winning group principle** instructs all other players to coordinate with that group.

The idea to coordinate by identifying a unique group, under the assumption of complete visibility, can be extended further to the concept of *anchor group*, or *focal group*. That is a group of players that have already coordinated with each other, being on some winning choice profile, and thereafter remains stationary while all others try

⁶ This principle leads to a deadlock if at some point all players are clustered in majority groups of the same size. In order to avoid this, one could modify this principle so that in such cases the players make a random choice to break the symmetry—possibly after their *patience* runs out.

to coordinate with that group. Various real life coordination scenarios are often based on such behaviour (e.g., a group that is in some special location, or that consists of the most ‘senior’ members, etc.). This idea leads to a family of principles using *unique focal groups*.

4.4 Gradual coordination by contact

Here we briefly discuss a variant of repeated coordination games where the players can gradually *establish communication contacts* with each other in the course of the play. This can be advantageous, especially in scenarios with limited visibility, e.g., local visibility. Making a contact could be achieved in several ways, but here it amounts to selecting choices on the same winning profile. After making contact, the players can communicate with each other, agree on some specific conventions, and thus thereafter synchronise their further actions by using a *collective protocol*. The communication allows players to select their joint actions freely and thus allows any group, that has made contact, to violate structurality in their protocols.

Suppose, e.g., a scenario of meeting in a city with the assumption that the players can see each other and make contact if and only if they meet at the same place. By making a contact, two (or more) players could agree for example to:

1. move together to look for the others, or
2. decide to meet again ‘at the same location’ after an agreed number of rounds, and then synchronise further again, or
3. decide that some of them ‘wait’ at the current location (by repeating their choices) while the others search for the remaining players, etc.

Assuming that the group of all players has a ‘leader,’ coordination by contact becomes quite easy: the leader chooses a winning profile and stays stationary, while everyone else searches for the leader until they make contact with him, and then finally everyone gathers on some chosen profile. The number of required rounds is not more than the total number of choices.

We next present another natural example where coordination by contact can be successful, under the additional assumption that, *once the players have established contact, they can maintain it* from a distance for the rest of the game (e.g., by exchanging mobile numbers).

Proposition 4.4. *Consider a repeated WLC-game G with n players, a maximum of m possible choices per player, and local visibility. Then any protocol for coordination by maintaining contacts that guarantees that some (any) two players will eventually meet on a winning profile within N rounds, can be modified to a protocol that guarantees coordination of all players in at most $N + m$ rounds.*

Proof. Suppose there is a protocol Σ_0 that guarantees that some two players will eventually meet on a winning profile. The final coordination protocol Σ is defined as follows. First Σ_0 is followed for N rounds. This guarantees that at least 2 players end up in the same winning profile. If there are several winning profiles where the group ends up, one is distinguished, called hereafter their ‘home base.’ There may be more than one such group, and groups can keep growing in the procedure. Once a group has formed, the players in it keep repeating their choices (stay in the ‘home base’) until N rounds have passed from the beginning of the game play.

After the first N rounds, every formed group starts collecting other players as follows. All but one of the players in the group remain in their ‘home base’ (i.e., start repeating their choices) as ‘stationary


players,’ whereas the remaining one, which we call the ‘searching player,’ starts going through all the other choices in some order. The searching player sends all players who he finds to his home base (i.e., instructs them to make the respective choices from that profile). Every player who after round N ends up ‘alone’ (not sharing any winning profile with other agents) must remain there (keep repeating their choice) until he is ‘visited’ by a searching player on some shared winning profile, and then follow the searching player’s instructions.

When two (or more) searching players, say A and B, meet on a winning profile, they merge their groups into one, e.g., into the group of A. Then player B instructs all players who are at his home base to move to the home base of A, whereas A continues visiting all other choices. Thus, an arbitrary winning profile is visited by at least one searching player within m rounds and all players from there are directed to a location where they join the others. \square

The case when the number of choices n is greater than the maximum number of choices m per player is particularly natural here. By the pigeonhole principle it follows that some players inevitably make contact in the first round and thus, by Proposition 4.4, it takes at most $m + 1$ rounds for everyone coordinate.

We also note that the protocol described in the proof of Proposition 4.4 can be improved when larger groups of players make a contact. The search can then be done more quickly, e.g., when only one player in the group remains stationary and the others divide the search in a suitable way.

We note that the assumption of maintaining contacts is essential for proving Proposition 4.4. Indeed, consider the 4-player game in the next figure. One possible way to guarantee contact of two players without communication is that players 1 and 2 apply symmetry-based reasoning to establish contact on the dashed-line profile on top, while players 3 and 4 likewise establish contact on the dashed-line profile in the bottom, already in the 1st round.



Thereafter, the searching players for both groups may continue acting in a completely symmetric way, thus never meeting each other, but each of them only meeting the stationary player of the other group. Then, if they cannot maintain contacts within the two groups, they can never agree on where to merge the groups.

We then present a coordination-by-contact protocol, using a priority hierarchy of players and complete visibility, that significantly improves the coordination times of Theorem 4.2. The speed-up could be crucial in scenarios involving large numbers of players.

Proposition 4.5. *In any WLC-game with n players with complete visibility and priority hierarchy the players can coordinate by contact in $1 + \lceil \log_2 n \rceil$ steps.*

Proof. The proof idea is to use the priority hierarchy in the following way, again assuming, w.l.o.g. that the hierarchy is $1 < \dots < n$. In the first round, the players choose randomly. If they coordinate, the game ends. Else, in the second round, every odd-number ranked player remains stationary and each of the other players with rank $2i$ coordinate with the stationary player $2i - 1$ ‘to the left’ of them, i.e., the player $2i$ chooses from some winning extension of the choice of $2i - 1$. This creates coordinated groups of two players; if n is odd, the

last group is the singleton $\{n\}$. In the third round, every second group (i.e., the groups $\{1, 2\}$, $\{5, 6\}$, $\{9, 10\}$, ...) remains stationary and the remaining groups coordinate with the stationary group 'to the left', i.e., $\{3, 4\}$ moves to some winning extension of $\{1, 2\}$, etc. Note that here the players 3, 4 are assumed to be able to proceed to the same winning extension of $\{1, 2\}$ based on their contact. This coordination of groups of pairs creates coordinated groups of four (with again possible exceptions at the right end of the hierarchy). Continuing this way, it is easy to see that all players can eventually coordinate in $1 + \lceil \log_2 n \rceil$ steps. \square

5 CONCLUDING REMARKS

We list some of the main conclusions of the present work:

- The study of repeated WLC-games extends essentially the study of one-step WLC-games in [8].
- Many repeated WLC-games can be solved by protocols agreed in advance but with no communication during the play, while their one-shot versions are unsolvable.
- Under some natural additional assumptions, such as a commonly known total hierarchy amongst the players or different patience thresholds and complete visibility, etc., all repeated WLC-games become solvable and optimising the time for guaranteed coordination becomes the main question.

Finally, we list some ongoing/future work:

- analysis of *expected* coordination times in repeated WLC-games where coordination cannot necessarily be guaranteed,
- search for protocols which give *optimal* guaranteed (or expected) coordination times in different scenarios,
- analysis of *boundedly repeated* WLC-games, where coordination must occur within a prescribed number of rounds.

ACKNOWLEDGEMENTS

The work of Valentin Goranko was partly supported by a research grant 2015-04388 of the Swedish Research Council. The work of Antti Kuusisto is funded by the Academy of Finland grants 324435 and 328987. We acknowledge Lauri Hella for pushing forward the idea of using our symmetry principles from [8] also in the repeated scenario for solving, in particular, the games $G_{2,m}$ in repeated attempts for odd m . We also thank the anonymous referees for valuable comments.

REFERENCES

- [1] Douglas Adams and John Lloyd, *The Meaning of Liff: The Original Dictionary Of Things There Should Be Words For*, Pan Books, 1983.
- [2] Robert J. Aumann and Michael B. Maschler, *Repeated Games with Incomplete Information*, MIT Press, 1995.
- [3] Cristina Bicchieri, *Rationality and Coordination*, Cambridge University Press, 1994.
- [4] Vincent P Crawford and Hans Haller, 'Learning How to Cooperate: Optimal Play in Repeated Coordination Games', *Econometrica*, **58**(3), 571–595, (May 1990).
- [5] David Gauthier, 'Coordination', *Dialogue*, **14**(02), 195–221, (1975).
- [6] Margaret Gilbert, 'Rationality, coordination, and convention', *Synthese*, **84**(1), 1–21, (1990).
- [7] Valentin Goranko, Antti Kuusisto, and Raine Rönholm, 'Rational coordination in games with enriched representations', in *Multi-Agent Systems and Agreement Technologies - 15th European Conference, EUMAS 2017, and 5th International Conference, AT 2017, Revised Selected Papers*, eds., Francesco Belardinelli and Estefania Argente, volume 10767 of *Lecture Notes in Computer Science*, pp. 323–338. Springer, (2017).
- [8] Valentin Goranko, Antti Kuusisto, and Raine Rönholm, 'Rational coordination with no communication or conventions', in *Proc. of LORI VI*, volume 10455 of *LNCS*, pp. 33–48. Springer, (2017).
- [9] Valentin Goranko, Antti Kuusisto, and Raine Rönholm, 'Rational coordination with no communication or conventions', Technical report, (2017). arXiv:1706.07412, under submission.
- [10] John C. Harsanyi and Reinhard Selten, *A General Theory of Equilibrium Selection in Games*, MIT Press, 1988.
- [11] D. Lewis, *Convention, A Philosophical Study*, Harvard University Press, 1969.
- [12] George J. Mailath and Larry Samuelson, *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press, 2006.
- [13] Thomas Schelling, *The strategy of conflict*, Harvard University Press, 1960.
- [14] Robert Sugden, 'Spontaneous order', *J. of Economic Perspectives*, **3**(4), 85–97, (1989).
- [15] Robert Sugden, 'Rational Choice: A Survey of Contributions from Economics and Philosophy', *Economic Journal*, **101**(407), 751–785, (July 1991).
- [16] Robert Sugden, 'A theory of focal points', *The Economic Journal*, 533–550, (1995).