Asynchronous Best-Reply Dynamics

Noam Nisan[1], Michael Schapira[2], and Aviv Zohar[2]

[1] Google Tel-Aviv and The School of Computer Science and Engineering, The
Hebrew University of Jerusalem, Israel.
[2] The School of Computer Science and Engineering, The Hebrew University of
Jerusalem, Israel.
{noam,mikesch,avivz}@cs.huji.ac.il

Abstract. In many real-world settings (e.g., interdomain routing in the
Internet) strategic agents are instructed to follow best-reply dynamics
in asynchronous environments. In such settings players learn of each
other's actions via update messages that can be delayed or even lost. In
particular, several players might update their actions *simultaneously*, or
make choices based on outdated information. In this paper we analyze
the convergence of best- (and better-)reply dynamics in asynchronous
environments. We provide sufficient conditions, and necessary conditions
for convergence in such settings, and also study the convergence-rate of
these natural dynamics.

1   Introduction

Many real-life protocols can be regarded as executions of best-reply dynamics,
i.e, players (computational nodes) are instructed to repeatedly best-reply to
the actions of other players. In many cases, like Internet settings, this occurs
in *asynchronous* environments: Think of the players as residing in a computer
network, where their best-replies are transmitted to other players and serve as
the basis for the other players' best-replies. These update messages that players
send to each other may be delayed or even lost, and so players may update
their actions simultaneously, and do so based on outdated information. Perhaps
the most notable example for this is the Border Gateway Protocol (BGP) that
handles interdomain routing in the Internet. As observed in [1], BGP can indeed
be seen as an execution of best-reply dynamics in asynchronous environments.

Asynchronous best-reply dynamics. The most fundamental question re-
garding best-reply dynamics in asynchronous settings is "When are such dy-
namics guaranteed to converge?". This will certainly not happen if a pure Nash
equilibrium does not exist, but is not guaranteed even in very simple and well-
structured games that have a pure Nash. We present a formal framework for the
analysis of best-reply dynamics in asynchronous environments. We then exhibit
a simple class of games for which convergence to a *unique* pure Nash equilib-
rium is guaranteed. We term this class, which contains all strictly-dominance-
solvable games (games where iterated elimination of strictly dominated strategies
leaves a single strategy profile [2]), "*max-solvable-games*". We also discuss the

convergence-rate of best-reply dynamics in asynchronous settings. We propose a notion of an *asynchronous phase*, and show that for max-solvable games convergence also happens quickly.

Theorem: *Best reply-dynamics converge within $\sum_i m_i$ phases for every max-solvable game, and in every asynchronous schedule. Here m_i is the size of the strategy space of the i'th player. In particular, this holds for all strictly-dominance-solvable games.*

This theorem shows that even though the "input" (a normal-form representation of a max-solvable game) is of exponential size (in the size of the strategy-spaces), best-reply dynamics converges in a linear number of phases.

We consider a generalization of max-solvable games, called "weakly-max-solvable games" that contains the class of weakly-dominance-solvable games (games where iterated elimination of weakly-dominated strategies leaves a single strategy profile [2]). For this class of games we show that no similar result holds; not only are best-reply dynamics not guaranteed to converge, but any procedure for finding a pure Nash equilibrium faces a severe obstacle.

Theorem: *Finding a pure Nash equilibrium in weakly-max-solvable games requires exponential communication in $\sum_i m_i$. This is even true for the more restricted class of weakly-dominance-solvable games.*

This result follows the line of research initiated by Conitzer and Sandholm [3], and further studied in the work of Hart and Mansour [4].

Asynchronous better-reply dynamics. At this point we turn our attention to better-reply dynamics. Now, players are not required to continuously best-reply to the strategies of the others, but merely to always choose strategies that are better replies than the ones they currently have. Once again, we are interested in figuring out when these dynamics converge in asynchronous settings. A natural starting point for this exploration is the well-known class of *potential games*, introduced by Monderer and Shapley [5], building on the seminal work of Rosenthal [6]. For these games, it is known that better-reply dynamics are guaranteed to converge (if players update their strategies one by one, and learn of each other's action immediately).

We show, in contrast, that even for these games asynchrony poses serious challenges and may even lead to persistent oscillations. We consider a restricted, yet expressive, form of asynchrony – settings in which players may update strategies simultaneously (and not necessarily one by one), but update messages arrive at their destinations immediately (no delay). We call such restricted asynchronous settings *"simultaneous settings"*. We prove the following theorem:

Theorem: *If every subgame of a potential game has a unique pure Nash equilibrium then better-reply dynamics are guaranteed to converge for every simultaneous schedule. (By subgame, we mean a game that is the result of elimination of players' strategies from the original game.)*

In fact, we show that this result is almost a characterization, in the sense that the uniqueness of pure Nash equilibria in every subgame is also a necessary condition for convergence in simultaneous settings for a large subclass of potential games.

Organization of the Paper: In Section 2 we present a model for analyzing best- and better-reply dynamics in asynchronous settings. In Section 3 we present and discuss max-solvable games. In Section 4 we explore potential games. Due to space constraints many of the proofs are omitted (see [7] for a full version).

2 Synchronous, Simultaneous, and Asynchronous Environments

We use standard game-theoretic notation: Let G be a *normal-form game* with n players $1, 2, \ldots, n$. We denote by S_i the (finite) *strategy space* of the i'th player. Let $S = S_1 \times \ldots \times S_n$, and let $S_{-i} = S_1 \times \ldots \times S_{i-1} \times S_{i+1} \times \ldots \times S_n$ be the cartesian product of all strategy spaces but S_i. Each player i has a *utility function* u_i that specifies i's *payoff* for any strategy-profile of the players. For any strategy $s_i \in S_i$, and every $(n-1)$-tuple of strategies $s_{-i} \in S_{-i}$, we shall denote by (s_i, s_{-i}) the strategy profile in which player i plays $s_i \in S_i$ and all other players play their strategies in s_{-i}. Given $s_{-i} \in S_{-i}$, $s_i \in S_i$ is said to be a best reply to s_{-i} if $u_i(s_i, s_{-i}) = max_{s_i' \in S_i} u_i(s_i', s_{-i})$. Given $s_{-i} \in S_{-i}$, $s_i' \in S_i$ is said to be a better-reply of player i than $s_i \in S_i$ if $u_i(s_i', s_{-i}) > u_i(s_i', s_{-i})$.

Consider the following best-reply dynamics procedure: We start with an initial strategy profile of the players $s \in S$. There is set of rounds $R = \{1, 2, \ldots\}$ of *infinite* size. In each round one *or more* players are chosen to participate. Every player chosen to participate must switch to a best-reply to his *most recent information* about the strategies of the other players, and send update messages to all other players announcing his strategy (a player must announce his strategy to all other players even if it did not change).

As in [1], there is an adversarial entity called the *Scheduler* that is in charge of making the following decisions: Choosing the initial strategy profile $s \in S$. Determining which players will participate in which round (a function f from R to subsets of the players). Determining when sent update messages reach their destinations (see below). The Scheduler must be restricted not to indefinitely starve any player from best-replying (that is, each player participates in infinitely many rounds). We shall name all the choices made by the Scheduler a *schedule*. We distinguish between three types of settings:

Synchronous settings: In these settings, the Scheduler can only choose *one* player to play in each round (that is, $|f(r)| = 1$ for any $r \in R$). In addition, update messages sent by players arrive at their destinations *immediately* (that is, at the end of the round in which they were sent). Hence, players' actions are *observable* to other players. Observe, that a game is a potential game iff for each of its subgames, better-reply dynamics are guaranteed to converge to a pure Nash equilibrium for any synchronous schedule.

Simultaneous settings: In simultaneous settings, the Scheduler can choose *any* number of players to play in each round ($|f(r)|$ can be any number in $1, 2, \ldots, n$ for any $r \in R$). As in synchronous settings, players' actions are observable (update messages sent by players arrive at their destinations immediately).

Asynchronous settings: As in simultaneous settings, the Scheduler can choose any number of players to play in each round. However, in asynchronous settings the Scheduler can also decide when each sent update message arrives at its destination (at the end of the round in which it was sent or in some subsequent round) subject to the limitation that messages that were sent earlier arrive before later ones. It can also decide to drop update messages. The Scheduler may not prevent all update messages of a player from reaching another player indefinitely.

Elementary examples (like the "Battle of the Sexes" game) show that even in very simple games, in which best-reply dynamics are guaranteed to converge in synchronous settings, they might not converge in simultaneous settings (and, in particular, in asynchronous settings). Similarly, it can be shown that convergence of best-reply dynamics in simultaneous settings does not imply convergence in asynchronous settings.

In an analogous way, we can now define synchronous, simultaneous, and asynchronous convergence of better-reply dynamics.

3 Max-Solvable Games

In this section we present a class of games called "max-solvable games" for which best-reply dynamics are guaranteed to converge to a pure Nash equilibrium even in asynchronous settings. We then discuss a generalization of these games, that contains all dominance-solvable games (games in which the iterated removal of dominated strategies results in a single strategy profile).

3.1 Max-Solvable Games - Definitions

We start by defining max-solvable games.

Definition 1. *A strategy $s_i \in S_i$ is max-dominated if for every strategy-profile of the other players $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$ there is a strategy s_i' such that $u_i(s_i', s_{-i}) > u_i(s_i, s_{-i})$.*

That is, a strategy of a player is max-dominated if it is not a best-reply to *any* strategy-profile of the other players. Observe, that every strictly dominated strategy is max-dominated. In fact, a strategy is max-dominated even if it is strictly dominated by a *mixed* strategy.

Informally, a max-solvable game is a game in which the iterated elimination of max-dominated strategies results in a single strategy-profile.

Definition 2. *A game G is said to be max-solvable if there is a sequence of games $G_0, \ldots, G_r$ such that:*

- $G_0 = G$
- *For every $k \in \{0, \ldots, r-1\}$, G_{k+1} is a subgame of G_k achieved by removing a max-dominated strategy from the strategy space of one player in G_k.*
- *The strategy space of each player in G_r is of size 1.*

The class of max-solvable-games contains all strictly-dominance-solvable ones. We shall refer to an elimination order of max-dominated strategies, that results in a single strategy-profile as an *elimination sequence* of a max-solvable game.

3.2   Asynchronous Best-Reply Dynamics and Max-Solvable Games

One of the helpful features of max-solvable games is the fact that such games always have a unique pure Nash equilibrium.

Proposition 1. *Any max-solvable game has a unique pure Nash equilibrium.*

We now show that in max-solvable games, best-reply dynamics always converge to the unique pure Nash equilibrium, even in asynchronous settings. How long does this take? Answering this question requires further clarifications as we must account for the fact that update messages can be arbitrarily delayed, and that players might be prevented from best-replying for long periods of time. We define an *asynchronous phase* to be a period of time in which *every* player is activated at least once, and every player receives at least one update message from *each* of his neighbours. We prove that, for any asynchronous schedule, best-reply dynamics converge to the unique pure Nash equilibrium in a number of asynchronous phases that is at most $\sum_i m_i$, where m_i is the size of the strategy space of the i'th player.

Theorem 1. *In any max-solvable game, best-reply dynamics converges for every asynchronous schedule within $\sum_i m_i$ asynchronous phases.*

Proof. Consider an elimination sequence of max-dominated strategies that results in a single strategy-profile. Let strategy s_1 of some player i be the first strategy to be eliminated. Player i is activated once during the first asynchronous phase. If he is playing s_1 then he will switch to another strategy since s_1 is max-dominated. Furthermore, no best-reply of player i in the future will *ever* cause him to choose strategy s_1. From this point onwards, the best-reply dynamics are effectively occurring in a game where s_1 does not exist. Let us now consider the next strategy in the elimination order s_2, which belongs to some player j (that can be i, or some other player). Given that player i never plays s_1, s_2 is now max-dominated. Player j is activated during the second asynchronous phase. If he is playing s_2 he will move to another strategy. No matter what, s_2 will never be played again. More generally, after k asynchronous phases the k'th strategy in the elimination order will never be played again. Therefore after $\sum_i (m_i - 1)$ asynchronous phases we are bound to reach the pure Nash equilibrium, which is the remaining strategy-profile.

3.3 Weakly-Max-Solvable-Games

The definition of max-dominated strategies required that, for any strategy-profile of the other players, a max-dominated strategy be *strictly* worse than another strategy. In this section we discuss the case of ties.

Definition 3. *A strategy $s_i \in S_i$ is weakly-max-dominated if for every strategy-profile of the other players $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$ there is another strategy s_i' such that $u_i(s_i', s_{-i}) \geq u_i(s_i, s_{-i})$.*

Now, we can define *weakly-max-solvable games* as games in which the iterative removal of weakly-max-dominated strategies results in a single strategy-profile. Observe that any weakly-dominance-solvable game is a weakly-max-solvable game. Unfortunately, as the following example demonstrates, best-reply dynamics are not guaranteed to converge even in weakly-dominance-solvable games.

Example 1. Consider the game depicted by the following matrix (the rows are player 1's strategies and the columns are player 2's strategies):

1,1	0,0
1,0	0,1
0,1	1,0

First, observe that this is indeed a weakly-dominance-solvable game. Observe that if the initial strategy-profile is the leftmost entry in the lower row (row 3) of the game-matrix, then the following best-reply dynamics is possible: Player one moves from row 3 to row 2, player 2 moves from the left column to the right one, player 1 moves from row 2 to row 3, player 2 moves from the right column to the left one, and so on.

Weakly-dominance-solvable games always have pure Nash equilibria. As we have just seen, best-reply dynamics are not guaranteed to converge to such an equilibrium. Is there a different procedure that can do so in reasonable time? We prove the following impossibility result:

Theorem 2. *Finding a pure Nash equilibrium in games that are weakly-dominance-solvable requires communicating exponentially many bits (in $\sum_i m_i$).*

4 Potential Games and Asynchrony

In this section we explore better-reply dynamics in the context of potential games. While it is easy to see that in potential games better-reply dynamics converge for any synchronous schedule, what happens in simultaneous and asynchronous environments? We study the structural properties of potential games for which convergence of better-reply dynamics in simultaneous settings is assured.

We prove the following theorem:

Theorem 3. *If every subgame of a potential game has a unique pure Nash equilibrium, then better-reply dynamics converge for any simultaneous schedule.*

We show that the uniqueness of pure Nash equilibria in every subgame of a potential game is almost a characterization of potential games for which better-reply dynamics always converge in simultaneous settings. We show this by proving that this is indeed also a necessary condition for a large subclass of potential games, we term "*strict* potential games".

Definition 4. *A game G is strict if for any two strategy profiles $s = (s_1, ..., s_n)$ and $s' = (s'_1, ..., s'_n)$, such that there is some $j \in [n]$ for which $s' = (s'_j, s_{-j})$, $u_j(s) \neq u_j(s')$.*

That is, a game is strict if for any player i, for any two strategies of that player $s_i, s'_i \in S_i$, and for any strategy-profile of the other players s_{-i}, i strictly prefers one strategy over the other. A strict potential game is a potential game that is strict.

Theorem 4. *If a strict potential game is such that better-reply dynamics converge for any simultaneous schedule, then every subgame of that games has a unique pure Nash equilibrium.*

Remark 1. One might hope that any strict game in which every subgame has a unique pure Nash equilibrium is a potential game. However, in the full version [7] of the paper we give an example that shows that this is not the case.

What about asynchronous settings? We now show that the property that guarantees the convergence of best-reply dynamics in a potential game (i.e., that every one of its subgames has a unique pure Nash equilibrium) does not necessarily guarantee convergence in asynchronous schedules.

Example 2. Consider the game described by Fig. 1. The arrows describe the better-replies of players from any strategy-profile (an arrow between strategy-profiles denotes the transition caused by a best-reply update of a single player).

Fig. 1. A game in which better-reply dynamics might diverge for some asynchronous schedule

The reader can verify that this is a potential game and that every subgame has a unique Nash equilibrium. Recall, that in asynchronous settings, the Scheduler may delay messages. We shall show that better-reply dynamics may never converge in such settings. Let us show such an oscillation (messages arrive immediately unless specifically noted): We begin with state A and allow the row player to update his strategy and notify everyone, thus arriving at state C. We then activate the column player and the matrix player simultaneously and arrive at state H. However, we delay the message sent to the row player by the matrix player so that the row player in fact believes we are in state D. We then activate the row player and allow him another update. He believes he moves to state B while in fact we arrive at state F. We then release the message to the row player and invoke the column player which updates his strategy from F to E. Then, the matrix player is activated and we return to state A. Repeating this over and over gives a permanent oscillation.

Acknowledgements

References

1. Levin, H., Schapira, M., Zohar, A.: Interdomain routing and games. In: Proceedings of STOC 08'
2. Osborne, M.J., Rubinstein, A.: A Course in Game Theory. MIT Press (1994)
3. Conitzer, V., Sandholm, T.: Communication complexity as a lower bound for learning in games
4. Hart, S., Mansour, Y.: The communication complexity of uncoupled Nash equilibrium procedures. In: Proceedings of STOC 2007
5. Monderer, D., Shapley, L.: Potential games. Games and Economic Behavior (14) (1996) 124–143
6. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. Int. J. Game Theory (2) (1973) 65–67
7. Nisan, N., Schapira, M., Zohar, A.: Asynchronous best-reply dynamics. Technical report, The Leibnitz Center for Research in Computer Science (2008)