Gradual guaranteed coordination in repeated win-lose coordination games

Valentin Goranko
Stockholm University

Antti Kuusisto
University of Helsinki

Raine Rönnholm
Université Paris-Saclay

ECAI 2020

What this work is about

Consider a scenario where a group of agents have the shared goal to coordinate but they cannot directly communicate during the task.

The main question: how can the agents learn to coordinate gradually in repeated unsuccessful coordination attempts?

Further issues: how can different natural features (such as agent hierarchies, different patience levels, etc.) help the coordination attempts.

Win-lose coordination games (**WLC**-games)

A **WLC**-game is an n-player strategic form game with a payoff matrix where every outcome is either

$(1, \ldots, 1)$ (winning for everyone) or
$(0, \ldots, 0)$ (losing for everyone).

An n-player **WLC**-game can be represented as an n-partite hypergraph where

- each player has a column of choice nodes.
- winning choice profiles are n-ary edges.

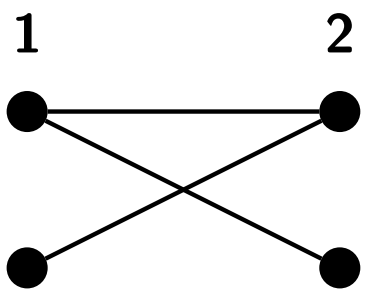

A two-player game with three winning choice profiles

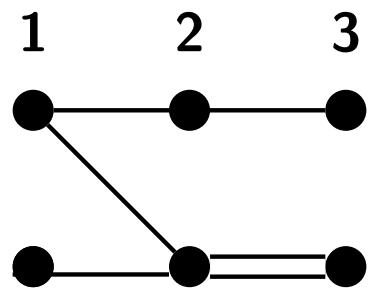

A three-player game with three winning choice profiles

Repeated **WLC**-games

Every **WLC**-game G generates a respective repeated **WLC**-game which consists of consecutive one-step plays of G until (if ever) the players reach a winning choice profile.

Example: Pavement Tango

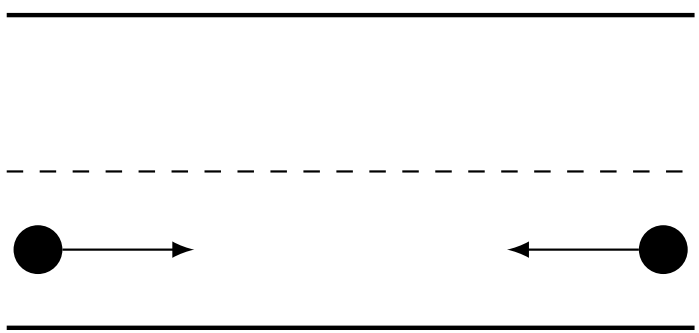

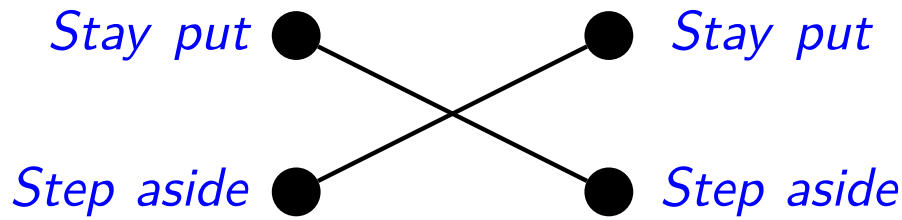

Protocols for repeated **WLC**-games

A protocol is a specification of a fixed nodeterministic strategy for all players in all repeated **WLC**-games.

We require that protocols must be structural. Intuitively, this means that protocols are independent of the *names of choices and player roles*.

For example, in the game below, all choices are *structurally symmetric* and thus no protocol can guarantee coordination in a single round.

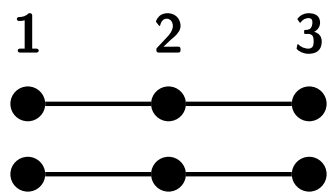

However, if the players can see and remember the earlier moves in a repeated game, then the structural symmetries between choices can be gradually reduced.

Gradual guaranteed coordination

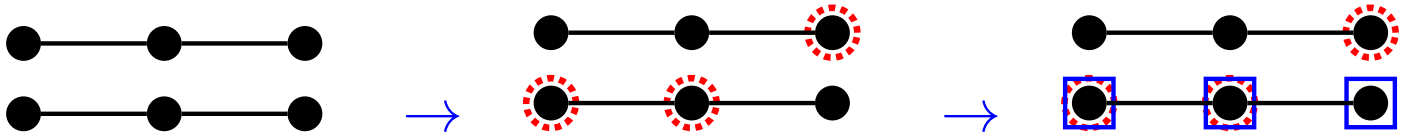

Here coordination could not be guaranteed in a single round, but the players could gradually guarantee coordination. The same happens below.

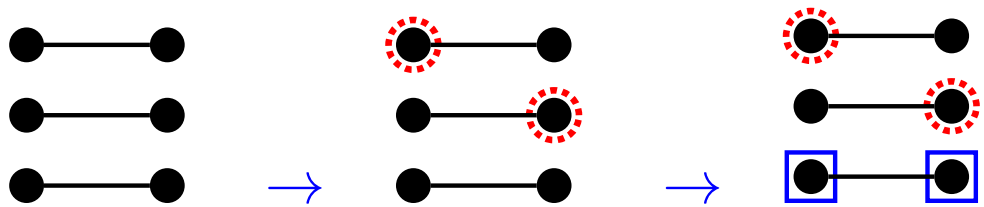

The Pavement Tango is a canonical example of a game where coordination can never be guaranteed with certainty (due to symmetries).

Coordination with additional features

We show that coordination can be gradually guaranteed in *every* repeated **WLC**-game with protocols that can utilize some natural extra features, such as:

- *priority hierarchies amongst players*,
- *different patience thresholds of players*,
- *communication links created via contacts*.

We also consider how different levels of visibility (of other players' previous moves) affect players' ability to coordinate.

Hierarchical coordination

Theorem

*Consider an n-player **WLC**-game with a commonly known total hierarchy amongst the players. Let m denote the maximum number of choices per player. Then there is a protocol guaranteeing coordination in*

1. *n rounds if the players have complete visibility of other players' moves,*
2. *$O(m^n)$ rounds under no visibility of other players' moves.*
3. *$O(mn)$ rounds under 'local visibility', an intermediate notion of visibility where players observe the other players' moves coordinating with their current move.*

Coordination based on patience thresholds

In many real-life scenarios, players take time to wait for their partners to move in order to learn their strategic behaviour.

The patience threshold of a player i is the number of rounds that the player is willing to repeat the same choice if no significant change occurs in the part of the game the player sees.

Proposition

Assuming complete visibility and different patience thresholds, there is a protocol solving every n-player game in $N + 1$ rounds, where N is the sum of the patience thresholds the players excluding the most patient one.

Gradual coordination by contact

Consider a group of agents trying to find each other in the city. Assume that whenever two or more agents meet at the same location (i.e., same winning profile), they can remain in contact (i.e., maintain a communication link) thereafter.

Proposition

*Consider a repeated **WLC**-game G with*

1. *n players that can make contacts,*
2. *max m choices per player,*
3. *and local visibility.*

Suppose there exists a protocol guaranteeing that at least two players will eventually meet on a winning profile within N rounds. Then this protocol can be modified to a protocol that guarantees coordination in $N + m$ rounds.

Proof idea. Use local 'home base' groups and search parties of agents for gradually gathering moves to the same edge.

Coordination by contact

Proposition

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

Proof idea. In successive rounds, use the player hierarchy to create increasing groups of size 2, 4, 8, 16...

Ongoing and future work

- analysis of *expected* coordination times in repeated **WLC**-games where coordination cannot necessarily be guaranteed.
- search for protocols giving *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.

Previous work

One-shot **WLC**-games introduced and studied in:

Valentin Goranko, Antti Kuusisto, Raine Rönnholm: *Rational Coordination with no Communication or Conventions*. LORI 2017: 33-48.

Valentin Goranko, Antti Kuusisto, Raine Rönnholm: *Rational Coordination in Games with Enriched Representations*. EUMAS/AT 2017: 323-338.

Valentin Goranko, Antti Kuusisto, Raine Rönnholm: *Rational Coordination with no Communication or Conventions*. Journal of Logic and Computation, to appear, 2020.

Thanks!