Equilibrium concepts in graph games

B Tech project report

by

Prathmesh Prabhu
Roll: 06005002

under the guidance of

Prof. Krishna S

Department of Computer Science and Engineering
Indian Institute of Technology, Bombay
November 2009

Contents

Acknowledgement

I would like to express my deep sense of gratitude to **Prof. Krishna S** for her invaluable help and guidance during the course of project. I am highly indebted to her for constantly encouraging me by her enthusiasm for my work. She has been of great help to boost my confidence in both my academic and other endeavors.

I also wish her a very enriching experience with the new journey she has embarked upon.

Prathmesh Prabhu
November 2009
Indian Institute of Technology, Bombay

Chapter 1

Introduction

Games have long been used to model interaction between selfish adversaries competing to gain the maximum possible benefit in a given situation. In the case of multiplayer games, objectives no longer remain necessarily opposing and each actor aims to do the best he can. The idea of strategies that are best for everyone is captured through Equilibria which have recently been elegantly extended to multiplayer games and treated in depth. We study the concepts behind different types of equilibria on finite graph games and try to solve some related problems that can help in deciding well paying stable strategies. Later, with the understanding gained from this study, we propose a test of stability of equilibria and try to evolve stable equilibria.
This theory is directly applicable in verification problems where multiple agents have their own goals and must work in a common system. Different results from this field have also been used in the past to solve other automata theoretical problems. The later part of the work is of independent theoretical interest, as the questions raised are foundational to the study of games.

Organization of the report

Chapter 2 sets up the preliminaries by introducing the reader to basic concepts related to (stochastic) graph games & various equilibria and motivates the difference through an example. Chapter 3 surveys existing results that are useful for later development. The main developments in this work are divided into two independent sections. First, chapter 4 deals with two problems recognized from the literature associated with equilibria in stochastic games. The two results in this chapter extend a known results to the stochastic case and improve an existing proof. Later, chapter 5 delves into the meaning of equilibrium and uses the familiarity gained with the concept to ask foundational questions about stability of strategy profiles. Towards the end it proposes partial solutions. We sum up the report with a short conclusion in chapter 7.

Chapter 2

Preliminaries

We assume throughout this report that the reader is familiar with basic concepts and terminology of automata and language theory. We now define games and some basic concepts related to them that will act as a base for the development through the report.

2.1 Graph Games

The games we study are turn based games of perfect information played on graphs.

Definition 1 *An* ***infinite*** (***turn-based, qualitative***) ***multiplayer game*** *is a tuple* $G = (\Pi, V, (V_i)_{i\in\Pi}, E, \chi, (Win_i)_{i\in\Pi})$ *where* Π *is the set of players,* $(V_i)_{i\in\Pi}$ *is a partition of* V *into the position sets for each player,* $\chi : V \to C$ *is a coloring of the positions by some set* C*, which is usually assumed to be finite, and* $Win_i \subseteq C^\omega$ *is the winning condition for player* i*.*

The structure G on which these games are played is called the *arena*. An *initialized* game begins with a token placed on an initial vertex v. The game proceeds by moving the token along graph nodes. The owner of the current vertex makes a *move* by moving the token to the next vertex along an edge. An infinite sequence of these moves gives us a *play*. In order to ensure that every play is infinite, it is assumed that the set $uE := \{v \in V : (u, v) \in E\}$ of vertices that are the successors of u is non empty. A play is winning for player

i if it belongs to Win_i. Beginning at vertex v_0, a typical play proceeds as $v_0, v_1, v_2, v_3, \ldots$ Any prefix of a play is called a *history* of the game. v_0, v_1, v_2 and v_0, v_1, v_2, v_3 are example histories.

finite graph games and non-terminating Σ-tree games: The games that we discuss in this report are played on finite graphs, i.e., V is of finite cardinality. Another important class of graph games that we use in several proofs are the Σ-tree games. The underlying graphs for these games are non terminating trees.

Zero sum games: Another important restriction to the types of games we consider is laid by requiring that for any feasible play on the graph, one and only one of the players wins. Such games are called zero sum. The name derives from the classical origins when two player zero sum games were considered and this condition implied that the sum of "value" for the two players was always zero.

Definition 2 *A strategy σ for player i is a probability distribution $V^*.V \rightarrow D(V)$ where $D(V)$ is the space of probability distributions on the set of vertices V. A play is said to be consistent with the strategy σ if after history h of the game, at the position v, $\sigma(h, v)$ is the probability distribution of the next move chosen by the player π who owns the vertex v. Note that $D(u) = 0$ if $(v, u) \notin E$*

- A strategy is called pure if σ is a function into V, i.e., if the player chooses the next node deterministically for each history.
- A strategy is called finite memory if it only uses finite number of previous nodes to determine the move at the current node
- A strategy is called stationary if the next move is determined based only on the current node.
- A strategy is called positional if it is both pure and stationary.

Definition 3 *A strategy profile is an n-tuple of strategies $(\sigma_1, \sigma_2, \sigma_3, \ldots \sigma_n)$ such that each σ_i is a strategy for player i. A play is said to be consistent* *with strategy profile* $(\sigma_1, \sigma_2, \sigma_3, \ldots \sigma_n)$ *if for a given history h and vertex v,* $D(u) = \sigma_i(h, v)$ *where* $v \in V_i$

We will use the notation $\langle(\sigma_i)_{i\in\Pi}\rangle$ to refer to a strategy profile where player i has the strategy σ_i and $\langle(\sigma_i)_{i\in\Pi\setminus\{j\}}, \pi\rangle$ to refer to a strategy profile where each player except j has the strategy σ_i and player j has the strategy π.

The very basic question that one is prompted to ask given a two player zero sum game is whether there exists a strategy for either player that guarantees a winning play. This idea is captured by determinacy. A game is said to be determined if there exists a strategy σ for one of the players such that that player wins the game for any strategy π of the other player.

A final distinction must be made here between concurrent and turn based games. In a concurrent game, all players make simultaneous moves at each vertex and the probability distribution for the next vertex is determined by all these actions together. We do not consider concurrent games in detail here, but use them briefly in quoting a known result. On the other hand, in turn based games, only one player makes a move at each vertex (and that player is said to own that vertex). Henceforth, unless otherwise mentioned we refer to turn based games as simply games.

Multiplayer games have found application in modeling systems where multiple agents interact with selfish goals. They have been recently applied in the study of driver verification where many independent components are involved.

The most studied graph games fall under the class of ω-regular games. ω-regular games are graph games in which each player wins on plays that form an ω-regular set.

Let α be a generic play of the game. Following are well studied special cases of ω-regular winning sets:

1. *Buchi:* Specified by a set $B \subseteq C$; α is winning if the set of colors occurring infinitely often in α, $\Omega(\alpha)$ is subset of B, $\Omega(\alpha) \subseteq B$

2. *co-Buchi:* Specified by a set co-B$\subseteq C$; α is winning if the set of colors occurring infinitely often in α, $\Omega(\alpha)$ is disjoint from co-B, $\Omega(\alpha) \bigcap$ co-B $= \phi$

3. *Parity:* Specified by a priority relation on the colors in C; α is winning if the least color occurring in $\Omega(\alpha)$ is odd

4. *Rabin:* Specified by a family of pair of sets (E_i, F_i), $E_i, F_i \subseteq C$; α is winning if $\exists(E_i, F_i)(\Omega(\alpha) \bigcap E_i \neq \phi \wedge \Omega(\alpha) \bigcap F_i = \phi)$

5. *Streett:* Specified by a family of pair of sets (E_i, F_i), $E_i, F_i \subseteq C$; α is winning if $\forall(E_i, F_i)(\Omega(\alpha) \bigcap E_i = \phi \vee \Omega(\alpha) \bigcap F_i \neq \phi)$

6. *Muller:* Specified by a family of sets F, $F \subseteq C$; α is winning if $\Omega(\alpha) = F$ for some F.

Given a play α it is possible to find out whether or not it is winning for each of the players by looking at their winning sets and checking the winning condition. A few points to note here are that ω-regular winning conditions are prefix independent, hence they do not depend on any finite history of the game. Also, Parity sets include Buchi sets, Rabin and Streett sets include Parity sets and are themselves a part of Muller sets. Hence, Muller sets are the most general form of ω-regular sets.

2.2 Equilibria

A very important concept that captures the stability of a strategy profile is the concept of equilibrium. Given a profile, a player would want to use a strategy that will give her the maximum benefit. Equilibria try to capture strategy profiles where, under certain assumptions, none of the players would want to change their strategies.

Definition 4 *A strategy profile $\langle(\sigma_i)_{i\in\Pi}\rangle$ is a **Nash Equilibrium** if no player can unilaterally change her strategy to some other strategy π and in-*

crease her payoff.

$$payoff_j \ \langle(\sigma_i)_{i\in\Pi}\rangle \geq payoff_j \ \langle(\sigma_i)_{i\in\Pi\setminus\{j\}}, \pi\rangle$$

for all player j strategies π

This means that under lack of communication, every player is doing her best against the strategy of the remaining players.
A stronger equilibrium is defined using the concept of Subgames. for a given game G = (Π, V, $(V_i)_{i\in\Pi}$, E, χ, $(Win_i)_{i\in\Pi}$, v_0), a subgame $G|_h$ = (Π, V, $(V_i)_{i\in\Pi}$, E, χ, $(Win_i|_h)_{i\in\Pi}$, v_0) is defined as the game from v such that $v_0.h = v$ and $\alpha \in Win_i|_h \equiv h.\alpha \in Win_i$. In effect, a subgame is a game continuing after an initial history h of the game. For any strategy σ of G, a natural restriction $\sigma|_h$ is given by $\sigma|_h(x, v) = \sigma(h.x, v)$, i.e., the player plays as if the game had actually begun at v_0.

Definition 5 *A strategy profile $\langle(\sigma_i)_{i\in\Pi}\rangle$ is a* ***Subgame Perfect Equilibrium*** *if for every feasible history h of the game, the profile $\langle(\sigma_i)_{i\in\Pi}\rangle|_h$ is a Nash Equilibrium in the game $G|_h$*

A subgame perfect equilibrium tries to capture situations where by taking a non optimal choice in the course of the game, a player can induce another player to change her strategy as well, leading to an increased payoff for the first player in the end. Note that every Subgame Perfect Equilibrium is a Nash Equilibrium

Definition 6 *A strategy profile σ is called* ***Secure*** *if for all players $i \neq j$ and for each strategy σ' of j it is the case that*

$$\langle(\sigma_i)_{i\in\Pi}\rangle \notin Win_j \vee \langle(\sigma_i)_{i\in\Pi\setminus\{j\}}, \sigma'\rangle \in Win_j$$

$$\Rightarrow \langle(\sigma_i)_{i\in\Pi}\rangle \notin Win_i \vee \langle(\sigma_i)_{i\in\Pi\setminus\{j\}}, \sigma'\rangle \in Win_i$$

A strategy profile is secure if none of the players can decrease some other player's payoff without decreasing their own payoff. A **Secure Equilibrium** is a Nash Equilibrium that is also Secure.

Subgame Perfect Equilibrium and Secure Equilibrium are two different extensions of the notion of Nash Equilibrium

2.3 Stochastic games

Definition 7 *A* ***two player stochastic game*** *is a tuple* $G = (\Pi,\ V_1, V_2, V_0,$ $E,\ \chi,\ p,\ Win_1,\ Win_2)$ *where* Π *is a directed graph with vertex set* $V = V_1 \bigcup V_2 \bigcup V_0$, $\chi : V \to C$ *is a coloring of the position by some set C, which is usually assumed to be finite,*$p : V_0 \to D(V)$ *is a function from vertices* V_0 *to* $D(V)$ *and* $Win_1 \subseteq C^\omega$ *and* $Win_2 \subseteq C^\omega$ *are the winning conditions for player 1 and 2 respectively.*

Here player 0 is *nature*. Nature chooses one of the successor nodes from its vertex probabilistically, as given by p. These games are often called $2\frac{1}{2}$ games because nature plays with a fixed stationary strategy. $n\frac{1}{2}$ games can be similarly defined with n players and the $n+1^{th}$ player being nature.

Definition 8 *The* ***value*** *of a game for a player i is the maximum payoff she can guarantee against any strategies played by her opponents. An* $n\frac{1}{2}$ *game is said to be determined if there exist strategies such that each player can achieve payoff equal to the value for that player.*

An example to explain equilibria

Example 1 *Let us see a* $2\frac{1}{2}$ *player game that clearly explains and distinguishes these equilibria.*

Consider the game shown in figure 2.1. The circles are player nodes owned by the player indicated. The diamonds are stochastic vertices with each outgoing edge marked with probability of that edge being taken. The squares are end

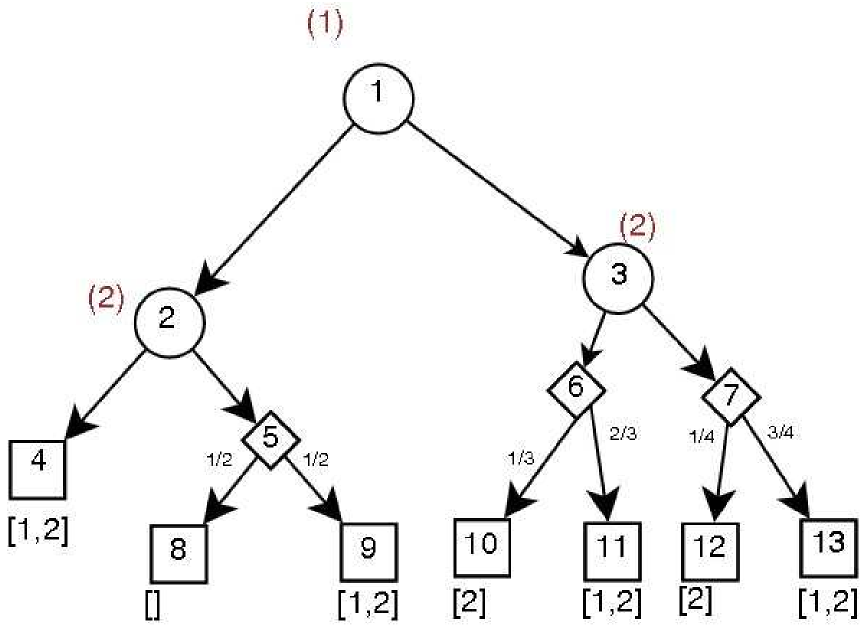

Figure 2.1: A two player game with different equilibria

vertices and are labeled with the players that win on reaching that vertex. Consider the following stationary positional strategy profiles:

P_1:	$\sigma_1(1) = 2,$	$\sigma_2(2) = 5,$	$\sigma_2(3) = 6$
P_2:	$\sigma_1(1) = 3,$	$\sigma_2(2) = 5,$	$\sigma_2(3) = 7$
P_3:	$\sigma_1(1) = 3,$	$\sigma_2(2) = 5,$	$\sigma_2(3) = 6$
P_4:	$\sigma_1(1) = 2,$	$\sigma_2(2) = 4,$	$\sigma_2(3) = 7$

We can make the following observations.

- P_1 is not a Nash Equilibrium because player 1 can increase her payoff by moving to node 3 instead of 2.
- P_2, P_3 and P_4 are all Nash Equilibria.
- P_2 is not a secure equilibrium since player 2 can move from node 3 to node 6 without decreasing her own payoff, while decreasing player 1s payoff. P_3, the profile thus obtained is a secure equilibrium.
- P_2 is not a subgame perfect equilibrium either. Since the profile is not a Nash Equilibrium after the history $1 \to 2$. P_4, on the other hand is

an SPE. Note here that P_2 failed to be an SPE because player 1 could change her strategy so that she could induce player 2 to change hers, and the final profile was better for both. We will dwell more on this interpretation of SPE later in the report.

Chapter 3

Existing results

3.1 Existence of equilibria

The following results are classical results for 2 player game determinacy. Later results build on these old results.

Theorem 1 [4] *Two player games with Borel winning conditions are determined.*

Theorem 2 [5] *Two player games with ω-regular winning conditions are positionally determined.*

Using the positional determinacy of 2 player games, Chaterjee et al showed that any multiplayer game with Borel winning conditions has a Nash Equilibrium. For n$\frac{1}{2}$ player games, the following is a parallel result.

Theorem 3 [1] *There exists a Nash Equilibrium with pure strategy profile in every turn-based stochastic game with parity payoff conditions.*

The existing proof for this result uses a much stronger result about discounted concurrent games. In treating other games for similar results, we hope to use alternative proof techniques than the one used here.
We state and prove the following two results about SPEs for two player and n player games originally from [2].

Theorem 4 [2] *Let (G,v_0) be an initialized two player game such that every subgame is determined. Then there exists a strategy profile (σ_1,σ_2) that is a subgame perfect equilibrium.*

Proof. Since we are concerned only with the existence of a subgame perfect equilibrium and not the complexity of finding one, we can consider any play on the graph to played on an equivalent Σ-tree generated by unrolling the graph on each node starting from the initialization vertex v_0. Since we already know that every subgame is determined, given a history h, either player 1 or 2 has a winning strategy from this position. Also, from theorem 2 we need to consider only positional strategies on the Σ-tree. Let the optimal strategy after history h for player i be denoted by σ_i^h.

If we allow both players to play their best strategies from every vertex in the Σ-tree, then either of the player must win. The problem in this simplistic approach is that in any play of the game, players might switch strategies to different σ_i^h infinitely many times. Thus no strategy is played for an infinite suffix of the play, and hence no player is really playing her winning strategy.

For a history h, define a partition h_1, h_2 of h to be a good partition for player i if

- $h = h_1 \circ h_2$
- $\sigma_i^{h_1}|_{h_2}$ is winning in $(G|_h, v_0)$
- h_1 is the minimal possible such h_1

Now, define the strategy for each player as follows:

$$\sigma_i(hv) = \sigma_i^{h_1}(h_2 v) \text{ where } h_1 \circ h_2 \text{ is good partition of h.}$$
$$\text{if no good partition exists, set } \sigma_i(hv) = \sigma_i^h(v)$$

Now consider any play after an initial history h. Let the strategy chosen as above be $\sigma_i^{h_1}$. We claim that if v_1, v_2, v_3, v_4 …, v_k be the subsequent vertices followed in the play according to this strategy, $\sigma_i^{h_1}$ is at least one permissible partition at v_k, since it is winning for player i from $h_1 \circ h_2$

against all strategies of player $\bar{i}$ and v_k is a state reached in accordance with this strategy and hence must still be winning. Secondly, assume that a better partition exists for h, v_1, v_2, v_3, v_4 …, v_k; since in the new partition $h'_1 \leq h_1$, this process must reach a fixed point. And hence on no path can there be an infinite number of switches in the strategy. So the above problem is solved. And σ_i as defined gives the required strategy profile.

Theorem 5 [2] *n player stochastic games with parity winning conditions have an SPE.*

3.2 The problems NE & SPE

Define the problem NE as follows:

Definition 9 *Given a multiplayer game, determine if there exists a Nash Equilibrium with payoff vector between $\bar{x}$ and $\bar{y}$.*

We say that the payoff vector $\bar{p}$ is between $\bar{x}$ and $\bar{y}$ when for each player i, $x_i \leq p_i \leq y_i$.
We discuss this problem in context of stochastic games. Since it is too difficult to talk about NE in general, the following restrictions of the problem are treated independently.

- PosNE: all players use positional strategies
- StatNE: all players use stationary strategies
- FinNE: all players use finite memory strategies
- PureNE:all players use pure strategies
- QualNE: the vectors $\bar{x}$ and $\bar{y}$ consist only of 0s and 1s.

It can be shown that these problems are indeed independent, i.e., there are games with one type of NE in a given range but not the other.

Definition 10 ***Simple Stochastic Multiplayer Games(SSMG)*** *are stochastic multiplayer games with payoff reachability objectives on terminal nodes. The only final vertices are terminal vertices where the game keeps looping forever and some players win while others lose.*

The motivation for using this restrictive type of games is that it is a subset of Buchi games and hence all ω-regular games. All the undecidability and hardness results for SSMG carry over to higher form of winning.
The following results summarize the current knowledge about NE in stochastic games

Theorem 6 [6] [7] PosNE is NP-complete for all ω-regular objectives as well as SSMGs.

Theorem 7 [6] [7] StatNE is in P-space for all ω-regular objectives as well as SSMGs.

Theorem 8 [6] PureNE and FinNE are undecidable for games with at least 9 players and 13 players respectively.

A similar problem can be defined for SPE as follows

Definition 11 *Given a multiplayer game, determine if there exists a Subgame Perfect Equilibrium with payoff vector between $\bar{x}$ and $\bar{y}$*

Nothing in known about SPE in stochastic games.

Chapter 4

Results on stochastic games

The first result is on the existence of subgame perfect equilibria in $2\frac{1}{2}$ player zero sum stochastic games.

4.1 Two player stochastic zero sum games

Since we are concerned only with the existence of equilibria and not the complexity of their computation, we argue about properties of Σ-tree games obtained from unraveling the game graph at every node from the start node.

Lemma 1 *For a game (G, v_0), let S he the set of all reachable histories h and let $\bar{\sigma}_0 = \langle(\sigma_i)_{i\in\{1,2\}}\rangle$ be a strategy profile optimal to both players. Then, either $val(\bar{\sigma}_0|_h) = val(G|_h)$ or h is not reachable in the strategy profile $\bar{\sigma}_0$*

Proof. Let h be reachable with strategy profile $\bar{\sigma}_0$. And assume that condition one does not hold. Hence,
$\exists$ history y such that val$(\bar{\sigma}_0|_y) <$ val$(G|_y)$
Then, in the game $G|_y$, $\exists$ a strategy σ_1^y for player 1 (against player 2) that gives him a higher value, viz, val$(G|_y)$.
Consider the strategy profile,
$\langle($ σ_1 for all histories whose prefix is not y, σ_1^y for suffixes of y$), \sigma_2\rangle$
then, player 1 has effectively increased his payoff to more than his value against an "optimal" strategy of player 2.
This is a contradiction.

Hence one of the two assumptions above must be wrong.

Next, we show with a proof similar to theorem 4 that a subgame perfect equilibrium exists for $2\frac{1}{2}$ player games.

Theorem 9 *$2\frac{1}{2}$ player zero sum games are subgame perfect determined.*

Proof. For a given history h, define a partition h_1, h_2 of h to be good for player i if

- $h = h_1 \circ h_2$
- $h_2 \in S_{\bar{\sigma}_i^{h_1}}$ in the game $G|_{h_1}$
- h_1 is the minimal possible such h_1

Now define the strategy for each player as follows:
$\sigma(hv) = \sigma^{h_1}(h_2 v)$ if a good partition h_1, h_2 exists,
$\sigma(hv) = \sigma^h(v)$ otherwise.

It remains to prove that this strategy profile is optimal for all histories h. Consider the game at history h. Let the strategy chosen as above be σ^{h_1}. We are assured that this strategy gives the optimal value to player 1 from lemma 1. Now consider any path consistent with the profile σ^{h_1} from h. With a reasoning similar to theorem 4 we can show that only a finite number of switches in strategy occur on any path, And hence, the players obtain their optimal value from history h.
This proves the result.

The next result is an effort towards reducing the gap in results associated with decidability of PureNE (theorem 8) We show with a proof technique very similar to the original suggested in [6] that PureNE is undecidable even with 5 players. The question of 2 player PureNE has proved elusive so far.

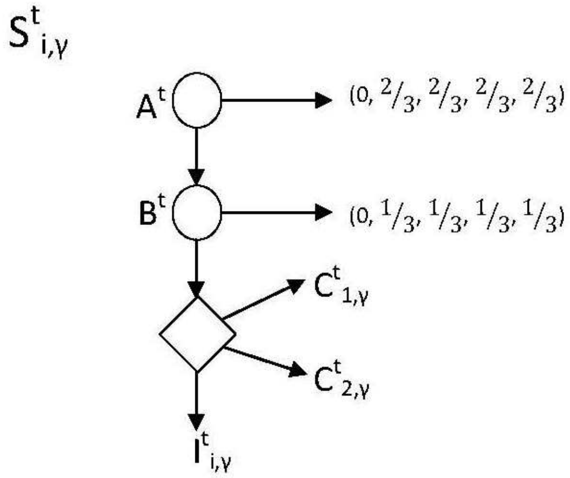

Figure 4.1: The I-gadget

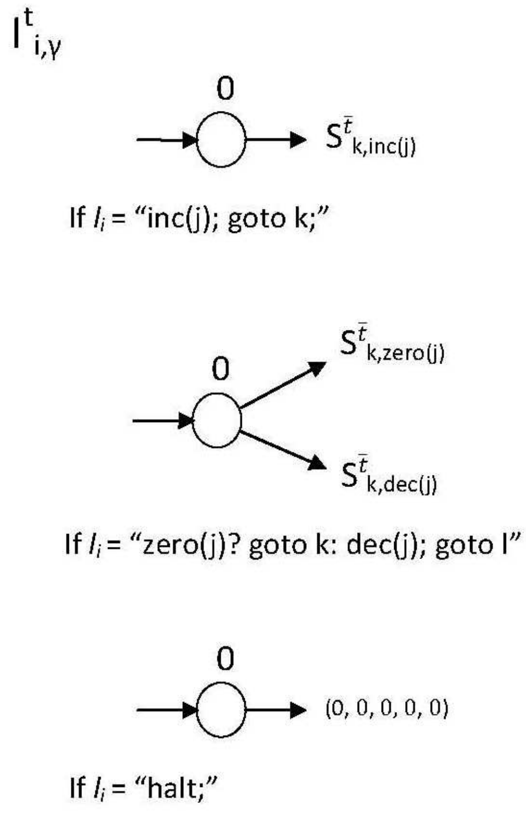

Figure 4.2: The S-gadget: Player 0 chooses the next I-gadget

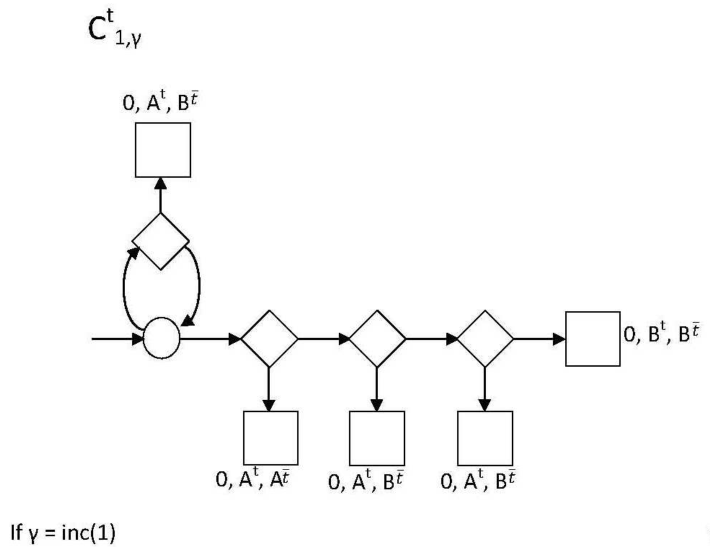

Figure 4.3: Terminal gadget for incrementing counter 1

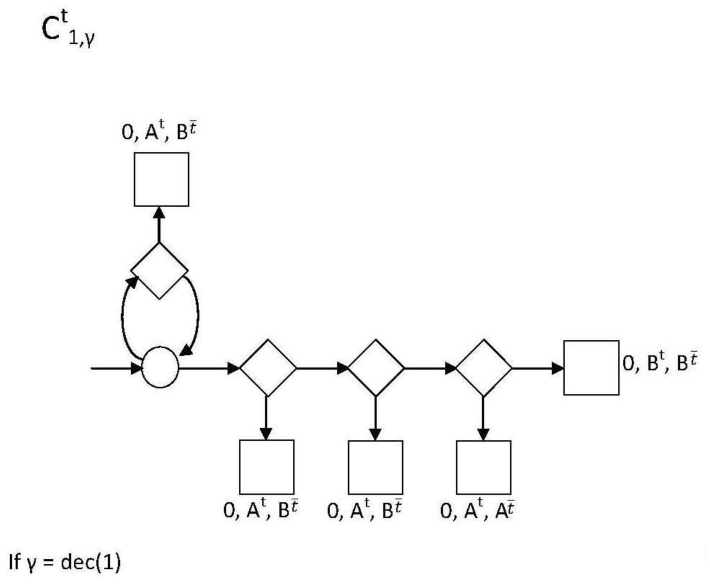

Figure 4.4: Terminal gadget for decrementing counter 1

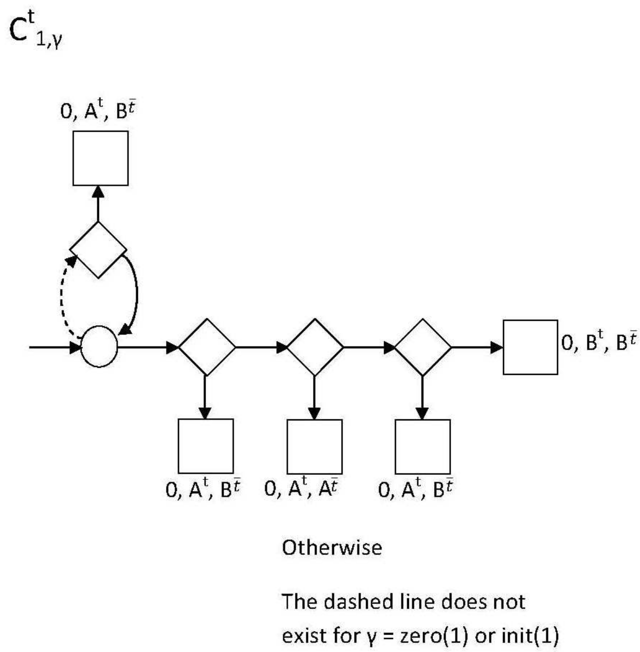

Figure 4.5: Terminal gadget for counter 1 in other cases

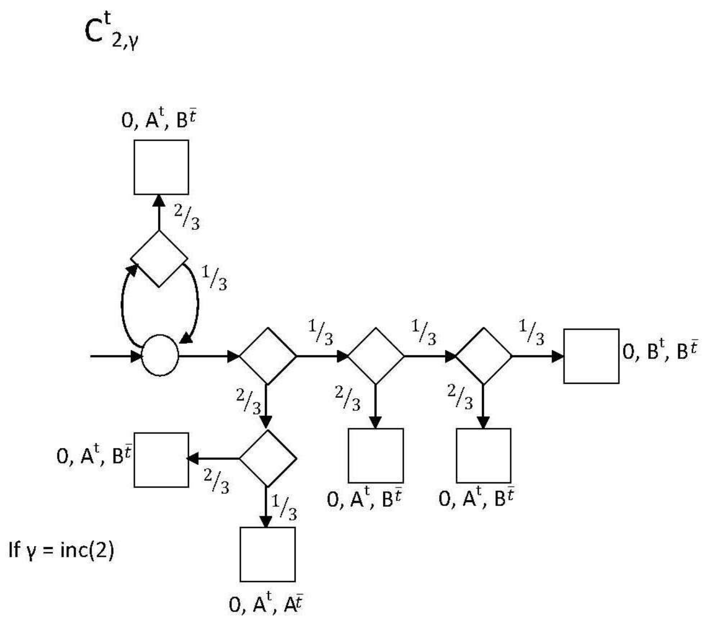

Figure 4.6: Terminal gadget for incrementing counter 2

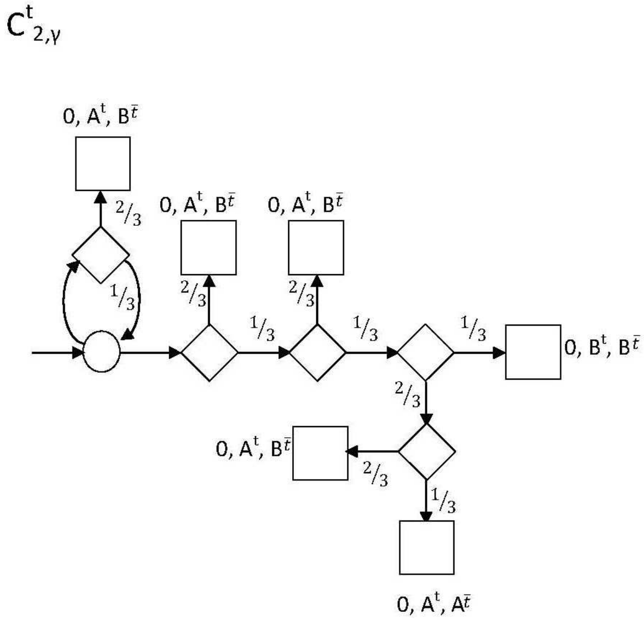

Figure 4.7: Terminal gadget for decrementing counter 2

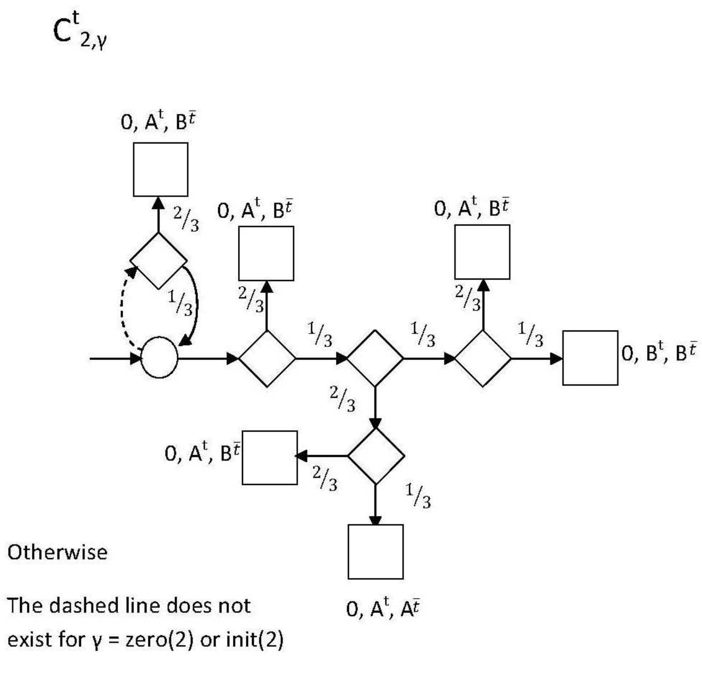

Figure 4.8: Terminal gadget for counter 2 in other cases

4.2 Undecidability of PureNE

In order to prove the undecidability of PureNE, we reduce the halting problem of two counter machine running on a program to a game such that the game has a Nash Equilibrium where a certain player wins iff the machine halts on the given program.

Definition 12 ***A two counter machine*** *consists of two counters that can be incremented, decremented and checked for zero/non-zero value. It takes as input the program text and beginning at the top of the program, follows the instructions manipulating the counters as directed. The machine stops when it encounters the instruction "halt".*

One possible instruction set for the machine is:

- "inc(j); goto k;"
- "zero(j)? goto k :dec(j); goto l;"
- "halt"

where k and l are positions in the program text, and j is either 1 or 2, refers to the two counters.
With this set of instructions, it is known that the halting problem is undecidable.

Lemma 2 *Given a program running on a counter machine as defined above, it is undecidable to determine whether the machine halts on executing the given program.*

In order to show our problem to be undecidable, we must encode a generic program for the counter machine as a game. The game G is played by 5 players, player 0, A^1, B^1, A^0 and B^0. *There exists a Nash Equilibrium such that player 0 wins iff the counter machine does not halt on the program encoded by this game.* Throughout this proof, A^t refers to either of A^1 or A^0 and $\bar{t} = (1 - t)$. We use the gadgets given in figures 4.1 through 4.8 to synthesize the game for a given program. For every line i of program, there is a gadget $S^1_{i,\gamma}$ (figure 4.2) where $\gamma \in$ {inc(j), dec(j), zero(j), init } for j $\in$ {0,1}; there are the end gadgets $C^t_{1,\gamma}$ and $C^t_{2,\gamma}$ associated with every S-gadget; and finally at the end of every S-gadget is an I-gadget, $I^t_{i,\gamma}$ (figure 4.1). Beginning with $S^1_{0,init}$. the gadgets naturally get connected in a network.

Following a program execution: Each $S^t_{i,\gamma}$ represents a program instruction being executed on the machine. To ensure that program execution is simulated faithfully, we must ensure that player zero makes the correct choices at all the I-gadgets, i.e, he chooses that edge which will get executed during the run. We do this by remembering the counter values at all times during the run and using them to force player zero to make correct moves.

Encoding the counter value: We will encode the counter values of the two counters at each gadget in the $C^t_{1,\gamma}$ and $C^t_{2,\gamma}$ respectively. The counter value at any point during program execution is exactly the same as the maximum number of times player zero plays the marked edge in the C-gadgets.

Maintaining the correct counter values: So the problem reduces to ensuring that the number of times player zero chooses the marked edge, call it c^j_n on the n^{th} step for counter j, is equal to the current value of the counters. We achieve this by ensuring the local relationships between c^1_n & c^1_{n+1} and c^2_n & c^2_{n+1} respectively.

Theorem 10 *Let $\bar{\sigma}$ be a strategy profile such that player zero almost surely wins. Then $\bar{\sigma}$ is a Nash Equilibrium if and only if*

$$c^j_n = \left\{ \begin{array}{cc} c^j_n + 1 & if \gamma_{n+1} = inc(j) \\ c^j_n - 1 & if \gamma_{n+1} = dec(j) \\ c^j_n = 0 & if \gamma_{n+1} = zero(j) \\ c^j_n & otherwise \\ & \end{array} \right\}$$

for j = 1 & 2

Proof. Towards this, we prove the following lemma:

Lemma 3 *The profile $\bar{\sigma}$ is a NE iff a_n, the probability of winning for player A^t at a node owned by A^t is equal to $\frac{2}{3}$.*

Proof. $\Rightarrow$ assume that $\bar{\sigma}$ is a Nash Equilibrium, then

$$a_n \geq \frac{2}{3}$$

since otherwise, player A^t would chose to leave the game at that point and increase her payoff and player 0 looses. Similarly, if b_n is the probability of winning at a node owned by B^t,

$$b_n \geq \frac{1}{3}$$

Also, at every vertex such at player 0 wins, exactly one of A^t and B^t win.

$$a_n + b_n = 1$$

$$a_n \leq \frac{2}{3}$$

$$a_n = \frac{2}{3}$$

$\Leftarrow$ assume that $\mathrm{a}_n = \frac{2}{3}$, then from reasoning similar to above, $\mathrm{b}_n = \frac{1}{3}$. Hence, neither A^t nor B^t can improve their payoff. Player zero already wins almost surely. Hence $\bar{\sigma}$ is a Nash Equilibrium.

Next, we show how these probabilities a_n are maintained locally.

Lemma 4 *$a_n = \frac{2}{3}$ iff p_n, the probability of winning for A^t in this and the next component, is $\frac{1}{2}$.*

Proof. $\Rightarrow$ let $\mathrm{a}_n = \frac{2}{3}$. Note that

$$a_n = p_n + \frac{1}{4} \times a_{n+2}$$

also

$$a_n = a_{n+2} = \frac{2}{3}$$

it follows:

$$p_n = \frac{1}{2}$$

$\Leftarrow$ if $p_n = \frac{1}{2}$ then we have,

$$a_n = \frac{1}{2} + \frac{1}{4} \times a_{n+2}$$

It is easy to see that the only value of a_0 and a_1 such that $0 \leq a_n \leq 1$ for all n is when $a_0 = a_1 = \frac{2}{3}$. But that means that $a_n = \frac{2}{3}$ for all n.

It remains to prove that p_n is $\frac{1}{2}$ iff the condition in the theorem is satisfied. Towards this we state the following lemma:

Lemma 5 *if*

$$\left(\frac{1}{2}\right)^{c_n^1} - \left(\frac{1}{2}\right)^{c_{n+1}^1} + \left(\frac{1}{3}\right)^{c_n^2} - \left(\frac{1}{3}\right)^{c_{n+1}^2} = 0$$

then

$$c_n^1 = c_{n+1}^1 \bigwedge c_n^2 = c_{n+1}^2$$

given c_t^j *are all non negative.*

Also note the following expressions derived from the arrangement of gadgets: If prob(n)(C) denotes the probability of winning of player A^t in the gadget C on the nth gadget, and prob(n+1)(C) the corresponding value for the n+1th gadget then we have the following:

$$p_n = \frac{1}{4} \cdot \left(\frac{1}{2}\right)^{c_n^1+3} + \frac{1}{4} \cdot prob(n, C_{2,\gamma}) + \frac{1}{2} \cdot (\frac{1}{4} \cdot prob(n, C_{1,\gamma}) + \frac{1}{4} \cdot prob(n, C_{2,\gamma}))$$

We can also derive the values of $prob(n, C)$ from the game graph as:
prob(n,C):

$$prob(n, C_{1,\gamma}^n) = 1 - \left(\frac{1}{2}\right)^{c_n^1+3}$$

$$prob(n, C_{2,\gamma}^n) = 1 - \left(\frac{1}{3}\right)^{c_n^2+3}$$

prob(n+1,C):

$$prob(n+1, C_{1,inc(1)}^n) = \left(\frac{1}{2}\right)^{c_{n+1}^1+1}$$

$$prob(n+1, C_{2,inc(2)}^n) = \left(\frac{1}{3}\right)^{c_{n+1}^2+2} \cdot 2$$

and so on. Hence, the above equation reduces to the form

$$\left(\frac{1}{2}\right)^{c_n^1} - \left(\frac{1}{2}\right)^{c_{n+1}^1} + \left(\frac{1}{3}\right)^{c_n^2} - \left(\frac{1}{3}\right)^{c_{n+1}^{21}} = 0$$

and using lemma 5, we get the required result.

Chapter 5

Capturing Equilibria

5.1 Philosophical basis

Equilibria were introduced to capture inherent stability in strategy profiles. Nash equilibrium capture this stability quite well in the case of two player zero-sum games. If (σ_1, σ_2) is a strategy profile, then player 1 is employing the best possible strategy against player 2's strategy, and vice-versa. Indeed, player 1 changing to a strategy σ_1' can not increase her payoff but may give player 2 a chance to change his strategy to σ_2' increasing his payoff against σ_1'. Since this is a zero sum game, this means that player 1's payoff reduces through this series of strategy changes. Employing a similar reasoning, neither players chooses to deviate from (σ_1, σ_2). Hence, this strategy profile is stable, in the sense that both players stick to this profile once decided.
In the case of non zero-sum games, It is possible that player 1 chooses σ_1' over σ_1. From the assumption that (σ_1, σ_2) is a nash equilbrium, $val_1(\sigma_1', \sigma_2) \leq val_1(\sigma_1, \sigma_2)$. As before, let player 2 now choose σ_2' increasing his payoff against σ_1'. In this case though, *it is possible to have* $val_1(\sigma_1', \sigma_2') \geq val_1(\sigma_1, \sigma_2)$. Thus, although (σ_1, σ_2) is a nash equilibrium, two rational players playing the game would not stick to the strategies but move on to the strategies (σ_1', σ_2'). Put down somewhat boldly, nash equilibria in this case are not truly equilibrium states. Hence, we see that in the non zero setting, it is possible to have in a sense non stable nash equilibria. A detailed discussion

of the stability of equilibria leads us to enquire into the process of determination of equilibria, and the precise meaning of moving-on / staying with a strategy. In the section on processes, we look at this process and finally try to abstract away once again to captures some of the insights gained. Below, we state a litmus test for equilibria that stems from similar thoughts, which we use to show the inadequacy of known concepts.

Concept 1 ***Litmus test:*** *We say that an equilibrium profile* (σ_1, σ_2) *fails the litmus test if the following sequence of actions is possible. player 1 changes her strategy to* σ_1'*, and then player 2 changes his profile to* σ_2'*. In the end, we find* $val_1(\sigma_1', \sigma_2') > val_1(\sigma_1, \sigma_2)$ *and* $val_2(\sigma_1', \sigma_2') > val_2(\sigma_1, \sigma_2)$

In the sequel, we say an equilibrium is stable if it passes the litmus test.

5.2 Treatment of known Equilibria

In this section, we show how subgame perfect equilibria fail in capturing stability with respect to the litmus test. We later touch upon the essential difference between these concepts and secure equilibria.
We claim that subgame perfect equilibria are neither sufficient nor necessary conditions for a nash equilibrium profile to pass the litmus test.

5.2.1 Sufficiency of SPE for stability

Following is a counter example of a subgame perfect equilibrium that fails the litmus test.
Consider the game graph in figure 5.2.1 with V_1 = {1, 3, 4, 6, 7, 8, 9, 10} and V_2= {2, 5}. Let the initial node be {1} and the Buchi winning sets F_1 = F_2 = {7, 10}.
We claim that there exists a subgame perfect equilibrium such that both players lose, and it is possible for one player to induce a change to a better equilibrium. Here is a strategy profile for the game:

$$\sigma_1(1) = 3, \quad \sigma_2(2) = 4, \quad \sigma_1(3) = 2, \quad \sigma_1(4) = 1,$$
$$\sigma_2(5) = 6, \quad \sigma_1(6) = 2, \quad \sigma_1(7) = 6, \quad \sigma_1(8) = 1,$$

$$\sigma_1(9) = 8, \qquad \sigma_1(10) = 8$$

This is a pure profile. In the diagram, chosen edges are shown in bold. (σ_1, σ_2) is a subgame perfect equilibrium, since after any finite history of the game, the strategy profile forces the game back to node 1 or 2 and no player can unilaterally force visits to 7 or 10 infinitely often. Hence, both players lose.

figure 5.2.1

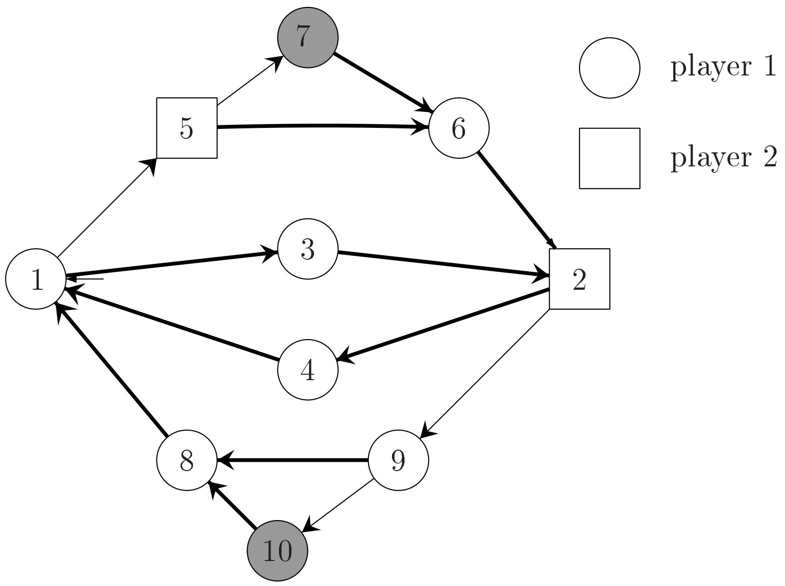

But (σ_1, σ_2) fails the litmus test. If player 1 changes his strategy at node 1 to $\sigma_1(1) = 5$, then player 2 can change her strategy at node 5 to $\sigma_2(5) = 7$. Then, the new strategy profile is a nash equilibrium where both players win. Note that, the first change ($\sigma_1(1) = 5$) does not make sense in the nash sense since it does not increase the payoff for player 1. On the other hand, the second change ($\sigma_2(5) = 7$) by player 2 is a completely selfish move, increasing her payoff from 0 to 1. In this process she ends up increasing player 1's payoff from 0 to 1 as well. This is exactly the instability we are interested in capturing.

5.2.2 Necessity of SPE for stability

It is easy to see that there exist games where nash equilibria that are not subgame perfect equilibria do pass the litmus test. For example, take the game in figure 5.2.2

figure 5.2.2

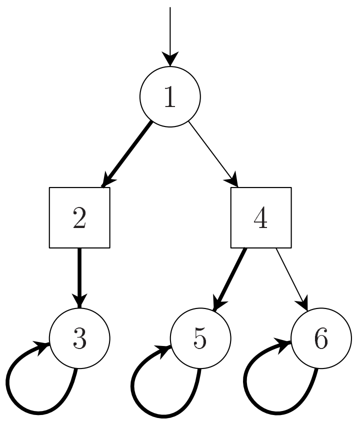

This is a game with Buchi winning conditions with,

$V_1 = \{1, 3, 5, 6\}$ $V_2 = \{2, 4\}$

$F_1 = \{3\}$ $F_2 = \{3, 6\}$

Then the strategy profile

$$\sigma_1(1) = 2, \qquad \sigma_2(2) = 3, \qquad \sigma_1(3) = 3,$$
$$\sigma_2(4) = 5, \qquad \sigma_1(5) = 5, \qquad \sigma_1(6) = 6$$

is not a subgame perfect equilibrium as the strategies restricted to the initial history $1 \rightarrow 4$ is not a nash equilibrium. But it passes the litmus test since both players win in the current profile and hence no strategy switch can be profitable to any player.

Hence we see that subgame perfect is neither necessary nor sufficient for stability defined with respect to our litmus test.

5.2.3 Semantics of subgame perfect equilibria

The above discussion forces us to question what subgame perfect equilibria actually capture. Following are the facts we know about SPE:

- It is a nash equilibrium
- Even for unreachable game states, a player has to play in a way that would ensure the best payoff against other players' moves from that state.
- In the words of the authors of subgame perfect equilibrium for graph games - *"respects the possibility of other player changing her strategy"* [2]

As the next example shows, the last point above is not entirely true.

figure 5.2.3

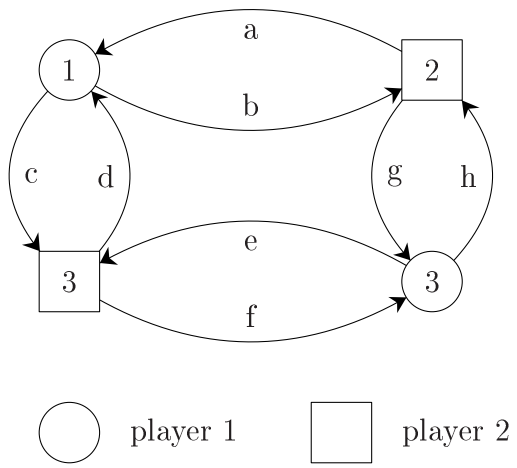

The above game is a two player non zero sum Muller game where the nodes are labeled with colors 1 thru 3. The winning sets are $F_1 = \{\{1, 2\}, \{3\}$

$\}$ and $F_2 = \{\{1, 2\}, \{2, 3\}\}$. The pure strategy profile (σ_1, σ_2) where the chosen edges are $\{a, b, d, h\}$ is a subgame perfect equilibrium with payoff $\{1, 1\}$. But if player 2 changes her strategy so that the active edges become $\{b, e, g, h\}$, then the new payoff vector is $\{0, 1\}$. Player 1 is now forced to change his strategy to ensure the best payoff.

figure 5.2.4

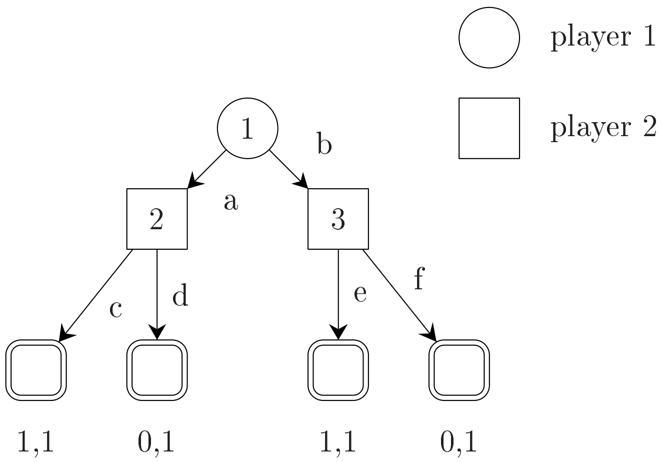

Indeed to "respect the possibility of other player changing her strategy" seems a very strong statement. This would require that for player 1, σ_1 is the best strategy against any strategy of player 2. The example in figure 5.2.4 shows that this not always possible, which is to be expected.
There does not exist σ_1 such that $val_1(\sigma_1', \sigma_2)$ is optimal for the following two player two strategies.

$$\begin{aligned} \sigma_2^1 &= d, e, \qquad \text{best } \sigma_1 = b \\ \sigma_2^2 &= c, f, \qquad \text{best } \sigma_1 = a \end{aligned}$$

Not knowing what player 2's strategy is reduces this game to an imperfect information game.
We can still say that subgame perfect equilibrium guarantees safety against a possible deviation from strategy by the opponent *for a finite history of the game.* As noted in the first example above, if the other player deviates from the declared strategy on an infinite number of terms in a play, the play might become non optimal for the first player.

Relation to secure equilibria

We must here note an important difference between secure equilibria and the stability properties we are interested in. Secure equilibria ensure that player 2 can not decrease player 1's payoff without incurring a cost herself, which is an added constraint put on the game structure and does not stem from the basic selfish game play. Since the games are non zero sum, there is no inherent advantage in reducing the opponent's payoff. On the other hand, the instability in our sense means that rational players would not consider the equilibrium advantageous from purely selfish motives, as it is possible to exploit this instability to increase one's payoff.

5.3 Maximal Nash Equilibria

We take a detour from the philosophical discourse on stability to discuss some results concerning maximal equilibria. We define two special types of nash equilibria, stric-max NE and max NE and try to answer some pertinent questions. The importance of these concepts becomes clearer in the following sections.

5.3.1 Strict-max NE for 2 players

We first define strict max NE for two players and look at some associated results before generalizing it to n players.

Definition 13 $(\sigma_1, \sigma_2) \in$ *strict-max NE iff*

$$\sim \exists (\pi_1, \pi_2) \in NE \circ \left\{ \begin{array}{c} val_1(\pi_1, \pi_2) > val_1(\sigma_1, \sigma_2) \\ \wedge val_2(\pi_1, \pi_2) > val_2(\sigma_1, \sigma_2) \\ \end{array} \right\}$$

Strictmax NE, then, is a nash equilibrium such that no other nash equilibrium gives both players better payoffs. Following result shows the existence of strict max NE.

Lemma 6 *Every 2 player non zero sum game with Borel winning conditions has a strict max NE.*

Proof.
We use the following related lemma in proving the above lemma:

Lemma 7 *For 2 player non zero sum games*

$$\forall \overline{\sigma} \in mixed \circ \{\{\overline{\sigma} \in NE \wedge val(\overline{\sigma}) = \{x_1, x_2\}\} \Longrightarrow$$

$$\exists \overline{\sigma}' \in pure \circ \{\overline{\sigma}' \in NE \wedge val(\overline{\sigma}') = \{y_1, y_2\} \wedge (y_1 \geq x_1) \wedge (y_2 \geq x_2)\}\}$$

With lemma 7 in place the result follows:

- For every n player non zero sum Borel game there exists a pure nash equilibrium profile $\overline{\sigma}$.
- From lemma 7 for any nash equilibrium with fractional payoffs, $\exists \overline{\sigma}' \in$NE such that the payoffs are strictly greater than earlier and are integral.
- If there exists $\overline{\sigma}''$ with payoffs greater than $\overline{\sigma}'$, then the above step may be repeated
- Since there are only finitely many integral payoffs, this repetition can not happen indefinitely.

Hence we must arrive at a nash equilibrium profile such that a profile with larger payoffs does not exist. This is the required strict-max NE profile.

figure 5.3.1

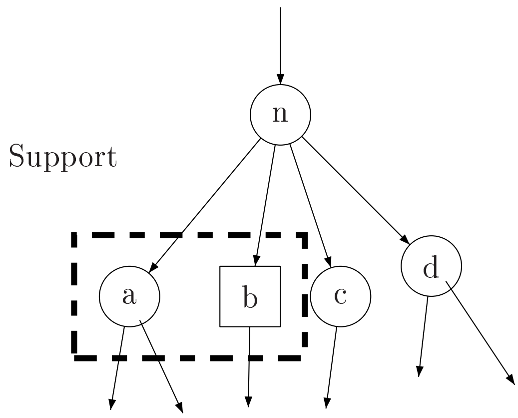

Proof of lemma 7

Let G = (V_1, V_2, E, F) be a game. We obtain game G′ be unrolling the game along the edges. Then, we know that every strategy in G can mapped to a positional strategy in G′ and vice-versa. Let $\overline{\sigma}'$ be the nash equilibrium profile obtained for G′. Then, starting at the root recursively apply the following transformation to $\overline{\sigma}'$ (fig 5.3.1):

Let a, b ∈ support(n), i.e.,

if $\sigma_{i|a \in V_i}(n)(a) > 0$ and

$\sigma_{i|b \in V_i}(n)(b) > 0$

Then indeed,

if $\sigma_{i|a \in V_i}(n)(a) = \sigma_{i|b \in V_i}(n)(b)$, since $\overline{\sigma}'$ is a nash equilibrium.

Otherwise, player i could have increased his payoff by choosing the node with higher value deterministically.

Without loss of generality, let

$$val_{1-i}(a) \geq val_{1-i}(b)$$

i.e. the other player has highest payoff if player i chooses node a.

Then modify $\overline{\sigma}'$ as follows:

1. Deterministically chose ‘a’ at n
2. Repeat this step for ‘a’ keeping $\overline{\sigma}'$ unchanged for the other subtrees.

Now, if the subsequent changes in $\overline{\sigma}'$ in the subtree under ‘a’ does not decrease the payoffs for both players then the payoffs at n have also not decreased for

either player.
Also, the new strategy profile is a nash equilibrium.

We know,

$$\overline{\sigma}' \in NE$$

hence

$$(val_i(a) = val_i(b)) > (val_i(c) = val_i(d))$$

and also, $val_i(a)$ has not decreased through our transformation.
$\Longrightarrow$ player i can not increase payoff changing strategy at n

Finally, starting at the root, a single play is deterministically followed, since at each node the choice made is deterministic.
$\Longrightarrow$ payoffs are natural numbers.
This proves Lemma 7.

Remark: The above proof of lemma 6 fails for n players because lemma 7 does not hold for n players as the following counterexample shows:

figure 5.3.2

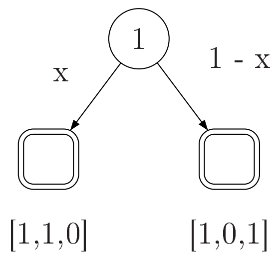

Shown in figure 5.3.2 is a 3 player game with node 1 where player 1 make a move and two terminal nodes. Here, player 1 wins no matter what strategy he employees. Now he may chose the edges with probabilities x and 1-x, for any x between 0 and 1. Every such strategy is a strictmax NE, since the payoffs of player 2 and 3 add up to 1 and increasing 2's payoff decreases 3's and vice versa. Hence, for the strategy with x=0.5, there is no "better" strategy profile with natural payoffs, as required by lemma 7.

Another pertinent question is whether strictmax NE for 2 players is unique. We answer this query in the negative with the example shown in figure 5.3.3.

figure 5.3.3

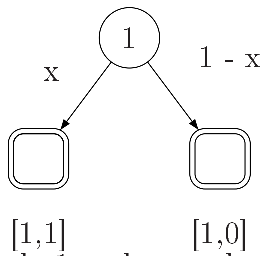

Here again, player 1 plays at node 1 and may chose to play the edges with any probability, every x gives a strictmaxNE.

Lemma 8 *There need not exist a unique strictmax NE profile for 2 player non zero sum games, and as a corollary for n player games.*

5.3.2 Strict-max NE for n players

We now extend the definition to n players:

Definition 14 $\overline{\sigma} \in$ *strict-max NE iff*

$$\sim \exists\ \overline{\pi} \in NE \circ (\forall\ i\ val_i(\overline{\pi}) > val_i(\overline{\sigma}))$$

The following lemma is the parallel of existence result in 2 player case.

Lemma 9 *Every n player non zero sum game with Borel winning conditions has a strict max NE.*

Proof:
Sketch: Note that it is enough to show that there exists a nash equilibrium where at least one player has the payoff 1, since no other equilibrium can then have strictly greater payoff for this player. We use the existence of subgame perfect equilibrium to get an initial pure nash equilibrium, and then if necessary, push the payoff of at least one player to 1.

Let G = (V,Π, O, E, F) be the game with Π set of players and O the ownership function from V to Π played on the unrolled game tree.
As usual, $G|_h$ denotes the game play restricted to history h, and $\overline{\sigma}|_h(x)$ gives the next move in the game $G|_h$ after an extend history h.x.
The existence proof for subgame perfect equilibrium in Borel games gives us a pure strategy profile $\overline{\sigma_0}$ such that
∀ h ∈ V* ($\overline{\sigma_0}|_h$ is NE)

if $val_i(\overline{\sigma_0}) = 1$ for some i

$\overline{\sigma_0}$ is strictmaxNE, as noted above

else

$\forall i\ val_i(\overline{\sigma_0}) = 1$, since $\forall i\ val_i(\overline{\sigma_0}) = 1$ or 0

Now, let V^ω denote the set of infinite plays.

if $\forall(\rho \in V^\omega)\ \forall i\ val_i(\overline{\sigma_0})(\rho) = 0$

There is no winning run for any player

$\Rightarrow \forall\overline{\sigma}\ \forall i\ val_i(\overline{\sigma}) = 0$

$\overline{\sigma}$ is strictmax NE.

else

$\exists\rho_0 \in V^\omega\ val_i(\overline{\sigma}) = 1$ for some i.

Hence, there is a run ρ where some player, say j, wins. But this run is not compatible with the strategy profile $\overline{(}\sigma)$.

if ∀ h ∈ V* $val_i(\overline{\sigma}) = 0\ \forall i$

if h $\leq \rho_0$

$\overline{\pi}(h) = \rho(len(h) + 1)$

else

$\overline{\pi}(h) = \overline{\sigma_0}(h)$

claim: $\overline{\pi} \in NE$

let $\sim(h \leq \rho_0)$

$\Rightarrow \forall x \sim$(h.x $\leq \rho_0$)

$\Rightarrow \forall x\ \overline{\pi}(h.x) = \overline{\sigma_0}(h.x)$

$\Rightarrow \overline{\pi}|_h(x) = \overline{\sigma_0}|_h(x)$

$\overline{\sigma_0}|_h$ is NE $\Rightarrow \overline{\pi}|_h$ is NE for $\sim(h \leq \rho_0)$

else $(h \leq \rho_0)$

then, $\forall x^+ \in$ extension(x),

$\sim (x^+ \leq \rho_0 \Rightarrow val_i(x^+) = 0$ as above $\wedge \overline{\pi}(x) = x^+ \Rightarrow X^+ \leq \rho_0$

$\Rightarrow \overline{\pi}(x) \geq \overline{\pi}(x)$

The only remaining case is when there indeed is some history h such that $val_i(\overline{\sigma_0}|_h) \neq 0$ for some i.

$\exists h\ (\overline{\sigma_0}|_h) \neq 0$ for some i.

let x $\in$ V* be such that $\forall v \leq x\ val_i(y) = 0$ for all i.

We will push this winning value up the history x.

We know $\overline{\sigma_0}$ is a subgame perfect equillibrium. Define:

$\overline{\pi} : \forall s \leq x$, let x = s.n.s$'$

$\overline{\pi}(s)$ = n

else

$\overline{\pi}(s) = \overline{\sigma_0}(s)$

then, $\overline{\pi}$ is a subgame perfect equilibrium.

$\forall h$, $\sim (h \leq x) \Rightarrow \overline{\pi}(s) = \overline{\sigma_0}(s)$

hence, $\overline{\pi}|_h$ is NE

$\forall h$ s.t. h = x.n.s,

let $\overline{\pi}|_{h.n} \in NE$, O(n)=j

since, $val_j(\overline{\sigma_0}|_h) = 0$

$\forall n' \in$ siblings(h) $\wedge\ n \neq n'$

$val_j(\overline{\pi}|_{h.n'} = val_j(\overline{\sigma}|_{h.n'}) = 0 \wedge \overline{\pi}|_{h.n'} \in NE$

Hence, $val_j(\overline{\pi}|_h) = val_j(\overline{\pi}|_{h.n}) \geq val_j(\overline{\pi}|_{h.n'})$

$\Rightarrow \overline{\pi}|_h$ is NE.

Hence, $\overline{\pi}$ is SPE with $val(\overline{\pi}) \neq 0$.

As a corollary,we get an alternate proof for the two player case. We have already seen that strictmax NE is not necessarily unique.

5.3.3 maxNE for n player games

A related definition is that of maxNE which relaxes the constraint that all the players get a better payoff to simply that no player gets a lower payoff. As it turns out, existence in the two player case is provable with a proof on lines of strictmax NE, but the n player proof breaks down.

Definition 15 $\overline{\sigma} \in max\ NE$ *iff*
$\sim \exists\ \overline{\pi} \in NE \circ (\forall i\ val_i(\overline{\pi}) \geq val_i(\overline{\sigma}) \wedge \exists i\ val_i(\overline{\pi}) > val_i(\overline{\sigma}))$

Lemma 10 *Every 2 player non zero sum game with Borel winning conditions has a max NE.*

Proof: The existence proof for 2 player strict max NE goes through.

figure 5.3.4

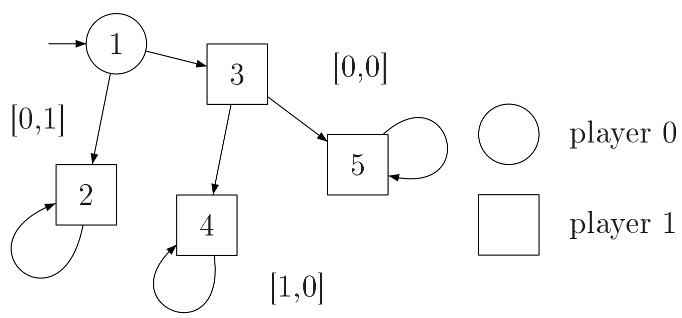

To the question whether max NE is unique, we again answer with a negative, albeit with a slightly more complex counter example of a Buchi game with indicated payoffs (fig 5.3.4):

maxNE$_1$: $\sigma_1(1) = 3$, $\sigma_2(3) = 4$, rest loops

maxNE$_2$: $\sigma_1(1) = 2$, $\sigma_2(3) = 5$, rest loops

Lemma 11 *There need not exist a unique max NE profile for 2 player non zero sum games, and as a corollary for n player games.*

The question of existence of maxNE profiles in n player games remains open. The earlier proof for strictmaxNE does not work because pushing the payoff of any one player to 1 is not enough to ensure that no other equilibria exist where some other player gets a better payoff. Nor could we directly extend the proof.

5.4 Process View

In order to understand the reason behind instability in equilibria we have to look at how a group of players actually reach an equilibrium profile. In the world of one shot complete information games, all players decide their strategies for all histories of the game before the game starts, and these strategies are public knowledge. Hence, a player decides his strategy having complete knowledge of what other players' strategies are, and in effect, knowing completely how the game will proceed (probabilistically, in case of mixed strategies) having decided his own strategy. Clearly, a strategy profile can be said to be stable if *no player would like to change to some other strategy given the strategy profile*, assuming that players are rational beings with selfish motives. We have seen that all the known equilibria fail the litmus test and hence are not stable.
But, how are strategies decided?
We look at a few models of strategy decisions. Since all players must know the strategies of all other players while deciding their own, we may consider a model where a complete nash equilibrium strategy profile is "announced" and the players choose to either except it or reject it. A profile is accepted if all the player accept it and rejected otherwise. We call a profile stable, if there is no other nash equilibrium which the players will accept replacing the current one. The pertinent question now is, when do players accept a new strategy profile?

- Case 1: Player i accepts the profile $\overline{\sigma}'$ over $\overline{\sigma}$ iff $val_i(\overline{\sigma}') > val_i(\overline{\sigma})$ i.e., the player are *lazy*, and will only accept a change in status quo when it increases their payoff. This is precisely the case of strictmax NE. We find that strictmax NE is the weakest kind of equilibrium that satisfies our meaning of stability. Other notions are extension of strictmax NE.

- Case 2: Player i accepts the profile $\overline{\sigma}'$ over $\overline{\sigma}$ iff $val_i(\overline{\sigma}') \geq val_i(\overline{\sigma})$ i.e., the player are *acquiescing*, and will accept any change in the profile as long as it does not decrease their payoff. This is the case of max NE. It can easily be seen that maxNE $\Rightarrow$ strictmaxNE.

The above scenario with parallel announcement of equilibrium strategies is a very artificial setting because, first, in a real world example there is no independent authority announcing strategies, and second, the announced profiles have to be nash equillibria, so the problem is assumed to be already half solved, before we begin. A more natural setting is the following: All players take turns in a circle. In his turn, a player may either choose to not change her strategy or announce a new strategy. The profile is said to be stable when all players have announced no change in one round. This is a much more natural setting and we can also describe nash equillibria in this setting.

Function characterization of process

If in the afore mentioned setting we assume that a player chooses to shift to a new strategy if it gives her a better payoff than the current profile, then we obtain a precise characterization of nash equilibrium.
For two players, consider that, a player does not change her strategy if she gets the best possible payoff with the current profile, or

$$\forall \overline{\sigma}' val_i(\overline{\sigma}) \geq val_i(\overline{\sigma}')$$

Now consider the function F : {(V*→D(V)) × (V*→ D(V)) × {0, 1}} ⟶ {(V*→D(V)) × (V* →D(V)) × {0, 1}}

$$F(\pi_0, \pi_1, i) = \left\{ \begin{array}{ll} \{\pi_0', \pi_1, 1-i\} \quad if & val_i(\pi_i', \pi_{1-i}) > val_i(\pi_i, \pi_{1-i}) \\ \{\pi_0, \pi_1, 1-i\} & otherwise \end{array} \right\}$$

This function captures the process of choosing a new strategy in response to the current profile by the players. Now, as mentioned, a stable profile in our process is obtained when two continuous applications of F lead back to the same profile. It can be shown that a profile is stable wrt this F iff it is a nash equilibrium. The proof is straightforward and we omit it here.

Let us now try to capture the "litmus test" through this function formalism. Define F as:

$$F(\pi_0,\pi_1, i) =$$

$$\{\pi_0', \pi_1, 1-i\} \text{ if}$$

$$\{val_i(\pi_i', \pi_{1-i}) > val_i(\pi_i, \pi_{1-i})\} \vee$$

$$\exists \pi_{1-i}' \left[\begin{array}{c} val_i(\pi_i', \pi_{1-i}') > val_i(\pi_i, \pi_{1-i}) \wedge \\ val_{1-i}(\pi_i', \pi_{1-i}') > val_{1-i}(\pi_i', \pi_{1-i}) \wedge \\ \forall \pi_i'' \{val_{1-i}(\pi_i', \pi_{1-i}'') \geq val_{1-i}(\pi_i', \pi_{1-i}') \rightarrow \\ val_i(\pi_i', \pi_{1-i}'') \geq val_i(\pi_i', \pi_{1-i}')\} \end{array} \right]$$

$$\{\pi_0, \pi_1, 1-i\} \text{ otherwise}$$

The "fix point" of F then as discussed above for nash equilibrium here captures the litmus test approximately. The last condition of safety has been put in to ensure that player (1-i) does play in a way that helps i. Since this function looks two steps into the future to decide on the best strategy, it can be called a ply-2 lookup. We can show that strictmaxNE and maxNE are subsumed by this formulation. In particular, we show how maxNE ensure ply-2 stability.

Lemma 12 *maxNE* ⟶ *ply-2 stability* and *strictmaxNE* ⟶ *ply-2 stability*

Proof:

let (π_0, π_1) be maxNE $\Rightarrow$

$$\forall \pi_i' \; val_i(\pi_i', \pi_{1-i}) \leq val_i(\pi_i, \pi_{1-i}) \cdots (1)$$

$$\wedge \; \forall \pi_{1-i}' \; val_i(\pi_i, \pi_{1-i}') \leq val_i(\pi_i, \pi_{1-i}) \cdots (2)$$

$$\wedge \; \forall \pi_i' \; \forall \pi_{1-i}'$$

$$[\; (val_i(\pi_i', \pi_{1-i}') \leq val_i(\pi_i, \pi_{1-i}) \vee val_{1-i}(\pi_i', \pi_{1-i}') < val_{1-i}(\pi_i, \pi_{1-i}))$$

$$\wedge \qquad \cdots (3)$$

$$(val_i(\pi_i', \pi_{1-i}') < val_i(\pi_i, \pi_{1-i}) \vee val_{1-i}(\pi_i', \pi_{1-i}') \leq val_{1-i}(\pi_i, \pi_{1-i}))]$$

From this follows:

$\forall \pi_i' \; \pi_{1-i}'$

$$val_i(\pi_i', \pi_{1-i}) \leq val_i(\pi_i, \pi_{1-i}) \text{ from } (1)$$

and

$$val_i(\pi_i', \pi_{1-i}') \leq val_i(\pi_i, \pi_{1-i}) \text{ from } (3)$$

which is enough to prove that F(π_0, π_1, 0) = {π_0, π_1, 1}. Similarly for F(π_0, π_1, 1).
A similar proof works for strictmaxNE.

Issues:
Note that we said above that we consider strictmaxNE to be the weakest equilibrium that captures stability, but ply-2 stability is implied by strictmaxNE. This is the case because ply-2 stability does not in fact capture stability completely. The safety condition explained above has been put in artificially and it is not clear if it is the best way to capture certainty that player 1-i does ends up helping player i while playing selfishly. But the preliminary results like capturing nash equilibrium and the clear emulation of the process of strategy determination motivate the study of these functions and their "fixed points" in detail. We do not delve deeper into these matters in this report.

5.5 collusion equilibrium

We end this report with an abstraction from the process view explained in the last section to introduce a new equilibrium concept that builds on the insights gained from elementary concepts like strictmaxNE and maxNE.

Definition 16 *A strategy profile $\overline{\sigma}$ in a game (V, O, Π, F) is a* ***collusion equilibrium*** *if no subset $\overline{\pi} \in \Pi$ can change their strategies in a way that increases the payoff of every $\pi \in \overline{\pi}$ i.e.,*

$$\forall \overline{\pi} \in \Pi \ \forall \overline{\sigma}' = (\sigma'_{\{i|\pi_i \in \overline{\pi}\}}, \sigma_{\{i, otherwise\}}) \\ \exists \{i|\pi_i \in \overline{\pi}\}\{val_i(\overline{\sigma}') \leq val_i(\overline{\sigma})\}$$

We note is collusion equilibrium $\Longrightarrow$ strictmaxNE

This definition is a process independent definition like nash equilibrium, although more complex. We hold the hope that results on existence of collusion equilibrium can be found parallel to those known for nash equilibrium. Another open problem is that of placing collusion equilibrium squarely in the process framework discussed in this report. We express our inability to treat them here for the lack of proper mathematical machinery.

Chapter 6

Conclusion & future directions

Following claims can be made at the end of this report:

- The undecidability result for pure-NE for SSMG was improved, the number of players required brought down to 5.
- Some advancement was made towards proving the existence of subgame perfect equilibria in stochastic games. We could prove the result for 2 player games, but the n player game remains open. We believe that the difficulty here lies in addressing the infinite tree zero sum stochastic game problem.
- It was felt the nash equilibrium and the other known equilibria do not capture the notion of rational stability of strategy profile adequately. This conviction was motivated through discriminating examples and an attempt was made towards trying to delve into the process of equilibrium reachability process, to gain insights into this foundational question. The problem still remains open.

Some important areas where we are left mid-stream by this work:

- Closing the undecidability gap for $n\frac{1}{2}$ games.
- Following up the process characterization in greater detail and trying to formulate the stability problem as a fixed point computaion.
- A complete treatment of collusion equilibrium.

Bibliography

[1] Krishnendu Chatterjee, Rupak Majumdar, and Marcin Jurdzinski. On nash equilibria in stochastic games. In *CSL*, pages 26 – 40, 2004.

[2] Erich Grädel and Michael 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.

[3] Chaterjee Krishnendu. *Stochastic Omega-regular games*. PhD thesis, University of California at Berkeley, 2007.

[4] Donald A. Martin. Borel determinacy. *Ann. Pure Appl. Logic*, 102:363 – 371, 1975.

[5] Michael Ummels. Rational behaviour and strategy construction in infinite multiplayer games. In *Proceedings of the 26th International Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2006*, volume 4337 of *LNCS*, pages 212 – 223. Springer, 2006.

[6] Michael Ummels and Dominik Wojtczak. The complexity of nash equilibria in simple stochastic multiplayer games. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming, ICALP 2009*, volume 5556 of *LNCS*, pages 297 – 308. Springer, 2009.

[7] Michael Ummels and Dominik Wojtczak. Decision problems for nash equilibria in stochastic games. In *Proceedings of the 18th Annual Conference of the European Association for Computer Science Logic, CSL '09*, volume 5771 of *LNCS*, pages 515 – 530. Springer, 2009.