Nash Equilibrium in Generalised Muller Games

Soumya Paul[1] **and Sunil Simon**[2]

[1] The Institute of Mathematical Sciences
C.I.T. Campus, Chennai 600 113, India.
soumya@imsc.res.in

[2] Institute for Logic, Language and Computation
University of Amsterdam.
s.e.simon@uva.nl

ABSTRACT. We suggest that extending Muller games with preference ordering for players is a natural way to reason about unbounded duration games. In this context, we look at the standard solution concept of Nash equilibrium for non-zero sum games. We show that Nash equilibria always exists for such generalised Muller games on finite graphs and present a procedure to compute an equilibrium strategy profile. We also give a procedure to compute a subgame perfect equilibrium when it exists in such games.

1 Introduction

Infinite two player games on graphs have been shown to have various applications in different branches of mathematics and computer science. These are games played on a directed graph where players take turns to move and trace out a path in the graph. The winning condition is given by a set of infinite paths. Such games are well studied in descriptive set theory and topology in the form of Banach-Mazur games. In computer science these are commonly used as models of games in verification and synthesis of open reactive systems. The key question of interest for such games is that of determinacy. That is, whether one of the players always has a winning strategy. It turns out that determinacy depends crucially on the topological properties of the winning set. A celebrated result by Martin [7] showed that all games with Borel winning conditions are determined.

Martin's result however, does not make any assertion as to whether it is possible to determine who the winner is or how "complex" the winning strategy is. These turn out to be the core questions in solving the verification and synthesis problems as well. Winning conditions for games which arise in computer science are typically specified as logical formulas in S1S, first order logic or LTL and are therefore regular conditions. A seminal result due to Büchi and Landweber [1] says that for games played on finite graphs where the wining condition is specified as a Muller set, the winner can be determined and that the winning strategy can be effectively synthesised in finite memory strategies.

A natural generalisation of two player zero sum games is multi-player games where each player has a win-lose objective. Players' objectives are allowed to overlap and therefore these define non-zero sum games. For non-zero sum games, determinacy is usually replaced

Foundations of Software Technology and Theoretical Computer Science (Kanpur) 2009.
Editors: Ravi Kannan and K. Narayan Kumar; pp 335–346
Leibniz International Proceedings in Informatics (LIPIcs), Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany.
Digital Object Identifier: 10.4230/LIPIcs.FSTTCS.2009.2330

by one of the most important solution concepts in game theory, that of Nash equilibrium: a strategy profile where none of the players have an incentive to deviate unilaterally. It has been shown in [3] that every multiplayer game with Borel winning condition has a Nash equilibrium (see [2] for a detailed exposition). The main idea here is the effective use of threat strategies whereby a player deviating from the equilibrium profile is punished by others to receive the outcome which she can guarantee on her own. For games where the win-lose objectives are regular, an equilibrium profile can be effectively synthesized as well.

For games of infinite duration, it is questionable whether Nash equilibrium defines a satisfactory notion of rational behaviour. A more refined concept is that of subgame perfect equilibrium which takes into account perturbations of players as well. The existence of subgame perfect equilibrium for multiplayer games with win-lose objectives was shown in [9]. [5] unifies both results and shows that the crucial requirement for the existence of equilibrium for such multiplayer games is the determinacy of the underlying two player games.

From a game theoretic perspective, it is quite natural to consider games where players have utilities associated with plays rather than just win-lose conditions. We suggest that in case we restrict our attention to classifying regular plays then this can be captured in terms of a generalised Muller game. These are Muller games where instead of interpreting the Muller table as defining win-lose conditions, we associate utilities over the sets in the Muller table. Such games define non-zero sum objectives for players and we can therefore ask the question whether Nash equilibrium always exists for this class of games. In this context we show the following results:

- Nash equilibrium always exists for generalised Muller games played on finite graphs.
- An equilibrium profile can be effectively synthesized.

One could employ threat strategies to show the existence of equilibrium. However, for infinite games with non-zero sum objectives, even coming up with rationality assumptions which justify the use of such strategies is a challenging task. On the other hand, backward inductive equilibrium profiles are known to be more versatile in the case of finite games. We show that the standard backward induction algorithm [10] can be effectively used to prove the existence of Nash equilibrium and to synthesize an equilibrium profile in generalised Muller games.

Subgame perfect equilibria in general need not exists for such games. However, we show that:

- It is decidable to check whether subgame perfect equilibrium exists in a generalised Muller game.
- It is possible to effectively synthesize a subgame perfect equilibrium profile (when it exists).

2 Preliminaries

We begin with a description of the game arena and the objectives of the players. We look at unbounded duration, turn based games played on finite graphs.

2.1 Game Arena

A game G consists of an arena $\mathcal{A}$ and an objective *Win*. For a directed graph $\mathcal{A} = (V, E)$ and for a node $v \in V$, let $vE = \{v' \mid (v, v') \in E\}$. Let $N = \{1, \dots, n\}$ be the set of players. A game arena is a finite graph $\mathcal{A} = (V, E)$ where V is the set of game positions and $E \subseteq V \times V$ is the move relation. V is partitioned into sets $V = V_1 \cup \dots \cup V_n$ where for all $i \in N$, V_i is the set of game positions of player i. For simplicity we assume that for all $v \in V$, the set vE is nonempty. An initialised game is a game G along with a starting vertex $v_0 \in V$. Henceforth when we use the notation (G, v_0), we will generally mean an initialised game with initial vertex v_0.

Given a game (G, v_0), a play in (G, v_0) can be viewed as follows: initially a token is placed at vertex v_0. At any point, if the token is at a vertex $v \in V_i$ (i.e. a player i vertex) then she moves the token to some $v' \in vE$. In this way an infinite path, $\pi = v_0 v_1 \dots$ where for all $j > 0$ we have $(v_{j-1}, v_j) \in E$, called a play is constructed in the arena.

For a finite sequence $\rho = v_0 v_1 \dots v_k$ let $first(\rho) = v_0$, $last(\rho) = v_k$ and for an infinite sequence $\pi = v_0 v_1 \dots$ let $\inf(\pi)$ denote the set of nodes that appear infinitely often in π. For any sequence $\pi = v_0 v_1 \dots$, let $\pi(i)$ denote the ith element of π, π_i denote the length i prefix of π, $|\pi|_v$ denote the number of v's occuring in π and $Occ(\pi) = \{v \mid |\pi|_v > 0\}$.

2.2 Strategies

A strategy for player i specifies at each game position of i which move to choose. It is a function $\sigma_i : V^* V_i \to V$ from the set of all finite plays (histories) ending in a player i node to the set of game positions which satisfies the condition:

- for all $\pi = v_0 \dots v_k$, such that $v_k \in V_i$, $\sigma_i(\pi) \in v_k E$.

Let $T_{\mathcal{A}}$ denote the tree unfolding of $\mathcal{A}$. A strategy σ_i can also be thought of as a subtree T_{σ_i} of $T_{\mathcal{A}}$ (called the strategy tree) where for each player i node there is a unique outgoing edge and for any other player node, we include all the edges.

A strategy σ is said to be bounded memory if there exists a finite state machine $\mathcal{M} = (M, g, h, m^I)$ where M is the memory of the strategy, $m^I \in M$ is the initial memory, $g : V \times M \to M$ the memory update function, and $h : V \times M \to V$ is the output function which specifies the choice of the player such that if $v_0 \dots v_k$ is a play and $m_0 \dots m_{k+1}$ is a sequence determined by $m_0 = m^I$ and $m_{i+1} = g(v_i, m_i)$ then $\sigma(v_0 \dots v_k) = h(v_k, m_{k+1})$. The strategy σ is said to be memoryless if M is a singleton.

Let Ω_i denote the set of all strategies for player i. A strategy profile $\bar{\sigma} = (\sigma_1, \dots, \sigma_n)$ defines a unique play in the game, we use $\pi_{\bar{\sigma}}$ to denote this play. We often use the notation $\bar{\imath}$ to denote the set $N \setminus \{i\}$ and $\bar{\sigma}_{-i}$ to denote the tuple $(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n)$.

2.3 Objectives

The arena specifies the rules of the game and the moves of the players. To describe a game fully, the objectives of the players have to be specified. The players play in a way that they can 'achieve/avoid' these objectives. The objective of a player is usually a subset ot the set of plays. However, for algorithmic analysis, the objectives need to be finitely presentable. The most widely studied of these presentations are ω-regular objectives, mean-payoff objectives and so on. These naturally arise in the specifications encountered in the verification and synthesis of reactive systems. In this paper we concentrate on a specific type of ω-regular objective, the Muller objective.

Binary Objectives: In this case the objective of each player i is an omega regular subset *Win* of plays. The game is not antagonistic since objectives of players are allowed to overlap. For instance, for Muller objectives, each player i has a collection $\mathcal{F}_i$ of Muller sets. She wins the game if and only if the game eventually settles down to some subset F of the set of vertices V such that $F \in \mathcal{F}_i$; otherwise she loses. We often call these objectives 'win-lose objectives'.

Generalised Objectives: In this paper we are concerned with games where players have preference orderings on the various Muller sets. Formally, player i has a total order $\sqsubseteq_i$ on the Muller sets. Such an ordering can also be viewed as an utility function $u_i : 2^V \to \mathbb{N}$. Since a strategy profile $\bar{\sigma} \in \Pi_{i=1}^n \Omega_i$ defines a unique Muller set $F = \inf(\pi_{\bar{\sigma}})$, we may also think of u_i to be a function from $u_i : \Pi_{i=1}^n \Omega_i \to \mathbb{N}$. We call such games generalised Muller games.

2.4 Best Response and Equilibrium

The notion of best response and equilibrium is defined as follows:

- A strategy σ_i of player i is said to be a best response for $\bar{\sigma}_{-i}$ if for all $\sigma_i' \in \Omega_i$, $u_i(\pi_{(\bar{\sigma}_{-i},\sigma_i')}) \leq u_i(\pi_{(\bar{\sigma}_{-i},\sigma_i)})$.
- A strategy tuple $\bar{\sigma} = (\sigma_1, \ldots, \sigma_n)$ is a Nash equilibrium if for all $i \in N$, σ_i is the best response for $\bar{\sigma}_{-i}$.
- A subgame perfect equilibrium (SPE) [8] can be defined in our setting as follows. Let ρ be a (finite) path in the arena $\mathcal{A}$. Given a strategy σ_i for player i, the strategy $\sigma_i(\rho)$ is defined to be a function: $\sigma_i(\rho) : \rho V^* V_i \to V$ such that $\sigma_i(\rho)(\rho') = \sigma_i(\rho\rho')$. Let $\bar{\sigma}(\rho)$ denote the tuple $(\sigma_1(\rho), \ldots, \sigma_n(\rho))$. A strategy tuple $\bar{\sigma}$ in the initialised game (G, v_0) is said to be an SPE if for every vertex v in $\mathcal{A}$ and for every path ρ from v_0 to v in $\mathcal{A}$, $\bar{\sigma}(\rho)$ is a Nash equilibrium for the initialised game (G, v).

2.5 Computing Nash Equilibrium

In [3], the authors show that n-player games with win-lose Borel objectives always have a Nash equilibrium. An equilibrium profile is where the players play 'threat' strategies in that, if any player i unilaterally deviates from her prescribed behaviour, all the other players punish her by playing a profile where she can never gain anything more than what she would have had she stuck to her prescribed strategy. The procedure can be appropriately modified to show that Nash equilibrium always exists in a generalised Muller game.

Threat strategies are naturally defined in the case of games with win-lose objectives. However, with general non-zero sum games, it is not clear whether threat strategies constitute 'efficient' solutions profiles. For finite games backward inductive solution profiles are known to preserve nice properties like Pareto efficiency [4]. Here we show that the standard backward induction procedure can be effectively utilised for computing Nash equilibria in generalised Muller games.

3 Solving Generalised Games

In this section, we develop our procedure for solving generalised Muller games and prove its correctness. The central idea of the procedure is to perform a finite unfolding of the game arena, making use of the 'latest appearance record' (LAR) data structure [6] and apply a backward induction on this unfolding.

3.1 The LAR Tree

Let $\mathcal{A} = (V, E)$ be a finite graph and $\sharp \notin V$. Let $\prec$ be a total order on the nodes of V. We Let

$$L_{\mathcal{A}} = \{l \in (V \cup \{\sharp\})^{|V|+1} \mid |l|_\sharp = 1 \wedge \forall v \in V(|l|_v = 1)\}$$

The set $L_{\mathcal{A}}$ is called the LAR memory. Henceforth we shall refer to elements from $L_{\mathcal{A}}$ as $x\sharp y$ where $x, y \in V^*$. We define a function *next* : $L_{\mathcal{A}} \times V \to L_{\mathcal{A}}$ as

$$next(x\sharp y, v) = \begin{cases} x'\sharp x'' yv & \text{iff } x\sharp y = x'vx''\sharp y \\ xy'\sharp y''v & \text{iff } x\sharp y = x\sharp y'vy'' \\ x\sharp y & \text{iff } x\sharp y = x\sharp y'v \end{cases}$$

For a finite play $\rho = v_0 v_1 \dots v_k$ in the arena we define LAR(ρ) inductively as:

- LAR(v_0)=$x\sharp v_0$ where x is ordered according to the total order $\prec$.
- LAR($v_0 \dots v_i$)=*next*(LAR($v_0 \dots v_{i-1}$),v_i), $i \geq 1$.

Given an arena $\mathcal{A} = (V, E)$ and an element $x\sharp y \in L_{\mathcal{A}}$ the (finite) LAR tree $T_{fin}(\mathcal{A}, x\sharp y)$ corresponding to $\mathcal{A}$ and $x\sharp y$, or just $T_{fin}(x\sharp y)$ when the arena $\mathcal{A}$ is fixed, is constructed as follows:

- $x\sharp y$ is the root of $T_{fin}(x\sharp y)$.
- For any node $x'\sharp y'v$ of $T_{fin}(x\sharp y)$, and for all $u \in vE$, $x''\sharp y'' = next(x'\sharp y'v, u)$ is a child of $x'\sharp y'v$ iff there is no node $x''\sharp y''$ in the unique path from the root to $x'\sharp y'v$, or $x''\sharp y''$ is the first node to repeat in the path.

That $T_{fin}(x\sharp y)$ is well defined follows from the fact that the function *next* is well defined. And the fact that $T_{fin}(x\sharp y)$ is finite can be ascertained by noting that along any sequence of the elements of $L_{\mathcal{A}}$ of length $(|V| + 1)! + 1$, at least one element is bound to repeat, by the pigeonhole principle.

3.2 Ensuring a Muller Set

Let $\mathcal{A} = (V, E)$ be an arena, v_0 be an initial vertex and $T_{fin}(x\sharp v_0)$ be the LAR tree of $\mathcal{A}$ corresponding to the LAR $x\sharp v_0$ where x is ordered according to the total order $\prec$. Let $\mathcal{F} \subseteq 2^V$ be a collection Muller sets and $M \subsetneq N, M \neq \emptyset$ be a subset of players. We label the leaf nodes of $T_{fin}(x\sharp v_0)$ with $\mathcal{F}$ or $\bar{\mathcal{F}}$ as follows. For a leaf node $x\sharp y$ of $T_{fin}(x\sharp v_0)$, let ρ be the unique path in $T_{fin}(x\sharp v_0)$ from the root to $x\sharp y$. Let ρ' be the least suffix of ρ such that *first*(ρ') = *last*$(\rho') = x\sharp y$ (note that such a suffix always exists by construction of $T_{fin}(x\sharp v_0)$).

Let Let $l_{max} = max\{|y| \mid x\sharp y \in Occ(\rho')\}$ and let $L_\rho = \{x\sharp y \mid |y| = l_{max}\}$. If there exists $x'\sharp y' \in L_\rho$ such that $\{y'\} \in \mathcal{F}$ then we label the leaf $x\sharp y$ with $\mathcal{F}$. Otherwise we label it with $\bar{\mathcal{F}}$.

We now label the entire $T_{fin}(x\sharp v_0)$ with $\mathcal{F}$ or $\bar{\mathcal{F}}$ and construct (memoryless) strategies $\mu_i : L_\mathcal{A} \to L_\mathcal{A}$, $i \in M$ using the following backward induction procedure.

Procedure 1

Suppose all children of node $x\sharp yv$ of $T_{fin}(x\sharp v_0)$ have been labelled. Let $T_{x\sharp yv}$ be the set of children of $x\sharp yv$ and let $T^{\mathcal{F}}_{x\sharp yv} \subseteq T_{x\sharp yv}$ be the nodes among these children that have been labelled with $\mathcal{F}$. Then

- $v \in V_i$ such that $i \in M$:
  - If $T^{\mathcal{F}}_{x\sharp yv} \neq \emptyset$ then let $x'\sharp y'v'$ be such that $v' \prec v''$ for all $x''\sharp y''v'' \in T^{\mathcal{F}}_{x\sharp yv}$. Label $x\sharp yv$ with $\mathcal{F}$ and put $\mu_i(x\sharp yv) = x'\sharp y'v'$.
  - If $T^{\mathcal{F}}_{x\sharp yv} = \emptyset$ then let $x'\sharp y'v' \in T_{x\sharp yv}$ be such that $v' \prec v''$ for all $x''\sharp y''v'' \in T_{x\sharp yv}$. Label $x\sharp yv$ with $\bar{\mathcal{F}}$ and put $\mu_i(x\sharp yv) = x'\sharp y'v'$.
- $v \in V_i$ such that $i \notin M$:
  - If $T^{\mathcal{F}}_{x\sharp yv} = T_{x\sharp yv}$, which means that every child of $x\sharp yv$ is labelled $\mathcal{F}$, then label $x\sharp yv$ with $\mathcal{F}$.
  - If $T^{\mathcal{F}}_{x\sharp yv} \subsetneq T_{x\sharp yv}$ then there exists a child $x'\sharp y'v'$ of $x\sharp yv$ such that $x'\sharp y'v'$ has label $\bar{\mathcal{F}}$. Label $x\sharp yv$ with $\bar{\mathcal{F}}$.

Note that, choosing the least v in the order $\prec$ in the above procedure ensures that the μ_i's constructed are well defined.

Players M are said to be able to ensure the Muller sets $\mathcal{F}$ by strategy μ_i, $i \in M$ on $T_{fin}(x\sharp v_0)$ if the root of $T_{fin}(x\sharp v_0)$ is labelled $\mathcal{F}$ by the above procedure and μ_i are the strategies constructed.

Given a memoryless strategy μ_i for player i on an LAR tree $T_{fin}(x\sharp y)$ we can construct the corresponding bounded memory strategy σ_i for player i on the arena $\mathcal{A}$ as follows:

- The memory M of σ_i is the set $L_\mathcal{A}$ and the initial memory m_I is $x\sharp y$.
- The memory update function $g_i : V \times M \to M$ is defined as $g_i(v, x'\sharp y') = next(x'\sharp y', v)$.
- The output function $h_i : V_i \times M \to V$ is defined as $h(v, x'\sharp y) = \mu_i(x'\sharp y)$.

For a word on notation, we denote memoryless strategies on $T_{fin}(\cdot)$ by μ and we denote the bounded memory strategies on the arena $\mathcal{A}$ by σ.

A strategy σ on the arena $\mathcal{A}$ is said to exist in $T_{fin}(x\sharp y)$ if it corresponds to some strategy μ on $T_{fin}(x\sharp y)$. A strategy μ is said to exist in $T_{fin}(x\sharp y)$ if it is some subtree of $T_{fin}(x\sharp y)$.

LEMMA 1. *If players $M \subsetneq N$, $M \neq \emptyset$ can ensure Muller sets $\mathcal{F}$ in $T_{fin}(x\sharp v_0)$ by strategies μ_i, $i \in M$, then they can ensure $\mathcal{F}$ in (G, v_0) by the bounded memory strategies σ_i corresponding to μ_i.*

PROOF. Suppose not and suppose that players M can ensure $\mathcal{F}$ in $T_{fin}(x\sharp v_0)$ by $\mu_i, i \in M$ but they cannot ensure $\mathcal{F}$ in (G, v_0) by the corresponding strategies σ_i. Then there exists a play π in (G, v_0) conforming to σ_i, $i \in M$ such that it settles down to a Muller set $F' \notin \mathcal{F}$. There are two cases to consider.

The first case is when there exists $v \in F'$ such that $v \notin F$ for any $F \in \mathcal{F}$. Let j be the first index such that $\pi(j) = v$ and $\pi(j-1) \in V_k$, $k \notin M$. Let ρ be the (finite) path in $T_{fin}(x\sharp v_0)$ corresponding to π. j must be greater than $|\rho|$; otherwise ρ couldn't have been labelled $\mathcal{F}$ and hence μ_i's couldn't have ensured $\mathcal{F}$. Let $x'\sharp y'$=LAR(π_{j-1}). By the construction of the LAR tree $T_{fin}(x\sharp v_0)$ there exists a node $x'\sharp y' \in \rho$ itself. But this means that player k had the option of playing v at the node $x'\sharp y'$ forcing $x'\sharp y'$ and hence the root to be labelled $\bar{\mathcal{F}}$. But this would contradict the fact that μ_i's ensure $\mathcal{F}$ in $T_{fin}(x\sharp v_0)$.

The other case is when there exists $v \in F \in \mathcal{F}$ such that $v \notin F'$. Let ρ be the (finite) path in $T_{fin}(x\sharp v_0)$ corresponding to π. Let l be the biggest index l such that $\pi(l) = v$ but $l < |\rho|$. Suppose $\pi(l-1) \in V_i$, $i \in M$. Then for all indices $l_1, l_2, \ldots$ such that $l < l_1 < l_2 < \ldots$ and LAR(π_{l_1})= LAR(π_{l_2})=… = LAR(π_{l-1}), player i has to play v as it is prescribed by the memoryless strategy μ_i, and hence in turn by the corresponding bounded memory strategy σ_i. But this contradicts the fact that the π settles down to F'.

Finally, suppose $\pi(l-1) \in V_k$, $k \notin M$. Then player k has the option of playing v at $\pi(l-1)$ and al all indices $l_1, l_2, \ldots$ such that $l < l_1 < l_2 < \ldots$ and LAR(π_{l_1})= LAR(π_{l_2})=… = LAR(π_{l-1}). Hence μ_i's could not have ensured $\mathcal{F}$ in $T_{fin}(x\sharp v_0)$ as the leaf node of ρ wouldn't have been labelled $\mathcal{F}$ and hence neither the root. ∎

LEMMA 2. *Let $\mathcal{F}$ be a collection of Muller sets. If players $M \subsetneq N$, $M \neq \emptyset$ have strategies σ_i, $i \in M$ to ensure $\mathcal{F}$ in the game (G, v_0), then they have strategies $\mu_i, i \in M$ to ensure $\mathcal{F}$ in $T_{fin}(x\sharp v_0)$.*

PROOF. Suppose players M do not have strategies $\mu_i, i \in M$ to ensure $\mathcal{F}$ in $T_{fin}(x\sharp v_0)$ then $T_{fin}(x\sharp v_0)$ being a finite tree (and hence a finite extensive form game) it follows that players $N \setminus M$ have strategies $\mu_i, i \in N \setminus M$ to ensure $2^V \setminus \mathcal{F}$ in $T_{fin}(x\sharp v_0)$, since finite games are determined. Then by Lemma 1, players $N \setminus M$ have bounded memory strategies $\sigma_i, i \in N \setminus M$ corresponding to the μ_i's to ensure $2^V \setminus \mathcal{F}$ in (G, v_0) as well. But this contradicts the assumption that players M have strategies to ensure $\mathcal{F}$ in (G, v_0). ∎

Combining the above two lemmata, we have the following theorem.

THEOREM 3. *Let (G, v_0) be an n-player game, N being the set of players. Let $M \subsetneq N$, $M \neq \emptyset$ be a subset of players and $\mathcal{F}$ be a collection of a Muller sets consisting of the nodes of the arena of the game. Then players M can ensure $\mathcal{F}$ in (G, v_0) if and only if they can ensure $\mathcal{F}$ in $T_{fin}(x\sharp v_0)$. Also, if players M can ensure $\mathcal{F}$ in $T_{fin}(x\sharp v_0)$ then the bounded memory strategies σ_i, $i \in M$ corresponding to the memoryless strategies μ_i, $i \in M$ computed by Procedure 1, ensures $\mathcal{F}$ in (G, v_0).*

3.3 Equilibrium Computation

Let (G, v_0) be a generalised Muller game with the set of players N and let u_i be the utility function of player i over the Muller sets. We label the leaf nodes of the LAR tree $T_{fin}(x\sharp v_0)$ consistently with tuples $(x_1, \ldots, x_n) \in \mathbb{N}^n$ as follows.

For a leaf node $x\sharp y$ of $T_{fin}(x\sharp v_0)$, let ρ be the unique path in $T_{fin}(x\sharp v_0)$ from the root to $x\sharp y$. Let ρ' be the least suffix of ρ such that $first(\rho') = last(\rho') = x\sharp y$. Let $l_{max} = max\{|y| \mid x\sharp y \in Occ(\rho')\}$ and let $L_\rho = \{x\sharp y \mid |y| = l_{max}\}$. Observe that, by the property of the LAR construction $y = y'$ for all $x\sharp y,\ x'\sharp y' \in L_\rho$. Let $Y = y'$ such that $x'\sharp y' \in L_\rho$. Label the leaf $x\sharp y$ with $(u_1(Y), \ldots, u_n(Y))$.

We now label the entire tree T_{fin} consistently, with tuples $(x_1, \ldots, x_n) \in \mathbb{N}^n$, and compute a strategy tuple $\bar{\mu} = (\mu_1, \ldots, \mu_n)$ as follows:

Procedure 2

Suppose all children of node $x\sharp yv$ have been labelled and $v \in V_i$. Let
$u_{x\sharp yv} = max\{u_i(x'\sharp y') \mid x'\sharp y' \text{ is a child of } x\sharp yv\}$,
$T_{x\sharp yv} = \{x'\sharp y' \mid x'\sharp y' \text{ is a child of } x\sharp yv \text{ and } u_i(x'\sharp y') = u_{x\sharp yv}\}$.
Put $\mu_i(x\sharp yv) = x'\sharp y'v' \in T_{x\sharp yv}$ such that $v' \prec v''$ for all $x''\sharp y''v'' \in T_{x\sharp yv}$. Label $x\sharp yv$ with $(u_1(x'\sharp y'v'), \ldots, u_n(x'\sharp y'v'))$.

THEOREM 4. *Every generalised Muller game (G, v_0) has a Nash equilibrium.*

The proof shows that the bounded memory strategy tuple $(\sigma_1, \ldots, \sigma_n)$ corresponding to the tuple $(\mu_1, \ldots, \mu_n)$ constructed by Procedure 2, is an equilibrium tuple in the game (G, v_0). The essence of the proof is the same as the one for Theorem 3: player i has an incentive to deviate from σ_i in (G, v_0) if and only if she has an incentive to deviate from μ_i in $T_{fin}(x\sharp yv_0)$. We omit the full proof due to space limitations.

Complexity. Let the number of vertices in the arena $\mathcal{A}$ be m. In Procedure 2, the number of permutations of the m vertices of the arena is equal to m. Thus the size of the LAR memory $L_{\mathcal{A}}$ may be as big as $(m+1)!$. This means that each path of the LAR tree might be $(m+1)!$ nodes long. As there are $O(m^{m!})$ such paths and the backward induction procedure runs in time linear in the size of the LAR tree, the running time of Procedure 2 is $O(m^{m!})$.

4 Subgame Perfection

Nash equilibrium, as a solution concept, has its limitations. One such limitation is that it does not take into account the sequential nature of the game. In an extensive form game, if a player deviates from equilibrium behaviour even for just one move, Nash equilibrium says nothing about the outcome of the game. One possible refinement to Nash equilibrium is to insist that strategies are optimal after every prefix. This is achieved by subgame perfect equilibrium [8]. Ummels [9] has shown that subgame perfect equilibria always exist for n-player infinite games on graphs for ω-regular win-lose objectives. The question therefore arises whether subgame perfect equilibria exist for n-player infinite games where the objectives are not win-lose but generalised.

For finite extensive form games, the backward induction procedure does indeed yield a subgame perfect equilibrium profile. Since our construction of the equilibrium profile for generalised Muller games employs a backward induction procedure (Procedure 2), it is natural to ask if the profile constructed is subgame perfect. The answer is affirmative for win-lose objectives as we show in the following proposition.

PROPOSITION 5. *Every generalised Muller game with binary objectives has a subgame perfect equilibrium.*

PROOF. We show that the strategy tuple $\bar{\sigma}$ corresponding to the tuple $\bar{\mu}$ constructed by Procedure 2 is a subgame perfect equilibrium of (G, v_0) when the objectives of the players are binary.

Suppose $\bar{\sigma}$ is not an SPE. Then there exists a vertex $v \in V, v \in V_i$ say, and a path ρ from v_0 to v such that $\bar{\mu}(\rho)$ is not an equilibrium tuple. Let ρ' be the (finite) path in $T_{fin}(x\sharp v_0)$ corresponding to ρ.

Suppose player j has an incentive to deviate from $\sigma_j(\rho)$. If $|\rho| < |\rho'|$, then player j has an incentive to deviate from μ_j as well. But this contradicts the fact that $\bar{\mu}$ is an equilibrium tuple (Theorem 3).

So assume that $|\rho| \geq |\rho'|$. Then by the property of the LAR tree $T_{fin}(x\sharp v_0)$, there exists ρ'' such that $|\rho''| < |\rho'|$ and LAR(ρ'') = LAR(ρ). Now, since $\bar{\sigma}_j$ corresponds to $\bar{\mu}_j$ which is a memoryless strategy constructed from $T_{fin}(x\sharp v_0)$, it prescribes the same action at ρ and ρ'' (since LAR(ρ'') = LAR(ρ)). Thus if player j has an incentive to deviate from $\sigma_j(\rho)$, she has an incentive to deviate from $\sigma_j(\rho'')$ as well which in turn means she has an incentive to deviate from σ_j in the first place. But this again contradicts fact that $\bar{\sigma}$ is an equilibrium tuple (Theorem 3). ∎

The argument for the above proof breaks down when the objectives of the players are not binary but generalised.

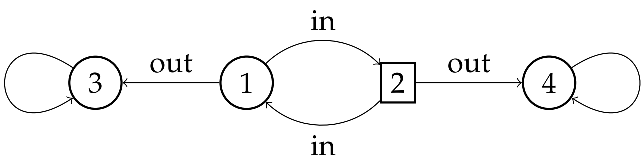

Figure 1: Non existence of subgame perfect equilibrium

Example 1 Consider the game arena shown in Figure 1. Player 1 nodes are denoted by ◯ and player 2 nodes are denoted by □. The game starts at node 1. The utilities of the players for the relevant Muller sets are as follows: $u_1(\{3\}) = 1, u_1(\{1,2\}) = 0, u_1(\{4\}) = 2$ and $u_2(\{3\}) = 0, u_2(\{1,2\}) = 2, u_2(\{4\}) = 1$. Procedure 2 gives the following strategies μ_1 and μ_2 for players 1 and 2 respectively. μ_1 prescribes that player 1 stays 'in' in her first move expecting player 2 to go 'out' and hence give 1 a better payoff. But if she plays 'in' then player 2 stays 'in' as prescribed by μ_2 because that gives her a better payoff. To this 1 assumes that player 2 will stay in forever and hence plays 'out' in her next move as prescribed by μ_1. The profile (μ_1, μ_2) is thus not subgame perfect. One can verify that the above game does not have a subgame perfect equilibrium.

The above example shows that in general subgame perfect equilibria need not exist for generalised Muller games. Thus an obvious question to ask would be: Is it decidable to check whether subgame perfect equilibrium exists in a given generalised Muller game? In this section, we develop a procedure to decide the existence of sub-game perfect equilibrium and to compute the equilibrium profile when it exists.

First, it is important to note that given an arena $\mathcal{A}$ and an initial vertex v, for any bounded memory strategy on $\mathcal{A}$ that uses memory $L_{\mathcal{A}}$ the initial element of $L_{\mathcal{A}}$ does not matter. In other words, no matter what element $x\sharp y$ such that $last(x\sharp y) = v$ of $L_{\mathcal{A}}$ we take as the root of the LAR tree, the backward induction procedure (Procedure 2) gives all the bounded memory strategies that are possible by using memory $L_{\mathcal{A}}$ and updating it as described in Section 3. This is because starting from any vertex v of $\mathcal{A}$, the tree explores all possible cycles reachable from v and a path of the LAR tree is terminated if and only if a cycle is completed.

Define the following property for a strategy tuple $\bar{\mu}$ on $T_{fin}(x\sharp v_0)$

Property 1 *For every $x\sharp y \in \mathcal{T}_{fin}(x\sharp v_0)$, there exists a strategy tuple $\bar{\mu}'$ on $\mathcal{T}_{fin}(x\sharp y)$ such that $\bar{\mu}'$ is derived by backward induction on $\mathcal{T}_{fin}(x\sharp y)$ and $\bar{\mu}'(x\sharp y) = \bar{\mu}(x\sharp y)$.*

Given an game (G, v_0), let $Path(G, v_0)$ be the set of all finite paths starting at v_0 in G. Define $P : L \to 2^{Path(G,v_0)}$ such that $P(x\sharp y) = \{\rho \in Path(G, v_0) \mid \text{LAR}(\rho) = x\sharp y\}$.

Given a strategy tuple $\bar{\sigma}$ on (G, v_0) define $Q_{\bar{\sigma}} : L \to 2^{2^V}$ as $Q_{\bar{\sigma}}(x\sharp y) = \{\inf(\pi_{\bar{\sigma}(\rho)}) \mid \rho \in P(x\sharp y)\}$ where $\pi_{\bar{\sigma}(\rho)}$ is the play conforming to $\bar{\sigma}(\rho)$. Let $C_{\bar{\sigma}}$ be a choice function $C_{\bar{\sigma}} : L \to 2^V$ such that

Property 2 *$x'\sharp y'$ is a child of $x\sharp y$ in $T_{fin}(x\sharp v_0)$ and $C_{\bar{\sigma}}(x\sharp y) \in Q_{\bar{\sigma}}(x'\sharp y')$ implies $C_{\bar{\sigma}}(x'\sharp y') = C_{\bar{\sigma}}(x\sharp y)$.*

It follows that

Property 3 *If $C_{\bar{\sigma}}(x\sharp yv) = F$, $v \in V_i$ then there actually exists a $\rho \in Path(G, v_0)$ such that LAR(ρ)= $x\sharp yv$, $\inf(\pi_{\bar{\sigma}(\rho)}) = F$, $\sigma_i(\rho) = v$ and $\inf(\pi_{\bar{\sigma}(\rho v)}) = F$.*

Assume for the moment that given any strategy tuple $\bar{\sigma}$, we have such a function $C_{\bar{\sigma}}$ satisfying Property 2. Now let $\bar{\sigma}$ be an SPE on (G, v_0). For every $i \in N$, construct σ_i' as follows: $\sigma_i' : V^*V_i \to V$ such that $\sigma_i'(uv) = uvv'$ iff $C_{\bar{\sigma}}(\text{LAR}(uv)) = C_{\bar{\sigma}}(\text{LAR}(uvv')), (v, v') \in E$

LEMMA 6. *$\bar{\sigma'}$ is an SPE on (G, v_0)*

PROOF. Suppose not. Then there exists $\rho \in Path(G, v_0)$ such that $\bar{\sigma'}(\rho)$ is not an NE. So suppose player i has an incentive to deviate from $\sigma_i'(\rho)$. Now by property 3 there exists a history $\rho' \in Path(G, v_0)$ such that $\sigma_i'(\rho) = \sigma_i(\rho')$. Then player i must have an incentive to deviate from $\sigma_i(\rho')$ itself. But this contradicts the subgame perfection of $\bar{\sigma}$. ∎

Now $\bar{\sigma'}$ exists on $T_{fin}(x\sharp v_0)$. Indeed, it is the strategy where $\sigma_i'(x\sharp y) = x'\sharp y'$ such that $x\sharp y$ is a parent of $x'\sharp y'$ in $T_{fin}(x\sharp v_0)$ and $C_{\bar{\sigma}}(x\sharp y) = C_{\bar{\sigma}}(x'\sharp y')$. Let $\bar{\mu'}$ denote this memoryless strategy tuple on $T_{fin}(x\sharp v_0)$ corresponding to $\bar{\sigma'}$.

LEMMA 7. *$\bar{\mu'}$ has Property 1*

PROOF. Suppose not. Then there exists a node $x\sharp y \in T_{fin}(x\sharp v_0)$ such that for any backward induction strategy profile $\bar{\mu}^+$ on $T_{fin}(x\sharp y)$, $\bar{\mu'}(x\sharp y) \neq \bar{\mu}^+(x\sharp y)$. Now we have that $\bar{\sigma'}$ is subgame perfect on (G, v_0) and bounded memory, the memory being $L_{\mathcal{A}}$. So $\bar{\mu'}(x\sharp y)$ must correspond to some equilibrium tuple $\bar{\sigma'}$ on $(G, last(x\sharp y))$ which exists in $T_{fin}(x\sharp y)$, as backward induction on $T_{fin}(x\sharp y)$ gives all the bounded memory equilibria starting at node $last(x\sharp y)$ with memory $L_{\mathcal{A}}$. But then the above cannot happen. ∎

LEMMA 8. *If a strategy tuple $\bar{\mu}$ on $\mathcal{T}_{fin}(x\sharp v_0)$ satisfies Property 1 then $\bar{\sigma}$ on (G, v_0) corresponding to $\bar{\mu}$ is an SPE.*

PROOF. Suppose not. Then there exists $\rho \in Path(G, v_0)$ such that $\bar{\sigma}(\rho)$ is not an equilibrium on $(G, last(\rho))$. Let LAR(ρ)=$x\sharp y$. Now $x\sharp y \in T_{fin}(x\sharp v_0)$ and $\bar{\sigma}(\rho)$ is bounded memory, the memory being $L_{\mathcal{A}}$. Thus $\bar{\mu}(\rho)$ corresponding to $\bar{\sigma}(\rho)$ cannot be a backward induction profile on $T_{fin}(x\sharp y)$ as backward induction on $T_{fin}(x\sharp y)$ gives all the bounded memory equilibria starting at node $last(x\sharp y)$ with memory $L_{\mathcal{A}}$. So $\bar{\mu}$ cannot satisfy Property 1. ∎

From the above set of lemmata we have the following theorem.

THEOREM 9. *A generalised Muller game (G, v_0) has a subgame perfect equilibrium if and only if there exists a strategy profile $\bar{\mu}$ on $T_{fin}(x\sharp v_0)$ that satisfies Property 1.*

PROOF. It only remains to construct the choice function $C_{\bar{\sigma}}$ satisfying Property 2 given a strategy profile $\bar{\sigma}$. We do that as follows: let $\lessdot$ be a breadth-first ordering on $T_{fin}(x\sharp v_0)$ and let $H = \emptyset$.
Till $H \neq T_{fin}(x\sharp v_0)$ do

- Let $x\sharp y$ be the minimum in the ordering $(T_{fin}(x\sharp v_0) \setminus H) \restriction \lessdot$.
- Let ρ be the path from the root to $x\sharp y$.
- Let $C_{\bar{\sigma}}(x\sharp y) = \inf(\pi_{\bar{\sigma}(\rho)}) = F$.
- There exists a path ρ' from $x\sharp y$ to a leaf node of $T_{fin}(x\sharp v_0)$ such that for all $x'\sharp y' \in \rho'$, $F \in Q_{\bar{\sigma}}(\text{LAR}(x'\sharp y'))$. Put $C_{\bar{\sigma}}(x'\sharp y') = F$ for all such $x'\sharp y' \in \rho'$. Let $H = H \cup \{x'\sharp y' \mid x'\sharp y' \in \rho'\}$.
- For all $x'\sharp y' \in T_{fin}(x\sharp v_0)$ such that $x'\sharp y' \notin \rho'$ and such that LAR($x'\sharp y'$) = LAR($x''\sharp y''$) for some $x''\sharp y'' \in \rho'$, put $C_{\bar{\sigma}}(x'\sharp y') = C_{\bar{\sigma}}(x''\sharp y'')$. Let $H = H \cup \{x'\sharp y'\}$.

It is immediate that the $C_{\bar{\sigma}}$ constructed above meets Property 2. ∎

The above theorem immediately gives us the following procedure to test if a generalised Muller game has a subgame perfect equilibrium.

Procedure:
For every backward induction strategy profile $\bar{\mu}$ on $T_{fin}(x\sharp v_0)$
    For all $x\sharp y \in T_{fin}(x\sharp v_0)$ such that $x\sharp y \neq x\sharp v_0$
        If $\bar{\mu}(x\sharp y) \neq \bar{\mu}^+(x\sharp y)$ for some backward induction strategy profile $\bar{\mu}^+$
        on $T_{fin}(x\sharp y)$, then return TRUE and exit
Return FALSE

Complexity: Let $|V| = m$. There are at most $1 + m + m^2 + \ldots + m^{m!} = (m^{m!+1} - 1)/(m-1)$ nodes in an LAR tree. There are atmost $m \cdot m^2 \cdot \ldots m^{m!} = m^{m!(m!+1)/2}$ strategy tuples in an LAR tree. Hence the complexity of the above procedure is $\mathcal{O}((m^{m!+1} - 1)/(m-1) \cdot m^{m!(m!+1)/2} \cdot (m^{m!+1} - 1)/(m-1)) = \mathcal{O}(m^{2m!} \cdot m^{(m!)^2})$.

5 Discussion

Nash equilibrium and subgame perfect equilibrium are well studied in finite games. In the setting of finite games, subgame perfection is justified under the trembling hand assumption and a subgame perfect profile is considered more robust than general Nash equilibrium profiles. When we move to nonzero sum games of infinite duration even coming up with an appropriate notion of rationality which justifies the trembling hand assumption is a challenging task. However, the equilibrium notions are mathematically well defined and deserves attention in their own right. In this paper rather than delve into issues concerning rationality, we have attempted to investigate equilibrium notions in the context of infinite games. We have shown that the standard technique of backward induction can be appropriately modified to compute equilibrium profile in generalised Muller games. Though the running time complexity of the procedures is not very encouraging, we would like to view this as a generic technique for solving games.

6 Acknowledgements

We are indebted to R. Ramanujam for suggesting this problem and for his constant support and encouragement. We are grateful to Dietmar Berwanger for the various discussions and for sharing his insights into rationality concerns associated with subgame perfect equilibria in infinite games. We thank the anonymous reviewers for their valuable comments and in particular for pointing out Example 1.

References

[1] J. R. Büchi and L. H. Landweber. Solving sequential conditions by finite-state strategies. *Transactions of the American Mathematical Society*, 138:295–311, 1969.

[2] K. Chatterjee. *Stochastic ω-Regular Games*. PhD thesis, University of California at Berkeley, 2007.

[3] K. Chatterjee, M. Jurdzinski, and R. Majumdar. On Nash equilibria in stochastic games. In *Proceedings of the 13th Annual Conference of the European Association for Computer Science Logic*, volume 3210 of *LNCS*, pages 26–40. Springer-Verlag, 2004.

[4] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 2005.

[5] E. Grädel and M. Ummels. Solution concepts and algorithms for infinite multiplayer games. In *New Perspectives on Games and Interaction*, volume 4 of *Texts in Logic and Games*, pages 151–178. Amsterdam University Press, 2008.

[6] Y. Gurevich and L. Harrington. Trees, automata and games. In *Proceedings of the 14th Annual Symposium on Theory of Computing*, pages 60–65. ACM Press, 1982.

[7] D. A. Martin. Borel determinacy. *Annals of Mathematics*, 102:363–371, 1975.

[8] R. Selten. Spieltheoretische behandlung eines oligopolmodells mit nachfrageträgheit. *Zeischrift Für Die Gesamte Staatswissenschaft*, 121:301–324, 1965.

[9] M. Ummels. Rational behaviour and strategy construction in infinite multiplayer games. In *Proceedings of the Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 4337 of *LNCS*, pages 212–223. Springer, 2006.

[10] E. Zermelo. Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels,. In *Proceedings of the Fifth Congress Mathematicians*, pages 501–504. Cambridge University Press, 1913.