

PhD Thesis from co-direction between
Université libre de Bruxelles
and **Université de Mons**



Equilibria in Multiplayer Games Played on Graphs

Aline Goeminne

A dissertation submitted in fulfilment of the requirements of
the degree of *Docteur en Sciences*

Jury

M. THOMAS BRIHAYE (Co-Directeur)

Université de Mons, Belgium

Mme. VÉRONIQUE BRUYÈRE (Présidente)

Université de Mons, Belgium

M. JÁNOS FLESCHE

Maastricht University, The Netherlands

M. GILLES GEERAERTS (Secrétaire)

Université libre de Bruxelles, Belgium

M. JEAN-FRANÇOIS RASKIN (Co-Directeur)

Université libre de Bruxelles, Belgium

M. OCAN SANKUR

Université de Rennes, CNRS, France

ACKNOWLEDGMENTS

First of all, I want to thank my supervisors, Thomas Brihaye and Jean-François Raskin, for their support and all their wise advice. Thomas, you introduced me to the world of research when I was still a master student. This idea for my master’s project led me here today, and I am very grateful for that. Jean-François, you accepted me in your team and you always provided a lot of ideas when we met. No doubt this was really beneficial for my research. I would also like to thank you both for giving me the independence I needed and for always letting me express my ideas.

I would like to thank each of the other members of my jury: Véronique Bruyère, János Flesch, Gilles Geeraerts and Ocan Sankur. I am happy that you accepted to review my work.

I will switch to French in order to thank my colleagues, my friends and my family with all the gratitude they deserve.

Avant toute chose, j’aimerais remercier l’Université de Mons de m’avoir octroyé une bourse partiellement financée qui m’a permis de réaliser ma thèse en co-tutelle avec l’Université libre de Bruxelles. Je remercie également Jean-François de m’avoir fait confiance et m’avoir fourni le financement nécessaire aux deux autres années de ma thèse, ce grâce au projet ARC “Non-Zero Sum Game Graphs: Applications to Reactive Synthesis and Beyond.”.

Je remercie mes autres co-auteurs: Véronique Bruyère, Nathan Thomasset

et Marie van den Bogaard. Véronique, tu m'as appris à écrire mes premiers papiers et prodigué de précieux conseils qui me seront, j'en suis sûre, bien utiles durant ma vie de chercheuse. Marie, je te remercie pour ces moments passés à faire chauffer nos neurones, pour tes accueils chaleureux lors de mes visites à l'ULB et pour les chocolats partagés lors des soumissions d'articles.

À mes collègues rencontrés à l'ULB, en particulier Marie, Nicolas, Ismaël, Léo et Damien, mais aussi à ceux que j'ai moins côtoyés. Les discussions lors des pauses thé ainsi que les fameux steaks-frites du mercredi me manqueront.

À tous ceux qui ont peuplé et embelli mes journées à l'UMONS, j'aimerais adresser quelques mots. Merci à Marion, Pierre C. et Quentin H. d'avoir été les premiers à m'accueillir parmi les doctorants. Merci également pour les nombreuses soirées jeux partagées et les explications de règles toujours très claires et passionnées de Pierre. Merci à mon coco-llègue Kévin, pour nos nombreuses années d'amitié et nos plus brèves années de partage de bureau. Merci pour les petites attentions qui rendaient les journées de travail plus douces. Merci à David et Quentin pour les soirées jeux de société en ligne et vos gentilles taquineries concernant mes oeuvres d'art. Merci mille fois pour vos encouragements et pour m'avoir écoutée me plaindre longuement. Merci à Céline pour les pauses café et les séances potins. Merci à Maximilien, Jérémy, Adrien et Pierre pour les séances de badminton. Merci à Pierre, Clément, James, Gaëtan et Horacio avec qui les pauses de midi ont rendu cette période moins solitaire. Merci pour votre bonne humeur ! Un merci particulier à Pierre et Clément et aux nombreux voyages effectués en votre compagnie. J'en garderai d'excellents souvenirs. À tous mes anciens professeurs et autres collègues avec qui je prenais plaisir à discuter au détour d'un couloir ou dans la cuisine.

Merci à Gilles et Mick pour leur écoute, leurs bons conseils et leur humour éclatant. Merci à toi, Thomas, de m'avoir guidée et permise, toujours avec bienveillance, à prendre confiance en mes ailes. Je serai toujours friande de nouvelles anecdotes relatives à Daniel Kahneman et des longues discussions qui en découlent.

Merci à Louise pour ses encouragements constants, son écoute et sa présence. J'ai hâte de pouvoir te retrouver autour d'un bon petit plat ! Merci à Laetitia pour ses délicieuses pâtisseries qui remontent le moral.

Enfin, mes derniers mots vont à ma famille, que mes longues journées de travail m'ont parfois poussée à négliger, je les prie de m'en excuser, et les remercie pour leur compréhension. À mes parents qui me supportent depuis toutes ces années, et ce, dans les deux sens du terme. À mon frère, ma soeur, ma belle-soeur et mon beau-frère, pour leur présence et tous ces beaux moments en famille hors des tracasseries de la thèse. À mes neveux et filleuls qui par un sourire, un éclat de rire, une parole n'ont de cesse de me rappeler la beauté du monde. À mes grands-parents pour qui tout ceci semble parfois bien mystérieux. Et en particulier, à toi papy, qui me demandes souvent ce qu'on peut encore chercher en mathématiques quand on a appris à résoudre une équation du second degré. À vous, qui avez toujours cru en moi, je vous remercie profondément.

À vous tous, qui par un mot, un sourire, une attention, un moment partagé, avez fait de ces années de thèse une aventure humaine et scientifique hautement enrichissante, encore une fois: merci !

ABSTRACT

Today, as *computer systems* are ubiquitous in our everyday life, there is no need to argue that their *correctness* is of capital importance. In order to prove (in a mathematical sense) that a given system *satisfies a given property*, *formal methods* have been introduced. They include concepts such as *model checking* and *synthesis*. Roughly speaking, when considering synthesis, we aim at building a model of the system which is correct by construction. In order to do so, models are mainly borrowed from *game theory*. During the last decades, there has been a shift from *two-player qualitative zero-sum games* (used to model antagonistic interactions between a system and its environment) to *multiplayer quantitative games* (used to model complex systems composed of several agents whose objectives are not necessarily antagonistic). In the latter setting, the solution concepts of interest include numerous *equilibria*, such as *Nash equilibrium* (NE) and *subgame perfect equilibrium* (SPE). While the existence of equilibria is widely studied, it is also well known that several equilibria may coexist in the same game. Nevertheless, some equilibria are more *relevant* than others. For example, if we consider a game in which each player aims at satisfying a given qualitative objective, it is possible to have both an equilibrium in which no player satisfies his objective and another one in which each player satisfies it. In this case one prefers the latter equilibrium which is more relevant.

In this thesis, we focus on *multiplayer turn-based games played on graphs* ei-

ther with qualitative or quantitative objectives. Our contributions are twofold: (i) we provide *equilibria characterizations* and (ii) we use these characterizations to solve *decision problems* related to the existence of relevant equilibria; and characterize their complexities.

Firstly, we provide a characterization of a weaker notion of SPE (weak SPE) in multiplayer games with ω -regular objectives based on the payoff profiles which are realizable by a weak SPE. We then adopt another point of view by characterizing the outcomes of equilibria instead of their payoff profiles. In particular we focus on weak SPE outcome characterization. As for some kinds of games (e.g., multiplayer quantitative Reachability games), weak SPEs and SPEs are equivalent, this characterization is useful in order to study SPEs in these games.

Secondly, we use those different equilibrium characterizations to provide the exact complexity classes of different decision problems related to the existence of relevant equilibria. We mainly focus on the *constrained existence problem*: if each player aims at maximizing his gain, this problem asks whether there exists an equilibrium such that each resulting player's gain is greater than a threshold (one per player). We also consider variants of relevant equilibria based on *the social welfare* and the *Pareto optimality* of the players' payoff. In this way, we prove the exact complexity classes for (i) the weak SPE constrained existence problem in multiplayer games with classical qualitative objectives such as Büchi, co-Büchi and Safety and (ii) the NE and SPE constrained existence problems (and variants) for qualitative and quantitative reachability games. In the latter case, the upper bounds on the required memory for such relevant equilibria are studied and proved to be finite. Studying memory requirements of strategies is important since with the synthesis process those strategies have to be implemented.

Finally, we consider multiplayer, non zero-sum, turn-based *timed games* with qualitative Reachability objectives together with the concept of SPE. We prove that the SPE constrained existence problem is EXPTIME-complete for qualitative Reachability timed games. In order to obtain an EXPTIME algorithm, we proceed in different steps. In the first step, we prove that the game variant of the classical region graph is a good abstraction for the SPE constrained existence problem. In fact, we identify conditions on bisimulations

under which the study of SPE in a given game can be reduced to the study of its quotient.

Les *systèmes informatiques* étant omniprésents dans notre vie quotidienne, il n'est pas nécessaire d'argumenter que leur *exactitude* est d'une importance cruciale. C'est pourquoi *les méthodes formelles* ont été introduites afin de prouver (dans le sens mathématique du terme) qu'un système donné *satisfait une certaine propriété*. Elles incluent des concepts tels que la *vérification de modèles* et la *synthèse*. Dans le cadre de la synthèse, le but est de construire un système qui est correct par construction. Les modèles utilisés à cet effet sont largement empruntés à la *théorie des jeux*. Durant les dernières décennies, une transition s'est effectuée entre les *jeux qualitatifs à deux joueurs et à somme-nulle* (utilisés pour modéliser les interactions antagonistes entre un système et son environnement) et les *jeux quantitatifs multijoueurs* (utilisés pour modéliser des systèmes complexes composés d'agents dont les objectifs ne sont pas nécessairement antagonistes). Dans ce dernier type de jeux, les concepts de solutions incluent de nombreux concepts d'équilibres tels que les équilibres de Nash (NEs) et les équilibres parfaits en sous-jeux (SPEs). Au-delà du fait que l'existence d'équilibres est déjà largement étudiée, il est aussi bien connu que plusieurs équilibres peuvent coexister dans le même jeu. Néanmoins, certains équilibres sont plus *pertinents* que d'autres. Par exemple, si on considère un jeu où tous les joueurs veulent satisfaire un certain objectif qualitatif, il est possible d'avoir à la fois un équilibre pour lequel aucun des joueurs ne satisfait son objectif et un autre équilibre pour lequel tous les joueurs le satisfont. Dans

ce cas, le dernier équilibre est préféré au premier.

Dans cette thèse, nous nous focalisons sur les *jeux multijoueurs joués sur des graphes* dans lesquels les joueurs jouent tour à tour et sont munis aussi bien d'objectifs qualitatifs que d'objectifs quantitatifs. Nos contributions sont doubles : (i) nous fournissons des caractérisations d'équilibres et (ii) nous utilisons ces caractérisations dans le but de résoudre des problèmes de décision relatifs à l'existence d'équilibres pertinents; et nous caractérisons leur complexité.

Premièrement, bien que nous donnons une caractérisation des weak SPEs basée sur les profils de gains réalisables par de tels équilibres et ce dans les jeux munis d'objectifs ω -réguliers, nous nous focalisons principalement sur une autre approche qui vise à caractériser l'ensemble des parties résultant de weak SPEs. Puisque dans certains types de jeux (*ex*: les jeux avec objectifs d'accessibilité) les notions de weak SPEs et de SPEs sont équivalentes, cette caractérisation est utile pour étudier les SPEs dans ces jeux particuliers.

Deuxièmement, nous utilisons ces différentes caractérisations d'équilibres afin de fournir les classes de complexité exactes de différents problèmes de décision relatifs à l'existence d'équilibres pertinents. Nous étudions principalement le *problème d'existence sous contraintes* : si chaque joueur a pour but de maximiser son gain, ce problème demande s'il existe un équilibre tel que le gain de chaque joueur, s'il se conforme à cet équilibre, soit plus grand qu'un certain seuil. Nous considérons aussi des variantes d'équilibres pertinents basés sur le *bien-être social* et la *Pareto optimalité* des profils de gains.

Finalement, nous étudions comment nos résultats peuvent être utilisés dans le cadre des jeux multijoueurs *temporisés* dans lesquels les joueurs jouent tour à tour et sont munis d'objectifs d'accessibilité qualitatifs.

1	Introduction	19
1.1	Context	19
1.2	Contributions	22
1.3	Related Works	26
1.4	Outline	28
 I	 Games Played on Graphs	 29
 2	 Background	 31
2.1	Arena and strategies	32
2.2	Objectives	38
2.2.1	Qualitative objectives	39
2.2.2	Quantitative objectives	43
2.2.3	Continuous and prefix-linear objectives	46
2.2.4	Play of a given gain profile in Boolean games	52
2.3	Zero-sum games	57
2.3.1	Qualitative zero-sum games	57
2.3.2	Quantitative zero-sum games	60
2.4	Multiplayer games	64
2.4.1	Nash equilibrium	64

2.4.2	Subgame perfect equilibrium	68
2.4.3	Weak subgame perfect equilibrium	71
II	Characterization of Equilibria Outcomes	77
3	Introduction	79
4	Coalitional Games and Extended Reachability Games	83
4.1	Coalitional games	83
4.2	Extended games	85
5	λ-consistent Plays	93
5.1	Labeling Function and λ -Consistency	93
6	NE Outcome Characterization	97
6.1	General Characterization	97
6.2	Characterization for Reachability Games	104
7	Weak SPE and SPE Outcome Characterizations	109
7.1	λ^k labeling functions for weak SPEs	110
7.1.1	(Good) local strongly prefix-linear constants set	110
7.1.2	Definition of the λ^k -labeling functions	112
7.1.3	Weak SPE outcome characterization	117
7.2	(Good) Symbolic witnesses for weak SPEs	119
7.2.1	(Good) Symbolic witnesses	120
7.2.2	Weak SPE outcome characterization	121
7.3	Proofs of the weak SPE outcome characterizations	123
7.4	Instantiations of the weak SPE outcome characterizations	130
7.4.1	Boolean games with prefix-independent gain functions	130
7.4.2	Safety and qualitative Reachability games	134
7.4.3	Quantitative and weighted Reachability games	136
7.4.4	λ -consistency vs Visit λ -consistency in Reachability games	142
7.5	SPE outcome characterization	143

III Relevant Equilibria	147
8 Introduction	149
9 Boolean Games with Prefix-Independent Gain Functions	153
9.1 Computation of λ^*	154
9.1.1 Reformulation of λ^k -Consistency	154
9.1.2 Algorithms	155
9.2 NP-Completeness for Multiplayer co-Büchi, Parity, Muller, Rabin, and Streett Games	160
9.3 P-Completeness for Büchi Games	163
9.3.1 Deciding the constrained existence problem of NEs	164
9.3.2 P-completeness of the constrained existence problem of weak SPEs in Büchi games	165
9.4 P-Completeness for Explicit Muller Games	169
10 Qualitative Reachability and Safety Games	171
10.1 PSPACE-membership	172
10.2 PSPACE-hardness	175
10.3 PSPACE-completeness of the constrained existence problem of SPEs	181
10.4 Memory requirements	181
11 Quantitative Reachability Games	185
11.1 Regions decomposition and regions based algorithm	187
11.1.1 Regions decomposition	187
11.1.2 Regions Based Algorithm	189
11.2 Counter graph	195
11.3 PSPACE-completeness	201
11.3.1 PSPACE-membership	201
11.3.2 PSPACE-hardness	209
12 Other Relevant Equilibria in Reachability Games	215
12.1 Studied problems	217
12.2 Existence problems	220
12.3 Solving decision problems	221

12.3.1 Sufficiency of lassoos	223
12.3.2 Algorithms and memory requirements	225
12.4 The qualitative setting	228
12.4.1 Qualitative Reachability games	229
12.4.2 Decision problems and complexity results	229
12.4.3 Existence problem	231
IV Timed Games	235
13 Introduction	237
14 Preliminaries	241
15 SPEs in a Game and its Quotient	247
15.1 Game bisimulation	247
15.2 Quotient game	248
15.3 SPE existence	249
16 Qualitative Reachability Games	251
16.1 Reachability games and quotient reachability games	251
16.2 Complexity results	252
16.2.1 Extended game	253
16.2.2 Complexity	255
17 Application to Timed Games	259
17.1 Timed automata and timed games	259
17.2 Regions and region games	261
17.3 Complexity results	263
17.4 Discussion and future works	265
V Discussion and Future Works	267
Bibliography	273

Appendices	285
Appendix A Characterization of The Gain Profiles of weak SPEs in Multiplayer Games with Prefix-Independent Qualitative Ob- jectives	287
A.1 Remove-Adjust procedure	288
A.2 Characterization and good symbolic witness	291
Appendix B Technical Proofs	299
B.1 Proofs of Chapter 11	299
B.1.1 Proofs of Section 11.1.2	299
B.1.2 Proofs of Section 11.2	305
B.2 Proofs of Chapter 12	324
B.2.1 Proof of Section 12.2	324
B.2.2 Proofs of Section 12.3	327
B.3 Proofs of Chapter 15	331
B.3.1 Proof of Proposition 15.3.3	331
B.3.2 Proof of Theorem 15.3.1	340

1.1 Context

Computer Systems and Formal Methods Nowadays *computer systems* are more and more involved in our everyday life. These systems become increasingly complex and interact either together or with humans. Moreover, some of these reactive systems are used by humans for critical tasks such as medicine, transport, etc. For this kind of *critical systems*, bugs may have dramatic consequences. For this reason it is crucial to ensure that a system is *correct* and *satisfies some properties*. One way to do so is program testing but as E. W. Dijkstra said: “Program testing can be used to show the presence of bugs, but never to show their absence”. *Formal methods* through concepts such as *model checking* and *synthesis* provide techniques to *mathematically prove* that a system is correct.

Model checking, proved relevant in companies such as IBM or Intel, allows one to systematically check if some properties hold in the system. In order to do so, we have to give formal descriptions of the system thanks to a mathematical model (for instance a transition system) and of the properties that the system has to satisfy (thanks to Linear Temporal Logic (LTL) formulae for example). Model checking intends to provide efficient algorithms to answer to the question “Does the system satisfy the specifications?”. The book authored by Baier and

Katoen [BK08] provides a comprehensive introduction to this broad subject. Another point of view is that of *synthesis* (see *e.g.*, [PR89]); with synthesis one wants to build a model of the system which is correct by construction: *game theory* is well-suited for this purpose.

Game Theory and Rationality *Game theory* is a branch of mathematics which allows one to model interactions between individuals. This field has been applied in many disciplines such as economics, biology or computer science. Von Neumann and Morgenstern are known as the fathers of game theory. They published in 1944 the book “Theory of Games and Economic Behavior” [vNM44]. Their work is based on an axiomatization of the *rational behavior* of a decision maker. Such a decision maker should maximize the expected value of his utility function if he is rational.

Two-player Games Played on Graphs *Two-player zero-sum games played on graphs* are commonly used to model *reactive systems* where a system interacts with its environment. In such setting the system wants to achieve a goal—to respect a certain property—and the environment acts in an antagonistic way. The underlying game is described as follows: the two players are the system and the environment, the vertices of the graph are all possible configurations in which the system can be and an infinite path in this graph, called a *play*, depicts a possible sequence of interactions between the system and its environment. In such a game, each player chooses a *strategy*: it is the way he plays given some information about the game and past actions of the other player. Following a strategy for each player results in a particular play in the game called the *outcome*. Finding how the system can ensure that a given property is satisfied amounts to finding a *winning strategy* for the system in this game. A winning strategy for the system is a strategy which ensures that the system achieves his goal whatever the strategy of the environment.

From Two-player Games to Multiplayer Games This previous model is not realistic since it assumes that the system is fully antagonistic and composed of two agents (the system and the environment). However the environment may have its own objective and the system may be composed of many components

which all have their own specifications. Thus, the model evolved from two-player zero-sum games to multiplayer games where all players have their own objectives which are not necessarily antagonistic.

From Qualitative to Quantitative Objectives Additionally, even if ω -regular objectives are widely studied, they only offer to express the satisfaction of a specification or not (qualitative objectives). Thus in some cases, for example if we want to measure the amount of energy used in a system, one needs to rely on quantitative specifications (quantitative objectives). These games are called multiplayer quantitative games. In this setting, all players are assumed to be rational in the sense that they all want to maximize their gains (or minimize their costs).

More details about these settings may be found in the following, non-exhaustive, references: *e.g.*, [Tho95, GTW02, AG11, ZP96, Bru17].

Equilibria In the setting of multiplayer quantitative game, the solution concept of winning strategy is not well-suited anymore. Equilibria are widely studied in multiplayer games: Nash equilibrium (NE) [Nas50], subgame perfect equilibrium (SPE) [Sel65], secure equilibrium (SE) [CHJ04, BBDPG13], ... Some of them are directly provided by classical game theory. Roughly speaking, an equilibrium is a contract between the players such that each player has no incentive to deviate from this contract if he assumes that the other players will follow it. The studied equilibrium depends on the kind of games which are of interest. For instance, SPEs take more accurately into account the sequential aspect of games played on graphs.

Relevant equilibria While the existence of equilibria is widely studied, it is also well known that several equilibria may coexist in the same game (*e.g.*, [Umm08]). Nevertheless, some equilibria are more *relevant* than others. For example, if we consider a game in which each player aims at satisfying a given qualitative objective, it is possible to have both an equilibrium in which no player satisfies his objective and another one in which each player satisfies it. In this case one prefers the latter equilibrium which is more relevant.

1.2 Contributions

Before providing our main contributions, we explain the philosophy of this thesis.

This thesis is mainly based on different joint works with Thomas Brihaye, Véronique Bruyère, Jean-François Raskin, Nathan Thomasset and Marie van den Bogaard [BBGR18, BBGR20, BBG⁺19, BBG⁺20, BBGT19, BG20], and on [Goe20].

In the previously cited papers, we have considered multiplayer games equipped with specific objective functions (including Boolean games with prefix-independent objective functions, qualitative Reachability games, Safety games and quantitative Reachability games). Most of these papers follow the same scheme: we provide an ad hoc characterization of equilibrium depending on the nature of the equilibrium considered, and the objective function; and then we exploit this characterization in order to obtain algorithms related to these equilibria. In this document, in order to obtain a uniform presentation of those characterizations, we decided to provide general characterizations. In this way the characterizations originally provided in our papers can be seen as instantiations of those general characterizations.

Let us now provide our main contributions, more details are given at the beginning of each part or chapter.

In this thesis, we consider multiplayer turn-based games played on graphs equipped with qualitative or quantitative objectives.

Characterizations of equilibria outcomes A natural question which can arise when studying equilibria is: “Does there exist a Nash equilibrium in this game such that, if each player complies with the equilibrium, each player satisfies his objective?”. In order to answer the latter question, it may be sufficient to check the existence of a play which is proved to be the outcome of a Nash equilibrium and such that each player satisfies his objective along it. Therefore, the following question arise “For a given kind of equilibrium, is it possible to exactly characterize the set of equilibrium outcomes in a game?”.

Equilibrium characterizations has already been investigated in the literature, for instance in [BBMR15, FKM⁺10, FP17, BMR14, Bru17]. In the thesis, we continue this line of research by providing characterizations based on the notion of λ -consistent play. A labeling function λ is a function that assigns a value to each vertex of the game graph. These values impose constraints on the play in such a way that, given a well-defined labeling function λ , a play which satisfies the constraints given by λ is the outcome of an equilibrium in the studied game. We formally introduce this notion in Part II.

In Chapter 6, we show how, from the notion of values in zero-sum games, we are able to obtain a labeling function Val^* which characterizes the set of outcomes of Nash equilibria in certain kinds of games with certain conditions on the players' cost functions. While this characterization was already known (*e.g.*, [Bru17]), it was not explained, up to our knowledge, how to adapt it to Reachability games. We provide such a characterization in Section 6.2. Moreover, if a play is a *lasso* and satisfies the criterion given by the characterization, we prove that there exists a finite-memory Nash equilibrium which has this play as outcome.

In Chapter 7, we provide a general approach to define a labeling function λ^* which allows, under some conditions on the game, to characterize the set of outcomes of weak SPEs in a given game. This labeling function λ^* is obtained from an iterative procedure that we assume reaches a fixpoint. We also introduce the notion of (*good*) *symbolic witness*. Roughly speaking, a good symbolic witness is a finite set of plays of the game which satisfy some “good properties” in such a way that if we have a good symbolic witness, we are able to build a weak SPE. Moreover, if each play in the good symbolic witness is a *lasso*, we provide an upper bound on the needed memory size of the weak SPE built from this good symbolic witness. Then we explain how from these characterizations, we are able to characterize the set of outcomes of weak SPEs in Boolean games with prefix-independent gain functions, in qualitative and quantitative Reachability games and in Safety games. Finally, since the notions of weak SPE and SPE are equivalent in qualitative and quantitative Reachability games, we obtain from these latter characterizations, an SPE outcome characterization for qualitative and quantitative Reachability games.

Our characterizations have allowed us to obtain several complexity results

of different decision problems related to the existence of relevant equilibria. We detail those results in the following paragraphs.

Deciding the existence of relevant equilibria As already mentioned, different equilibria may coexist in the same game. It is the reason why in Part III we study different decision problems related to the existence of relevant equilibria in a game.

In our works we have mainly focused on the *constrained existence problem* (CEP) but in the particular case of Reachability games we have chosen to go further by considering the *upper threshold decision problem*, the *social welfare decision problem* and the *Pareto optimal decision problem*. The constrained existence problem is the following: given a game, an upper threshold and a lower threshold for each player, decide if there exists an equilibrium such that the cost of each player lies between the lower and the upper thresholds. The upper threshold decision problem is the same as the constrained existence problem but without lower thresholds. The main idea behind the social welfare decision problem is to maximize the number of players that reach their target set and to minimize the sum of their costs. With the Pareto optimal decision problem our aim is to decide the existence of an equilibrium with a cost profile that is Pareto optimal in the set of all possible cost profiles in the studied game.

In Chapter 9, we consider Boolean games with prefix-independent objective functions, we begin by providing a naive algorithm that decides the constrained existence problem of weak SPEs in these games. Then we prove that the CEP is (i) NP-complete for games with co-Büchi, Parity, Muller, Streett and Rabin objectives and (ii) P-complete for games with Explicit Muller and Büchi objectives. In Chapter 10, we consider qualitative Reachability games and Safety games. In these settings, we prove that the CEP of weak SPEs and SPEs is PSPACE-complete. In Chapter 11, we prove that the CEP of weak SPEs and SPEs in quantitative Reachability games is PSPACE-complete. Finally, in Chapter 12, we mainly focus on quantitative Reachability and we prove that: (i) for NEs: the upper threshold decision problem and the social welfare decision problem are NP-complete, while the Pareto optimal decision problem is NP-hard and belongs to Σ_2^P ; (ii) for SPEs: the upper threshold, social

welfare and Pareto optimal decision problems are PSPACE-complete. We also briefly mention the qualitative setting and provide upper bounds on the needed size memory of equilibria. All these results are summarized in Tables 1.1-1.4.

Table 1.1: Complexity classes of the CEP of weak SPEs for classical prefix-independent qualitative objectives

weak SPE	Expl. Muller	Büchi	Co-Büchi	Parity	Muller	Rabin	Streett
P-complete	×	×					
NP-complete			×	×	×	×	×

Table 1.2: Complexity class of the CEP of weak SPEs and SPEs in qualitative Reachability, Safety and quantitative Reachability games

	qual. Reach		Safety		quant. Reach	
	weak SPE	SPE	weak SPE	SPE	weak SPE	SPE
PSPACE-c	×	×	×	×	×	×

Table 1.3: Complexity classes for upper threshold, social welfare and Pareto optimal decision problems.

Complexity	Qualitative Reach.		Quantitative Reach.	
	NE	SPE	NE	SPE
Upper threshold	NP-c [CFGR16, Umm05]	PSPACE-c	NP-c	PSPACE-c
Social welfare	NP-c	PSPACE-c	NP-c	PSPACE-c
Pareto opti.	NP-h/ Σ_2^P	PSPACE-c	NP-h/ Σ_2^P	PSPACE-c

Table 1.4: Memory results for upper threshold, social welfare and Pareto optimal decision problems.

Memory	Qualitative Reach.		Quantitative Reach.	
	NE	SPE	NE	SPE
Upper threshold	Polynomial [CFGR16]	Exponential	Polynomial	Exponential
Social welfare	Polynomial	Exponential	Polynomial	Exponential
Pareto opti.	Polynomial	Exponential	Polynomial	Exponential

Application to multiplayer timed games In Part IV, we consider multiplayer turn-based timed games with qualitative Reachability objectives to-

gether with the concept of SPE. We focus on the constrained existence problem of SPEs: given a timed game, we want to decide whether there exists an SPE where some players have to win and some other ones have to lose. The main result of this part is that the constrained existence problem of SPEs is EXPTIME-complete for qualitative Reachability timed games.

1.3 Related Works

We present in this section some related works. This list is obviously not exhaustive.

(Very) Weak SPEs The restricted class of deviating strategies used in very weak SPEs is a well-known notion that for instance appears with the one-step deviation property in the Folk theorem for infinitely repeated game [OR94]. Weak SPEs and very weak SPEs are equivalent notions, but there are games for which there exists a weak SPE but no SPE [BBMR15]. Nevertheless, (very) weak SPEs and SPEs are equivalent for quantitative and qualitative Reachability games, an important property used in the proofs of [BBMR15] and of this document. The equivalence between SPEs and very weak SPEs is also implicitly used as a proof technique in a continuous setting in [FL83] and in a lower-semicontinuous setting in [FKM⁺10].

In [BRPR17], general conditions are given that guarantee the existence of a weak SPE. It follows that there always exists a weak SPE for games where players use a prefix-independent or finite range cost function.

Characterizations of equilibria and fixpoint techniques In [Bre12], Brenguier provides an outcome characterization of Nash equilibria in concurrent games, while in [BMR14] a characterization of outcomes of secure equilibria in the restricted case of two-player games can be found. Notice that in this latter paper, they also consider the constrained existence problem of secure equilibria and prove it is decidable for games equipped with some kinds of objective functions. Recently, authors in [BRvdB21] provide a characterization of all the subgame perfect equilibria in games with mean-payoff objectives.

Fixpoint techniques are used in several papers to establish the existence

of (weak) SPEs in some classes of games like [FKM⁺10, BBMR15, BRPR17]. However they cannot be used in our context to get the PSPACE complexity result of the constrained existence problem of (weak) SPEs in quantitative Reachability games.

Existence of relevant equilibria Regarding the constrained existence problem and the upper threshold decision problem, for NEs, it is shown to be NP-complete in qualitative Reachability games in [CFGR16] and in weighted concurrent Reachability games [KLST12], notice that null weights on the edges are allowed in [KLST12]. Notice that in [Umm08], variants of these problems for games with Streett, Parity or co-Büchi winning conditions are shown NP-complete and decidable in polynomial time for Büchi conditions. In [GU08, Umm06], a tree automata-based algorithm is given to decide the constrained existence problem of SPEs in graph games with ω -regular objectives defined by Parity conditions. A complexity gap is left open: this algorithm executes in EXPTIME and NP-hardness of the decision problem is proved. Quantitative Reachability objectives are investigated in [BBMR15] where it is proved that the constrained existence problem for weak SPEs and SPEs is decidable, but its exact complexity is left open.

Regarding the social welfare decision problem, in the setting of games played on matrices, deciding the existence of an NE such that the expected social welfare is at most k is NP-hard [CS03]. Moreover, in [BMS14] it is shown that deciding the existence of an NE which maximizes the social welfare is undecidable in concurrent games in which a cost profile is associated only with terminal nodes.

Regarding the Pareto optimal problem, in the setting of zero-sum two-player multidimensional mean-payoff games, the *Pareto-curve* (the set of maximal thresholds that a player can force) is studied in [BR15] by giving some properties on the geometry of this set. The authors provide a Σ_2^P algorithm to decide if this set intersects a convex set defined by linear inequations.

Regarding the memory, it is shown in [BDS13] that there always exists an NE with polynomial memory in quantitative Reachability games, without any constraint on the cost of the NE. It is shown in [Umm06] that, in multiplayer games with ω -regular objectives, there exists an SPE with a given gain profile

if and only if there exists an SPE with the same gain profile but with finite memory.

Notions of rationality Other notions of rationality and their use for reactive synthesis have been studied in the literature: rational synthesis in cooperative [FKL10] and adversarial [KPV14] setting, and their algorithmic complexity has been studied in [CFGR16]. Extensions with imperfect information have been investigated in [FGR18]. Synthesis rules based on the notion of admissible strategies have been studied in [Ber07, BRS14, BRS15, BPRS16, BJP⁺18].

Notice that the related works for timed games are provided in Part IV.

1.4 Outline

This document is mainly divided into six parts.

Part I gives a brief overview of necessary definitions and some results about games played on graphs.

Part II provides the characterizations of outcomes of equilibria.

Part III provides the complexity classes of the problems we have considered in order to decide the existence of a relevant equilibrium in a game.

Part IV considers the model of multiplayer turn-based timed games with qualitative Reachability objectives. Since this model is different from the model studied in the rest of this document, we aim to keep this part as self-contained as possible to allow its reading independently from the others.

Part V briefly concludes this thesis by providing some possible future works.

Appendices The purpose of these appendices is twofold: *(i)* in Appendix A, we provide the characterization of Boolean games with prefix-independent objective functions, based on gain profiles of weak SPEs, that we have presented in [BBGR18]; and *(ii)* in order to ease the reading of Part III and Part IV, we have chosen to relegate some of the (technical) proofs in Appendix B.

Part I:

GAMES PLAYED ON GRAPHS

CHAPTER 2

BACKGROUND

This chapter is devoted to introducing the general background we use throughout this document.

We are interested in *games played on graphs*. Such a game is equipped with an *arena*: a finite set of *players*, a (finite) set of *vertices* partitioned between the players and a set of *edges*. A *play* in such a game is as follows: a token is placed on a vertex, the player who owns this vertex moves the token through an outgoing edge. The token is in a new vertex and this procedure is repeated infinitely often leading to an infinite path in the game graph.

Additionally, each player has an *objective* that he wants to achieve. One may consider *qualitative objectives* or *quantitative objectives*. With a qualitative objective either a player achieves his objective or not. For example, if a player wants to reach a given vertex in the game graph, either he reaches it or not, no matter the number of vertices he has to visit before reaching it. With a quantitative objective, a player has a cost (resp. gain) function that he wants to minimize (resp. maximize). In the context of the previous example, the player does not only want to reach the given vertex, he wants to do it as soon as possible.

The behavior of the players may be adversarial or not. In the setting of *two-player zero-sum games*, the objectives of the two players are antagonistic while in the setting of *multiplayer games* the objectives of the players (potentially

more than two players) are not necessarily antagonistic. Once a game is fixed we aim at studying the players’ “rational behavior”, this might be done through the principle of *solution concepts*. Depending on the kind of studied game, different solutions may be considered, for example *winning strategies*, *optimal strategies*, *Nash equilibria*, and so on.

We formalize all these notions in the next sections. The notations and definitions roughly follow those of [BBGR18, BBGT19, BBG⁺19, Goe20] some of them are inspired by those of [De 13].

2.1 Arena and strategies

In this section we introduce the main definitions and notations related to the notion of *arena*, *plays*, *strategies* and *game played on a graph*.

Definition 2.1.1 (Arena). An (finite) *arena* A is a tuple $A = (\Pi, V, E, (V_i)_{i \in \Pi})$ where

- Π is a finite set of players;
- $G = (V, E)$ is a (finite) *directed graph* with V the set of vertices and $E \subseteq V \times V$ the set of edges. Moreover, for each $v \in V$ there exists $v' \in V$ such that $(v, v') \in E$ (*i.e.*, each vertex has at least one outgoing edge);
- $(V_i)_{i \in \Pi}$ is a partition of V between the players. If $v \in V_i$ for some $i \in \Pi$ it means that the vertex v is *owned/controlled* by Player i .

Since the set of players is finite we can number them and assume that $\Pi = \{1, \dots, n\}$, for some $n \in \mathbb{N}$. Thus we often call them Player 1, Player 2, ... Moreover, the set of vertices is partitioned between the players and they play alternately. This kind of game is called a *turn-based game*.

Example 2.1.2. In Figure 2.1 we consider the arena $A = (\Pi, V, E, (V_i)_{i \in \Pi})$ where:

- $\Pi = \{1, 2\}$;

- $V = \{v_0, v_1, v_2, v_3, v_4\}$;
- $E = \{(v_0, v_1), (v_0, v_2), (v_2, v_3), (v_2, v_4), (v_3, v_0), (v_4, v_0), (v_1, v_0)\}$;
- $V_1 = \{v_1, v_2, v_3, v_4\}$ (the rounded vertices) and $V_2 = \{v_0\}$ (the rectangular vertices).

When we will consider a game with two players, we will keep the convention that the vertices owned by Player 1 (resp. Player 2) are the rounded vertices (resp. rectangular vertices).




Figure 2.1: An arena A where Player 1 (resp. Player 2) owns rounded vertices (resp. rectangular vertices).

Remark 2.1.3. Notice that in Definition 2.1.1 we assume that the *game graph* is finite, this will be the case throughout this document except in Part IV where the number of vertices may be infinite.

Moreover, we assume, except in Part IV, that the graph game is not a *directed multigraph*. In a directed multigraph, if E is the set of edges, then the function $e : E \rightarrow V \times V$ which assigns to each edge its source and its target vertices is not necessarily injective.

We will explicitly exhibit these differences when it is needed.

Remark 2.1.4 (Weighted arena). When we want to express quantitative specifications (see Section 2.2.2), it is sometimes useful to equip the arena A with a *weight function* $w_i : E \rightarrow \mathbb{Z}$ for each player $i \in \Pi$. In this case, we write $A = (\Pi, V, E, (V_i)_{i \in \Pi}, (w_i)_{i \in \Pi})$ and we call A a *weighted arena*.

Plays and histories When a token is moved by the players along the edges of the arena $A = (\Pi, V, E, (V_i)_{i \in \Pi})$, it leads to an infinite sequence of vertices called a *play*. Formally, a play $\rho \in V^\omega$ is such that if $\rho = \rho_0 \rho_1 \dots$, then for all $k \in \mathbb{N}$ $(\rho_k, \rho_{k+1}) \in E$. A *history* $h \in V^*$ is either the empty sequence ϵ or a

finite sequence $h = h_0 h_1 \dots h_k$ for some $k \in \mathbb{N}$ such that for each $0 \leq n \leq k-1$, $(h_n, h_{n+1}) \in E$. The length $|h|$ of $h \neq \epsilon$ is the number k of its edges. We depict the set of plays (resp. histories) in A by Plays (resp. Hist). When it is not clear from the context which arena is considered, we highlight it by Plays_A (resp. Hist_A). If $h = h_0 \dots h_k$ for some $k \in \mathbb{N}$ is a history, we write $\text{First}(h)$ ¹ to denote the first vertex h_0 of h , in the same way $\text{Last}(h)$ denotes the last vertex h_k of h . The set Hist_i depicts the set of histories which end with a vertex owned by Player i , *i.e.*, $\text{Hist}_i = \{h \in \text{Hist} \setminus \{\epsilon\} \mid \text{Last}(h) \in V_i\}$. Moreover, if $h = h_0 \dots h_k$, for some $k \in \mathbb{N}$, is an history and $v \in V$ is a vertex such that $(h_k, v) \in E$, we denote by hv the history $h_0 \dots h_k v$. Finally, given a play $\rho \in \text{Plays}$, if there exist $h, \ell \in \text{Hist}$ such that $h\ell \in \text{Hist}$ and $\rho = h\ell^\omega$, we say that ρ is a *lasso*. Notice that ℓ is not necessarily a simple cycle. Moreover, in the rest of this document, we sometimes refer to the *length* L of a lasso $\rho \in \text{Plays}(v)$ or to the fact that this length is bounded by a value L . By that we mean that there exist $h \in \text{Hist}(v)$ and $\ell \in \text{Hist}$ such that (i) $h\ell \in \text{Hist}(v)$; (ii) $\rho = h\ell^\omega$ and (iii) $|h\ell| = L$ (resp. $\leq L$). In particular, it means that the lasso ρ has a finite representation of size L (resp. at most L).

We denote by $\text{Succ}(v) = \{v' \mid (v, v') \in E\}$ the set of successors of v , for $v \in V$, and by Succ^* the transitive closure of Succ . If $\rho = \rho_0 \rho_1 \dots$ is a play in A , then for each $k \in \mathbb{N}$, $\rho_{\leq k}$ denotes the history $\rho_0 \dots \rho_k$ and $\rho_{\geq k}$ denotes the play $\rho_k \rho_{k+1} \dots$. A *prefix* (resp. *proper prefix*) h' of a history $h = h_0 \dots h_k$, with $k \in \mathbb{N}$ and $h \neq \epsilon$, is a history $h_0 \dots h_\ell$, with $\ell \leq k$ (resp. $\ell < k$), denoted by $h' \leq h$ (resp. $h' < h$). In the same way, a prefix h of a play ρ is denoted by $h < \rho$.

Let $S \subseteq V$ be a subset of vertices of the arena A and let $\rho \in \text{Plays}$ be a play in A , we say that the play $\rho = \rho_0 \rho_1 \dots$ *reaches/visits* S if there exists $k \in \mathbb{N}$ such that $\rho_k \in S$. If $S = \{v\}$ for some $v \in V$, we say that ρ *reaches/visits* v (instead of $\{v\}$) if there exists $k \in \mathbb{N}$ such that $\rho_k = v$. Moreover, we denote by $\text{Occ}(\rho)$ the set of vertices visited along ρ , *i.e.*, $\text{Occ}(\rho) = \{v \in V \mid \exists k \in \mathbb{N}, \rho_k = v\}$. These notions may be easily adapted for histories. We also define the set $\text{Inf}(\rho)$ as the set of vertices which are visited infinitely often, that is $\text{Inf}(\rho) = \{v \in V \mid \forall k \in \mathbb{N}, \exists n \geq k, \rho_n = v\}$.

¹We can use $\text{First}(\rho)$ for all plays $\rho \in \text{Plays}$ in the same way as for histories.

Strategies A *strategy* of Player i , $i \in \Pi$, is a function $\sigma_i : \text{Hist}_i \rightarrow V$ which assigns to each history $h \in \text{Hist}_i$ a successor $v \in V$ such that $(\text{Last}(h), v) \in E$. Intuitively, it represents the choice of Player i when it is his turn to play and he can choose the next move taking into account the past history h . The set of strategies of Player i , $i \in \Pi$, is denoted by Σ_i .

A *strategy automaton* [Umm05] for a strategy σ_i of Player i in the arena $A = (\Pi, V, E, (V_i)_{i \in \Pi})$ is a tuple $\mathcal{M}_i = (M, m_0, \delta, \nu)$ such that:

- M is a (non-empty) finite set of *memory states*;
- $m_0 \in M$ is the *initial memory state*;
- $\delta : M \times V \rightarrow M$ is a *memory transition function*;
- $\nu : M \times V_i \rightarrow V$ with $(v, \nu(m, v)) \in E$ for all $m \in M$ and $v \in V$ is the *next-choice function*.

The transition function δ may be extended to a function $\delta^* : M \times \text{Hist} \rightarrow M$ defined by $\delta^*(m, \epsilon) = m$ for all $m \in M$ and $\delta^*(m, hv) = \delta(\delta^*(m, h), v)$ for all $m \in M$ and $hv \in \text{Hist}$. In this way, the strategy $\sigma_{\mathcal{M}_i}$ encoded by \mathcal{M}_i is defined as follows: $\sigma_{\mathcal{M}_i}(hv) = \nu(\delta^*(m_0, h), v)$ for all $hv \in \text{Hist}_i$.

A strategy σ_i of Player i is a *finite-memory strategy* if there exists a strategy automaton \mathcal{M}_i such that $\sigma_{\mathcal{M}_i} = \sigma_i$. Moreover, we say that the *memory size* of σ_i is at most equal to $|M|$ where M is the set of memory states of \mathcal{M}_i . A particular kind of finite-memory strategies are the strategies for which $|M| = 1$, these strategies are called *memoryless strategies* (or *positional strategies*). A strategy σ_i of Player i is memoryless if for all $hv, h'v' \in \text{Hist}$, $(v = v' \implies \sigma_i(hv) = \sigma_i(h'v'))$. Thus with a memoryless strategy σ_i , Player i only needs to consider the current vertex to make his choice, that is we can assume that $\sigma_i : V_i \rightarrow V$.

A play $\rho = \rho_0\rho_1\dots$ is said *consistent* with a strategy σ_i of Player i if for all $k \in \mathbb{N}$ such that $\rho_k \in V_i$, we have that $\rho_{k+1} = \sigma_i(\rho_0\dots\rho_k)$. This notion can easily be extended to histories.

A *strategy profile* $\sigma = (\sigma_i)_{i \in \Pi}$ is a tuple of strategies, one for each player. It is called *positional* (resp. *finite-memory*) if for each $i \in \Pi$, σ_i is positional (resp. finite-memory). A play ρ is said consistent with a strategy profile σ , if ρ is consistent with each strategy σ_i , $i \in \Pi$. When one fixes a strategy profile

σ and an initial vertex $v \in V$, there exists only one play, beginning in v , which is consistent with σ from v . This play is called the *outcome* of σ from v and is denoted by $\langle \sigma \rangle_v$. If we want to highlight the strategy of Player i , $i \in \Pi$, in a strategy profile σ , we write (σ_i, σ_{-i}) where $\sigma_{-i} = (\sigma_j)_{j \in \Pi \setminus \{i\}}$ and $-i = \Pi \setminus \{i\}$ denotes the set of all players except Player i . If $\sigma'_i \neq \sigma_i$ for some $i \in \Pi$, we say that σ'_i is a *deviating* strategy of Player i w.r.t σ .

Example 2.1.5. Let us come back to the arena A depicted in Example 2.1.2. We define a strategy σ_1 of Player 1 and a strategy σ_2 of Player 2. The strategy $\sigma_1 : V_1 \rightarrow V$ is defined by

$$\sigma_1(v) = \begin{cases} v_0 & \text{if } v \in \{v_1, v_3, v_4\} \\ v_3 & \text{if } v = v_2 \end{cases}.$$

This strategy is a memoryless strategy.


The strategy $\sigma_2 : \text{Hist}_2 \rightarrow V$ is defined by

$$\sigma_2(hv) = \begin{cases} v_1 & \text{if } |hv|_{v_0} \text{ is even} \\ v_2 & \text{otherwise} \end{cases}$$

where for all $h' = h'_0 \dots h'_k$ for some $k \in \mathbb{N}$ and for all $v' \in V$, $|h'|_{v'} = \{n \in \mathbb{N} \mid 0 \leq n \leq q \text{ and } h_n = v'\}$ denotes the number of occurrences of v' along h' . The idea is that when a even number of v_0 is visited along hv , Player 2 chooses to go to v_1 and he chooses to go to v_2 otherwise. This strategy is a finite-memory strategy with memory size at most equal to 2. We describe hereunder a strategy automaton \mathcal{M}_2 such that $\sigma_{\mathcal{M}_2} = \sigma_2$. See Figure 2.2 for a graphical representation of the memory transition function δ of \mathcal{M}_2 . Formally, $\delta : M \times V \rightarrow M$ is defined as follows:

$$\delta(m, v) = \begin{cases} m_1 & \text{if } m = m_0 \text{ and } v = v_0 \\ m_0 & \text{if } m = m_0 \text{ and } v \neq v_0 \\ m_0 & \text{if } m = m_1 \text{ and } v = v_0 \\ m_1 & \text{if } m = m_1 \text{ and } v \neq v_0 \end{cases}.$$

The next-choice function $\nu : M \times V_2 \rightarrow V$ is defined as follows:

Figure 2.2: The strategy automaton $\sigma_{\mathcal{M}_2}$ of Example 2.1.5.

$$\nu(m, v) = \begin{cases} v_1 & \text{if } m = m_0 \text{ and } v = v_0 \\ v_2 & \text{if } m = m_1 \text{ and } v = v_0 \end{cases}.$$

The outcome of (σ_1, σ_2) from v_0 is the play $\langle \sigma_1, \sigma_2 \rangle_{v_0} = (v_0 v_1 v_0 v_2 v_3)^\omega$. Indeed, if $\langle \sigma_1, \sigma_2 \rangle_{v_0} = \rho_0 \rho_1 \dots$, we have that: $\rho_0 = v_0$, $\rho_1 = \sigma_2(v_0) = \nu(\delta^*(m_0, \epsilon), v_0) = \nu(m_0, v_0) = v_1$, $\rho_3 = \sigma_2(v_0 v_1 v_0) = \nu(\delta^*(m_0, v_0 v_1), v_0) = \nu(m_1, v_0) = v_2$ where $\delta^*(m_0, v_0 v_1) = \delta(\delta^*(m_0, v_0), v_1) = \delta[\delta(\delta^*(m_0, \epsilon), v_0), v_1] = \delta(\delta(m_0, v_0), v_1) = \delta(m_1, v_1) = m_1$, and so on.

Games played on graphs To complete the definition of a *game played on a graph*, we need to define what is the *objective* of each player. In this document, all the objectives that we consider may be expressed thanks to an *objective function* $\text{Obj} : \text{Plays} \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$. If Obj_i is the objective function which represents the objective of Player i , then Obj_i assigns a value to each play in the arena A and Player i aims at *optimizing* this value: either *maximizing* or *minimizing*. We detail this notion of objective function in Section 2.2.

A game played on a graph is an arena A equipped with a profile of objective functions $\text{Obj} = (\text{Obj}_i)_{i \in \Pi}$, one for each player.

Definition 2.1.6 (Game played on a graph). A *game played on a graph* \mathcal{G} is a tuple $\mathcal{G} = (A, \text{Obj})$ where

- $A = (\Pi, V, E, (V_i)_{i \in \Pi})$ is an arena (see Definition 2.1.1);
- $\text{Obj} = (\text{Obj}_i)_{i \in \Pi}$ is a profile of objective functions, such that for each $i \in \Pi$, $\text{Obj}_i : \text{Plays} \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ represents the objective of Player i .

Since we only consider games played on graphs in this document, we refer

to the notion of a game played on a graph by the term “game”.

When one considers games played on graphs, one often fixes a particular *initial vertex* and one assumes that the token is placed in this particular vertex at the beginning of all plays. This kind of games is called an *initialized game* and is denoted by (\mathcal{G}, v_0) where $v_0 \in V$ is the fixed initial vertex in the arena A . In this document, the initial vertex is often denoted by v_0 .

Definition 2.1.7 (Initialized game). An *initialized game* (\mathcal{G}, v_0) is a tuple $(\mathcal{G}, v_0) = (A, \text{Obj})$ where:

- \mathcal{G} is a game as in Definition 2.1.6;
- $v_0 \in V$ is a fixed *initial vertex*.

In an initialized game (\mathcal{G}, v_0) , the plays begin in v_0 and we denote this particular set of plays by $\text{Plays}(v_0) = \{\rho \in \text{Plays} \mid \text{First}(\rho) = v_0\}$. We can define the set $\text{Hist}(v_0)$ and $\text{Hist}_i(v_0)$ in the same way. Additionally, the strategies of the players can be defined only on the histories beginning in v_0 . That is, for a strategy σ_i of Player i : $\sigma_i : \text{Hist}_i(v_0) \rightarrow V$. In particular, for a memoryless strategy we can define σ_i only on the reachable vertices of Player i from v_0 , that is $\sigma_i : V_i \cap \text{Succ}^*(v_0) \rightarrow V$. In this case, for all $i \in \Pi$, as the definition of a strategy $\sigma_i \in \Sigma_i$ (resp. a strategy profile σ) depends on the initial vertex, when it is not clear from the context we sometimes highlight that σ_i is a strategy (resp. σ is the strategy profile) in (\mathcal{G}, v_0) .

When a strategy profile σ is fixed in the initialized game (\mathcal{G}, v_0) , we called the outcome of σ from v_0 , *i.e.*, $\langle \sigma \rangle_{v_0}$, the outcome of σ in (\mathcal{G}, v_0) .

Remark 2.1.8. In all this document, when we write $\mathcal{G} = (A, \text{Obj})$ without any precision, we assume that $A = (\Pi, V, E, (V_i)_{i \in \Pi})$. We make the same assumption for an initialized game (\mathcal{G}, v_0) and for an arena A .

2.2 Objectives

In the previous section (Section 2.1), we have defined a game \mathcal{G} as (i) an arena A equipped with (ii) an objective function $\text{Obj}_i : \text{Plays} \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ for each Player i , $i \in \Pi$. If Player i , $i \in \Pi$, aims at maximizing his objective

function, this objective function can be seen as a *gain* for Player i . Conversely, if Player i , $i \in \Pi$, aims at minimizing his objective function, this objective function can be seen as a *cost* for Player i . We sometimes write Gain_i (resp. Cost_i) if the function Obj_i represents a gain (resp. a cost), we call Gain_i a *gain function* (resp. *cost function*).

Remark 2.2.1. Notice that if the objective of Player i is to maximize his gain function Gain_i , it is equivalent to say that he wants to minimize his cost function $\text{Cost}_i = -\text{Gain}_i$. And vice versa if Player i wants to minimize his cost function Cost_i .

Let us fix a game $\mathcal{G} = (A, \text{Obj})$. If for all $i \in \Pi$, Obj_i is a gain function Gain_i , we often write $\mathcal{G} = (A, \text{Gain})$ with $\text{Gain} = (\text{Gain}_i)_{i \in \Pi}$. For all $\rho \in \text{Plays}$, $\text{Gain}_i(\rho)$ represents the amount that Player i *earns* for the play ρ . Moreover, we say that $(\text{Gain}_i(\rho))_{i \in \Pi}$ is the *gain profile* of the play ρ and is denoted by $\text{Gain}(\rho)$.

In the same way, if for all $i \in \Pi$, Obj_i is a cost function Cost_i , we often write $\mathcal{G} = (A, \text{Cost})$ with $\text{Cost} = (\text{Cost}_i)_{i \in \Pi}$. For all $\rho \in \text{Plays}$, the value $\text{Cost}_i(\rho)$ represents the amount that Player i *loses* for the player ρ . Additionally, $(\text{Cost}_i(\rho))_{i \in \Pi}$ is the *cost profile* of the play ρ and is denoted by $\text{Cost}(\rho)$.

Finally, for all $x \in (\mathbb{R} \cup \{-\infty, +\infty\})^{|\Pi|}$, for all $i \in \Pi$, x_i depicts the i th component of x . For all $x \in (\mathbb{R} \cup \{-\infty, +\infty\})^{|\Pi|}$ and $y \in (\mathbb{R} \cup \{-\infty, +\infty\})^{|\Pi|}$, we write $x \leq y$ (resp. $x < y$) if and only if for all $i \in \Pi$, $x_i \leq y_i$ (resp. $x_i < y_i$).

2.2.1 Qualitative objectives

When the objective of a player is to satisfy (or to achieve) a property without any quantitative consideration, this objective is called a *qualitative objective*. In other words, with a qualitative objective either the player achieves his objective or not. Commonly [GTW02] these objectives are described by a subset of plays $\text{Win} \subseteq \text{Plays}$ called *the set of winning plays*. Given $\text{Win}_i \subseteq \text{Plays}$ the set of winning plays of Player i , $i \in \Pi$, if a play ρ belongs to Win_i we say that ρ is *winning* for Player i and that this latter player *wins*. Otherwise, if the play ρ does not belong to Win_i , we say that Player i *loses* along ρ .

Some classical qualitative objectives studied in the literature are: (*qualitative*) *Reachability, Safety, Büchi, co-Büchi, Muller, Explicit Muller, Parity,*

Streett and *Rabin* objectives.

Definition 2.2.2 (Classical qualitative objectives). Let $A = (\Pi, V, E, (V_i)_{i \in \Pi})$ be an arena. The set $\text{Win}_i \subseteq \text{Plays}$ describes a (*qualitative*) *Reachability*, *Safety*, *Büchi*, *Co-Büchi*, *Parity*, *Explicit Muller*, *Muller*, *Rabin*, or *Streett* objective for Player i if and only if Win_i is composed of the plays ρ satisfying:

- (*qualitative*) *Reachability*: given a *target set* $F \subseteq V$, $\text{Occ}(\rho) \cap F \neq \emptyset$;
- *Safety*: given $F \subseteq V$, $\text{Occ}(\rho) \cap F = \emptyset$;
- *Büchi*: given a *target set* $F \subseteq V$, $\text{Inf}(\rho) \cap F \neq \emptyset$;
- *Co-Büchi*: given $F \subseteq V$, $\text{Inf}(\rho) \cap F = \emptyset$;
- *Parity*: given a coloring function $c : V \rightarrow \{1, \dots, d\}$, $\max(\text{Inf}(c(\rho)))^a$ is even;
- *Explicit Muller*: given $\mathcal{F} \subseteq 2^V$, $\text{Inf}(\rho) \in \mathcal{F}$;
- *Muller*: given a coloring function $c : V \rightarrow \{1, \dots, d\}$, and $\mathcal{F} \subseteq 2^{c(V)}$, $\text{Inf}(c(\rho)) \in \mathcal{F}$;
- *Rabin*: given $(G_j, R_j)_{1 \leq j \leq k}$ a family of pair of sets $G_j, R_j \subseteq V$, there exists $j \in 1, \dots, k$ such that $\text{Inf}(\rho) \cap G_j \neq \emptyset$ and $\text{Inf}(\rho) \cap R_j = \emptyset$;
- *Streett*: given $(G_j, R_j)_{1 \leq j \leq k}$ a family of pair of sets $G_j, R_j \subseteq V$, for all $j \in 1, \dots, k$, $\text{Inf}(\rho) \cap G_j = \emptyset$ or $\text{Inf}(\rho) \cap R_j \neq \emptyset$.

^aWhere $c(\rho) = c(\rho_0)c(\rho_1) \dots c(\rho_n) \dots$

From the set of winning plays to the objective function We now explain how, for each $i \in \Pi$, we obtain a gain function $\text{Gain}_i : \text{Plays} \rightarrow \{0, 1\}$ from Win_i . For each player $i \in \Pi$, let $\text{Win}_i \subseteq \text{Plays}$ be his set of winning plays, the gain function $\text{Gain}_i : \text{Plays} \rightarrow \{0, 1\}$ is defined such that $\text{Gain}_i(\rho) = 1$ (resp. $\text{Gain}_i(\rho) = 0$) if and only if $\rho \in \text{Win}_i$ (resp. $\rho \notin \text{Win}_i$).

In the rest of this document we use interchangeably either Win_i or Gain_i

(associated with Win_i) to denote the qualitative objective of Player i . Notice that when we equip a player i with a gain function Gain_i , we implicitly assume that the informations needed to define the gain function are provided. That is:

- For (qualitative) Reachability, Safety, Büchi, co-Büchi objectives, the set $F_i \subseteq V$ is given;
- For Parity and Muller objectives, the coloring function c_i is given;
- For Explicit Muller and Muller objectives, the set $\mathcal{F}_i \subseteq 2^V$ is given;
- For Rabin and Stress objectives, the family $(G_j^i, R_j^i)_{1 \leq j \leq k}$ such that $G_j^i, R_j^i \subseteq V$ is given.

Remark 2.2.3. In the rest of this document, we sometimes need to refer to the gain functions associated with particular qualitative objective functions. We define:

- $\text{qR} : \text{Plays} \rightarrow \{0, 1\}$ to denote the gain function associated with a qualitative reachability objective;
- $\text{Safe} : \text{Plays} \rightarrow \{0, 1\}$ to denote the gain function associated with a safety objective;
- $\text{Buchi} : \text{Plays} \rightarrow \{0, 1\}$ to denote the gain function associated with a Büchi objective.

Remark 2.2.4. Even if it is less intuitive, in Section 2.3.1, we sometimes represent a qualitative objective thanks to a cost function rather than a gain function. With this convention, $\text{Cost}_i : \text{Plays} \rightarrow \{0, 1\}$ is such that $\text{Cost}_i(\rho) = 1$ (resp. $\text{Cost}_i(\rho) = 0$) if and only if $\rho \notin \text{Win}_i$ (resp. $\rho \in \text{Win}_i$).

Boolean Games A (initialized) game is called a (initialized) *Boolean game*, if each player has a qualitative objective.

Definition 2.2.5 (Boolean game). Let $\mathcal{G} = (A, \text{Obj})$ be a game, if for each $i \in \Pi$, Obj_i is a gain function $\text{Gain}_i : \text{Plays} \rightarrow \{0, 1\}$ then \mathcal{G} (resp. (\mathcal{G}, v_0))

is a *Boolean game*.

Thus, in a Boolean game, each player aims at maximizing his gain. For $O \in \{\text{qualitative Reachability, Safety, Büchi, co-Büchi, Parity, Explicit Muller, Muller, Rabin, Streett}\}$, we say that the Boolean game \mathcal{G} is an O game if each player has an O objective. For example, if $O = \text{Büchi}$, a Büchi game is a game where all players have a Büchi objective.

In the rest of this document, a qualitative reachability game is denoted $\mathcal{G} = (A, \text{qR}, (F_i)_{i \in \Pi})$ with $\text{qR} = (\text{qR}_i)_{i \in \Pi}$ and for all $i \in \Pi$, $F_i \subseteq V$ is the target set of Player i . In the same way, a safety game is denoted $\mathcal{G} = (A, \text{Safe}, (F_i)_{i \in \Pi})$ with $\text{Safe} = (\text{Safe}_i)_{i \in \Pi}$ and for all $i \in \Pi$, $F_i \subseteq V$ is the set of vertices that Player i wants to avoid. Moreover, a Büchi game is denoted $\mathcal{G} = (A, \text{Buchi}, (F_i)_{i \in \Pi})$ with $\text{Buchi} = (\text{Buchi}_i)_{i \in \Pi}$ and for all $i \in \Pi$, $F_i \subseteq V$ is the target set of Player i .

Example 2.2.6. We consider the initialized Boolean game $(\mathcal{G}, v_0) = (A, \text{Gain})$ with the arena A given in Example 2.1.2 and $\text{Gain} = (\text{qR}_1, \text{Buchi}_2)$. The objective of Player 1 is a qualitative objective given by $F_1 = \{v_1, v_3\}$ and the objective of Player 2 is a Büchi objective given by $F_2 = \{v_1\}$. Thus, Player 1 wants to reach (at least once) either the vertex v_1 or the vertex v_3 while Player 2 wants to reach vertex v_1 infinitely often.

We illustrate this game on Figure 2.3, the gray vertex represents F_1 and double circled vertices represent F_2 .




Figure 2.3: Initialized Boolean game (\mathcal{G}, v_0) where $F_1 = \{v_1, v_3\}$ and $F_2 = \{v_1\}$.

The play $v_0 v_2 v_4 (v_0 v_2 v_3)^\omega$ is winning for Player 1 but not for Player 2. Indeed, $\text{Occ}(\rho) = \{v_0, v_2, v_3, v_4\}$ and $\text{Inf}(\rho) = \{v_0, v_2, v_3\}$, thus $\text{Occ}(\rho) \cap F_1 \neq \emptyset$ and $\text{Occ}(\rho) \cap F_2 = \emptyset$. It follows that $\text{Gain}(\rho) = (\text{qR}_1(\rho), \text{Buchi}_2(\rho)) = (1, 0)$.

2.2.2 Quantitative objectives

In the previous section, we have seen how, thanks to qualitative objectives, we are able to express if a property is satisfied or not. Sometimes this binary point of view is too restrictive and we want to express quantitative specifications. For example, we not only want to know that a player reaches a given subset of vertices (qualitative reachability), we want to count the amount of energy it takes (quantitative reachability). In order to express quantitative specifications, we associate a *quantitative objective* to each player. That is, for each $i \in \Pi$: $\text{Obj}_i : \text{Plays} \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ ².

There exists a lot of different quantitative objectives in the literature, in this document we choose to mainly focus on *quantitative reachability*, *weighted reachability*, *discounted-sum* and *mean-payoff* objectives.

When a player has a quantitative/weighted reachability objective he aims at reaching a subset of vertices F (his *target set*) as soon as possible. In the case of quantitative reachability objective, given a play ρ we only count the number of edges between the initial vertex and the first vertex v along ρ such that $v \in F$. If such a vertex does not exist, we assume that the player has to pay a cost of $+\infty$. In the case of weighted reachability objective, given a play ρ , we aggregate the sum of the weights on the edges from the initial vertex to the first vertex v along ρ such that $v \in F$. As for quantitative reachability, if this vertex v does not exist, the player has a cost equal to $+\infty$.

With a reachability objective (quantitative or weighted), when the player reaches a vertex of his target he does not care about what happens in the future. Indeed their next actions (and those of the other players) do not affect his cost. With a discounted-sum objective, there is no such moment when the cost of the player will not change anymore. The weights on the edges along the play are infinitely aggregated but additionally these weights are discounted by a discount factor strictly between 0 and 1 which varies along the play.

A mean-payoff objective considers the limit of the means of the weights along the play. Since this limit may not exist, we consider the limsup of this mean.

² Notice that a qualitative objective can be seen as a quantitative objective.

Definition 2.2.7. Given an (weighted) arena $A = (\Pi, V, E, (V_i)_{i \in \Pi}, (w_i)_{i \in \Pi})$ where, for each $i \in \Pi$, w_i is the weight function of Player i (see Remark 2.1.4). Let $i \in \Pi$ be a player, we define the following cost functions for Player i : for all plays $\rho = \rho_0 \rho_1 \dots$ in A ,

- *Quantitative reachability* objective: given a subset of vertices $F_i \subseteq V$ called *the target set of Player i* , $\text{QR}_i : \text{Plays} \rightarrow \mathbb{N} \cup \{+\infty\}$ is such that

$$\text{QR}_i(\rho) = \begin{cases} k & \text{if } k \text{ is the least index such that } \rho_k \in F \\ +\infty & \text{if such } k \text{ does not exist} \end{cases}.$$

- *Weighted reachability* objective: given a subset of vertices $F_i \subseteq V$ called *the target set of Player i* and the weight function $w_i : E \rightarrow \mathbb{N}_0^a$ of Player i , $\text{WR}_i : \text{Plays} \rightarrow \mathbb{N} \cup \{+\infty\}$ is such that

$$\text{WR}_i(\rho) = \begin{cases} \sum_{n=0}^{k-1} w_i(\rho_n, \rho_{n+1}) & \text{if } k \text{ is the least index st. } \rho_k \in F \\ +\infty & \text{if such } k \text{ does not exist} \end{cases}.$$

- *Discounted-sum* objective: given a discount factor $\lambda \in]0, 1[$, $\text{DS}_i^\lambda : \text{Plays} \rightarrow \mathbb{R}$ is such that

$$\text{DS}_i^\lambda(\rho) = (1 - \lambda) \cdot \sum_{n=0}^{+\infty} \lambda^n \cdot w_i(\rho_n, \rho_{n+1}).$$

- *Mean-payoff* objective: $\overline{\text{MP}}_i : \text{Plays} \rightarrow \mathbb{R}$

$$\overline{\text{MP}}_i(\rho) = \limsup_{k \rightarrow +\infty} \frac{\sum_{n=0}^{k-1} w_i(\rho_n, \rho_{n+1})}{k}.$$

^aIn this document, when we consider a weighted reachability objective, we assume that the weights on the edges are in $\mathbb{N}_0 = \{1, 2, \dots\}$. This objective may also be defined with negative weights, but we will explicitly precise it when it is the case in our examples.

Remark 2.2.8. A quantitative reachability objective is in particular a weighted reachability objective with $w_i(e) = 1$ for all $e \in E$.

As for qualitative objectives, when we equip a player i with a cost function

Cost_{*i*}, we implicitly assume that:

- For quantitative reachability and weighted reachability objectives, the target set $F_i \subseteq V$ is given;
- For weighted reachability, discounted-sum and mean-payoff objectives, the weight function w_i is given;
- For the discounted-sum objective, the discount factor $\lambda \in]0, 1[$ is given.

In the same way as for qualitative objectives, in a game \mathcal{G} with arena A , if each player has an objective $O \in \{\text{quantitative reachability, weighted reachability, discounted-sum, mean-payoff}\}$, we say that the game is a *O game*. For example, if $O = \text{quantitative reachability}$, it means that for all $i \in \Pi$, $\text{Obj}_i = \text{QR}_i$ and we call the game a quantitative Reachability game.

In this case we often explicit the target sets and the weight functions as follows:

- A quantitative Reachability game is denoted $\mathcal{G} = (A, \text{QR}, (F_i)_{i \in \Pi})$ with $\text{QR} = (\text{QR}_i)_{i \in \Pi}$ and for all $i \in \Pi$, $F_i \subseteq V$ is the target set of Player i ;
- A weighted Reachability game is denoted $\mathcal{G} = (A, \text{WR}, (F_i)_{i \in \Pi})$ with a weighted arena $A = (\Pi, V, E, (V_i)_{i \in \Pi}, (w_i)_{i \in \Pi})$, $\text{WR} = (\text{WR}_i)_{i \in \Pi}$ and for all $i \in \Pi$, $F_i \subseteq V$ is the target set of Player i ;

For the rest of this document, we generalize qualitative, quantitative and weighted Reachability games thanks to the notion of *Reachability games*.

Definition 2.2.9 (Reachability game). A *Reachability game* is a game $\mathcal{G} = (A, \text{Reach}, (F_i)_{i \in \Pi})$ with $\text{Reach} = (\text{Reach}_i)_{i \in \Pi}$ and $F_i \subseteq V$ is the target set of Player i for all $i \in \Pi$ and where Reach refers to qR, QR or WR if \mathcal{G} is a qualitative, quantitative or weighted reachability game respectively.

Example 2.2.10. Let us come back to the arena depicted in Figure 2.1 on which we have added a weight function for each player (see Figure 2.4). A tuple $t \in \mathbb{N}^{|\Pi|}$ on an edge $e \in E$ means that for each $i \in \Pi$, $w_i(e) = t_i$, if there is no weight on an edge $e \in E$, we assume that $w_i(e) = (0, 0)$ for all $i \in \Pi$.




Figure 2.4: Initialized game (\mathcal{G}, v_0) with quantitative objectives.

Let $\rho = (v_0 v_2 v_3)^\omega$ be a play in $A = (\Pi, V, E, (V_i)_{i \in \Pi}, (w_i)_{i \in \Pi})$.

- Let us assume that both players have a quantitative/weighted reachability objective and that the target set of Player 1 (resp. Player 2) is $F_1 = \{v_3\}$ (resp. $F_2 = \{v_1\}$). Then, we have that $\text{QR}(\rho) = (\text{QR}_1(\rho), \text{QR}_2(\rho)) = (2, +\infty)$ and $\text{WR}(\rho) = (\text{WR}_1(\rho), \text{WR}_2(\rho)) = (3, +\infty)$.

2.2.3 Continuous and prefix-linear objectives

In Section 2.2.1 and Section 2.2.2 we define what is a qualitative or a quantitative objective and we provide some classical examples. In this section we are interested in some properties that these objective functions may have. In particular, we consider *continuous* and *prefix-linear* objective functions.


Continuous objective functions

We assume that the reader is familiar with the notion of topology on V^ω : V is endowed with the discrete topology and V^ω with the product topology (see e.g., [PP04]). The following definition explains what is a *continuous objective function*.

Definition 2.2.11 (Continuous objective function). An objective function Obj_i is *continuous* if whenever $\lim_{n \rightarrow +\infty} \rho_n = \rho$, we have that $\lim_{n \rightarrow +\infty} \text{Obj}_i(\rho_n) = \text{Obj}_i(\rho)$.

Remark 2.2.12. Let us comment the continuity of some of the objective functions introduced in this section.

1. For all $\lambda \in]0, 1[$, the objective function $DS^\lambda : \text{Plays} \rightarrow \mathbb{R}$ is a continuous function.
2. We explicit an example which shows that the objective function \overline{MP} is not a continuous objective function. Let us consider the game with only one player and such that its arena is depicted by the figure just bellow. We assume that $w_1 : E \rightarrow \mathbb{Z}$ is such that for all $e \in E$ such that $e \neq (v_1, v_1)$, we have that $w_1(e) = 0$ and $w_1(v_1, v_1) = 1$.



Let $(\rho^n)_{n \in \mathbb{N}}$ be a sequence of plays in this arena such that for each $n \in \mathbb{N}$, $\rho^n = v_0^{n+1}v_1^\omega$, that is the vertex v_0 is repeated $n + 1$ times and then the vertex v_1 is repeated infinitely often. This sequence converges to the play v_0^ω but $\overline{MP}_1(\rho^n) = 1$ for all $n \in \mathbb{N}$ and $\overline{MP}_1(v_0^\omega) = 0$.

3. With the same example and by assuming that $F_1 = \{v_1\}$, we can prove that the qualitative reachability objective function is not a continuous function.

When we deal with a weighted/quantitative reachability objective, we can transform the corresponding objective functions into real-valued cost functions³ which are continuous. These new objective functions will be equivalent for our purpose.

Proposition 2.2.13. *Given a weighted arena $A = (\Pi, V, E, (V_i)_{i \in \Pi}, (w_i)_{i \in \Pi})$ with $w_i : E \rightarrow \mathbb{N}_0$, the function WR_i can be transformed into a real-valued*

³A *real-valued objective function* is an objective function which assigns a real value to each play.

and continuous function $WR'_i : \text{Plays} \rightarrow [0, 1]$:

$$WR'_i(\rho) = \begin{cases} 1 - \frac{1}{WR_i(\rho)+1} & \text{if } WR_i(\rho) < +\infty \\ 1 & \text{otherwise} \end{cases}.$$

Prefix-linear objectives

In Part II, we state some results which hold for a particular class of games: games with *prefix-linear objectives*. Thus in this section, we explain what is a *prefix-independent objective* and a *prefix-linear objective*.

Let A be an arena, a prefix-independent objective is an objective such that for all plays $\rho \in \text{Plays}$, the value of the objective function for ρ remains the same even if we add/delete a finite prefix h to ρ .

Definition 2.2.14 (Prefix-independent objective). Let A be an arena, an objective function Obj_i is *prefix-independent* in A if for all $hv \in \text{Hist}_A$ and $\rho \in \text{Plays}_A(v)$:

$$\text{Obj}_i(h\rho) = \text{Obj}_i(\rho).$$

If the objective of Player i is given by a prefix-independent objective function, we say that the objective is *prefix-independent*.

A prefix-linear objective function of Player i is an objective Obj_i such that for all plays ρ in the arena A and all prefixes $hv < \rho$ such that one can write ρ as $\rho = h\rho'$ with $\rho' \in \text{Plays}(v)$, we can split the value of the objective function between the value of the prefix h and the value of the suffix ρ' . That is, there exist $a \in \mathbb{R}$ and $b \in \mathbb{R}^+$ such that $\text{Obj}_i(\rho) = a + b \cdot \text{Obj}_i(\rho')$.

Definition 2.2.15 (Prefix-linear objective). Let A be an arena, an objective function Obj_i is *prefix-linear* in A if for all $hv \in \text{Hist}_A$: there exists $a(h, v) \in \mathbb{R}$ and $b(h, v) \in \mathbb{R}^+$ such that for all $\rho \in \text{Plays}_A(v)$:

$$\text{Obj}_i(h\rho) = a(h, v) + b(h, v) \cdot \text{Obj}_i(\rho).$$

If the objective of Player i is given by a prefix-linear objective function, we

say that the objective is *prefix-linear*.

If an objective function Obj_i is prefix-linear and for all $hv \in \text{Hist}$, $b(h, v)$ is strictly greater than 0, then we say that Obj_i is *strongly prefix-linear*.

Definition 2.2.16 (Strongly prefix-linear objective). Let A be an arena, an objective function Obj_i is *strongly prefix-linear* in A if for all $hv \in \text{Hist}_A$: there exists $a(h, v) \in \mathbb{R}$ and $b(h, v) \in \mathbb{R}_0^+$ such that for all $\rho \in \text{Plays}_A(v)$:

$$\text{Obj}_i(h\rho) = a(h, v) + b(h, v) \cdot \text{Obj}_i(\rho).$$

If the objective of Player i is given by a strongly prefix-linear objective function, we say that the objective is *strongly prefix-linear*.

Remark 2.2.17.

- Notice that a prefix-independent objective in A is a strongly prefix-linear objective in A with $a(h, v) = 0$ and $b(h, v) = 1$ for all $hv \in \text{Hist}$.
- When an objective (function) is prefix-independent (resp. (strongly) prefix-linear) in all A , we only say that this objective (function) is prefix-independent (resp. (strongly) prefix-linear).

Example 2.2.18. We give some examples of prefix-independent and prefix-linear objective functions.

- Büchi, co-Büchi, Parity, Explicit Muller, Muller, Rabin and Streett objectives are prefix-independent objectives.
- A mean-payoff objective is a prefix-independent objective.
- The qualitative reachability objectives and the safety objectives are prefix-linear objectives.
- A discounted-sum objective is a prefix-linear objective. Indeed, given $\lambda \in]0, 1[$, and DS_i^λ the discounted-sum objective of Player i , let $hv \in \text{Hist}$ be such that $hv = h_0 \dots h_k$, we take $a(h, v) = (1 - \lambda) \sum_{n=0}^{k-1} \lambda^n w_i(h_n, h_{n+1})$

and $b(h, v) = \lambda^k$. We have for all $\rho \in \text{Plays}(v)$ that $\text{DS}_i^\lambda(h\rho) = a(h, v) + b(h, v) \cdot \text{DS}_i^\lambda(\rho)$.

- A quantitative/weighted reachability objective is a prefix-linear objective. Given a weighted reachability objective WR_i , let $hv = h_0 \dots h_k$,
 - if there exists $0 \leq \ell \leq k$ such that $h_\ell \in F_i$, where F_i is the target set of Player i , and if we assume that ℓ is the least index such that $h_\ell \in F_i$, we take $a(h, v) = \sum_{n=0}^{\ell-1} w_i(h_n, h_{n+1})$ and $b(h, v) = 0$ with the convention that $0 \cdot +\infty = 0$;
 - otherwise we take $a(h, v) = \sum_{n=0}^{k-1} w_i(h_n, h_{n+1})$ and $b(h, v) = 1$.

We have for all $\rho \in \text{Plays}(v)$, $\text{WR}_i(h\rho) = a(h, v) + b(h, v) \cdot \text{WR}_i(\rho)$.

Remark 2.2.19. Notice that qualitative, quantitative and weighted reachability objectives as well as safety objective are prefix-linear objectives but not strongly prefix-linear objectives.

We conclude this section by providing some properties of strongly prefix-linear objectives.

Lemma 2.2.20. *Given a game $\mathcal{G} = (A, \text{Obj})$ such that Obj_i is strongly prefix-linear in A for all $i \in \Pi$.*

For all $i \in \Pi$, all $hv \in \text{Hist}$ and for all $\rho, \rho' \in \text{Plays}(v)$:

$$(C_1) \quad \text{Obj}_i(\rho) \geq \text{Obj}_i(\rho') \implies \text{Obj}_i(h\rho) \geq \text{Obj}_i(h\rho');$$

$$(C_2) \quad \text{Obj}_i(\rho) > \text{Obj}_i(\rho') \implies \text{Obj}_i(h\rho) > \text{Obj}_i(h\rho').$$

Proof. Let $\mathcal{G} = (A, \text{Obj})$ be a game such that Obj_i is strongly prefix-linear for all $i \in \Pi$.

Let $i \in \Pi$ be a player, $hv \in \text{Hist}$ be a history and ρ, ρ' be two plays beginning in v . Since Obj_i is strongly prefix-linear, we have that there exist $a(h, v) \in \mathbb{R}$ and $b(h, v) \in \mathbb{R}_0^+$ such that for all $\tilde{\rho} \in \text{Plays}(v)$, we have:

$$\text{Obj}_i(h\tilde{\rho}) = a(h, v) + b(h, v) \cdot \text{Obj}_i(\tilde{\rho}). \quad (2.1)$$

1. First, let us assume that $\text{Obj}_i(\rho) \geq \text{Obj}_i(\rho')$. Since $b(h, v) \in \mathbb{R}_0^+$,

we have that $a(h, v) + b(h, v) \cdot \text{Obj}_i(\rho) \geq a(h, v) + b(h, v) \cdot \text{Obj}_i(\rho')$.
From (2.1) comes the result.

2. Second, let us assume that $\text{Obj}_i(\rho) > \text{Obj}_i(\rho')$. Since $b(h, v) \in \mathbb{R}_0^+$, we have that $a(h, v) + b(h, v) \cdot \text{Obj}_i(\rho) > a(h, v) + b(h, v) \cdot \text{Obj}_i(\rho')$.
From (2.1) comes the result. \square

Even if qualitative, quantitative and weighted reachability objectives are not strongly prefix-linear, it is clear that the property (C₁) also holds for those objectives. Moreover, if the target set of Player i is not reached along h , property (C₂) is also satisfied by reachability objectives.

Lemma 2.2.21. *Given a Reachability game $\mathcal{G} = (\text{A}, \text{Reach}, (F_i)_{i \in \Pi})$.*

(C₁) *For all $i \in \Pi$, all $hv \in \text{Hist}$ and for all $\rho, \rho' \in \text{Plays}(v)$: $\text{Reach}_i(\rho) \geq \text{Reach}_i(\rho') \implies \text{Reach}_i(h\rho) \geq \text{Reach}_i(h\rho')$;*

(C₂) *For all $i \in \Pi$, all $hv \in \text{Hist}$ such that h does not visit F_i and for all $\rho, \rho' \in \text{Plays}(v)$: $\text{Reach}_i(\rho) > \text{Reach}_i(\rho') \implies \text{Reach}_i(h\rho) > \text{Reach}_i(h\rho')$.*

Proof. Let $\mathcal{G} = (\text{A}, \text{Reach}, (F_i)_{i \in \Pi})$ be either a qualitative, quantitative or weighted Reachability game.

The proof of (C₁) is the same as the one provided in Lemma 2.2.20. We thus only prove (C₂). Let $i \in \Pi$.

- If $\text{Reach}_i = \text{qR}_i$: let $hv \in \text{Hist}$ be such that h does not visit F_i and let $\rho, \rho' \in \text{Plays}(v)$ be such that $\text{qR}_i(\rho) > \text{qR}_i(\rho')$, then $\text{qR}_i(h\rho) = \text{qR}_i(\rho) > \text{qR}_i(\rho') = \text{qR}_i(h\rho')$.
- If $\text{Reach}_i = \text{WR}_i$: let $hv \in \text{Hist}$ be such that h does not visit F_i and let $\rho, \rho' \in \text{Plays}(v)$ be such that $\text{WR}_i(\rho) > \text{WR}_i(\rho')$, then, by assuming that $hv = h_0 \dots h_k$ and $S = \sum_{n=0}^{k-1} w_i(h_n, h_{n+1})$, $\text{WR}_i(h\rho) = S + \text{WR}_i(\rho) > S + \text{WR}_i(\rho') = \text{WR}_i(h\rho')$.

- If $\text{Reach}_i = \text{QR}_i$, it is a particular case $\text{Reach}_i = \text{WR}_i$. \square

2.2.4 Play of a given gain profile in Boolean games

In the remaining part of this document, given a Boolean game \mathcal{G} , we sometimes need to know what is the complexity of checking the existence of a play ρ in the game arena such that $\text{Gain}(\rho) = p$ for some gain profile $p \in \{0, 1\}^{|\Pi|}$. This section considers this question for classical qualitative objectives.

Lemma 2.2.22 ([BBGR18]). *Let \mathcal{G} be a Boolean game. Let $p \in \{0, 1\}^{|\Pi|}$ and $v \in V$. Determining whether there exists a play with gain profile p from v is*

- *in polynomial time for Büchi, co-Büchi, Explicit Muller, and Parity objectives,*
- *in $\mathcal{O}(2^{|\Pi|}(|V| + |E|))$ time for Reachability and Safety objectives, and*
- *in $\mathcal{O}(2^L \cdot M + (L^L \cdot |V|)^5)$ time for Rabin, Streett, and Muller objectives, where $L = 2^\ell$ and*
 - $\ell = \sum_{i=1}^{|\Pi|} 2 \cdot k_i$ and $M = \mathcal{O}(\sum_{i=1}^{|\Pi|} 2 \cdot k_i)$ *such that for each player $i \in \Pi$, k_i is the number of his pairs $(G_j^i, R_j^i)_{1 \leq j \leq k_i}$ in the case of Rabin and Streett objectives, and*
 - $\ell = \sum_{i=1}^{|\Pi|} d_i$ and $M = \mathcal{O}(\sum_{i=1}^{|\Pi|} |\mathcal{F}_i| \cdot d_i)$ *such that for each player $i \in \Pi$, d_i (resp. $|\mathcal{F}_i|$), is the number of his colors (the size of his family of subsets of colors) in the case of Muller objectives.*

The general approach to prove this lemma is the following one. A play with gain profile p from v in a Boolean game \mathcal{G} is a play satisfying an objective O equal to the conjunction of objectives Win_i (when $p_i = 1$) and of objectives $V^\omega \setminus \text{Win}_i$ (when $p_i = 0$). It is nothing else than an infinite path in the underlying graph $G = (V, E)$ satisfying some particular ω -regular objective O . The existence of such paths is a well studied problem; we gather in the next proposition the known results that we need for proving Lemma 2.2.22. Recall

that a *Generalized Reachability* (resp. *Generalized Büchi*) objective O is a conjunction of Reachability (resp. Büchi) objectives. Moreover, an objective O equal to a Boolean combination of Büchi objectives, called a *BC Büchi objective*, is defined as follows. Let F_1, \dots, F_ℓ be ℓ subsets of V , and ϕ be a Boolean formula over variables f_1, \dots, f_ℓ . We say that an infinite path ρ in G satisfies $(\phi, F_1, \dots, F_\ell)$ if the truth assignment

$$f_i = 1 \text{ if and only if } \text{Inf}(\rho) \cap F_i \neq \emptyset, \text{ and } f_i = 0 \text{ otherwise}$$

satisfies ϕ . All operators \vee, \wedge, \neg are allowed in a BC Büchi objective. However we denote by $|\phi|$ the size of ϕ equal to the number of disjunctions and conjunctions inside ϕ , and we say that the BC Büchi objective $(\phi, F_1, \dots, F_\ell)$ is of size $|\phi|$ and with ℓ variables.

Proposition 2.2.23 ([BBGR18]). *Let $G = (V, E)$ be a graph, $v \in V$ be one of its vertices, and $O \subseteq V^\omega$ be an objective. Then deciding the existence of an infinite path from v in G that satisfies O is*

- *in polynomial time when O is either a Streett objective, or an Explicit Muller objective, or the complement of an Explicit Muller objective, or a conjunction of a Generalized Büchi objective and a co-Büchi objective,*
- *in $\mathcal{O}(2^\ell(|V| + |E|))$ time when O is a conjunction of a Generalized Reachability objective and a Safety objective, where ℓ is the number of reachability objectives,*
- *in $\mathcal{O}(2^L \cdot |\phi| + (L^L \cdot |V|)^5)$ time for a BC Büchi objective $O = (\phi, F_1, \dots, F_\ell)$, where $L = 2^\ell$.*

Proof. Let O be an objective. If it is a Streett objective, then the result is proved in [EL87].

For the other objectives, we use known results about *two-player zero-sum games*, where player A aims at satisfying a certain objective O whereas player B tries to prevent him to satisfy it. A classical problem is to decide whether player A has a winning strategy that allows him to satisfy O against any strategy of player B , see for instance [Bru17, GTW02]. When player A is

the only one to play, the existence of a winning strategy for him is equivalent to the existence of a path satisfying O (see [Bru17, Section 3.1]). This is exactly the problem that we want to solve. In the rest of the proof, we mean by (G, O) a two-player zero-sum game, where player A (resp. player B) aims at satisfying O (resp. $V^\omega \setminus O$).

If O is an Explicit Muller objective, then deciding the existence of a winning strategy for player A (resp. player B) (G, O) can be done in polynomial time by [Hor08]. Thus the case where O is the complement of an Explicit Muller objective is also proved (by exchanging players A and B).

Suppose that O is the conjunction of a Generalized Büchi objective and a co-Büchi objective. By a classic reduction (see [CHVB18, Theorem 12]), the game (G, O) can be polynomially transformed into a game (G', O') with an objective O' equal to the conjunction of a Büchi objective and a co-Büchi objective. The existence of a winning strategy for player A in the latter game can be tested in polynomial time [dAF07].

Suppose that O is the conjunction of a Generalized Reachability objective and a Safety objective, such that ℓ is the number of Reachability objectives and F is the set of vertices to be avoided in the Safety objective. We first treat separately the Safety objective by removing from G all the vertices of F . This can be done in $\mathcal{O}(|V| + |E|)$ time. In the resulting graph G' , we then test the existence of a winning strategy for player A in the game (G', O') with O' being the Generalized Reachability objective. This can be done in $\mathcal{O}(2^\ell(|V| + |E|))$ time [FH13].

If O is a BC Büchi objective $(\phi, F_1, \dots, F_\ell)$, then deciding the existence of a winning strategy for player A in the game (G, O) can be done in $\mathcal{O}(2^L \cdot |\phi| + (L^L \cdot |V|)^5)$ time with $L = 2^\ell$ by [BHR18]. \square

Proof of Lemma 2.2.22. A play with gain profile p from v in \mathcal{G} is a play satisfying an objective O equal to the conjunction of objectives Win_i (when $p_i = 1$) and of objectives $V^\omega \setminus \text{Win}_i$ (when $p_i = 0$). For each type of Boolean objectives Win_i , we first explain what kind of objective O we obtain and we then apply Proposition 2.2.23.

- Consider a Boolean game \mathcal{G} with Parity objectives. In this case, as

Win_i is a Parity objective for all $i \in \Pi$, each $V^\omega \setminus \text{Win}_i$ is again a Parity objective, and O is thus a conjunction of Parity objectives which is a Streett objective [CHP07]. Therefore the existence of a play with gain profile p in \mathcal{G} can be tested in polynomial time by Proposition 2.2.23.

- Consider the case of Büchi objectives. Then, the intersection of Büchi objectives Win_i (when $p_i = 1$) is a Generalized Büchi objective and the intersection of co-Büchi objectives $V^\omega \setminus \text{Win}_i$ (when $p_i = 0$) is again a co-Büchi objective. Hence O is the conjunction of a Generalized Büchi objective and a co-Büchi objective. The existence of a play with gain profile p in \mathcal{G} can be tested in polynomial time by Proposition 2.2.23.

Notice that the case of Boolean games with co-Büchi objectives is solved exactly in the same way. Indeed we have the same kind of objective O since Win_i is a co-Büchi objective if and only if $V^\omega \setminus \text{Win}_i$ is a Büchi objective.

- Consider a Boolean game with Reachability objectives. The intersection of Reachability objectives Win_i (when $p_i = 1$) is a Generalized Reachability objective and the intersection of Safety objectives $V^\omega \setminus \text{Win}_i$ (when $p_i = 0$) is again a Safety objective. The existence of a play with gain profile p in \mathcal{G} can be tested in $\mathcal{O}(2^{|\Pi|}(|V| + |E|))$ time by Proposition 2.2.23 as there are at most $|\Pi|$ Reachability objectives. The case of Boolean games with Safety objectives is solved in the same way.
- Consider a Boolean game with Rabin objectives (with the related families $(G_j^i, R_j^i)_{1 \leq j \leq k_i}$, $i \in \Pi$). In this case, the objective O is the conjunction of Rabin objectives (when $p_i = 1$) and of Streett objectives (when $p_i = 0$), that is, O is a BC Büchi objective $(\phi, (G_j^i, R_j^i)_{1 \leq j \leq k_i}, i \in \Pi)$ such that

$$\phi = \bigwedge_{i|p_i=1} \bigvee_{j=1}^{k_i} (g_j^i \wedge \neg r_j^i) \quad \wedge \quad \bigwedge_{i|p_i=0} \bigwedge_{j=1}^{k_i} (\neg g_j^i \vee r_j^i) \quad (2.2)$$

In this formula, the variable g_j^i (resp. r_j^i) is associated with the set G_j^i

(resp. R_j^i), and ϕ has size $\mathcal{O}(\sum_{i=1}^{|\Pi|} 2 \cdot k_i)$ and has $\sum_{i=1}^{|\Pi|} 2 \cdot k_i$ variables. The announced complexity for deciding the existence of a play with gain profile p follows from Proposition 2.2.23.

The case of Boolean games with Streett objectives is solved in the same way.

- The case of Boolean games with Muller objectives (with the related coloring functions $\Omega_i : V \rightarrow \{1, \dots, d_i\}$ and families $\mathcal{F}_i \subseteq 2^{\Omega_i(V)}$, $i \in \Pi$) is treated as in the previous item. Indeed a play satisfies the Muller objective Win_i if there exists an element F of \mathcal{F}_i such that all colors of F are seen infinitely often along the play while no other color is seen infinitely often. Therefore, as the objective O is a conjunction of Muller objectives and of the complement of Muller objectives, O is a BC Büchi objective $(\phi, (F_c^i)_{c \in \{1, \dots, d_i\}, i \in \Pi})$ described by the following formula ϕ

$$\phi = \bigwedge_{i|p_i=1} \bigvee_{F \in \mathcal{F}_i} \left(\bigwedge_{c \in F} f_c^i \wedge \bigwedge_{c \notin F} \neg f_c^i \right) \wedge \bigwedge_{i|p_i=0} \bigwedge_{F \in \mathcal{F}_i} \left(\bigvee_{c \in F} \neg f_c^i \vee \bigvee_{c \notin F} f_c^i \right) \quad (2.3)$$

In this formula, the variable f_c^i is associated with the subset $F_c^i = \{v \in V \mid \Omega_i(v) = c\}$ of vertices colored by color $c \in \{1, \dots, d_i\}$, $i \in \Pi$. This formula has size $\mathcal{O}(\sum_{i=1}^{|\Pi|} |\mathcal{F}_i| \cdot d_i)$ and has $\sum_{i=1}^{|\Pi|} d_i$ variables.

- It remains to treat the case of Boolean games with Explicit Muller objectives (with the related families $\mathcal{F}_i \subseteq 2^V$, $i \in \Pi$). The approach is a little different in a way to get a polynomial algorithm. By definition, there exists a play ρ with gain profile p if and only if for all i , $F = \text{Inf}(\rho) \in \mathcal{F}_i$ exactly when $p_i = 1$.

If $p \neq (0, \dots, 0)$, such potential sets F can be computed as follows. Initially let O be an empty set. Then for each $F \in \cup_{i \in \Pi} \mathcal{F}_i$, we compute $q \in \{0, 1\}^{|\Pi|}$ such that $q_i = 1$ if and only if $F \in \mathcal{F}_i$, and if $p = q$ we add F to O . Notice that O can be computed in polynomial time. Hence to test the existence of a play with gain profile p in \mathcal{G} , we test the existence of a path in G satisfying the Explicit Muller objective O . This can be done in polynomial time by Proposition 2.2.23.

If $p = (0, \dots, 0)$, there exists a play ρ with gain profile p if and only if no $F \in O' = \cup_{i \in \Pi} \mathcal{F}_i$ is equal to $\text{Inf}(\rho)$, *i.e.*, if and only if there exists a path in G satisfying the complement of the Explicit Muller objective O' . This can be tested in polynomial time by Proposition 2.2.23. \square

2.3 Zero-sum games

Zero-sum games are games with only two players ($|\Pi| = 2$). This kind of games allows to model a situation in which a player wants to ensure a given property whatever the behaviour of the other player. Thus the second player can be viewed as the *opponent* of the first player. In term of objectives that means that the objectives of both players are *antagonistic*. It can be modeled by $\text{Obj}_2 = -\text{Obj}_1$. The term “zero-sum” comes from this point of view, since $\text{Obj}_1 + \text{Obj}_2 = 0$. That means that what Player 1 earns (resp. pays) Player 2 pays (resp. earns) and vice versa. Thus, we can assume, for example, that the objective of Player 1 is given by a cost function Cost_1 that he aims at *minimizing* and the objective of Player 2 is given by a gain function Gain_2 such that $\text{Gain}_2 = \text{Cost}_1$ that he aims at *maximizing*. It is the reason why Player 1 (resp. Player 2) is also called Player Min (resp. Player Max).

Definition 2.3.1 (Zero-sum game). A *zero-sum game* is a game $\mathcal{G} = (A, (\text{Gain}_{\text{Max}}, \text{Cost}_{\text{Min}}))$ where

- The arena is given by $A = (\{\text{Max}, \text{Min}\}, V, E, (V_{\text{Max}}, V_{\text{Min}}))$
- The objective functions are such that $\text{Gain}_{\text{Max}} = \text{Cost}_{\text{Min}}$.

2.3.1 Qualitative zero-sum games

One particular case of zero-sum games are qualitative zero-sum games. In this setting, (i) Player Max has a qualitative objective $\text{Win} \subseteq \text{Plays}$ also given by the corresponding gain function $\text{Gain}_{\text{Max}} : \text{Plays} \rightarrow \{0, 1\}$ (as explained in Section 2.2.1) and (ii) Player Min has a qualitative objective $\text{Plays} \setminus \text{Win}$ also given by a cost function $\text{Cost}_{\text{Min}} : \text{Plays} \rightarrow \{0, 1\}$ (as explained in Re-

mark 2.2.4). The purpose of Player Max is thus to achieve his qualitative objective, he wants to win, while Player Min wants to prevent that. Recall that due to the convention explained in Remark 2.2.4, Player Min wins if and only if Player Max loses. Conversely, Player Min loses if and only if Player Max wins.

Since the game is entirely characterized by A and the objective function of Player Max, we denote it $\mathcal{G} = (A, \text{Gain}_{\text{Max}})$ where $A = (\{\text{Max}, \text{Min}\}, V, E, (V_{\text{Max}}, V_{\text{Min}}))$. Moreover, if Player Max has an O objective with $O \in \{\text{qualitative Reachability, Safety, Büchi, co-Büchi, Parity, Explicit Muller, Muller, Rabin, Streett}\}$, we say that that \mathcal{G} is an O zero-sum game.

Definition 2.3.2 (Qualitative zero-sum game). A *qualitative zero-sum game* \mathcal{G} is a zero-sum game such that Gain_{Max} is a qualitative objective.

For example, if the objective of Player Max is a reachability objective and his target set is $F \subseteq V$, he aims at reaching F while Player Min tries to keep the play outside F . Notice that it means that the objective of Player Min is thus a safety objective.

Winning Strategy In qualitative zero-sum games, since the players are rational they both want to win. We wonder whether a player is able to ensure to win whatever the strategy of the other player. This behavior is formalized through the notion of *winning strategy*.

Definition 2.3.3 (Winning strategy). Given a qualitative zero-sum game \mathcal{G} , a strategy σ_i of Player i , $i \in \{\text{Max}, \text{Min}\}$, is a *winning strategy* for Player i from vertex $v \in V$ if for all strategies σ_{-i} of Player $-i$, the outcome of the strategy profile (σ_i, σ_{-i}) from v is winning for Player i .

If Player i , $i \in \{\text{Max}, \text{Min}\}$, has a winning strategy from $v_0 \in V$, we say that Player i *wins* the initialized game (\mathcal{G}, v_0) . For both players, we can consider the set of vertices from which they win the corresponding initialized

game: this set is called the *winning region*.

Definition 2.3.4 (Winning region). Given a qualitative zero-sum game \mathcal{G} .

- The *winning region* of Player Max is the set $W_{\text{Max}} = \{v \in V \mid \exists \sigma_{\text{Max}} \in \Sigma_{\text{Max}} \text{ st } \forall \sigma_{\text{Min}} \in \Sigma_{\text{Min}} \text{ Gain}_{\text{Max}}(\langle \sigma_{\text{Max}}, \sigma_{\text{Min}} \rangle_v) = 1\}$.
- The *winning region* of Player Min is the set $W_{\text{Min}} = \{v \in V \mid \exists \sigma_{\text{Min}} \in \Sigma_{\text{Min}} \text{ st } \forall \sigma_{\text{Max}} \in \Sigma_{\text{Max}} \text{ Cost}_{\text{Min}}(\langle \sigma_{\text{Max}}, \sigma_{\text{Min}} \rangle_v) = 0\}$.

Definition 2.3.5 (Memoryless winning strategy). Given a qualitative zero-sum game \mathcal{G} , a memoryless strategy σ_i of Player i , $i \in \{\text{Max}, \text{Min}\}$, is a *memoryless winning strategy* for Player i if, for all vertices $v \in W_i$, for all strategies σ_{-i} of Player $-i$, the outcome of the strategy profile (σ_i, σ_{-i}) from v is winning for Player i .

Determinacy Since a player wins if and only if the other player loses, we can easily notice that $W_{\text{Max}} \cap W_{\text{Min}} = \emptyset$. A more complex question is to determine whether, for each vertex in the game graph, there exists a winning strategy for one player from this vertex. That is: “Is the following equality satisfied: $W_{\text{Max}} \cup W_{\text{Min}} = V$?”. A game for which this property is satisfied is called a *determined game*.

Definition 2.3.6 (Determined game). Given a qualitative zero-sum game \mathcal{G} , the game \mathcal{G} is *determined* if the following equality holds:

$$W_{\text{Max}} \cup W_{\text{Min}} = V.$$

Martin [Mar75] proved that every qualitative zero-sum game equipped with a *Borel qualitative objective* – i.e., Win is a Borel set – is determined.

Theorem 2.3.7 ([Mar75]). *Every qualitative zero-sum game equipped with a Borel qualitative objective is determined.*

It follows that if Player Max has one of the classical qualitative objective function defined in Definition 2.2.2, the zero-sum game is determined.

Proposition 2.3.8 ([GTW02]). *Every qualitative Reachability, Safety, Büchi, co-Büchi, Parity, Explicit Muller, Muller, Rabin and Streett zero-sum game is determined.*

Moreover, for qualitative Reachability, Safety, Büchi, co-Büchi, Parity and Rabin zero-sum games, Player Max has a memoryless winning strategy.

2.3.2 Quantitative zero-sum games

Quantitative zero-sum games are the generalization of qualitative zero-sum games to quantitative objective functions. Since the gain function of Player Max is entirely determined by the cost function of Player Min, we choose to take the point of view of Player Min. If $\mathcal{G} = (A, (\text{Gain}_{\text{Max}}, \text{Cost}_{\text{Min}}))$ is a quantitative zero-sum game, we denote it $\mathcal{G} = (A, \text{Cost}_{\text{Min}})$ to ease the notation. Moreover in the rest of this document a O zero-sum game with $O \in \{\text{quantitative Reachability, weighted Reachability, discounted-sum, mean-payoff}\}$ is a quantitative zero-sum game such that Cost_{Min} represents a O objective.

In the general case of quantitative zero-sum games, the notions of “winning” and “winning strategies” are not appropriate. We are rather interested in the *best cost or gain* a player can *ensure* from a given vertex. This is related with the concepts of *values* and *optimal strategies* that we now define.

In order to ease the notation, in a quantitative zero-sum game \mathcal{G} , Player Min (resp. Player Max) is also called Player 1 (resp. Player 2).

Values and Determinacy In a quantitative zero-sum game \mathcal{G} , given a vertex $v \in V$, the *upper value* represents the lowest cost which can be ensured by Player Min from v while the *lowest value* represents the greatest gain which can be ensured by Player Max from v .

Definition 2.3.9 (Lower and upper value). Given a quantitative zero-sum game \mathcal{G} , for all $v \in V$, the *lower value* $\underline{\text{Val}}(v)$ and the *upper value* $\overline{\text{Val}}(v)$ are defined as:

$$\underline{\text{Val}}(v) = \sup_{\sigma_{\text{Max}} \in \Sigma_{\text{Max}}} \inf_{\sigma_{\text{Min}} \in \Sigma_{\text{Min}}} \text{Cost}_{\text{Min}}(\langle \sigma_{\text{Min}}, \sigma_{\text{Max}} \rangle_v)$$

and

$$\overline{\text{Val}}(v) = \inf_{\sigma_{\text{Min}} \in \Sigma_{\text{Min}}} \sup_{\sigma_{\text{Max}} \in \Sigma_{\text{Max}}} \text{Cost}_{\text{Min}}(\langle \sigma_{\text{Min}}, \sigma_{\text{Max}} \rangle_v).$$

For all $v \in V$, the inequality $\underline{\text{Val}}(v) \leq \overline{\text{Val}}(v)$ holds but the reverse inequality is not always true. When for all $v \in V$, the lower and the upper values coincide, the game is said *determined*.

Definition 2.3.10 (Determined game). Given a quantitative zero-sum game \mathcal{G} , if for all $v \in V$, the following equality holds, the game is *determined*:

$$\underline{\text{Val}}(v) = \overline{\text{Val}}(v).$$

In this case, we say that the game \mathcal{G} has a *value* from v . This value $\text{Val}(v)$ is such that $\text{Val}(v) = \underline{\text{Val}}(v) = \overline{\text{Val}}(v)$.

Optimal Strategies Given a quantitative zero-sum game \mathcal{G} , a strategy of Player Min is an *optimal strategy* from a vertex $v \in V$, if following this strategy from v ensures a cost less than or equal to $\overline{\text{Val}}(v)$ to Player Min whatever the strategy of the other player. On the other hand, a strategy of Player Max is an *optimal strategy* from a vertex $v \in V$, if following this strategy from v ensures a gain greater than or equal to $\underline{\text{Val}}(v)$ to Player Max whatever the strategy of Player Min.

Definition 2.3.11 (Optimal strategy). Let \mathcal{G} be a quantitative zero-sum game.

Let σ_{Min}^* be a strategy of Player Min.

- Given $v \in V$, σ_{Min}^* is an *optimal strategy* for Player Min from v if for all strategies σ_{Max} of Player Max:

$$\text{Cost}_{\text{Min}}(\langle \sigma_{\text{Min}}^*, \sigma_{\text{Max}} \rangle_v) \leq \overline{\text{Val}}(v).$$

- σ_{Min}^* is an *optimal strategy* for Player Min if, for all $v \in V$ and for all strategies σ_{Max} of Player Max:

$$\text{Cost}_{\text{Min}}(\langle \sigma_{\text{Min}}^*, \sigma_{\text{Max}} \rangle_v) \leq \overline{\text{Val}}(v).$$

Let σ_{Max}^* be a strategy of Player Max.

- Given $v \in V$, σ_{Max}^* is an *optimal strategy* for Player Max from v if for all strategies σ_{Min} of Player Min:

$$\text{Cost}_{\text{Min}}(\langle \sigma_{\text{Min}}, \sigma_{\text{Max}}^* \rangle_v) \geq \underline{\text{Val}}(v).$$

- σ_{Max}^* is an *optimal strategy* for Player Max if, for all $v \in V$ and for all strategies σ_{Min} of Player Min:

$$\text{Cost}_{\text{Min}}(\langle \sigma_{\text{Min}}, \sigma_{\text{Max}}^* \rangle_v) \geq \underline{\text{Val}}(v).$$

If the game is determined, some links with the notion of optimal strategies can easily be done. The following lemma states two of them.

Lemma 2.3.12. *Given a determined quantitative zero-sum game \mathcal{G} , a vertex $v \in V$ of the game graph, $\sigma_1^* \in \Sigma_{\text{Min}}$ an optimal strategy for Player Min from v and $\sigma_2^* \in \Sigma_{\text{Max}}$ an optimal strategy for Player Max from v . The following assertions hold.*

1. $\text{Cost}_{\text{Min}}(\langle \sigma_1^*, \sigma_2^* \rangle_v) = \text{Val}(v)$;

$$2. \quad \inf_{\sigma_1 \in \Sigma_{\text{Min}}} \text{Cost}_{\text{Min}}(\langle \sigma_1, \sigma_2^* \rangle_v) = \text{Val}(v) = \sup_{\sigma_2 \in \Sigma_{\text{Max}}} \text{Cost}_{\text{Min}}(\langle \sigma_1^*, \sigma_2 \rangle_v).$$

Quantitative Reachability, weighted Reachability, discounted-sum and mean-payoff zero-sum games are determined and both players have memoryless optimal strategies.

Proposition 2.3.13 ([De 13]). *Let \mathcal{G} be a quantitative Reachability, weighed Reachability, discounted-sum or mean-payoff zero-sum game. Then:*

- \mathcal{G} is determined;
- There exist memoryless optimal strategies for both players.

Remark 2.3.14. At the beginning of this section we claim that quantitative zero-sum games are a generalization of qualitative zero-sum games. We now highlight how the concept of values and optimal strategies extend those of winning regions and winning strategies.

Let us first notice that if \mathcal{G} is a qualitative zero-sum game, for all $v \in V$, we have that $v \in W_{\text{Max}}$ if and only if $\underline{\text{Val}}(v) = 1$ and $v \in W_{\text{Min}}$ if and only if $\overline{\text{Val}}(v) = 0$. It follows that $W_{\text{Min}} \cup W_{\text{Max}} = V$ if and only if $\underline{\text{Val}}(v) = \overline{\text{Val}}(v)$ for all $v \in V$. Thus, the notions of determinacy coincide (see Definition 2.3.6 and Definition 2.3.10).

Moreover, given $v \in V$, if σ_{Max}^* (resp. σ_{Min}^*) is a winning strategy for Player Max (resp. Player Min) from v , then σ_{Max}^* (resp. σ_{Min}^*) is an optimal strategy for Player Max (resp. Player Min) from v . In this case, all strategies of Player Min (resp. Player Max) are optimal strategies for him from v .

Finally, in a determined game, given $v \in V$, if $\overline{\text{Val}}(v) = 0$ (then $\text{Val}(v) = 0$) and σ_{Min}^* is an optimal strategy for Player Min from v , then σ_{Min}^* is a winning strategy for Player Min from v . In the same way, given $v \in V$, if $\underline{\text{Val}}(v) = 1$ (then $\text{Val}(v) = 1$) and σ_{Max}^* is an optimal strategy for Player Max from v , then σ_{Max}^* is a winning strategy for Player Max from v .

2.4 Multiplayer games

A (initialized) multiplayer game is a (initialized) game for which we assume that its arena A is such that $|\Pi| \geq 2$. In the setting of multiplayer games, the objectives of the players are not necessarily antagonistic. Thus the solution concepts of winning strategies and optimal strategies are not well suited anymore. It is the reason why we are interested in solution concepts that are *equilibria*: *Nash equilibria*, *subgame perfect equilibria*, *weak subgame perfect equilibria*, ... Notice that with equilibria it is no longer question of optimality but rather of stability. Moreover, there exist games in which there exist (at least) two equilibria: one which is a “good” equilibrium for both players and one which is a “bad” equilibrium for both players. We explicit an example in Example 2.4.2 with a multiplayer quantitative Reachability games.

In this section, we give all the definitions by assuming that the objective function of each player is a cost function (see Remark 2.2.1).⁴ We define the concepts of *Nash equilibria* (Section 2.4.1), *subgame perfect equilibria* (Section 2.4.2) and *weak subgame equilibria* (Section 2.4.3). For this latter notion, we also need to introduce the notion of *weak Nash equilibrium*.

2.4.1 Nash equilibrium

The famous notion of *Nash equilibrium* [Nas50] is one of the most studied solution concept. A strategy profile is a Nash equilibrium if no player has an incentive to unilaterally change his strategy. It is like a contract between the players in which no player can improve alone his cost by deviating from his strategy.

Definition 2.4.1 (Nash equilibrium). Given an arena A and an initialized multiplayer game $(\mathcal{G}, v_0) = (A, \text{Cost})$, the strategy profile σ is a *Nash equilibrium* (NE) in (\mathcal{G}, v_0) if for all $i \in \Pi$ and all $\sigma'_i \in \Sigma_i$,

$$\text{Cost}_i(\langle \sigma \rangle_{v_0}) \leq \text{Cost}_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}).$$

⁴ To obtain the definitions for a gain function, it suffices to replace the \leq (resp. $<$) symbols by \geq (resp. $>$) symbols and Cost_i by Gain_i .

A related notion to Nash equilibria is the notion of *profitable deviation*. Given an initialized game (\mathcal{G}, v_0) and a strategy profile σ in this game, we have that σ'_i is a *profitable deviation* for Player i w.r.t. σ in (\mathcal{G}, v_0) if

$$\text{Cost}_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}) < \text{Cost}_i(\langle \sigma \rangle_{v_0}).$$

Thus a Nash equilibrium is a strategy profile such that no player has a profitable deviation.

Example 2.4.2. Let us consider the initialized quantitative Reachability game depicted in Figure 2.5. In this game the target set of both player is $F_1 = F_2 = \{v_3\}$.




Figure 2.5: Initialized game (\mathcal{G}, v_0) with quantitative objectives and different Nash equilibria.

- The memoryless strategy profile σ depicted by double arrows in Figure 2.5 is not a Nash equilibrium. Formally,

$$\sigma_1(v) = \begin{cases} v_0 & \text{if } v \in \{v_1, v_3, v_4\} \\ v_4 & \text{if } v = v_2 \end{cases} \quad \text{and} \quad \sigma_2(v) = v_2 \text{ when } v = v_0.$$

The outcome of this strategy profile is $(v_0v_2v_4)^\omega$ and thus its cost profile is $(+\infty, +\infty)$ since no player visits his target set. Clearly the deviating strategy σ'_1 of Player 1 defined as follows $\sigma'_1(v) = v_3$ if $v = v_2$ and $\sigma'_1(v) = \sigma_1(v)$ otherwise is a profitable deviation for Player 1 w.r.t. σ in (\mathcal{G}, v_0) . Indeed if both players follow the strategy profile (σ'_1, σ_2) the cost of Player 1 is equal to 2.

- In fact the strategy profile (σ'_1, σ_2) is a Nash equilibrium with outcome $(v_0v_2v_3)^\omega$ and with cost profile $(2, 2)$. In particular, by following this equilibrium both players visit their target set.

- Another Nash equilibrium in this game is the strategy profile (σ_1, σ'_2) where σ'_2 is a memoryless strategy such that $\sigma'_2(v) = v_1$ when $v = v_0$. The outcome of this equilibrium is $(v_0v_1)^\omega$ with cost profile $(+\infty, +\infty)$. In particular, by following this equilibrium no player visit his target set.

Let us now mention some results about the existence of Nash equilibria. Other various results about the existence of Nash equilibria and the needed amount of memory may be found in [De 13].

Theorem 2.4.3 ([De 13, Theorem 4.3.2]). *In every initialized multiplayer weighted Reachability game, there exists a finite-memory Nash equilibrium.*

Theorem 2.4.4 ([De 13, Proposition 4.4.11]). *In every initialized multiplayer game where*

- *each cost function is prefix-independent and;*
- *for each player i , the coalitional game^a \mathcal{G}_i is determined and both players have optimal strategies,*

there exists a Nash equilibrium.

^aThe notion of coalitional game is defined in Section 4.1.

Even if the existence of a Nash equilibrium is guaranteed in several kinds of multiplayer games, there are games in which no Nash equilibrium exists. We illustrate this phenomenon in the following example.

Example 2.4.5. Let us consider the multiplayer weighted Reachability game (\mathcal{G}, v_0) depicted in Figure 2.6. Notice that in this example we allow **negative weights** on edges. In this example, there are two players: Player 1 and Player 2. The only vertex owned by Player 1 (resp. Player 2) is v_0 (resp. v_1) and its target set is $\{v_1\}$ (resp. $\{v_1\}$).

Let us show that there is no NE in (\mathcal{G}, v_0) . In this game Player 2 has only one strategy σ_2 : always cycling in v_1 . Thus, we only consider the different strategies of Player 1. There are two kinds of strategies for Player 1: one

strategy that consists in cycling infinitely often in v_0 and one family of strategies which consist in cycling a given number of times in v_0 and then going in v_1 . Formally, the different possible strategies are $\sigma_1^\infty : \text{Hist}_1(v_0) \rightarrow V$ such that $\sigma_1^\infty(hv_0) = v_0$ for all $hv_0 \in \text{Hist}_1(v_0)$ and, for all $k \in \mathbb{N}_0$, $\sigma_1^k : \text{Hist}_1(v_0) \rightarrow V$ such that $\sigma^k(hv_0) = v_1$ if the number of occurrences of v_0 along hv_0 is equal to k and $\sigma^k(hv_0) = v_0$ otherwise.

Obviously, the strategy profile $(\sigma_1^\infty, \sigma_2)$ is not an NE since Player 1 does not reach his target set and he only has to choose to go in v_1 in order to obtain a cost less than $+\infty$.

Let us prove that for all $k \in \mathbb{N}_0$, the strategy profile (σ_1^k, σ_2) is not an NE. Let $k \in \mathbb{N}_0$, we have that $\langle \sigma_1^k, \sigma_2 \rangle_{v_0} = v_0^k v_1^\omega$ (where v_0^k is the k times concatenation of vertex v_0). Hence $\text{Cost}_1(\langle \sigma_1^k, \sigma_2 \rangle_{v_0}) = (-1) \cdot (k - 1) + 1$. It follows that the strategy σ_1^{k+1} is a profitable deviation for Player 1 w.r.t. σ_1^k .




Figure 2.6: A multiplayer game without Nash equilibrium.

To conclude this section, let us mention that, unfortunately Nash equilibria have non desirable drawbacks: they are subject to *uncredible threats*: decisions in subgames that are irrational and used to threaten the other players and force them to follow a given behavior. Example 2.4.2 illustrates this phenomenon. If we consider the Nash equilibrium (σ_1, σ_2') , we observe that in v_2 Player 1 does not have a rational behaviour. Actually, in v_2 he should prefer to go in v_3 where he visits his target instead to choose to go in v_4 .

Example 2.4.2 illustrates two weaknesses of NEs: (i) some equilibria may appear not relevant and (ii) NEs do not take into account the sequential aspect of games played on graphs and so they do not handle irrational behaviours in subgames. To avoid this latter problem, the concept of *subgame perfect equilibria* has been proposed. We define formally this solution concept and the notion of subgames in the next section. We come back to (i) in Part III.

2.4.2 Subgame perfect equilibrium

To prevent from uncredible threats and take into account the sequential aspect of games played on graphs, a more suitable notion of equilibria than NEs is the notion of *subgame perfect equilibrium* [Sel65, OR94]. A strategy profile is a subgame perfect equilibrium if it is a Nash equilibrium in each subgame.

Before formally defining the concept of subgame perfect equilibrium, we need to introduce the notion of *subgame*. Given an initialized game $(\mathcal{G}, v_0) = (A, \text{Cost})$ and a history $hv \in \text{Hist}(v_0)$, the initialized game $(\mathcal{G}_{|h}, v)$ is called a *subgame* of (\mathcal{G}, v_0) and is such that $\mathcal{G}_{|h} = (A, \text{Cost}_{|h})$ and $\text{Cost}_{i|h}(\rho) = \text{Cost}_i(h\rho)$ for all $i \in \Pi$ and $h\rho \in \text{Plays}$. Notice that (\mathcal{G}, v_0) is a subgame of itself.

Moreover if σ_i is a strategy of Player i in (\mathcal{G}, v_0) , then $\sigma_{i|h}$ denotes the strategy in $(\mathcal{G}_{|h}, v)$ such that for all histories $h' \in \text{Hist}_i(v)$, $\sigma_{i|h}(h') = \sigma_i(hh')$. Similarly, from a strategy profile σ in (\mathcal{G}, v_0) , we derive the strategy profile $\sigma_{|h}$ in $(\mathcal{G}_{|h}, v)$. Let (\mathcal{G}, v_0) be an initialized multiplayer game, following this formalism, a strategy profile σ is a *subgame perfect equilibrium* in (\mathcal{G}, v_0) if for all $hv \in \text{Hist}(v_0)$, $\sigma_{|h}$ is an NE in $(\mathcal{G}_{|h}, v)$.

Definition 2.4.6 (Subgame perfect equilibrium). Given an arena A and an initialized multiplayer game $(\mathcal{G}, v_0) = (A, \text{Cost})$, the strategy profile σ is a *subgame perfect equilibrium* (SPE) in (\mathcal{G}, v_0) if for all histories $hv \in \text{Hist}(v_0)$, $\sigma_{|h}$ is a Nash equilibrium in $(\mathcal{G}_{|h}, v)$.

Remark 2.4.7. A subgame perfect equilibrium σ in (\mathcal{G}, v_0) is in particular a Nash equilibrium in (\mathcal{G}, v_0) .

Example 2.4.8. Let us come back to the initialized multiplayer game of Example 2.4.2.

- The strategy profile (σ_1, σ'_2) is not an SPE in (\mathcal{G}, v_0) . The memoryless strategy σ'_1 is a profitable deviation for Player 1 w.r.t. $(\sigma_{1|v_0}, \sigma'_{2|v_0})$ in

$(\mathcal{G}_{|v_0}, v_2)$. Indeed, we have that

$$\begin{aligned} \text{QR}_{1|v_0}(\langle \sigma'_1, \sigma'_{2|v_0} \rangle_{v_2}) &= \text{QR}_1(v_0 \langle \sigma'_1, \sigma'_{2|v_0} \rangle_{v_2}) \\ &= \text{QR}_1(v_0 v_2 v_3 (v_0 v_1)^\omega) = 2 \end{aligned}$$

while

$$\begin{aligned} \text{QR}_{1|v_0}(\langle \sigma_{1|v_0}, \sigma'_{2|v_0} \rangle_{v_2}) &= \text{QR}_1(v_0 \langle \sigma_{1|v_0}, \sigma'_{2|v_0} \rangle_{v_2}) \\ &= \text{QR}_1(v_0 v_2 v_4 (v_0 v_1)^\omega) = +\infty. \end{aligned}$$

Let us recall that (σ_1, σ'_2) is an NE in (\mathcal{G}, v_0) with an uncredible threat.

- The strategy profile (σ'_1, σ_2) is an SPE in (\mathcal{G}, v_0) .

Let us now mention some results about the existence of subgame perfect equilibria.

Theorem 2.4.9 ([De 13, Theorem 6.2.1]). *In every initialized multiplayer weighted Reachability game, there exists a subgame perfect equilibrium.*

A more general result is the following.


Theorem 2.4.10 ([FL83, Har85]). *Given an initialized multiplayer game $(\mathcal{G}, v_0) = (A, \text{Cost})$, if the cost functions Cost_i , for each $i \in \Pi$, are continuous and real-valued, then there exists a subgame perfect equilibrium in (\mathcal{G}, v_0) .*

In the case of qualitative objectives, the following result holds.

Theorem 2.4.11 ([GU08]). *Let (\mathcal{G}, v_0) be an initialized multiplayer game with qualitative Borel objectives, there always exists a subgame perfect equilibrium in (\mathcal{G}, v_0) .*

Even if subgame perfect equilibria prevent uncredible threats, some strange phenomenon may also appear with subgame perfect equilibria. We explicit that in Example 2.4.12.

Example 2.4.12. We consider the multiplayer initialized quantitative Reachability game (\mathcal{G}, v_0) illustrates in the following figure and such that the target sets are $F_1 = \{v_3\}$ and $F_2 = \{v_3, v_5\}$. Notice that there is one additional vertex compared to our running example.



Let us consider the memoryless strategy profile σ depicted by the double arrows in this figure. We can prove that σ is an SPE in (\mathcal{G}, v_0) . The outcome of this strategy profile is $v_0(v_1v_5)^\omega$ and thus Player 1 does not visit his target set whereas Player 2 does. We informally provide the ideas of why σ is an SPE.

Player 2 visits his target in two steps and he cannot do better, this is the value of a shortest path between v_0 and one of the vertices of F_2 . It means that Player 2 has no profitable deviation. This argument holds in all subgames, indeed in each subgame Player 2 always visits his target set as soon as possible. It remains to show that Player 1 does not have any profitable deviation in any subgame. In the main game (\mathcal{G}, v_0) , Player 1 cannot do anything because even if he changes his strategy Player 2 forces the play to stay in the left part of the arena. In a subgame $(\mathcal{G}_{|h}, v)$ with $v \in \{v_0, v_1, v_4, v_5\}$ for some $hv \in \text{Hist}(v_0)$ this argument remains true. In a subgame $(\mathcal{G}_{|h}, v_3)$ for some $hv_3 \in \text{Hist}(v_0)$ Player 1 has obviously no incentive to deviate from his strategy and in the subgame $(\mathcal{G}_{|h}, v_2)$ for some $hv_2 \in \text{Hist}(v_0)$, Player 2 chooses to go directly in v_3 (an element of his target set), thus he cannot make a better choice.

We find the SPE σ a little bit counter-intuitive because if Player 1 is really clever and selfish he should disagree to follow this contract. Indeed, he has a way to force Player 2 to visit both F_1 and F_2 . He only has to always choose to go in v_0 from v_1 and follows σ_1 from the other vertices. In this way, if Player 2 wants to visit his target set he has to choose to go in v_2 at a moment. Hopefully an SPE with this behavior of Player 1 exists, Player 2 has to play following the strategy σ'_2 defined as follows: for all $v \in V_2$, $\sigma'_2(v) = v_2$. In this

way the outcome of the SPE is $(v_0v_2v_3)^\omega$ and both players reach their target set with a cost equal to 2.

2.4.3 Weak subgame perfect equilibrium

While a subgame perfect equilibrium does not always exist, Kuhn’s theorem asserts that every *finite extensive game with perfect information* has a subgame perfect equilibrium (see *e.g.*, [OR94]). The proof of this theorem is based on the central notion of *one deviation property*. Motivated by the purpose of extending this concept, authors in [BBMR15] introduced the notion of *weak Nash equilibrium* and *weak subgame perfect equilibrium*.

Before defining this kinds of equilibria we introduce the notion of *deviation step*, *finitely deviating strategy* and *one-shot deviating strategy*. This section is inspired by [BBMR15].

Deviation Step, Finitely Deviating Strategy, One-shot Deviating Strategy

When we fix a strategy profile σ in an initialized multiplayer game and a strategy σ'_i of Player i , we can look at the number of times that Player i changes his behavior along the outcome of (σ'_i, σ_{-i}) with respect to σ_i . The prefixes of the outcome of (σ'_i, σ_{-i}) where this phenomenon occurs are called *deviation steps* from σ .

Definition 2.4.13 (Deviation step). Let (\mathcal{G}, v_0) be an initialized multiplayer game, σ be a strategy profile and σ'_i be a strategy of Player i . We say that σ'_i has an *hv-deviation step* from σ for some $hv \in \text{Hist}_i(v_0)$ if

$$hv < \langle \sigma'_i, \sigma_{-i} \rangle_{v_0} \text{ and } \sigma'_i(hv) \neq \sigma_i(hv).$$

If the number of deviation steps of σ'_i is finite this deviating strategy is called a *finitely deviating strategy* while if σ_i and σ'_i differs only in v_0 , σ'_i is called a *one-shot deviating strategy*.

Definition 2.4.14 (Finitely/one-shot deviating strategy). Let (\mathcal{G}, v_0) be an initialized multiplayer game and σ be a strategy profile in (\mathcal{G}, v_0) .

- A strategy σ'_i in (\mathcal{G}, v_0) is a *finitely deviating strategy* from σ if it has a finite number of deviation steps from σ .
- A strategy σ'_i in (\mathcal{G}, v_0) is a *one-shot deviating strategy* from σ if it has a v_0 -deviation step from σ and it is the only one.

(Very) Weak Nash Equilibria and (Very) Weak Subgame Perfect Equilibria

A Nash equilibrium is a strategy profile which is resistant against any deviating strategy of any player. In particular, the considered strategies may have an infinite number of deviation steps. With *(very) weak Nash equilibria* it is not the case, the strategy profile must only be resistant against finitely (resp. one-shot) deviating strategies of any player.

Definition 2.4.15 (Weak Nash equilibrium). Let (\mathcal{G}, v_0) be an initialized multiplayer game, the strategy profile σ is a *weak Nash equilibrium* (weak NE) in (\mathcal{G}, v_0) if for all $i \in \Pi$ and all finitely deviating strategy σ'_i from σ , we have:

$$\text{Cost}_i(\langle \sigma \rangle_{v_0}) \leq \text{Cost}_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}).$$

Definition 2.4.16 (Very weak Nash equilibrium). Let (\mathcal{G}, v_0) be an initialized multiplayer game, the strategy profile σ is a *very weak Nash equilibrium* (very weak NE) in (\mathcal{G}, v_0) if for all $i \in \Pi$ and all one-shot deviating strategy σ'_i from σ , we have:

$$\text{Cost}_i(\langle \sigma \rangle_{v_0}) \leq \text{Cost}_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}).$$

As for Nash equilibria and subgame perfect equilibria, a (very) weak NE


may be a (very) weak NE in each subgame of the initialized multiplayer game (\mathcal{G}, v_0) .

Definition 2.4.17 ((Very) weak subgame perfect equilibrium). Let (\mathcal{G}, v_0) be an initialized multiplayer game, the strategy profile σ is a *weak (resp. very weak) subgame perfect equilibrium* (weak (resp. very weak) SPE) in (\mathcal{G}, v_0) if for all history $hv \in \text{Hist}(v_0)$, $\sigma|_h$ is a weak (resp. very weak) Nash equilibrium in $(\mathcal{G}|_h, v)$.

The following proposition asserts that both notions are equivalent.

Proposition 2.4.18 ([BBMR15]). Let (\mathcal{G}, v_0) be an initialized multiplayer game and σ be strategy profile, σ is a weak SPE in (\mathcal{G}, v_0) if and only if σ is a very weak SPE in (\mathcal{G}, v_0) .

Example 2.4.19. We consider the initialized multiplayer Büchi game $(\mathcal{G}, v_0) = (\text{A}, \text{Buchi})$ such that its arena is depicted by the following figure and $\text{Buchi} = (\text{Buchi}_1, \text{Buchi}_2)$. In this example the target set of Player 1 (resp. Player 2) is $F_1 = \{v_4\}$ (resp. $F_2 = \{v_5\}$). Thus in this game Player 1 (resp. Player 2) wants to reach infinitely often vertex v_4 (resp. v_5).




The memoryless strategy profile (σ_1, σ_2) depicted by double arrows is not a weak NE in (\mathcal{G}, v_0) and in particular not a weak SPE (\mathcal{G}, v_0) . We have that

σ_1 is defined as follows, for all $v \in V_1$, $\sigma_1(v) = \begin{cases} v_5 & \text{if } v = v_1 \\ v_3 & \text{if } v = v_2 \\ v_0 & \text{if } v \in \{v_3, v_4\} \end{cases}$ and σ_2 is

defined as follows, for all $v \in V_2$ (that is for $v = v_0$), $\sigma_2(v) = v_2$. We have that $\langle \sigma_1, \sigma_2 \rangle_{v_0} = (v_0 v_2 v_3)^\omega$ and $\text{Buchi}(\langle \sigma_1, \sigma_2 \rangle_{v_0}) = (0, 0)$.

We can easily see that Player 2 has a finitely deviating strategy τ_2 from (σ_1, σ_2) such that τ_2 is a profitable deviation for Player 2 w.r.t. (σ_1, σ_2) . We define τ_2 as follows, for all $hv \in \text{Hist}_2(v_0)$, in particular $\text{Last}(hv) = v_0$, $\tau_2(hv) = \begin{cases} v_1 & \text{if } h = \epsilon \\ \sigma_2(v) & \text{otherwise} \end{cases}$. We have that τ_2 is a one-shot deviating strategy from (σ_1, σ_2) since it only differs from σ_2 in v_0 . Moreover, $\langle \sigma_1, \tau_2 \rangle_{v_0} = v_0(v_1v_5)^\omega$ and $\text{Gain}(\langle \sigma_1, \tau_2 \rangle_{v_0}) = (0, 1)$. Thus τ_2 is a profitable deviation for Player 2 from (σ_1, σ_2) . It concludes the proof that (σ_1, σ_2) is not a weak NE.

On the contrary, if we consider the same game with a different memoryless strategy profile (σ'_1, σ_2) depicted by double arrows in the following picture, this strategy profile is a weak NE. Notice that the only difference is that from v_1 Player 1 always chooses to go in v_0 instead of v_5 .



That is σ'_1 is defined as follows, for all $v \in V_1$, $\sigma_1(v) = \begin{cases} v_0 & \text{if } v = v_1 \\ v_3 & \text{if } v = v_2 \\ v_0 & \text{if } v \in \{v_3, v_4\} \end{cases}$.

Let us assume that τ_1 is a profitable deviation of Player 1 from (σ'_1, σ_2) . The only way for Player 1 to obtain a gain of 1 is to change his choice from v_1 an infinite number of times by choosing to go in v_4 instead of v_3 . In this way, the outcome of the strategy profile (τ_1, σ_2) visits v_4 infinitely often. But it means that $\{v_0v_2(v_3v_0v_2)^k \mid k \in \mathbb{N} \text{ and } \tau_1(v_0v_2(v_3v_0v_2)^k) \neq \sigma'_1(v_0v_2(v_3v_0v_2)^k)\}$ is infinite and so τ_1 cannot be a finitely deviating strategy of Player 1 from (σ'_1, σ_2) .

For Player 2, the main idea is that since Player 1 always chooses to go in v_0 from v_1 , whatever the strategy τ'_2 of Player 2, the outcome $\langle \sigma'_1, \tau'_2 \rangle_{v_0}$ never visits v_5 . Thus Player 2 has no profitable deviation w.r.t. (σ'_1, σ_2) .

Notice that (σ'_1, σ_2) is also a weak SPE.

Let us now mention some results about existence of weak subgame perfect

equilibria. We first mention that as an SPE is a weak SPE, all results provided in the previous section about the existence of SPEs also hold for weak SPEs.

Theorem 2.4.20 ([BRPR17]). *Let $(\mathcal{G}, v_0) = (A, \text{Cost})$ be an initialized multiplayer game such that*

- *either each cost function Cost_i , $i \in \Pi$, is prefix-independent;*
- *or each Cost_i , $i \in \Pi$, has a finite range.*

Then there exists a weak SPE in (\mathcal{G}, v_0) .

Remark 2.4.21. Sometimes the notion of finitely deviating strategy is defined slightly differently (see [BRPR17] for example). With this other definition we do not only count the number of deviations along the outcome. Given an initialized multiplayer game (\mathcal{G}, v_0) , a strategy profile σ and a strategy σ'_i of Player i both in this game, we consider the number of histories (without any restriction) such that σ'_i differs from σ_i . That is σ'_i is a finitely deviating strategy from σ if the set $\{h \in \text{Hist}_i(v_0) \mid \sigma_i(h) \neq \sigma'_i(h)\}$ is finite.

Obviously these two notions are not equivalent but when we deal with weak NEs and weak SPEs we can use interchangeably both definitions. Indeed, each strategy σ'_i such that $\{h \in \text{Hist}_i(v_0) \mid \sigma_i(h) \neq \sigma'_i(h)\}$ is finite is a finitely deviating strategy from σ in the sense of Definition 2.4.14. Moreover, if we have a finitely deviating strategy σ'_i from σ in the sense of Definition 2.4.14, we can build a strategy τ'_i such that: (i) $\langle \sigma'_i, \sigma_{-i} \rangle_{v_0} = \langle \tau'_i, \sigma_{-i} \rangle_{v_0}$ and thus $\text{Cost}_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}) = \text{Cost}_i(\langle \tau'_i, \sigma_{-i} \rangle_{v_0})$ and (ii) $\{h \in \text{Hist}_i(v_0) \mid \sigma_i(h) \neq \tau'_i(h)\}$ is finite. The strategy τ'_i has only to follow σ'_i along the outcome $\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$ and to mimic σ_i otherwise.

Equivalence with the notion of SPE

Even if it is well known that there exist games with weak SPEs but no SPE (see *e.g.*, [BBMR15]), in some classes of games the notions of very weak SPE, weak SPE and SPE are equivalent. For example this is the case with games such that the cost functions are continuous.

Proposition 2.4.22 ([BBMR15]). *Let $(\mathcal{G}, v_0) = (A, \text{Cost})$ be an initialized multiplayer game and let σ be a strategy profile in this game. If for all $i \in \Pi$, Cost_i is a continuous, then: σ is a weak SPE in (\mathcal{G}, v_0) if and only if σ is a very weak SPE in (\mathcal{G}, v_0) if and only if σ is an SPE in (\mathcal{G}, v_0) .*

Corollary 2.4.23. *Let $(\mathcal{G}, v_0) = (A, \text{Cost})$ be an initialized multiplayer quantitative or weighted Reachability game, let σ be a strategy profile in this game, then: σ is a weak SPE in (\mathcal{G}, v_0) if and only if σ is a very weak SPE in (\mathcal{G}, v_0) if and only if σ is an SPE in (\mathcal{G}, v_0) .*

As shown by Example 2.4.19, a weak Nash equilibrium has to be seen as a contract where no one has an incentive to finitely deviate alone. In such a contract a player may have a profitable deviation which needs infinitely many changes. Notice that the winning strategy of Player 1 σ_1^* in Example 2.4.19 is a profitable deviation for Player 1 w.r.t. (σ_1', σ_2) . Although playing σ_1^* is optimal for Player 1 this cannot be achieved from σ_1' by finitely deviating. As in the classical setting of game theory, the study of equilibria concerns stability rather than optimality. In this spirit, weak Nash equilibria can be used to model stability under finitely deviating strategies.

Beyond this weakness of weak SPEs, Proposition 2.4.22 shows a real strength of weak SPEs. Indeed, it means that for games with continuous objective functions if we want to study SPEs in these games we only have to study weak SPEs which is a more simple solution concept since they only have to be resistant against finitely deviating strategies. This property is one of the keys of some of our results in Part III.

Part II:

CHARACTERIZATION OF
EQUILIBRIA OUTCOMES

Looking for an equilibrium in an initialized multiplayer game (\mathcal{G}, v_0) amounts to (i) fixing a strategy for each player—a function which assigns a next vertex to each history such that its last vertex is a vertex owned by this player—and (ii) verifying that this strategy profile is resistant against a class of deviating strategies either in (\mathcal{G}, v_0) or in each of its subgames $(\mathcal{G}_{|h}, v)$ (with $hv \in \text{Hist}(v_0)$). It is a huge amount of work especially if we are only interested in the equilibrium outcome. For example, one might ask if there exists an equilibrium with a given cost profile (see Part III). In this case we only want to decide if there exists a play such that (i) this play can be an equilibrium outcome and (ii) the cost profile of this play is the one we desire.

In order to consider this kind of questions, we aim at finding a way to rephrase the existence of an equilibrium in terms of equilibrium outcome instead of strategy profile.

“Is there a play in the game which satisfies some “good” properties, which characterize equilibria outcomes?”

Thus when these “good” properties are well defined for a particular kind of equilibrium (NE, SPE or weak SPE), the set of plays which satisfy these properties is exactly set of outcomes of this particular kind of equilibrium. Different ways to characterize the set of outcomes of equilibria may be considered. Our

approach relies on the notion of λ -consistent play.

λ -Consistent Plays Our purpose with the notion of λ -consistent plays is to impose some constraints on the plays which are equilibrium outcomes. These constraints concern the cost (resp. gain) of each player along a given play. We define a *labeling function* $\lambda : V \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ which labels the vertices of the game graph with values. A play is a λ -consistent play if from each vertex v along this play, let us assume that $v \in V_i$, the cost (or gain) of Player i from v is less than or equal to (resp. greater than or equal to) the value $\lambda(v)$ of this vertex.

For example, let us consider an initialized Boolean game (\mathcal{G}, v_0) and assume that we have built a labeling function $\lambda : V \rightarrow \{0, 1\}$. Let $\rho = \rho_0\rho_1\dots$ be a play in (\mathcal{G}, v_0) , if ρ_n is a vertex owned by Player i and $\lambda(\rho_n)$ is equal to 0, then the labeling function λ imposes no constraint on ρ . On the other hand, if $\lambda(\rho_n) = 1$ it means that from ρ_n Player i should achieve his qualitative objective along ρ .

It is important to notice that the notion of λ -consistent plays does not allow us to characterize the outcomes of equilibria in games equipped with any cost (resp. gain) functions. We explicit the conditions that the cost (resp. gain) have to satisfy in Chapter 6 and Chapter 7. Moreover, it seems natural that the same labeling function cannot characterize both NEs outcomes, weak SPEs outcomes and SPEs outcomes.

The formal definition of λ -consistent play is provided in Chapter 5.

Characterization of Nash Equilibria Outcomes The philosophy of the characterization of Nash equilibria outcomes that we consider is the same as the one of the well-known Folk theorem for infinitely repeated game [OR94].

Let (\mathcal{G}, v_0) be an initialized multiplayer game. We first study what should be the behaviour of a player if he plays against the coalition of the other players. In order to do so, for each player $i \in \Pi$, we consider the zero-sum game \mathcal{G}_i in which Player i is one of the players and the coalition of the other players $-i$ becomes the other player. This zero-sum game is called the *coalitional game* of Player i . The value $\text{Val}_i(v)$ of a vertex v in this game represents the cost

(resp. gain) that Player i can ensure from v in \mathcal{G} if all the other players play against him.

From the values of the vertices in all coalitional games $(\mathcal{G}_i)_{i \in \Pi}$, the labeling function $\lambda = \text{Val}^* : V \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ is defined as follows: if $v \in V_i$ then $\text{Val}^*(v) = \text{Val}_i(v)$. We claim that the labeling function Val^* allows to characterize Nash equilibria outcomes in multiplayer games equipped with strongly prefix-linear cost (resp. gain) functions and such that the coalitional games are determined with memoryless optimal strategies for both players.

In few words, if $\rho \in \text{Plays}(v_0)$ is not a Val^* -consistent play that means that there exists a vertex v along ρ , belonging to some player $i \in \Pi$, such that the cost (resp. gain) of Player i from v is worse than what he could ensure from v even if he played against the coalition of the other players. That means that he can do better from v and so he has a profitable deviation. On the other hand, if ρ is Val^* -consistent, then following the play ρ until some player deviates and punishing the deviator by forming a coalition against him once a deviation occurs leads to a Nash equilibrium.

All the details are given in Chapter 6.

Characterization of Weak Subgame Perfect Equilibria Outcomes

Since some NEs are not weak SPEs, we have to find another labeling function to characterize weak SPEs outcomes. The labeling function λ that characterizes weak SPEs outcomes is obtained from an iterative procedure. We begin with an initial labeling function that imposes no constraint on the plays and then, by iterating an operator, we reinforce the constraints step after step, up to obtaining a fixpoint which is the required function λ .

Roughly speaking, at the initial step λ^0 is such that all plays are λ^0 -consistent. Then, if we assume that λ^k is known for some $k \in \mathbb{N}$, for each $v \in V$, by assuming that $v \in V_i$ for some $i \in \Pi$, the value $\lambda^{k+1}(v)$ represents the best cost (resp. gain) that Player i can ensure for himself from v with a “one-shot” choice. Moreover, this choice takes into account that only plays which are λ^k -consistent from each $v' \in \text{Succ}(v)$, have to be considered.

If this procedure reaches a fixpoint λ^* and with some additional conditions on the cost functions (resp. gain functions) involved in the multiplayer game,

we are able to build a weak SPE from the sets of λ^* -consistent plays.

All details are given in Section 7.3.

In Section 7.4, we explain how we obtain weak SPEs outcomes characterizations from the general characterizations provided in Section 7.3. In this way, we obtain characterizations for Boolean games with prefix-independent gain functions, for qualitative and quantitative Reachability games and for Safety games.

Characterization of Subgame Perfect Equilibria Outcomes Since the notion of weak SPE and SPE is equivalent for multiplayer games with continuous cost functions (resp. gain functions). It is possible to obtain a characterization of SPEs outcomes for some particular multiplayer games. We discuss this in Section 7.5.

Organization of the part In all this part except in Chapter 4, in the definitions and results we assume that all the objective functions are cost functions. The corresponding definitions and results for gain functions may be obtained by replacing the \leq (resp. $<$) symbols by \geq (resp. $>$) symbols and vice versa. In the same way, \min (resp. \inf) have to be replaced by \max (resp. \sup) and vice versa. Notice that we sometimes consider gain functions in our examples.

In Chapter 4, we begin by introducing two particular kinds of games which are needed to obtain our different characterizations: the *coalitional games* and the *extended reachability games*. In Chapter 5, we define the notion of λ -consistent play for a given labeling function λ . This notion is at the heart of the characterizations that we provide in the other chapters of this part. In Chapter 6, we provide an NE outcome characterization based on the values of the players in the different coalitional games. In Chapter 7, we first give a characterizations of weak SPEs outcomes for multiplayer games such that the players' objective functions satisfy some conditions and then we explicit that these characterizations also work to characterize SPEs outcomes if the objective functions are continuous in addition to satisfying the previous conditions.

CHAPTER 4

COALITIONAL GAMES AND EXTENDED REACHABILITY GAMES

In order to properly define the labeling functions which exactly characterize NEs outcomes and weak SPEs outcomes, we need to refer to some related games: *coalitional games* and *extended (Reachability) games*.

4.1 Coalitional games

Given a multiplayer game \mathcal{G} , for each player $i \in \Pi$, we study the game in which Player i plays against the coalition of the other players. This game \mathcal{G}_i is a zero-sum game and is called the *coalitional game* for Player i .

Definition 4.1.1 (Coalitional game). Let $\mathcal{G} = (A, \text{Obj})$ be a multiplayer game. For each $i \in \Pi$, we build the zero-sum game $\mathcal{G}_i = (A_i, (\text{Cost}_{\text{Min}}, \text{Gain}_{\text{Max}}))$ such that:

Gain function: If the objective function of Player i is a gain function Gain_i , then Player Max is Player i and Player Min is the coalition of the other players $-i$. Thus, in particular:

- the arena $A_i = (V, E, (V_{\text{Min}}, V_{\text{Max}}))$ is such that $V_{\text{Max}} = V_i$ and

$$V_{\text{Min}} = \bigcup_{j \in \Pi \setminus \{i\}} V_j;$$

- $\text{Gain}_{\text{Max}} = \text{Cost}_{\text{Min}} = \text{Gain}_i$.

To ease the notation, we also write $\mathcal{G}_i = (A_i, \text{Gain}_i)$ to denote this coalitional game.

Cost function: If the objective function of Player i is a cost function Cost_i , then Player Min is Player i and Player Max is the coalition of the other players $-i$. Thus, in particular:

- the arena $A_i = (V, E, (V_{\text{Min}}, V_{\text{Max}}))$ is such that $V_{\text{Min}} = V_i$ and $V_{\text{Max}} = \bigcup_{j \in \Pi \setminus \{i\}} V_j$;
- $\text{Cost}_{\text{Min}} = \text{Gain}_{\text{Max}} = \text{Cost}_i$.

To ease the notation, we also write $\mathcal{G}_i = (A_i, \text{Cost}_i)$ to denote this coalitional game.

The zero-sum game \mathcal{G}_i is called the *coalitional game* for Player i .

Since a coalitional game is a zero-sum game, all the notions and results provided in Section 2.3 apply. In particular, given a multiplayer game \mathcal{G} and the coalitional game \mathcal{G}_i for Player i , if \mathcal{G}_i is determined, we may consider the value of a vertex $v \in V$ in \mathcal{G}_i . We denote this value $\text{Val}_i(v)$ in order to highlight that it is the zero-sum game \mathcal{G}_i which is considered.

Definition 4.1.2 (Value in a coalitional game). Given a multiplayer game \mathcal{G} , a player $i \in \Pi$ and \mathcal{G}_i the coalitional game for Player i . If \mathcal{G}_i is determined, for all $v \in V$, we write $\text{Val}_i(v)$ to denote the value of vertex v in \mathcal{G}_i .

Example 4.1.3. A multiplayer quantitative Reachability game $\mathcal{G} = (A, (\text{QR}_1, \text{QR}_2), (F_1, F_2))$ with two players is depicted in Figure 4.1. The round vertices are owned by Player 1 whereas the square vertices are owned by Player 2. The target sets of the players are respectively equal to $F_1 = \{v_2\}$

(grey vertex), $F_2 = \{v_2, v_5\}$ (double circled vertices).




Figure 4.1: A multiplayer quantitative Reachability game: Player 1 (resp. Player 2) owns round (resp. square) vertices and $F_1 = \{v_2\}$ and $F_2 = \{v_2, v_5\}$.

Since this is a game with two players, we consider the two corresponding coalitional game: the coalitional game for Player 1 and the coalitional game for Player 2.

Coalitional game for Player 1:

Let $\mathcal{G}_1 = (A_1, \text{QR}_1)$ with $F_1 = \{v_2\}$ be the coalitional game for Player 1, this is a quantitative Reachability zero-sum game and it is depicted in Figure 4.2a. The rounded vertices (in green) are owned by Player 1 (aka Player Min) and the rectangular vertices (in red) are owned by Player 2 (aka Player Max).

For all $v \in V$, the value $\text{Val}_1(v)$ is depicted in bold near the corresponding vertex and are summarized in Table 4.1.

Coalitional game for Player 2:

Let $\mathcal{G}_2 = (A_2, \text{QR}_2)$ with $F_2 = \{v_2, v_5\}$ be the coalitional game for Player 2, this is a quantitative Reachability zero-sum game and it is depicted in Figure 4.2b. The rectangular vertices (in green) are owned by Player 2 (aka Player Min) and the rounded vertices (in red) are owned by Player 1 (aka Player Max).

For all $v \in V$, the value $\text{Val}_2(v)$ is depicted in bold near the corresponding vertex and are summarized in Table 4.1.

4.2 Extended games

In this section we mainly focus on Reachability games: qualitative Reachability games, quantitative Reachability games or weighted Reachability games. For each kind of Reachability games, we may consider its extended version. The



(a) Coalitional game for Player 1. This is a quantitative Reachability zero-sum game where $F_1 = \{v_2\}$.

(b) Coalitional game for Player 2. This is a quantitative Reachability zero-sum game where $F_2 = \{v_2, v_5\}$.

Figure 4.2: Coalitional games related with the game depicted in Figure 4.1.

Table 4.1: Values in the coalitional games of the multiplayer quantitative Reachability game illustrated in Figure 4.1.

	v_0	v_3	v_1	v_6	v_7	v_2	v_4	v_5
Val ₁	$+\infty$	$+\infty$	3	2	1	0	$+\infty$	$+\infty$
Val ₂	$+\infty$	$+\infty$	$+\infty$	2	1	0	$+\infty$	0

vertices (v, I) of the extended game store a vertex $v \in V$ as well as a subset $I \subseteq \Pi$ of players that have already visited their target sets.

Let us recall that in order to provide a generic definition of a Reachability game whatever if it is a qualitative, quantitative or weighted Reachability game, we assume that a Reachability game is a game $\mathcal{G} = (A, \text{Reach}, (F_i)_{i \in \Pi})$ with $\text{Reach} = (\text{Reach}_i)_{i \in \Pi}$ and $F_i \subseteq V$ is the target set of Player i for all $i \in \Pi$ and where Reach refers to qR, QR or WR if \mathcal{G} is a qualitative, quantitative or weighted Reachability game respectively. It allows us to obtain a generic definition of its extended game.

Definition 4.2.1 (Extended game). Let $\mathcal{G} = (A, \text{Reach}, (F_i)_{i \in \Pi})$ be a Reachability game with an (weighted) arena $A = (\Pi, V, E, (V_i)_{i \in \Pi}, (w_i)_{i \in \Pi})$, and let v_0 be an initial vertex. The *extended game* of \mathcal{G} is equal to $\mathcal{X} = (X, \text{Reach}^X, (F_i^X)_{i \in \Pi})$ with the (weighed) arena $X = (\Pi, V^X, E^X, (V_i^X)_{i \in \Pi}, (w_i^X)_{i \in \Pi})$, such that:

- $V^X = V \times 2^\Pi$
- $((v, I), (v', I')) \in E^X$ if and only if $(v, v') \in E$ and $I' = I \cup \{i \in \Pi \mid$

$$v' \in F_i\}$$

- $(v, I) \in V_i^X$ if and only if $v \in V_i$
- $(v, I) \in F_i^X$ if and only if $i \in I$
- for all $i \in \Pi$, $w_i^X : E^X \rightarrow \mathbb{N}_0$ is such that for all $((v, I), (v', I')) \in E^X$, $w_i^X((v, I), (v', I')) = w_i(v, v')$;
- $\text{Reach}^X = (\text{Reach}_i^X)_{i \in \Pi}$ and for all $i \in \Pi$, Reach_i^X is the objective function Reach_i used in the game \mathcal{X}^a .

The initialized extended game (\mathcal{X}, x_0) associated with the initialized Reachability game (\mathcal{G}, v_0) is such that $x_0 = (v_0, I_0)$ with $I_0 = \{i \in \Pi \mid v_0 \in F_i\}$.

^aWe use Reach^X (resp. Reach_i^X) when we want to highlight that the objective functions are considered in the extended game with the sets F_i^X . Nevertheless, we could only write Reach (resp. Reach_i).

Remark 4.2.2. Given a safety game $\mathcal{G} = (\mathbf{A}, \text{Safe}, (F_i)_{i \in \Pi})$, we can also consider its extended game $\mathcal{X} = (X, \text{Safe}^X, (F_i^X)_{i \in \Pi})$ defined as in Definition 4.2.1 except that (i) there is no weight functions w_i and w_i^X , and (ii) the objective functions Reach_i and Reach_i^X , for all $i \in \Pi$, are replaced by Safe_i and Safe_i^X respectively. For all $i \in \Pi$, the gain function Safe_i^X is a safety objective associated with the sets $(F_i^X)_{i \in \Pi}$ in the extended game \mathcal{X} .

Remark 4.2.3. To ease the notations, if $x \in V^X$ is such that $x = (v, I)$ for some $v \in V$ and $I \subseteq \Pi$, we sometimes write $I(x)$ to denote the set I of players that have already visited their target sets in x .

Firstly, notice that if Reach is qR or QR the weight functions $(w_i)_{i \in \Pi}$ and so $(w_i^X)_{i \in \Pi}$ are omitted.

Secondly, if \mathcal{G} is a qualitative, quantitative, weighed Reachability game, then its extended game is also a qualitative, quantitative, weighted Reachability game. In the same way, if \mathcal{G} is a Safety game, then its extended game is a Safety game.

Thirdly, remark the way each target set F_i^X is defined: if $v \in F_i$, then $(v, I) \in F_i^X$ but also $(v', I') \in F_i^X$ for all $(v', I') \in \text{Succ}^*(v, I)$.

Example 4.2.4. The extended game (\mathcal{X}, x_0) of the multiplayer quantitative Reachability drawn in Figure 4.1 is depicted in Figure 4.3 (only the part reachable from the initial vertex $x_0 = (v_0, \emptyset)$ is drawn). The gray vertices represent the target set F_1^X of Player 1 as the double circled vertices represent the target set F_2^X of Player 2.

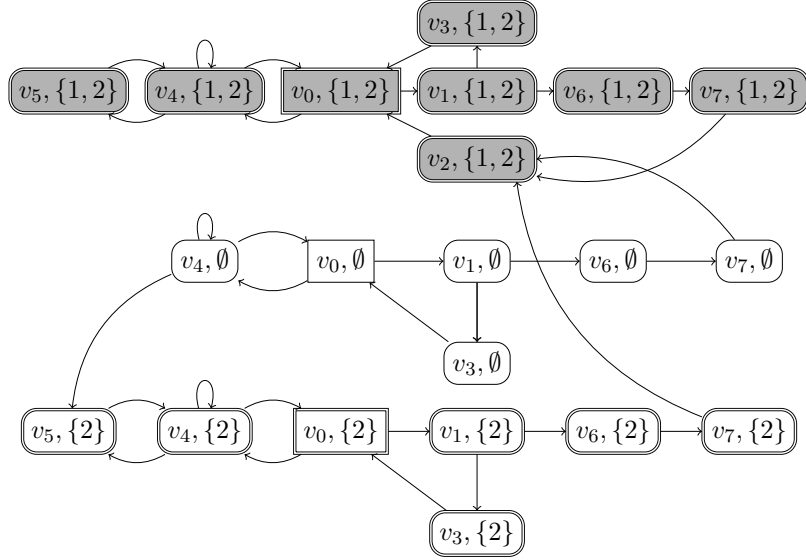


Figure 4.3: The extended game (\mathcal{X}, x_0) of the initialized quantitative Reachability game (\mathcal{G}, v_0) of Figure 4.1.

Remark 4.2.5. Even if the construction of (\mathcal{X}, x_0) from (\mathcal{G}, v_0) causes an exponential blow-up of the number of states, the extended game of a Reachability or Safety game enjoys some useful properties.

I-monotonicity For each $\rho = (v_0, I_0)(v_1, I_1) \dots \in \text{Plays}_X(x_0)$, we have the next property called *I-monotonicity*:

$$I_k \subseteq I_{k+1} \quad \text{for all } k \in \mathbb{N}. \quad (4.1)$$

One-to-one correspondence 1 Given an initialized Reachability or Safety game (\mathcal{G}, v_0) and its extended game (\mathcal{X}, x_0) , there is a one-to-one correspondence between plays in $\text{Plays}_A(v_0)$ and plays in $\text{Plays}_X(x_0)$:

- from $\rho = \rho_0\rho_1 \dots \in \text{Plays}_A(v_0)$, we derive $\rho^X = (\rho_0, I_0)(\rho_1, I_1) \dots \in \text{Plays}_X(x_0)$ such that I_k is the set of Players i such that F_i is visited along $\rho_{\leq k}$;
- from $\rho = (v_0, I_0)(v_1, I_1) \dots \in \text{Plays}_X(x_0)$, we derive $\rho^A = v_0v_1 \dots \in \text{Plays}_A(v_0)$ such that the second components $I_k, k \in \mathbb{N}$, are omitted.

One-to-one correspondence 2 Third, given $\rho \in \text{Plays}_A(v_0)$, we have that $\text{Reach}^X(\rho^X) = \text{Reach}(\rho)$, and conversely given $\rho \in \text{Plays}_X(x_0)$, we have that $\text{Reach}(\rho^A) = \text{Reach}^X(\rho)$. The same property holds by replacing Reach by Safe and Reach^X by Safe^X . It follows that outcomes of SPEs (resp. NEs) can be equivalently studied in (\mathcal{G}, v_0) and in (\mathcal{X}, x_0) , as stated in the next lemma.

Lemma 4.2.6.

Reachability game *Given an initialized Reachability game (\mathcal{G}, v_0) and its initialized extended game (\mathcal{X}, x_0) , if ρ is the outcome of an SPE (resp. NE) in (\mathcal{G}, v_0) , then ρ^X is the outcome of an SPE (resp. NE) in (\mathcal{X}, x_0) with $\text{Reach}(\rho) = \text{Reach}^X(\rho^X)$. Conversely, if ρ is the outcome of an SPE (resp. NE) in (\mathcal{X}, x_0) , then ρ^A is the outcome of an SPE (resp. NE) in (\mathcal{G}, v_0) with $\text{Reach}(\rho^A) = \text{Reach}^X(\rho)$.*

Safety game *Given an initialized Safety game (\mathcal{G}, v_0) and its initialized extended game (\mathcal{X}, x_0) , if ρ is the outcome of an SPE (resp. NE) in (\mathcal{G}, v_0) , then ρ^X is the outcome of an SPE (resp. NE) in (\mathcal{X}, x_0) with $\text{Safe}(\rho) = \text{Safe}^X(\rho^X)$. Conversely, if ρ is the outcome of an SPE (resp. NE) in (\mathcal{X}, x_0) , then ρ^A is the outcome of an SPE (resp. NE) in (\mathcal{G}, v_0) with $\text{Safe}(\rho^A) = \text{Safe}^X(\rho)$.*

Notice that in the same way there is a one-to-one correspondence between strategies in (\mathcal{G}, v_0) and in its extended game (\mathcal{X}, x_0) . Thus Lemma 4.2.6 could be rephrased in term of SPEs (resp. NEs) instead of SPEs outcomes (resp. NEs outcomes).

Another property of extended games is that in the extended game the objective functions QR_i and WR_i become strongly prefix-linear and qR_i and

Safe_{*i*} become prefix-independent (and so strongly prefix-linear).

Proposition 4.2.7. *Let $\mathcal{G} = (\mathbf{A}, \text{Obj}, (F_i)_{i \in \Pi})$ be either a Reachability game or a Safety game and let $\mathcal{X} = (X, \text{Obj}^X, (F_i^X)_{i \in \Pi})$ be its extended game.*

1. *If $\text{Obj} = \text{qR}$, then for all $i \in \Pi$, qR_i^X is prefix-independent in X ;*
2. *If $\text{Obj} = \text{Safe}$, then for all $i \in \Pi$, Safe_i^X is prefix-independent in X ;*
3. *If $\text{Obj} = \text{QR}$ (resp. $\text{Obj} = \text{WR}$), then for all $i \in \Pi$, QR_i^X (resp. WR_i^X) is strongly prefix-linear in X .*

Proof. Let $\mathcal{G} = (\mathbf{A}, \text{Obj}, (F_i)_{i \in \Pi})$ be either a Reachability game or a Safety game and let $\mathcal{X} = (X, \text{Obj}^X, (F_i^X)_{i \in \Pi})$ be its extended game.

1. Let us assume that $\text{Obj} = \text{qR}$ (resp. Safe), let $i \in \Pi$ be a player, we have to prove that qR_i^X (resp. Safe_i^X) is prefix-independent in X . Let $hv \in \text{Hist}^X$ be a history such that $hv = h_0 \dots h_k$ for some $k \in \mathbb{N}$. We have to prove that for all $\rho \in \text{Plays}_X(v)$, $\text{qR}_i^X(h\rho) = \text{qR}_i^X(\rho)$ (resp. $\text{Safe}_i^X(h\rho) = \text{Safe}_i^X(\rho)$).
 - If there exists $0 \leq n \leq k$ such that $i \in I(h_n)$ that means that for all $x \in \text{Succ}^*(h_n)$, $i \in I(x)$ (by I -monotonicity). Thus in particular $i \in I(v)$. It implies that $\text{qR}_i^X(h\rho) = 1$ (resp. $\text{Safe}_i^X(h\rho) = 0$) and $\text{qR}_i^X(\rho) = 1$ (resp. $\text{Safe}_i^X(\rho) = 0$).
 - Otherwise we have that the gain of Player i only depends whether F_i^X is visited along ρ or not. Hence we obtain $\text{qR}_i^X(h\rho) = \text{qR}_i^X(\rho)$ (resp. $\text{Safe}_i^X(h\rho) = \text{Safe}_i^X(\rho)$).
2. Let us assume that $\text{Obj} = \text{QR}$ or $\text{Obj} = \text{WR}$. Let $i \in \Pi$ be a player, since QR_i is a particular case of WR_i , we only have to prove that WR_i^X is strongly prefix-linear in X . Let $hv \in \text{Hist}_X$ be a history such that $hv = h_0 \dots h_k$ for some $k \in \mathbb{N}$:
 - If there exists $0 \leq \ell \leq k$ such that $i \in I(h_\ell)$, i.e., hv visits F_i^X , let us assume that ℓ is the least such index: we choose $a(h, v) =$

$\sum_{n=0}^{\ell-1} w_i(h_n, h_{n+1})^a$ and $b(h, v) = 1$. Indeed, if $i \in I(h_\ell)$, then $i \in I(v)$ (by I -monotonicity). Hence for all $\rho \in \text{Plays}_X(v)$, $\text{WR}_i^X(\rho) = 0$. It follows that for all $\rho \in \text{Plays}_X(v)$, $\text{WR}_i^X(h\rho) = a(h, v) + b(h, v) \cdot \text{WR}_i^X(\rho) = a(h, v) = \sum_{n=0}^{\ell-1} w_i(h_n, h_{n+1})$.

- Otherwise: we choose $a(h, v) = \sum_{n=0}^{k-1} w_i(h_n, h_{n+1})$ and $b(h, v) = 1$.

For all $\rho \in \text{Plays}_X(v)$, we have that $\text{WR}_i^X(h\rho) = a(h, v) + b(h, v) \cdot$

$$\text{WR}_i^X(\rho) = \sum_{n=0}^{k-1} w_i(h_n, h_{n+1}) + \text{WR}_i^X(\rho). \quad \square$$

^aWith the assumption that if $\ell = 0$, $\sum_{n=0}^{\ell-1} w_i(h_n, h_{n+1}) = 0$.

CHAPTER 5

λ -CONSISTENT PLAYS

In order to characterize equilibria outcomes, we need to define a criterion which has to be satisfied by the plays which are effectively outcomes of equilibria. Both for Nash equilibria, weak subgame perfect equilibria and subgame perfect equilibria, the criterion in this document is based on the notion of λ -consistent play where λ is a function which labels the vertices of the game graph with some appropriate values. We explain how to obtain these values in Chapter 6 and Chapter 7.

In this chapter, we only define the notions of *labeling function* and λ -consistent plays. Since the notion of λ -consistent play is one of the central notions of our approach in the rest of the document, we have chosen to devote an entire chapter to only introduce this concept.

5.1 Labeling Function and λ -Consistency

Given an arena A , a labeling function λ is a function which assigns a value to each vertex of this arena.

Definition 5.1.1 (Labeling function). Given an arena A , the function $\lambda : V \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ which assigns a value to each vertex of A is a *labeling*

function.

Given a game $\mathcal{G} = (A, \text{Cost})$, a labeling function $\lambda : V \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ and a play $\rho \in \text{Plays}$ we say that ρ is consistent with λ if all the suffixes of ρ satisfy the constraints given by the labeling function λ .

Definition 5.1.2 (λ -consistent play). Given a game $\mathcal{G} = (A, \text{Cost})$, a labeling function $\lambda : V \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ and a plays $\rho = \rho_0\rho_1 \dots \in \text{Plays}$, ρ is λ -consistent if

$$\forall i \in \Pi, \forall n \in \mathbb{N}, (\rho_n \in V_i \implies \text{Cost}_i(\rho_{\geq n}) \leq \lambda(\rho_n)).$$

If ρ is λ -consistent, we write $\rho \models \lambda$. Otherwise, we write $\rho \not\models \lambda$.

Example 5.1.3. We come back to the multiplayer quantitative Reachability game depicted in Figure 4.1. A labeling function $\lambda : V \rightarrow \mathbb{N} \cup \{+\infty\}$ is given in Figure 5.1 where for all $v \in V$, $\lambda(v)$ is written near the vertex v . These values are summarized in the following table.

	v_0	v_1	v_2	v_3	v_4	v_5	v_6	v_7
λ	$+\infty$	3	0	$+\infty$	$+\infty$	$+\infty$	2	1




Figure 5.1: Multiplayer quantitative Reachability game of Figure 4.1 labeled by a labeling function λ .

The play $(v_0v_4)^\omega$ is λ -consistent. Indeed, if we write $(v_0v_4)^\omega = \rho_0\rho_1 \dots$, let $n \in \mathbb{N}$, if n is even, then $\rho_n = v_0$ and $\rho_n \in V_2$. In this case, $\text{Cost}_2(\rho_{\geq n}) = \text{Cost}_2((v_0v_4)^\omega) = +\infty$ and $\lambda(\rho_n) = +\infty$ so the inequality $\text{Cost}_2(\rho_{\geq n}) \leq \lambda(\rho_n)$ is satisfied. In the same way, if n is odd, then $\rho_n = v_4$ and $\rho_n \in V_1$. In this

case, $\text{Cost}_1(\rho_{\geq n}) = \text{Cost}_1((v_4v_0)^\omega) = +\infty$ and $\lambda(v_4) = +\infty$ so the inequality $\text{Cost}_1(\rho_{\geq n}) \leq \lambda(\rho_n)$ is satisfied.

The play $(v_0v_1v_3)^\omega$ is not λ -consistent. Let us assume that $(v_0v_1v_3)^\omega = \rho_0\rho_1\dots$, we have that $\text{Cost}_1(\rho_{\geq 1}) = \text{Cost}_1((v_1v_3v_0)^\omega) = +\infty$ and $\lambda(\rho_1) = \lambda(v_1) = 3$, it follows that $\text{Cost}_1(\rho_{\geq 1}) > \lambda(\rho_1)$.

CHAPTER 6

NE OUTCOME CHARACTERIZATION

This chapter provides a Nash equilibrium outcome characterization for some classes of games. In Section 6.1, we first formally define the labeling function Val^* . Secondly, we prove that if the cost functions satisfy some hypotheses and if the associated coalitional games are determined with positional optimal strategies for both players, then a play is the outcome of a Nash equilibrium if and only if it is Val^* -consistent. This characterization holds in particular for games with strongly prefix-linear cost functions. Since reachability objectives (qualitative, quantitative and weighted) are not strongly prefix-linear, we explain in Section 6.2 how to adapt this characterization to Reachability games.

6.1 General Characterization

Let \mathcal{G} be a multiplayer game such that, for all $i \in \Pi$, the coalitional game \mathcal{G}_i is determined, we define the labeling function Val^* in the following way: if $v \in V$ is a vertex of Player i , then $\text{Val}^*(v) = \text{Val}_i(v)$.

Definition 6.1.1 (Val^* -labeling function). Given a multiplayer game \mathcal{G} such that for all $i \in \Pi$, the coalitional game \mathcal{G}_i for Player i is determined, we define

the labeling function $\text{Val}^* : V \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ such that for all $v \in V$, if $v \in V_i$, then $\text{Val}^*(v) = \text{Val}_i(v)$.

If the initialized multiplayer game $(\mathcal{G}, v_0) = (A, \text{Cost})$ satisfies some conditions, this labeling function allows to determine whether a play $\rho \in \text{Plays}(v_0)$ is the outcome of a Nash equilibrium in (\mathcal{G}, v_0) .

Theorem 6.1.2 (Nash equilibrium outcome characterization). *Given an initialized multiplayer game $(\mathcal{G}, v_0) = (A, \text{Cost})$. If,*

(C₁) *For all $i \in \Pi$, all $hv \in \text{Hist}$ and for all $\rho, \rho' \in \text{Plays}(v)$: $\text{Cost}_i(\rho) \geq \text{Cost}_i(\rho') \implies \text{Cost}_i(h\rho) \geq \text{Cost}_i(h\rho')$;*

(C₂) *For all $i \in \Pi$, all $hv \in \text{Hist}$ and for all $\rho, \rho' \in \text{Plays}(v)$: $\text{Cost}_i(\rho) > \text{Cost}_i(\rho') \implies \text{Cost}_i(h\rho) > \text{Cost}_i(h\rho')$;*

(C₃) *For all $i \in \Pi$, the coalitional game \mathcal{G}_i for Player i is determined and there exist a memoryless optimal strategy σ_i^* for Player i (aka Player Min) and a memoryless optimal strategy σ_{-i}^* for Player $-i$ (aka Player Max);*

then, given a play $\rho \in \text{Plays}(v_0)$, the following assertions are equivalent:

1. *there exists an NE σ in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = \rho$;*
2. *the play ρ is Val^* -consistent.*

Remark 6.1.3. Notice that conditions (C₁) and (C₂) are satisfied if the cost functions are strongly prefix-linear in A (Lemma 2.2.20). It follows, by Proposition 2.3.8 and Proposition 2.3.13, that the conditions of Theorem 6.1.2 are satisfied by

- multiplayer Büchi, co-Büchi and Parity games;
- multiplayer discounted-sum games and mean-payoff games.

Thus, for these games, the set of plays beginning in v_0 which are Val^* -

consistent is exactly the set of plays which are NEs outcomes in (\mathcal{G}, v_0) . That is $\{\rho \in \text{Plays}(v_0) \mid \rho \models \text{Val}^*\} = \{\rho \in \text{Plays}(v_0) \mid \exists \sigma \text{ an NE in } (\mathcal{G}, v_0) \text{ st. } \langle \sigma \rangle_{v_0} = \rho\}$.

Remark 6.1.4. Some other characterizations rely on other very similar conditions (e.g., [Bru17]) but our aim is not to provide an exhaustive characterization since the one enunciated in Theorem 6.1.2 is sufficient to obtain our results in the remaining part of this document.

For the sake of completeness and in order to ease the proof of Corollary 6.1.5 we provide a detailed proof of Theorem 6.1.2. The main idea is that if the second assertion is false, then there exists a player i who has an incentive to deviate along ρ . Indeed, if there exists $k \in \mathbb{N}$ such that $\text{Cost}_i(\rho_{\geq k}) > \text{Val}_i(\rho_k)$ ($\rho_k \in V_i$) it means that Player i can ensure a better cost for him even if the other players play in coalition and in an antagonistic way. Thus, Player i has a profitable deviation. For the second implication, the Nash equilibrium σ is defined as follows: all players follow the outcome ρ but if one player, assume it is Player i , deviates from ρ the other players form a coalition $-i$ and punish the deviator by playing the optimal strategy of Player $-i$ in the coalitional game \mathcal{G}_i .

Proof of Theorem 6.1.2. Let (\mathcal{G}, v_0) be an initialized multiplayer game which satisfies conditions (C₁), (C₂) and (C₃). Let $\rho \in \text{Plays}(v_0)$ be a play such that $\rho = \rho_0\rho_1 \dots$

Let us begin with some remarks. Since for all $i \in \Pi$, \mathcal{G}_i is determined and there exist memoryless optimal strategies σ_i^* and σ_{-i}^* for Player i and Player $-i$ respectively. For all $v \in V$, by Lemma 2.3.12, we have:

$$\inf_{\sigma_i \in \Sigma_{\text{Min}}} \text{Cost}_i(\langle \sigma_i, \sigma_{-i}^* \rangle_v) = \text{Val}_i(v) = \sup_{\sigma_{-i} \in \Sigma_{\text{Max}}} \text{Cost}_i(\langle \sigma_i^*, \sigma_{-i} \rangle_v). \quad (6.1)$$

From the memoryless optimal strategy σ_{-i}^* in \mathcal{G}_i we can extract a memoryless strategy $\sigma_{j,i}^*$ of Player j in \mathcal{G} . Notice also that even if σ_i^* is a strategy in \mathcal{G}_i , we can use it as a strategy of Player i in \mathcal{G} .

Let us prove the equivalence between the two assertions of Theorem 6.1.2.

(1 \Rightarrow 2): Let σ be a Nash equilibrium in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = \rho$. Let us

assume, by contradiction, that ρ is not Val^* -consistent:

there exist $i \in \Pi$ and $k \in \mathbb{N}$ such that $\rho_k \in V_i$ and:

$$\text{Cost}_i(\rho_{\geq k}) > \text{Val}_i(\rho_k). \quad (6.2)$$

Let $h = \rho_0 \dots \rho_{k-1}$, we can write:

$$\text{Cost}_i(\rho_{\geq k}) = \text{Cost}_i(\langle \sigma_{\uparrow h} \rangle_{\rho_k}). \quad (6.3)$$

Additionally, by Equation (6.1) and thanks to the fact that the optimal strategies are memoryless:

$$\begin{aligned} \text{Val}_i(\rho_k) &= \sup_{\tau_{-i} \in \Sigma_{Max}} \text{Cost}_i(\langle \sigma_i^*, \tau_{-i} \rangle_{\rho_k}) \\ &\geq \text{Cost}_i(\langle \sigma_i^*, \sigma_{-i|h} \rangle_{\rho_k}) \end{aligned} \quad (6.4)$$

where σ_i^* is the optimal strategy of Player i in \mathcal{G}_i and σ_{-i} is an abuse of notation to depict the strategy of the coalition $-i = \Pi \setminus \{i\}$ which follows strategies σ_j for all $j \neq i$.

By (6.2), (6.3) and (6.4), it follows that:

$$\text{Cost}_i(\langle \sigma_i^*, \sigma_{-i|h} \rangle_{\rho_k}) < \text{Cost}_i(\langle \sigma_{\uparrow h} \rangle_{\rho_k}).$$

By condition (C_2) , we can conclude that:

$$\text{Cost}_i(h \langle \sigma_i^*, \sigma_{-i|h} \rangle_{\rho_k}) < \text{Cost}_i(h \langle \sigma_{\uparrow h} \rangle_{\rho_k}) = \text{Cost}_i(\rho). \quad (6.5)$$

This means that following σ_i along h and then σ_i^* once he reaches ρ_k is a profitable deviation for Player i . This concludes the proof.

(2 \Rightarrow 1): Let us assume that ρ is Val^* -consistent and let us prove that there exists a Nash equilibrium σ in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = \rho$. Let τ be a strategy profile such that $\langle \tau \rangle_{v_0} = \rho$. From τ we aim to construct a Nash equilibrium with the same outcome. The main idea is the following one:

first, all players play according to τ . But if a player, let us call him Player i deviates from τ_i , the other players form a coalition and each of them plays his strategy obtained from the strategy σ_{-i}^* in \mathcal{G}_i .

In order to define properly the Nash equilibrium that we are looking for, we have to define a punishment function $P : \text{Hist}(v_0) \rightarrow \Pi \cup \{\perp\}$ which allows us to know who is the player who has deviated for the first time from τ . So, for all $h \in \text{Hist}(v_0)$, $P(h) = \perp$ if no player has yet deviated and $P(h) = i$ for some $i \in \Pi$ if Player i is the first player who has deviated along h . We can define P as follows: for the initial vertex $P(v_0) = \perp$ and then for all history $hv \in \text{Hist}(v_0)$ with $v \in V$:

$$P(hv) = \begin{cases} \perp & \text{if } P(h) = \perp \text{ and } hv \text{ is a prefix of } \rho, \\ i & \text{if } P(h) = \perp, hv \text{ is not a prefix of } \rho \text{ and } h \in \text{Hist}_i, \\ P(h) & \text{otherwise.} \end{cases}$$

We now define σ . For all $i \in \Pi$ and for all $h \in \text{Hist}_i(v_0)$:

$$\sigma_i(h) = \begin{cases} \tau_i(h) & \text{if } P(h) = \perp, \\ \sigma_i^*(h) & \text{if } P(h) = i, \\ \sigma_{i,P(h)}^*(h) & \text{otherwise.} \end{cases}$$

It is clear that $\langle \sigma \rangle_{v_0} = \rho$. It remains to prove that σ is a Nash equilibrium in (\mathcal{G}, v_0) . Let us assume that σ is not an NE. It means that there exists a profitable deviation depicted by $\tilde{\sigma}_i$ for some player $i \in \Pi$. Let $\tilde{\rho} = \langle \tilde{\sigma}_i, \sigma_{-i} \rangle_{v_0}$ the outcome such that Player i plays his profitable deviation. As $\tilde{\sigma}_i$ is a profitable deviation we have:

$$\text{Cost}_i(\tilde{\rho}) < \text{Cost}_i(\rho). \quad (6.6)$$

Moreover as ρ and $\tilde{\rho}$ both begin in v_0 , they have a common prefix. Let $hv \in \text{Hist}_i$ the longest common prefix of ρ and $\tilde{\rho}$. We have that: $\rho = h\langle \sigma_{|h} \rangle_v$ and $\tilde{\rho} = h\langle \tilde{\sigma}_{i|h}, \sigma_{-i|h} \rangle_v$. But, by definition of σ and

as the optimal strategies in \mathcal{G}_i are memoryless, we can rewrite these two equalities as follows: $\rho = h\langle\tau_{|h}\rangle_v$ and $\tilde{\rho} = h\langle\tilde{\sigma}_{i|h}, (\sigma_{j,i}^*)_{j \in \Pi \setminus \{i\}}\rangle_v$. Additionally, by Equation (6.1):

$$\begin{aligned} \text{Val}_i(v) &= \inf_{\mu_i \in \Sigma_{Min}} \text{Cost}_i(\langle\mu_i, \sigma_{-i}^*\rangle_v) \\ &\leq \text{Cost}_i(\langle\tilde{\sigma}_{i|h}, \sigma_{-i}^*\rangle_v) \\ &= \text{Cost}_i(\langle\tilde{\sigma}_{i|h}, (\sigma_{j,i}^*)_{j \in \Pi \setminus \{i\}}\rangle_v) \\ &= \text{Cost}_i(\langle\tilde{\sigma}_{i|h}, \sigma_{-i|h}\rangle_v) \end{aligned} \tag{6.7}$$

By hypothesis, ρ is Val^* -consistent, thus we have that

$$\text{Val}_i(v) \geq \text{Cost}_i(\langle\tau_{|h}\rangle_v) \tag{6.8}$$

Thus by (6.7), it follows that:

$$\text{Cost}_i(\langle\tilde{\sigma}_{i|h}, \sigma_{-i|h}\rangle_v) \geq \text{Cost}_i(\langle\tau_{|h}\rangle_v).$$

And thanks to condition (C_1) , we have that:

$$\begin{aligned} \text{Cost}_i(\tilde{\rho}) &= \text{Cost}_i(h\langle\tilde{\sigma}_{i|h}, \tau_{-i|h}\rangle_v) \\ &\geq \text{Cost}_i(h\langle\tau_{|h}\rangle_v). \\ &= \text{Cost}_i(\rho). \end{aligned}$$

Which leads to a contradiction with (6.6) and concludes the proof. \square

Theorem 6.1.2 and its previous proof can be easily adapted to lassoes as follows.

Corollary 6.1.5 (of Theorem 6.1.2). *Let (\mathcal{G}, v_0) be a multiplayer game such that conditions (C_1) , (C_2) and (C_3) of Theorem 6.1.2 are satisfied. Given $\rho = h\ell^\omega \in \text{Plays}(v_0)$ a lasso, the following assertions are equivalent:*

1. *there exists an NE σ with memory in $\mathcal{O}(|h\ell| + |\Pi|)$ and such that $\langle\sigma\rangle_{v_0} =$*

ρ .

2. the play ρ is Val^* -consistent.

Proof sketch. Let us now assume that $\rho = h\ell^\omega$ is a lasso. The implication $\mathbf{1} \Rightarrow \mathbf{2}$ is the same as in the previous proof. Thus we only have to prove that, in the implication $\mathbf{2} \Rightarrow \mathbf{1}$, the previously built strategy σ has memory in $\mathcal{O}(|h\ell| + |\Pi|)$. The intuition is the following. If $\rho = h\ell^\omega$, a player has to remember: (i) $h\ell$ to know both what he has to play and if someone has deviated and (ii) who is the deviator. Once a deviation has occurred, both players play memoryless strategies. \square

Example 6.1.6. We consider the initialized multiplayer Büchi game $(\mathcal{G}, v_0) = (A, \text{Buchi}, (F_1, F_2))$ such that its arena is depicted in Figure 6.1 where Player 1 (resp. Player 2) owns rounded vertices (resp. rectangular vertices) and with the target sets $F_1 = \{v_1\}$ and $F_2 = \{v_3, v_5\}$.




Figure 6.1: Multiplayer Büchi game. In this example, Player 1 (resp. Player 2) owns rounded (resp. rectangular) vertices and $F_1 = \{v_1\}$ (resp. $F_2 = \{v_3, v_5\}$).

The labeling function Val^* is depicted in Figure 6.1, for each $v \in V$, $\text{Val}^*(v)$ is written in bold near the vertex v and these values are summarized in the following table.

	v_0	v_1	v_2	v_3	v_4	v_5	v_6
Val^*	1	1	1	1	1	0	0

The play $v_0(v_1v_2)^\omega$ is not Val^* -consistent thus it cannot be an NE outcome. Indeed, $\text{Buchi}_2(v_0(v_2v_1)^\omega) = 0 < 1 = \text{Val}^*(v_0)$. The play $v_0v_4v_5^\omega$ is Val^* -consistent. Indeed, $\text{Buchi}_2(v_0v_4v_5^\omega) = 1 \geq 1 =$

$\text{Val}^*(v_0)$, $\text{Buchi}_2(v_4v_5^\omega) = 1 \geq 1 = \text{Val}^*(v_4)$ and if we write $v_0v_4v_5^\omega$ as $\rho_0\rho_1\dots$, for all $n \geq 2$, $\text{Buchi}_1(\rho_{\geq n}) = \text{Buchi}_1(v_5^\omega) = 0 \geq 0 = \text{Val}^*(v_5)$. Thus $v_0v_4v_5^\omega$ is an NE outcome in (\mathcal{G}, v_0) : for example, it is the outcome of the memoryless strategy profile (σ_1, σ_2) from v_0 . This strategy profile is depicted by double arrows in Figure 6.1. We can easily prove that (σ_1, σ_2) is an NE in (\mathcal{G}, v_0) .

6.2 Characterization for Reachability Games

The NE outcome characterization provided in Theorem 6.1.2 does not apply to multiplayer Reachability games (qualitative, quantitative or weighted). Indeed, they do not fulfill condition (C_2) of this theorem.

Nevertheless, let $\mathcal{G} = (A, \text{Reach}, (F_i)_{i \in \Pi})$ be a multiplayer Reachability game and let $\mathcal{X} = (X, \text{Reach}^X, (F_i^X)_{i \in \Pi})$ be its associated extended game, we know that for all $i \in \Pi$, Reach_i^X is strongly prefix-linear in X (Proposition 4.2.7). It follows, by Lemma 2.2.20, that conditions (C_1) and (C_2) are satisfied in the extended game. Additionally, by Proposition 2.3.8 and Proposition 2.3.13, condition (C_3) is also satisfied in the extended game.

From a theoretical point of view, that means that, thanks to Lemma 4.2.6, if we want to study the outcomes of Nash equilibria in an initialized multiplayer Reachability game (\mathcal{G}, v_0) we can study the Val^* -consistent plays of the associated extended game. But from a practical point of view, that means that one has to consider the associated extended game (\mathcal{X}, x_0) which has a size exponential in the size of \mathcal{G} and, in particular, to compute an exponential number of values to obtain Val^* .

Corollary 6.2.1 (of Theorem 6.1.2). *Let $(\mathcal{G}, v_0) = (A, \text{Reach}, (F_i)_{i \in \Pi})$ be an initialized multiplayer Reachability game and let $(\mathcal{X}, x_0) = (X, \text{Reach}^X, (F_i^X)_{i \in \Pi})$ be its associated extended game. Given a play $\rho \in \text{Plays}_X(x_0)$. The following assertions are equivalent:*

1. *there exists an NE σ in (\mathcal{X}, x_0) such that $\langle \sigma \rangle_{x_0} = \rho$;*
2. *ρ is Val^* -consistent.*

To avoid this exponential blow-up, we explain in the remaining part of this

section, how the definition of Val^* -consistency may be slightly modified to be directly applied to the Reachability game.

The difference is that, for a given player $i \in \Pi$, instead of checking that the constraints given by the labeling function λ are satisfied by all the suffixes of the play beginning in a vertex owned by Player i , we proceed in this way until ρ visits the target set F_i . Indeed, once Player i has reached his target set along ρ , his cost will no longer change.

In the following definition, given a history $h \in \text{Hist}$, the set $\text{Visit}(h)$ depicts the set of players that reach their target set along h , *i.e.*, if $h = h_0 \dots h_k$ for some $k \in \mathbb{N}$, $\text{Visit}(h) = \{i \in \Pi \mid \exists 0 \leq n \leq k \text{ st. } h_n \in F_i\}$.

Definition 6.2.2 (Visit λ -consistent play). Let $\mathcal{G} = (A, \text{Reach}, (F_i)_{i \in \Pi})$ be a multiplayer Reachability game. Let $\lambda : V \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ be a labeling function and let $\rho \in \text{Plays}$. If for all $i \in \Pi$, for all $n \in \mathbb{N}$:

$$(\rho_n \in V_i \text{ and } i \notin \text{Visit}(\rho_0 \dots \rho_n)) \implies \text{Reach}_i(\rho_{\geq n}) \leq \lambda(\rho_n)^a$$

then we say that ρ is Visit λ -consistent.

^aNotice that, here we assume that Reach_i is a cost function (either QR_i or WR_i). The counterpart for qR_i is thus $\text{qR}_i(\rho_{\geq n}) \geq \lambda(\rho_n)$ since that is a gain function.

This notion of Visit λ -consistency allows to obtain the counterpart of Theorem 6.1.2 for Reachability games.

Theorem 6.2.3 (NE outcome characterization for Reachability games). *Let $(\mathcal{G}, v_0) = (A, \text{Reach}, (F_i)_{i \in \Pi})$ be an initialized multiplayer Reachability game. Given a play $\rho \in \text{Plays}(v_0)$, the following assertions are equivalent:*

1. *There exists an NE σ in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = \rho$;*
2. *The play ρ is Visit Val^* -consistent.*

Roughly speaking, the proof of this theorem is the same proof as the proof of Theorem 6.1.2. Indeed, conditions (C_1) and (C_3) are satisfied by multiplayer

Reachability games (Lemma 2.2.21, Proposition 2.3.8 and Proposition 2.3.13). Hence, we only have to pay attention to (i) the use of the Visit Val^{*}-consistency concept instead of Val^{*}-consistency; and (ii) how condition (C₂) is used since this latter condition is satisfied in Reachability games only if the target set of the considered player is not visited along the prefix h .

Proof. Due to Lemma 2.2.21, Proposition 2.3.8 and Proposition 2.3.13, the proof is the same as the proof of Theorem 6.1.2 with the exception of the arguments that we have highlighted by boxes. Let us comment those slight modifications.

- (1 \Rightarrow 2) :
1. Equation (6.2) and the definition of ρ is not Val^{*}-consistent has to be replaced by: ρ is not Visit Val^{*}-consistent: there exists $i \in \Pi$ and $k \in \mathbb{N}$ such that $\rho_k \in V_i$ and $i \notin \text{Visit}(\rho_0 \dots \rho_k)$ and $\text{Cost}_i(\rho_{\geq k}) > \text{Val}^*(\rho_k) = \text{Val}_i(\rho_k)$.
 2. Since $h = \rho_0 \dots \rho_{k-1}$, we also have that $i \notin \text{Visit}(h)$, thus by Lemma 2.2.21, condition (C₂) is also satisfied for Reachability objectives and the inequality (6.5) still holds.
 3. Other arguments to prove this implication do not change.
- (2 \Rightarrow 1) :
1. The only modification for this implication is in Equation (6.8). The Val^{*}-consistency of play ρ has to be changed by Visit Val^{*}-consistency of play ρ . Notice that, since hv is a common prefix of ρ and ρ' and $\text{Cost}_i(\tilde{\rho}) < \text{Cost}_i(\rho)$, $i \notin \text{Visit}(hv)$. And thus, one can also state that $\text{Val}_i(v) \geq \text{Cost}_i(\langle \tau|_h \rangle v)$.
 2. Other arguments to prove this implication do not change. \square

Corollary 6.2.4 (of Theorem 6.2.3). *et* $(\mathcal{G}, v_0) = (\mathbf{A}, \text{Reach}, (F_i)_{i \in \Pi})$ be an initialized multiplayer Reachability game. Given $\rho = h\ell^\omega \in \text{Plays}(v_0)$ a lasso, the following assertions are equivalent:

1. there exists an NE σ with memory in $\mathcal{O}(|h\ell| + |\Pi|)$ and such that $\langle \sigma \rangle_{v_0} = \rho$.
2. the play ρ is Visit Val^{*}-consistent.

We conclude this section with an example.

Example 6.2.5. We come back to Example 4.1.3. Let us recall it is a multiplayer quantitative Reachability game $\mathcal{G} = (\mathbf{A}, (\mathbf{QR}_1, \mathbf{QR}_2), (F_1, F_2))$ with two players. The game arena is depicted in Figure 4.1, the round vertices are owned by Player 1 whereas the square vertices are owned by Player 2. The target sets of the players are respectively equal to $F_1 = \{v_2\}$ (grey vertex), $F_2 = \{v_2, v_5\}$ (double circled vertices). The value of the labeling function Val^* may be obtained from Table 4.1. For the sake of clarity, we provide again the arena of the game in Figure 6.2 and add the values of Val^* in bold near the corresponding vertex.




Figure 6.2: The multiplayer quantitative Reachability game of Example 4.1.3 enhanced with the values of Val^* .

Let us consider the play $\rho = (v_0v_4)^\omega$ and let us prove that this play is Visit Val^* -consistent. First of all, we have that for all $n \in \mathbb{N}$, $\mathbf{QR}(\rho_{\geq n}) = (+\infty, +\infty)$. Moreover, $\text{Val}^*(v_0) = \text{Val}^*(v_4) = +\infty$. It easily follows that, given $i \in \Pi$ and $n \in \mathbb{N}$ such that $\rho_n \in V_i$ and $i \notin \text{Visit}(\rho_0 \dots \rho_n)$, we have that $\mathbf{QR}_i(\rho_{\geq n}) \leq \text{Val}^*(\rho_n) = +\infty$. In particular it means that ρ is the outcome of a Nash equilibrium in (\mathcal{G}, v_0) .

In the same way, let us consider the play $\rho' = v_0v_1v_6v_7v_2(v_0v_4)^\omega$, let us notice that both players reach their target set along ρ' (in v_2) and thus $\mathbf{QR}(\rho') = (4, 4)$. That means that we only have to prove these following inequalities:

$$\begin{aligned}
\text{QR}_2(v_0v_1v_6v_7v_2(v_0v_4)^\omega) &\leq \text{Val}^*(v_0) \\
\text{QR}_1(v_1v_6v_7v_2(v_0v_4)^\omega) &\leq \text{Val}^*(v_1) \\
\text{QR}_1(v_6v_7v_2(v_0v_4)^\omega) &\leq \text{Val}^*(v_6) \\
\text{QR}_1(v_7v_2(v_0v_4)^\omega) &\leq \text{Val}^*(v_7).
\end{aligned}$$

It is easy to check that those inequalities are true. Therefore, we conclude that ρ' is Visit Val^{*}-consistent and the outcome of a Nash equilibrium in (\mathcal{G}, v_0) .

Let us remark that ρ' seems to be a more relevant equilibrium outcome than ρ since both players visit their target set along ρ' and not along ρ .

Finally, we prove that the play $\pi = (v_0v_1v_3)^\omega$ cannot be the outcome of a Nash equilibrium in (\mathcal{G}, v_0) since it is not Visit Val^{*}-consistent. Indeed, we have that $v_1 \in V_1$ and $1 \notin \text{Visit}(v_0v_1)$ but $\text{QR}_1(v_1v_3(v_0v_1v_3)^\omega) = +\infty > \text{Val}^*(v_1) = 3$.

CHAPTER 7

WEAK SPE AND SPE OUTCOME CHARACTERIZATIONS

In Chapter 6, we have shown that for a wide class of multiplayer games, it is possible to exactly characterize the set of outcomes of Nash equilibria. This characterization relies on (i) the values of the vertices in the coalitional games and (ii) the notion of Val^* -consistent plays.

This chapter is mainly devoted to providing (i) a *weak SPE outcome characterization* for some class of multiplayer games based on the same spirit as the one of Chapter 6; (ii) a *compact representation* of weak SPEs; and (iii) identify multiplayer games such that, as an interesting by-product, this weak SPE outcome characterization allows also to characterize the set of outcomes of SPEs. It is divided as follows.

Firstly, in Section 7.1, in order to obtain our weak SPE outcome characterization, we have to define a labeling function which exactly characterizes if a play is the outcome of a weak SPE in the game. This labeling function, called λ^* , is obtained as the fixpoint of an iterative procedure.

Secondly, in Section 7.2, we define the notion of *good symbolic witness*. Roughly speaking, a good symbolic witness is a set of plays which respect some “good” properties such that given such a good symbolic witness we are able to build a weak SPE from those plays.

Thirdly, in Section 7.3, we prove that there exists a weak SPE with a given outcome if and only if this outcome is λ^* -consistent if and only if there exists a good symbolic witness that contains this outcome.

Fourthly, in Section 7.4, we show how from the general characterizations provided in the previous sections, we obtain characterizations of outcomes of weak SPEs in the multiplayer games of our interest in the rest of this document: multiplayer Boolean games with prefix-independent objective functions, multiplayer qualitative Reachability games, multiplayer Safety games, multiplayer quantitative Reachability games and multiplayer weighted Reachability games.

Finally, in Section 7.5, we explain how from the characterizations obtained for weak SPEs outcomes, we obtain characterizations of outcomes of SPEs in some kinds of games (*e.g.*, in Reachability games).

7.1 λ^k labeling functions for weak SPEs

In this section, we explain how from a sequence of well-defined labeling functions $(\lambda^k)_{k \in \mathbb{N}}$, we obtain a labeling function λ^* such that the set of weak SPEs outcomes exactly corresponds to the set of λ^* -consistent plays.

In Section 7.1.1, we begin by introducing some technical definitions and notations needed to define the sequence of labeling functions $(\lambda^k)_{k \in \mathbb{N}}$. Those labeling functions are properly defined in Section 7.1.2. Finally, in Section 7.1.3, we prove, under some hypotheses, a weak SPE outcome characterization based on the fixpoint of the sequence $(\lambda^k)_{k \in \mathbb{N}}$ of labeling functions.

7.1.1 (Good) local strongly prefix-linear constants set

Before being able to properly define the sequence of labeling functions $(\lambda^k)_{k \in \mathbb{N}}$, we have to introduce some technical notions and notations.

Those notions are related to the notion of strongly prefix-linear cost function in an arena (Definition 2.2.16) but with a more local point of view. Given a multiplayer game $\mathcal{G} = (A, \text{Cost})$, for each $i \in \Pi$ and each $(v, v') \in E$, we

consider the existence of a pair $(a^i(v, v'), b^i(v, v')) \in \mathbb{R} \times \mathbb{R}_0^+$ such that for all plays ρ beginning in v' , we have that $\text{Cost}_i(v\rho) = a^i(v, v') + b^i(v, v') \cdot \text{Cost}_i(\rho)$.

If such a pair exists for all $i \in \Pi$ and $(v, v') \in E$, one is fixed and added in a set of tuples of the form $(i, v, v', a^i(v, v'), b^i(v, v')) \in \Pi \times V \times V \times \mathbb{R} \times \mathbb{R}_0^+$, which is called a *good local strongly prefix-linear constants set*. We now define this concept formally.

Definition 7.1.1. Let $\mathcal{G} = (A, \text{Cost})$ be a multiplayer game. Given $i \in \Pi$, $(v, v') \in E$, $a^i(v, v') \in \mathbb{R}$ and $b^i(v, v') \in \mathbb{R}_0^+$.

If the following property is satisfied,

$$\forall \rho \in \text{Plays}(v'), \text{Cost}_i(v\rho) = a^i(v, v') + b^i(v, v') \cdot \text{Cost}_i(\rho) \quad (7.1)$$

then, we say that the tuple $(i, v, v', a^i(v, v'), b^i(v, v'))$ satisfies Property (7.1).

Definition 7.1.2 ((Good) local strongly prefix-linear constants set). Let $\mathcal{G} = (A, \text{Cost})$ be a multiplayer game.

1. A *local strongly prefix-linear constants set* (LSPLC set) associated with \mathcal{G} , denoted by $\mathcal{C}_{\mathcal{G}}$, is a subset of $\Pi \times V \times V \times \mathbb{R} \times \mathbb{R}_0^+$ such that, for each $i \in \Pi$ and each $(v, v') \in E$:
 - one tuple $(i, v, v', a, b) \in \Pi \times V \times V \times \mathbb{R} \times \mathbb{R}_0^+$ is fixed and added to $\mathcal{C}_{\mathcal{G}}$;
 - by convention, once a $\mathcal{C}_{\mathcal{G}}$ is built, $\alpha^i(v, v')$ and $\beta^i(v, v')$ refer respectively to a and b in the tuple (i, v, v', a, b) fixed for i and (v, v') .
2. An LSPLC set $\mathcal{C}_{\mathcal{G}}$ associated with \mathcal{G} is a *good local strongly prefix-linear constants set* if each $(i, v, v', \alpha^i(v, v'), \beta^i(v, v')) \in \mathcal{C}_{\mathcal{G}}$ satisfies Property (7.1).

If the cost functions of \mathcal{G} are strongly prefix-linear in A , there exists a good local strongly prefix-linear constants set associated with \mathcal{G} .

Lemma 7.1.3. *Let $\mathcal{G} = (A, \text{Cost})$ be a multiplayer game. If for all $i \in \Pi$, Cost_i is strongly prefix-linear in A , then there exists a good LSPLC set $\mathcal{C}_{\mathcal{G}}$ associated with \mathcal{G} . \square*

7.1.2 Definition of the λ^k -labeling functions

We are looking for a labeling function $\lambda : V \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ which imposes constraints on the plays and such that the set of plays which satisfy those constraints, that is the plays which are λ -consistent, is exactly the set of weak SPEs outcomes.

Roughly speaking, our algorithm works as follows: the labeling function λ that characterizes the set of SPEs outcomes is obtained from an initial labeling function that imposes no constraint and by iterating an operator that reinforces the constraints step after step, up to obtaining a fixpoint which is the required function λ . Thus, if λ^k is the labeling function computed at step k and $\Lambda^k(v)$, $v \in V$, the related sets of λ^k -consistent plays beginning in v , initially we have $\Lambda^0(v) = \text{Plays}(v)$, and step by step, the constraints imposed by λ^k become stronger and the sets $\Lambda^k(v)$ become smaller, until a fixpoint is reached.


Definition 7.1.4. Let $\mathcal{G} = (A, \text{Cost})$ be a multiplayer game, let $\lambda^k : V \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ for some $k \in \mathbb{N}$ be a labeling function. For all $v \in V$, we denote by $\Lambda^k(v)$ the set of plays which begin in v and which are λ^k -consistent. That is, for all $v \in V$:

$$\Lambda^k(v) = \{\rho \in \text{Plays}(v) \mid \rho \models \lambda^k\}.$$

We describe in Definition 7.1.5 how the labeling functions λ^k are iteratively defined and provide some explanations just after.

General definition of λ^k

Definition 7.1.5. Given $\mathcal{G} = (A, \text{Cost})$ a multiplayer game and given a good LSPLC set $\mathcal{C}_{\mathcal{G}}$ associated with \mathcal{G} . The labeling functions $\lambda^k : V \rightarrow$



Initially, we want a labeling function λ^0 that imposes no constraint on the plays, thereby we define λ^0 as the constant function $+\infty$. When it is updated, for each $v \in V$, if $v \in V_i$ for some $i \in \Pi$, the value $\lambda^{k+1}(v)$ represents what is the best cost that Player i can ensure for himself from v with a “one-shot” choice by only taking into account plays of $\Lambda^k(v')$ with $v' \in \text{Succ}(v)$.

Remark 7.1.6. Given $v \in V$, such that $v \in V_i$ for some $i \in \Pi$, if there exists an upper bound $B \in \mathbb{R}$ on the possible costs of plays beginning in v , *i.e.*, $\forall \rho \in \text{Plays}(v), \text{Cost}_i(\rho) \leq B$, then $\lambda^0(v)$ could be initialized as $\lambda^0(v) = B$. Indeed, with this convention λ^0 does not impose any constraint on plays either. As a result, each play beginning in v is also λ^0 -consistent.

Remark 7.1.7. Notice that to define the labeling functions λ^k for all $k \in \mathbb{N}$, we

previously fix a good local strongly prefix-linear constants set $\mathcal{C}_{\mathcal{G}}$ associated with the game \mathcal{G} of interest. The reason is that, if for all $i \in \Pi$, Cost_i is strongly prefix-linear, then such a good LSPLC set exists but is not always unique. Indeed, for example, if for a given $i \in \Pi$ and a given $(v, v') \in E$, we have that for all $\rho \in \text{Plays}(v')$, $\text{Cost}_i(\rho) = 0$ and if there exists $a \in \mathbb{R}$ and $b \in \mathbb{R}_0^+$ such that for all $\rho \in \text{Plays}(v')$, $\text{Cost}_i(v\rho) = a + b \cdot \text{Cost}_i(\rho)$ then every $b' \in \mathbb{R}_0^+$ satisfies that for all $\rho \in \text{Plays}(v')$, $\text{Cost}_i(v\rho) = a + b' \cdot \text{Cost}_i(\rho)$.

However, given $i \in \Pi$ and $v' \in V$, if there exist $\rho, \rho' \in \text{Plays}(v')$ such that $\text{Cost}_i(\rho) = C$ and $\text{Cost}_i(\rho') = C'$ for some C, C' such that (i) $C \neq C'$ and (ii) $C \neq 0$ and $C' \neq 0$, then there exists a unique $a \in \mathbb{R}$ and $b \in \mathbb{R}_0^+$ such that for all $\rho \in \text{Plays}(v')$, $\text{Cost}_i(v\rho) = a + b \cdot \text{Cost}_i(\rho)$.

Given $i \in \Pi$ and $(v, v') \in \text{Plays}(v')$ such that there exists $\rho \in \text{Plays}(v')$, $\text{Cost}_i(\rho) = C$ such that $C \neq 0$ and there exists $\rho' \in \text{Plays}(v')$ such that $\text{Cost}_i(\rho') = C'$ with $C' \neq 0$ and $C' \neq C$. If there exist $a \in \mathbb{R}$ and $b \in \mathbb{R}_0^+$ such that for all $\pi \in \text{Plays}(v')$, $\text{Cost}_i(v\pi) = a + b \cdot \text{Cost}_i(\pi)$, then $(\exists a' \in \mathbb{R} \wedge \exists b' \in \mathbb{R}_0^+ \text{ st. } \forall \pi \in \text{Plays}(v'), \text{Cost}_i(v\pi) = a' + b' \cdot \text{Cost}_i(\pi)) \implies a = a' \wedge b = b'$.

Let us assume that:

1. there exist $a \in \mathbb{R}$ and $b \in \mathbb{R}_0^+$ such that for all $\pi \in \text{Plays}(v')$, $\text{Cost}_i(v\pi) = a + b \cdot \text{Cost}_i(\pi)$; and
2. there exist $a' \in \mathbb{R}$ and $b' \in \mathbb{R}_0^+$ such that for all $\pi \in \text{Plays}(v')$, $\text{Cost}_i(v\pi) = a' + b' \cdot \text{Cost}_i(\pi)$.

That means in particular that:

$$\begin{cases} \text{Cost}_i(v\rho) = a + b \cdot C \\ \text{Cost}_i(v\rho) = a' + b' \cdot C \end{cases} \quad \begin{cases} \text{Cost}_i(v\rho') = a + b \cdot C' \\ \text{Cost}_i(v\rho') = a' + b' \cdot C' \end{cases}$$

Thus we obtain:

$$\begin{cases} (a - a') = (b' - b) \cdot C \\ (a - a') = (b' - b) \cdot C' \end{cases} .$$

and thus, we have that $(a - a') \cdot (C' - C) = 0$. Since $C' \neq C$ and $C \neq 0$, we conclude that $a' = a$ and $b' = b$.

Moreover, in Part III, we use these labeling functions in particular kinds of games thus we fix particular values to each $\alpha^i(v, v')$ and $\beta^i(v, v')$. More details and examples of the computation of $(\lambda^k)_{k \in \mathbb{N}}$ are provided in Section 7.4.

We conclude this section with two properties on the sequence of labeling functions $(\lambda_k)_{k \in \mathbb{N}}$. Lemma 7.1.8 states that the set of plays beginning in v which are λ^0 -consistent exactly corresponds to the set of plays beginning in v and Lemma 7.1.9 asserts that for all $v \in V$, the sequences $(\lambda^k(v))_{k \in \mathbb{N}}$ and $(\Lambda^k(v))_{k \in V}$ are non-increasing.

Lemma 7.1.8. *Let $\mathcal{G} = (A, \text{Cost})$ be a multiplayer game and let $(\lambda^k)_{k \in \mathbb{N}}$ a sequence of labeling functions as defined in Definition 7.1.5. For all $v \in V$, we have:*

$$\Lambda^0(v) = \text{Plays}(v).$$

Proof. Let $\mathcal{G} = (A, \text{Cost})$ be a multiplayer game and let $(\lambda^k)_{k \in \mathbb{N}}$ a sequence of labeling functions as defined in Definition 7.1.5. Let $v \in V$ be a vertex. In view of the definition of $\Lambda^0(v)$, one obviously have that $\Lambda^0(v) \subseteq \text{Plays}(v)$. Let $\rho \in \text{Plays}(v)$, we have to prove that $\rho \models \lambda^0$. That is:

$$\forall i \in \Pi, \forall n \in \mathbb{N}, (\rho_n \in V_i \implies \text{Cost}_i(\rho_{\geq n}) \leq \lambda^0(\rho_n)).$$

The fact that $\lambda^0(v) = +\infty$ for all $v \in V$ concludes the proof. □

Lemma 7.1.9. *Let $\mathcal{G} = (A, \text{Cost})$ be a multiplayer game and let $(\lambda^k)_{k \in \mathbb{N}}$ be a sequence of labeling functions as defined in Definition 7.1.5. For all $v \in V$, the sequences $(\lambda^k(v))_{k \in \mathbb{N}}$ and $(\Lambda^k(v))_{k \in V}$ are non-increasing.*

In order to prove Lemma 7.1.9, we prove the following intermediate result.

Lemma 7.1.10. *Let $\mathcal{G} = (A, \text{Cost})$ be a multiplayer game and let $(\lambda^k)_{k \in \mathbb{N}}$ be a sequence of labeling functions as defined in Definition 7.1.5. Given $k \in \mathbb{N}$, if for all $v \in V$, $\lambda^{k+1}(v) \leq \lambda^k(v)$, then for all $v \in V$, $\Lambda^{k+1}(v) \subseteq \Lambda^k(v)$.*

Proof. Let $\mathcal{G} = (A, \text{Cost})$ be a multiplayer game and let $(\lambda^k)_{k \in \mathbb{N}}$ be a sequence of labeling functions as defined in Definition 7.1.5.

Let $k \in \mathbb{N}$. Let us assume that for all $v \in V$, $\lambda^{k+1}(v) \leq \lambda^k(v)$. Let $v \in V$ and $\rho \in \Lambda^{k+1}(v)$ be a λ^{k+1} -consistent play. That means that:

$$\forall i \in \Pi, \forall n \in \mathbb{N}, \rho_n \in V_i \implies \text{Cost}_i(\rho_{\geq n}) \leq \lambda^{k+1}(\rho_n). \quad (7.2)$$

Since for all $v \in V$, we have that $\lambda^{k+1}(v) \leq \lambda^k(v)$, from (7.2) we obtain:

$$\forall i \in \Pi, \forall n \in \mathbb{N}, \rho_n \in V_i \implies \text{Cost}_i(\rho_{\geq n}) \leq \lambda^k(\rho_n).$$

That proves that $\rho \in \Lambda^k(v)$. □

Proof of Lemma 7.1.9. Let $\mathcal{G} = (A, \text{Cost})$ be a multiplayer game and let $(\lambda^k)_{k \in \mathbb{N}}$ be a sequence of labeling functions as defined in Definition 7.1.5.

We will prove that for all $k \in \mathbb{N}$, for all $v \in V$, $\lambda^{k+1}(v) \leq \lambda^k(v)$. We prove this property by induction on k .

If $k = 0$: let $v \in V$, then the inequality $\lambda^1(v) \leq \lambda^0(v) = +\infty$ is satisfied.

Let us now assume that for all $n \in \mathbb{N}$ such that $0 \leq n \leq k$: for all $v \in V$, $\lambda^{n+1}(v) \leq \lambda^n(v)$ is satisfied.

Let us prove it remains true for $k + 1$. Let $v \in V$, we have to prove that $\lambda^{k+2}(v) \leq \lambda^{k+1}(v)$.

Let us assume that $v \in V_i$ for some $i \in \Pi$. By definition of $\lambda^{k+1}(v)$, there exists $v' \in \text{Succ}(v)$ such that $\lambda^{k+1}(v) = \alpha^i(v, v') + \beta^i(v, v') \cdot \sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v)\}$.

From the definition of λ^{k+2} follows:

$$\begin{aligned}\lambda^{k+2}(v) &= \min_{u \in \text{Succ}(v)} \left\{ \alpha^i(v, u) + \beta^i(v, u) \cdot \sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^{k+1}(u)\} \right\} \\ &\leq \alpha^i(v, v') + \beta^i(v, v') \cdot \sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^{k+1}(v')\}\end{aligned}\quad (7.3)$$

By induction hypothesis, for all $u \in V$, $\lambda^{k+1}(u) \leq \lambda^k(u)$. Thereby, by Lemma 7.1.10, we have that $\Lambda^{k+1}(v') \subseteq \Lambda^k(v')$. Which allows to conclude, from (7.3):

$$\begin{aligned}\lambda^{k+2}(v) &\leq \alpha^i(v, v') + \beta^i(v, v') \cdot \sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^{k+1}(v')\} \\ &\leq \alpha^i(v, v') + \beta^i(v, v') \cdot \sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v')\} \quad (\beta^i(v, v') \geq 0) \\ &= \lambda^{k+1}(v).\end{aligned}$$

It remains to prove that for all $k \in \mathbb{N}$ and all $v \in V$, $\Lambda^{k+1}(v) \subseteq \Lambda^k(v)$. That result follows from the previous proof and Lemma 7.1.10. \square

7.1.3 Weak SPE outcome characterization

In the previous section, we have explained how a sequence of labeling functions $(\lambda^k)_{k \in \mathbb{N}}$ is defined. We have also mentioned that this sequence of labeling functions has to reach a *fixpoint* to exactly characterize the set of weak SPEs outcomes. This notion of *fixpoint* is properly defined in the following definition.

Definition 7.1.11 (Existence of a fixpoint). Let $\mathcal{G} = (A, \text{Cost})$ be a multiplayer game and let $(\lambda^k)_{k \in \mathbb{N}}$ be a sequence of labeling function as defined in Definition 7.1.5. If there exists $k \in \mathbb{N}$ such that for all $m \in \mathbb{N}$ and for all $v \in V$:

$$\lambda^{k+m}(v) = \lambda^k(v). \quad (7.4)$$

We say that $(\lambda^k)_{k \in \mathbb{N}}$ reaches a *fixpoint*.

Moreover, the least natural number which satisfies the equality (7.4) is called *the fixpoint* of $(\lambda^k)_{k \in \mathbb{N}}$ and is denoted by k^* . In order to ease the notation

we often write λ^* (resp. $\Lambda^*(v)$ for $v \in V$) instead of λ^{k^*} (resp. $\Lambda^{k^*}(v)$ for $v \in V$).

Remark 7.1.12. Notice that in practice, we only have to find $k^* \in \mathbb{N}$ (as small as possible) such that for all $v \in V$:

$$\lambda^{k^*+1}(v) = \lambda^{k^*}(v).$$

Indeed, if such a natural number exists, it implies that for all $m \in \mathbb{N}$ and for all $v \in V$:

$$\lambda^{k^*+m}(v) = \lambda^{k^*}(v).$$

Additionally, for all $i \in \Pi$, for all $k \in \mathbb{N}$ and all $v \in V$ such that $\Lambda^k(v) \neq \emptyset$, we have to impose that the supremum of the Player i 's costs of the plays beginning in v which are λ^k -consistent is a maximum. If it is the case, we say that $(\lambda^k)_{k \in \mathbb{N}}$ satisfies the *existence of maxima property*.

Definition 7.1.13 (Existence of maxima property). Let $\mathcal{G} = (\mathbf{A}, \text{Cost})$ be a multiplayer game and let $(\lambda^k)_{k \in \mathbb{N}}$ be a sequence of labeling functions as defined in Definition 7.1.5. If for all $i \in \Pi$, for all $v \in V$, for all $k \in \mathbb{N}$ such that $\Lambda^k(v) \neq \emptyset$, there exists $\rho \in \Lambda^k(v)$ such that $\text{Cost}_i(\rho) = \sup\{\text{Cost}_i(\rho') \mid \rho' \in \Lambda^k(v)\}$ then we say that $(\lambda^k)_{k \in \mathbb{N}}$ satisfies the *existence of maxima property*.

Moreover, for all $i \in \Pi$, for all $v \in V$ and for all $k \in \mathbb{N}$, we can write:

$$\sup\{\text{Cost}_i(\rho') \mid \rho' \in \Lambda^k(v)\} = \max\{\text{Cost}_i(\rho') \mid \rho' \in \Lambda^k(v)\}.$$

We are now able to enunciate our weak SPE outcome characterization. Given an initialized multiplayer game $(\mathcal{G}, v_0) = (\mathbf{A}, \text{Cost})$ and a sequence of labeling functions $(\lambda^k)_{k \in \mathbb{N}}$ as defined in Definition 7.1.5, if (i) for each $i \in \Pi$, Cost_i is strongly prefix-linear in \mathbf{A} ; (ii) $(\lambda^k)_{k \in \mathbb{N}}$ satisfies the *existence of maxima property* and (iii) $(\lambda^k)_{k \in \mathbb{N}}$ reaches a fixpoint and its fixpoint is $k^* \in \mathbb{N}$, then the set of plays beginning in v_0 which are λ^* -consistent is equal to the set of plays which are outcomes of weak SPEs in (\mathcal{G}, v_0) .

Proposition 7.1.14 (weak SPE outcome characterization). *Let $(\mathcal{G}, v_0) = (A, \text{Cost})$ be an initialized multiplayer game and let $(\lambda^k)_{k \in \mathbb{N}}$ be a sequence of labeling functions as defined in Definition 7.1.5 such that:*

- (H₁) *For all $i \in \Pi$, Cost_i is strongly prefix-linear in A ;*
- (H₂) *$(\lambda^k)_{k \in \mathbb{N}}$ satisfies the existence of maxima property;*
- (H₃) *$(\lambda^k)_{k \in \mathbb{N}}$ reaches a fixpoint and its fixpoint is $k^* \in \mathbb{N}$.*

Then, given $\rho \in \text{Plays}(v_0)$, the following assertions are equivalent:

1. *There exists a weak SPE σ in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = \rho$;*
2. *ρ is λ^* -consistent.*

The proof of Proposition 7.1.14 is provided in Section 7.3. Examples of utilizations of this characterization of the outcomes of weak SPEs may be found in Section 7.4.

7.2 (Good) Symbolic witnesses for weak SPEs

In this section, we introduce the notion of (good) symbolic witness and explain how from this concept we can characterize the set of weak SPEs outcomes.

Roughly speaking, a *symbolic witness* is a finite set of plays and a *good symbolic witness* is a symbolic witness such that the plays respect some good property. The intuition behind this property is the following: each time you consider a play ρ in the good symbolic witness no player has an incentive to deviate and to follow another play of the good symbolic witness. We will see that a good symbolic witness is all we need to build a weak SPE. Moreover, if each play of the good symbolic witness is a lasso, the weak SPE requires finite memory and this good symbolic witness provides a finite representation of this equilibrium.

In Section 7.2.1, we formally define the notion of good symbolic witness and, in Section 7.2.2, we provide a weak SPE outcome characterization based

on the notion of good symbolic witness.

7.2.1 (Good) Symbolic witnesses

Given an initialized multiplayer game (\mathcal{G}, v_0) , we depict by \mathcal{I} the set of pairs (player, vertex) such that $(i, v) \in \mathcal{I}$ means that v is reachable in one step from a vertex of Player i which is itself reachable from v_0 . An additional pair $(0, v_0)$ is added for the initial vertex, notice that we assume that 0 does not represent a player.

Definition 7.2.1. Let (\mathcal{G}, v_0) be an initialized multiplayer game. We define the set \mathcal{I} as follows:

$$\mathcal{I} = \{(0, v_0)\} \cup \{(i, v') \in \Pi \times V \mid \exists v \in \text{Succ}^*(v_0) \cap V_i \text{ st. } (v, v') \in E\}.$$

To each pair $(i, v) \in \mathcal{I}$, a play $\rho^{(i,v)}$ is associated in order to obtain a set of plays. This set contains at most $|\Pi| \cdot |V| + 1$ plays and is called a *symbolic witness*.

Definition 7.2.2 (Symbolic witness). Given an initialized game (\mathcal{G}, v_0) , a *symbolic witness* $\mathcal{P} \subseteq \text{Plays}$ is a set of plays such that for all $(i, v) \in \mathcal{I}$, a play $\rho^{(i,v)} \in \text{Plays}(v)$ beginning in v is fixed and added to \mathcal{P} . In particular, $|\mathcal{P}| \leq |\Pi| \cdot |V| + 1$.

By convention, when we take a play $\rho^{(i,v)} \in \mathcal{P}$ for some $(i, v) \in \mathcal{I}$, we assume that this is the play which has been previously fixed for the pair (i, v) when the symbolic witness \mathcal{P} was built.

A symbolic witness \mathcal{P} is a representation of some strategy profile σ in (\mathcal{G}, v_0) . It is a finite set of plays that represent some subgame outcomes of σ : the play $\rho^{(0,v_0)}$ of \mathcal{P} represents the outcome $\langle \sigma \rangle_{v_0}$, and each other play $\rho^{(i,v')}$ represents the subgame outcome $\langle \sigma_{|h} \rangle_{v'}$ for some particular histories $hv' \in \text{Hist}(v_0)$. The index i records that Player i can move from v (the last vertex of h) to v' (with the convention that $i = 0$ for the outcome $\langle \sigma \rangle_{v_0}$).

When σ is a weak SPE, the related symbolic witness \mathcal{P} is good, that is, its plays avoid profitable one-shot deviations between them.

Definition 7.2.3 (Good symbolic witness). Let \mathcal{P} be a symbolic witness. If for all $\rho^{(j,w)} \in \mathcal{P}$, for all $i \in \Pi$, for all $k \in \mathbb{N}$, for all $v' \in \text{Succ}(\rho_k^{(j,w)})$,

$$\left(\rho_k^{(j,w)} \in V_i \implies \text{Cost}_i(\rho_{\geq k}^{(j,w)}) \leq \text{Cost}_i(\rho_k^{(j,w)} \rho^{(i,v')}) \right), \quad (7.5)$$

then \mathcal{P} is called a *good symbolic witness*.

In Figure 7.1, we provide an illustration of the condition which should be satisfied to be a good symbolic witness (see Definition 7.2.3).




Figure 7.1: Illustration of condition (7.5) of Definition 7.2.3.

7.2.2 Weak SPE outcome characterization

Given an initialized multiplayer game $(\mathcal{G}, v_0) = (A, \text{Cost})$ such that for all $i \in \Pi$, the cost function Cost_i of Player i is strongly prefix-linear, given a play $\rho \in \text{Plays}(v_0)$, there is an equivalence between the existence of a good symbolic witness \mathcal{P} such that $\rho^{(0,v_0)} = \rho$ and the existence of a weak SPE σ in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = \rho$.

Proposition 7.2.4. *Let $(\mathcal{G}, v_0) = (A, \text{Cost})$ be an initialized multiplayer game and let $(\lambda^k)_{k \in \mathbb{N}}$ be a sequence of labeling functions as defined in Definition 7.1.5 such that:*

(H₁) *For all $i \in \Pi$, Cost_i is strongly prefix-linear in A ;*

(H₂) *$(\lambda^k)_{k \in \mathbb{N}}$ satisfies the existence of maxima property;*

(H₃) *$(\lambda^k)_{k \in \mathbb{N}}$ reaches a fixpoint and its fixpoint is $k^* \in \mathbb{N}$.*

Then, given $\rho \in \text{Plays}(v_0)$, the following assertions are equivalent:

1. *There exists a weak SPE σ in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = \rho$;*
2. *There exists a good symbolic witness \mathcal{P} such that the play $\rho^{(0, v_0)} = \rho$ with $\rho^{(0, v_0)} \in \mathcal{P}$.*

The proof of Proposition 7.2.4 is provided in Section 7.3. An example of good symbolic witness associated with its related weak SPE may be found in Section 7.4.

If we impose that each play of a symbolic witness is a lasso, we obtain a *finite symbolic witness*: a finite good symbolic witness \mathcal{P} is thus a *compact representation* of a weak SPE with finite memory.

Definition 7.2.5 (Finite symbolic witness). Let \mathcal{P} be a symbolic witness. If for all $\rho^{(i, v)} \in \mathcal{P}$, $\rho^{(i, v)}$ is a lasso (there exist $h \in \text{Hist}(v)$, $\ell \in \text{Hist}$ such that $h\ell \in \text{Hist}(v)$ and $\rho^{(i, v)} = h\ell^\omega$), we say that \mathcal{P} is a *finite symbolic witness*.

Notice that we can consider in the same way finite good symbolic witnesses.

Proposition 7.2.6. *Let $(\mathcal{G}, v_0) = (A, \text{Cost})$ be an initialized multiplayer game such that:*

(H₁) *For all $i \in \Pi$, Cost_i is strongly prefix-linear in A ;*

If there exists a finite good symbolic witness \mathcal{P} such that, for some $L \in \mathbb{N}$: for

all $\rho^{(i,v)} \in \mathcal{P}$, there exist $h \in \text{Hist}(v)$ and $\ell \in \text{Hist}$ such that: $h\ell \in \text{Hist}(v)$, $|h\ell| \leq L$ and $\rho^{(i,v)} = h\ell^\omega$, then, there exists a finite-memory weak SPE σ in (\mathcal{G}, v_0) with memory size in $\mathcal{O}(|\Pi| \cdot |V| \cdot L)$ and such that $\langle \sigma \rangle_{v_0} = \rho^{(0,v_0)}$.

The proof of Proposition 7.2.6 is provided in Section 7.3.

Remark 7.2.7. One could wonder if hypotheses (H_2) and (H_3) may be omitted in the statement of Proposition 7.2.4.

7.3 Proofs of the weak SPE outcome characterizations

In order to prove the weak SPE outcome characterizations of Section 7.1.3 and Section 7.2.2, we prove the following theorem.

Theorem 7.3.1. *Let $(\mathcal{G}, v_0) = (A, \text{Cost})$ be an initialized multiplayer game and let $(\lambda^k)_{k \in \mathbb{N}}$ be a sequence of labeling functions as defined in Definition 7.1.5 such that:*

- (H_1) *For all $i \in \Pi$, Cost_i is strongly prefix-linear in A ;*
- (H_2) *$(\lambda^k)_{k \in \mathbb{N}}$ satisfies the existence of maxima property;*
- (H_3) *$(\lambda^k)_{k \in \mathbb{N}}$ reaches a fixpoint and its fixpoint is $k^* \in \mathbb{N}$.*

Then, given $\rho^ \in \text{Plays}(v_0)$, the following assertions are equivalent:*

1. *There exists a weak SPE σ in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = \rho^*$;*
2. *ρ^* is λ^* -consistent, i.e., $\rho^* \in \Lambda^*(v_0)$;*
3. *There exists a good symbolic witness \mathcal{P} such that the play $\rho^{(0,v_0)} = \rho^*$ with $\rho^{(0,v_0)} \in \mathcal{P}$.*

Before this proof, we need to prove a technical result about the non emptiness of $\Lambda^*(v_0)$.

Lemma 7.3.2. *Let $(\mathcal{G}, v_0) = (A, \text{Cost})$ be an initialized multiplayer game and let $(\lambda^k)_{k \in \mathbb{N}}$ be a sequence of labeling functions as defined in Definition 7.1.5.*

If $(\lambda^k)_{k \in \mathbb{N}}$ reaches a fixpoint and its fixpoint is $k^ \in \mathbb{N}$, then we have that:*

$$\rho \in \Lambda^*(v_0) \iff (\rho \in \Lambda^*(v_0) \wedge \forall v \in \text{Succ}^*(v_0), \Lambda^*(v) \neq \emptyset).$$

Proof. Let $(\mathcal{G}, v_0) = (A, \text{Cost})$ be an initialized multiplayer game and let $(\lambda^k)_{k \in \mathbb{N}}$ be a sequence of labeling functions as defined in Definition 7.1.5 and let us assume that $(\lambda^k)_{k \in \mathbb{N}}$ reaches a fixpoint and its fixpoint is $k^* \in \mathbb{N}$. We only prove the \implies implication.

Let $\rho \in \Lambda^*(v_0)$ be a λ^* -consistent play beginning in v_0 . We assume by contradiction that there exists $v \in \text{Succ}^*(v_0)$ such that $\Lambda^*(v) = \emptyset$.

We prove, by induction on n , that for all $n \in \mathbb{N}$, if $u \in \text{Pred}^n(v)$, then $\Lambda^*(u) = \emptyset$. With for all $v \in V$, $\text{Pred}^n(v)$ is defined by induction as follows:

$$\begin{cases} \text{Pred}^0(v) = \{v\} \\ \text{Pred}^{n+1}(v) = \{u \in V \mid \exists (u, v') \in E \text{ st. } v' \in \text{Pred}^n(v)\} \end{cases}.$$

For $n = 0$, $\text{Pred}^0(v) = \{v\}$, thus for all $u \in \text{Pred}^0(v)$, $\Lambda^*(u) = \emptyset$ is true.

Let us assume the property is true for all $\ell \in \mathbb{N}$ such that $0 \leq \ell \leq n$, and let us prove it remains true for $n + 1$.

Let $u \in \text{Pred}^{n+1}(v)$, that means that there exists $(u, v') \in E$ such that $v' \in \text{Pred}^n(v)$. By induction hypothesis, we have that $\Lambda^*(v') = \emptyset$.

By definition of λ^* and by assuming that $u \in V_i$ for some $i \in \Pi$, we have:

$$\begin{aligned} \lambda^*(u) &= \lambda^{k^*+1}(u) \\ &= \min_{u' \in \text{Succ}(u)} \left\{ \alpha^i(u, u') + \beta^i(u, u') \cdot \sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^{k^*}(u')\} \right\} \\ &= -\infty \end{aligned}$$

The last equality holds because $v' \in \text{Succ}(u)$ and since $\Lambda^*(v') = \emptyset$, we have that $\alpha^i(u, v') + \beta^i(u, v') \cdot \sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^{k^*}(v')\} = -\infty$.

We can conclude that $\Lambda^*(u) = \emptyset$. Indeed, otherwise it would mean that there exists $\rho \in \text{Plays}(u)$ such that $\rho \models \lambda^*$. Thus in particular: $\text{Cost}_i(\rho) \leq \lambda^*(u) = -\infty$. That leads to a contradiction.

To conclude the proof, since $v \in \text{Succ}^*(v_0)$, there exists $n \in \mathbb{N}$ such that $v_0 \in \text{Pred}^n(v)$. Thus, $\Lambda^*(v_0) = \emptyset$. That is the contradiction we expected to find. \square

Proof of Theorem 7.3.1. Let $(\mathcal{G}, v_0) = (A, \text{Cost})$ be an initialized multiplayer game and let $(\lambda^k)_{k \in \mathbb{N}}$ be a sequence of labeling functions as defined in Definition 7.1.5 such that:

- (H₁) For all $i \in \Pi$, Cost_i is strongly prefix-linear in A;
- (H₂) $(\lambda^k)_{k \in \mathbb{N}}$ satisfies the *existence of maxima property*;
- (H₃) $(\lambda^k)_{k \in \mathbb{N}}$ reaches a fixpoint and its fixpoint is $k^* \in \mathbb{N}$.

Let $\rho^* \in \text{Plays}(v_0)$ be a play in (\mathcal{G}, v_0) . We will prove $1 \implies 2 \implies 3 \implies 1$.

(1 \implies 2): Let σ be a weak SPE in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = \rho^*$. We prove the following property:

$$\forall k \in \mathbb{N}, \forall hv \in \text{Hist}(v_0), \langle \sigma_{\uparrow h} \rangle_v \in \Lambda^k(v).$$

It will be in particular true for $k = k^*$. Thus, it will follow that $\langle \sigma \rangle_{v_0} \in \Lambda^*(v_0)$ and for all $hv \in \text{Hist}(v_0)$, $\langle \sigma_{\uparrow h} \rangle_v \in \Lambda^*(v)$. In particular, this latter result proves that $\Lambda^*(v) \neq \emptyset$ for all $v \in \text{Succ}^*(v_0)$.

We proceed by induction on k .

If $k = 0$: let $hv \in \text{Hist}(v_0)$, the property is true since $\langle \sigma_{\uparrow h} \rangle_v \in \text{Plays}(v) = \Lambda^0(v)$ by Lemma 7.1.8.

Let us assume that the property is true for all $n \in \mathbb{N}$ such that $0 \leq n \leq k$.

Let us prove it remains true for $k + 1$:

By contradiction, let us assume that:

$$\exists hv \in \text{Hist}(v_0) \text{ such that } \langle \sigma_{\uparrow h} \rangle_v \notin \Lambda^{k+1}(v). \quad (7.6)$$

Let $\rho = \langle \sigma_{|h} \rangle_v$.

That means by definition of λ^{k+1} -consistency:

$$\exists i \in \Pi, \exists n \in \mathbb{N}, \text{ such that } \rho_n \in V_i \text{ and } \text{Cost}_i(\rho_{\geq n}) > \lambda^{k+1}(\rho_n). \quad (7.7)$$

By definition of λ^{k+1} ,

$$\lambda^{k+1}(\rho_n) = \min_{v' \in \text{Succ}(\rho_n)} \left\{ \alpha^i(\rho_n, v') + \beta^i(\rho_n, v') \cdot \sup\{\text{Cost}_i(\rho') \mid \rho' \in \Lambda^k(v')\} \right\}$$

Then, there exists $v' \in \text{Succ}(\rho_n)$ such that:

$$\lambda^{k+1}(\rho_n) = \alpha^i(\rho_n, v') + \beta^i(\rho_n, v') \cdot \sup\{\text{Cost}_i(\rho') \mid \rho' \in \Lambda^k(v')\}. \quad (7.8)$$

Let $h' = h\rho_{\leq n-1}^a$.

Let us prove that there exists a strategy τ_i of Player i in $(\mathcal{G}_{|h'}, \rho_n)$ which is a one-shot deviating strategy from $\sigma_{|h'}$ and a profitable deviation for Player i w.r.t. $\sigma_{|h'}$ in $(\mathcal{G}_{|h'}, \rho_n)$. That proves that σ is not a very weak SPE and thus not a weak SPE by Proposition 2.4.18. That is a contradiction.

Let τ_i be a strategy of Player i in $(\mathcal{G}_{|h'}, \rho_n)$ which is a one-shot deviating strategy from $\sigma_{|h'}$ such that $\tau_i(\rho_n) = v'$.

Let us prove that τ_i is a profitable deviation for Player i w.r.t. $\sigma_{|h'}$ in $(\mathcal{G}_{|h'}, \rho_n)$. That is :

$$\text{Cost}_i(h' \langle \tau_i, \sigma_{-i|h'} \rangle_{\rho_n}) < \text{Cost}_i(h' \langle \sigma_{|h'} \rangle_{\rho_n}).$$

By induction hypothesis, we have that $\langle \sigma_{|h' \rho_n} \rangle_{v'} \in \Lambda^k(v')$. It follows:

$$\text{Cost}_i(\langle \sigma_{|h' \rho_n} \rangle_{v'}) \leq \sup\{\text{Cost}_i(\rho') \mid \rho' \in \Lambda^k(v')\}.$$

Since $\beta^i(\rho_n, v') \in \mathbb{R}_0^+$:

$$\begin{aligned} & \alpha^i(\rho_n, v') + \beta^i(\rho_n, v') \cdot \text{Cost}_i(\langle \sigma_{|h' \rho_n} \rangle_{v'}) \\ & \leq \\ & \alpha^i(\rho_n, v') + \beta^i(\rho_n, v') \cdot \sup\{\text{Cost}_i(\rho') \mid \rho' \in \Lambda^k(v')\}. \end{aligned}$$

Thus:

$$\begin{aligned} \text{Cost}_i(\rho_n \langle \sigma_{|h' \rho_n} \rangle_{v'}) &\leq \lambda^{k+1}(\rho_n) && \text{(by (7.8))} \\ &< \text{Cost}_i(\rho_{\geq n}) && \text{(by (7.7)).} \end{aligned}$$

As Cost_i is strongly prefix-linear and by Lemma 2.2.20, it follows that:

$$\text{Cost}_i(h' \rho_n \langle \sigma_{|h' \rho_n} \rangle_{v'}) < \text{Cost}_i(h' \rho_{\geq n}). \quad (7.9)$$

Since τ_i is a one-shot deviating strategy, we have that $\text{Cost}_i(h' \langle \tau_i, \sigma_{-i|h'} \rangle_{\rho_n}) = \text{Cost}_i(h' \rho_n \langle \sigma_{|h' \rho_n} \rangle_{v'})$. Moreover, $\rho_{\geq n} = \langle \sigma_{|h'} \rangle_{\rho_n}$.

These two facts and (7.9) conclude the proof.

(2 \implies 3): Let us assume that $\rho^* \in \Lambda^*(v_0)$ and that for all $v \in \text{Succ}^*(v_0)$, $\Lambda^*(v) \neq \emptyset$ (Lemma 7.3.2).

We have to build a symbolic witness \mathcal{P} and then to prove it is good.

We choose $\rho^{(0, v_0)} = \rho^*$.

Then for each $(i, v) \in \mathcal{I}$, we choose $\rho^{(i, v)} \in \Lambda^*(v)$ such that

$$\text{Cost}_i(\rho^{(i, v)}) = \max\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^*(v)\} \quad (7.10)$$

Notice that such a play $\rho^{(i, v)}$ exists by hypothesis.

Now we have to prove that \mathcal{P} is a good symbolic witness.

Let $\rho^{(i, v)} \in \mathcal{P}$, let $j \in \Pi$, let $k \in \mathbb{N}$, let $v' \in \text{Succ}(\rho_k^{(i, v)})$. We assume that $\rho_k^{(i, v)} \in V_j$, we have to prove that:

$$\text{Cost}_j(\rho_{\geq k}^{(i, v)}) \leq \text{Cost}_j(\rho_k^{(i, v)} \rho^{(j, v')}).$$

To ease the notations hereunder, we set $a = \alpha^j(\rho_k^{(i, v)}, v')$ and $b = \beta^j(\rho_k^{(i, v)}, v')$. By construction $\rho^{(i, v)} \in \Lambda^*(v)$ and k^* is the fixpoint of $(\lambda^k)_{k \in \mathbb{N}}$, thus:

$$\begin{aligned}
\text{Cost}_j(\rho_{\geq k}^{(i,v)}) &\leq \lambda^*(\rho_k^{(i,v)}) = \lambda^{k^*+1}(\rho_k^{(i,v)}) \\
&= \min_{w \in \text{Succ}(\rho_k^{(i,v)})} \left\{ \alpha^j(\rho_k^{(i,v)}, w) + \beta^j(\rho_k^{(i,v)}, w) \cdot \max\{\text{Cost}_j(\pi) \mid \pi \in \Lambda^*(w)\} \right\} \\
&\leq a + b \cdot \max\{\text{Cost}_j(\pi) \mid \pi \in \Lambda^*(v')\} && (v' \in \text{Succ}(\rho_k^{(i,v)})) \\
&= a + b \cdot \text{Cost}_j(\rho^{(j,v')}) && (\text{By (7.10) and } (j, v') \neq (0, v_0)) \\
&= \text{Cost}_j(\rho_k^{(i,v)} \rho^{(j,v')}).
\end{aligned}$$

(3) \implies 1): Let us assume that \mathcal{P} is a good symbolic witness with $\rho^{(0,v_0)} = \rho^*$. We will build from \mathcal{P} a weak SPE σ in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = \rho^{(0,v_0)} = \rho^*$.

We define σ by induction on the subgames of (\mathcal{G}, v_0) . First, we define σ such that $\langle \sigma \rangle_{v_0} = \rho^{(0,v_0)}$. Then, let $hvv' \in \text{Hist}(v_0)$ such that $\langle \sigma_{\upharpoonright h} \rangle_v$ is already properly defined but not yet $\langle \sigma_{\upharpoonright hv} \rangle_{v'}$. By assuming that $v \in V_i$, for some $i \in \Pi$, we extend the definition of σ in the following way:

$$\langle \sigma_{\upharpoonright hv} \rangle_{v'} = \rho^{(i,v')}.$$

We prove that σ is a (very) weak SPE in (\mathcal{G}, v_0) . Let $hv \in \text{Hist}(v_0)$, by assuming that $v \in V_i$ for some $i \in \Pi$, let τ_i be a one-shot deviating strategy from $\sigma_{\upharpoonright h}$ in $(\mathcal{G}_{\upharpoonright h}, v)$. Thus, in particular $\tau_i(v) = v' \neq \sigma_{i\upharpoonright h}(v)$ for some $v' \in \text{Succ}(v)$. We have to prove:

$$\text{Cost}_i(h\langle \sigma_{\upharpoonright h} \rangle_v) \leq \text{Cost}_i(h\langle \tau_i, \sigma_{-i\upharpoonright h} \rangle_v).$$

By construction, since $v' \neq \sigma_{i\upharpoonright h}(v)$, $\langle \sigma_{\upharpoonright hv} \rangle_{v'} = \rho^{(i,v')}$.

There exists $h' \leq h$, such that $h' \in \text{Hist}_j(v_0)$ for some $j \in \Pi$ and $h\langle \sigma_{\upharpoonright h} \rangle_v = h'\rho^{(j,w)}$ for some $w \in V$. Moreover, there exists $k \in \mathbb{N}$ such that $\rho_k^{(j,w)} = v$.

Thus, by the property of good symbolic witness:

$$\begin{aligned}
\text{Cost}_i(\rho_{\geq k}^{(j,w)}) &\leq \text{Cost}_i(\rho_k^{(j,w)} \rho^{(i,v')}) \\
&= \text{Cost}_i(v \rho^{(i,v')}) \\
&= \text{Cost}_i(v \langle \sigma_{\upharpoonright hv} \rangle_{v'})
\end{aligned}$$

Since, Cost_i is strongly prefix-linear by hypothesis and thanks to Lemma 2.2.20, it follows that:

$$\text{Cost}_i(h\rho_{\geq k}^{(j,w)}) \leq \text{Cost}_i(hv\langle\sigma_{|hv}\rangle_{v'}).$$

As τ_i is a one-shot deviating strategy from $\sigma_{|h}$, we have that $\text{Cost}_i(hv\langle\sigma_{|hv}\rangle_{v'}) = \text{Cost}_i(h\langle\tau_i, \sigma_{-i|h}\rangle_v)$. This latter equality allows us to conclude:

$$\text{Cost}_i(h\langle\sigma_{|h}\rangle_v) \leq \text{Cost}_i(h\langle\tau_i, \sigma_{-i|h}\rangle_v).$$

□

^aIf $n = 0$, let $h' = h$.

Proof of Proposition 7.1.14. Follows directly from Theorem 7.3.1. □

Proof of Proposition 7.2.4. Follows directly from Theorem 7.3.1. □

Proof of Proposition 7.2.6. Let $(\mathcal{G}, v_0) = (A, \text{Cost})$ be an initialized multi-player game such that:

(H₁) For all $i \in \Pi$, Cost_i is strongly prefix-linear in A;

Let us assume that there exists a finite good symbolic witness \mathcal{P} such that, for some $L \in \mathbb{N}$: for all $\rho^{(i,v)} \in \mathcal{P}$, there exist $h \in \text{Hist}(v)$ and $\ell \in \text{Hist}$ such that: $h\ell \in \text{Hist}(v)$, $|h\ell| \leq L$ and $\rho^{(i,v)} = h\ell^\omega$.

In view of Theorem 7.3.1, we already know that it is possible to build from \mathcal{P} a weak SPE σ in (\mathcal{G}, v_0) such that $\langle\sigma\rangle_{v_0} = \rho^{(0,v_0)}$. It remains to prove that if \mathcal{P} is a finite good symbolic witness then σ is a finite-memory strategy profile.

Let us prove that σ as built in the proof (3 \implies 1) of Theorem 7.3.1 is finite-memory with memory size in $\mathcal{O}(|\Pi| \cdot |V| \cdot L)$.

For each $\rho^{(i,v)} \in \mathcal{P}$, we know that there exist $h^{(i,v)} \in \text{Hist}(v)$ and $\ell^{(i,v)} \in \text{Hist}$ such that: $h^{(i,v)}\ell^{(i,v)} \in \text{Hist}(v)$, $|h^{(i,v)}\ell^{(i,v)}| \leq L$ and $\rho^{(i,v)} = h^{(i,v)}(\ell^{(i,v)})^\omega$.

Having (j, u) in memory (the last deviating player j and the vertex u where he moved), the strategy automaton \mathcal{M}_i , $i \in \Pi$, which represents the strategy σ_i , has to produce the lasso $\rho^{(j,u)} = h^{(j,u)}(\ell^{(j,u)})^\omega$ with the length of $h^{(j,u)}\ell^{(j,u)}$ bounded by L . There are at most $|\Pi| \cdot |V| + 1$ such lassoes. It leads to a memory in $\mathcal{O}(|\Pi| \cdot |V| \cdot L)$. \square

Remark 7.3.3. All along this chapter we follow some implicit conventions:

- As in all this part we assume that the considered game $\mathcal{G} = (A, \text{Cost})$ is a multiplayer game such that the objective functions of the players are cost functions;
- $\sup \emptyset = -\infty$ and $\max \emptyset = -\infty$;
- For all $a \in \mathbb{R}$: $a + (+\infty) = +\infty$ and $a + (-\infty) = -\infty$;
- For all $b \in \mathbb{R}_0^+$: $b \cdot (+\infty) = +\infty$ and $b \cdot (-\infty) = -\infty$.
- Additionally, we assume that for all $i \in \Pi$: $\text{Cost}_i : \text{Plays} \rightarrow \mathbb{R} \cup \{+\infty\}$. If Cost_i may take $-\infty$ as value, we should define $\sup \emptyset$ and $\max \emptyset$ in another way. We could, for example, define an additional value \perp and assuming that for all $x \in \mathbb{R} \cup \{-\infty, +\infty\}$, $\perp < x$. Then we could take the convention that $\sup \emptyset = \max \emptyset = \perp$.

7.4 Instantiations of the weak SPE outcome characterizations

The point of this section is to illustrate how the general definition of the sequence of λ^k labeling functions (Definition 7.1.5) and thus the resulting general weak SPE outcome characterizations (Theorem 7.3.1) may be applied to some particular classes of games.

7.4.1 Boolean games with prefix-independent gain functions

Let us assume that $\mathcal{G} = (A, \text{Gain})$ is a multiplayer Boolean game such that for all $i \in \Pi$, Gain_i is a gain function which is prefix-independent in A .

For all $i \in \Pi$ and for all $(v, v') \in E$, we fix:

$$\begin{aligned}\alpha^i(v, v') &= 0 \\ \beta^i(v, v') &= 1\end{aligned}$$

The obtained local strongly prefix-linear constants set $\mathcal{C}_{\mathcal{G}}$ is good. Indeed, let $i \in \Pi$, let $(v, v') \in E$ and let $\rho \in \text{Plays}(v')$, we have that:

$$\begin{aligned}\text{Gain}_i(v\rho) &= \text{Gain}_i(\rho) && (\text{Gain}_i \text{ is prefix-independent}) \\ &= \alpha^i(v, v') + \beta^i(v, v') \cdot \text{Gain}_i(\rho).\end{aligned}$$

Additionally, we know that for all $i \in \Pi$ and for all $\rho \in \text{Plays}$: $\text{Gain}_i(\rho) \geq 0$. In this way, we obtain the following definition of the λ^k labeling functions.

Definition of λ^k for Boolean games with prefix-independent gain functions

Definition 7.4.1. Given $\mathcal{G} = (A, \text{Gain})$ a multiplayer Boolean game such that for all $i \in \Pi$, Gain_i is a prefix-independent gain function in A . The labeling functions $\lambda^k : V \rightarrow \{0, 1\}$ for $k \in \mathbb{N}$ are defined by induction on k .

Let $v \in V$, if $v \in V_i$ for some $i \in \Pi$:

- $\lambda^0(v) = 0$
- $\lambda^{k+1}(v) = \max_{v' \in \text{Succ}(v)} \min\{\text{Gain}_i(\rho) \mid \rho \in \Lambda^k(v')\}.$

In order to use the weak SPE outcome characterizations based on this sequence of labeling functions, it remains to prove that $(\lambda^k)_{k \in \mathbb{N}}$ reaches a fixpoint $k^* \in \mathbb{N}$ and (ii) satisfies the existence of minima property. Since, the image of Gain_i is $\{0, 1\}$, (ii) is obviously satisfied. It remains to prove that $(\lambda^k)_{k \in \mathbb{N}}$ reaches a fixpoint.

Lemma 7.4.2. *There exists $k^* \in \mathbb{N}$ such that for all $m \in \mathbb{N}$ and for all $v \in V$, $\lambda^{k^*}(v) = \lambda^{k^*+m}(v)$ and k^* is at most equal to $|V|$.*

Proof sketch. In the initialization step, all the vertex values are equal to 0. Then at each iteration, (i) if the value of a vertex was equal to 1 in the previous step, then it stays equal to 1 all along the procedure and (ii) if the value of the vertex was equal to 0 then it either stays equal to 0 (for this iteration step) or it becomes equal to 1 (for all the next steps thanks to (i)). At each step, at least one vertex value changes and when no value changes the procedure has reached a fixpoint which corresponds to the values of λ^* . Thus, it means that λ^* is obtained in at most $|V|$ steps. \square

We are now able to obtain a corollary of Theorem 7.3.1 in the particular case of multiplayer Boolean games with prefix-independent gain functions.

Corollary 7.4.3 (of Theorem 7.3.1). *Let $(\mathcal{G}, v_0) = (A, \text{Gain})$ be an initialized multiplayer Boolean game such that for all $i \in \Pi$, Gain_i is prefix-independent in A and let $(\lambda^k)_{k \in \mathbb{N}}$ the sequence of labeling functions provided in Definition 7.4.1.*

Given $\rho^ \in \text{Plays}(v_0)$, the following assertions are equivalent:*

1. *There exists a weak SPE σ in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = \rho^*$;*
2. *ρ^* is λ^* -consistent;*
3. *There exists a good symbolic witness \mathcal{P} such that the play $\rho^{(0, v_0)} = \rho^*$ with $\rho^{(0, v_0)} \in \mathcal{P}$.*

We now apply those weak SPE outcome characterizations on an example.

Example 7.4.4. Let us consider the multiplayer Büchi game depicted by Figure 7.2. In this game, Player 1 (resp. Player 2) owns round (resp. square) vertices and $F_1 = \{v_1\}$ (resp. $F_2 = \{v_3, v_5\}$). The numbers in bold are the values of λ^* . Their iterative computation is given by the following table:

	v_0	v_1	v_2	v_3	v_4	v_5	v_6
λ^0	0	0	0	0	0	0	0
λ^1	0	0	0	1	1	0	0
$\lambda^2 = \lambda^*$	1	0	0	1	1	0	0

Computation of λ^* : First, by definition of λ^0 : for all $v \in V$, $\lambda^0(v) = 0$. Then we have to compute $\lambda^1(v)$ for each $v \in V$. Notice that all plays are λ^0 -consistent. We only focus on the computation of $\lambda^1(v_0)$. In fact, $\lambda^1(v_0)$ is still 0. Indeed, v_0 has v_1 and v_4 as successors and we take the maximum between $\min\{\text{Buchi}_2(\rho) \mid \rho \in \text{Plays}(v_1)\} = 0$ (since $\text{Buchi}_2((v_1v_2)^\omega) = 0$) and $\min\{\text{Buchi}_2(\rho) \mid \rho \in \text{Plays}(v_4)\} = 0$ (since $\text{Buchi}_2(v_4v_6^\omega) = 0$): this is 0. By assuming that we have computed $\lambda^1(v)$ for each $v \in V$, we now compute $\lambda^2(v)$ for each $v \in V$. Let us look at v_0 : we compute $\min\{\text{Buchi}_2(\rho) \mid \rho \in \Lambda^1(v_1)\} = 0$ (since $(v_1v_2)^\omega$ is λ^1 -consistent) and $\min\{\text{Buchi}_2(\rho) \mid \rho \in \Lambda^1(v_4)\} = 1$ (since from v_4 the only λ^1 -consistent play is $v_4v_5^\omega$). Notice that here there is a restriction on the considered plays. To obtain $\lambda^2(v_0)$, it remains to take the maximum between 0 and 1: this is 1. Other values do not change. Finally, by computing λ^3 we can easily see that a fixpoint is reached.




Figure 7.2: Example of a Boolean game with Büchi objectives. In this example, Player 1 (resp. Player 2) owns round (resp. square) vertices and $F_1 = \{v_1\}$ (resp. $F_2 = \{v_3, v_5\}$).

Weak SPEs outcomes: The play $v_0v_4v_6^\omega$ is not the outcome of a weak SPE in (\mathcal{G}, v_0) . Indeed, this play is not λ^* -consistent since $\text{Buchi}_2(v_0v_4v_6^\omega) = 0 < 1 = \lambda^*(v_0)$. On the contrary, the play $v_0v_1v_2v_3^\omega$ is λ^* -consistent and so there exists a weak SPE σ in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = v_0v_1v_2v_3^\omega$. Such a memoryless weak SPE σ is depicted by the double arrows in Figure 7.2.

Notice that ρ is not the outcome of a Nash equilibrium: one can prove thanks to Example 6.1.6 that ρ is not Val^* -consistent.

Good symbolic witness: A finite good symbolic witness \mathcal{P} is depicted

in Table 7.1. Let us illustrate the property (7.5) of Definition 7.2.3 that the plays in \mathcal{P} have to satisfy. For example, let us consider $\rho^{(0,v_0)} = v_0v_1v_2v_3^\omega$.

The successors of v_0 are v_1 and v_4 . Thus we have to check if $\text{Buchi}(v_0v_1v_2v_3) \geq \text{Buchi}(v_0\rho^{(2,v_1)}) = \text{Buchi}(v_0v_1v_2v_3^\omega)$ and $\text{Buchi}(v_0v_1v_2v_3^\omega) \geq \text{Buchi}(v_0\rho^{(2,v_4)}) = \text{Buchi}(v_0v_4v_5^\omega)$. Then we proceed as before from v_1 . The only successor of v_1 is v_2 , thus we only have to check if $\text{Buchi}(v_1v_2v_3) \geq \text{Buchi}(v_1\rho^{(1,v_2)}) = \text{Buchi}(v_1v_2v_3^\omega)$. We proceed in that way for all vertices along $\rho(0, v_0)$ and for all $\rho^{(i,v)} \in \mathcal{P}$.

Notice that the chosen lassoes in \mathcal{P} are subgame outcomes of σ in the subgames of (\mathcal{G}, v_0) .

Table 7.1: An example of finite good symbolic witness

(i, v)	$(0, v_0)$	$(2, v_1)$	$(1, v_2)$	$(1, v_1)$	$(1, v_3)$	$(2, v_3)$
lasso	$v_0v_1v_2v_3^\omega$	$v_1v_2v_3^\omega$	$v_2v_3^\omega$	$v_1v_2v_3^\omega$	v_3^ω	v_3^ω
gain profile	$(0, 1)$	$(0, 1)$	$(0, 1)$	$(0, 1)$	$(0, 1)$	$(0, 1)$
(i, v)	$(2, v_4)$	$(2, v_5)$	$(2, v_6)$	$(1, v_5)$	$(1, v_6)$	
lasso	$v_4v_5^\omega$	v_5^ω	v_6^ω	v_5^ω	v_6^ω	
gain profile	$(0, 1)$	$(0, 1)$	$(0, 0)$	$(0, 1)$	$(0, 0)$	

7.4.2 Safety and qualitative Reachability games

Let $\mathcal{G} = (A, \text{Gain}, (F_i)_{i \in \Pi})$ be either a qualitative Reachability game or a Safety game and let $\mathcal{X} = (X, \text{Gain}^X, (F_i^X)_{i \in \Pi})$ be its extended game. Since Safe and qR are not necessarily strongly prefix-linear in A , we cannot apply the weak SPE outcome characterization on \mathcal{G} . Nevertheless, we have proved that Safe^X and qR^X become prefix-independent in X (Proposition 4.2.7), thus we define the labeling functions λ^k on the extended game.

Let us first fix a local strongly prefix-linear constants set. For all $i \in \Pi$ and for all $((v, I), (v', I')) \in E^X$, we fix:

$$\alpha^i((v, I), (v', I')) = 0$$

$$\beta^i((v, I), (v', I')) = 1$$

The obtained local strongly prefix-linear constants set $\mathcal{C}_{\mathcal{G}}$ is good. Indeed, let $i \in \Pi$, let $((v, I), (v', I')) \in E^X$ and let $\rho \in \text{Plays}_X(v', I')$, we have that:

$$\begin{aligned} \text{Gain}_i^X((v, I)\rho) &= \text{Gain}_i^X(\rho) \quad (\text{Gain}_i^X \text{ is prefix-independent in } X) \\ &= \alpha^i((v, I), (v', I')) + \beta^i((v, I), (v', I')) \cdot \text{Gain}_i^X(\rho) \end{aligned}$$

Additionally, we know that for all $i \in \Pi$ and for all $\rho \in \text{Plays}$: $\text{Gain}_i^X(\rho) \geq 0$. In this way, we obtain the following definition of the λ^k labeling functions.

Definition of λ^k for Safety or qualitative Reachability games

Definition 7.4.5. Given $\mathcal{G} = (A, \text{Gain})$ a multiplayer Safety or qualitative Reachability game. Let $\mathcal{X} = (X, \text{Gain}^X)$ be its associated extended game. The labeling functions $\lambda^k : V^X \rightarrow \{0, 1\}$ for $k \in \mathbb{N}$ are defined by induction on k .

Let $x \in V^X$, if $x \in V_i^X$ for some $i \in \Pi$:

- $\lambda^0(x) = 0$
- $\lambda^{k+1}(x) = \max_{x' \in \text{Succ}(x)} \min\{\text{Gain}_i^X(\rho) \mid \rho \in \Lambda^k(x')\}.$

In order to use the weak SPE outcome characterizations based on this sequence of labeling functions, it remains to prove that $(\lambda^k)_{k \in \mathbb{N}}$ (i) reaches a fixpoint $k^* \in \mathbb{N}$ and (ii) satisfies the existence of minima property. Since, the image of Gain_i^X is $\{0, 1\}$, (ii) is obviously satisfied. It remains to prove that $(\lambda^k)_{k \in \mathbb{N}}$ reaches a fixpoint. The same argument as for Lemma 7.4.2 holds.

Lemma 7.4.6. *There exists $k^* \in \mathbb{N}$ such that for all $m \in \mathbb{N}$ and for all $v \in V^X$, $\lambda^{k^*}(v) = \lambda^{k^*+m}(v)$ and k^* is at most equal to $|V^X| = |V| \cdot 2^{|\Pi|}$. \square*

Remark 7.4.7. Notice that we can also obtain all these results from the fact that Safe^X and qR^X becomes prefix-independent in X and from the instantiation to Boolean games with prefix-independent gain functions (Section 7.4.1).

We are now able to obtain a corollary of Theorem 7.3.1 in the particular

case of multiplayer Safety and qualitative Reachability games.

Corollary 7.4.8 (of Theorem 7.3.1). *Given $(\mathcal{G}, v_0) = (A, \text{Gain})$ an initialized multiplayer Safety or qualitative Reachability game. Let $(\mathcal{X}, x_0) = (X, \text{Gain}^X)$ be its associated extended game. Let $(\lambda^k)_{k \in \mathbb{N}}$ the sequence of labeling functions provided in Definition 7.4.5.*

Given $\rho^ \in \text{Plays}_X(x_0)$, the following assertions are equivalent:*

1. *There exists a weak SPE σ in (\mathcal{X}, x_0) such that $\langle \sigma \rangle_{x_0} = \rho^*$;*
2. *ρ^* is λ^* -consistent;*
3. *There exists a good symbolic witness \mathcal{P} in (\mathcal{X}, x_0) such that the play $\rho^{(0, x_0)} = \rho^*$ with $\rho^{(0, x_0)} \in \mathcal{P}$.*

7.4.3 Quantitative and weighted Reachability games

Since a multiplayer quantitative Reachability game is a particular case of weighted Reachability game (for all $i \in \Pi$, for all $(v, v') \in E$, $w_i(v, v') = 1$), we only explain how to instantiate the definition of the labeling functions λ^k to weighted Reachability games.

Let $\mathcal{G} = (A, \text{WR}, (F_i)_{i \in \Pi})$ be a weighted Reachability game and let $\mathcal{X} = (X, \text{WR}^X, (F_i^X)_{i \in \Pi})$ be its associated extended game. As for multiplayer Safety games and qualitative Reachability games we cannot directly apply the weak SPE outcome characterizations on \mathcal{G} . But by Proposition 4.2.7, we know that for all $i \in \Pi$, WR_i^X is strongly prefix-linear in X . Therefore the labeling functions λ^k are defined on the extended game.

As previously, we begin by fixing a local strongly prefix-linear constants set. Let $i \in \Pi$, let $((v, I), (v', I')) \in E^X$, we fix:

$$\alpha^i((v, I), (v', I')) = \begin{cases} w_i^X((v, I), (v', I')) & \text{if } i \notin I \\ 0 & \text{otherwise} \end{cases}$$

$$\beta^i((v, I), (v', I')) = 1$$

The obtained local strongly prefix-linear constants set $\mathcal{C}_{\mathcal{G}}$ is good. Indeed,

let $i \in \Pi$, let $((v, I), (v', I')) \in E^X$ and let $\rho \in \text{Plays}_X(v', I')$, we have that:

- If $i \in I$, then $i \in I'$ by I-monotonicity. It follows that $\text{WR}_i^X((v, I)\rho) = \text{WR}_i^X(\rho) = 0$.
Moreover, in this case $\alpha^i((v, I), (v', I')) = 0$ and $\beta^i((v, I), (v', I')) = 1$, thus we obtain $\text{WR}_i^X((v, I)\rho) = \alpha^i((v, I), (v', I')) + \beta^i((v, I), (v', I')) \cdot \text{WR}_i^X(\rho)$.
- If $i \notin I$, then $\alpha^i((v, I), (v', I')) = w_i^X((v, I), (v', I'))$ and $\beta^i((v, I), (v', I')) = 1$.
It follows that $\text{WR}_i^X((v, I)\rho) = \alpha^i((v, I), (v', I')) + \alpha^i((v, I), (v', I')) \cdot \text{WR}_i^X(\rho)$.

Additionally, we know that for all $i \in \Pi$, for all $(v, I) \in V^X$ such that $i \in I$, for all $\rho \in \text{Plays}(v, I)$: $\text{WR}_i(\rho) = 0$. In this way, we obtain the following definition of the λ^k labeling functions.

Definition of λ^k for quantitative and weighted Reachability games

Definition 7.4.9. Given $\mathcal{G} = (A, \text{WR}, (F_i)_{i \in \Pi})$ a weighted reachability game. Let $\mathcal{X} = (X, \text{WR}^X, (F_i^X)_{i \in \Pi})$ be its associated extended game. The labeling functions $\lambda^k : V^X \rightarrow \mathbb{N} \cup \{+\infty\}$ for $k \in \mathbb{N}$ are defined by induction on k .

Let $x \in V^X$, if $x \in V_i^X$ for some $i \in \Pi$:

$$\bullet \lambda^0(x) = \begin{cases} 0 & \text{if } i \in I(x) \\ +\infty & \text{otherwise} \end{cases}.$$

$$\bullet \lambda^{k+1}(x) = \begin{cases} 0 & \text{if } i \in I(x) \\ \min_{x' \in \text{Succ}(x)} \left\{ w_i^X(x, x') + \sup\{\text{WR}_i^X(\rho) \mid \rho \in \Lambda^k(x')\} \right\} & \text{otherwise} \end{cases}.$$

We now prove that $(\lambda^k)_{k \in \mathbb{N}}$ reaches a fixpoint $k^* \in \mathbb{N}$.

Proposition 7.4.10. *There exists $k^* \in \mathbb{N}$ such that for all $m \in \mathbb{N}$ and for all $v \in V^X$:*

$$\lambda^{k^*}(v) = \lambda^{k^*+m}(v).$$

Proof. For all $v \in V^X$, the sequences $(\lambda^k(v))_{k \in \mathbb{N}}$ are non-increasing (Lemma 7.1.9) and the component-wise ordering is a well quasi-ordering on $(\mathbb{N} \cup \{+\infty\})^{|V^X|}$ thus there exists $k \in \mathbb{N}$ such that for all $v \in V^X$ and for all $m \in \mathbb{N}$: $\lambda^k(v) = \lambda^{k+m}(v)$. \square

We now need to establish an important property satisfied by the sets $\Lambda^k(v)$: when for some player i , the costs $\text{WR}_i^X(\rho)$ associated with the plays ρ in $\Lambda^k(v)$ are unbounded, there actually exists some play in this set that has an infinite cost. In other terms, either $\Lambda^k(v)$ contains at least one play ρ with an infinite cost $\text{WR}_i^X(\rho)$ or there exists a constant $c \in \mathbb{N}$ such that $\text{WR}_i^X(\rho) \leq c$ for all $\rho \in \Lambda^k(v)$.

Proposition 7.4.11. *For every $k \in \mathbb{N}$, for every $v \in V^X$ and for every $i \in \Pi$, the following implication holds: if $\sup\{\text{WR}_i^X(\rho) \mid \rho \in \Lambda^k(v)\} = +\infty$, then there exists a play $\rho \in \Lambda^k(v)$ such that $\text{WR}_i^X(\rho) = +\infty$.*

Proof. Let $k \in \mathbb{N}$, let $v \in V^X$ and let $i \in \Pi$. Let us assume that $\sup\{\text{WR}_i^X(\rho) \mid \rho \in \Lambda^k(v)\} = +\infty$. It follows that for all $n \in \mathbb{N}$, there exists $\rho^n \in \Lambda^k(v)$ such that $\text{WR}_i^X(\rho^n) > n$.

In particular, $(\rho^n)_{n \in \mathbb{N}}$ is a sequence of plays in $\text{Plays}(v)$. Since V^ω is a compact set and $\text{Plays}(v)$ is a closed set in V^ω , $\text{Plays}(v)$ is compact. It follows that there exists a subsequence $(\rho^{n_\ell})_{\ell \in \mathbb{N}}$ of $(\rho^n)_{n \in \mathbb{N}}$ and $\rho \in \text{Plays}(v)$ such that:

$$\rho = \lim_{\ell \rightarrow +\infty} \rho^{n_\ell}.$$

As WR_i^X is a continuous function, we obtain that $\text{WR}_i^X(\rho) = +\infty$. In order to conclude the proof, we show that $\rho \in \Lambda^k(v)$.

Let us assume that $\rho \notin \Lambda^k(v)$: there exist $j \in \Pi$, $t \in \mathbb{N}$ such that $\rho_t \in V_j$ and $\text{WR}_j^X(\rho_{\geq t}) > \lambda^k(\rho_t)$. In particular, it implies that $\lambda^k(\rho_t) < +\infty$.

Let $h = \rho_0 \dots \rho_t \dots \rho_{t+\lambda^k(\rho_t)}$ be a prefix of ρ .

Since $\rho = \lim_{\ell \rightarrow +\infty} \rho^{n_\ell}$, we can choose n_ℓ large enough that ρ^{n_ℓ} and ρ share h as a common prefix. We have to consider two cases:

- If there exists $t' \in \mathbb{N}$ such that $t \leq t' \leq t + \lambda^k(\rho_t)$ and $j \in I(\rho_{t'})$ (Player j visits F_j^X along $\rho_t \dots \rho_{t+\lambda^k(\rho_t)}$), then $\text{WR}_j^X(\rho_{\geq t}) = \text{WR}_j^X(\rho_{\geq t}^{n_\ell})$ and $\text{WR}_j^X(\rho_{\geq t}^{n_\ell}) > \lambda^k(\rho_t) = \lambda^k(\rho_t^{n_\ell})$. It is a contradiction with $\rho^{n_\ell} \in \Lambda^k(v)$.
- Otherwise: $\text{WR}_j^X(\rho_{\geq t}^{n_\ell}) > |\rho_t \dots \rho_{t+\lambda^k(\rho_t)}| = \lambda^k(\rho_t) = \lambda^k(\rho_t^{n_\ell})$ (because $w_j^X(e) \geq 1$ for all $e \in E^X$). It follows that $\rho^{n_\ell} \notin \Lambda^k(v)$, that is contradiction. \square

The next corollary is a direct consequence of Proposition 7.4.11.

Corollary 7.4.12. *For every $k \in \mathbb{N}$, for every $v \in V^X$ and for every $i \in \Pi$:*

$$\sup\{\text{WR}_i^X(\rho) \mid \rho \in \Lambda^k(v)\} = \max\{\text{WR}_i^X(\rho) \mid \rho \in \Lambda^k(v)\}.$$

We are now able to obtain a corollary of Theorem 7.3.1 in the particular case of multiplayer quantitative and weighted Reachability games.

Corollary 7.4.13 (of Theorem 7.3.1). *Given $(\mathcal{G}, v_0) = (A, \text{WR}, (F_i)_{i \in \Pi})$ an initialized multiplayer quantitative or weighted Reachability game. Let $(\mathcal{X}, x_0) = (X, \text{WR}^X, (F_i^X)_{i \in \Pi})$ be its associated extended game. Let $(\lambda^k)_{k \in \mathbb{N}}$ the sequence of labeling functions provided in Definition 7.4.9.*

Given $\rho^ \in \text{Plays}_X(x_0)$, the following assertions are equivalent:*

1. *There exists a weak SPE σ in (\mathcal{X}, x_0) such that $\langle \sigma \rangle_{x_0} = \rho^*$;*
2. *ρ^* is λ^* -consistent;*
3. *There exists a good symbolic witness \mathcal{P} in (\mathcal{X}, x_0) such that the play*

$$\rho^{(0,x_0)} = \rho^* \text{ with } \rho^{(0,x_0)} \in \mathcal{P}.$$

We conclude this section by providing an example of the computation of $(\lambda^k)_{k \in \mathbb{N}}$ and of utilization of Corollary 7.4.13.

Example 7.4.14. We come back to Example 4.1.3. Let us recall it is a multi-player quantitative reachability game $\mathcal{G} = (A, (QR_1, QR_2), (F_1, F_2))$ with two players. The game arena is depicted in Figure 4.1, the round vertices are owned by Player 1 whereas the square vertices are owned by Player 2. The target sets of the players are respectively equal to $F_1 = \{v_2\}$ (grey vertex), $F_2 = \{v_2, v_5\}$ (double circled vertices). Its associated extended game is provided in Figure 4.3. In Figure 7.3, we provide this extended game enhanced with the value of λ^* , in bold, near the corresponding vertex.

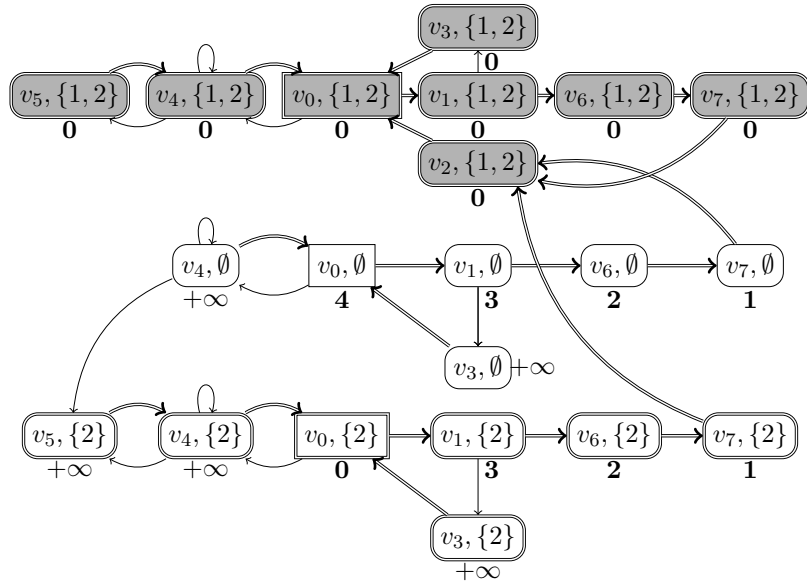


Figure 7.3

Computation of λ^* : The different computation steps of λ^* are summarized in Table 7.2. The columns indicate the vertices according to their related set I of players that have already visited their target set, respectively Π , $\{2\}$, and \emptyset . Notice that for $I = \Pi$, we only write one column v as for all vertices

(v, Π) the value of λ is equal to 0 all along the procedure.

For all $v \in V^X$, $\lambda^0(v) = 0$ by definition.

Let us first consider the vertices $v \in V^X$ with $I(v) = \Pi$. Since all $v' \in \text{Succ}(v)$ are such that $I(v') = I$, we obtain that $\lambda^k(v) = 0$ for all $k \in \mathbb{N}$ and in particular $\lambda^*(v) = 0$.

Let us now explain how to compute λ^1 from λ^0 for the other vertices. For $v = (v_7, \{2\})$, we have that $\lambda^1(v) = 1 + \min_{(v, v') \in E^X} \sup\{\text{QR}_1^X(\rho) \mid \rho \in \Lambda^0(v')\}$. As the unique successor of v is $(v_2, \{1, 2\})$ and all λ^0 -consistent plays beginning in this successor have cost 0 for Player 1, we have that $\lambda^1(v) = 1$. For the computation of $\lambda^1(v_6, \{2\})$, the same argument holds since $(v_6, \{2\})$ has the unique successor $(v_7, \{2\})$. The vertex $(v_1, \{2\})$ has two successors: $(v_6, \{2\})$ and $(v_3, \{2\})$. Again, we know that all λ^0 -consistent plays beginning in $(v_6, \{2\})$ have cost 2 for Player 1. From $(v_3, \{2\})$ however, the play $(v_3, \{2\})(v_0, \{2\})(v_4, \{2\})^\omega$ is λ^0 -consistent and has cost $+\infty$ for Player 1. Thus, we obtain that $\lambda^1(v_1, \{2\}) = 3$. For the other vertices $v \in V^X$ with $I(v) = \{2\}$, one can see that $\lambda^1(v) = \lambda^0(v)$. The same arguments hold for vertices $v \in V^X$ such that $I(v) = \emptyset$.

Now we compute λ^2 from λ^1 . For all vertices $(v, I) \in V^X$ such that $I = \{2\}$ those values does not evolve anymore. But for $(v, I) = (v_0, \emptyset)$, $\lambda^2(v_0, \emptyset)$ becomes 4. Indeed, this vertex has two successors: (v_4, \emptyset) with the play $((v_4, \emptyset))^\omega$ which is λ^1 -consistent and (v_1, \emptyset) such that the only λ^1 -consistent plays from this vertex are those with prefix $(v_1, \emptyset)(v_6, \emptyset)(v_7, \emptyset)(v_2, \{2\})$. It follows that $1 + \min_{v' \in \text{Succ}(v_0, \emptyset)} \sup\{\text{QR}_i^X(\rho) \mid \rho \in \Lambda^1(v')\} = 4$.

Finally, one can see that $\lambda^3 = \lambda^2$ and thus $(\lambda^k)_{k \in \mathbb{N}}$ has reached a fixpoint with $k^* = 2$.

Weak SPEs outcomes: The play $((v_0, \emptyset)(v_4, \emptyset))^\omega$ is not λ^* -consistent, thus $((v_0, \emptyset)(v_4, \emptyset))^\omega$ is not the outcome of a weak SPE in (\mathcal{X}, x_0) . It follows that $(v_0 v_4)^\omega$ is not the outcome of a weak SPE in (\mathcal{G}, v_0) . Notice that we have proved in Example 6.2.5 that it is the outcome of Nash equilibrium.

On the contrary, the play

$$\rho = (v_0, \emptyset)(v_1, \emptyset)(v_6, \emptyset)(v_7, \emptyset)((v_2, \Pi)(v_0, \Pi)(v_1, \Pi)(v_6, \Pi)(v_7, \Pi))^\omega$$

is λ^* -consistent. In particular ρ is the outcome of a weak SPE in (\mathcal{X}, x_0) . For example, the memoryless strategy profile (σ_1, σ_2) depicted by double arrows in Figure 7.3 is a weak SPE in (\mathcal{X}, x_0) with outcome ρ . It follows that the corresponding strategy profile σ' in (\mathcal{G}, v_0) is a weak SPE in (\mathcal{G}, v_0) with outcome $(v_0 v_1 v_6 v_7 v_2)^\omega$.

I	{1, 2}	{2}							\emptyset					
	v	v_0	v_1	v_6	v_7	v_3	v_4	v_5	v_0	v_1	v_6	v_7	v_3	v_4
λ^0	0	0	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$
λ^1	0	0	3	2	1	$+\infty$	$+\infty$	$+\infty$	$+\infty$	3	2	1	$+\infty$	$+\infty$
$\lambda^2 = \lambda^*$	0	0	3	2	1	$+\infty$	$+\infty$	$+\infty$	4	3	2	1	$+\infty$	$+\infty$

Table 7.2: The different steps of the computation of λ^* for the extended game of Figure 7.3

7.4.4 λ -consistency vs Visit λ -consistency in Reachability games

In this section, we prove that given a multiplayer Reachability game $\mathcal{G} = (\mathbf{A}, \text{Reach}, (F_i)_{i \in \Pi})$, its associated extended game and a labeling function λ^k as defined in Definition 7.4.5 or Definition 7.4.9: for all play in the extended game, the play is λ^k -consistent if and only if it is Visit λ^k -consistent (Definition 6.2.2).

Proposition 7.4.15. *Let $\mathcal{G} = (\mathbf{A}, \text{Reach}, (F_i)_{i \in \Pi})$ be a multiplayer Reachability game and let $\mathcal{X} = (X, \text{Reach}^X, (F_i^X)_{i \in \Pi})$ be its associated extended game. Given a play $\rho \in \text{Plays}_X$ and a labeling function λ^k as defined in Definition 7.4.5 or Definition 7.4.9:*

ρ is λ^k -consistent if and only if ρ is Visit λ^k -consistent.

Proof. Let $\rho \in \text{Plays}_X$ and a labeling function λ^k as defined in Definition 7.4.5 if $\text{Reach}^X = \text{qR}^X$ or as in Definition 7.4.9 if $\text{Reach}^X = \text{QR}^X$ or WR^X .

- If $\text{Reach}^X = \text{qR}^X$:

(\Leftarrow) Let us assume that for all $i \in \Pi$ and for all $n \in \mathbb{N}$, we have

that $(i \notin \text{Visit}(\rho_0 \dots \rho_n) \wedge \rho_n \in V_i^X \Rightarrow \text{qR}_i^X(\rho_{\geq n}) \geq \lambda^k(\rho_n))$.
 Let $i \in \Pi$ and $k \in \mathbb{N}$ such that $\rho_k \in V_i^X$, we have to prove that $\text{qR}_i^X(\rho_{\geq k}) \geq \lambda^k(\rho_k)$. If $i \notin \text{Visit}(\rho_0 \dots \rho_k)$, it is true by $\text{Visit } \lambda^k$ -consistency. Otherwise, it means that $i \in I(\rho_k)$ and thus $\text{qR}_i(\rho_{\geq k}) = 1$ and $\lambda^k(\rho_k) = 1$.

(\Rightarrow) Follows directly from the definitions of λ^k -consistency and $\text{Visit } \lambda^k$ -consistency.

- If $\text{Reach}^X = \text{QR}^X$ or WR^X :

(\Leftarrow) Let us assume that for all $i \in \Pi$ and for all $n \in \mathbb{N}$, we have that $(i \notin \text{Visit}(\rho_0 \dots \rho_n) \wedge \rho_n \in V_i^X \Rightarrow \text{WR}_i^X(\rho_{\geq n}) \leq \lambda^k(\rho_n))$.
 Let $i \in \Pi$ and $k \in \mathbb{N}$ such that $\rho_k \in V_i^X$, we have to prove that $\text{WR}_i^X(\rho_{\geq k}) \leq \lambda^k(\rho_k)$. If $i \notin \text{Visit}(\rho_0 \dots \rho_k)$, it is true by $\text{Visit } \lambda^k$ -consistency. Otherwise, it means that $i \in I(\rho_k)$ and thus $\text{WR}_i(\rho_{\geq k}) = 0$ and $\lambda^k(\rho_k) = 0$.

(\Rightarrow) Follows directly from the definitions of λ^k -consistency and $\text{Visit } \lambda^k$ -consistency. \square

7.5 SPE outcome characterization

From the characterizations provided in the previous sections (Theorem 7.3.1), we are able to obtain the same characterizations for SPEs by adding an hypothesis on the cost functions of the game: the cost functions have to be continuous. Indeed, with this additional hypothesis the notions of weak SPE and SPE are equivalent (Proposition 2.4.22).

Theorem 7.5.1 (SPE outcome characterizations). *Let $(\mathcal{G}, v_0) = (A, \text{Cost})$ be an initialized multiplayer game and let $(\lambda^k)_{k \in \mathbb{N}}$ be a sequence of labeling functions as defined in Definition 7.1.5 such that:*

(H_1) *For all $i \in \Pi$, Cost_i is strongly prefix-linear in A and **continuous**;*

(H_2) *$(\lambda^k)_{k \in \mathbb{N}}$ satisfies the existence of maxima property;*

(H₃) $(\lambda^k)_{k \in \mathbb{N}}$ reaches a fixpoint and its fixpoint is $k^* \in \mathbb{N}$.

Then, given $\rho^* \in \text{Plays}(v_0)$, the following assertions are equivalent:

1. There exists **an SPE** σ in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = \rho^*$;
2. $\rho^* \in \Lambda^*(v_0)$ and for all $v \in \text{Succ}^*(v_0)$, $\Lambda^*(v) \neq \emptyset$;
3. There exists a good symbolic witness \mathcal{P} such that the play $\rho^{(0, v_0)} = \rho^*$ with $\rho^{(0, v_0)} \in \mathcal{P}$.

In particular this result may be applied to quantitative and weighted Reachability games.

Corollary 7.5.2 (of Theorem 7.5.1). *Given $(\mathcal{G}, v_0) = (A, \text{WR}, (F_i)_{i \in \Pi})$ an initialized multiplayer quantitative or weighted Reachability game. Let $(\mathcal{X}, x_0) = (X, \text{WR}^X, (F_i^X)_{i \in \Pi})$ be its associated extended game. Let $(\lambda^k)_{k \in \mathbb{N}}$ the sequence of labeling functions provided in Definition 7.4.9.*

Given $\rho^ \in \text{Plays}_X(x_0)$, the following assertions are equivalent:*

1. *There exists an SPE σ in (\mathcal{X}, x_0) such that $\langle \sigma \rangle_{x_0} = \rho^*$;*
2. *ρ^* is λ^* -consistent;*
3. *There exists a good symbolic witness \mathcal{P} in (\mathcal{X}, x_0) such that the play $\rho^{(0, x_0)} = \rho^*$ with $\rho^{(0, x_0)} \in \mathcal{P}$.*

Since, we can also prove that in multiplayer qualitative Reachability games, the concepts of weak SPE and SPE are equivalent, in the same way we also obtained SPE outcome characterizations for qualitative Reachability games.

Corollary 7.5.3 (of Theorem 7.3.1). *Given $(\mathcal{G}, v_0) = (A, \text{qR}, (F_i)_{i \in \Pi})$ an initialized multiplayer qualitative Reachability game. Let $(\mathcal{X}, x_0) = (X, \text{qR}^X, (F_i^X)_{i \in \Pi})$ be its associated extended game. Let $(\lambda^k)_{k \in \mathbb{N}}$ the sequence of labeling functions provided in Definition 7.4.5.*

Given $\rho^* \in \text{Plays}_{\mathcal{X}}(x_0)$, the following assertions are equivalent:

1. There exists an SPE σ in (\mathcal{X}, x_0) such that $\langle \sigma \rangle_{x_0} = \rho^*$;
2. ρ^* is λ^* -consistent;
3. There exists a good symbolic witness \mathcal{P} in (\mathcal{X}, x_0) such that the play $\rho^{(0, x_0)} = \rho^*$ with $\rho^{(0, x_0)} \in \mathcal{P}$.

Proposition 7.5.4. Let $(\mathcal{G}, v_0) = (A, \text{qR}, (F_i)_{i \in \Pi})$ be an initialized multiplayer qualitative Reachability game and let σ be a strategy profile in (\mathcal{G}, v_0) . Then σ is an SPE in (\mathcal{G}, v_0) if and only if σ is a weak SPE in (\mathcal{G}, v_0) .

Proof. Let $(\mathcal{G}, v_0) = (A, \text{qR}, (F_i)_{i \in \Pi})$ be an initialized multiplayer qualitative Reachability game and let σ be a strategy profile in (\mathcal{G}, v_0) .

(\Rightarrow) This implication is a consequence of the definitions of SPE and weak SPE.

(\Leftarrow) Let σ be a weak SPE in (\mathcal{G}, v_0) . Assume that σ is not an SPE, *i.e.*, there exists $hv \in \text{Hist}(v_0)$ such that $\sigma_{\upharpoonright h}$ is not an NE in $(\mathcal{G}_{\upharpoonright h}, v)$. Then some player i has a profitable deviation σ'_i in the subgame $(\mathcal{G}_{\upharpoonright h}, v)$. As qR_i takes its values in $\{0, 1\}$, this means that

$$0 = \text{qR}_i(h\rho) < \text{qR}_i(h\rho') = 1$$

with $\rho = \langle \sigma_{\upharpoonright h} \rangle_v$ and $\rho' = \langle \sigma'_i, \sigma_{-i \upharpoonright h} \rangle_v$. We consider the first occurrence of a vertex of F_i along $h\rho'$ (which appears in ρ' and not in h as $\text{qR}_i(h\rho) = 0$): let g' of minimal length such that $hg' < h\rho'$ and g' ends in some $v' \in F_i$. Let us define a strategy τ_i that is finitely deviating from $\sigma_{i \upharpoonright h}$ and profitable for Player i in $(\mathcal{G}_{\upharpoonright h}, v)$. This will be in contradiction with our hypothesis. For all $g \in \text{Hist}_i(v)$, let

$$\tau_i(g) = \begin{cases} \sigma'_i(g) & \text{if } g \leq g' \\ \sigma_{i \upharpoonright h}(g) & \text{otherwise} \end{cases} .$$

By definition of τ_i , we have that $\text{qR}_i(h\langle\tau_i, \sigma_{-i|h}\rangle_v) = \text{qR}_i(h\rho') = 1$ and τ_i is finitely deviating from $\sigma_{i|h}$ since $|g'|$ is finite. \square

Part III:

RELEVANT EQUILIBRIA

Motivations In order to introduce this part and our motivations, we come back to an example that we have already discussed (Example 6.2.5). Let us first recall this example.

It is a multiplayer quantitative Reachability game $\mathcal{G} = (A, (QR_1, QR_2), (F_1, F_2))$ with two players. The game arena is depicted again in Figure 8.1, the round vertices are owned by Player 1 whereas the square vertices are owned by Player 2. The target sets of the players are respectively equal to $F_1 = \{v_2\}$ (grey vertex), $F_2 = \{v_2, v_5\}$ (double circled vertices).




Figure 8.1: The multiplayer quantitative Reachability game of Example 4.1.3 and Example 6.2.5

In Example 6.2.5, we have proved that the plays $\rho = (v_0v_4)^\omega$ and $\rho' = v_0v_1v_6v_7v_2(v_0v_4)^\omega$ are both outcomes of Nash equilibria. That means that in the same initialized game different Nash equilibria coexist. Moreover, we recall that $QR(\rho) = (+\infty, +\infty)$ and no player reaches his target set while

$\text{QR}(\rho') = (4, 4)$ and both players reach their target set.

In view of that, synthesizing a Nash equilibrium with outcome ρ' instead of ρ seems more *relevant*. Obviously, this phenomenon of coexistence of equilibria in the same game does not only occur with Nash equilibria in Reachability game. A natural question which can arise is thus “What is a relevant equilibrium in a given game?”, for some notion of equilibrium (NE, SPE or weak SPE) and in some kind of games (multiplayer Boolean games, multiplayer Reachability games, . . .). We are also interested in answering this other question “Given a notion of relevant equilibrium, what is the complexity of deciding if such a relevant equilibrium exists in a given game?”.

Constrained existence problem One way to obtain a relevant equilibrium is to minimize (resp. maximize) the cost (resp. the gain) of each player. The *constrained existence problem* (CEP) is the decision problem associated with this optimization problem.

Even if we could give a general definition of this problem, we choose to provide a definition for multiplayer Boolean games and another for multiplayer games with cost functions¹.

Problem 1 ((Boolean) Constrained existence problem). Let $(\mathcal{G}, v_0) = (A, \text{Gain})$ be a multiplayer Boolean game, let $x, y \in \{0, 1\}^{|\Pi|}$ be two thresholds, decide whether there exists an equilibrium^a σ in (\mathcal{G}, v_0) such that $x \leq \text{Gain}(\langle \sigma \rangle_{v_0}) \leq y$.

^aIn this document, an equilibrium is either an NE, an SPE or a weak SPE.

Notice that with Problem 1, if for some player $i \in \Pi$, $x_i = 1$ then, it implies that Player i has to win along the outcome of the equilibrium. On the contrary, if for some player $i \in \Pi$, $y_i = 0$ then, it implies that Player i has to loose along the outcome of the equilibrium.

¹Once again, the counterpart for multiplayer games with gain functions may be easily obtained from the one for games with cost functions.

Problem 2 ((Quantitative) Constrained existence problem). Let $(\mathcal{G}, v_0) = (A, \text{Cost})$ be a multiplayer game, let $x, y \in (\mathbb{R} \cup \{-\infty, +\infty\})^{|\Pi|}$ be two thresholds, decide whether there exists an equilibrium σ in (\mathcal{G}, v_0) such that $x \leq \text{Cost}_i(\langle \sigma \rangle_{v_0}) \leq y$.

Organization of the part In Chapter 9, we consider multiplayer Boolean games with prefix-independent gain functions and solve the constrained existence problem of weak SPEs in co-Büchi, Parity, Muller, Rabin, Streett, Büchi and Explicit Muller games. In Chapter 10, we study the constrained existence problem of weak SPEs and SPEs in qualitative Reachability games and Safety games. In Chapter 11, we prove that the constrained existence problem of weak SPEs (and SPEs) in quantitative Reachability games is PSPACE-complete. Finally, in Chapter 12, we focus on quantitative Reachability games and consider variants of the constrained existence problem in this particular setting.

CHAPTER 9

BOOLEAN GAMES WITH PREFIX-INDEPENDENT GAIN FUNCTIONS

In this chapter, based on [BBGR18, Goe20], we study the complexity classes of the constrained existence problem of weak SPEs in multiplayer Boolean games with classical prefix-independent gain functions. Our results are summarized in Table 9.1.

Table 9.1: Complexity classes of the constrained existence problem of weak SPEs for classical prefix-independent qualitative objectives

	Expl. Muller	Büchi	Co-Büchi	Parity	Muller	Rabin	Streett
P-complete	×	×					
NP-complete			×	×	×	×	×

Those results rely on the weak SPE outcome characterizations (Corollary 7.4.3) provided in Section 7.4.1 and based on the labeling functions $(\lambda^k)_{k \in \mathbb{N}}$ defined in Definition 7.4.1. Therefore, in this chapter, when we refer to a labeling function λ^k this is the function defined in Definition 7.4.1.

In Section 9.1, we provide a naive algorithm to decide the constrained existence problem of weak SPEs in multiplayer Boolean games with prefix-independent gain functions. This algorithm is based on a reformulation of the notion of λ -consistency in the particular setting of Boolean games with

prefix-independent gain functions.

In Section 9.2, we prove that the CEP of weak SPEs in multiplayer Boolean games with co-Büchi, Parity, Muller, Rabin, and Streett objectives is NP-complete. The NP-membership is based on the notion of finite good symbolic witness (see Corollary 7.4.3).

In Section 9.3, the P-completeness of the CEP of weak SPEs in multiplayer Büchi games is proved. The P-membership relies on the reformulation of λ -consistency provided in Section 9.1 and on an algorithm already used by Ummels [Umm08] to prove that the CEP of NEs in multiplayer Büchi games belongs to P.

Finally, in Section 9.4, the naive algorithm explained in Section 9.1 allows to obtain the P-membership of the CEP of weak SPEs in multiplayer Explicit Muller games.

9.1 Computation of λ^*

9.1.1 Reformulation of λ^k -Consistency

In the context of multiplayer Boolean games with prefix-independent gain functions, by defining for all $i \in \Pi$ and for all $k \in \mathbb{N}$ the set of vertices $W_i^k = \{v \in V_i \mid \lambda^k(v) = 1\}$, one can rephrase the λ^k -consistency in the following way.

Lemma 9.1.1. *Let $(\mathcal{G}, v_0) = (A, \text{Gain})$ be a multiplayer Boolean game with prefix-independent gain functions in A .*

Let $\rho \in \text{Plays}$ be a play and $k \in \mathbb{N}$, $\rho \models \lambda^k$ if and only if for each $i \in \Pi$, $(\text{Gain}_i(\rho) = 0 \implies \forall n \in \mathbb{N}$ such that $\rho_n \in V_i$, $\rho_n \notin W_i^k$).

The idea is that a play ρ is λ^k -consistent if and only if for each player $i \in \Pi$ either this play has gain 1 for Player i (i.e., Player i achieves his objective along ρ) or ρ does not cross the set W_i^k .

Proof. Let $(\mathcal{G}, v_0) = (A, \text{Gain})$ be a multiplayer Boolean game with prefix-independent gain functions.

Let $\rho \in \text{Plays}$ be a play and $k \in \mathbb{N}$.

(\implies) Let us assume that $\rho \models \lambda^k$.

Let $i \in \Pi$, we assume that $\text{Gain}_i(\rho) = 0$. We have to prove that for all $n \in \mathbb{N}$ such that $\rho_n \in V_i$, $\rho_n \notin W_i^k$.

Let us assume by contradiction that there exists $n \in \mathbb{N}$ such that $\rho_n \in V_i$ and $\rho_n \in W_i^k$.

By definition of W_i^k , it follows that $\lambda^k(\rho_n) = 1$. Additionally, we have:

$$\begin{aligned} \text{Gain}_i(\rho) &= \text{Gain}_i(\rho_{\geq n}) && (\text{Gain}_i \text{ is prefix-independent}) \\ &\geq \lambda^k(\rho_n) && \rho \models \lambda^k \\ &= 1 \end{aligned}$$

That is a contradiction with the assumption that $\text{Gain}_i(\rho) = 0$.

(\impliedby) Let us assume that for each $i \in \Pi$, $\text{Gain}_i(\rho) = 0 \implies \forall n \in \mathbb{N}$ such that $\rho_n \in V_i$, $\rho_n \notin W_i^k$.

We have to prove that ρ is λ^k -consistent.

Let $i \in \Pi$ and let $n \in \mathbb{N}$ such that $\rho_n \in V_i$. If $\text{Gain}_i(\rho) = 1$, then $\text{Gain}_i(\rho_{\geq n}) = 1$ because Gain_i is prefix-independent. Thus $\text{Gain}_i(\rho_{\geq n}) \geq \lambda^k(\rho_n)$ whatever the value of $\lambda^k(\rho_n)$. If $\text{Gain}_i(\rho) = 0$, then it implies by hypothesis that $\rho_n \notin W_i^k$. It follows that $\lambda^k(\rho_n) = 0$ and $\text{Gain}_i(\rho_{\geq n}) = 0 \geq \lambda^k(\rho_n)$ by prefix-independency of Gain_i . That concludes the proof. \square

For all $k \in \mathbb{N}$, we write \overline{W}^k to depict $\{W_i^k \mid i \in \Pi\}$. Notice that one can recover λ^k from \overline{W}^k and vice versa.

9.1.2 Algorithms

In view of the reformulation of λ^k -consistency (Lemma 9.1.1), given (i) a gain profile $p \in \{0, 1\}^{|\Pi|}$, (ii) \overline{W}^k for some $k \in \mathbb{N}$ and (iii) a vertex $v \in V$, we are able to provide an algorithm that decides if there exists a play $\rho \in \text{Plays}(v)$ such that ρ is λ^* -consistent with $\text{Gain}(\rho) = p$.

The main idea is the following: for each player i such that $p_i = 1$, there is nothing to check since Player i wins; on the contrary for each player i such

that $p_i = 0$ the play ρ cannot cross W_i^k otherwise ρ is not λ^k -consistent. Thus, for all players $i \in \Pi$ such that $p_i = 0$, all vertices W_i^k are removed from the arena A . Then we check if there exists a play with gain profile p in this new arena.

With this algorithm we consider a *sub-arena* of the arena A . Given $V' \subseteq V$, the associated sub-arena of A , denoted by $A_{|V'}$ is given by $A_{|V'} = (\Pi, V', E', (V'_i)_{i \in \Pi})$ where (i) $E' = \{(v, v') \in E \mid v, v' \in V'\}$ and (ii) for each $i \in \Pi$, $V'_i = V_i \cap V'$.

Let us assume that the algorithm $\text{EXISTSPLAY}(A, p, v)$ returns true if there exists a play beginning in v with gain profile p in A ; false otherwise (see Lemma 2.2.22). We obtain, the algorithm $\text{EXISTENCEOFLAMBDACONSISTENTPLAY}(\mathcal{G}, \overline{W}^k, p, v)$ (see Algorithm 1) which returns true if there exists a play $\rho \in \text{Plays}(v)$ such that $\rho \models \lambda^k$ and $\text{Gain}(\rho) = p$; false otherwise.

Algorithm 1: $\text{EXISTENCEOFLAMBDACONSISTENTPLAY}(\mathcal{G}, \overline{W}^k, p, v)$

- 1 $V' = V \setminus \bigcup_{i \in \Pi | p_i = 0} W_i^k$;
- 2 $A' = A_{|V'}$;
- 3 **return** $\text{EXISTSPLAY}(A', p, v)$

Lemma 9.1.2 (Correctness and complexity of the algorithm $\text{EXISTENCEOFLAMBDACONSISTENTPLAY}$). *Given a multiplayer Boolean game \mathcal{G} with prefix-independent gain functions, a gain profile $p \in \{0, 1\}^{|\Pi|}$, \overline{W}^k the values of λ^k and a vertex $v \in V$:*

- $\text{EXISTENCEOFLAMBDACONSISTENTPLAY}(\mathcal{G}, \overline{W}^k, p, v)$ returns true if and only if there exists a play $\rho \in \text{Plays}(v)$ such that $\rho \models \lambda^k$ and $\text{Gain}(\rho) = p$.
- This algorithm has a time complexity polynomial in the size of \mathcal{G} and in path where path is the complexity of the algorithm EXISTSPLAY .

Proof. Let \mathcal{G} be a multiplayer Boolean game with prefix-independent gain functions, $p \in \{0, 1\}^{|\Pi|}$, \overline{W}^k and $v \in V$:

We define $A' = A|_{V'}$ with $V' = V \setminus \bigcup_{i \in \Pi | p_i = 0} W_i^k$.

(\Rightarrow) : Let us assume that the algorithm `EXISTENCEOFLAMBDACONSISTENTPLAY`($\mathcal{G}, \overline{W}^k, p, v$) returns true. That means that there exists $\rho \in \text{Plays}_{A'}(v)$ such that $\text{Gain}(\rho) = p$. We now prove that $\rho \models \lambda^k$. In order to do so, we use Lemma 9.1.1. Thus we have to prove that for all $i \in \Pi$, $(\text{Gain}_i(\rho) = 0 \implies \forall n \in \mathbb{N} \text{ st. } \rho_n \in V_i, \rho_n \notin W_i^k)$.

Let $i \in \Pi$. Let us assume that $\text{Gain}_i(\rho) = 0$. Let $n \in \mathbb{N}$ such that $\rho_n \in V_i$, $\rho_n \notin W_i^k$ since W_i^k has been removed from A .

(\Leftarrow) : Let us now assume that there exists $\rho \in \text{Plays}(v)$ such that $\text{Gain}(\rho) = p$ and $\rho \models \lambda^k$. We prove that $\rho \in \text{Plays}_{A'}(v)$.

Otherwise that there exists $n \in \mathbb{N}$ such that $\rho_n \notin V'$. It follows that, by assuming $\rho_n \in V_i$ for some $i \in \Pi$, $p_i = 0$ and $\lambda^k(\rho_n) = 1$. By prefix-independence of Gain_i , we have that $\text{Gain}_i(\rho_{\geq n}) = 0$ and thus $\text{Gain}_i(\rho_{\geq n}) < \lambda^k(\rho_n)$. That is a contradiction with $\rho \models \lambda^k$. \square

Now that we have this algorithm at our disposal, we can easily provide an algorithm that computes \overline{W}^{k+1} from \overline{W}^k for some $k \in \mathbb{N}$.

By assuming that we have \overline{W}^k as an input, we thus have for all $v \in V$ the value $\lambda^k(v)$. In order to compute $\lambda^{k+1}(v)$ we have to compute

$$\max_{v' \in \text{Succ}(v)} \min\{\text{Gain}_i(\rho) \mid \rho \in \Lambda^k(v')\}$$

(if $v \in V_i$ for some $i \in \Pi$). Since we are considering a Boolean game, this maximum is equal to 1 if and only if there exists a vertex $v' \in V$ such that $\min\{\text{Gain}_i(\rho) \mid \rho \in \Lambda^k(v')\} = 1$ and this minimum is equal to 1 if and only if for all $\rho \in \Lambda^k(v')$, $\text{Gain}_i(\rho) = 1$. In other words, this minimum is equal to 0 if and only if there exists a play $\rho \in \Lambda^k(v')$ such that $\text{Gain}_i(\rho) = 0$. In particular, if we check for all gain profil $p \in \{0, 1\}^{|\Pi|}$ such that $p_i = 0$, if there exists a λ^k -consistent play beginning in v' and with gain profile p , we can decide if this minimum is equal to 0.

From these observations follows the NEXTSTEPLABELINGFUNCTIONS algorithm (Algorithm 2). In this algorithm, the set $\mathbf{P}(v)$ depicts the set of all possible gain profiles from v . That is $\mathbf{P}(v) = \{\text{Gain}(\rho) \mid \rho \in \text{Plays}(v)\}$. The complexity of the NEXTSTEPLABELINGFUNCTIONS relies on the cardinality of this set. Its maximal cardinality is denoted by m : $m = \max_{v \in V} |\mathbf{P}(v)|$. It follows that the running time of this algorithm is polynomial in the size of \mathcal{G} , in *path* and in m .

Algorithm 2: NEXTSTEPLABELINGFUNCTIONS($\mathcal{G}, \overline{W}^k$)	
1	foreach $i \in \Pi$ do
2	$W_i^{k+1} = \emptyset$
3	foreach $v \in V$ do
4	Max 1 = False ;
5	foreach $v' \in \text{Succ}(v)$ do
6	Min 0 = False;
7	foreach $p \in \mathbf{P}(v')$ such that $p_i = 0$ for $i \in \Pi$ such that $v \in V_i$ do
8	$res =$
	EXISTENCEOFLAMBDACONSISTENTPLAY($\mathcal{G}, \overline{W}^k, p, v'$);
9	if res is true then
10	Min 0 = True
11	if Min 0 is False then
12	Max 1 = True
13	if Max 1 is True then
14	Add v to W_i^{k+1} for $i \in \Pi$ such that $v \in V_i$
15	return \overline{W}^{k+1}

Lemma 9.1.3. *Given a multiplayer Boolean game \mathcal{G} with prefix-independent gain functions and \overline{W}^k for some $k \in \mathbb{N}$. The set \overline{W}^{k+1} can be computed in time complexity polynomial in the size of \mathcal{G} , *path* and m . \square*

Lemma 7.4.2 states that the fixpoint of $(\lambda^k)_{k \in \mathbb{N}}$ is reached in at most $|V|$ steps. It follows that $\overline{W}^* = \overline{W}^{k^*}$ can also be computed thanks to an algorithm

that runs in time complexity which is polynomial in the size of \mathcal{G} , $path$ and m .

Proposition 9.1.4. *Given a multiplayer Boolean game \mathcal{G} with prefix-independent gain functions, the set \overline{W}^* can be computed in time complexity which is polynomial in the size of \mathcal{G} , $path$ and m .*

Proof. We begin with \overline{W}^0 which is easily built since for all $v \in V$, $\lambda^0(v) = 0$. Thus, for all $i \in \Pi$, $W_i^0 = \emptyset$. Then we apply the NEXTSTEP LABELING-FUNCTIONS algorithm at most $|V|$ times (Lemma 7.4.2) to obtain \overline{W}^* . The resulting time complexity of this procedure is thus also polynomial in the size of \mathcal{G} , $path$ and m . \square

Given two thresholds $x, y \in \{0, 1\}^{|\Pi|}$, a naive approach to decide the constrained existence problem would be to (i) compute \overline{W}^* (Proposition 9.1.4) and (ii) for each gain profile $p \in \{0, 1\}^{|\Pi|}$ such that $x \leq p \leq y$ check if there exists a λ^* -consistent play beginning in v_0 and with gain profile equal to p (Lemma 9.1.2).

Proposition 9.1.5. *The constrained existence problem of weak SPEs in Boolean games with prefix-independent gain functions can be decided thanks to an algorithm which has a time complexity that is polynomial in the size of \mathcal{G} , $path$ and m .* \square

Remark 9.1.6. Notice that to decide the constrained existence problem of weak SPEs in Boolean games with prefix-independent gain functions, since $m \leq 2^{|\Pi|}$, Proposition 9.1.5 implies a time complexity at least exponential, if no smaller upper bound on m is known.

9.2 NP-Completeness for Multiplayer co-Büchi, Parity, Muller, Rabin, and Streett Games

In this section we prove that the constrained existence problem of weak SPEs in multiplayer Boolean games with co-Büchi, Parity, Muller, Rabin, and Streett objectives is NP-complete.

Theorem 9.2.1. *The constrained existence problem of weak SPEs in multiplayer Boolean games with co-Büchi, Parity, Muller, Rabin, and Streett objectives is NP-complete.*

Before proving this result, one needs to prove that there exists a weak SPE in (\mathcal{G}, v_0) with some gain profile p if and only if there exists a *finite* good symbolic witness \mathcal{P} such that $\rho^{(0, v_0)}$ has a gain profile equal to p . This will imply, thanks to Proposition 7.2.6, that there exists a finite-memory weak SPE in (\mathcal{G}, v_0) with the same gain profile.

Proposition 9.2.2. *Let $(\mathcal{G}, v_0) = (A, \text{Gain})$ be a multiplayer Boolean game with Büchi, co-Büchi, Parity, Muller, Explicit Muller, Rabin or Streett objectives. Given a gain profile $p \in \{0, 1\}^{|\Pi|}$, the following assertions are equivalent:*

1. *There exists a weak SPE σ in (\mathcal{G}, v_0) such that $\text{Gain}(\langle \sigma \rangle_{v_0}) = p$;*
2. *There exists a finite good symbolic witness \mathcal{P} such that:*
 - $\text{Gain}(\rho^{(0, v_0)}) = p$;
 - *For all $\rho^{(i, v)} \in \mathcal{P}$, there exist $h \in \text{Hist}(v)$ and $\ell \in \text{Hist}$ such that $h\ell \in \text{Hist}(v)$, $|h\ell| \leq 2 \cdot |V|^2$ and $\rho^{(i, v)} = h\ell^\omega$.*
3. *There exists a finite-memory weak SPE τ in (\mathcal{G}, v_0) with memory size in $\mathcal{O}(|\Pi| \cdot |V|^3)$ and $\text{Gain}(\langle \tau \rangle_{v_0}) = p$.*

Proof. Let $(\mathcal{G}, v_0) = (A, \text{Gain})$ be a multiplayer Boolean game with Büchi, co-Büchi, Parity, Muller, Explicit Muller, Rabin or Streett objectives and a gain profile $p \in \{0, 1\}^{|\Pi|}$.

We successively prove $1 \implies 2$, $2 \implies 3$ and $3 \implies 1$.

(1 \implies 2) : Let us assume that there exists a weak SPE σ in (\mathcal{G}, v_0) such that $\text{Gain}(\langle \sigma \rangle_{v_0}) = p$. Since Gain_i is prefix-independent for all $i \in \Pi$ and the labeling functions defined in Definition 7.4.1 for Boolean games with prefix-independent objectives functions fulfill the hypotheses of Theorem 7.3.1, we know that $\langle \sigma \rangle_{v_0} \in \Lambda^*(v_0)$ and for all $v \in \text{Succ}^*(v_0)$, $\Lambda^*(v) \neq \emptyset$ (Lemma 7.3.2). Therefore we can use the same kind of construction as for the proof of $(2 \implies 3)$ to obtain a finite good symbolic witness.

In this proof for each $\rho^{(i,v)}$ we chose a play $\pi \in \text{Plays}(v)$ such that $\pi \in \Lambda^*(v)$ and $\text{Gain}_i(\pi) = \min\{\text{Gain}_i(\rho') \mid \rho' \in \Lambda^*(v)\}^a$. We proceed in the same way except that we impose that the play π is a lasso such that there exist $h \in \text{Hist}(v)$ and $\ell \in \text{Hist}$ such that $|h\ell| \leq 2 \cdot |V|^2$ and $\pi = h\ell^\omega$.

Let $\pi \in \Lambda^*(v)$ such that $\text{Gain}_i(\pi) = \min\{\text{Gain}_i(\rho') \mid \rho' \in \Lambda^*(v)\}$, we claim that there exists a play π' that is a lasso such that there exist $h \in \text{Hist}(v)$ and $\ell \in \text{Hist}$ such that (i) $|h\ell| \leq 2 \cdot |V|^2$, (ii) $\pi' = h\ell^\omega$ and (iii) $\pi' \in \Lambda^*(v)$. Indeed it is proved in [BBMU15, Proposition 3.1] that given a play ρ , one can construct a lasso ρ' which respects (i) and (ii) and such that $\text{First}(\rho) = \text{First}(\rho')$, $\text{Occ}(\rho) = \text{Occ}(\rho')$, and $\text{Inf}(\rho) = \text{Inf}(\rho')$ (this construction eliminates some cycles of ρ in a clever way). We thus apply this construction on π in order to obtain π' .

Since $\text{Inf}(\pi) = \text{Inf}(\pi')$ we have that $\text{Gain}(\pi) = \text{Gain}(\pi')$. It remains to prove that $\pi' \in \Lambda^*(v)$. Let $j \in \Pi$, let $n \in \mathbb{N}$, and let us assume that $\pi'_n \in V_j$. We have to prove that $\text{Gain}_j(\pi'_{\geq n}) \geq \lambda^*(\pi'_n)$. Since $\text{Occ}(\pi') = \text{Occ}(\pi)$, there exists $m \in \mathbb{N}$ such that $\pi'_n = \pi_m$. We obtain :

$$\begin{aligned}
\text{Gain}_j(\pi'_{\geq n}) &= \text{Gain}_j(\pi') && (\text{Gain}_j \text{ is prefix-independent}) \\
&= \text{Gain}_j(\pi) && (\text{Inf}(\pi) = \text{Inf}(\pi')) \\
&= \text{Gain}_j(\pi_{\geq m}) && (\text{Gain}_j \text{ is prefix-independent}) \\
&\geq \lambda^*(\pi_m) && (\pi \in \Lambda^*(v)) \\
&= \lambda^*(\pi'_n) && (\pi'_n = \pi_m).
\end{aligned}$$

Finally, we proceed exactly as explained before to obtain a lasso $\rho^{(0,v_0)} = h\ell^\omega$ with $|h\ell| \leq 2 \cdot |V|^2$ from $\langle \sigma \rangle_{v_0}$. Moreover, $\rho^{(0,v_0)} \in \Lambda^*(v_0)$ and $\text{Gain}(\rho^{(0,v_0)}) = \text{Gain}(\langle \sigma \rangle_{v_0})$.

By construction the obtained finite symbolic witness is a good symbolic witness.

(2 \implies 3) : It is a direct consequence of Proposition 7.2.6.

(3 \implies 1) : Obviously true. \square

^a Recall that we change the max by a min since we here consider gain functions instead of cost functions.

We are now able to prove Theorem 9.2.1.

Proof of Theorem 9.2.1. We begin with the NP membership. The objectives considered in Theorem 9.2.1 are such that we can apply Proposition 9.2.2. Given thresholds $x, y \in \{0, 1\}^{|\Pi|}$, there exists a weak SPE in (\mathcal{G}, v_0) with gain profile p such that $x \leq p \leq y$ if and only if there exists a finite good symbolic witness \mathcal{P} that contains a lasso $\rho^{(0,v_0)}$ with gain profile p and such that for all $\rho^{(i,v)} \in \mathcal{P}$, there exist $h \in \text{Hist}(v)$ and $\ell \in \text{Hist}$ such that $h\ell \in \text{Hist}(v)$, $|h\ell| \leq 2 \cdot |V|^2$ and $\rho^{(i,v)} = h\ell^\omega$. Hence a nondeterministic polynomial algorithm works as follows: guess a set \mathcal{P} composed of at most $|\Pi| \cdot |V| + 1$ lassoes that may be represented thanks to a prefix of length at most $2 \cdot |V|^2$ and check that \mathcal{P} is a good symbolic witness that contains a lasso $\rho^{(0,v_0)}$ with gain profile p such that $x \leq p \leq y$. Clearly checking that \mathcal{P} is a finite symbolic witness can be done in polynomial time. Checking that it is good requires to compute the gain profiles of its lassoes and to compare them. This can also be done in

polynomial time for co-Büchi, Parity, Muller, Rabin, and Streett objectives. We now proceed to the NP-hardness. It is obtained thanks to a polynomial reduction from SAT. In [Umm08] is provided a polynomial reduction from SAT to the constrained existence problem of NEs in Boolean games with co-Büchi objectives. Due to the structure of the game constructed in this approach, the same reduction holds for the constrained existence problem of weak SPEs. As co-Büchi objectives can be polynomially translated into Parity, Muller, Rabin, and Streett objectives (see [GTW02]), the constrained existence problem for Boolean games with those objectives is also NP-hard. \square

9.3 P-Completeness for Büchi Games

This section is devoted to the proof of the following result.

Theorem 9.3.1. *The constrained existence problem of weak SPEs in multi-player Büchi games is P-complete.*

Recall that for all $i \in \Pi$ and for all $k \in \mathbb{N}$ the set of vertices $W_i^k = \{v \in V_i \mid \lambda^k(v) = 1\}$ and, for all $k \in \mathbb{N}$, $\overline{W}^k = \{W_i^k \mid i \in \Pi\}$ (Section 9.1.1).

In order to prove Theorem 9.3.1, we prove that: (i) for some $k \in \mathbb{N}$, $x, y \in \{0, 1\}^{|\Pi|}$ and $w \in V$, if the set \overline{W}^k is given, deciding if there exists a play $\rho \in \text{Plays}(w)$ such that $x \leq \text{Gain}(\rho) \leq y$ and $\rho \models \lambda^k$ can be done in polynomial time (Proposition 9.3.3) and (ii) the set \overline{W}^* can be computed in polynomial time (Corollary 9.3.5). From (i) and (ii), we obtain that the constrained existence problem of weak SPEs is decidable in P.

In [Umm08], Ummels provides an algorithm to decide a problem similar to (i), we detail this algorithm in Section 9.3.1 and in Section 9.3.2 we explain how we use it to decide the constrained existence problem of weak SPEs in P.

9.3.1 Deciding the constrained existence problem of NEs

In [Umm08], Ummels studies the constrained existence problem but for NEs instead of weak SPEs. Since we use the same kind of approach to prove that the constrained existence problem of weak SPEs in Büchi games belongs to P, we summarize the approach of [Umm08] in this section.

Constrained existence problem of NEs Let (\mathcal{G}, v_0) be a Boolean game, given x and y in $\{0, 1\}^{|\Pi|}$, we want to decide if there exists an NE σ in (\mathcal{G}, v_0) such that $x \leq \text{Gain}(\langle \sigma \rangle_{v_0}) \leq y$. We call this decision problem the constrained existence problem of NEs.

NE outcome characterization Ummels proved that this problem belongs to P [Umm08] for Büchi games. To prove this result, he uses an outcome characterization of NEs. This characterization is based on the notion of *winning vertices*. A winning vertex $v \in V$ for a player $i \in \Pi$ is a vertex owned by Player i such that he has a strategy, called a winning strategy, which ensures the achievement of his objective whatever the strategy of the coalition of the other players. This set of vertices is denoted by W_i and it is well known that this set can be computed in polynomial time [Tho95]. The sets W_i for all $i \in \Pi$ provide a kind of labeling of the vertices which imposes constraints on the plays which can be outcomes of NEs. In the same spirit as Lemma 9.1.1 for weak SPEs, the main idea is that a play ρ is the outcome of an NE if and only if for each player $i \in \Pi$ either this play has gain 1 for Player i or ρ does not cross the set W_i . Indeed, if Player i achieves his objective along ρ , he has no incentive to deviate from his strategy. On the contrary, if Player i does not achieve his objective along ρ and there exists $n \in \mathbb{N}$ such that $\rho_n \in W_i$, since ρ_n is a winning vertex for Player i , Player i has a profitable deviation by following his winning strategy from ρ_n . The characterization can be rephrased in the following way:

Proposition 9.3.2 (NE outcome characterization [Umm08]). *Let (\mathcal{G}, v_0) be a Büchi game and $\rho \in \text{Plays}(v_0)$, there exists an NE σ in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = \rho$ if and only if for each $i \in \Pi$, $\text{Gain}_i(\rho) = 0 \Rightarrow \forall n \in \mathbb{N}$ st. $\rho_n \in$*

$V_i, \rho_n \notin W_i.$

Deciding the constrained existence problem of NEs Let (\mathcal{G}, v_0) be a Büchi game, let $x, y \in \{0, 1\}^{|\Pi|}$, we want to decide the constrained existence problem of NEs in (\mathcal{G}, v_0) . Due to Proposition 9.3.2, we only have to find $\rho \in \text{Plays}(v_0)$ such that $x \leq \text{Gain}(\rho) \leq y$ and ρ is the outcome of an NE in (\mathcal{G}, v_0) .

Since $\text{Inf}(\rho)$ is a strongly connected set of vertices, it amounts to finding a strongly connected set of vertices from which it is possible to build a play ρ which satisfies the constraints and the condition to be the outcome of an NE. It is done by first finding some strongly connected sets in the graph and then refine it. The algorithm developed by Ummels runs in polynomial time and works as follows: (i) it computes W_i for each $i \in \Pi$, (ii) since $y_i = 0$ implies that Player i has to obtain a gain of 0, it removes beforehand the set F_i from V by computing $X = V \setminus (\bigcup_{i \in \Pi: y_i = 0} F_i)$ and (iii) it runs Algorithm 3 on X . The inputs of this procedure are a Boolean game and two thresholds $x, y \in \{0, 1\}^{|\Pi|}$. The output is the set Z of vertices such that for each $v \in Z$, there exists a play $\rho \in \text{Plays}(v)$ such that: (i) $x \leq \text{Gain}(\rho) \leq y$ and (ii) ρ is the outcome of an NE in (\mathcal{G}, v) . Notice that even if Algorithm 3 has the set X in input, in line 8 computing the set of vertices from which the SCC C is reachable is done in $G_{\uparrow V \setminus \bigcup_{i \in L} W_i}$.

9.3.2 P-completeness of the constrained existence problem of weak SPEs in Büchi games

Thanks to the outcome characterization provided in Section 7.4.1 and the ideas of the algorithm explained for NEs in Section 9.3.1, we prove that the constrained existence problem of weak SPEs in multiplayer Büchi games belongs to P.

We recall that in the same way as the sets W_i for all $i \in \Pi$ for NEs, the sets W_i^k for all $i \in \Pi$ impose some constraints on the candidate plays to be outcomes of weak SPEs. Thus, let $k \in \mathbb{N}$, thanks to Lemma 9.1.1, if we know \overline{W}^k , we may use the Ummels algorithm to decide the existence of a play $\rho \in \text{Plays}(v)$

¹A non-trivial SCC is an SCC with at least one edge.

Algorithm 3: SolveSubgame(X)	
1	$Z = \emptyset$;
2	Decompose $\mathcal{G}_{\uparrow X}$ into strongly connected components ;
3	foreach <i>non-trivial SCC</i> C do
4	$L = \{i \in \Pi \mid C \cap F_i = \emptyset\}$;
5	if $i \notin L$ for each i such that $x_i = 1$ then
6	$Y = C \setminus (\bigcup_{i \in L} W_i)$; // we check if $\forall i \in L, C \cap W_i = \emptyset$
7	if $Y = C$ then
8	$Z = Z \cup \{v \in V \mid C \text{ is reachable from } v \text{ in } G_{\uparrow V \setminus (\bigcup_{i \in L} W_i)}\}$
9	else // C is not a good candidate
10	$Z = Z \cup \text{SolveSubgame}(Y)$;
11	return Z

such that $x \leq \text{Gain}(\rho) \leq y$ and $\rho \models \lambda^k$ for some $x, y \in \{0, 1\}^{|\Pi|}$. Indeed, we only have to (i) assume that the set \overline{W}^k is given in input, (ii) compute $X = V \setminus (\bigcup_{i \in \Pi: y_i = 0} F_i)$ (iii) replace the sets W_i by W_i^k and (iv) run Algorithm 3 on X . Since the only differences with Algorithm 3 are the assumption that \overline{W}^k is given in input and the utilization of the sets W_i^k instead of the sets W_i , the correctness and the polynomial running time of this approach hold from those of Algorithm 3. All these reasons provide a proof sketch of Proposition 9.3.3.

Proposition 9.3.3. *Let $k \in \mathbb{N}$, $x, y \in \{0, 1\}^{|\Pi|}$ and $w \in V$, given \overline{W}^k , deciding if there exists a play $\rho \in \text{Plays}(w)$ such that (i) $x \leq \text{Gain}(\rho) \leq y$ and (ii) $\rho \models \lambda^k$ can be done in polynomial time. \square*

If we prove that we can obtain the set \overline{W}^* in polynomial time, we are done. Since we know that k^* is polynomial and that $\lambda^0(v) = 0$ for all $v \in V$, we only have to prove that we can compute \overline{W}^{k+1} from \overline{W}^k in polynomial time.

Proposition 9.3.4. *Given \overline{W}^k for some $k \in \mathbb{N}$, the set \overline{W}^{k+1} can be computed in polynomial time.*

Proof. First, we recover from \overline{W}^k the labeling function λ^k . Then from λ^k we compute the function λ^{k+1} : for all $v \in V$, $\lambda^{k+1}(v) = \max_{v' \in \text{Succ}(v)} \min\{\text{Gain}_i(\rho) \mid \rho \in \Lambda^k(v')\}$ if $v \in V_i$ for some $i \in \Pi$. The computation of $\min\{\text{Gain}_i(\rho) \mid \rho \in \Lambda^k(v')\}$ is done as follows: (i) this min = 0 if and only if there exists $\rho \in \text{Plays}(v')$ such that $\rho \models \lambda^k$ and $\text{Gain}_i(\rho) = 0$, thus (ii) we use Proposition 9.3.3 with x such that for all $j \in \Pi$ $x_j = 0$ and y such that $y_i = 0$ and $y_j = 1$ for all $j \neq i$. If the answer to this decision problem is yes, then min = 0, otherwise min = 1. This is done in polynomial time.

Due to $\max_{v' \in \text{Succ}(v)}$, we have to do it at most $|E| \leq |V|^2$ times. Since we proceed as previously for each $v \in V$, we do this whole procedure at most $|V| \cdot |V|^2$ times. From λ^{k+1} we build the set \overline{W}^{k+1} . It follows that we can compute \overline{W}^{k+1} from \overline{W}^k in polynomial time. \square

Corollary 9.3.5. *The set \overline{W}^* can be computed in polynomial time.*

Proof. We begin by computing W_i^0 for all $i \in \Pi$. Since for all $v \in V$, $\lambda^0(v) = 0$, these sets can be computed in polynomial time. Then, thanks to Proposition 9.3.4, we can compute iteratively each set \overline{W}^{k+1} from the set \overline{W}^k in polynomial time until we obtain \overline{W}^* . Due to Lemma 7.4.2, we know that we have to proceed in this way a polynomial number of times. \square

Proof of Theorem 9.3.1. P-membership Let $x, y \in \{0, 1\}^{|\Pi|}$, we want to decide if there exists a weak SPE σ such that $x \leq \text{Gain}(\langle \sigma \rangle_{v_0}) \leq y$. Thanks to Theorem 7.4.3, it amounts to deciding the existence of a play $\rho \in \text{Plays}(v_0)$ such that $\rho \models \lambda^*$ and $x \leq \text{Gain}(\rho) \leq y$. Therefore we only have to use Proposition 9.3.3 with x, y and the set \overline{W}^* which is computable in polynomial time (Corollary 9.3.5). If the answer to this decision problem is yes, then there exists a weak SPE σ which satisfies the constraints given by the constrained existence problem, otherwise there does not exist such a weak SPE.

P-hardness The P-hardness is due to a logspace reduction from the Circuit Value Problem which is known to be P-complete [Sip13, Theorem 10.44]. The main idea is that a Boolean circuit can be reduced to a Boolean game with Büchi objectives and with two players. In this corresponding game, loops are added to the leaves and Player 1 (resp. Player 2) owns OR-gates (resp. AND-gates). The Büchi objective of Player 1 (resp. Player 2) are positive leaves (resp. negative leaves) of the circuit. The circuit evaluates to true if and only if there exists a weak SPE in the corresponding game with gain profile $(1, 0)$. \square

Remark 9.3.6. Notice that the major difference with the naive algorithm of Section 9.1 is when we want to check the existence of a λ^k -consistent play such that its gain profile lies between two thresholds. In the algorithm for Büchi games, we are able to directly check existence of a λ^k -consistent play such that its gain profile lies between two thresholds in polynomial time. While in the naive approach, we only check the existence of a λ^k -consistent play with a given gain profile and thus we have to check all possible gain profiles between the two thresholds. This phenomenon appears when:

- We have to compute the min to obtain \overline{W}^{k+1} from \overline{W}^k . In the algorithm for Büchi games, we only have to check the existence of a λ^k -consistent play with gain profile p and such that $x \leq p \leq y$ with $x = (0, \dots, 0, x_i = 0, 0, \dots, 0)$ and $y = (1, \dots, 1, y_i = 0, 1, \dots, 1)$ while in the naive algorithm we have to check one by one all $p \in \{0, 1\}^{|\Pi|}$ such that $x \leq p \leq y$ and $p_i = 0$. That is potentially an exponential number of possible gain profiles.
- In the same way, when \overline{W}^* is build and we want to decide the constrained existence problem with thresholds $x, y \in \{0, 1\}^{|\Pi|}$. In the algorithm for Büchi games, we only have to check once the existence of a λ^* -consistent play with gain profile between x and y . While with the naive approach, one have to check the existence of a λ^* -consistent play with gain profile p for all $p \in \{0, 1\}^{|\Pi|}$ and $x \leq p \leq y$.

9.4 P-Completeness for Explicit Muller Games

This section is devoted to the proof of Theorem 9.4.1 which states that the constrained existence problem of weak SPEs in multiplayer Explicit Muller games is P-complete.

The proof relies on the naive algorithm explained in Section 9.1. Indeed, in the case of multiplayer Explicit Muller the number of realizable gain profiles in the game is polynomially bounded and the complexity of the algorithm which checks if there exists a play with a given gain profile runs in polynomial time.

Theorem 9.4.1. *The constrained existence problem of weak SPEs in multiplayer Explicit Muller games is P-complete.*

Proof. Let $(\mathcal{G}, v_0) = (A, \text{Gain})$ be an Explicit Muller game and, for all $i \in \Pi$, let $\mathcal{F}_i \subseteq 2^V$ be the set of subsets of vertices associated with the Explicit Muller objective of Player i .

P-membership: In this case $m = \max_{v \in V} |\mathbf{P}(v)|$ is polynomial. Indeed, for all $v \in V$, if there exists a play ρ with gain profile p from v , then either $\text{Inf}(\rho) \in \cup_{i \in \Pi} \mathcal{F}_i$ or $p = (0, \dots, 0)$. It follows:

$$\mathbf{P}(v) \subseteq \{(0, \dots, 0)\} \cup \{p \in \{0, 1\}^{|\Pi|} \mid \exists F \in \cup_{i \in \Pi} \mathcal{F}_i, p_i = 1 \Leftrightarrow F \in \mathcal{F}_i\}.$$

It follows that $|\mathbf{P}(v)| \leq |\cup_{i \in \Pi} \mathcal{F}_i| + 1$ and m are polynomial.

Additionally, in Explicit Muller game, deciding if there exists a play with a given gain profile from a given vertex can be done in polynomial time (Lemma 2.2.23). Thus *path* is polynomial.

Proposition 9.1.5 allows us to conclude that the constrained existence problem of weak SPEs in Explicit Muller games belongs to P.

P-hardness: The P-hardness is obtained thanks to a reduction from the AND-OR graph reachability problem that is P-complete [Imm81]. Indeed, the P-hardness of the constrained existence problem for SPEs (instead of weak SPEs) in Boolean games with Reachability objectives is proved

in [Umm05, Corollary 6.22] thanks to such a reduction, and it is not difficult to see that the same reduction also holds for weak SPEs and Explicit Muller objectives. □

CHAPTER 10

QUALITATIVE REACHABILITY AND SAFETY GAMES

Theorem 10.0.1. *The constrained existence problem of weak SPEs in Boolean games with qualitative Reachability or Safety objectives is PSPACE-complete.*

Recall that weak SPEs and SPEs are equivalent notions for qualitative Reachability objectives (Proposition 7.5.4). It follows from Theorem 10.0.1 that the constrained existence problem of SPEs (instead of weak SPEs) for Boolean games with qualitative Reachability objectives is PSPACE-complete. We will see later (in Section 10.3, from the proof of Theorem 10.0.1) that the constrained existence problem of SPEs is also PSPACE-complete for Safety objectives.

Corollary 10.0.2. *The constrained existence problem of SPEs in Boolean games with qualitative Reachability and Safety objectives is PSPACE-complete.*

We detail the proof of Theorem 10.0.1 in the next two sections for qualitative Reachability objectives, and we also show how to adapt it for Safety objectives. To get the PSPACE-easiness, we transform the Boolean game (\mathcal{G}, v_0)

with qualitative Reachability objectives (which are not prefix-independent) into its associated extended game (\mathcal{X}, x_0) (Definition 4.2.1). In this way, it is possible to use the concept of good symbolic witness as done before in Section 9.2. Even if the size of the extended game (\mathcal{X}, x_0) is exponential in the size of the initial game (\mathcal{G}, v_0) , we manage to get a PSPACE-membership thanks to the classical complexity result $\text{PSPACE} = \text{APTIME}$. The PSPACE-hardness is obtained with a polynomial reduction from QBF. The reduction is more involved than the one in Theorem 9.2.1. Indeed the reduction for NP-hardness already works for NEs whereas the reduction for PSPACE-hardness really exploits the subgame perfect aspects.

10.1 PSPACE-membership

We here prove that the constrained existence problem of weak SPEs is in PSPACE for qualitative Reachability objectives, and we then explain how to adapt the proof for Safety objectives.

Proposition 10.1.1. *The constrained existence problem of weak SPEs in qualitative Reachability games is in PSPACE.*

Proof. First, we transform the qualitative Reachability game (\mathcal{G}, v_0) into its associated extended game (\mathcal{X}, x_0) .

Let us recall that there is a one-to-one correspondence between plays $\rho = v_0v_1 \dots v_k \dots$ in \mathcal{G} from v_0 and $\rho' = (v_0, I_0)(v_1, I_1) \dots (v_k, I_k) \dots$ in \mathcal{X} from (v_0, I_0) , with the important property that

$$I_k \subseteq I_{k+1} \text{ for all } k \geq 1. \quad (10.1)$$

In particular, there exists a weak SPE with gain profile p in (\mathcal{G}, v_0) if and only if there is one in (\mathcal{X}, x_0) with gain profile p if and only if there is a good symbolic witness \mathcal{P} containing a play $\rho^{(0, x_0)}$ with gain profile p in (\mathcal{X}, x_0) (by Corollary 7.4.8).

Second, let us prove that one may consider that the plays in the symbolic

witness are lassoes of the form $h\ell^\omega$ with $|h\ell|$ polynomially bounded by

$$(|\Pi| + 1) \cdot |V|. \quad (10.2)$$

Indeed recall that a play $\rho^{(i,v)}$ in \mathcal{P} is (i) a λ^* -consistent play $\rho^{(0,v_0)}$ with gain profile p if $(i,v) = (0,v_0)$; or (ii) if $(i,v) \neq (0,v_0)$, $\rho^{(i,v)}$ is a λ^* -consistent play with $\text{Gain}_i(\rho^{(i,v)}) = \min\{\text{Gain}_i(\rho') \mid \rho' \in \Lambda^*(v)\}$ (see the proof $2 \Rightarrow 3$ of Theorem 7.3.1)

Consider such a play $\rho = (v_1, I_1)(v_2, I_2) \dots (v_k, I_k) \dots$ from v_1 . By (10.1), there exists $I \subseteq \Pi$ and $k \in \mathbb{N}$ such that for all $k' \geq k$, $I_{k'} = I$. Hence from ρ , we can construct a lasso ρ' of length bounded by (10.2) such that $\text{First}(\rho') = \text{First}(\rho)$, $\text{Occ}(\rho') \subseteq \text{Occ}(\rho)$, and $\text{Gain}(\rho') = \text{Gain}(\rho)$,

- by eliminating all cycles in the history $(v_1, I_1)(v_2, I_2) \dots (v_{k-1}, I_{k-1})$ (leading to a history of length at most $|\Pi| \cdot |V|$), and
- by detecting in the play $(v_k, I)(v_{k+1}, I) \dots$ the first repeated vertex $(v_{k'}, I) = (v_{k'+\ell+1}, I)$ and replacing this play by the lasso

$$(v_k, I)(v_{k+1}, I) \dots ((v_{k'}, I) \dots (v_{k'+\ell}, I))^\omega$$

of length at most $|V|$.

In this way, if ρ is a λ^* -consistent play with gain profile p from v , then the constructed lasso ρ' is also a λ^* -consistent play with gain profile p from v .

Third we prove PSPACE-membership of the constrained existence problem by proving that it is in APTIME. Given the extended game (\mathcal{X}, x_0) and two thresholds $x, y \in \{0, 1\}^{|\Pi|}$, the alternating Turing machine works as follows. Existential and universal states (respectively controlled by Player \vee and Player \wedge) alternate along an execution of the machine. Player \vee proposes a lasso $\rho^{(j,u)}$ of length bounded by $(|\Pi| + 1) \cdot |V|$ (in the initial state, he proposes a lasso $\rho^{(0,x_0)}$). Then Player \wedge chooses a vertex $w \in V_i^X$, for some $i \in \Pi$, of $\rho^{(j,u)}$ and proposes to move to v such that $(w, v) \in E^X$. Player \vee reacts by proposing a lasso $\rho^{(i,v)}$ of length bounded by $(|\Pi| + 1) \cdot |V|$, and so on. The execution stops after

$$2 \cdot |\Pi|^2 \cdot |V| + 1 \text{ turns.} \quad (10.3)$$

Such an execution is accepting if:

- for the gain profile p of the initial lasso $\rho^{(0,x_0)}$, we have $x \leq p \leq y$;
- for each lasso $\rho^{(j,u)}$ proposed by Player \vee , for the corresponding move $(w, v) \in E^X$ with $w \in V_i^X$ made by Player \wedge , and the answer $\rho^{(i,v)}$ of Player \vee , we have $\text{Gain}_i^X(\rho_{\geq k}^{(j,u)}) \geq \text{Gain}_i^X(\rho_k^{(j,u)} \rho^{(i,v)})$, if $w = \rho_k^{(j,u)}$.

The intuition is that if there exists in (\mathcal{X}, x_0) a finite good symbolic witness \mathcal{P} containing a lasso $\rho^{(0,v_0)}$ with gain profile p such that $x \leq p \leq y$, then Player \vee will play with the lassoes of \mathcal{P} according to Definition 7.2.3. Notice that along an execution of the Turing machine, Player \wedge has no interest to choose twice the same pair (i, v) since Player \vee will react with the same lasso $\rho^{(i,v)}$. Remembering property (10.1), the maximum number of times that Player \wedge has to play is

$$|\Pi|^2 \cdot |V|. \quad (10.4)$$

Indeed for a fixed $I \subseteq \Pi$, Player \wedge can choose at most $|\Pi| \cdot |V|$ different pairs (i, v') with v' of the form (v, I) , and the size of I can only increase. This explains the number of turns of any execution of the machine (see (10.3)): an initial lasso proposed by Player \vee followed by $|\Pi|^2 \cdot |V|$ alternations between moves of both Players \vee and \wedge .

If we stop the proof here, the latter procedure is in fact not in APTIME since it works on the extended game which has an exponential size in the size of (\mathcal{G}, v_0) . It is not a problem since, all this procedure may be done “on-the-fly” without explicitly building the extended game. Let us briefly explain how it is possible.

1. At the beginning of the procedure Player \vee proposes a lasso $\rho^{(0,v_0)}$ in (\mathcal{G}, v_0) beginning in v_0 such that $\rho^{(0,v_0)} = h\ell^\omega$ and $|h\ell| \leq (|\Pi| + 1) \cdot |V|$. Then, one can polynomially obtain the corresponding lasso $\rho^{(0,x_0)^X}$ in (\mathcal{X}, x_0) by enhancing each vertex of $\rho^{(0,v_0)}$ by the set of players that have reached their target set in the prefix. That is: $\rho^{(0,x_0)^X} = (v_0, I_0)(v_1, I_1) \dots (v_k, I_k) \dots$ with $I_0 = \{i \in \Pi \mid v_0 \in F_i\}$ and $I_k = I_{k-1} \cup \{i \in \Pi \mid v_k \in F_i\}$.

2. Then Player \wedge chooses a vertex (v_k, I_k) along $\rho^{(0, x_0)^X}$ and a successor $w \in \text{Succ}(v_k)$ in (\mathcal{G}, v_0) . The corresponding successor in (\mathcal{X}, x_0) is (w, I) with $I = I_k \cup \{i \in \Pi \mid w \in F_i\}$.
3. If $v_k \in V_i$, Player \vee has now to (i) remember that w is enhanced by the set of players I ; (ii) proposes a lasso $\rho^{(i, w)}$ in (\mathcal{G}, v_0) beginning in w such that $\rho^{(i, w)} = h\ell^\omega$ and $|h\ell| \leq (|\Pi| + 1) \cdot |V|$ and (iii) building on-the-fly the corresponding lasso $\rho^{(i, w)^X}$ in (\mathcal{X}, x_0) beginning in (w, I) in the same way as before.

In this way, checking whether an execution is accepting is done in polynomial time since Player \vee proposes lassoes of polynomial size by (10.2), there is a polynomial numbers of turns by (10.3), and computing and comparing gain profiles of those lassoes is done in polynomial time. Indeed, given $\rho^X = (v_1, I_1)(v_2, I_2) \dots$, it is not necessary to explicitly know F_i^X to compute $\text{Gain}_i^X(\rho^X)$: $\text{Gain}_i^X(\rho^X) = 1$ if and only if there exists $k \in \mathbb{N}$ such that $i \in I_k$. So the constrained existence problem of weak SPEs in qualitative Reachability games is in $\text{APTIME} = \text{PSPACE}$. \square

The constrained existence problem of weak SPEs in multiplayer Safety games is solved similarly.

Proposition 10.1.2. *The constrained existence problem of weak SPEs in multiplayer Safety games is in PSPACE.*

10.2 PSPACE-hardness

We now prove that the constrained existence problem of weak SPEs is PSPACE-hard for multiplayer qualitative Reachability games, and we then show how to adapt the proof for multiplayer Safety games.

Proposition 10.2.1. *The constrained existence problem of weak SPEs in qualitative Reachability games is PSPACE-hard.*

To prove this proposition, we give a polynomial reduction from the QBF problem that is PSPACE-complete. This problem is to decide whether a fully quantified Boolean formula ψ is true. The formula ψ can be assumed to be in prenex Conjunctive Normal Form (CNF) $Q_1x_1Q_2x_2\dots Q_mx_m\phi(X)$ such that the quantifiers are alternating existential and universal quantifiers ($Q_1 = \exists$, $Q_2 = \forall$, $Q_3 = \exists, \dots$), $X = \{x_1, x_2, \dots, x_m\}$ is the set of quantified variables, and $\phi(X) = C_1 \wedge \dots \wedge C_n$ is an unquantified Boolean formula over X equal to the conjunction of the clauses C_1, \dots, C_n .

Such a formula ψ is true if there exists a value of x_1 such that for all values of x_2 , there exists a value of $x_3 \dots$, such that the resulting valuation ν of all variables of X evaluates $\phi(X)$ to true. Formally, for each odd (resp. even) k , $1 \leq k \leq m$, let us denote by $f_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$ (resp. $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$) a valuation of variable x_k given a valuation of previous variables x_1, \dots, x_{k-1} ¹. Given these sequences $f = f_1, f_3, \dots$ and $g = g_2, g_4, \dots$, let us denote by $\nu = \nu_{(f,g)}$ the valuation of all variables of X such that $\nu(x_1) = f_1$, $\nu(x_2) = g_2(\nu(x_1))$, $\nu(x_3) = f_3(\nu(x_1)\nu(x_2))$, \dots . Then

$$\psi = Q_1x_1Q_2x_2\dots Q_mx_m\phi(X) \text{ is true} \\ \text{if and only if}$$

there exist $f = f_1, f_3, \dots$ such that for all $g = g_2, g_4, \dots$, the valuation $\nu_{f,g}$ evaluates $\phi(X)$ to true.

Let us detail a polynomial reduction from the QBF problem to the constrained existence problem of weak SPEs in Boolean games with qualitative Reachability objectives. Let $\psi = Q_1x_1Q_2x_2\dots Q_mx_m\phi(X)$ with $\phi(X) = C_1 \wedge \dots \wedge C_n$ be a fully quantified Boolean formula in prenex Conjunctive Normal Form. We build the following Boolean game $\mathcal{G}_\psi = (\Pi, V, (V_i)_{i \in \Pi}, E, (\text{Gain}_i)_{i \in \Pi})$ (see Figure 10.1):

- the set V of vertices:
 - for each variable $x_k \in X$ under quantifier Q_k , there exist vertices x_k , $\neg x_k$ and q_k ;
 - for each clause C_k , there exist vertices c_k and t_k ;

¹Notice that $f_1 : \emptyset \rightarrow \{0, 1\}$.

- there exists an additional vertex t_{n+1} ;
- the set E of edges:
 - from each vertex q_k there exist an edge to x_k and an edge to $\neg x_k$;
 - from each vertex x_k and $\neg x_k$, there exists an edge to q_{k+1} , except for $k = m$ where this edge is to c_1 ;
 - from each vertex c_k , there exist an edge to t_k and an edge to c_{k+1} , except for $k = n$ where there exist an edge to t_n and an edge to t_{n+1} ;
 - there exists a loop on each t_k ;
- the set Π of $n + 2$ players:
 - each player i , $1 \leq i \leq n$, owns vertex c_i ;
 - Player $n + 1$ (resp. $n + 2$) is the player who owns the vertices q_i for each existential (resp. universal) quantifier Q_i ;
 - as all other vertices have only one outgoing edge, it does not matter which player owns them;
- each function Gain_i is associated with the objective of visiting the set F_i defined as follows:
 - for all i , $1 \leq i \leq n$, $F_i = \{\ell \in V \mid \ell \text{ is a literal of clause } C_i\} \cup \{t_i\}$;
 - $F_{n+1} = \{t_{n+1}\}$;
 - $F_{n+2} = \{t_1, \dots, t_n\}$.




Figure 10.1: Reduction from the formula ψ to the Boolean game \mathcal{G}_ψ

Remark 10.2.2. (1) Notice that a sequence f of functions $f_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k odd, $1 \leq k \leq m$, as presented above, can be translated into a strategy σ_{n+1} of Player $n+1$ in the initialized game (\mathcal{G}_ψ, q_1) , and conversely. Similarly, a sequence g of functions $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$ is nothing else than a strategy σ_{n+2} of Player $n+2$. (2) Notice also that if ρ is a play in (\mathcal{G}_ψ, q_1) , then $\text{Gain}_{n+1}(\rho) = 1$ if and only if $\text{Gain}_{n+2}(\rho) = 0$. Moreover, suppose that ρ visits t_{n+1} , then for all i , $1 \leq i \leq n$, $\text{Gain}_i(\rho) = 1$ if and only if for all i , $1 \leq i \leq n$, ρ visits a vertex that is a literal of C_i if and only if there is a valuation of all variables of X that evaluates $\phi(X)$ to true.

Proof of Proposition 10.2.1. The game \mathcal{G}_ψ can be constructed from ψ in polynomial time. Let us now show that ψ is true if and only if there exists a weak SPE in (\mathcal{G}_ψ, q_1) with a gain profile $p \geq (0, \dots, 0, 1, 0)$ (that is, such that $p_{n+1} = 1$).

(\Rightarrow) Suppose that ψ is true. Then there exists a sequence f of functions $f_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k odd, $1 \leq k \leq m$, such that for all sequences g of functions $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$, the valuation $\nu_{f,g}$ evaluates $\phi(X)$ to true. We define a strategy profile σ as follows:

- for Player $n+1$, his strategy σ_{n+1} is the strategy corresponding to the sequence f (by Remark 10.2.2);
- for Player $n+2$, his strategy is an arbitrary strategy σ_{n+2} ; we denote by g the corresponding sequence $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$ (by Remark 10.2.2);
- for each player i , $1 \leq i \leq n$,
 - if $hv \in \text{Hist}_i(q_1)$ with $v = c_i$, is consistent with σ_{n+1} , then $\sigma_i(hv) = c_{i+1}$ if $i \neq n$ and t_{n+1} otherwise
 - else $\sigma_i(hv) = t_i$.

Let us prove that σ is a weak SPE, that is, for each history $hv \in \text{Hist}(q_1)$, there is no one-shot deviating strategy in the subgame $(\mathcal{G}_{\psi|h}, v)$ that is profitable to the player who owns vertex v (by Proposition 2.4.18). This is clearly true for all $v = t_i$, $1 \leq i \leq n+1$, since t_i has only one outgoing edge. For the other vertices v , we study two cases:

- hv is consistent with σ_{n+1} : First notice that $\text{Gain}(\rho) = (1, 1, \dots, 1, 0)$ with $\rho = h\langle\sigma_{\upharpoonright h}\rangle_v$. Indeed by hypothesis, the valuation $\nu_{f,g}$ evaluates $\phi(X)$ to true. Hence by Remark 10.2.2, the play ρ visits a vertex of F_i for all i , $1 \leq i \leq n$, and by definition of σ , ρ eventually loops on t_{n+1} . Second, as σ_{n+2} is arbitrary in the definition of σ , using another strategy σ'_{n+2} in place of σ_{n+2} will lead to a play ρ' such that $\text{Gain}(\rho') = \text{Gain}(\rho) = (1, 1, \dots, 1, 0)$.

Now if hv is consistent with σ_{n+1} , it is maybe not consistent with σ_{n+2} , but with another arbitrary strategy σ'_{n+2} , and thus $\text{Gain}(h\langle\sigma_{\upharpoonright h}\rangle_v) = (1, 1, \dots, 1, 0)$ as explained previously. Thus only Player $n + 2$ has an incentive to deviate in the subgame $(\mathcal{G}_{\psi \upharpoonright h}, v)$ to increase his gain. Nevertheless, using another strategy σ''_{n+2} will not change his gain (again by the same argument).

- hv is not consistent with σ_{n+1} : Suppose that $v = c_k$. Then by definition of σ , the play $h\langle\sigma_{\upharpoonright h}\rangle_v$ eventually loops on t_k leading to a gain of 1 for Player k . This player has thus no incentive to deviate with a one-shot deviation in the subgame $(\mathcal{G}_{\psi \upharpoonright h}, v)$.

Suppose that $v = q_k$. Then by definition of σ , the play $\rho = h\langle\sigma_{\upharpoonright h}\rangle_v$ eventually loops on t_1 . It follows that $\text{Gain}_{n+1}(\rho) = 0$ and $\text{Gain}_{n+2}(\rho) = 1$. As we only have to consider one-shot deviating strategies, if $q_k \in V_{n+2}$, Player $n + 2$ has no incentive to deviate, and if $q_k \in V_{n+1}$, Player $n + 1$ could try to use a one-shot deviating strategy, however the resulting play still eventually loops on t_1 .

This proves that σ is a weak SPE. Its gain profile is equal to $p = (1, 1, \dots, 1, 0)$ as explained previously. Therefore it satisfies the constraint $p \geq (0, \dots, 0, 1, 0)$.

(\Leftarrow) Suppose that there exists a weak SPE σ in $(\mathcal{G}_{\psi}, q_1)$ with outcome ρ and gain profile $\text{Gain}(\rho) \geq (0, \dots, 0, 1, 0)$, that is, $\text{Gain}_{n+1}(\rho) = 1$. By Remark 10.2.2, it follows that $\text{Gain}_{n+2}(\rho) = 0$. We have to prove that ψ is true. To this end, consider the sequence f of functions $f_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k odd, $1 \leq k \leq m$, that corresponds to strategy σ_{n+1} of Player $n + 1$

by Remark 10.2.2. Let us show that for all sequences g of functions $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$, the valuation $\nu_{f,g}$ evaluates $\phi(X)$ to true.

By contradiction assume that it is not the case for some sequence g' and consider the related strategy σ'_{n+2} of Player $n+2$ by Remark 10.2.2. Notice that σ'_{n+2} is a finitely deviating strategy. Let us consider the outcome ρ' of the strategy profile $(\sigma'_{n+2}, \sigma_{-(n+2)})$ from q_1 . As $\text{Gain}_{n+2}(\rho) = 0$, we must have $\text{Gain}_{n+2}(\rho') = 0$, otherwise σ'_{n+2} is a profitable deviation for Player $n+2$ whereas σ is a weak SPE. It follows that $\text{Gain}_{n+1}(\rho') = 1$ by Remark 10.2.2, that is, ρ' eventually loops on t_{n+1} .

Now recall that the valuation $\nu_{f,g'}$ evaluates $\phi(X)$ to false, which means that it evaluates some clause C_k of $\phi(X)$ to false. Consider the history $hc_k < \rho'$. As strategy σ'_{n+2} only acts on the left part of the underlying graph of \mathcal{G}_ψ , we have $\rho' = \langle \sigma'_{n+2}, \sigma_{-(n+2)} \rangle_{q_1} = h \langle \sigma_{\uparrow h} \rangle_{c_k}$. In the subgame $(\mathcal{G}_{\psi \uparrow h}, c_k)$, the outcome of $\sigma_{\uparrow h}$ gives a gain of 0 to Player k because $\rho' = h \langle \sigma_{\uparrow h} \rangle_{c_k}$ does not visit t_k and $\nu_{f,g'}$ evaluates C_k to false. In this subgame, Player k has thus a profitable one-shot deviation: to move to t_k . It follows that σ is not a weak SPE which is impossible. Therefore ψ is true. \square

For multiplayer Safety games, we can use the same reduction and the same kind of arguments as for qualitative Reachability objectives.

Proposition 10.2.3. *The constrained existence problem of weak SPEs in multiplayer Safety games is PSPACE-hard.*

Proof sketch. Given a fully quantified Boolean formula ψ , we construct the same game as in the proof of Proposition 10.2.1 (see Figure 10.1), except that each player i , $1 \leq i \leq n+2$, aims at avoiding the set F'_i (instead of visiting the set F_i) defined as follows:

- for all i , $1 \leq i \leq n$, $F'_i = \{\ell \in V \mid \ell \text{ is a literal of clause } C_i\} \cup \{t_{n+1}\}$;
- $F'_{n+1} = \{t_1, t_2, \dots, t_n\}$;

- $F'_{n+2} = \{t_{n+1}\}$.

Recall how the sets F_i were defined: $F_{n+1} = \{t_{n+1}\}$, $F_{n+2} = \{t_1, t_2, \dots, t_n\}$, and for all i , $1 \leq i \leq n$, $F_i = \{\ell \in V \mid \ell \text{ is a literal of clause } C_i\} \cup \{t_i\}$. Hence we have a clear duality for players $n+1$ and $n+2$: a play ρ visits F_{n+1} (resp. F_{n+2}) if and only if ρ avoids F'_{n+1} (resp. F'_{n+2}). This is not the case for the other players, but one can check that the proof works in the same way as for qualitative Reachability games. \square

10.3 PSPACE-completeness of the constrained existence problem of SPEs

In this section we prove Corollary 10.0.2 stating that the constrained existence problem of SPEs (instead of weak SPEs) in multiplayer Boolean games with qualitative Reachability and Safety objectives is PSPACE-complete.

Proof of Corollary 10.0.2. As weak SPEs and SPEs are equivalent notions for qualitative Reachability objectives (Proposition 7.5.4), by Theorem 10.0.1, the constrained existence problem of SPEs in multiplayer qualitative Reachability games is PSPACE-complete.

We need to use other arguments for the case of Safety objectives. The reduction from QBF proposed in the proof of Proposition 10.2.3 uses the game \mathcal{G}_ψ of Figure 10.1. Due to the structure of the underlying graph, all weak SPEs of \mathcal{G}_ψ are SPEs since any deviating strategy from a given strategy is necessarily finitely deviating. This shows that the constrained existence problem of SPEs is PSPACE-hard for Safety objectives. It is proved in [Umm05] that this problem is in PSPACE. \square

10.4 Memory requirements

We conclude this chapter by proving that given an initialized multiplayer qualitative Reachability or Safety game (\mathcal{G}, v_0) and its associated extended game (\mathcal{X}, x_0) , there exists a weak SPE σ in (\mathcal{X}, x_0) with gain profile $p \in \{0, 1\}^{|\Pi|}$ if

and only if there exists a weak SPE τ in (\mathcal{X}, x_0) with gain profile p and which is finite-memory.

Proposition 10.4.1. *Let (\mathcal{G}, v_0) be either a initialized multiplayer qualitative Reachability or Safety game and let (\mathcal{X}, x_0) be its associated extended game. Given a gain profile $p \in \{0, 1\}^{|\Pi|}$, the following assertions are equivalent.*

1. *There exists a weak SPE σ in (\mathcal{X}, x_0) such that $\text{Gain}(\langle \sigma \rangle_{x_0}) = p$;*
2. *There exists a finite good symbolic witness \mathcal{P} in (\mathcal{X}, x_0) such that:*
 - *there is $\rho^{(0, x_0)} \in \mathcal{P}$ such that $\text{Gain}(\rho^{(0, x_0)}) = p$;*
 - *for each lasso $\rho^{(i, x)} \in \mathcal{P}$, this lasso has a length bounded by $2 \cdot |V^X|^2 = 2 \cdot 2^{2 \cdot |\Pi|} \cdot |V|^2$.*
3. *There exists a weak SPE τ in (\mathcal{X}, x_0) such that $\text{Gain}(\langle \tau \rangle_{x_0}) = p$ and which has a memory size in $\mathcal{O}(|\Pi| \cdot |V^X|^3) = \mathcal{O}(2^{3 \cdot |\Pi|} \cdot |\Pi| \cdot |V|^3)$.*

Proof sketch. One way to be convinced by this result is the following. In the extended game, the gain functions qR_i and Safe_i , for some $i \in \Pi$, are prefix-independent (Proposition 4.2.7). Thus, one can hope that the same proof as the one used for Proposition 9.2.2 is also correct.

(1 \Rightarrow 2) The same arguments applied on the extended to hold. Indeed, the gain functions in the extended game are also prefix-independent. Moreover, the lassoes obtained by the procedure are such that, if π' is obtained from π : (i) if $\pi' = h\ell^\omega$, then $|h\ell| \leq 2 \cdot |V^X|^2$ (ii) $\text{First}(\pi') = \text{First}(\pi)$, $\text{Occ}(\pi') = \text{Occ}(\pi)$ and $\text{Inf}(\pi') = \text{Inf}(\pi)$. It follows that $\text{qR}^X(\pi') = \text{qR}^X(\pi)$ (resp. $\text{Safe}^X(\pi') = \text{Safe}^X(\pi)$). Finally, in the same way as in the proof of Proposition 9.2.2, by replacing $\text{Inf}(\pi') = \text{Inf}(\pi)$ by $\text{Occ}(\pi') = \text{Occ}(\pi)$ when needed, we have that if π is λ^* -consistent then π' is also λ^* -consistent. Hence, the obtained finite symbolic witness in (\mathcal{X}, x_0) is good by construction and composed of lassoes with length bounded by $2 \cdot |V^X|^2 = 2 \cdot 2^{2 \cdot |\Pi|} \cdot |V|^2$.

(2 \Rightarrow 3): This implication follows once again from Corollary 7.2.6 applied to

the extended game. In this way, we obtain a weak SPE τ in (\mathcal{X}, x_0) such that $\text{Gain}(\langle \tau \rangle_{x_0}) = \text{Gain}(\rho^{(0, x_0)}) = p$ and with memory size in $\mathcal{O}(|\Pi| \cdot |V^X| \cdot L)$ with $L = 2 \cdot |V^X|^2$. Therefore, we obtain a memory size in $\mathcal{O}(|\Pi| \cdot |V|^3 \cdot 2^{3 \cdot |\Pi|})$.

(3 \Rightarrow 1): Obvious. \square

Remark 10.4.2. Notice that we cannot use this result in the proof of the PSPACE-membership of the constrained existence problem (Proposition 10.1.1) since the obtained lassoes do not have a polynomial length.

Remark 10.4.3. On the contrary, we could use the elimination of unnecessary cycles explained in the proof of Proposition 10.1.1 in the proof of Proposition 10.4.1 to obtain a better bound on the length of the lassoes. But the resulting memory size would also be exponential in the size of (\mathcal{G}, v_0) , since the memory size would be in $\mathcal{O}(|V^X| \cdot |\Pi| \cdot L)$ where L is an upper bound on the length of the lassoes.

Remark 10.4.4. Even if this result provide an upper bound on the memory size needed for a weak SPE in the extended game associated with a qualitative Reachability or Safety game (\mathcal{G}, v_0) , the corresponding weak SPE in (\mathcal{G}, v_0) needs the same amount of memory.

CHAPTER 11

QUANTITATIVE REACHABILITY GAMES

In this chapter, based on [BBG⁺19, BBG⁺20], we prove that the CEP of SPEs in multiplayer quantitative Reachability games is PSPACE-complete. In fact, we prove that the CEP of weak SPEs in this kind of games is PSPACE-complete, but since the concepts of SPE and weak SPE are equivalent in quantitative Reachability games (Corollary 2.4.23), this latter result allows us to obtain the same complexity result for SPEs.

Theorem 11.0.1. *The constrained existence problem of SPEs in multiplayer quantitative Reachability games is PSPACE-complete.*

Even if the outcome characterization of weak SPEs in quantitative Reachability games (Corollary 7.4.13) provided in Section 7.4.3 allows to exactly characterize the set of weak SPEs (and so of SPEs) outcomes in a multiplayer quantitative Reachability game, the way the labeling functions $(\lambda^k)_{k \in \mathbb{N}}$, used for this characterization, are defined does not allow us to obtain the PSPACE algorithm to decide the constrained existence problem.

This is the reason why we need a more subtle approach. Roughly speaking, the labeling function λ^k are still defined on the associated extended game but this latter game is divided into *regions*. A region I in the extended game

is the sub-arena in which all players belonging to I have already visited their target set. The intuition is that these regions correspond to strongly connected components (SCCs) in the extended game, and the algorithm which computes the labeling functions λ^k , for all $k \in \mathbb{N}$, proceeds in a bottom-up fashion: it begins with the bottom strongly connected components (BSCCs) and once the BSCCs are processed, the algorithm computes the values of all λ^k in the SCCs which are a predecessor of a BSCC, and so on. For the same arguments as those of the characterization which relies on the labeling functions given in Definition 7.4.9, the plays which are λ^* -consistent for the labeling function λ^* obtained by this new procedure are exactly the plays which are outcomes of weak SPEs. The PSPACE membership of the CEP is obtained due to a fine-grained analysis of this procedure.

In Section 11.1, we formally define the notion of region and this bottom-up algorithm. Our PSPACE algorithm to decide the CEP relies both on this procedure and on the notion of *counter graph*: a graph which keeps track of the constraints imposed by a labeling function along plays in a game. This concept is formally defined in Section 11.2 and we provide some useful results for the proof of PSPACE-easiness. Finally, Section 11.3 provides the proofs of PSPACE-hardness and PSPACE-membership of the constrained existence problem of weak SPEs (and SPEs) in multiplayer quantitative Reachability games.

Since the approach is quite the same as the one followed in Chapter 7, the proofs of Section 11.1 are relegated in Appendix B.1.1.

All along this chapter a game $\mathcal{G} = (A, \text{Cost})$ depicts a multiplayer **quantitative** Reachability games, *i.e.*, $\text{Cost}_i = \text{QR}_i$ for all $i \in \Pi$. Moreover, as already done in some parts of this document, we do not always write Cost_i^X to refer to the cost function of Player i in the extended game. We prefer to only write Cost_i when it is clear from the context that we are considering the extended game.

11.1 Regions decomposition and regions based algorithm

11.1.1 Regions decomposition

By construction, the arena X of the initialized extended game is divided into different regions according to the players who have already visited their target set. Let us provide some useful notions with respect to this decomposition. (An illustrative example is given hereafter.) We will often use them in the following sections. Let $\mathcal{I}^1 = \{I \subseteq \Pi \mid \text{there exists } v \in V \text{ such that } (v, I) \in \text{Succ}^*(x_0)\}$ be the set of sets I accessible from the initial state x_0 , and let $N = |\mathcal{I}|$ be its size. For $I, I' \in \mathcal{I}$, if there exists $((v, I), (v', I')) \in E^X$, we say that I' is a *successor* of I and we write $I' \in \text{Succ}(I)$. Given $I \in \mathcal{I}$, $X^I = (V^I, E^I)$ refers to the sub-arena of X restricted to the vertices $\{(v, I) \in V^X \mid v \in V\}$. Hence X^I has all its vertices with the same second component I . We say that X^I is the *region*² associated with I . Such a region X^I is called a *bottom region* whenever $\text{Succ}(I) = I$.

There exists a partial order on \mathcal{I} such that $I < I'$ if and only if $I' \in \text{Succ}^*(I) \setminus \{I\}$. We fix an arbitrary *total order* on \mathcal{I} that extends this partial order $<$ as follows:

$$J_1 < J_2 < \dots < J_N. \quad (11.1)$$

(with X^{J_N} a bottom region).³ With respect to this total order, given $n \in \{1, \dots, N\}$, we denote by $X^{\geq J_n} = (V^{\geq J_n}, E^{\geq J_n})$ the sub-arena of X restricted to the vertices $\{(v, I) \in V^X \mid I \geq J_n\}$.

The total order given in (11.1) together with the I -monotonicity (see (4.1)) leads to the following lemma.

¹This notation should not be confused with the one chosen to denote the set of tuples (player, vertex) used in the definition of a symbolic witness in Section 7.2.1.

²In the rest of this part, we indifferently call region either X^I , or V^I , or I .

³We use notation J_n , $n \in \{1, \dots, N\}$, to avoid any confusion with the sets I_k appearing in a play $\rho = (v_0, I_0)(v_1, I_1) \dots$

Lemma 11.1.1 (Region decomposition and section). *Let π be a path in the arena X of the extended game \mathcal{X} . Then there exists a region decomposition of π as*

$$\pi[\ell]\pi[\ell + 1] \dots \pi[m]$$

with $1 \leq \ell \leq m \leq N$, such that for each n , $\ell \leq n \leq m$:

- $\pi[n]$ is a (possibly empty) path in X ,
- every vertex of $\pi[n]$ is of the form (v, J_n) for some $v \in V$.

Each path $\pi[n]$ is called a section. The last section $\pi[m]$ is infinite if and only if π is infinite.

Example 11.1.2. Let us come back to the initialized game (\mathcal{G}, v_0) of Figure 4.1. Its extended game (\mathcal{X}, x_0) is depicted in Figure 4.3 (only the part reachable from the initial vertex $x_0 = (v_0, \emptyset)$ is depicted). As we can see, the extended game is divided into three different regions: one region associated with $I = \emptyset$ that contains the initial vertex x_0 , a second region associated with $I = \{2\}$, and a third bottom region associated with $I = \{1, 2\} = \Pi$. Hence the set $\mathcal{I} = \{\emptyset, \{2\}, \Pi\}$ is totally ordered as $J_1 = \emptyset < J_2 = \{2\} < J_3 = \Pi$.

For all vertices (v, I) of the region associated with $I = \{2\}$, we have $(v, I) \notin F_1^X$ and $(v, I) \in F_2^X$, and for those of the region associated with $I = \Pi$, we have $(v, I) \in F_1^X \cap F_2^X$. The sub-arena $X^{\geq J_2}$ of X is composed of all vertices (v, I) such that $I = \{2\}$ or $I = \Pi$.

From the SPE σ given in Example 7.4.14 with outcome $\rho = (v_0 v_1 v_6 v_7 v_2)^\omega \in \text{Plays}_A(v_0)$ and cost $(4, 4)$, we derive the SPE outcome $\rho^X \in \text{Plays}_X(x_0)$ equal to

$$\rho^X = (v_0, \emptyset)(v_1, \emptyset)(v_6, \emptyset)(v_7, \emptyset)((v_2, \Pi)(v_0, \Pi)(v_1, \Pi)(v_6, \Pi)(v_7, \Pi))^\omega$$

with the same cost $(4, 4)$.

The region decomposition of ρ^X is equal to $\rho^X[1]\rho^X[2]\rho^X[3]$ such that its second section $\rho^X[2]$ is empty, and its two other sections $\rho^X[1]$ and $\rho^X[3]$ are respectively equal to

$$(v_0, \emptyset)(v_1, \emptyset)(v_6, \emptyset)(v_7, \emptyset), \text{ and } ((v_2, \Pi)(v_0, \Pi)(v_1, \Pi)(v_6, \Pi)(v_7, \Pi))^\omega.$$

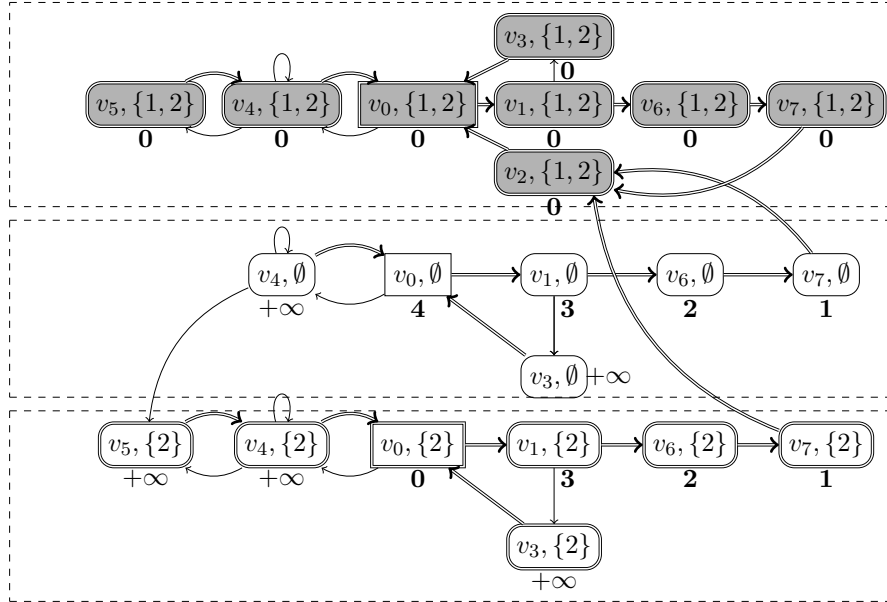


Figure 11.1: The extended game (\mathcal{X}, x_0) for the initialized game (\mathcal{G}, v_0) of Figure 4.1. The values of a labeling function λ^* are indicated near each vertex. The dashed rectangles correspond to the different regions reachable from x_0 .

11.1.2 Regions Based Algorithm

Definition 11.1.3 (Initial labeling). For all $v \in V^X$, let $i \in \Pi$ be such that $v \in V_i^X$,

$$\lambda^0(v) = \begin{cases} 0 & \text{if } i \in I(v) \\ +\infty & \text{otherwise.} \end{cases}$$

This labeling function λ^0 does not impose any constraint on the plays.

Lemma 11.1.4. $\rho \in \Lambda_0(v)$ if and only if $\rho \in \text{Plays}_X(v)$. □

Let us now explain how our algorithm computes the labeling functions λ^k , $k \geq 1$, and the related sets $\Lambda^k(v)$, $v \in V^X$. It works in a *bottom-up* manner,

according to the total order $J_1 < J_2 < \dots < J_N$ of \mathcal{I} given in (11.1): it first iteratively updates the labeling function for all vertices v of the arena X^{J_N} until reaching a fixpoint in this arena, it then repeats this procedure in $X^{\geq J_{N-1}}, X^{\geq J_{N-2}}, \dots, X^{\geq J_1} = X$. Hence, suppose that we currently treat the arena $X^{\geq J_n}$ and that we want to compute λ^{k+1} from λ^k . We define the updated function λ^{k+1} as follows (we use the convention that $1 + (+\infty) = +\infty$).

Definition 11.1.5 (Labeling update). Let $k \geq 0$ and suppose that we treat the arena $X^{\geq J_n}$, with $n \in \{1, \dots, N\}$. For all $v \in V^X$,

- if $v \in V^{\geq J_n}$, let $i \in \Pi$ be such that $v \in V_i^X$, then

$$\lambda^{k+1}(v) = \begin{cases} 0 & \text{if } i \in I(v) \\ 1 + \min_{(v,v') \in E^X} \sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v')\} & \text{otherwise} \end{cases}$$

- if $v \notin V^{\geq J_n}$, then

$$\lambda^{k+1}(v) = \lambda^k(v).$$

Let us provide some explanations. As this update concerns the arena $X^{\geq J_n}$, we keep $\lambda^{k+1} = \lambda^k$ outside of this arena. Suppose now that v belongs to $X^{\geq J_n}$ and let i be such that $v \in V_i^X$. We define $\lambda^{k+1}(v) = 0$ whenever $i \in I(v)$ (as already explained for the definition of λ^0). When it is updated, the value $\lambda^{k+1}(v)$ represents what is the best cost that player i can ensure for himself from v with a “one-shot” choice by only taking into account plays of $\Lambda^k(v')$ with $v' \in \text{Succ}(v)$.

Notice that it makes sense to run the algorithm in a bottom-up fashion according to the total ordering $J_1 < \dots < J_N$ since given a play $\rho = \rho_0 \rho_1 \dots$, if ρ_0 is a vertex of $V^{\geq J_n}$, then for all $k \in \mathbb{N}$, ρ_k is a vertex of $V^{\geq J_n}$ (by I -monotonicity). Moreover running the algorithm in this way is essential to prove that the constrained existence problem in quantitative Reachability games is in PSPACE.

Example 11.1.6. We consider again the extended game (\mathcal{X}, x_0) of Figure 11.1 with the total order $J_1 = \emptyset < J_2 = \{2\} < J_3 = J_N = \Pi$ of its set \mathcal{I} .

Let us illustrate Definition 11.1.5 on the arena $X^{\geq J_2}$. Let $k \geq 0$ and suppose that the labeling function λ^k has been computed such that $\lambda^k(v_0, J_2) = 0$, $\lambda^k(v, J_2) = +\infty$ for every other vertex in region X^{J_2} (notice that λ^k is not the labeling function indicated in Figure 11.1). Let us show how to compute $\lambda^{k+1}(v_1, J_2)$. We need to compute $\sup\{\text{Cost}_1(\rho) \mid \rho \in \Lambda^k(v')\}$ for the two successors v' of (v_1, J_2) , that is, respectively $v' = (v_3, J_2)$ and $v' = (v_6, J_2)$. Recall that $\Lambda^k(v')$ is the set of all plays λ^k -consistent from v' . All the plays in $\Lambda^k(v_6, J_2)$ have cost 2 for player 1 and player 2: indeed, they all first follow the history $(v_6, J_2), (v_7, J_2), (v_2, J_N)$. Thus, $\sup\{\text{Cost}_1(\rho) \mid \rho \in \Lambda^k(v_6, J_2)\} = 2$. On the other hand, as $\lambda^k(v_3, J_2) = +\infty$, the play $(v_3, J_2)(v_0, J_2)(v_4, J_2)^\omega$ is λ^k -consistent and belongs to $\Lambda^k(v_3, J_2)$. Thus, $\sup\{\text{Cost}_1(\rho) \mid \rho \in \Lambda^k(v_3, J_2)\} = +\infty$. Hence, the minimum is attained with successor (v_6, J_2) , and $\lambda^{k+1}(v_1, J_2) = 3$.

We can now provide our algorithm (see Algorithm 4) that computes the sequence $(\lambda^k(v))_{k \in \mathbb{N}}$ until a fixpoint is reached (see Proposition 11.1.7 below). Initially, the labeling function is λ^0 (see Definition 11.1.3). For the next steps $k > 0$, we begin with the bottom region X^{J_N} of \mathcal{X} and update λ^{k-1} to λ^k as described in Definition 11.1.5. At some point, the values of λ^k do not change anymore in X^{J_N} and $(\lambda^k)_{k \in \mathbb{N}}$ reaches locally (on X^{J_N}) a fixpoint (see again Proposition 11.1.7). Now, we consider the arena $X^{\geq J_{N-1}}$ and in the same way, we continue to update locally the values of λ^k in $X^{\geq J_{N-1}}$. We repeat this procedure with arenas $X^{\geq J_{N-2}}, X^{\geq J_{N-3}}, \dots$ until the arena $X^{\geq J_1} = X$ is completely processed. From the last computed λ^k , we derive the sets $\Lambda^k(v)$, $v \in V^X$, that we need for the characterization of outcomes of SPEs (see Theorem 11.1.9 below). An example illustrating the execution of this algorithm is given at the end of this section.

We now state that the sequence $(\lambda^k)_{k \in \mathbb{N}}$ computed by this algorithm reaches a fixpoint - locally on each arena $X^{\geq J_n}$ and globally on X - in the following meaning:

Proposition 11.1.7. *There exists a sequence of integers $0 = k_N^* < k_{N-1}^* < \dots < k_1^* = k^*$ such that*

Algorithm 4: Fixpoint

```

1 compute  $\lambda^0$ 
2  $k \leftarrow 0$ 
3  $n \leftarrow N$ 
4 while  $n \neq 0$  do
5   repeat
6      $k \leftarrow k + 1$ 
7     compute  $\lambda^k$  from  $\lambda^{k-1}$  with respect to  $X^{\geq J_n}$ 
8   until  $\lambda^k = \lambda^{k-1}$ 
9    $n \leftarrow n - 1$ 
10 return  $\lambda^k$ 

```

- **Local fixpoint:** for all $J_n \in \mathcal{I}$, all $m \in \mathbb{N}$ and all $v \in V^{\geq J_n}$,

$$\lambda^{k_n^*+m}(v) = \lambda^{k_n^*}(v), \quad (11.2)$$

- **Global fixpoint:** in particular, with $k^* = k_1^*$, for all $m \in \mathbb{N}$ and all $v \in V^X$,

$$\lambda^{k^*+m}(v) = \lambda^{k^*}(v).$$

The global fixpoint λ^{k^*} is also simply denoted by λ^* , and each set $\Lambda^{k^*}(v), v \in V^X$, is denoted by $\Lambda^*(v)$.

This proposition indicates that Algorithm 4 terminates. Indeed for each $J_n \in \mathcal{I}$, taking the *least index* k_n^* which makes Equality (11.2) true shows that the repeat loop is broken and the variable n decremented by 1. The value $n = 0$ is eventually reached and the algorithm stops with the global fixpoint λ^* . Notice that the first local fixpoint is reached with $k_N^* = 0$ because X^{J_N} is a bottom region. To prove that the algorithm stops in a finite number of steps, we show that each region requires a finite number of steps to be treated. To do so, we rely on the fact that constraint values on vertices cannot increase from one step to another and that the set of constraint values can be seen as a well-quasi ordering. Thus a decreasing sequence in this set is stationary, and

this means our algorithm reaches a fixpoint and terminates.

Proposition 11.1.7 also shows that when a local fixpoint is reached in the arena $X^{\geq J_{n+1}}$ and the algorithm updates the labeling function λ^k in the arena $X^{\geq J_n}$, the values of $\lambda^k(v)$ do not change anymore for any $v \in V^{\geq J_{n+1}}$ but can still be modified for some $v \in V^{J_n}$. Recall also that outside of $X^{\geq J_n}$, the values of $\lambda^k(v)$ are still equal to the initial values $\lambda^0(v)$. These properties will be useful when we will prove that the constrained existence problem of SPEs in quantitative Reachability games is in PSPACE. They are summarized in the next lemma.

Lemma 11.1.8. *Let $k \in \mathbb{N}$ be a step of the algorithm and let J_n with $n \in \{1, \dots, N\}$. For all $v \in V^{J_n}$:*

- if $k \leq k_{n+1}^*$, then $\lambda^{k+1}(v) = \lambda^k(v) = \lambda^0(v)$,
- if $k_n^* \leq k$, then $\lambda^{k+1}(v) = \lambda^k(v) = \lambda^{k_n^*}(v)$,

Hence the values of $\lambda^k(v)$ and $\lambda^{k+1}(v)$ may be different only when $k_{n+1}^* < k < k_n^*$. □

We are ready to state how we characterize plays that are outcomes of SPEs. This is possible with the global fixpoint computed by Algorithm 4.

Theorem 11.1.9 (Characterization). *Let (\mathcal{G}, v_0) be an initialized quantitative game and (\mathcal{X}, x_0) be its initialized extended game. Let ρ^0 be a play in $\text{Plays}_X(x_0)$. Then ρ^0 is the outcome of an SPE in (\mathcal{X}, x_0) if and only if $\rho^0 \in \Lambda^*(x_0)$.*

Notice that this theorem also provides a characterization of the outcomes of SPEs in (\mathcal{G}, v_0) by Lemma 4.2.6.

As in Proposition 7.4.11, when for some player i , the costs $\text{Cost}_i(\rho)$ associated with the plays ρ in $\Lambda^k(v)$ are unbounded, there actually exists some play in this set that has an infinite cost. In other terms, either $\Lambda^k(v)$ contains at least one play ρ with an infinite cost $\text{Cost}_i(\rho)$ or there exists a constant $c \in \mathbb{N}$ such that $\text{Cost}_i(\rho) \leq c$ for all $\rho \in \Lambda^k(v)$.

Proposition 11.1.10. *For every $k \in \mathbb{N}$, for every $v \in V^X$ and for every $i \in \Pi$, the following implication holds: if $\sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v)\} = +\infty$, then there exists a play $\rho \in \Lambda^k(v)$ such that $\text{Cost}_i(\rho) = +\infty$.*

The next corollary is a direct consequence of Proposition 11.1.10.

Corollary 11.1.11. $\sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v)\} = \max\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v)\}$.

Example 11.1.12. Let us come back to the running example of Figure 11.1 and illustrate Proposition 11.1.7. The different steps of Algorithm 4 are given in Table 11.1. The columns indicate the vertices according to their region, respectively Π , $\{2\}$, and \emptyset . Notice that for the region Π , we only write one column v as for all vertices (v, Π) the value of λ is equal to 0 all along the algorithm.

Recall that $J_1 = \emptyset < J_2 = \{2\} < J_3 = \Pi = \{1, 2\}$. The algorithm begins with the arena X^{J_3} . A local fixpoint ($\lambda^1 = \lambda^0$) is immediately reached because all vertices belong to the target set of both players in X^{J_3} . Thus the first local fixpoint is reached with $k_3^* = 0$.

The algorithm then treats the arena $X^{\geq J_2}$. By Lemma 11.1.8, it is enough to consider the region X^{J_2} . Let us explain how to compute $\lambda^2(v)$ from $\lambda^1(v)$ on this region. For $v = (v_7, \{2\})$, we have that $\lambda^2(v) = 1 + \min_{(v, v') \in E^X} \sup\{\text{Cost}_1(\rho) \mid \rho \in \Lambda^1(v')\}$. As the unique successor of v is $(v_2, \{1, 2\})$, all λ^1 -consistent plays beginning in this successor have cost 0 for player 1. So, we have that $\lambda^2(v) = 1$. For the computation of $\lambda^2(v_6, \{2\})$, the same argument holds since $(v_6, \{2\})$ has the unique successor $(v_7, \{2\})$. The vertex $(v_1, \{2\})$ has two successors: $(v_6, \{2\})$ and $(v_3, \{2\})$.⁴ Again, we know that all λ^1 -consistent plays beginning in $(v_6, \{2\})$ have cost 2 for player 1. From $(v_3, \{2\})$ however, one can easily check that the play $(v_3, \{2\})(v_0, \{2\})((v_4, \{2\}))^\omega$ is λ^1 -consistent and has cost $+\infty$ for player 1. Thus, we obtain that $\lambda^2(v_1, \{2\}) = 3$. For the other vertices of X^{J_2} , one can see that $\lambda^2(v) = \lambda^1(v)$.

Finally, we can check that the local fixed point is reached in the arena $X^{\geq J_2}$

⁴The computation of $\lambda^2(v_1, \{2\})$ was already explained in Example 11.1.6.

Table 11.1: The different steps of the algorithm computing λ^* for the extended game of Figure 11.1

Region	{1, 2}	{2}							\emptyset					
	v	v_0	v_1	v_6	v_7	v_3	v_4	v_5	v_0	v_1	v_6	v_7	v_3	v_4
$\lambda^0 = \lambda^1$	0	0	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$
$\lambda^2 = \lambda^3$	0	0	3	2	1	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$
λ^4	0	0	3	2	1	$+\infty$	$+\infty$	$+\infty$	$+\infty$	3	2	1	$+\infty$	$+\infty$
$\lambda^5 = \lambda^*$	0	0	3	2	1	$+\infty$	$+\infty$	$+\infty$	4	3	2	1	$+\infty$	$+\infty$

(resp. $X^{\geq J_1}$) with $\lambda^3 = \lambda^2$ (resp. $\lambda^6 = \lambda^5 = \lambda^*$). Therefore the respective fixpoints are reached with $k_2^* = 2$ and $k_1^* = 5$. The labeling function indicated in Figure 11.1 is the one of λ^* .

11.2 Counter graph

In the previous section, we have proposed an algorithm that computes a sequence of labeling functions $(\lambda^k)_{k \in \mathbb{N}}$ until reaching a fixpoint λ^* such that the plays that are λ^* -consistent are exactly the SPE outcomes. In this section, given a labeling function λ , we introduce the concept of *counter graph* such that its infinite paths coincide with the plays that are λ -consistent. We then show that the counter graph associated with the fixpoint function λ^* has an exponential size, an essential step to prove PSPACE membership of the constrained existence problem of SPEs in quantitative Reachability games. Technical proofs of this section are provided in Appendix B.1.2.

For the entire section, we fix a quantitative reachability game $\mathcal{G} = (\mathbb{A}, (\text{Cost}_i)_{i \in \Pi}, (F_i)_{i \in \Pi})$ with an arena $\mathbb{A} = (\Pi, V, (V_i)_{i \in \Pi}, E)$, and v_0 an initial vertex. Let $\mathcal{X} = (X, (\text{Cost}_i)_{i \in \Pi}, (F_i^X)_{i \in \Pi})$ with the arena $X = (\Pi, V^X, (V_i^X)_{i \in \Pi}, E^X)$ be its associated extended game. Furthermore, when we speak about a play ρ we always mean a play in the extended game \mathcal{X} .

A labeling function λ give constraints on costs of plays from each vertex in X , albeit only for the player that owns this vertex. However, by the property of λ -consistence, constraints for a player carry over all the successive vertices, whether they belong to him or not. In order to check efficiently this property, we introduce the counter graph to keep track *explicitly* of the accumulation of constraints for *all* players at each step of a play. Let us first fix some notation.

Definition 11.2.1 (Restriction and maximal finite range). Let $\lambda : V^X \rightarrow \mathbb{N} \cup \{+\infty\}$ be a labeling function.

- We consider *restrictions* of λ to sub-arenas of V^X as follows. Let $n \in \{1, \dots, N\}$, we denote by $\lambda_n : V^{J_n} \rightarrow \mathbb{N} \cup \{+\infty\}$ the restriction of λ to V^{J_n} . Similarly we denote by $\lambda_{\geq n}$ (resp. $\lambda_{>n}$) the restriction of λ to $V^{\geq J_n}$ (resp. $V^{>J_n}$).
- The *maximal finite range* of λ , denoted by $\text{mR}(\lambda)$, is equal to

$$\text{mR}(\lambda) = \max\{c \in \mathbb{N} \mid \lambda(v) = c \text{ for some } v \in V^X\}$$

with the convention that $\text{mR}(\lambda) = 0$ if λ is the constant function $+\infty$. We also extend this notion to restrictions of λ with the convention that $\text{mR}(\lambda_{>n}) = 0$ if J_n is a bottom region.

Notice that in the definition of maximal finite range, we only consider the *finite* values of λ (and not the value $+\infty$).

Definition 11.2.2 (Counter Graph). Let $\lambda : V^X \rightarrow \mathbb{N} \cup \{+\infty\}$ be a labeling function. Let $\mathcal{K} := \{0, \dots, K\} \cup \{+\infty\}$ with $K = \text{mR}(\lambda)$. The *counter graph* $\mathbb{C}(\lambda)$ for \mathcal{G} and λ is equal to $\mathbb{C}(\lambda) = (\Pi, V^C, (V_i^C)_{i \in \Pi}, E^C)$, such that:

- $V^C = V^X \times \mathcal{K}^{|\Pi|}$
- $(v, (c_i)_{i \in \Pi}) \in V_j^C$ if and only if $v \in V_j^X$
- $((v, (c_i)_{i \in \Pi}), (v', (c'_i)_{i \in \Pi})) \in E^C$ if and only if:
 - $(v, v') \in E^X$, and
 - for every $i \in \Pi$

$$c'_i = \begin{cases} 0 & \text{if } i \in I(v') \\ c_i - 1 & \text{if } i \notin I(v'), v' \notin V_i^X \text{ and } c_i > 1 \\ \min(c_i - 1, \lambda(v')) & \text{if } i \notin I(v'), v' \in V_i^X \text{ and } c_i > 1. \end{cases}$$

Intuitively, the counter graph is constructed such that once a value $\lambda(v)$ is finite for a vertex $v \in V_i^X$ along a play in \mathcal{X} , the corresponding path in $\mathbb{C}(\lambda)$ keeps track of the induced constraint by (i) decrementing the counter value c_i for the concerned player i by 1 at every step, (ii) updating this counter if a stronger constraint for player i is encountered by visiting a vertex v' with a smaller value $\lambda(v')$, and (iii) setting the counter c_i to 0 if player i has reached his target set.

Note that in the counter graph, there may be some vertices with *no outgoing edges*. Indeed, consider a vertex $(v, (c_i)_{i \in \Pi}) \in V^C$ such that $c_j = 1$ for some player j . By construction of $\mathbb{C}(\lambda)$, the only outgoing edges from $(v, (c_i)_{i \in \Pi})$ must link to vertices $(v', (c'_i)_{i \in \Pi})$ such that $(v, v') \in E^X$, $c'_j = 0$ and $j \in I(v')$ (as in Definition 11.2.2 the two other cases require that $c_j > 1$). However, it may be the case that no successor v' of v in X is such that $j \in I(v')$.

Note as well that for each vertex $v \in V^X$, there exist many different vertices $(v, (c_i)_{i \in \Pi})$ in $\mathbb{C}(\lambda)$, one for each counter values profile. However, the intended goal of the counter graph is to monitor explicitly the constraints accumulated by each player along a play in \mathcal{X} regarding the function λ . Thus, we will only consider paths in $\mathbb{C}(\lambda)$ that start in vertices $(v, (c_i)_{i \in \Pi})$ such that the counter values correspond indeed to the constraint at the *beginning of a play* in \mathcal{X} regarding λ :

Definition 11.2.3 (Starting vertex in $\mathbb{C}(\lambda)$). Let $v \in V^X$. We distinguish one vertex $v^C = (v, (c_i)_{i \in \Pi})$ in V^C , such that for every $i \in \Pi$:

$$c_i = \begin{cases} 0 & \text{if } i \in I(v) \\ \lambda(v) & \text{if } i \notin I(v) \text{ and } v \in V_i^X \\ +\infty & \text{otherwise.} \end{cases}$$

We call v^C the *starting vertex* associated with v , and denote by $\text{SV}(\lambda)$ the set of all starting vertices in $\mathbb{C}(\lambda)$.

Example 11.2.4. Recall the extended game (\mathcal{X}, x_0) of Figure 11.1, and the labeling function λ whose values are indicated under each vertex. In Figure 11.2, we illustrate a part of the counter graph $\mathbb{C}(\lambda)$. In Example 7.4.14, we have




Figure 11.2: Part of the counter graph $\mathbb{C}(\lambda)$ associated with the game of Figure 4.1

shown that the play

$$\rho = (v_0, \emptyset)(v_1, \emptyset)(v_6, \emptyset)(v_7, \emptyset)((v_2, \Pi)(v_0, \Pi)(v_1, \Pi)(v_6, \Pi)(v_7, \Pi))^\omega$$

was λ -consistent. Let us show that there is a corresponding infinite path π that starts in $(v_0, \emptyset)^C = (v_0, \emptyset, (+\infty, 4))$ in $\mathbb{C}(\lambda)$. Following Definition 11.2.2, we see that in $\mathbb{C}(\lambda)$, there exists an edge between $(v_0, \emptyset, (+\infty, 4))$ and $(v_1, \emptyset, (3, 3))$ and that

$$\pi = (v_0, \emptyset)^C(v_1, \emptyset, (3, 3))(v_6, \emptyset, (2, 2))(v_7, \emptyset, (1, 1))\pi'^\omega$$

with

$$\pi' = (v_2, \Pi, (0, 0))(v_0, \Pi, (0, 0))(v_1, \Pi, (0, 0))(v_6, \Pi, (0, 0))(v_7, \Pi, (0, 0)).$$

Come back now to the play $\rho' = ((v_0, \emptyset)(v_4, \emptyset))^\omega$ described in Example 7.4.14, which is not λ -consistent. From $(v_0, \emptyset)^C = (v_0, \emptyset, (+\infty, 4))$, there is an edge to $(v_4, \emptyset, (+\infty, 3))$, then to $(v_0, \emptyset, (+\infty, 2))$ and to $(v_4, \emptyset, (+\infty, 1))$. For the latter vertex, there is no outgoing edge back to $(v_0, \emptyset, (+\infty, 0))$ because $2 \notin I(v_0, \emptyset)$. Therefore there is no infinite path starting in $(v_0, \emptyset)^C$ in $\mathbb{C}(\lambda)$ that corresponds to ρ' .

There exists a correspondence between λ -consistent plays in \mathcal{X} and *infinite* paths from starting vertices in $\mathbb{C}(\lambda)$ in the following way. On one hand, every play ρ in \mathcal{X} that is not λ -consistent does not appear in the counter graph: the first constraint regarding λ that is violated along ρ is reflected by a vertex in $\mathbb{C}(\lambda)$ with a counter value getting to 1 and no outgoing edges. On the other hand, λ -consistent plays in \mathcal{X} have a corresponding infinite path in the counter graph $\mathbb{C}(\lambda)$. We call *valid paths* the infinite paths of $\mathbb{C}(\lambda)$. This correspondence is formalized in the two following lemmas:

Lemma 11.2.5. *Let $\rho = \rho_0\rho_1\dots$ be a λ -consistent play in $\text{Plays}_X(v)$. Then there exists an associated infinite path $\pi = \pi_0\pi_1\dots$ in $\mathbb{C}(\lambda)$ such that:*

- $\pi_0 = v^C$,
- ρ is the projection of π on V^X , that is, π_n is of the form $(v', (c'_i)_{i \in \Pi})$ with $v' = \rho_n$, for every $n \in \mathbb{N}$.

Lemma 11.2.6. *Let $v^C = (v, (c_i)_{i \in \Pi})$ be a starting vertex in $\text{SV}(\lambda)$. Let $\pi = \pi_0\pi_1\dots$ be an infinite path in $\mathbb{C}(\lambda)$ such that $\pi_0 = v^C$. Then there exists a corresponding play $\rho = \rho_0\rho_1\dots$ in \mathcal{X} such that:*

- ρ is λ -consistent,
- ρ is the projection of π on V^X , that is, $\rho_n = v'$ with $\pi_n = (v', (c'_i)_{i \in \Pi})$, for every $n \in \mathbb{N}$.

Since the edge relation E^C in the counter graph respects the edge relation E^X in the extended game, the region decomposition of path in \mathcal{X} given in Lemma 11.1.1 can also be applied to a path in $\mathbb{C}(\lambda)$. We will often use such path region decompositions in the proofs of this section.

In order to prove the PSPACE membership for the constrained existence problem, we need to show that the counter graph $\mathbb{C}(\lambda^*)$, with λ^* the fixpoint function computed by Algorithm 4, has an exponential size. To this end, as the size of $|\mathbb{C}(\lambda^*)|$ of $\mathbb{C}(\lambda^*)$ is equal to $|V| \cdot 2^{|\Pi|} \cdot (K + 2)^{|\Pi|}$ with $K = \text{mR}(\lambda^*)$ defined in Definition 11.2.1, it is enough to show an exponential upper bound on K . We proceed in two steps: First, with the next proposition, given a labeling function λ and its restriction $\lambda_{\geq \ell}$ to $V^{\geq \ell}$, we exhibit a bound on the supremum of the cost of λ -consistent plays for each player, in terms of the maximal finite range $\text{mR}(\lambda_{\geq \ell})$. Second, we consider the actual sequence of functions $(\lambda^k)_{k \in \mathbb{N}}$ defined in Definitions 11.1.3 and 11.1.5, as implemented by Algorithm 4. With Theorem 11.2.8, we show that $\text{mR}(\lambda^*)$ is bounded by an exponential in the size of the original game \mathcal{G} . The proof is by induction on k and uses Proposition 11.2.7.

Proposition 11.2.7 (Bound on finite supremum). *Let λ be a labeling function. Let $v \in V^X$ such that $I(v) = J_\ell$ with $\ell \in \{1, \dots, N\}$. Let $c \in \mathbb{N} \cup \{+\infty\}$ be such that $\sup \{\text{Cost}_i(\rho) \mid \rho \in \Lambda(v)\} = c$. If $c < +\infty$, then*

$$c \leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \sum_{r=|J_\ell|+1}^{|\Pi|} |V| + 2 \cdot \max_{\substack{J_j > J_\ell \\ |J_j|=r}} \text{mR}(\lambda_j).$$

Moreover, in both cases $c = +\infty$ and $c < +\infty$, there exists a valid path π in $\mathbb{C}(\lambda)$ starting in v^C that is a lasso hg^ω with the length of hg bounded by $2 \cdot |\mathbb{C}(\lambda)|$ and such that its corresponding play ρ in \mathcal{X} belongs to $\Lambda(v)$ and has its cost $\text{Cost}_i(\rho)$ equal to c .

We now come back to the labeling functions $(\lambda^k)_{k \in \mathbb{N}}$ computed by Algorithm 4. Recall that this algorithm works in a bottom-up manner (see Proposition 11.1.7): it first computes the local fixpoint $\lambda^{k_N^*}$ on region V^{J_N} , then the local fixpoint $\lambda^{k_{N-1}^*}$ on $V^{\geq J_{N-1}}$, \dots , until finally computing the global fixpoint λ^* on $V^{\geq J_1} = V^X$. Recall also that when the algorithm computes the local fixpoint in the arena $X^{\geq J_n}$, the values of $\lambda^k(v)$ may only change in the region V^{J_n} (Lemma 11.1.8). We are now ready to show an exponential bound (in the size of \mathcal{G}) on the maximal finite ranges $\text{mR}(\lambda_\ell^{k_\ell^*})$ for each region X^{J_ℓ} and $\text{mR}(\lambda_{\geq \ell}^{k_\ell^*})$ for each arena $X^{\geq J_\ell}$. With $\ell = 1$, we get that $\text{mR}(\lambda^*)$ is of exponential size, and thus also the size of the counter graph.

Theorem 11.2.8 (Bound on maximal finite range). *For each $\ell \in \{1, \dots, N\}$, we have*

$$\text{mR}(\lambda_\ell^{k_\ell^*}) \leq \mathcal{O}(|V|^{(|V|+3) \cdot (|\Pi \setminus J_\ell|+2)})$$

and also:

$$\text{mR}(\lambda_{\geq \ell}^{k_\ell^*}) \leq \mathcal{O}(|V|^{(|V|+3) \cdot (|\Pi|+2)}).$$

In particular for the global fixpoint λ^ we have*

$$\text{mR}(\lambda^*) \leq \mathcal{O}(|V|^{(|V|+3) \cdot (|\Pi|+2)}).$$

The next corollary is a direct consequence of Theorem 11.2.8.

Corollary 11.2.9. *The counter graph $\mathbb{C}(\lambda^*)$ has a size $|\mathbb{C}(\lambda^*)| = |V| \cdot 2^{|\Pi|} \cdot (\text{mR}(\lambda^*) + 2)^{|\Pi|}$ that is exponential in the size of the game \mathcal{G} .*

We conclude this section with two other corollaries that will be useful in the next section.

Corollary 11.2.10. *For every $k \in \mathbb{N}$ and region V^{J_ℓ} , we have*

$$\text{mR}(\lambda_\ell^k) \leq \mathcal{O}\left(|V|^{(|V|+3)(|\Pi|+2)}\right)$$

and also

$$\text{mR}(\lambda_{\geq \ell}^k) \leq \mathcal{O}\left(|V|^{(|V|+3)(|\Pi|+2)}\right).$$

Corollary 11.2.11. *Let $v \in V^X$ with $I(v) = J_\ell$ with $\ell \in \{1, \dots, N-1\}$. Let $k \in \mathbb{N}$. Suppose there exists $c \in \mathbb{N}$ such that $\sup \{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v)\} = c$. Then, the following holds:*

$$c \leq \mathcal{O}\left(|V|^{(|V|+3)(|\Pi|+2)}\right).$$

11.3 PSPACE-completeness

In this section, we prove Theorem 11.0.1. We first prove that the constrained existence problem of SPEs in quantitative Reachability games is in PSPACE and then that it is PSPACE-hard.

11.3.1 PSPACE-membership

The purpose of this section is to prove that determining if, given a reachability game (\mathcal{G}, v_0) and two thresholds $x, y \in (\mathbb{N} \cup \{+\infty\})^{|\Pi|}$, there exists an SPE σ in this game such that for all $i \in \Pi$, $x_i \leq \text{Cost}_i(\langle \sigma \rangle_{v_0}) \leq y_i$ can be done in PSPACE.

Proposition 11.3.1. *The constrained existence problem of SPEs in quantitative Reachability games is in PSPACE.*

Let us provide a high level sketch of the proof of our PSPACE procedure for this constrained existence problem. Thanks to Theorem 11.1.9, solving the constrained existence problem for a given game (\mathcal{G}, v_0) reduces in finding a λ^* -consistent play ρ in (\mathcal{X}, x_0) satisfying the constraints. By Lemmas 11.2.5 and 11.2.6, the latter problem reduces in finding a valid path π in the counter graph $\mathbb{C}(\lambda^*)$ that satisfies the constraints. We will see that it suffices to decide the existence of such a valid path that is a lasso hg^ω . As $\mathbb{C}(\lambda^*)$ is exponential in the size of the input \mathcal{G} (by Corollary 11.2.9), classical arguments using Savitch’s Theorem can thus be used to prove the PSPACE membership. Nevertheless, the detailed proof is more intricate for two reasons. The first reason is that the counter graph is constructed from the labeling function λ^* . We thus also have to prove that λ^* can be computed in PSPACE. The second reason is that, a priori, although we know that the counter graph is of exponential size, we do not know explicitly its size. This is problematic when using classical NPSpace algorithms that guess, vertex by vertex, some finite path in a graph of exponential size, where a counter bounded by the size of the graph is needed to guarantee the termination of the procedure. In order to overcome this, we also need a PSPACE procedure to obtain the actual size of $\mathbb{C}(\lambda^*)$. Recall that the size $|\mathbb{C}(\lambda^*)|$ is equal to $|V| \cdot 2^{|\Pi|} \cdot (K + 2)^{|\Pi|}$ where $K = \text{mR}(\lambda^*)$. Hence to compute the size of the counter graph, we have to compute the actual value of $\text{mR}(\lambda^*)$.

The PSPACE procedure to compute λ^* and $\text{mR}(\lambda^*)$ works by induction on k , the steps in the computation of the labeling function λ^* . Moreover, it exploits the structural evolution of the local fixpoints formalized in Proposition 11.1.7 and Lemma 11.1.8. Indeed recall that these local fixpoints are computed region by region, from X^{J_N} to X^{J_1} , and that if X^{J_ℓ} is the currently treated region, then the values of $\lambda^{k+1}(v)$ are computed from $\lambda^k(v)$ for all $v \in V^{J_\ell}$ (those values remain unchanged for all v outside of V^{J_ℓ}). For the moment, suppose that we have at our disposal a PSPACE procedure to compute $\{\lambda^k(v) \mid v \in V^{J_\ell}\}$ and the maximal finite range $\text{mR}(\lambda_{\geq \ell}^k)$:

Proposition 11.3.2. *Given an initialized reachability game (\mathcal{G}, v_0) , for all $k \in \mathbb{N}$ and for all J_ℓ , $\ell \in \{1, \dots, N\}$, the set $\{\lambda^k(v) \mid v \in V^{J_\ell}\}$ and the maximal finite range $\text{mR}(\lambda_{\geq \ell}^k)$ can be computed in PSPACE.*

Let us prove Proposition 11.3.1. The proof of Proposition 11.3.2 will be given just after.

Proof of Proposition 11.3.1. Let (\mathcal{G}, v_0) be an initialized reachability game and let $x, y \in (\mathbb{N} \cup \{+\infty\})^{|\Pi|}$ be two thresholds. Let (\mathcal{X}, x_0) be the extended game of (\mathcal{G}, v_0) , $\mathbb{C}(\lambda^*)$ the counter graph constructed from the labeling function λ^* and its maximal finite range $\text{mR}(\lambda^*)$.

We first prove that there exists an SPE in (\mathcal{G}, v_0) such that its outcome ρ satisfies the constraints $x_i \leq \text{Cost}_i(\rho) \leq y_i$ for all $i \in \Pi$, if and only if, there exists a valid path in $\mathbb{C}(\lambda^*)$ starting from the starting vertex $x_0^{\mathcal{C}}$ associated with x_0 such that it is a lasso hg^ω with the length $|hg|$ bounded by $d + 2 \cdot |\mathbb{C}(\lambda^*)|$ with

$$d := \max\{x_i \mid x_i < +\infty\}$$

and such that it also satisfies these constraints.

We already know by Theorem 11.1.9 and Lemmas 11.2.5-11.2.6 that the existence of an SPE outcome ρ in \mathcal{G} satisfying the constraints is equivalent to the existence of a valid path π in $\mathbb{C}(\lambda^*)$ satisfying these constraints. It remains to show that the latter path can be chosen as a lasso hg^ω with the announced length of hg . This lasso is constructed as follows. Consider the suffix $\pi_{\geq d}$ of π and its region decomposition $\pi_{\geq d}[m]\pi_{\geq d}[m+1] \dots \pi_{\geq d}[n]$. For all $\ell \in \{m, \dots, n-1\}$, we remove all the cycles in section $\pi_{\geq d}[\ell]$ to get a simple path π'_ℓ , and from the last section $\pi_{\geq d}[n]$, we derive the infinite path π'_n formed of the first vertices of $\pi_{\geq d}[n]$ until a cycle is reached and then repeated forever. Notice that there is no other cycle in $\pi'_m \pi'_{m+1} \dots \pi'_n$ by the I -monotonicity property (4.1). The required lasso hg^ω is equal to the concatenation of the prefix $\pi_{\leq d}$ with the modified suffix $\pi'_m \pi'_{m+1} \dots \pi'_n$. By construction $\pi'_m \dots \pi'_n$ is itself a lasso $h'g'^\omega$ with $|h'g'|$ bounded by $2 \cdot |\mathbb{C}(\lambda^*)|$ and thus $|hg|$ is bounded by $d + 2 \cdot |\mathbb{C}(\lambda^*)|$. Moreover, for all i , π visits the target set of player i if and only if hg^ω visits this set (maybe earlier if the

visit is inside $\pi'_m \dots \pi'_n$). By definition of d , the constraints imposed by x_i, y_i , $i \in \Pi$, are simultaneously satisfied by π and hg^ω .

Let us now show how to get a PSPACE procedure for the constrained existence problem. As just explained, we have to guess a lasso $\pi = hg^\omega$ in $\mathbb{C}(\lambda^*)$ that starts in x_0^C , satisfies the constraints, and such that $|hg|$ is bounded by

$$L = d + 2 \cdot |\mathbb{C}(\lambda^*)|.$$

We cannot guess π entirely and we have to proceed region by region. Suppose that $I(x_0) = J_m$ for some $m \in \{1, \dots, N\}$, and consider the region decomposition $\pi[m]\pi[m+1] \dots \pi[n]$ of π , where some sections $\pi[\ell]$ may be empty.

We guess successively the sections $\pi[m]$, $\pi[m+1]$ and so on. To guess $\pi[\ell]$ with $\ell \in \{m, \dots, n\}$, assuming it is not empty, we guess one by one its vertices that all belong to the same region V^{J_ℓ} . To guess such a vertex $(v', (c'_i)_{i \in \Pi})$, we only have to keep its predecessor $(v, (c_i)_{i \in \Pi})$ in memory and to know the value $\lambda^*(v')$ (in a way to compute each c'_i from c_i). So, we need to know $\{\lambda^*(v) \mid v \in V^{J_\ell}\}$. By Proposition 11.3.2, we can compute this set in PSPACE. Once we move to another region in a way to guess the next section of π , we can forget this set and compute the new one. We also need to guess which vertex will be the first vertex of g .

Notice that any vertex of the counter graph can be encoded in polynomial size memory. Indeed it is composed of a vertex of V , a subset I of Π , and $|\Pi|$ counter values that belong to $\{0, \dots, \text{mR}(\lambda^*)\} \cup \{+\infty\}$ ($\text{mR}(\lambda^*)$ is at most exponential in the input by Theorem 11.2.8). Moreover the set $\{\lambda^*(v) \mid v \in V^{J_\ell}\}$ can also be encoded in polynomial size memory since it is composed of $|V|$ values that belong to $\{0, \dots, \text{mR}(\lambda^*)\} \cup \{+\infty\}$.

Recall that the length $|hg|$ for the guessed lasso $\pi = hg^\omega$ cannot exceed constant L . We thus have to compute and store L . The computation in PSPACE of L requires the computation of $|\mathbb{C}(\lambda^*)|$, and thus in particular the computation of $\text{mR}(\lambda^*)$. This is possible thanks to Proposition 11.3.2. And it follows by Corollary 11.2.9 that L can be stored in polynomial size memory. In addition to L , during the guessing of π , we also have a counter C_L to count the current length of π , and for each player $i \in \Pi$ a counter C_i

keeping track of the current cost of player i along π . As $C_L \leq L$ and $C_i \leq L$ for all i , all these counters can be also encoded in polynomial size memory. Finally, we stop guessing π when either its length exceeds L or when its currently guessed vertex is equal to the first vertex of cycle g for the second time. In the latter case, we check whether π satisfies the constraints, that is, $x_i \leq C_i \leq y_i$ for all i . This completes the proof that the given procedure works in PSPACE. \square

Let us now prove Proposition 11.3.2. We proceed by induction on the steps in the computation of the labeling function λ^* and, once a step k is fixed, we proceed region by region, beginning with the bottom region J_N and then proceeding bottom-up by following the total order $J_1 < \dots < J_N$. Let X^{J_ℓ} be a region, we aim at proving that the set $\{\lambda^{k+1}(v) \mid v \in V^{J_\ell}\}$ and the value $\text{mR}(\lambda_{\geq \ell}^{k+1})$ are both computable in PSPACE (from the previous step k). To this aim consider Proposition 11.1.7 and especially Lemma 11.1.8. Let k_ℓ^* (resp. $k_{\ell+1}^*$) be the step where the local fixpoint is reached for region X^{J_ℓ} (resp. $X^{J_{\ell+1}}$). Recall that $k_{\ell+1}^* < k_\ell^*$ and that when $k \leq k_{\ell+1}^*$ (resp. $k \geq k_\ell^*$), we have that $\lambda^{k+1}(v) = \lambda^k(v)$, for each $v \in V^{J_\ell}$. The tricky case is when $k_{\ell+1}^* < k < k_\ell^*$. In the latter case, the computation of $\lambda^{k+1}(v)$ from $\lambda^k(v)$, for all $v \in V^{J_n}$, relies on the computation of the maximal cost, for the player who owns vertex v , of the plays of $\Lambda^k(v')$, with $v' \in \text{Succ}(v)$ (see Definition 11.1.5). This will be possible with the same approach as in the proof of Proposition 11.3.1: to guess a lasso in the counter graph $\mathbb{C}(\lambda^k)$ that realizes this maximal cost.

Proof of Proposition 11.3.2. We proceed by induction on $k \in \mathbb{N}$.

Base case. For $k = 0$ we have to prove that for all $J_\ell, \ell \in \{1, \dots, N\}$, the set $\{\lambda^0(v) \mid v \in V^{J_\ell}\}$ and the value $\text{mR}(\lambda_{\geq \ell}^0)$ can be both computed in PSPACE. Given J_ℓ , thanks to Definition 11.1.3, we have that either $\lambda^0(v) = 0$ if $v \in V_i^X$ and $i \in I(v) = J_\ell$ or $\lambda^0(v) = +\infty$ otherwise. Thus $\text{mR}(\lambda_{\geq \ell}^0) = 0$. So, we clearly have a PSPACE procedure in this case.

General case. Now, assume that for all $n, n \in \{0, \dots, k\}$, and for all $J_\ell, \ell \in \{1, \dots, N\}$, the set $\{\lambda^n(v) \mid v \in V^{J_\ell}\}$ and the value $\text{mR}(\lambda_{\geq \ell}^n)$ can be computed in PSPACE. Let us prove that it remains true for $n = k + 1$, that

is:

$$\{\lambda^{k+1}(v) \mid v \in V^{J_\ell}\} \text{ and } \text{mR}(\lambda_{\geq \ell}^{k+1}) \text{ can be computed in PSPACE} \quad (11.3)$$

We proceed by induction on the region X^{J_ℓ} (we thus use a double induction, one on the computation steps and the other one on the regions).

If $J_\ell = J_N$, then it is a bottom region. Then for all $v \in V^{J_N}$, $\lambda^{k+1}(v) = \lambda^0(v)$ and so $\{\lambda^{k+1}(v) \mid v \in V^{J_N}\} = \{\lambda^0(v) \mid v \in V^{J_N}\}$ (local fixpoint $k_N^* = 0$ in Proposition 11.1.7). By induction hypothesis we know that this latter set can be computed in PSPACE. Moreover, we have that $\text{mR}(\lambda_{\geq N}^{k+1}) = \text{mR}(\lambda_N^{k+1}) = 0$. Therefore Assertion (11.3) holds.

Now, let J_ℓ be a region different from J_N . If it is a bottom region, we have $\{\lambda^{k+1}(v) \mid v \in V^{J_\ell}\} = \{\lambda^0(v) \mid v \in V^{J_\ell}\}$ as for J_N . Moreover $\text{mR}(\lambda_\ell^{k+1}) = \text{mR}(\lambda_\ell^0) = 0$ and then $\text{mR}(\lambda_{\geq \ell}^{k+1}) = \max\{\text{mR}(\lambda_\ell^{k+1}), \text{mR}(\lambda_{\geq \ell+1}^{k+1})\} = \text{mR}(\lambda_{\geq \ell+1}^{k+1})$. Thus by induction hypothesis, $\{\lambda^{k+1}(v) \mid v \in V^{J_\ell}\}$ and $\text{mR}(\lambda_{\geq \ell}^{k+1})$ can be computed in PSPACE and Assertion (11.3) holds.

Let us now suppose that J_ℓ is not a bottom region. We first recall Proposition 11.1.7 that states that Algorithm 4 reaches a local fixpoint for each region. Let k_ℓ^* (resp. $k_{\ell+1}^*$), the step after which the region X^{J_ℓ} (resp. $X^{J_{\ell+1}}$) has reached its local fixpoint. Recall that $k_{\ell+1}^* < k_\ell^*$. Let us now consider the three cases of Lemma 11.1.8.

- If $k \leq k_{\ell+1}^*$, then by Lemma 11.1.8, the region $X^{J_{\ell+1}}$ has not reached its local fixpoint yet. So, it implies that the labeling of the vertices of V^{J_ℓ} has not change since initialization. More formally, for all $v \in V^{J_\ell}$, $\lambda^{k+1}(v) = \lambda^0(v)$. Thus $\text{mR}(\lambda_\ell^{k+1}) = 0$ and then $\text{mR}(\lambda_{\geq \ell}^{k+1}) = \text{mR}(\lambda_{\geq \ell+1}^{k+1})$. So by induction hypothesis, both $\{\lambda^{k+1}(v) \mid v \in V^{J_\ell}\}$ and $\text{mR}(\lambda_{\geq \ell}^{k+1})$ can be computed in PSPACE showing (11.3).
- If $k > k_\ell^*$, then by Lemma 11.1.8, the local fixpoint of region X^{J_ℓ} is reached, that is, $\{\lambda^{k+1}(v) \mid v \in V^{J_\ell}\} = \{\lambda^k(v) \mid v \in V^{J_\ell}\}$ and $\text{mR}(\lambda_{\geq \ell}^{k+1}) = \text{mR}(\lambda_{\geq \ell}^k)$ and (11.3) holds by induction hypothesis.

(Notice the little difference with the inequalities given in Lemma 11.1.8: we here consider case $k > k_\ell^*$ instead of case $k \geq k_\ell^*$ of Lemma 11.1.8.)

Indeed when $k = k_\ell^*$, we still need to compute λ^{k+1} to realize that the fixpoint is effectively reached. This is thus postponed in the next case.)

- It remains to consider the case $k_{\ell+1}^* < k \leq k_\ell^*$, which is the most difficult one. In this case, either the values of $\lambda^k(v)$ and $\lambda^{k+1}(v)$ differ for some $v \in V^{J_\ell}$, or $k = k_\ell^*$ and we realize that the local fixpoint is effectively reached on X^{J_ℓ} .

Let us first show that the set $\{\lambda^{k+1}(v) \mid v \in V^{J_\ell}\}$ can be computed in PSPACE. Given $v \in V^{J_\ell}$, if $v \in V_i^X$, then by Definition 11.1.5, $\lambda^{k+1}(v)$ is either equal to 0 (if $i \in J_\ell$) or it is computed from the values $\sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v')\}$, $v' \in \text{Succ}(v)$. Thus we have to show that each value $\sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v')\}$ can be computed in PSPACE.

To this aim, we use Proposition 11.2.7 stating that if

$$\sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v')\} = c,$$

then there exists a valid path $\pi = hg^\omega$ in $\mathbb{C}(\lambda^k)$ starting in v'^C that is a lasso with $|hg|$ bounded by $2 \cdot |\mathbb{C}(\lambda^k)|$ and such that its corresponding play ρ in \mathcal{X} belongs to $\Lambda^k(v)$ and has its cost $\text{Cost}_i(\rho)$ equal to c . Notice that we can restrict λ^k to $\lambda_{\geq \ell}^k$ since any path beginning in V^{J_ℓ} only visits vertices of $V^{\geq J_\ell}$. We thus work in the counter graph $\mathbb{C}(\lambda_{\geq \ell}^k)$ restricted to $V^{\geq \ell}$ (see Remark B.1.5). We will guess such a lasso π as done in the proof of Proposition 11.3.1. More precisely, as the value of c is unknown, we will first test in PSPACE whether there exists such a lasso π with cost $c = +\infty$. If yes, we are done, otherwise by considering increasing values $d \in \mathbb{N}$, we will test in PSPACE whether there exists a lasso π with cost $\geq d$. The last value d for which the answer is yes is the required cost c .

- Let us detail the case $c = +\infty$. Similarly to the proof of Proposition 11.3.1, we guess the lasso $\pi = hg^\omega$ by on one hand guessing the first vertex of its cycle g and on the other hand guessing the sections $\pi[\ell], \pi[\ell + 1], \dots$ of π , one by one, from the starting vertex $v'^C \in V^{J_\ell}$. Given $m \geq \ell$, section $\pi[m]$ is guessed vertex by

vertex, where each vertex belongs to V^{J_m} . By induction hypothesis, we can compute the set $\{\lambda^k(u) \mid u \in V^{J_m}\}$ and the value $\text{mR}(\lambda_{\geq m}^k)$ in PSPACE, and by Corollary 11.2.10 the previous set and each vertex of the counter graph $\mathbb{C}(\lambda_{\geq \ell}^k)$ can be encoded in polynomial size memory. Once $\pi[m]$ is guessed we can forget the set $\{\lambda^k(u) \mid u \in V^{J_m}\}$ and guess the next section of π .

Additionally, to check that the length of hg does not exceed the constant $L = 2 \cdot |\mathbb{C}(\lambda_{\geq \ell}^k)|$, we have to compute $\text{mR}(\lambda_{\geq \ell}^k)$ that can be done in PSPACE by induction hypothesis. Thus we can also compute L in PSPACE and by Corollary 11.2.10 we can encode it in polynomial size memory. We also store a counter $C_L \leq L$ (which is initialized to 0 and incremented by 1 each time we guess a new vertex of π) and a boolean C_i (which is equal to 0 as long as player i does not visit his target set and is equal to 1 after this visit).

We stop if either if $C_L > L$ or if we have found the lasso. In the latter case we check whether $C_i = 0$ or not.

- Let us now proceed to the case $c < +\infty$. We check whether there exists a lasso $\pi = hg^\omega$ with $|hg|$ bounded by $L = 2 \cdot |\mathbb{C}(\lambda_{\geq \ell}^k)|$ with cost (for player i) $\geq d$, for values $d = 0, d = 1, \dots$, until the answer is no. The last value d for which the answer is yes is equal to c .

We know by Corollary 11.2.11, that the size of c cannot exceed an exponential in the size of the input. Thus each tested value d can be encoded in polynomial size memory.

To check in PSPACE the existence of such a lasso with length of hg bounded by L with cost $\geq d$, we proceed exactly as for the case $c = +\infty$ except that instead of the boolean C_i we keep a counter $C_i \leq L$ that keeps track of the cost of player i for the lasso that we are guessing.

Notice that the depth of the recursion of the procedure is at most $|\Pi|$ by the I -monotonicity property (4.1) (any path crosses at most $|\Pi|$ regions). And at each recursion level, only a polynomial size information

is stored. This concludes the proof that the set $\{\lambda^{k+1}(v) \mid v \in V^{J_\ell}\}$ can be computed in PSPACE.

To conclude the case $k_{\ell+1}^* < k \leq k_\ell^*$, it remains to prove that the value $\text{mR}(\lambda_{\geq \ell}^{k+1})$ can also be computed in PSPACE. Clearly we can compute $\text{mR}(\lambda_\ell^{k+1})$ in PSPACE as we now have $\{\lambda^{k+1}(v) \mid v \in V^{J_\ell}\}$ in memory and we can compute $\text{mR}(\lambda_{\geq \ell+1}^{k+1})$ by induction hypothesis. Notice that both values can be encoded in polynomial memory size by Corollary 11.2.10. So, as $\text{mR}(\lambda_{\geq \ell}^{k+1}) = \max\{\text{mR}(\lambda_\ell^{k+1}), \text{mR}(\lambda_{\geq \ell+1}^{k+1})\}$, we can compute $\text{mR}(\lambda_{\geq \ell}^{k+1})$ in PSPACE and Assertion (11.3) is proved. \square

11.3.2 PSPACE-hardness

We now prove that the constrained existence problem is PSPACE-hard for quantitative Reachability games.

Proposition 11.3.3. *The constrained existence problem of SPEs in multi-player quantitative Reachability games is PSPACE-hard.*

The proof of this proposition is based on a polynomial reduction from the QBF problem which is PSPACE-complete. It is close to the proof given in Section 10.2 for the PSPACE-hardness of the constrained existence problem of weak SPEs in qualitative Reachability games. The main difference is to manipulate costs instead of considering qualitative reachability.

The QBF problem is to decide whether a fully quantified Boolean formula ψ is true. The formula ψ can be assumed to be in prenex normal form $Q_1x_1Q_2x_2 \dots Q_mx_m \phi(X)$ such that the quantifiers are alternating existential and universal quantifiers ($Q_1 = \exists, Q_2 = \forall, Q_3 = \exists, \dots$), $X = \{x_1, x_2, \dots, x_m\}$ is the set of quantified variables, and $\phi(X) = C_1 \wedge \dots \wedge C_n$ is an unquantified Boolean formula over X equal to the conjunction of the clauses C_1, \dots, C_n .

Such a formula ψ is true if there exists a value of x_1 such that for all values of x_2 , there exists a value of $x_3 \dots$, such that the resulting valuation ν of all variables of X evaluates $\phi(X)$ to true. Formally, for each odd

(resp. even) k , $1 \leq k \leq m$, let us denote by $f_k : \{0,1\}^{k-1} \rightarrow \{0,1\}$ (resp. $g_k : \{0,1\}^{k-1} \rightarrow \{0,1\}$) a valuation of variable x_k given a valuation of previous variables x_1, \dots, x_{k-1} ⁵. Given these sequences $f = f_1, f_3, \dots$ and $g = g_2, g_4, \dots$, let us denote by $\nu = \nu_{(f,g)}$ the valuation of all variables of X such that $\nu(x_1) = f_1$, $\nu(x_2) = g_2(\nu(x_1))$, $\nu(x_3) = f_3(\nu(x_1)\nu(x_2))$, \dots . Then

$$\psi = Q_1x_1Q_2x_2\dots Q_mx_m\phi(X) \text{ is true} \\ \text{if and only if}$$

there exist $f = f_1, f_3, \dots$ such that for all $g = g_2, g_4, \dots$, the valuation $\nu_{f,g}$ evaluates $\phi(X)$ to true.

Let us first detail a polynomial reduction from the QBF problem to the constraint problem for quantitative reachability games.

Let $\psi = Q_1x_1Q_2x_2\dots Q_mx_m\phi(X)$ with $\phi(X) = C_1 \wedge \dots \wedge C_n$ be a fully quantified Boolean formula in prenex normal form. We build the following quantitative Reachability game $\mathcal{G}_\psi = (A, (\text{Cost}_i)_{i \in \Pi}, (F_i)_{i \in \Pi})$ with $A = (\Pi, V, (V_i)_{i \in \Pi}, E)$ (see Figure 11.3):

- the set V of vertices:
 - for each variable $x_k \in X$ under quantifier Q_k , there exist vertices $x_k, \neg x_k$ and q_k ;
 - for each clause C_k , there exist vertices c_k and t_k ;
 - there exists an additional vertex t_{n+1} ;
- the set E of edges:
 - from each vertex q_k there exist an edge to x_k and an edge to $\neg x_k$;
 - from each vertex x_k and $\neg x_k$, there exists an edge to q_{k+1} , except for $k = m$ where this edge is to c_1 ;
 - from each vertex c_k , there exist an edge to t_k and an edge to c_{k+1} , except for $k = n$ where there exist an edge to t_n and an edge to t_{n+1} ;
 - there exists a loop on each t_k ;

⁵Notice that $f_1 : \emptyset \rightarrow \{0,1\}$.

- the set Π of $n + 2$ players:
 - each player i , $1 \leq i \leq n$, owns vertex c_i ;
 - player $n + 1$ (resp. $n + 2$) is the player who owns the vertices q_i for each existential (resp. universal) quantifier Q_i ;
 - as all other vertices have only one outgoing edge, it does not matter which player owns them;
- each function Cost_i is associated with the target set F_i defined as follows:
 - for all i , $1 \leq i \leq n$, $F_i = \{\ell \in V \mid \ell \text{ is a literal of clause } C_i\} \cup \{t_i\}$;
 - $F_{n+1} = \{t_{n+1}\}$;
 - $F_{n+2} = \{t_1, \dots, t_n\}$.

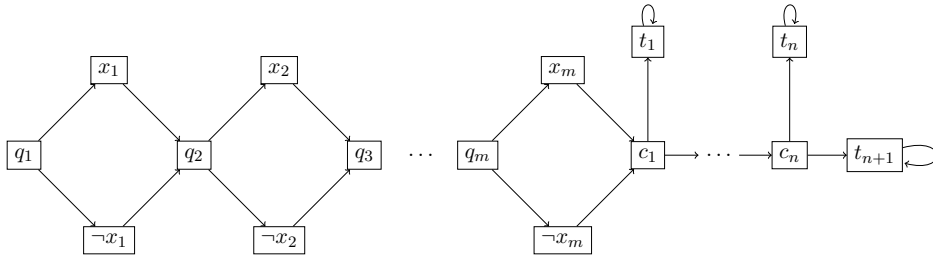


Figure 11.3: Reduction from the formula ψ to the quantitative Reachability game \mathcal{G}_ψ

Remark 11.3.4. (1) Notice that a sequence f of functions $f_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k odd, $1 \leq k \leq m$, as presented above, can be translated into a strategy σ_{n+1} of player $n + 1$ in the initialized game (\mathcal{G}_ψ, q_1) , and conversely. Similarly, a sequence g of functions $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$ is nothing else than a strategy σ_{n+2} of player $n + 2$. (2) Notice also that if ρ is a play in (\mathcal{G}_ψ, q_1) , then $\text{Cost}_{n+1}(\rho) < +\infty$ if and only if $\text{Cost}_{n+2}(\rho) = +\infty$. Moreover, suppose that ρ visits t_{n+1} , then for all i , $1 \leq i \leq n$, $\text{Cost}_i(\rho) \leq 2 \cdot m$ if and only if for all i , $1 \leq i \leq n$, ρ visits a vertex that is a literal of C_i , and that is the case if and only if there is a valuation of all variables of X that evaluates $\phi(X)$ to true.

Proof of Proposition 11.3.3. Consider the game \mathcal{G}_ψ and the bound $y = (2 \cdot m, \dots, 2 \cdot m, 2 \cdot m + n, +\infty)$. Both can be constructed from ψ in polynomial time. Let us now show that ψ is true if and only if there exists an SPE in (\mathcal{G}_ψ, q_1) with $\text{cost} \leq y$.

(\Rightarrow) Suppose that ψ is true. Then there exists a sequence f of functions $f_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k odd, $1 \leq k \leq m$, such that for all sequences g of functions $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$, the valuation $\nu_{f,g}$ evaluates $\phi(X)$ to true. We define a strategy profile σ as follows:

- for player $n + 1$, his strategy σ_{n+1} is the strategy corresponding to the sequence f (by Remark 11.3.4);
- for player $n + 2$, his strategy is an arbitrary strategy σ_{n+2} ; we denote by g the corresponding sequence $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$ (by Remark 11.3.4);
- for each player i , $1 \leq i \leq n$,
 - if $hv \in \text{Hist}_i(q_1)$ with $v = c_i$, is consistent with σ_{n+1} , then $\sigma_i(hv) = c_{i+1}$ if $i \neq n$ and t_{n+1} otherwise
 - else $\sigma_i(hv) = t_i$.

Let us first prove that the play $\rho = \langle \sigma \rangle_{q_1}$ has a $\text{cost} \leq y = (2 \cdot m, \dots, 2 \cdot m, 2 \cdot m + n, +\infty)$. By hypothesis, the valuation $\nu_{f,g}$ evaluates $\phi(X)$ to true, that is, it evaluates all clauses C_i to true. Hence by Remark 11.3.4, ρ visits a vertex of F_i for all i , $1 \leq i \leq n$, and by definition of σ , ρ eventually loops on t_{n+1} . It follows that $\text{Cost}_i(\rho) \leq 2 \cdot m$ for all i , $1 \leq i \leq n$, $\text{Cost}_{n+1}(\rho) \leq 2 \cdot m + n$, and $\text{Cost}_{n+2}(\rho) = +\infty$. Hence $\text{Cost}(\rho) \leq y$.

Let us now prove that σ is an SPE, that is, for each history $hv \in \text{Hist}(q_1)$, there is no one-shot deviating strategy in the subgame $(\mathcal{G}_{\psi|_h}, v)$ that is profitable to the player who owns vertex v (by Corollary 2.4.23). This is clearly true for all $v = t_i$, $1 \leq i \leq n + 1$, since t_i has only one outgoing edge. For the other vertices v , we study two cases:

- hv is consistent with σ_{n+1} : Notice that hv is maybe not consistent with σ_{n+2} , but with another arbitrary strategy σ'_{n+2} . Let g' be the sequence

corresponding to σ'_{n+2} by Remark 11.3.4. By hypothesis, the valuation $\nu_{f,g'}$ evaluates $\phi(X)$ to true. Hence as explained previously for $\langle \sigma \rangle_{q_1}$, the cost of play $\rho = h\langle \sigma|_h \rangle_v$ is such that $\text{Cost}_i(\rho) \leq 2 \cdot m$ for all i , $1 \leq i \leq n$, $\text{Cost}_{n+1}(\rho) \leq 2 \cdot m + n$, and $\text{Cost}_{n+2}(\rho) = +\infty$. If v belongs to player i , $1 \leq i \leq n$, this player has no incentive to deviate since he has already visited his target set along h and thus cannot decrease his cost. If v belongs to player $n+1$, a one-shot deviation will lead to a play eventually looping on t_1 by definition of σ , thus leading to a cost $+\infty$ which is not profitable for player $n+1$. Finally if v belongs to player $n+2$, a one-shot deviation will not decrease his cost by definition of σ (σ_{n+2} is arbitrary).

- hv is not consistent with σ_{n+1} : Suppose that $v = c_k$. Then by definition of σ , the play $h\langle \sigma|_h \rangle_v$ eventually loops on t_k leading to a cost $\leq 2 \cdot m + k$ for player k . In fact, if player k has already seen his target set along hv , using a one-shot deviation in the subgame $(\mathcal{G}_{\psi|_h}, v)$ leads to the same cost for him. Otherwise, it leads to a cost equal to $+\infty$: indeed, deviating here means going to the state c_{k+1} (or t_{n+1} if $k = n$, which leads to a cost of $+\infty$ for player n), and since hv is not consistent with σ_{n+1} , by definition of σ_{k+1} , player $k+1$ will choose to go to t_{k+1} . This player has thus no incentive to deviate.

Suppose that $v = q_k$. Then by definition of σ , the play $\rho = h\langle \sigma|_h \rangle_v$ eventually loops on t_1 . It follows that $\text{Cost}_{n+1}(\rho) = +\infty$ and $\text{Cost}_{n+2}(\rho) = 2 \cdot m + 1$. Due to the structure of the game graph, $2 \cdot m + 1$ is the smallest cost that player $n+2$ is able to obtain. So if $q_k \in V_{n+2}$, player $n+2$ has no incentive to deviate. And if $q_k \in V_{n+1}$, player $n+1$ could try to use a one-shot deviating strategy, however the resulting play still eventually loops on t_1 .

This proves that σ is an SPE and we already showed that its cost was bounded by y .

(\Leftarrow) Suppose that there exists an SPE σ in $(\mathcal{G}_{\psi}, q_1)$ with outcome ρ such that $\text{Cost}(\rho) \leq y$. In particular $\text{Cost}_{n+1}(\rho) < +\infty$. By Remark 11.3.4, it follows that $\text{Cost}_{n+2}(\rho) = +\infty$. We have to prove that ψ is true. To this end,

consider the sequence f of functions $f_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k odd, $1 \leq k \leq m$, that corresponds to strategy σ_{n+1} of player $n+1$ by Remark 11.3.4. Let us show that for all sequences g of functions $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$, the valuation $\nu_{f,g}$ evaluates $\phi(X)$ to true.

By contradiction assume that it is not the case for some sequence g' and consider the related strategy σ'_{n+2} of player $n+2$ by Remark 11.3.4. Notice that σ'_{n+2} is a finitely deviating strategy. Let us consider the outcome ρ' of the strategy profile $(\sigma'_{n+2}, \sigma_{-(n+2)})$ from q_1 . As $\text{Cost}_{n+2}(\rho) = +\infty$, we must have $\text{Cost}_{n+2}(\rho') = +\infty$, otherwise σ'_{n+2} is a profitable deviation for player $n+2$ whereas σ is an SPE. It follows that $\text{Cost}_{n+1}(\rho') < +\infty$ by Remark 11.3.4, that is, ρ' eventually loops on t_{n+1} .

Now recall that the valuation $\nu_{f,g'}$ evaluates $\phi(X)$ to false, which means that it evaluates some clause C_k of $\phi(X)$ to false. Consider the history $hc_k < \rho'$. As strategy σ'_{n+2} only acts on the left part of the underlying graph of \mathcal{G}_ψ , we have $\rho' = \langle \sigma'_{n+2}, \sigma_{-(n+2)} \rangle_{q_1} = h \langle \sigma_{\uparrow h} \rangle_{c_k}$. In the subgame $(\mathcal{G}_{\psi \uparrow h}, c_k)$, the outcome of $\sigma_{\uparrow h}$ gives a cost of $+\infty$ to player k because $\rho' = h \langle \sigma_{\uparrow h} \rangle_{c_k}$ does not visit t_k and $\nu_{f,g'}$ evaluates C_k to false. In this subgame, player k has thus a profitable one-shot deviation that consists to move to t_k . It follows that σ is not an SPE which is impossible. Then ψ is true. \square

CHAPTER 12

OTHER RELEVANT EQUILIBRIA IN REACHABILITY GAMES

In this chapter, based on [BBGT19], we focus on multiplayer Reachability games: qualitative or quantitative ones. In those particular settings, our purpose was to identify what is a relevant equilibrium. In the previous chapters, we have already started to answer this question by considering the constrained existence problem of SPEs in multiplayer qualitative and quantitative Reachability games (Chapters 10 - 11) that we proved PSPACE-complete in both cases. In this chapter, our aim is to pursue the study of relevant equilibria in quantitative and qualitative Reachability games by considering other kinds of relevant equilibria and the complexity of their related decision problems.

With the constrained existence problem, we impose both a lower and an upper bound on costs of equilibria. It is much more natural to only impose upper bounds and to try to obtain an equilibrium with the smallest possible costs. This is the reason why, we consider a variant of the constrained existence problem: *the upper threshold decision problem* (Problem 3).

The constrained existence problem and the upper threshold problem consider the players' costs purely individually. One may also aim at finding an equilibrium such that the *social welfare* of the players is optimized. That is (i) as many players as possible reach their target sets; (ii) the sum of the costs

of those players is minimized. We consider the *social welfare decision problem* (Problem 4).

Finally, when one consider the set of cost profiles which are realizable in a game, *i.e.*, there exists a play with this cost profile, some of them are *Pareto optimal*. Thus we want to decide if there exists an equilibrium such that its cost profile is pareto optimal in the set of all realizable cost profiles. We call this decision problem the *Pareto optimal decision problem* (Problem 5).

When considering equilibria synthesis, those strategies have to be implemented. Thus knowing the memory requirements of such equilibria and obtaining finite-memory ones may be crucial. This is the reason why, in addition to the study of the complexity classes, in case of a positive answer to any of the three decision problems, we prove that finite-memory strategies are sufficient.

Our results gathered with previous works are summarized in Table 12.1 for complexity results and in Table 12.2 for memory results.

Table 12.1: Complexity classes for Problems 3-5.

Complexity	Qualitative Reach.		Quantitative Reach.	
	NE	SPE	NE	SPE
Prob. 3	NP-c [CFGR16, Umm05]	PSPACE-c [Chap.10]	NP-c	PSPACE-c [Chap.11]
Prob. 4	NP-c	PSPACE-c	NP-c	PSPACE-c
Prob. 5	NP-h/ Σ_2^P	PSPACE-c	NP-h/ Σ_2^P	PSPACE-c

Table 12.2: Memory results for Problems 3-5.

Memory	Qualitative Reach.		Quantitative Reach.	
	NE	SPE	NE	SPE
Prob. 3	Polynomial [CFGR16]	Exponential [Chap.10]	Polynomial	Exponential
Prob. 4	Polynomial	Exponential	Polynomial	Exponential
Prob. 5	Polynomial	Exponential	Polynomial	Exponential

In this chapter we mainly focus on the results for multiplayer quantitative Reachability games (Section 12.1 to Section 12.3) and then briefly explain results for qualitative Reachability games (Section 12.4). More precisely, in Section 12.1, we explain the problems of our interest in the quantitative setting.

In Section 12.2, we show that for particular families of Reachability games and requirements, there is no need to solve the related decision problems because they always have a positive answer in this case. In Section 12.3, we state our complexity and memory results in the quantitative setting and provide the material necessary to prove them. In Section 12.4, we briefly discuss the qualitative setting.

12.1 Studied problems

In this section, we define the decision problems of our interest in the quantitative setting. Let us first recall the concepts of social welfare and Pareto optimality. Let (\mathcal{G}, v_0) be an initialized multiplayer quantitative Reachability game with $\mathcal{G} = (A, (\text{Cost}_i)_{i \in \Pi}, (F_i)_{i \in \Pi})$. Given $\rho = \rho_0 \rho_1 \dots \in \text{Plays}(v_0)$, let us recall that we denote by $\text{Visit}(\rho)$ the set of players that visit their target set along ρ , *i.e.*, $\text{Visit}(\rho) = \{i \in \Pi \mid \text{there exists } n \in \mathbb{N} \text{ st. } \rho_n \in F_i\}$.¹ The *social welfare* of ρ , denoted by $\text{SW}(\rho)$, is the pair $(|\text{Visit}(\rho)|, \sum_{i \in \text{Visit}(\rho)} \text{Cost}_i(\rho))$ if $\text{Visit}(\rho) \neq \emptyset$ and the pair $(0, +\infty)$ otherwise. Note that it takes into account both the number of players who visit their target set and their accumulated cost to reach those sets. Finally, let $P = \{(\text{Cost}_i(\rho))_{i \in \Pi} \mid \rho \in \text{Plays}(v_0)\} \subseteq (\mathbb{N} \cup \{+\infty\})^{|\Pi|}$. A cost profile $p \in P$ is *Pareto optimal in* $\text{Plays}(v_0)$ if it is minimal in P with respect to the componentwise ordering \leq on P .²

Let us now state the studied decision problems. The first two problems are classical: they ask whether there exists a solution (NE or SPE) σ satisfying certain requirements that impose bounds on either $(\text{Cost}_i(\langle \sigma \rangle_{v_0}))_{i \in \Pi}$ or on $\text{SW}(\langle \sigma \rangle_{v_0})$.

The most natural requirements are to impose upper bounds on the costs that the players have to pay and no lower bounds.

Problem 3 (Upper threshold decision problem). Given an initialized multiplayer quantitative Reachability game (\mathcal{G}, v_0) , given a threshold $y \in (\mathbb{N} \cup \{+\infty\})^{|\Pi|}$, decide whether there exists a solution σ such that

¹We can easily adapt this definition to histories.

²For convenience, we prefer to say that p is Pareto optimal in $\text{Plays}(v_0)$ rather than in P .

$$(\text{Cost}_i(\langle \sigma \rangle_{v_0}))_{i \in \Pi} \leq y.$$

Let us recall that one might also be interested in imposing an interval $[x_i, y_i]$ in which the cost paid by Player i must lie and that we have already solved this latter problem for SPEs in quantitative Reachability games (Chapter 11). In fact, the complexity class of the upper threshold decision problem is the same, as claimed in the following theorem.

Theorem 12.1.1. *For SPEs, Problem 3 with upper (and lower) bounds is PSPACE-complete.*

In the second problem, constraints are imposed on the social welfare, with the aim to maximize it. We use the lexicographic ordering on $\mathbb{N} \times (\mathbb{N} \cup \{+\infty\})$ such that $(k, c) \succeq (k', c')$ if and only if (i) $k > k'$ or (ii) $k = k'$ and $c \leq c'$.

Problem 4 (Social welfare decision problem). Given an initialized multiplayer quantitative Reachability game (\mathcal{G}, v_0) , given two thresholds $k \in \{0, \dots, |\Pi|\}$ and $c \in \mathbb{N} \cup \{+\infty\}$, decide whether there exists a solution σ such that $\text{SW}(\langle \sigma \rangle_{v_0}) \succeq (k, c)$.

Notice that with the lexicographic ordering, we want to first maximize the number of players who visit their target set, and then to minimize the accumulated cost to reach those sets. Let us now state the last studied problem.

Problem 5 (Pareto optimal decision problem). Given an initialized multiplayer quantitative Reachability game (\mathcal{G}, v_0) decide whether there exists a solution σ in (\mathcal{G}, v_0) such that $(\text{Cost}_i(\langle \sigma \rangle_{v_0}))_{i \in \Pi}$ is Pareto optimal in $\text{Plays}(v_0)$.

Remark 12.1.2. Problems 3 and 4 impose constraints with *non-strict* inequalities. We could also impose strict inequalities or even a mix of strict and non-strict inequalities. The results of this chapter can be easily adapted to




Figure 12.1: A two-player quantitative Reachability game with $F_1 = \{v_3, v_4\}$ and $F_2 = \{v_1, v_4\}$.

those variants.

We conclude this section with an illustrative example.

Example 12.1.3. Consider the quantitative Reachability game (\mathcal{G}, v_0) of Figure 12.1. We have two players such that the vertices of Player 1 (resp. Player 2) are the rounded (resp. rectangular) vertices. For the moment, the reader should not consider the value indicated on the right of the vertices' labeling. Moreover $F_1 = \{v_3, v_4\}$ and $F_2 = \{v_1, v_4\}$. In this figure, an edge (v, v') labeled by x should be understood as a path from v to v' with length x . Observe that F_1 and F_2 are both reachable from the initial vertex v_0 . Moreover the two Pareto optimal cost profiles are $(3, 3)$ and $(2, 6)$: take a play with prefix $v_0v_2v_4$ in the first case, and a play with prefix $v_0v_2v_3v_0v_1$ in the second case.

For this example, we claim that there is no NE (and thus no SPE) such that its cost profile is Pareto optimal (see Problem 5). Assume the contrary and suppose that there exists an NE σ such that its outcome ρ has cost profile $(3, 3)$, meaning that ρ begins with $v_0v_2v_4$. Then Player 1 has a profitable deviation such that after history v_0v_2 he goes to v_3 instead of v_4 in a way to pay a cost of 2 instead of 3, which is a contradiction. Similarly assume that there exists an NE σ such that its outcome ρ has cost profile $(2, 6)$, meaning that ρ begins with $v_0v_2v_3v_0v_1$. Then Player 2 has a profitable deviation such that after history v_0 he goes to v_1 instead of v_2 , again a contradiction. So there is no NE σ in (\mathcal{G}, v_0) such that $(\text{Cost}_i(\langle \sigma \rangle_{v_0}))_{i \in \Pi}$ is Pareto optimal in $\text{Plays}(v_0)$.

The previous discussion shows that there is no NE σ such that $(0, 0) = x \leq (\text{Cost}_i(\langle \sigma \rangle_{v_0}))_{i \in \Pi} \leq y = (3, 3)$ (see Problem 3). This is no longer true with $y = (6, 3)$. Indeed, one can construct an NE τ whose outcome has prefix $v_0v_1v_0v_2v_3$ and cost profile $(6, 3)$. This also shows that there exists an NE σ

(the same τ as before) that satisfies $\text{SW}(\langle \sigma \rangle_{v_0}) \succeq (k, c) = (2, 9)$ (with τ both players visit their target set and their accumulated cost to reach it equals 9).

12.2 Existence problems

In this section, we show that for particular families of Reachability games and requirements, there is no need to solve the related decision problems because they always have a positive answer in this case.

We begin with the family constituted by all Reachability games with a *strongly connected* arena. The next theorem then states that there always exists a solution that visits all non-empty target sets.

Theorem 12.2.1. *Let (\mathcal{G}, v_0) be an initialized multiplayer quantitative Reachability game such that its arena A is strongly connected. There exists an SPE σ (and thus an NE) such that its outcome $\langle \sigma \rangle_{v_0}$ visits all target sets F_i , $i \in \Pi$, that are non-empty.*

Let us comment on this result. For this family of games, the answer to Problem 3 is always positive for particular thresholds. In case of quantitative Reachability, take strict constraints $< +\infty$ if $F_i \neq \emptyset$ and non-strict constraints $\leq +\infty$ otherwise. We will see later that the strict constraints $< +\infty$ can be replaced by the non-strict constraints $\leq |V| \cdot |\Pi|$ (see Theorem 12.3.2). We will also see that, in this setting, the answer to Problem 4 is also always positive for thresholds $k = |\{i \mid F_i \neq \emptyset\}|$ and $c = |\Pi|^2 \cdot |V|$ (see Theorem 12.3.2). In order to ease the reading, we relegate the proof of Theorem 12.2.1 to Appendix B.2.1.

In the statement of Theorem 12.2.1, as the arena is strongly connected, F_i is non-empty if and only if F_i is reachable from v_0 . Also notice that the hypothesis that the arena is strongly connected is necessary. Indeed, it is easy to build an example with two players (Player 1 and Player 2) such that from v_0 it is not possible to reach both F_1 and F_2 . This is illustrated in Example 12.2.2.

Example 12.2.2. Consider the initialized quantitative Reachability game (\mathcal{G}, v_0) of Figure 12.2. There are two players, Player 1 who owns round vertices and Player 2 who owns square vertices, and $F_1 = \{v_1\}$, $F_2 = \{v_2\}$.

Clearly there is a unique NE $\sigma = (\sigma_1, \sigma_2)$ in (\mathcal{G}, v_0) such that $\sigma_1(v_0) = v_1$ and $\sigma_2(v_1) = v_1$, $\sigma_2(v_2) = v_2$. Its outcome only visits F_1 (and not F_2).




Figure 12.2: A two-player quantitative Reachability game with $F_1 = \{v_1\}$ and $F_2 = \{v_2\}$ where one target set is not reached in equilibrium.

We now turn to the second result of this section. The next theorem states that even with only two players there exists an initialized multiplayer quantitative Reachability game that has no NE with a cost profile which is Pareto optimal. To prove this result, we only have to come back to the quantitative Reachability game of Figure 12.1. We explained in Example 12.1.3 that there is no NE in this game such that its cost profile is Pareto optimal.

Theorem 12.2.3. *There exists an initialized multiplayer quantitative Reachability game with $|\Pi| = 2$ that has no NE with a cost profile which is Pareto optimal in $\text{Plays}(v_0)$.*

12.3 Solving decision problems

In this section, we present our main results concerning the three decision problems defined in Section 12.1. In Theorem 12.3.1 we provide our complexity results and in Theorem 12.3.2 the memory requirements for the equilibria.

Theorem 12.3.1. *Let (\mathcal{G}, v_0) be a multiplayer quantitative Reachability game.*

- *For NEs: Problem 3 and Problem 4 are NP-complete while Problem 5 is NP-hard and belongs to Σ_2^P .*
- *For SPEs: Problems 3, 4 and 5 are PSPACE-complete.*

Theorem 12.3.2. *Let (\mathcal{G}, v_0) be a multiplayer quantitative Reachability game.*

- *For NEs: for each decision problem, if the answer is positive, then there exists a strategy profile σ with memory in $\mathcal{O}(|\Pi| \cdot |V|)$ which satisfies the conditions.*
- *For SPEs: for each decision problem, if the answer is positive, then there exists a strategy profile σ with memory in $\mathcal{O}(|\Pi| \cdot 2^{|\Pi|} \cdot |V|^{(|\Pi|+2) \cdot (|V|+3)+1})$ which satisfies the conditions.*

Moreover, for both NEs and SPEs:

- *for Problem 3 and Problem 5, σ is such that: if $i \in \text{Visit}(\langle \sigma \rangle_{v_0})$, then $\text{Cost}_i(\langle \sigma \rangle_{v_0}) \leq |\Pi| \cdot |V|$;*
- *for Problem 4, if $\text{Visit}(\langle \sigma \rangle_{v_0}) \neq \emptyset$, then $\sum_{i \in \text{Visit}(\langle \sigma \rangle_{v_0})} \text{Cost}_i(\langle \sigma \rangle_{v_0}) \leq |\Pi|^2 \cdot |V|$.*

Notice that no assumption is made on the arena of the game. Even if we provide complexity lower bounds in Theorem 12.3.1, the main part is to give the upper bounds. Roughly speaking the decision algorithms work as follows: they guess a path and check that it is the outcome of an equilibrium satisfying the relevant property (such as Pareto optimality). In order to verify that a path is an equilibrium outcome, we rely on the outcome characterizations of equilibria, presented in Section 6.2 for NEs (resp. Section 7.4.3 for SPEs). These characterizations rely themselves on the notion of Visit λ -consistent play (resp. λ -consistent play). As the guessed path should be finitely representable, we show that it is sufficient to consider Visit λ -consistent lassoes (resp. λ -consistent lassoes), in Section 12.3.1. In order to make the presentation of the results uniform and since the notions of λ -consistent play and Visit λ -consistent plays are equivalent (Proposition 7.4.15) in the associated extended game of quantitative Reachability game, we enunciate our results using the Visit λ -consistency both for NEs and SPEs. We then expose in Section 12.3.2 the philosophy of the algorithms providing the upper bounds on the complexity of the three problems. Finally, all the technical details and proofs are relegated

to Appendix B.2.

12.3.1 Sufficiency of lassoes

In this section, we provide technical results which given a Visit λ -consistent play produce an associated Visit λ -consistent lasso. In the rest of this document, we show that working with these lassoes is sufficient for the algorithms.

The associated lassoes are built by eliminating some *unnecessary cycles* and then identifying a prefix $h\ell$ such that ℓ can be repeated infinitely often. An unnecessary cycle is a cycle inside of which no new player visits his target set. More formally, let $\rho = \rho_0\rho_1 \dots \rho_k \dots \rho_{k+\ell} \dots$ be a play in \mathcal{G} , if $\rho_k = \rho_{k+\ell}$ and $\text{Visit}(\rho_0 \dots \rho_k) = \text{Visit}(\rho_0 \dots \rho_{k+\ell})$ then the cycle $\rho_k \dots \rho_{k+\ell}$ is called an unnecessary cycle.

We call:

- (P1) the procedure which eliminates an unnecessary cycle, *i.e.*, let $\rho = \rho_0\rho_1 \dots \rho_k \dots \rho_{k+\ell} \dots$ such that $\rho_k \dots \rho_{k+\ell}$ is an unnecessary cycle, ρ becomes $\rho' = \rho_0 \dots \rho_k \rho_{k+\ell+1} \dots$
- (P2) the procedure which turns ρ into a lasso $\rho' = h\ell^\omega$ by copying ρ long enough for all players to visit their target set and then to form a cycle after the last player has visited his target set. If no player visits his target set along ρ , then (P2) only copies ρ long enough to form a cycle.

Notice that, given $\rho \in \text{Plays}$, applying (P1) or (P2) may involve a decreasing of the costs but for (P1) and (P2) $\text{Visit}(\rho) = \text{Visit}(\rho')$. Additionally, after applying (P2) we have that $\text{Visit}(h) = \text{Visit}(\rho')$. Moreover, applying (P1) until it is no longer possible and then (P2) leads to a lasso with length at most $(|\Pi| + 1) \cdot |V|$ and cost less than or equal to $|\Pi| \cdot |V|$ for players who have visited their target set.

Lemma 12.3.3. *Let (\mathcal{G}, v_0) be a multiplayer quantitative Reachability game and $\rho \in \text{Plays}$ be a play.*

- *If ρ' is obtained by applying (P1) on ρ , then $(\text{Cost}_i(\rho'))_{i \in \Pi} \leq (\text{Cost}_i(\rho))_{i \in \Pi}$*

- If ρ' is obtained by applying (P2) on ρ , then $(\text{Cost}_i(\rho'))_{i \in \Pi} = (\text{Cost}_i(\rho))_{i \in \Pi}$.
- Applying (P1) until it is no longer possible and then (P2), leads to a lasso ρ' with length at most $(|\Pi| + 1) \cdot |V|$ and $\text{Cost}_i(\rho') \leq |V| \cdot |\Pi|$ for each $i \in \text{Visit}(\rho')$.

□

Remark 12.3.4 (about Lemma 12.3.3). Notice that, given a quantitative Reachability game (\mathcal{G}, v_0) , as its extended game (\mathcal{X}, x_0) is also a quantitative Reachability game, all statements of Lemma 12.3.3 also hold for the latter game. Notice that the third assertion applied to (\mathcal{X}, x_0) leads to upper bounds where $|V|$ must be replaced by $|V^{\mathcal{X}}|$ which is exponential in $|\Pi|$ (see Definition 4.2.1).

In fact, even in the extended game (\mathcal{X}, x_0) we can obtain the same result: applying (P1) until it is no longer possible and then (P2), leads to a lasso ρ' with size at most $(|\Pi| + 1) \cdot |V|$ and $\text{Cost}_i(\rho') \leq |V| \cdot |\Pi|$ for each $i \in \text{Visit}(\rho')$. This is because along a play ρ in the extended game, the second components of the vertices of ρ form a non-decreasing sequence.

Additionally, applying (P1) or (P2) on Visit λ -consistent plays preserves this property. This is stated in Lemma 12.3.5 which is in particular true for extended games.

Lemma 12.3.5. *Let (\mathcal{G}, v_0) be a multiplayer quantitative Reachability game and $\rho \in \text{Plays}$ be a Visit λ -consistent play for a given labeling function λ . If ρ' is the play obtained by applying (P1) or (P2) on ρ , then ρ' is Visit λ -consistent.*

Lemmas 12.3.3 and 12.3.5 allow us to claim that it is sufficient to deal with lassoes with polynomial length to solve Problems 3, 4 and 5. Moreover, it yields some bounds on the needed memory and the costs for each problem as stated in the next two propositions.

The first proposition is used to solve Problems 3 and 4.

Proposition 12.3.6. *Let σ be an NE (resp. SPE) in a multiplayer quantitative Reachability game (\mathcal{G}, v_0) (resp. (\mathcal{X}, x_0) its extended game). Let $w_0 = v_0$ (resp. $w_0 = x_0$). Then there exists τ an NE (resp. SPE) in (\mathcal{G}, v_0) (resp. (\mathcal{X}, x_0)) such that:*

- $\langle \tau \rangle_{w_0}$ is a lasso $h\ell^\omega$ such that $|h\ell| \leq (|\Pi| + 1) \cdot |V|$;
- for each $i \in \text{Visit}(\langle \tau \rangle_{w_0})$, $\text{Cost}_i(\langle \tau \rangle_{w_0}) \leq |\Pi| \cdot |V|$;
- τ has memory in $\mathcal{O}(|\Pi| \cdot |V|)$ (resp. $\mathcal{O}(|\Pi| \cdot 2^{|\Pi|} \cdot |V|^{(|\Pi|+2) \cdot (|V|+3)+1})$).

Moreover, given $y \in (\mathbb{N} \cup \{+\infty\})^{|\Pi|}$, $k \in \{0, \dots, |\Pi|\}$ and $c \in \mathbb{N} \cup \{+\infty\}$:

- If $(\text{Cost}_i(\langle \sigma \rangle_{w_0}))_{i \in \Pi} \leq y$, then $(\text{Cost}_i(\langle \tau \rangle_{w_0}))_{i \in \Pi} \leq y$;
- If $\text{SW}(\langle \sigma \rangle_{w_0}) \succeq (k, c)$, then $\text{SW}(\langle \tau \rangle_{w_0}) \succeq (k, c)$.

The following proposition is used to solve Problem 5.

Proposition 12.3.7. *Let σ be an NE (resp. SPE) in a multiplayer quantitative Reachability game (\mathcal{G}, v_0) (resp. (\mathcal{X}, x_0) its extended game). Let $w_0 = v_0$ (resp. $w_0 = x_0$). If we have that $(\text{Cost}_i(\langle \sigma \rangle_{w_0}))_{i \in \Pi}$ is Pareto optimal in $\text{Plays}(w_0)$, then:*

- for all $i \in \text{Visit}(\langle \sigma \rangle_{w_0})$, $\text{Cost}_i(\langle \sigma \rangle_{w_0}) \leq |V| \cdot |\Pi|$;
- there exists τ an NE (resp. SPE) such that $\langle \tau \rangle_{w_0} = h\ell^\omega$, $|h\ell| \leq (|\Pi| + 1) \cdot |V|$ and $(\text{Cost}_i(\langle \sigma \rangle_{w_0}))_{i \in \Pi} = (\text{Cost}_i(\langle \tau \rangle_{w_0}))_{i \in \Pi}$;
- τ has memory in $\mathcal{O}(|\Pi| \cdot |V|)$ (resp. $\mathcal{O}(|\Pi| \cdot 2^{|\Pi|} \cdot |V|^{(|\Pi|+2) \cdot (|V|+3)+1})$).

For the sake of clarity we relegate the proofs of Propositions 12.3.6 and 12.3.7 to Appendix B.2.

12.3.2 Algorithms and memory requirements

In this section, we provide the main ideas behind the results stated in Theorems 12.3.1 and 12.3.2.

Algorithm for NEs

We first focus on Theorem 12.3.1 for NEs, *i.e.*, Problem 3 and Problem 4 are NP-complete while Problem 5 is NP-hard and belongs to Σ_2^P . We only provide algorithms to solve these problems and their related complexity, since the proof for the NP-hardness is very similar to the one given in [CFGR16]. Recall that Σ_2^P is by definition the class NP^{NP} , it is also equal to $\text{NP}^{\text{co-NP}}$. The algorithm for each problem works as follows:

1. it guesses a lasso of polynomial length;
2. it verifies that the cost profile of this lasso satisfies the conditions³ given by the problem;
3. it verifies that the lasso is the outcome of an NE.

Let us comment on the different steps of these algorithms.

- Step 1: For Problem 3 and Problem 4 (resp. Problem 5), it is sufficient to consider plays which are lassoes with polynomial length thanks to Proposition 12.3.6 (resp. Proposition 12.3.7).
- Step 3: This property is verified thanks to Theorem 6.2.3. This is done in polynomial time as the lasso has a polynomial length and the values of the coalitional games are computed in polynomial time [BGHM17].
- Step 2: For Problem 3 and Problem 4, this verification can be obviously done in polynomial time. For Problem 5, we need to have an oracle allowing us to know if the cost profile of the lasso is Pareto optimal. As a consequence, we study Problem 6 which lies in co-NP.

Problem 6. Given a multiplayer quantitative Reachability game (\mathcal{G}, v_0) (resp. its extended game (\mathcal{X}, x_0)) and a lasso $\rho \in \text{Plays}(v_0)$ (resp. $\rho \in \text{Plays}(x_0)$), we want to decide if $(\text{Cost}_i(\rho))_{i \in \Pi}$ is Pareto optimal in $\text{Plays}(v_0)$ (resp. $\text{Plays}(x_0)$).

³Satisfying the conditions is either satisfying the constraints (Problem 3 and Problem 4) or having a cost profile which is Pareto optimal (Problem 5).

Proposition 12.3.8. *Problem 6 lies in co-NP.*

Proof. Let (\mathcal{G}, v_0) be a multiplayer quantitative Reachability game (resp. (\mathcal{X}, x_0) be its extended game) and let $\rho \in \text{Plays}(v_0)$ (resp. $\rho \in \text{Plays}(x_0)$) be a lasso. If ρ is not Pareto optimal, there exists a play ρ' such that $(\text{Cost}_i(\rho))_{i \in \Pi} \geq (\text{Cost}_i(\rho'))_{i \in \Pi}$ and $(\text{Cost}_i(\rho))_{i \in \Pi} \neq (\text{Cost}_i(\rho'))_{i \in \Pi}$. Moreover, thanks to Lemma 12.3.3, one may assume that ρ' is a lasso with size at most $(|\Pi| + 1) \cdot |V|$. So, we only have to guess such a lasso ρ' and to verify that $(\text{Cost}_i(\rho))_{i \in \Pi} \geq (\text{Cost}_i(\rho'))_{i \in \Pi}$ and $(\text{Cost}_i(\rho))_{i \in \Pi} \neq (\text{Cost}_i(\rho'))_{i \in \Pi}$. This can be done in polynomial time. \square

Algorithm for SPEs

We now focus on Theorem 12.3.1 for SPEs, *i.e.*, Problems 3, 4 and 5 are PSPACE-complete. The PSPACE-completeness of Problem 3 is already solved (see Theorem 12.1.1). We thus provide algorithms to solve Problems 4 and 5 and their related complexity. We do not provide the proof for the PSPACE-hardness as it is very similar to the one given in Section 11.3.2.

The algorithm for Problem 4 and 5 works as follows:

1. it guesses a lasso of polynomial length;
2. it verifies that the cost profile c of this lasso satisfies the conditions given by the problem;
3. it checks, whether there exists an SPE with cost profile equal to c .

The explanations for the first and the second steps are the same as for the algorithms for NEs. Finally, we know that the third step can be done in PSPACE by Theorem 12.1.1.

Memory requirements

We now turn to Theorem 12.3.2 that provides memory requirements for the equilibria in case of positive answer to the studied decision problems. Its

proof directly follows from Proposition 12.3.6 (resp. Proposition 12.3.7) for Problems 3 and 4 (resp. Problem 5).

Notice nevertheless that these results provide memory requirements for SPEs in the extended game and not in the quantitative Reachability game itself. In fact the corresponding SPE in the quantitative Reachability game needs the same amount of memory.

We also present a useful result, although also interesting in its own right, to prove memory requirements for SPEs.

Proposition 12.3.9. *Let (\mathcal{G}, v_0) be an initialized multiplayer quantitative Reachability game and let (\mathcal{X}, x_0) be its associated extended game. Given a cost profile $c \in (\mathbb{N} \cup \{+\infty\})^{|\Pi|}$, we set $M = \max_{i \in \Pi} \{c_i \mid c_i < +\infty\}$ if this maximum exists and $M = 0$ otherwise. Then, the following assertions are equivalent:*

1. *There exists an SPE σ in (\mathcal{X}, x_0) such that $\text{Cost}(\langle \sigma \rangle_{x_0}) = c$;*
2. *There exists a finite good symbolic witness \mathcal{P} in (\mathcal{X}, x_0) such that:*
 - $\text{Cost}(\rho^{(0, x_0)}) = c$;
 - *the length of the lasso $\rho^{(0, x_0)}$ is bounded by $M + |V|$;*
 - *the length of each lasso $\rho^{(i, x)}$ (with $(i, x) \neq (0, x_0)$) is bounded by a value in $\mathcal{O}(|V|^{(|\Pi|+2) \cdot (|V|+3)} + |\Pi| \cdot |V|)$.*
3. *There exists an SPE τ in (\mathcal{X}, x_0) such that $\text{Cost}(\langle \tau \rangle_{x_0}) = c$ and the memory size of τ is in $\mathcal{O}(M + |\Pi| \cdot 2^{|\Pi|} \cdot |V|^{(|\Pi|+2) \cdot (|V|+3)+1})$.*

A proof of this result is given in Appendix B.2.

12.4 The qualitative setting

In the previous sections, we have considered quantitative Reachability problems. In this section we consider the qualitative variant and investigate the difference with the previously obtained results.

12.4.1 Qualitative Reachability games

All along this section we focus on *qualitative Reachability games*.

Let us recall that in this particular setting, players only aim at reaching their target set but do not take into account the number of steps it takes. Player i receives a gain of 1 if ρ visits his target set F_i , and a gain of 0 otherwise. Thus each player i wants to maximize his gain.

Lemma 12.4.1 ([De 13, Proposition 4.1.4]). *Let (\mathcal{G}, v_0) be an initialized multiplayer quantitative Reachability game and σ be a strategy profile in this game. Consider the related qualitative Reachability game \mathcal{G}' with the same arena A and target sets $(F_i)_{i \in \Pi}$, but the gain functions $(\text{Gain}_i)_{i \in \Pi}$. Then if σ is an NE (resp. SPE) in (\mathcal{G}, v_0) , then σ is also an NE (resp. SPE) in (\mathcal{G}', v_0) .*

Thus, as it is proved that there always exists an SPE (and thus an NE) in a quantitative Reachability game, there always exists one in a qualitative Reachability game.

Theorem 12.4.2. *In every initialized multiplayer qualitative Reachability game, there always exists an SPE, and thus also an NE.*

12.4.2 Decision problems and complexity results

In case of qualitative Reachability, as for quantitative Reachability games, we are interested in a solution that fulfills certain requirements. For example, we would like to know whether there exists a solution such that a maximum number of players visit their target sets.

Let (\mathcal{G}, v_0) be an initialized multiplayer qualitative Reachability game with $\mathcal{G} = (A, (\text{Gain}_i)_{i \in \Pi}, (F_i)_{i \in \Pi})$. Given $\rho \in \text{Plays}(v_0)$, we denote by $\text{Visit}(\rho)$ the set of players i such that ρ visits F_i , that is, $\text{Visit}(\rho) = \{i \in \Pi \mid \text{Gain}_i(\rho) = 1\}$. The *social welfare* $\text{SW}(\rho)$ of ρ is the size of $\text{Visit}(\rho)$. Let $P \subseteq \{0, 1\}^{|\Pi|}$ be the set of all gain profiles $p = (\text{Gain}_i(\rho))_{i \in \Pi}$, with $\rho \in \text{Plays}(v_0)$. A cost profile $p \in P$ is called *Pareto optimal in $\text{Plays}(v_0)$* if it is maximal in P with respect

to the componentwise ordering \leq on P . Notice that if there exists ρ with $\text{Visit}(\rho) = \Pi$, then its social welfare is the largest possible and there exists a unique Pareto optimal gain profile equal to $(1, 1, \dots, 1)$. Notice also that certain target sets F_i might be empty or not reachable from the initial vertex v_0 . Hence in this case, the best that we can hope is a (unique) Pareto optimal gain profile p such that $p_i = 1$ if and only if F_i is reachable⁴ from v_0 .

Qualitative variant of Problem 3 Given an initialized multiplayer qualitative Reachability game (\mathcal{G}, v_0) , given two thresholds $x, y \in \{0, 1\}^{|\Pi|}$, decide whether there exists a solution σ such that $x \leq (\text{Gain}_i(\rho))_{i \in \Pi} \leq y$.

Let us notice that the Qualitative variant of Problem 3 is in fact the (Boolean) constrained existence problem (Problem 1). We recall that imposing a lower bound $x_i = 1$ means that Player i has to visit his target set whereas imposing an upper bound $y_i = 0$ means that Player i cannot visit his target set.

Unlike quantitative Reachability, social welfare in qualitative Reachability games only aims to maximize the number of players who visit their target set.

Qualitative variant of Problem 4 Given an initialized multiplayer qualitative Reachability game (\mathcal{G}, v_0) , given a threshold $k \in \{0, \dots, |\Pi|\}$, decide whether there exists a solution σ such that $\text{SW}(\langle \sigma \rangle_{v_0}) \geq k$.

Let us now state the last studied problem for qualitative Reachability games.

Qualitative variant of Problem 5 Given an initialized multiplayer qualitative Reachability game (\mathcal{G}, v_0) decide whether there exists a solution σ in (\mathcal{G}, v_0) such that $(\text{Gain}_i(\langle \sigma \rangle_{v_0}))_{i \in \Pi}$ is Pareto optimal in $\text{Plays}(v_0)$.

The latter problem has some connections with the two previous ones. For instance in case of qualitative Reachability, suppose there exists a play in

⁴Notice that if F_i is reachable from v_0 , then it is necessarily not empty.

Plays(v_0) that visits all target sets. As already explained, there is only one Pareto optimal gain $(1, \dots, 1)$. Asking for the existence of a solution σ such that $(\text{Gain}_i(\langle \sigma \rangle_{v_0}))_{i \in \Pi}$ is Pareto optimal is equivalent to asking for the existence of a solution σ such that $\text{Gain}_i(\langle \sigma \rangle_{v_0})_{i \in \Pi} \geq (1, \dots, 1)$ (see Qualitative variant of Problem 3), or such that $\text{SW}(\langle \sigma \rangle_{v_0}) \geq |\Pi|$ (see Qualitative variant of Problem 4).

We can now state the qualitative variant of Theorem 12.3.1 whose results are obtained thanks to similar arguments.

Theorem 12.4.3. *Let (\mathcal{G}, v_0) be a multiplayer qualitative Reachability game.*

- *For NEs: the Qualitative variants of Problem 3 and Problem 4 are NP-complete while the Qualitative variant of Problem 5 is NP-hard and belongs to Σ_2^P .*
- *For SPEs: the Qualitative variants of Problems 3, 4 and 5 are PSPACE-complete.*

12.4.3 Existence problem

Theorem 12.4.4 is a direct consequence of Theorem 12.2.1 and Lemma 12.4.1.

Theorem 12.4.4. *Let (\mathcal{G}, v_0) be an initialized multiplayer qualitative Reachability game such that its arena A is strongly connected. Then there exists an SPE σ (and thus an NE) such that its outcome $\langle \sigma \rangle_{v_0}$ visits all target sets F_i , $i \in \Pi$, that are non-empty.*

Let us comment on this result. For this family of games, the answer to the Qualitative variant of Problem 3 is always positive for particular thresholds. Take thresholds x, y such that $x_i = 1$ (and thus $y_i = 1$) if and only if $F_i \neq \emptyset$. The answer to the Qualitative variant of Problem 4 is also always positive for threshold $k = |\{i \mid F_i \neq \emptyset\}|$. Finally, the answer to the Qualitative variant of Problem 5 is also always positive since there exists a unique Pareto optimal gain profile p such that $p_i = 1$ if and only if $F_i \neq \emptyset$.

Recall that we explained before why it was enough to prove Theorem 12.2.1 for SPEs and for quantitative reachability games only. Notice that in case of qualitative Reachability games, there exists a simpler construction of the required NE or SPE. Indeed, as the arena is strongly connected, there exists a play $\rho \in \text{Plays}(v_0)$ that visits all non-empty target sets. (i) Hence to get an NE, construct a strategy profile σ in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = \rho$. As the gain profile of σ is the best that each player can hope, no player has an incentive to deviate and σ is then an NE. (ii) The construction is a little more complex to get an SPE. We again construct a strategy profile σ in (\mathcal{G}, v_0) such that $\langle \sigma \rangle_{v_0} = \rho$, and inductively extend its construction to all subgames $(\mathcal{G}_{\upharpoonright h}, v)$ as follows. Assume that $\sigma_{\upharpoonright h}$ is not yet constructed, then extend the construction of σ such that $\sigma_{\upharpoonright h} = g\rho$ for some gv_0 starting in v and ending in v_0 (such a history gv_0 exists because the arena is strongly connected). In this way, the outcome of $\sigma_{\upharpoonright h}$ in each subgame (\mathcal{G}, v) has gain profile $(1, \dots, 1)$ and no player has an incentive to deviate. It follows that σ is an SPE.

The next theorem states that the Qualitative variant of Problem 5 has a positive answer for all *qualitative* Reachability games with a number of players *limited to two*, and that this existence result cannot be extended to three players.

Theorem 12.4.5. *Let (\mathcal{G}, v_0) be an initialized multiplayer qualitative Reachability game,*

- *Let (\mathcal{G}, v_0) be an initialized multiplayer qualitative Reachability game such that $|\Pi| = 2$, there exists an SPE σ (and thus an NE) with a gain profile that is Pareto optimal in $\text{Plays}(v_0)$.*
- *There exists an initialized multiplayer qualitative Reachability games with $|\Pi| = 3$ that has no NE with a gain profile that is Pareto optimal in $\text{Plays}(v_0)$.*

Let us focus on the proof of Theorem 12.4.5 which is based on the next lemma, which is interesting in its own right.

Lemma 12.4.6. *Let (\mathcal{G}, v_0) be an initialized multiplayer qualitative Reachability game. Let p be a gain profile equal to $(0, 0, \dots, 0)$ or $(1, 1, \dots, 1)$. If p is Pareto optimal^a in $\text{Plays}(v_0)$, then there exists an SPE σ with gain profile p .*

^a $(1, 1, \dots, 1)$ is trivially Pareto optimal.

Proof. The case $p = (0, 0, \dots, 0)$ is easy to solve. By Pareto optimality, all plays in $\text{Plays}(v_0)$ have gain profile p . Hence every strategy profile σ is trivially an SPE with gain profile p . Let us turn to case $p = (1, 1, \dots, 1)$ and let $\rho = \rho_0 \rho_1 \dots \in \text{Plays}(v_0)$ with gain profile p . By Theorem 12.4.2^a, there exists an SPE σ in (\mathcal{G}, v_0) . If $(\text{Gain}_i(\langle \sigma \rangle_{v_0}))_{i \in \Pi} = p$, we are done. Otherwise let us show how to modify σ into another SPE τ with outcome ρ and thus with gain profile p . Let $h \in \text{Hist}_i(v_0)$, $i \in \Pi$,

- if h is a prefix of ρ , then $\tau_i(h) = \rho_{|h|+1}$,
- otherwise, $\tau_i(h) = \sigma_i(h)$.

Let us prove that τ is an SPE. Clearly for each history hv that is not a prefix of ρ , $\tau_{\upharpoonright h} = \sigma_{\upharpoonright h}$ is an NE in the subgame $(\mathcal{G}_{\upharpoonright h}, v)$. So let $hv = \rho_0 \dots \rho_k$. As $\langle \tau_{\upharpoonright h} \rangle_v$ has gain profile $(1, 1, \dots, 1)$ in $(\mathcal{G}_{\upharpoonright h}, v)$, player i such that $v \in V_i$ has no incentive to deviate, and then $\tau_{\upharpoonright h}$ is also an NE in $(\mathcal{G}_{\upharpoonright h}, v)$. \square

^aNotice that we cannot apply Theorem 12.2.1 since the arena is not necessarily strongly connected.




Figure 12.3: A multiplayer qualitative Reachability game that has no NE with a gain profile that is Pareto optimal

Proof of Theorem 12.4.5. We begin with the first item. There are three cases to study: either the unique Pareto optimal gain profile of $\text{Plays}(v_0)$ is equal to $(0, 0)$, or it is equal to $(1, 1)$, or there are one or two Pareto optimal gain profiles that belong to $\{(0, 1), (1, 0)\}$. In the first two cases, we get the required SPE by Lemma 12.4.6. Hence it remains to treat the last case. From Lemma B.2.1, we know that there exists an SPE in (\mathcal{G}, v_0) whose outcome ρ visits a least one target set F_i , $i \in \{1, 2\}$. Therefore the gain profile of ρ is either equal to $(0, 1)$ or $(1, 0)$ as required.

For the second item, consider the initialized multiplayer qualitative Reachability game (\mathcal{G}, v_0) of Figure 12.3. We have three players such that player 3 owns diamond vertices. Moreover, $F_1 = \{v_4, v_5\}$, $F_2 = \{v_3, v_5\}$, and $F_3 = \{v_4, v_6\}$. There are four plays in $\text{Plays}(v_0)$ whose gain profile is indicated below each of them. The set of Pareto optimal gain profiles in $\text{Plays}(v_0)$ is equal to $\{(1, 0, 1), (1, 1, 0)\}$. Consider a strategy profile σ with outcome $v_0 v_1 v_4^\omega$ and gain profile $(1, 0, 1)$. Then it is not an NE because player 2 has a profitable deviation by going from v_1 to v_3 (instead of v_4). Similarly the strategy profile σ with outcome $v_0 v_2 v_5^\omega$ and gain profile $(1, 1, 0)$ is not an NE. Therefore there is no NE in (\mathcal{G}, v_0) with a gain profile that is Pareto optimal. \square

Part IV:

TIMED GAMES

In this last part, based on [BG20], we investigate how the characterizations and the studied problems of the previous parts may be considered in other models. For that purpose we choose to investigate multiplayer turn-based timed games, and in particular *multiplayer turn-based qualitative Reachability timed games*.

Even if we are conscious that our choice leads to redundancy, we want that this part can be read independently of the rest of this document. Hence all the needed background and previous results is recalled when needed. Moreover, since in the previous parts the arena of the games are assumed to be finite and without alphabet on the edge, this choice allows to avoid any ambiguity on the models of interest.

Remark 13.0.1. Let us mention the connections with the previous parts.

- In Section 16.2.1:
 - the notion of extended game of a qualitative Reachability game is defined: see Section 4.2;
 - The characterization of outcomes of SPEs in qualitative Reachability games based on λ^* -consistent plays is provided: see Section 7.4.2.
- In Section 16.2.2: we prove that the constrained existence problem of SPEs in (finite) qualitative Reachability game can be solved by an algorithm whose time complexity is at most exponential in the number of

players and polynomial in the size of its transition system. The arguments used in this proof are the same as those used in Section 9.1.2.

- Notice also that, in this part, when we consider Reachability objectives it is always a **qualitative** Reachability objective.

Games

In the context of *reactive systems*, *two-player zero-sum games played on graphs* are commonly used to model the purely antagonistic interactions between a system and its environment [PR89]. The system and the environment are the two players of a game played on a graph whose vertices represent the configurations. Finding how the system can ensure the achievement of his objective amounts to finding, if it exists, a *winning strategy* for the system.

When modeling complex systems with several agents whose objectives are not necessarily antagonistic, the two-player zero-sum framework is too restrictive and we rather rely on *multiplayer non zero-sum games*. In this setting, the notion of winning strategy is replaced by various notions of *equilibria* including the famous concept of *Nash equilibrium* (NE) [Nas50]. When considering games played on graphs, the notion of *subgame perfect equilibrium* (SPE) is often preferred to the classical Nash equilibrium [Os04]. Indeed, Nash equilibrium does not take into account the sequential structure of the game and may allow irrational behaviors in some subgames.

Timed games

Timed automata [AD94] is now a well established model for complex systems including real time features. Timed automata have been naturally extended into two-player zero-sum *timed games* [AM99, CDF⁺05, BCD⁺07, JT07]. Multiplayer non zero-sum extensions have also been considered [BBM10, Bre12, KNP19]. In these models both time and multiplayer aspects coexist. In this non zero-sum timed framework, the main focus has been on NE, and, to our knowledge, not on SPE.

Main contributions and organization of the part.

In this part, we consider *multiplayer, non zero-sum, turn-based timed games with reachability objectives* together with the concept of SPE. We focus on the constrained existence problem (for SPE): given a timed game, we want to decide whether there exists an SPE where some players have to win and some other ones have to lose. The main result of this part is a proof that the SPE constrained existence problem is EXPTIME-complete for reachability timed games. Let us notice that the NE constrained existence problem for reachability timed games is also EXPTIME-complete [Bre12]. This may look surprising as often, there is a complexity jump when going from NE to SPE, for example the constrained existence problem on qualitative reachability game is NP-complete for NE [CFGR16] and PSPACE-complete for SPE [BBGR18] (see also Chapter 10). Intuitively, the complexity jump is avoided because the exponential blow up due to the transition from SPE to NE is somehow absorbed by the classical exponential blow up due to the classical region graph used for the analysis of timed systems.

In order to obtain an EXPTIME algorithm, we proceed in different steps. In the first step, we prove that the game variant of the *classical region graph* is a good abstraction for the SPE constrained existence problem. In fact, we identify conditions on *bisimulations* under which the study of SPE of a given (potentially infinite game) can be reduced to the study of its quotient. This is done in Chapter 15 for (untimed) games with general objectives. In Chapter 16, we then focus on (untimed) finite reachability game and provide an EXPTIME algorithm to solve the constrained existence problem. Proving this result may look surprising, as we already know from [BBGR18] (Chapter 10) that this problem is indeed PSPACE-complete for (untimed) finite games. However the PSPACE algorithm provided in [BBGR18] (Chapter 10) did not allow us to obtain the EXPTIME algorithm for timed games. The latter EXPTIME algorithm is discussed in Chapter 17.

Related works

There are many results on games played on graphs, we refer the reader to [Bru17] for a survey and an extended bibliography. Here we focus on the re-

sults directly related to our contributions. The constrained existence of SPEs is studied in finite multiplayer turn-based games with different kinds of objectives, for example: (qualitative) reachability and safety objectives [BBGR18] (Chapter 10), ω -regular winning conditions [Umm06], quantitative reachability objectives [BBG⁺19],... In [BBM10], they prove that the constrained existence problem for Nash equilibria in concurrent timed games with reachability objectives is EXPTIME-complete. This same problem in the same setting is studied in [Bre12] with others qualitative objectives.

Transition systems, bisimulations and quotients

A *transition system* is a tuple $T = (\Sigma, V, E)$ where (i) Σ is a finite alphabet; (ii) V a set of *states* (also called *vertices*) and (iii) $E \subseteq V \times \Sigma \times V$ a set of *transitions* (also called *edges*). To ease the notation, an edge $(v_1, a, v_2) \in E$ is sometimes denoted by $v_1 \xrightarrow{a} v_2$. Notice that V may be uncountable. We said that the transition system is *finite* if V and E are finite.

Given two transition systems on the same alphabet $T_1 = (\Sigma, V_1, E_1)$ and $T_2 = (\Sigma, V_2, E_2)$, a *simulation of T_1 by T_2* is a binary relation $\mathbf{R} \subseteq V_1 \times V_2$ which satisfies the following conditions: (i) $\forall v_1, v'_1 \in V_1, \forall v_2 \in V_2$ and $\forall a \in \Sigma$: $((v_1, v_2) \in \mathbf{R} \text{ and } v_1 \xrightarrow{a} v'_1) \Rightarrow (\exists v'_2 \in V_2, v_2 \xrightarrow{a} v'_2 \text{ and } (v'_1, v'_2) \in \mathbf{R})$ and (ii) for each $v_1 \in V_1$ there exists $v_2 \in V_2$ such that $(v_1, v_2) \in \mathbf{R}$. We say that T_2 *simulates* T_1 . It implies that any transition $v_1 \xrightarrow{a} v'_1$ in T_1 is simulated by a corresponding transition $v_2 \xrightarrow{a} v'_2$ in T_2 .

Given two transition systems on the same alphabet $T_1 = (\Sigma, V_1, E_1)$ and $T_2 = (\Sigma, V_2, E_2)$, a *bisimulation between T_1 and T_2* is a binary relation $\mathbf{R} \subseteq V_1 \times V_2$ such that \mathbf{R} is a simulation of T_1 by T_2 and the converse relation \mathbf{R}^{-1} is a simulation of T_2 by T_1 where $\mathbf{R}^{-1} = \{(v_2, v_1) \in V_2 \times V_1 \mid (v_1, v_2) \in \mathbf{R}\}$. When \mathbf{R} is a bisimulation between two transition systems, we write β instead of \mathbf{R} . If $T = (\Sigma, V, E)$ is a transition system, a bisimulation on $V \times V$ is called

a bisimulation on T .

Given a transition system $T = (\Sigma, V, E)$ and an equivalence relation \sim on V , we define the *quotient of T by \sim* , denoted by $\tilde{T} = (\Sigma, \tilde{V}, \tilde{E})$, as follows: (i) $\tilde{V} = \{[v]_\sim \mid v \in V\}$ where $[v]_\sim = \{v' \in V \mid v \sim v'\}$ and (ii) $[v_1]_\sim \xrightarrow{a} [v_2]_\sim$ if and only if there exist $v'_1 \in [v_1]_\sim$ and $v'_2 \in [v_2]_\sim$ such that $v'_1 \xrightarrow{a} v'_2$. When the equivalence relation is clear from the context, we write $[v]$ instead of $[v]_\sim$.

Given a transition system $T = (\Sigma, V, E)$, a bisimulation \sim on T which is also an equivalence relation is called a *bisimulation equivalence*. In this context, the following result holds.

Lemma 14.0.1. *Given a transition system T and a bisimulation equivalence \sim , there exists a bisimulation \sim_q between T and its quotient \tilde{T} . This bisimulation is given by the function $\sim_q: V \rightarrow \tilde{V} : v \mapsto [v]_\sim$*

Turn-based games

Arenas, plays and histories An *arena* $A = (\Sigma, V, E, \Pi, (V_i)_{i \in \Pi})$ is a tuple where (i) $T = (\Sigma, V, E)$ is a transition system such that for each $v \in V$, there exists $a \in \Sigma$ and $v' \in V$ such that $(v, a, v') \in E$; (ii) $\Pi = \{1, \dots, n\}$ is a finite set of players and (iii) $(V_i)_{i \in \Pi}$ is a partition of V between the players. An arena is finite if its transition system T is finite.

A *play* in A is an infinite path in its transition system, *i.e.*, $\rho = \rho_0 \rho_1 \dots \in V^\omega$ is a play if for each $i \in \mathbb{N}$, there exists $a \in \Sigma$ such that $(\rho_i, a, \rho_{i+1}) \in E$. A *history* h in A can be defined in the same way but $h = h_0 \dots h_k \in V^*$ for some $k \in \mathbb{N}$ is a finite path in the transition system. We denote the set of plays by Plays and the set of histories by Hist . When it is necessary, we use the notation Plays_A and Hist_A to recall the underlying arena A . Moreover, the set Hist_i is the set of histories such that their last vertex v is a vertex of Player i , *i.e.*, $v \in V_i$. A play (resp. a history) in (\mathcal{G}, v_0) is then a play (resp. a history) in \mathcal{G} starting in v_0 . The set of such plays (resp. histories) is denoted by $\text{Plays}(v_0)$ (resp. $\text{Hist}(v_0)$). We also use the notation $\text{Hist}_i(v_0)$ when these histories end in a vertex $v \in V_i$.

Given a play $\rho \in \text{Plays}$ and $k \in \mathbb{N}$, its suffix $\rho_k \rho_{k+1} \dots$ is denoted by $\rho_{\geq k}$.

We denote by $\text{Succ}(v) = \{v' \mid (v, a, v') \in E \text{ for some } a \in \Sigma\}$ the set of successors of v , for $v \in V$, and by Succ^* the transitive closure of Succ . Given a play $\rho = \rho_0\rho_1\dots$, the set $\text{Occ}(\rho) = \{v \in V \mid \exists k, \rho_k = v\}$ is the set of vertices *visited* along ρ .

Remark 14.0.2. When we consider a play in an arena $A = (\Sigma, V, E, \Pi, (V_i)_{i \in \Pi})$, we do not care about the alphabet letter associated with each edge of the play. It is the reason why two different infinite paths in $T = (\Sigma, V, E)$ $v_0 \xrightarrow{a} v_1 \xrightarrow{a} \dots \xrightarrow{a} v_n \xrightarrow{a} \dots$ and $v_0 \xrightarrow{b} v_1 \xrightarrow{b} \dots \xrightarrow{b} v_n \xrightarrow{b} \dots$ correspond to only one play $\rho = v_0v_1\dots v_n\dots$ in A . The same phenomenon appears with finite paths and histories. We explain later why this is not a problem for our purpose.

Multiplayer turn-based game An (*initialized multiplayer Boolean turn-based*) game is a tuple $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \Pi})$ such that: (i) $A = (\Sigma, V, E, \Pi, (V_i)_{i \in \Pi})$ is an arena; (ii) $v_0 \in V$ is the *initial vertex* and (iii) for each $i \in \Pi$, $\text{Gain}_i : \text{Plays} \rightarrow \{0, 1\}$ is a gain function for Player i . In this setting, each player $i \in \Pi$ is equipped with a set $\Omega_i \subseteq \text{Plays}$ that we call the *objective* of Player i . Thus, for each $i \in \Pi$, for each $\rho \in \text{Plays}$: $\text{Gain}_i(\rho) = 1$ if and only if $\rho \in \Omega_i$. If $\text{Gain}_i(\rho) = 1$ (resp. $= 0$), we say that Player i *wins* (resp. *loses*) along ρ . In the remaining part of this document, we refer to the notion of initialized multiplayer Boolean turn-based game by the term “game”. For each $\rho \in \text{Plays}$, we write $\text{Gain}(\rho) = p$ for some $p \in \{0, 1\}^{|\Pi|}$ to depict $\text{Gain}_i(\rho) = p_i$ for each $i \in \Pi$.

Strategies and outcomes Given a game (\mathcal{G}, v_0) , a *strategy* of Player i is a function $\sigma_i : \text{Hist}_i(v_0) \rightarrow V$ with the constraint that for each $hv \in \text{Hist}_i(v_0)$, $\sigma_i(hv) \in \text{Succ}(v)$. A play $\rho = \rho_0\rho_1\dots$ is *consistent* with σ_i if for each ρ_k such that $\rho_k \in V_i$, $\rho_{k+1} = \sigma_i(\rho_0\dots\rho_k)$. A *strategy profile* $\sigma = (\sigma_i)_{i \in \Pi}$ is a tuple of strategies, one for each player. Given a game (\mathcal{G}, v_0) and a strategy profile σ , there exists a unique play from v_0 consistent with each strategy σ_i . We call this play the *outcome* of σ and denote it by $\langle \sigma \rangle_{v_0}$.

Remark 14.0.3. We follow up Remark 14.0.2. The objectives we consider are of the form $\Omega \subseteq \text{Plays}$. These objectives only depend on the sequence of visited states along a play (for example: visiting infinitely often a given state) regardless of the sequence of visited alphabet letters. This is why defining the

strategy of a player with a choice of the next vertex instead of a couple of an alphabet letter and a vertex is not a problem. Actually, in all this part one may consider that the alphabet is $\Sigma = \{a\}$. The reason why we allow alphabet letters on edges is to be able to consider *synchronous products* of (timed) automata [BK08, AD94]. In this way, we could consider wider class of objectives (see Section 17.4).

Subgame perfect equilibria

In the multiplayer game setting, the solution concepts usually studied are *equilibria* (see [GU08]). We here recall the concepts of Nash equilibrium and subgame perfect equilibrium.

Let $\sigma = (\sigma_i)_{i \in \Pi}$ be a strategy profile in a game (\mathcal{G}, v_0) . When we highlight the role of Player i , we denote σ by (σ_i, σ_{-i}) where σ_{-i} is the profile $(\sigma_j)_{j \in \Pi \setminus \{i\}}$. A strategy $\sigma'_i \neq \sigma_i$ is a *deviating* strategy of Player i , and it is a *profitable deviation* for him if $\text{Gain}_i(\langle \sigma \rangle_{v_0}) < \text{Gain}_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0})$. A strategy profile σ in a game (\mathcal{G}, v_0) is a *Nash equilibrium* (NE) if no player has an incentive to deviate unilaterally from his strategy, *i.e.*, no player has a profitable deviation.

A refinement of NE is the concept of *subgame perfect equilibrium* (SPE) which is a strategy profile being an NE in each subgame. Formally, given a game $(\mathcal{G}, v_0) = (\mathbf{A}, (\text{Gain}_i)_{i \in \Pi})$ and a history $hv \in \text{Hist}(v_0)$, the game $(\mathcal{G}_{\upharpoonright h}, v)$ is called a *subgame* of (\mathcal{G}, v_0) such that $\mathcal{G}_{\upharpoonright h} = (\mathbf{A}, (\text{Gain}_{i \upharpoonright h})_{i \in \Pi})$ and $\text{Gain}_{i \upharpoonright h}(\rho) = \text{Gain}_i(h\rho)$ for all $i \in \Pi$ and $\rho \in V^\omega$. Notice that (\mathcal{G}, v_0) is subgame of itself. Moreover if σ_i is a strategy for Player i in (\mathcal{G}, v_0) , then $\sigma_{i \upharpoonright h}$ denotes the strategy in $(\mathcal{G}_{\upharpoonright h}, v)$ such that for all histories $h' \in \text{Hist}_i(v)$, $\sigma_{i \upharpoonright h}(h') = \sigma_i(hh')$. Similarly, from a strategy profile σ in (\mathcal{G}, v_0) , we derive the strategy profile $\sigma_{\upharpoonright h}$ in $(\mathcal{G}_{\upharpoonright h}, v)$. Let (\mathcal{G}, v_0) be a game, following this formalism, a strategy profile σ is a *subgame perfect equilibrium* in (\mathcal{G}, v_0) if for all $hv \in \text{Hist}(v_0)$, $\sigma_{\upharpoonright h}$ is an NE in $(\mathcal{G}_{\upharpoonright h}, v)$.

Studied problem

Given a game (\mathcal{G}, v_0) , several SPEs may coexist. It is the reason why we are interested in the *constrained existence* of an SPE in this game: some players have to win and some other ones have to lose. The related decision problem is the following one:

Definition 14.0.4 (Constrained existence problem). Given a game (\mathcal{G}, v_0) and two gain profiles $x, y \in \{0, 1\}^{|\Pi|}$, does there exist an SPE σ in (\mathcal{G}, v_0) such that $x \leq \text{Gain}(\langle \sigma \rangle_{v_0}) \leq y$.

CHAPTER 15

SPEs IN A GAME AND ITS QUOTIENT

In this section, we first define the concepts of *bisimulation between games* and of *bisimulation on a game*. Then, we explain how given such bisimulations we can obtain a new game, called the *quotient game*, thanks to a quotient of the initial game. Finally, we prove that if there exists an SPE in a game with a given gain profile, there exists an SPE in its associated quotient game with the same gain profile, and vice versa.

15.1 Game bisimulation

We extend the notion of bisimulation between transition systems (resp. on a transition system) to the one of bisimulation between games (resp. on a game). In this document, by bisimulation between games (resp. on a game) we mean:

Definition 15.1.1 (Game bisimulation). Given two games $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \Pi})$ and $(\mathcal{G}', v'_0) = (A', (\text{Gain}'_i)_{i \in \Pi})$ with the same alphabet and the same set of players, we say that $\sim \subseteq V \times V'$ is a *bisimulation between* (\mathcal{G}, v_0) and (\mathcal{G}', v'_0) if (i) \sim is a bisimulation between $T = (\Sigma, V, E)$ and $T' = (\Sigma, V', E')$ and (ii) $v_0 \sim v'_0$. In the same way, if $\sim \subseteq V \times V$ we say that \sim is a *bisimulation on* (\mathcal{G}, v_0) if \sim is a bisimulation on $T = (\Sigma, V, E)$.

The notion of bisimulation equivalence on a transition system is extended in the same way to games. In the rest of this document, we use the following notations: (1) If $\sim \subseteq V \times V'$ is a bisimulation between $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \Pi})$ and $(\mathcal{G}', v'_0) = (A', (\text{Gain}'_i)_{i \in \Pi})$, for each $\rho \in \text{Plays}_A$ and for all $\rho' \in \text{Plays}_{A'}$, we write $\rho \sim \rho'$ if and only if for each $n \in \mathbb{N}$: $\rho_n \sim \rho'_n$. (2) If $\sim \subseteq V \times V$ is a bisimulation on $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \Pi})$, for each $\rho \in \text{Plays}_A$ and for all $\rho' \in \text{Plays}_A$, we write $\rho \sim \rho'$ if and only if for each $n \in \mathbb{N}$: $\rho_n \sim \rho'_n$. (3) Notations 1 and 2 can be naturally adapted to histories¹.

A natural property that should be satisfied by a bisimulation on a game is the respect of the vertices partition. It means that if a vertex bisimulates another vertex, then these vertices should be owned by the same player.

Definition 15.1.2 (\sim respects the partition). Given a game $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \Pi})$ and a bisimulation \sim on (\mathcal{G}, v_0) , we say that \sim *respects the partition* if for all $v, v' \in V$ such that $v \sim v'$, if $v \in V_i$ then $v' \in V_i$.

15.2 Quotient game

Given a game (\mathcal{G}, v_0) and a bisimulation equivalence \sim on it which respects the partition, one may consider its associated *quotient game* $(\tilde{\mathcal{G}}, [v_0])$ such that its transition system is defined as the quotient of the transition system of (\mathcal{G}, v_0) .

Definition 15.2.1 (Quotient game). Given a game $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \Pi})$ such that $A = (\Sigma, V, E, \Pi, (V_i)_{i \in \Pi})$, if \sim is a bisimulation equivalence on (\mathcal{G}, v_0) which respects the partition, the associated *quotient game* $(\tilde{\mathcal{G}}, [v_0]) = (\tilde{A}, (\tilde{\text{Gain}}_i)_{i \in \Pi})$ is defined as follows: (i) $\tilde{A} = (\Sigma, \tilde{V}, \tilde{E}, (\tilde{V}_i)_{i \in \Pi})$ is such that $\tilde{T} = (\Sigma, \tilde{V}, \tilde{E})$ is the quotient of T and, for each $i \in \Pi$, $[v] \in \tilde{V}_i$ if and only if $v \in V_i$ and (ii) for each $i \in \Pi$, $\tilde{\text{Gain}}_i : \text{Plays}_{\tilde{A}} \rightarrow \{0, 1\}$ is the gain function

¹Once again, with this convention it is possible that two plays (or histories) such that $\rho \sim \rho'$ do not preserve the sequence of alphabet letters as it should be when we classically consider bisimulated paths in two bisimulated transitions systems. Remark 14.0.3 explains why it is not a problem for us.

of Player i .

In order to preserve some equivalent properties between a game and its quotient game, the equivalence relation on the game should respect the gain functions in both games. It means that if we consider two bisimulated plays either both in the game itself or one in the game and the other one in its quotient game, the gain profile of these plays should be equal.

Definition 15.2.2 (\sim respects the gain functions). Given an initialized game $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \pi})$ such that $A = (\Sigma, V, E, \Pi, (V_i)_{i \in \Pi})$ and a bisimulation equivalence \sim on (\mathcal{G}, v_0) , we say that \sim *respects the gain functions* if the following properties hold: (i) for each ρ and ρ' in Plays, if $\rho \sim \rho'$ then $\text{Gain}(\rho) = \text{Gain}(\rho')$ and (ii) for each $\rho \in \text{Plays}_A$ and $\tilde{\rho} \in \text{Plays}_{\tilde{A}}$, if $\rho \sim_q \tilde{\rho}$ then $\text{Gain}(\rho) = \tilde{\text{Gain}}(\tilde{\rho})$.

15.3 SPE existence

The aim of this section is to prove that, if there exists an SPE in a game equipped with a bisimulation equivalence which respects the partition and the gain functions, there exists an SPE in its associated quotient game with the same gain profile, and vice versa.

Theorem 15.3.1. *Let $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \Pi})$ be a game and $(\tilde{\mathcal{G}}, [v_0]) = (\tilde{A}, (\tilde{\text{Gain}}_i)_{i \in \Pi})$ its associated quotient game where \sim is a bisimulation equivalence on (\mathcal{G}, v_0) . If \sim respects the partition and the gain functions, we have that: there exists an SPE σ in (\mathcal{G}, v_0) such that $\text{Gain}(\langle \sigma \rangle_{v_0}) = p$ if and only if there exists an SPE τ in $(\tilde{\mathcal{G}}, [v_0])$ such that $\tilde{\text{Gain}}(\langle \tau \rangle_{[v_0]}) = p$.*

The key idea is to prove that: if there exists an SPE in a game equipped with a bisimulation equivalence, there exists an SPE in this game which is *uniform* and with the same gain profile (see Proposition 15.3.3). If σ_i is a uniform strategy, each time we consider two histories $h \sim h'$, the choices of Player i taking into account h or h' are in the same equivalence class (see

Definition 15.3.2).

Definition 15.3.2. Let (\mathcal{G}, v_0) be a game and \sim a bisimulation on it, we say that the strategy σ_i is *uniform* if for all $h, h' \in \text{Hist}_i(v_0)$ such that $h \sim h'$, we have that $\sigma_i(h) \sim \sigma_i(h')$. A strategy profile σ is *uniform* if for all $i \in \Pi$, σ_i is uniform.

Proposition 15.3.3. Let $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \pi})$ be a game and \sim be a bisimulation equivalence on (\mathcal{G}, v_0) which respects the partition and such that for each ρ and ρ' in Plays, if $\rho \sim \rho'$ then $\text{Gain}(\rho) = \text{Gain}(\rho')$, there exists an SPE σ in (\mathcal{G}, v_0) such that $\text{Gain}(\langle \sigma \rangle_{v_0}) = p$ if and only if there exists an SPE τ in (\mathcal{G}, v_0) which is uniform and such that $\text{Gain}(\langle \tau \rangle_{v_0}) = p$.

In order to ease the reading, the proofs of this section are provided in Appendix B.3.

CHAPTER 16

QUALITATIVE REACHABILITY GAMES

In this section we focus on a particular kind of game called (*qualitative*) *reachability game*. In these games, each player has a subset of vertices that he wants to reach. First, we formally define the concepts of *reachability games* and *reachability quotient games*. Then, we provide an algorithm which solves the constrained existence problem in finite reachability games in time complexity at most exponential in the number of players and polynomial in the size of the transition system of the game.

16.1 Reachability games and quotient reachability games

Definition 16.1.1. A *reachability game* $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \Pi}, (F_i)_{i \in \Pi})$ is a game where each player $i \in \Pi$ is equipped with a *target set* F_i that he wants to reach. Formally, the objective of Player i is $\Omega_i = \{\rho \in \text{Plays} \mid \text{Occ}(\rho) \cap F_i \neq \emptyset\}$ where $F_i \subseteq V$. This is a *reachability objective*.

Given a reachability game $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \Pi}, (F_i)_{i \in \Pi})$ and a bisimulation equivalence \sim on this game which respects the partition, one may consider its quotient game $(\tilde{\mathcal{G}}, [v_0]) = (\tilde{A}, (\tilde{\text{Gain}}_i)_{i \in \Pi}, (\tilde{F}_i)_{i \in \Pi})$ where for each $i \in \Pi$, $\tilde{F}_i \subseteq \tilde{V}$. In attempts to ensure the respect of the gain functions by \sim ,

we add a natural property on \sim (see Definition 16.1.2) and define the sets \tilde{F}_i in a proper way. In the rest of this document, we assume that this property is satisfied and that the quotient game of a reachability game is defined as in Definition 16.1.3.

Definition 16.1.2 (\sim respects the target sets). Let (\mathcal{G}, v_0) be a reachability game and \sim be a bisimulation equivalence on this game, we say that \sim respects the target sets if for all $v \in V$ and for all $v' \in V$ such that $v \sim v'$: $v \in F_i \Leftrightarrow v' \in F_i$.

Definition 16.1.3 (Reachability quotient game). Given a reachability game $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \Pi}, (F_i)_{i \in \Pi})$ and a bisimulation equivalence \sim on this game which respects the partition and the target sets, its quotient game is the reachability game $(\tilde{\mathcal{G}}, [v_0]) = (\tilde{A}, (\tilde{\text{Gain}}_i)_{i \in \Pi}, (\tilde{F}_i)_{i \in \Pi})$ where $\tilde{F}_i = \{[v]_{\sim} \mid v \in F_i\}$ for each $i \in \Pi$. We call this game the *reachability quotient game*.

Lemma 16.1.4. *Let (\mathcal{G}, v_0) be a reachability game and let \sim be a bisimulation equivalence. If \sim respects the target sets in this game, then \sim respects the gain functions.*

16.2 Complexity results

It is proved that the constrained existence problem is PSPACE-complete in finite reachability games [BBGR18] (Chapter 10). Our final purpose is to obtain an EXPTIME algorithm for the constrained existence problem on *reachability timed games* (see Chapter 17). Naively applying the PSPACE algorithm of [BBGR18] (Chapter 10) to the region games would lead to an EXPSPACE algorithm. That is why we provide here an alternative EXPTIME algorithm to solve the constrained existence problem on (untimed) finite games. This new algorithm will have the advantage to have a running time at most exponential only in the number of players (and polynomial in the size of its transition system). This feature will be crucial to obtain the EXPTIME algorithm on timed

games.

Theorem 16.2.1. *Given a finite reachability game (\mathcal{G}, v_0) , the constrained existence problem can be solved by an algorithm whose time complexity is at most exponential in $|\Pi|$ and polynomial in the size of its transition system.*

This approach follows the proof for quantitative reachability games in [BBG⁺19] (Chapter 11). This latter proof relies on two key ingredients: (i) the *extended game* of a reachability game and (ii) an *SPE outcome characterization based on a fixpoint computation of a labeling function of the states*. Those two key ingredients will be defined below.

16.2.1 Extended game

Let (\mathcal{G}, v_0) be a finite reachability game, its associated *extended game* $(\mathcal{X}, x_0) = (X, (\text{Gain}_i^X)_{i \in \Pi}, (F_i^X)_{i \in \Pi})$ is the reachability game such that the vertices are enriched with the set of players that have already visited their target sets along a history. The arena $X = (\Sigma, V^X, E^X, \Pi, (V_i^X)_{i \in \Pi})$ is defined as follows: (i) $V^X = V \times 2^\Pi$; (ii) $((v, I), a, (v', I')) \in E^X$ if and only if $(v, a, v') \in E$ and $I' = I \cup \{i \in \Pi \mid v' \in F_i\}$; (iii) $(v, I) \in V_i^X$ if and only if $v \in V_i$; (iv) $(v, I) \in F_i^X$ if and only if $i \in I$ and $(v, x_0) = (v_0, I_0)$ where $I_0 = \{i \in \Pi \mid v_0 \in F_i\}$.

The construction of (\mathcal{X}, x_0) from (\mathcal{G}, v_0) causes an exponential blow-up of the number of states. The main idea of this construction is that if you consider a play $\rho = (v_0, I_0)(v_1, I_1) \dots (v_n, I_n) \dots \in \text{Plays}_X(x_0)$, the set I_n means that each player $i \in I_n$ has visited his target set along $\rho_0 \dots \rho_n$. The important points are that there is a one-to-one correspondence between plays in $\text{Plays}_A(v_0)$ and $\text{Plays}_X(x_0)$ and that the gain profiles of two corresponding plays beginning in the initial vertices are equal. From these observations, we have:

Proposition 16.2.2. *Let (\mathcal{G}, v_0) be a reachability game and (\mathcal{X}, x_0) be its associated extended game, let $p \in \{0, 1\}^{|\Pi|}$ be a gain profile, there exists an SPE σ in (\mathcal{G}, v_0) with gain profile p if and only if there exists an SPE τ in (\mathcal{X}, x_0) with gain profile p .*

In the rest of this section, we will write $v \in V^X$ (instead of (u, I)) and we depict by $I(v)$ the set I of the players who have already visited their target set.

Outcome characterization

Once this extended game is built, we want a way to decide whether a play in this game corresponds to the outcome of an SPE or not: we want an *SPE outcome characterization*. The vertices of the extended game are labeled thanks to a *labeling function* $\lambda^* : V^X \rightarrow \{0, 1\}$. For a vertex $v \in V^X$ such that $v \in V_i^X$, the value 1 imposes that Player i should reach his target set if he follows an SPE from v and the value 0 does not impose any constraint on the gain of Player i from v .

The labeling function λ^* is obtained thanks to an iterative procedure such that each step k of the iteration provides a λ^k -labeling function. This procedure is based on the notion of λ -consistent play: that is a play which satisfies the constraints given by λ all along it.

Definition 16.2.3. Let $\lambda : V^X \rightarrow \{0, 1\}$ be a labeling function and $\rho \in \text{Plays}_X$, we say that ρ is λ -consistent if for each $i \in \Pi$ and for each $n \in \mathbb{N}$ such that $\rho_n \in V_i^X$: $\text{Gain}_i^X(\rho_{\geq n}) \geq \lambda(\rho_n)$. We write $\rho \models \lambda$.

The iterative computation of the sequence $(\lambda^k)_{k \in \mathbb{N}}$ works as follows: (i) at step 0, for each $v \in V^X$, $\lambda^0(v) = 0$, (ii) at step $k + 1$, for each $v \in V^X$, by assuming that $v \in V_i^X$, $\lambda^{k+1}(v) = \max_{v' \in \text{Succ}(v)} \min\{\text{Gain}_i^X(\rho) \mid \rho \in \text{Plays}_X(v') \wedge \rho \models \lambda^k\}$ and (iii) we stop when we find $n \in \mathbb{N}$ such that for each $v \in V^X$, $\lambda^{n+1}(v) = \lambda^n(v)$. The least natural number k^* which satisfies (iii) is called the *fixpoint* of $(\lambda^k)_{k \in \mathbb{N}}$ and λ^* is defined as λ^{k^*} . The following lemma states that this natural number exists and so that the iterative procedure stops.

Lemma 16.2.4. The sequence $(\lambda^k)_{k \in \mathbb{N}}$ reaches a fixpoint in $k^* \in \mathbb{N}$. Moreover, k^* is at most equal to $|V| \cdot 2^{|\Pi|}$.

Proof sketch. In the initialization step, all the vertex values are equal to 0. Then at each iteration, (i) if the value of a vertex was equal to 1 in the previous step, then it stays equal to 1 all along the procedure and (ii) if the value of the vertex was equal to 0 then it either stays equal to 0 (for this iteration step) or it becomes equal to 1 (for all the next steps thanks to (i)). At each step, at least one vertex value changes and when no value changes the procedure has reached a fixpoint which corresponds to the values of λ^* . Thus, it means that λ^* is obtained in at most $|V| \times 2^{|\Pi|}$ steps. \square

As claimed in the following proposition, the labeling function λ^* exactly characterizes the set of SPE outcomes. The proof is quite the same as for the quantitative setting.

Proposition 16.2.5. *Let (\mathcal{X}, x_0) be the extended game of a finite reachability game (\mathcal{G}, v_0) and let $\rho^X \in \text{Plays}_{\mathcal{X}}(x_0)$ be a play, there exists an SPE σ with outcome ρ^X in (\mathcal{X}, x_0) if and only if ρ^X is λ^* -consistent.*

16.2.2 Complexity

Proposition 16.2.5 allows us to prove Theorem 16.2.1. Indeed, we only have to find a play in the extended game which is λ^* -consistent and with a gain profile which satisfies the constrained given by the decision problem.

Proof sketch of Theorem 16.2.1. Let $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \Pi}, (F_i)_{i \in \Pi})$ be a reachability game and let $(\mathcal{X}, x_0) = (X, (\text{Gain}_i^X)_{i \in \Pi}, (F_i^X)_{i \in \Pi})$ be its associated extended game. The game (\mathcal{X}, x_0) is built from (\mathcal{G}, v_0) in time at most exponential in the number of players and polynomial in the size of the transition system of A .

The proof will be organized in three steps whose respective proofs will rely on the previous step(s): (i) given a gain profile $p \in \{0, 1\}^{|\Pi|}$, given λ^k for some $k \in \mathbb{N}$ and given some $v \in V^X$, we show that we can decide in the required complexity the existence of a play which is λ^k -consistent, beginning in v and with gain profile p ; (ii) given λ^k for some $k \in \mathbb{N}$, we show that the computation of λ^{k+1} can be performed within the required complexity;

and finally (iii) given $x, y \in \{0, 1\}^{|\Pi|}$, we show that the existence of a λ^* -consistent play beginning in x_0 with a gain profile p such that $x \leq p \leq y$ can be decided within the required complexity.

- **Proof of (i):** Given λ^k , $v \in V^X$ and $p \in \{0, 1\}^{|\Pi|}$, we want to know if there exists a play $\rho \in \text{Plays}_X(v)$ which is λ^k -consistent and with gain profile p . If a play ρ is such that $\text{Gain}^X(\rho) = p$, then for each $i \in \Pi$ such that $p_i = 1$, the condition of being a λ^k -consistent play is satisfied. For those such that $p_i = 0$, for each $n \in \mathbb{N}$ such that $\rho_n \in V_i^X$, $\text{Gain}_i^X(\rho_{\geq n}) = 0$ should be greater than $\lambda^k(\rho_n)$. This condition is satisfied if and only if for each $n \in \mathbb{N}$ such that $\rho_n \in V_i^X$, $\lambda^k(\rho_n) \neq 1$. Thus, we remove from (\mathcal{X}, x_0) all vertices (and all related edges) $v \in V_i^X$ such that $\lambda^k(v) = 1$, for each player i such that $p_i = 0$. Then, we only have to check if there exists a play ρ which begins in v and with gain profile p in this modified extended reachability game. This can be done in $O(2^{|\Pi|} \cdot (|V^X| + |E^X|))$ (Lemma 2.2.22), thus this procedure runs in time at most exponential in the number of players and polynomial in the size of the transition system of A .
- **Proof of (ii):** Given λ^k , we want to compute λ^{k+1} . For each $v \in V^X$, $\lambda^{k+1}(v) = \max_{v' \in \text{Succ}(v)} \min\{\text{Gain}_i^X(\rho) \mid \rho \in \text{Plays}_X(v') \wedge \rho \models \lambda^k\}$ (by assuming that $v \in V_i^X$). Thus for each $v' \in \text{Succ}(v)$, we have to compute $\min = \min\{\text{Gain}_i^X(\rho) \mid \rho \in \text{Plays}_X(v') \wedge \rho \models \lambda^k\}$. But $\min = 0$ if and only if there exists $\rho \in \text{Plays}_X(v')$ which is λ^k -consistent and such that $\text{Gain}_i^X(\rho) = 0$. Thus for each $p \in \{0, 1\}^{|\Pi|}$ such that $p_i = 0$, we use point (i) to decide if $\min = 0$. From that follows a procedure which runs in $O(|V^X| \cdot |V^X| \cdot 2^{|\Pi|} \cdot 2^{|\Pi|} \cdot (|V^X| + |E^X|))$ (running time at most exponential in the number of players and polynomial in the size of the transition system A).
- **Proof of (iii):** It remains to prove that the existence of a λ^* -consistent play beginning in x_0 with a gain profile p such that $x \leq p \leq y$ can be decided within the required complexity. In order to do so, we evaluate the complexity to obtain λ^* . First, we build λ^0 such that $\lambda^0(v) = 0$ for all $v \in V^X$ in $O(|V^X|)$ time. Then, we apply point (ii) at most

$|V| \cdot 2^{|\Pi|}$ times (by Lemma 16.2.4) to obtain λ^* . Given $x, y \in \{0, 1\}^{|\Pi|}$, we consider each $p \in \{0, 1\}^{|\Pi|}$ such that $x \leq p \leq y$ (at most $2^{|\Pi|}$ such ones) and we use point (i) to check if there exists a play which begins in x_0 with gain profile p and which is λ^* -consistent. This can be done in running time at most exponential in the number of players and polynomial in the size of the transition system of A .

We conclude the proof by applying Proposition 16.2.5. □

CHAPTER 17

APPLICATION TO TIMED GAMES

In this section, we are interested in models which are enriched with *clocks* and *clock guards* in order to consider time elapsing. *Timed automata* [AD94] are well known among such models. We recall some of their classical concepts, then we explain how (*turn-based*) *timed games* derive from timed automata.

17.1 Timed automata and timed games

In this section, we use the following notations. The set $C = \{c_1, \dots, c_k\}$ denotes a set of k *clocks*. A *clock valuation* is a function $\nu : C \rightarrow \mathbb{R}^+$. The set of clock valuations is depicted by C_V . Given a clock valuation ν , for $i \in \{1, \dots, k\}$, we sometimes write ν_i instead of $\nu(c_i)$. Given a clock valuation ν and $d \in \mathbb{R}^+$, $\nu + d$ denote the clock valuation $\nu + d : C \rightarrow \mathbb{R}^+$ such that $(\nu + d)(c_i) = \nu(c_i) + d$ for each $c_i \in C$. A *guard* is any finite conjunctions of expressions of the form $c_i \diamond x$ where c_i is a clock, $x \in \mathbb{N}$ is a natural number and \diamond is one of the symbols $\{\leq, <, =, >, \geq\}$. We denote by G the set of guards. Let g be a guard and ν be a clock valuation, notation $\nu \models g$ means that (ν_1, \dots, ν_k) satisfies g . A *reset* $Y \in 2^C$ indicates which clocks are reset to 0. We denote by $[Y \leftarrow 0]\nu$ the valuation ν' such that for each $c \in Y$, $\nu'(c) = 0$ and for each $c \in C \setminus Y$, $\nu'(c) = \nu(c)$.

A *timed automaton* (TA) is a tuple $(\mathcal{A}, \ell_0) = (\Sigma, L, \rightarrow, C)$ where: (i) Σ is a finite alphabet; (ii) L is a finite set of *locations*; (iii) C is a finite set of *clocks*; (iv) $\rightarrow \subseteq L \times \Sigma \times G \times 2^C \times L$ a finite set of transitions; and (v) $\ell_0 \in L$ an initial location. Additionally, we may equip a timed automaton with a set of players and partition the locations between them. It results in a *players partitioned timed automaton*.

Definition 17.1.1 ((Reachability) Players partitioned timed automaton). A *players partitioned timed automaton* (PPTA) $(\mathcal{A}, \ell_0) = (\Sigma, L, \rightarrow, C, \Pi, (L_i)_{i \in \Pi})$ is a timed automaton equipped with: (i) Π a finite set of players and (ii) $(L_i)_{i \in \Pi}$ a partition of the locations between the players. If (\mathcal{A}, ℓ_0) is equipped with a target set $\text{Goal}_i \subseteq L$ for each player $i \in \Pi$, we call it a *reachability PPTA*.

The semantic of a timed automaton (\mathcal{A}, ℓ_0) is given by its associated transition system $T_{\mathcal{A}} = (\Sigma, V, E)$ where: (i) $V = L \times C_V$ is a set of vertices of the form (ℓ, ν) where ℓ is a location and $\nu : C \rightarrow \mathbb{R}^+$ is a clock valuation; and (ii) $E \subseteq V \times \Sigma \times V$ is such that $((\ell, \nu), a, (\ell', \nu')) \in E$ if $(\ell, a, g, Y, \ell') \in \rightarrow$ for some $g \in G$ and some $Y \in 2^C$, and there exists $d \in \mathbb{R}^+$ such that: (1) for each $x \in X \setminus Y$: $\nu'(x) = \nu(x) + d$ (**time elapsing**); (2) for each $x \in Y$: $\nu'(x) = 0$ (**clocks resetting**); (3) $\nu + d \models g$ (**respect of the guard**).

In the same way, the semantic of a PPTA (\mathcal{A}, ℓ_0) is given by its associated game $(\mathcal{G}_{\mathcal{A}}, v_0)$.

Definition 17.1.2 ((Reachability) Timed games $\mathcal{G}_{\mathcal{A}}$). Let $(\mathcal{A}, \ell_0) = (\Sigma, L, \rightarrow, C, \Pi, (L_i)_{i \in \Pi})$ be a PPTA, its associated game $(\mathcal{G}_{\mathcal{A}}, v_0) = (\mathbb{A}_{\mathcal{A}}, (\text{Gain}_i)_{i \in \Pi})$, called *timed game*, is such that: (i) $\mathbb{A}_{\mathcal{A}} = (\Sigma, V, E, \Pi, (V_i)_{i \in \Pi})$ where $T_{\mathcal{A}} = (\Sigma, V, E)$ is the associated transition system of (\mathcal{A}, ℓ_0) and, for each $i \in \Pi$, $(\ell, \nu) \in V_i$ if and only if $\ell \in L_i$; (ii) for each $i \in \Pi$, $\text{Gain}_i : \text{Plays}_{\mathbb{A}_{\mathcal{A}}} \rightarrow \{0, 1\}$ is a gain function; (iii) $v_0 = (\ell_0, \mathbf{0})$ where $\mathbf{0}$ is the clock valuation such that for all $c \in C$, $\mathbf{0}(c) = 0$. If (\mathcal{A}, ℓ_0) is a reachability PPTA, its associated timed game is a reachability game $(\mathcal{G}_{\mathcal{A}}, v_0) = (\mathbb{A}_{\mathcal{A}}, (\text{Gain}_i)_{i \in \Pi}, (F_i)_{i \in \Pi})$ such that for each $i \in \Pi$, $(\ell, \nu) \in$

F_i if and only if $\ell \in \text{Goal}_i$. We call this game a *reachability timed game*.

Thus, in a timed game, when it is the turn of Player i to play, if the play is in location ℓ , he has to choose a delay $d \in \mathbb{R}^+$ and a next location ℓ' such that $(\ell, a, g, Y, \ell') \in \rightarrow$ for some $a \in \Sigma$, $g \in G$ and $Y \in 2^C$. If the choice of d respects the guard g , then the choice of Player i is valid: the clock valuation evolves according to the past clock valuation, d and Y and location ℓ' is reached. Then, the play continues.

17.2 Regions and region games

In this section, we consider a bisimulation equivalence on $T_{\mathcal{A}}$ (the classical time-abstract bisimulation from [AD94]) which allows us to solve the constrained existence in the quotient of the original timed game (the region game). All along this section we use the following notations. We denote by x_i the maximum value in the guards for clock c_i . For all positive number $d \in \mathbb{R}^+$, $\lfloor d \rfloor$ is the integral part of d and \bar{d} is the fractional part of d .

Definition 17.2.1 (\approx and region).

- Two clock valuations ν and ν' are equivalent (written $\nu \approx \nu'$) iff: (i) $\lfloor \nu_i \rfloor = \lfloor \nu'_i \rfloor$ or $\nu_i, \nu'_i > x_i$, for all $i \in \{1, \dots, k\}$; (ii) $\bar{\nu}_i = 0$ iff $\bar{\nu}'_i = 0$, for all $i \in \{1, \dots, k\}$ with $\nu_i \leq x_i$ and (iii) $\bar{\nu}_i \leq \bar{\nu}_j$ iff $\bar{\nu}'_i \leq \bar{\nu}'_j$ for all $i \neq j \in \{1, \dots, k\}$ with $\nu_j \leq x_j$ and $\nu_i \leq x_i$.
- We extend the equivalence relation to the states ($\approx \subseteq V \times V$): $(\ell, \nu) \approx (\ell', \nu')$ iff $\ell = \ell'$ and $\nu \approx \nu'$;
- A *region* r is an equivalence class for some $v \in V$: $r = [v]_{\approx}$.

This equivalence relation on clocks and its extension to states of $T_{\mathcal{A}}$ is usual and the following result is well known [AD94].

Lemma 17.2.2 ([AD94]). *Let (\mathcal{A}, ℓ_0) be a TA, $\approx \subseteq V \times V$ is a bisimulation equivalence on $T_{\mathcal{A}}$.*

It means that if $(\mathcal{G}_{\mathcal{A}}, v_0)$ is a (reachability) timed game, \approx is a bisimulation equivalence on it. Moreover, it respects the partition. Thus, we can consider the (reachability) quotient game of this game. We call this game the *(reachability) region game*. Notice that \approx respects the target sets, so the reachability quotient game is defined as in Definition 16.1.3.

Definition 17.2.3 ((Reachability) region game). *Let $(\mathcal{G}_{\mathcal{A}}, v_0)$ be a (reachability) timed game and $\approx \subseteq V \times V$ be the bisimulation equivalence defined in Definition 17.2.1, its associated *(reachability) region game* is its associated (reachability) quotient game $(\tilde{\mathcal{G}}_{\mathcal{A}}, [v_0])$.*

We recall [AD94] that the size of $\tilde{T}_{\mathcal{A}}$, *i.e.*, its number of states (regions) and edges, is in $O((|V| + |\rightarrow|) \cdot 2^{|\delta(\mathcal{A})|})$ where $\delta(\mathcal{A})$ is the binary encoding of the constants (guards and costs) appearing in \mathcal{A} . Thus $|\tilde{T}_{\mathcal{A}}|$ is in $O(2^{|\mathcal{A}|})$ where $|\mathcal{A}|$ takes into account the locations, edges and constants of \mathcal{A} . From this follows the following lemma.

Lemma 17.2.4. *The (reachability) region game $(\tilde{\mathcal{G}}_{\mathcal{A}}, [v_0])$ is a finite (reachability) game.*

Finally, in light of the construction of the reachability region game, the bisimulation equivalence \approx respects the gain functions of the reachability timed game and of the reachability region game.

Lemma 17.2.5. *Given $(\mathcal{G}_{\mathcal{A}}, v_0) = (\mathbb{A}_{\mathcal{A}}, (\text{Gain}_i)_{i \in \Pi}, (F_i)_{i \in \Pi})$ be a reachability timed game and $(\tilde{\mathcal{G}}_{\mathcal{A}}, [v_0]) = (\tilde{\mathbb{A}}_{\mathcal{A}}, (\tilde{\text{g}}_i)_{i \in \Pi}, (\tilde{F}_i)_{i \in \Pi})$ its associated region game, \approx respects the gain functions.*

Remark 17.2.6. Let $\mathcal{A} = (\Sigma, L, \rightarrow, C)$ be a timed automaton, $T_{\mathcal{A}} = (\Sigma, V, E)$

be its associated transition system and \approx be the bisimulation equivalence on $T_{\mathcal{A}}$ as defined in Definition 17.2.1, we have that $((\ell, \nu), a, (\ell', \nu')) \in E$ if and only if there exist $g \in G$, $Y \in 2^C$ and $d \in \mathbb{R}^+$ such that $(\ell, a, g, Y, \ell') \in \rightarrow$, $\nu' = [Y \leftarrow 0](\nu + d)$ and $\nu + d \models g$. Thus, we abstract the notion of time elapsing in the edges of the transition system.

Then, since \approx is a bisimulation equivalence on $T_{\mathcal{A}}$, for all $((\ell_1, \nu_1), a, (\ell'_1, \nu'_1)) \in E$ and for all $(\ell_2, \nu_2) \in V$ such that $(\ell_1, \nu_1) \approx (\ell_2, \nu_2)$, there exists $(\ell'_2, \nu'_2) \in V$ such that $((\ell_2, \nu_2), a, (\ell'_2, \nu'_2)) \in E$ and $(\ell'_1, \nu'_1) \approx (\ell'_2, \nu'_2)$. The time elapsing between ν_1 and ν'_1 is not necessarily the same as between ν_2 and ν'_2 . Thus, \approx is a timed abstract bisimulation in the classical way [AD94].

17.3 Complexity results

Theorem 17.3.1. *Given a reachability PPTA (\mathcal{A}, ℓ_0) and $x, y \in \{0, 1\}^{|\Pi|}$, the constrained existence problem in reachability timed games is EXPTIME-complete.*

The EXPTIME-hardness is due to a reduction from *countdown games* and is inspired by the one provided in [Bre12, Section 6.3.3]. Thus, we only prove the EXPTIME-easiness and provide a proof sketch for the EXPTIME-hardness.

EXPTIME-easiness. Given a PPTA (\mathcal{A}, ℓ_0) with target sets $(\text{Goal}_i)_{i \in \Pi}$ and given $x, y \in \{0, 1\}^{|\Pi|}$. Thanks to Theorem 15.3.1, it is equivalent to solve this problem in the reachability region game. Moreover, the size of the reachability region game is exponential, because its transition system $\tilde{\tilde{T}}_{\mathcal{A}}$ is exponential in the size of \mathcal{A} , but not in the number of players. Then, since the reachability region game is a finite reachability game (Lemma 17.2.4), we can apply Theorem 16.2.1. It causes an exponential blow-up in the number of players but is polynomial in the size of transition system $\tilde{\tilde{T}}_{\mathcal{A}}$. Thus, this entire procedure runs in (simple) exponential time in the size of the PPTA (\mathcal{A}, ℓ_0) .

□

EXPTIME-hardness (Proof sketch). In [Bre12], Proposition 6.12 asserts that the value problem for timed games with Büchi objectives and only two clocks is EXPTIME-hard. The proof relies on the notion of countdown game [JLS07] which is known to be EXPTIME-complete. When reading the proof of the latter proposition, one can easily be convinced that it is also proved that the value problem for timed games with reachability objectives and only two clocks is EXPTIME-hard. Indeed, the only accepting state is a deadlock with a self-loop (named w_{\exists}). Moreover, one can also notice that although the results of [Bre12] concern concurrent games, the proof of [Bre12, Proposition 6.12] relies on turn-based games.

The proof of Proposition 6.12 can be slightly modified in order to prove that the constrained existence problem in reachability timed games is EXPTIME-hard with two clocks. The problem in the original proof being that Adam does not have a reachability, but a safety objective. Given a countdown game \mathcal{C} , we build a reachability timed games by using nearly the same construction as the one presented in the proof of [Bre12, Proposition 6.12]. The difference are the following ones.

- We replace all the guards $y \neq c_0$ by the guards $y < c_0$.
- We add a winning state for Adam w_{\forall} .
- From every state belonging to Eve, we add a transition to w_{\forall} with guard $x = 0 \wedge y > c_0$.

The proposed transformations does not really affect the behaviors of the timed game, in the sense that it still bisimulates closely the countdown game. The only difference is discussed below. In the original encoding, Eve was winning if and only if she is able to reach w_{\exists} . This could happen only when the clock y is equal to c_0 . As the game is zero-sum, Adam was winning when w_{\exists} is never reached. In practice, as the timed game of the encoding is strongly non-zeno, in every winning play of Adam, the clock value y eventually overtakes c_0 . In our new encoding, every winning play of Adam ends up in w_{\forall} . That is the only difference. This is important, as we can now see the timed game as a reachability timed game where both players have a reachability

objective. One can be convinced that Eve has a winning strategy (in the original timed game proposed in [Bre12]) if and only if there exists an SPE where only Eve achieves her objective (in the variant of the timed game proposed above). \square

Notice that, since there always exists an SPE in a finite reachability game [Umm06], there always exists an SPE in the region game and so in the reachability timed game (Theorem 15.3.1).

17.4 Discussion and future works

In this part, we focus on (qualitative) reachability timed games, and ignore the effect of Zeno behaviors¹. Nevertheless we believe that our approach is rather robust and can be extended to richer objectives and take into account Zeno behaviors. In the following paragraphs, we try to briefly explain how this could be achieved.

Time-bounded reachability. A natural extension of our framework would be to equip the objective of each player with a time-bound. Player i aims at visiting F_i within TB_i time units. We believe that this time-bound variant of our constrained problem is decidable. Indeed, for each player, his time-bound reachability objective can easily be encoded via a deterministic timed automaton (on finite timed words) \mathcal{A}_i . Given a timed game \mathcal{G}_b equipped with a time-bounded objective for each player (described via \mathcal{A}_i), we could, via standard product construction build a new reachability timed game (without time-bound) \mathcal{G} . Solving the constrained existence problem (with time-bound) in \mathcal{G}_b is equivalent to solving the constrained existence problem (of Definition 14.0.4) in \mathcal{G} (the constrained being encoded in the \mathcal{A}_i 's). This approach could extend to any property that can be expressed via a deterministic timed automaton.

¹ A run $\rho = (\ell_0, \nu_0) \xrightarrow{d_1, a_1} (\ell_1, \nu_1) \xrightarrow{d_2, a_2} \dots$ in a timed automaton is said *timed-divergent* if the sequence $(\sum_{j \leq i} d_j)_i$ diverges. A timed automaton is *non-Zeno* if any finite run can be extended into a time-divergent run [BFL⁺18].

Towards ω -regular objectives. Let us briefly explain how our approach could be adapted to prove the decidability of the constrained existence problem for timed games with ω -regular objectives. For the sake of clarity, we here focus on parity objectives. First, let us notice that the results of Chapter 15 (including Theorem 15.3.1) apply to a general class of games, including infinite games with classical ω -regular objectives such as parity. An algorithm to decide the constrained existence problem (Definition 14.0.4) on parity on finite games can be found in [Umm06] via translation into tree automata. Equipped with these two tools, we believe that we could adapt the definitions and results of Chapter 17 to obtain the decidability of the constrained existence problem for parity timed games. Notice that, in order to obtain our complexity results for finite reachability games, we use other simpler tools than tree automata.

About Zenoness. In the present document, we allow a player to win (or to prevent other players to win) even if his strategy is responsible of Zeno behaviors. In [dAFH⁺03], the authors propose an elegant approach to *blame* a player that would prevent divergence of time. The main idea is to transform the ω -regular objective of each player into another one which will make him lose if he blocks the time. We believe that this idea could be exploited in our framework in order to prevent from winning a “blocking time player”.

Part V:

DISCUSSION AND FUTURE
WORKS

Let us conclude this document with a brief discussion about some potential future works.

About discounted-sum games In Section 7, we have provided a characterization of outcomes of weak SPEs that can be applied to games which respect some properties. We have instantiated this characterization to multiplayer Boolean games with prefix-independent gain functions, multiplayer qualitative and quantitative Reachability games and multiplayer Safety games.

One can wonder if this characterization may be instantiated to multiplayer discounted-sum games. Indeed, a discounted-sum objective is strongly prefix-linear and thus at least one of the properties required by the characterization is satisfied. Moreover, the continuity of the objective functions should make possible to prove the maxima existence property referred to in the statement of the characterization.

This would be particularly interesting as it would also imply a characterization of outcomes of SPEs since the notions of weak SPE and SPE are equivalent when the objective functions are continuous.

Nevertheless, let us mention that, concerning multiplayer discounted-sum games, the constrained existence problem of equilibria is closely related to the *target discounted-sum problem*. This problem is still an open problem. More details about this problem are provided in [Bru17].

Other kinds of considered equilibria, relevant equilibria or objectives

In this thesis, we have mainly focused on weak SPEs (and SPEs). In the same way the characterization of NEs has been adapted to secure equilibria in two player games [BMR14], we can wonder whether it is possible to adapt our characterization of weak SPE (resp. SPE) to some variants of secure equilibria or other notion of equilibrium. Notice that subgame perfect secure equilibria has already been studied in [De 13].

Concerning the existence of relevant equilibria, we have principally considered the constrained existence problem. We have also studied some variants in Reachability games: the social welfare decision problem and the Pareto optimal decision problem. We could consider these two variants on other kinds of games than Reachability games. We could also wonder whether other types of

relevant, more ad hoc, equilibria could be considered depending on the nature of the considered game, *e.g.*, relevant equilibria that takes into account the time aspect in multiplayer timed games.

In this document, the objective functions of our interest were classical prefix-independent qualitative objective (Büchi, co-Büchi, Parity, ...), qualitative and quantitative Reachability objectives and Safety objectives. One can be interested in considering other objective functions. For example, one may consider *multi-objectives* [Ran14]: each player has several objectives that he wants to satisfy.

Beyond determinism In this thesis, we consider multiplayer games without probability distributions on the edges of the game graph: if a player chooses to go in vertex v , he is sure that he will reach vertex v . A potential future work could be to consider *stochastic games* and to wonder how a characterization of equilibria in such a setting may be designed.

Another potential direction is to consider *multi-strategies*. In the classical definition of strategy, a strategy of a player is a function that assigns exactly one next vertex to each history in the game graph. In [BMOU11], authors consider the notion of multi-strategy to deal with the error-prone nature of computer systems. A multi-strategy of a player is a function that assigns a set of next vertices to each history in the game graph. In a prospective perspective, in order to solve problems related to the existence of equilibria in this setting, one can consider how equilibrium outcome characterizations, in the same spirit of those developed in this thesis, may be obtained and used with respect to the notion of multi-strategy.

Towards effective algorithms In this thesis, our results are of theoretical nature, hence a natural question arises: “How can we design effective algorithms to find relevant equilibria in games?”. If we want to directly exploit the equilibria outcomes characterizations, we should compute the values of some labeling function λ and then find a λ -consistent play. The algorithms presented in the previous chapters (in order to obtain the complexity results) compute the values of the labeling function λ , globally, in a monolithic fashion. One could wonder whether, the computation of λ can be done more locally. For instance,

one could decompose the game graph into strongly connected components, compute λ in each component, hoping that the global λ could be recovered from these local informations. One could also abandon exact algorithms, and consider metaheuristics in order to find quickly a relevant λ -consistent play in the game.

Multiplayer timed games Up to our knowledge, few results are known about multiplayer timed games ([Bre12, BRS17]), it seems to be a promising direction for future works. In Section 17, we provide some potential future works directly related to the results presented in Part IV.

BIBLIOGRAPHY

- [AD94] Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126:183–235, 1994.
6 citations in pages 238, 244, 259, 261, 262 and 263
- [AG11] Krzysztof R Apt and Erich Grädel. *Lectures in game theory for computer scientists*. Cambridge University Press, 2011.
Cited in page 21
- [AM99] Eugene Asarin and Oded Maler. As soon as possible: Time optimal control for timed automata. In Frits W. Vaandrager and Jan H. van Schuppen, editors, *HSCC 1999*, volume 1569 of *Lecture Notes in Computer Science*, pages 19–30. Springer, 1999.
Cited in page 238
- [BBDPG13] Thomas Brihaye, Véronique Bruyère, Julie De Pril, and Hugo Gimbert. On (Subgame Perfect) Secure Equilibrium in Quantitative Reachability Games. *Logical Methods in Computer Science*, Volume 9, Issue 1, February 2013.
Cited in page 21
- [BBG⁺19] Thomas Brihaye, Véronique Bruyère, Aline Goeminne, Jean-François Raskin, and Marie van den Bogaard. The complexity of subgame perfect equilibria in quantitative reachability games. In Wan J. Fokkink and Rob van Glabbeek, editors, *30th Interna-*

- tional Conference on Concurrency Theory, CONCUR 2019, August 27-30, 2019, Amsterdam, the Netherlands*, volume 140 of *LIPICs*, pages 13:1–13:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. *5 citations in pages 22, 32, 185, 240 and 253*
- [BBG⁺20] Thomas Brihaye, Véronique Bruyère, Aline Goeminne, Jean-François Raskin, and Marie van den Bogaard. The complexity of subgame perfect equilibria in quantitative reachability games. *Log. Methods Comput. Sci.*, 16(4), 2020. *2 citations in pages 22 and 185*
- [BBGR18] Thomas Brihaye, Véronique Bruyère, Aline Goeminne, and Jean-François Raskin. Constrained existence problem for weak subgame perfect equilibria with ω -regular boolean objectives. In Andrea Orlandini and Martin Zimmermann, editors, *Proceedings Ninth International Symposium on Games, Automata, Logics, and Formal Verification, GandALF 2018, Saarbrücken, Germany, 26-28th September 2018*, volume 277 of *EPTCS*, pages 16–29, 2018. *10 citations in pages 22, 28, 32, 52, 53, 153, 239, 240, 252 and 287*
- [BBGR20] Thomas Brihaye, Véronique Bruyère, Aline Goeminne, and Jean-François Raskin. Constrained existence problem for weak subgame perfect equilibria with ω -regular boolean objectives. *Information and Computation*, page 104594, 2020. *Cited in page 22*
- [BBGT19] Thomas Brihaye, Véronique Bruyère, Aline Goeminne, and Nathan Thomasset. On relevant equilibria in reachability games. In Emmanuel Filiot, Raphaël M. Jungers, and Igor Potapov, editors, *Reachability Problems - 13th International Conference, RP 2019, Brussels, Belgium, September 11-13, 2019, Proceedings*, volume 11674 of *Lecture Notes in Computer Science*, pages 48–62. Springer, 2019. *3 citations in pages 22, 32 and 215*
- [BBM10] Patricia Bouyer, Romain Brenguier, and Nicolas Markey. Nash equilibria for reachability objectives in multi-player timed games.

In *CONCUR 2010*, pages 192–206, 2010.

2 citations in pages 238 and 240

- [BBMR15] Thomas Brihaye, Véronique Bruyère, Noémie Meunier, and Jean-François Raskin. Weak subgame perfect equilibria and their application to quantitative reachability. In Stephan Kreutzer, editor, *24th EACSL Annual Conference on Computer Science Logic, CSL 2015, September 7-10, 2015, Berlin, Germany*, volume 41 of *LIPICs*, pages 504–518. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015.

7 citations in pages 23, 26, 27, 71, 73, 75 and 76

- [BBMU15] Patricia Bouyer, Romain Brenguier, Nicolas Markey, and Michael Ummels. Pure Nash equilibria in concurrent deterministic games. *Logical Methods in Computer Science*, 11(2), 2015.

2 citations in pages 161 and 295

- [BCD⁺07] Gerd Behrmann, Agnès Cougnard, Alexandre David, Emmanuel Fleury, Kim Guldstrand Larsen, and Didier Lime. Uppaal-tiga: Time for playing games! In Werner Damm and Holger Hermanns, editors, *CAV 2007*, volume 4590 of *Lecture Notes in Computer Science*, pages 121–125. Springer, 2007. *Cited in page 238*

- [BDS13] Thomas Brihaye, Julie De Pril, and Sven Schewe. Multiplayer cost games with simple Nash equilibria. In *LFCS Proceedings*, volume 7734 of *Lecture Notes in Computer Science*, pages 59–73. Springer, 2013. *Cited in page 27*

- [Ber07] Dietmar Berwanger. Admissibility in infinite games. In *STACS*, volume 4393 of *LNCS*, pages 188–199. Springer, 2007.

Cited in page 28

- [BFL⁺18] Patricia Bouyer, Uli Fahrenberg, Kim Guldstrand Larsen, Nicolas Markey, Joël Ouaknine, and James Worrell. Model checking real-time systems. In Clarke et al. [CHVB18], pages 1001–1046.

Cited in page 265

- [BG20] Thomas Brihaye and Aline Goeminne. On subgame perfect equilibria in turn-based reachability timed games. In Nathalie Bertrand and Nils Jansen, editors, *Formal Modeling and Analysis of Timed Systems - 18th International Conference, FORMATS 2020, Vienna, Austria, September 1-3, 2020, Proceedings*, volume 12288 of *Lecture Notes in Computer Science*, pages 94–110. Springer, 2020. *2 citations in pages 22 and 237*
- [BGHM17] Thomas Brihaye, Gilles Geeraerts, Axel Haddad, and Benjamin Monmege. Pseudopolynomial iterative algorithm to solve total-payoff games and min-cost reachability games. *Acta Inf.*, 54(1):85–125, 2017. *Cited in page 226*
- [BHR18] Véronique Bruyère, Quentin Hautem, and Jean-François Raskin. Parameterized complexity of games with monotonically ordered omega-regular objectives. In *CONCUR 2018*, volume 118 of *LIPICs*, pages 29:1–29:16. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018. *Cited in page 54*
- [BJP⁺18] Nicolas Basset, Ismaël Jecker, Arno Pauly, Jean-François Raskin, and Marie van den Bogaard. Beyond admissibility: Dominance between chains of strategies. In *CSL*, volume 119 of *LIPICs*, pages 10:1–10:22. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018. *Cited in page 28*
- [BK08] Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT press, 2008. *2 citations in pages 20 and 244*
- [BMOU11] Patricia Bouyer, Nicolas Markey, Jörg Olschewski, and Michael Ummels. Measuring permissiveness in parity games: Mean-payoff parity games revisited. In Tefik Bultan and Pao-Ann Hsiung, editors, *Automated Technology for Verification and Analysis, 9th International Symposium, ATVA 2011, Taipei, Taiwan, October 11-14, 2011. Proceedings*, volume 6996 of *Lecture Notes in Computer Science*, pages 135–149. Springer, 2011. *Cited in page 270*
- [BMR14] Véronique Bruyère, Noémie Meunier, and Jean-François Raskin. Secure equilibria in weighted games. In Thomas A. Henzinger and

Dale Miller, editors, *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, Vienna, Austria, July 14 - 18, 2014*, pages 26:1–26:26. ACM, 2014.

3 citations in pages 23, 26 and 269

[BMS14] Patricia Bouyer, Nicolas Markey, and Daniel Stan. Mixed Nash equilibria in concurrent terminal-reward games. In *FSTTCS Proceedings*, volume 29 of *Leibniz International Proceedings in Informatics*, pages 351–363. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2014. *Cited in page 27*

[BPRS16] Romain Brenguier, Guillermo A. Pérez, Jean-Francois Raskin, and Ocan Sankur. Admissibility in Quantitative Graph Games. In *FSTTCS*, volume 65 of *LIPICs*, pages 42:1–42:14, 2016.

Cited in page 28

[BR15] Romain Brenguier and Jean-François Raskin. Pareto curves of multidimensional mean-payoff games. In *CAV Proceedings*, volume 9207 of *Lecture Notes in Computer Science*, pages 251–267. Springer, 2015. *Cited in page 27*

[Bre12] Romain Brenguier. *Nash equilibria in concurrent games : application to timed games*. Theses, École normale supérieure de Cachan - ENS Cachan, November 2012.

8 citations in pages 26, 238, 239, 240, 263, 264, 265 and 271

[BRPR17] Véronique Bruyère, Stéphane Le Roux, Arno Pauly, and Jean-François Raskin. On the existence of weak subgame perfect equilibria. In Javier Esparza and Andrzej S. Murawski, editors, *Foundations of Software Science and Computation Structures - 20th International Conference, FOSSACS 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings*, volume 10203 of *Lecture Notes in Computer Science*, pages 145–161, 2017. *5 citations in pages 26, 27, 75, 288 and 297*

- [BRS14] Romain Brenguier, Jean-François Raskin, and Mathieu Sassolas. The complexity of admissibility in omega-regular games. In *CSL-LICS*, pages 23:1–23:10. ACM, 2014. *Cited in page 28*
- [BRS15] Romain Brenguier, Jean-François Raskin, and Ocan Sankur. Assume-admissible synthesis. In *CONCUR*, LIPIcs 42, pages 100–113. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2015. *Cited in page 28*
- [BRS17] Nicolas Basset, Jean-François Raskin, and Ocan Sankur. Admissible strategies in timed games. In Luca Aceto, Giorgio Bacci, Giovanni Bacci, Anna Ingólfssdóttir, Axel Legay, and Radu Mardare, editors, *Models, Algorithms, Logics and Tools - Essays Dedicated to Kim Guldstrand Larsen on the Occasion of His 60th Birthday*, volume 10460 of *Lecture Notes in Computer Science*, pages 403–425. Springer, 2017. *Cited in page 271*
- [Bru17] Véronique Bruyère. Computer aided synthesis: A game-theoretic approach. In *DLT*, volume 10396 of *LNCS*, pages 3–35. Springer, 2017. *7 citations in pages 21, 23, 53, 54, 99, 239 and 269*
- [BRvdB21] Léonard Brice, Jean-François Raskin, and Marie van den Bogaard. Subgame-perfect equilibria in mean-payoff games. *CoRR*, abs/2101.10685, 2021. *Cited in page 26*
- [CDF⁺05] Franck Cassez, Alexandre David, Emmanuel Fleury, Kim Guldstrand Larsen, and Didier Lime. Efficient on-the-fly algorithms for the analysis of timed games. In Martín Abadi and Luca de Alfaro, editors, *CONCUR 2005*, volume 3653 of *Lecture Notes in Computer Science*, pages 66–80. Springer, 2005. *Cited in page 238*
- [CFGR16] Rodica Condurache, Emmanuel Filiot, Raffaella Gentilini, and Jean-François Raskin. The complexity of rational synthesis. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *ICALP 2016*, volume 55 of *LIPIcs*, pages 121:1–121:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. *6 citations in pages 25, 27, 28, 216, 226 and 239*

- [CHJ04] Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdzinski. Games with secure equilibria. In Frank S. de Boer, Marcello M. Bonsangue, Susanne Graf, and Willem P. de Roever, editors, *Formal Methods for Components and Objects, Third International Symposium, FMCO 2004, Leiden, The Netherlands, November 2 - 5, 2004, Revised Lectures*, volume 3657 of *Lecture Notes in Computer Science*, pages 141–161. Springer, 2004.
Cited in page 21
- [CHP07] Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. Generalized parity games. In *FOSSACS*, volume 4423 of *LNCS*, pages 153–167. Springer, 2007.
Cited in page 55
- [CHVB18] Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, editors. *Handbook of Model Checking*. Springer, 2018.
2 citations in pages 54 and 275
- [CS03] Vincent Conitzer and Tüomas Sandholm. Complexity results about Nash equilibria. In *IJCAI Proceedings*, pages 765–771. Morgan Kaufmann Publishers Inc., 2003.
Cited in page 27
- [dAF07] Luca de Alfaro and Marco Faella. An accelerated algorithm for 3-color parity games with an application to timed games. In *CAV*, volume 4590 of *LNCS*, pages 108–120. Springer, 2007.
Cited in page 54
- [dAFH⁺03] Luca de Alfaro, Marco Faella, Thomas A. Henzinger, Rupak Majumdar, and Mariëlle Stoelinga. The element of surprise in timed games. In *CONCUR 2003*, pages 142–156, 2003.
Cited in page 266
- [De 13] Julie De Pril. *Equilibria in Multiplayer Cost Games*. PhD thesis, UMONS, 2013.
6 citations in pages 32, 63, 66, 69, 229 and 269
- [EL87] E. Allen Emerson and Chin-Laung Lei. Modalities for model checking: Branching time logic strikes back. *Science of Computer Programming*, 8(3):275–306, 1987.
Cited in page 53

- [FGR18] Emmanuel Filiot, Raffaella Gentilini, and Jean-François Raskin. Rational synthesis under imperfect information. In *LICS*, pages 422–431. ACM, 2018. *Cited in page 28*
- [FH13] Nathanaël Fijalkow and Florian Horn. Les jeux d’accessibilité généralisée. *Technique et Science Informatiques*, 32(9-10):931–949, 2013. *Cited in page 54*
- [FKL10] Dana Fisman, Orna Kupferman, and Yoad Lustig. Rational synthesis. In *TACAS*, volume 6015 of *LNCS*, pages 190–204. Springer, 2010. *Cited in page 28*
- [FKM⁺10] János Flesch, Jeroen Kuipers, Ayala Mashiah-Yaakovi, Gijs Schoenmakers, Eilon Solan, and Koos Vrieze. Perfect-information games with lower-semicontinuous payoffs. *Mathematics of Operation Research*, 35:742–755, 2010. *3 citations in pages 23, 26 and 27*
- [FL83] Drew Fudenberg and David Levine. Subgame-perfect equilibria of finite- and infinite-horizon games. *Journal of Economic Theory*, 31:251–268, 1983. *2 citations in pages 26 and 69*
- [FP17] János Flesch and Arkadi Predtetchinski. A characterization of subgame-perfect equilibrium plays in borel games of perfect information. *Math. Oper. Res.*, 42(4):1162–1179, 2017. *Cited in page 23*
- [Goe20] Aline Goeminne. Constrained existence of weak subgame perfect equilibria in multiplayer büchi games. *Information Processing Letters*, 163:105996, 2020. *3 citations in pages 22, 32 and 153*
- [GTW02] Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *LNCS*. Springer, 2002. *5 citations in pages 21, 39, 53, 60 and 163*
- [GU08] Erich Grädel and Michael Ummels. Solution Concepts and Algorithms for Infinite Multiplayer Games. In *New Perspectives*

- on Games and Interaction*, volume 4, pages 151–178. Amsterdam University Press, 2008. *3 citations in pages 27, 69 and 244*
- [Har85] Christopher Harris. Existence and characterization of perfect equilibrium in games of perfect information. *Econometrica: Journal of the Econometric Society*, pages 613–628, 1985. *Cited in page 69*
- [Hor08] Florian Horn. Explicit Muller games are PTIME. In *FSTTCS*, volume 2 of *LIPICs*, pages 235–243. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2008. *Cited in page 54*
- [Imm81] Neil Immerman. Number of quantifiers is better than number of tape cells. *Journal of Computer and System Sciences*, 22:384–406, 1981. *Cited in page 169*
- [JLS07] Marcin Jurdzinski, François Laroussinie, and Jeremy Sproston. Model checking probabilistic timed automata with one or two clocks. In *TACAS 2007*, pages 170–184, 2007. *Cited in page 264*
- [JT07] Marcin Jurdzinski and Ashutosh Trivedi. Reachability-time games on timed automata. In Lars Arge, Christian Cachin, Tomasz Jurdzinski, and Andrzej Tarlecki, editors, *ICALP 2007*, volume 4596 of *Lecture Notes in Computer Science*, pages 838–849. Springer, 2007. *Cited in page 238*
- [KLST12] Miroslav Klimos, Kim G. Larsen, Filip Stefanak, and Jeppe Thaarup. Nash equilibria in concurrent priced games. In Adrian-Horia Dediu and Carlos Martín-Vide, editors, *Language and Automata Theory and Applications - 6th International Conference, LATA 2012, A Coruña, Spain, March 5-9, 2012. Proceedings*, volume 7183 of *Lecture Notes in Computer Science*, pages 363–376. Springer, 2012. *Cited in page 27*
- [KNP19] Marta Kwiatkowska, Gethin Norman, and David Parker. Verification and control of turn-based probabilistic real-time games. In *The Art of Modelling Computational Systems: A Journey from*

- Logic and Concurrency to Security and Privacy - Essays Dedicated to Catuscia Palamidessi on the Occasion of Her 60th Birthday*, pages 379–396, 2019. *Cited in page 238*
- [KPV14] Orna Kupferman, Giuseppe Perelli, and Moshe Y. Vardi. Synthesis with rational environments. In *EUMAS*, LNCS 8953, pages 219–235. Springer, 2014. *Cited in page 28*
- [Mar75] Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, September 1975. *Cited in page 59*
- [Nas50] J. F. Nash. Equilibrium points in n -person games. In *PNAS*, volume 36, pages 48–49. National Academy of Sciences, 1950. *3 citations in pages 21, 64 and 238*
- [OR94] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*, volume 1 of *MIT Press Books*. The MIT Press, December 1994. *5 citations in pages 26, 68, 71, 80 and 291*
- [Osb04] Martin J. Osborne. *An introduction to game theory*. Oxford Univ. Press, 2004. *Cited in page 238*
- [PP04] Dominique Perrin and Jean-Eric Pin. Infinite words: Automata, semigroups, logic and games. 01 2004. *Cited in page 46*
- [PR89] Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *Proceedings of the 16th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 179–190, 1989. *2 citations in pages 20 and 238*
- [Ran14] Mickael Randour. *Synthesis in Multi-Criteria Quantitative Games*. PhD thesis, UMONS, 2014. *Cited in page 270*
- [Sel65] Reinhard Selten. Spieltheoretische behandlung eines oligopolmodells mit nachfrageträgheit. *Zeitschrift für die gesamte Staatswissenschaft*, 1965. *2 citations in pages 21 and 68*
- [Sip13] Michael Sipser. *Introduction to the Theory of Computation*. Course Technology, Boston, MA, third edition, 2013. *Cited in page 168*

- [Tho95] Wolfgang Thomas. On the synthesis of strategies in infinite games. In Ernst W. Mayr and Claude Puech, editors, *STACS 95, 12th Annual Symposium on Theoretical Aspects of Computer Science, Munich, Germany, March 2-4, 1995, Proceedings*, volume 900 of *Lecture Notes in Computer Science*, pages 1–13. Springer, 1995. *2 citations in pages 21 and 164*
- [Umm05] Michael Ummels. Rational Behaviour and Strategy Construction in Infinite Multiplayer Games. Diploma thesis, RWTH Aachen, 2005. *5 citations in pages 25, 35, 170, 181 and 216*
- [Umm06] Michael Ummels. Rational behaviour and strategy construction in infinite multiplayer games. In S. Arun-Kumar and Naveen Garg, editors, *FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science, 26th International Conference, Kolkata, India, December 13-15, 2006, Proceedings*, volume 4337 of *Lecture Notes in Computer Science*, pages 212–223. Springer, 2006. *4 citations in pages 27, 240, 265 and 266*
- [Umm08] Michael Ummels. The complexity of Nash equilibria in infinite multiplayer games. In *FOSSACS*, volume 4962 of *LNCS*, pages 20–34. Springer, 2008. *5 citations in pages 21, 27, 154, 163 and 164*
- [vNM44] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, New Jersey, 1944. *Cited in page 20*
- [ZP96] Uri Zwick and Mike Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158(1-2):343–359, 1996. *Cited in page 21*

APPENDICES

APPENDIX A

CHARACTERIZATION OF THE GAIN PROFILES OF WEAK SPES IN MULTIPLAYER GAMES WITH PREFIX-INDEPENDENT QUALITATIVE OBJECTIVES

In this appendix, we provide a characterization of weak SPEs in multiplayer games with prefix-independent qualitative objectives. Contrary to the characterizations provided in Part II, this one is based on the gain profiles realizable by a weak SPE in the game. Initially, this approach allowed us to obtain the results presented in [BBGR18].

We choose to present these other characterizations and related results as they were presented in [BBGR18] even if it seems quite redundant with some notions already introduced in Part II. Notice that the notion of *good symbolic witness* is slightly different of that provided in Section 7.2 since in this appendix we only consider Boolean games with prefix-independent objectives.

In this section our aim is twofold: first, we characterize the set of possible gain profiles of weak SPEs and second, we show how it is possible to build a weak SPE given a set of lassoes with some “good properties”. Those characterizations work for Boolean games with *prefix-independent gain functions*. We make this hypothesis all along Appendix A.

A.1 Remove-Adjust procedure

Let (\mathcal{G}, v_0) be an initialized Boolean game with prefix-independent gain functions. The computation of the set of all the gain profiles of weak SPEs in (\mathcal{G}, v_0) is inspired by a fixpoint procedure explained in [BRPR17]. Each vertex v is *labeled* by a set of gain profiles $p \in \{0, 1\}^{|\Pi|}$. Initially, these gain profiles are those for which there exists a play in $\text{Plays}(v)$ with gain profile p . Then step by step, some gain profiles are removed for the labeling of v as soon as we are sure they cannot be the gain profile of $\sigma_{|h}$ in a subgame $(\mathcal{G}_{|h}, v)$ for some weak SPE σ .¹ When a fixpoint is reached, the labeling of the initial vertex v_0 exactly contains all the gain profiles of weak SPEs in (\mathcal{G}, v_0) . Hence, at each step of this procedure, the gain profiles labeling a vertex v are gain profiles of *potential* subgame outcomes of a weak SPE. Their number decreases until reaching a fixpoint.

We formally proceed as follows. For all $v \in V$, we define the initial labeling of v as:

$$\mathbf{P}_0(v) = \{p \in \{0, 1\}^{|\Pi|} \mid \text{there exists } \rho \in \text{Plays}(v) \text{ such that } \text{Gain}(\rho) = p\}.$$

Then for each step $k \in \mathbb{N} \setminus \{0\}$, we compute the set $\mathbf{P}_k(v)$ by alternating between two operations: *Remove* and *Adjust*. To this end, we need to introduce the notion of (p, k) -labeled play. Let p be a gain profile and k be a step, a play $\rho = \rho_0\rho_1\rho_2 \dots$ is (p, k) -labeled if for all $j \in \mathbb{N}$ we have $p \in \mathbf{P}_k(\rho_j)$, that is, ρ visits only vertices that are labeled by p at step k . We first give some intuition about the Remove-Adjust procedure and then give the definition.

We start with the Remove operation. Let p that labels vertex v . This means that it is the gain profile of a potential subgame outcome of a weak SPE that starts in v . Suppose that v is a vertex of Player i and v has a successor v' such that $p_i < p'_i$ for all p' labeling v' . Then p cannot be the gain profile of $\sigma_{|h}$ in the subgame $(\mathcal{G}_{|h}, v)$ for some weak SPE σ and some history h , otherwise Player i would have a profitable (one-shot) deviation by moving from v to v' in this subgame.

¹The value of h is not important since the gain functions are prefix independent. This is why we only focus on v and not on hv .

Let us now explain the Adjust operation. It may happen that for another vertex u having p in its labeling, all potential subgame outcomes of a weak SPE from u with gain profile p necessarily visit vertex v . As p has been removed from the labeling of v , these potential plays do no longer survive and p is also removed from the labeling of u by the Adjust operation.

Let us now formally define the Remove-Adjust procedure.

Definition A.1.1 (Remove-Adjust procedure). Let $k \in \mathbb{N} \setminus \{0\}$.

- If k is odd, process the *Remove* operation:
 - If for some $v \in V_i$ there exists $p \in \mathbf{P}_{k-1}(v)$ and $v' \in \text{Succ}(v)$ such that $p_i < p'_i$ for all $p' \in \mathbf{P}_{k-1}(v')$, then $\mathbf{P}_k(v) = \mathbf{P}_{k-1}(v) \setminus \{p\}$ and for all $u \neq v$, $\mathbf{P}_k(u) = \mathbf{P}_{k-1}(u)$.
 - If such a vertex v does not exist, then $\mathbf{P}_k(u) = \mathbf{P}_{k-1}(u)$ for all $u \in V$.
- If k is even, process the *Adjust* operation:
 - If some gain profile p was removed from $\mathbf{P}_{k-2}(v)$ (that is, $\mathbf{P}_{k-1}(v) = \mathbf{P}_{k-2}(v) \setminus \{p\}$), then
 - * For all $u \in V$ such that $p \in \mathbf{P}_{k-1}(u)$, check whether there still exists a $(p, k-1)$ -labeled play with gain profile p from u . If it is the case, then $\mathbf{P}_k(u) = \mathbf{P}_{k-1}(u)$, otherwise $\mathbf{P}_k(u) = \mathbf{P}_{k-1}(u) \setminus \{p\}$.
 - * For all $u \in V$ such that $p \notin \mathbf{P}_{k-1}(u)$: $\mathbf{P}_k(u) = \mathbf{P}_{k-1}(u)$.
 - Otherwise $\mathbf{P}_k(u) = \mathbf{P}_{k-1}(u)$ for all $u \in V$.

We can state the existence of a fixpoint of the sequences $(\mathbf{P}_k(v))_{k \in \mathbb{N}}$, $v \in V$, in the following meaning:

Proposition A.1.2 (Existence of a fixpoint). *There exists an even natural number $k^* \in \mathbb{N}$ such that for all $v \in V$, $\mathbf{P}_{k^*}(v) = \mathbf{P}_{k^*+1}(v) = \mathbf{P}_{k^*+2}(v)$.*

Proof. For all $v \in V$, the sequence $(\mathbf{P}_k(v))_{k \in \mathbb{N}}$ is nonincreasing because the Remove and Adjust operations never add a new gain profile. As each $\mathbf{P}_0(v)$ is finite (it contains at most $2^{|\Pi|}$ gain profiles), there exists a natural odd number $k^* + 1$ such that for all $v \in V$, $\mathbf{P}_{k^*}(v) = \mathbf{P}_{k^*+1}(v)$ during the Remove operation, and thus for all $v \in V$, $\mathbf{P}_{k^*+1}(v) = \mathbf{P}_{k^*+2}(v)$ during the Adjust operation. \square

Example A.1.3. We illustrate the different steps of the Remove-Adjust procedure on the example depicted in Figure 7.2, and we display the result of this computation in Table A.1. Initially, the sets $\mathbf{P}_0(v)$, $v \in V$, contains all gain profiles p such that there exists a play $\rho \in \text{Plays}(v)$ with $\text{Gain}(\rho) = p$. At step $k = 1$, we apply a Remove operation to $v = v_4$ (this is the only possible v): v is a vertex of Player $i = 2$ and v has a successor $v' = v_5$ such that $(0, 1) \in \mathbf{P}_0(v_5)$. Therefore $(0, 0)$ is removed from $\mathbf{P}_0(v_4)$ to get $\mathbf{P}_1(v_4)$. By definition of the Remove operation, the other sets $\mathbf{P}_0(u)$ are not modified and are thus equal to $\mathbf{P}_1(u)$. At step $k = 2$, we apply an Adjust operation. The only way to obtain the gain profile $(0, 0)$ from v_0 is by visiting v_4 with the play $v_0v_4v_6^\omega$. As there does not exist a $((0, 0), 1)$ -labeled play with gain profile $(0, 0)$ anymore, we have to remove $(0, 0)$ from $\mathbf{P}_1(v_0)$. The other sets $\mathbf{P}_1(v)$ remain unchanged. At step $k = 3$, the Remove operation removes gain profile $(1, 0)$ from $\mathbf{P}_2(v_0)$ due to the unique gain profile $(0, 1)$ in $\mathbf{P}_2(v_4)$. At step $k = 4$, the Adjust operation leaves all sets $\mathbf{P}_3(v)$ unchanged. Finally at step $k = 5$, the Remove operation also leaves all sets $\mathbf{P}_4(v)$ unchanged, and the fixpoint is reached. Therefore, we have $k^* = 4$.

Table A.1: Computation of the fixpoint on the example of Figure 7.2

	v_0	v_1	v_2	v_3	v_4	v_5	v_6
$\mathbf{P}_0(v)$	$\{(0, 0), (1, 0), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(0, 1)\}$	$\{(\mathbf{0}, \mathbf{0}), (0, 1)\}$	$\{(0, 1)\}$	$\{(0, 0)\}$
$\mathbf{P}_1(v)$	$\{(\mathbf{0}, \mathbf{0}), (1, 0), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 0)\}$
$\mathbf{P}_2(v)$	$\{(\mathbf{1}, \mathbf{0}), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 0)\}$
$\mathbf{P}_3(v)$	$\{(0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 0)\}$
$\mathbf{P}_4(v)$	$\{(0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 0)\}$

A.2 Characterization and good symbolic witness

The fixpoint $\mathbf{P}_{k^*}(v)$, $v \in V$, provides a characterization of the gain profiles of all weak SPEs as described in the following theorem. This result is in the spirit of the classical *Folk Theorem* which characterizes the gain profiles of all NEs in infinitely repeated games (see for instance [OR94]).

Theorem A.2.1 (Characterization). *Let (\mathcal{G}, v_0) be an initialized Boolean game with prefix-independent gain functions. Then there exists a weak SPE σ with gain profile p_0 in (\mathcal{G}, v_0) if and only if $p_0 \in \mathbf{P}_{k^*}(v_0)$.^a*

^aWe use notation $p_0 \in \{0, 1\}^{|\Pi|}$ to highlight that this is the gain profile of σ from vertex v_0 . It should not be confused with any component p_i , $i \in \Pi$, of a gain profile p .

The rest of this section is devoted to the proof of Theorem A.2.1. We begin with a lemma that states that if a given gain profile p survives at step k (where k is even) in the labeling of v , this means that there exists a play with gain profile p from v that only visits vertices also labeled by p .

Lemma A.2.2. *For all even k and in particular for $k = k^*$, p belongs to $\mathbf{P}_k(v)$ if and only if there exists a (p, k) -labeled play $\rho \in \text{Plays}(v)$ such that $\text{Gain}(\rho) = p$.*

Proof. (\Leftarrow) Suppose that there exists a (p, k) -labeled play $\rho = \rho_0\rho_1 \dots \in \text{Plays}(v)$ such that $\text{Gain}(\rho) = p$. By definition of a (p, k) -labeled play, we have $p \in \mathbf{P}_k(\rho_j)$ for all j , and so in particular for $j = 0$.

(\Rightarrow) Let us prove that if p belongs to $\mathbf{P}_k(v)$, then there exists a (p, k) -labeled play $\rho \in \text{Plays}(v)$ such that $\text{Gain}(\rho) = p$. We proceed by induction on k . For $k = 0$, the assertion is satisfied by definition of $\mathbf{P}_0(v)$ and because Gain_i is prefix-independent for all $i \in \Pi$.

Suppose that the assertion is true for an even k and let us prove that it remains true for $k + 2$. Let $p \in \mathbf{P}_{k+2}(v)$. As $\mathbf{P}_{k+2}(v) \subseteq \mathbf{P}_{k+1}(v) \subseteq \mathbf{P}_k(v)$, we have $p \in \mathbf{P}_k(v)$ and there exists a (p, k) -labeled play $\rho \in \text{Plays}(v)$ such

that $\text{Gain}(\rho) = p$ by induction hypothesis. In other words $p \in \mathbf{P}_k(\rho_j)$ for all j .

We suppose that there exists v' such that $\mathbf{P}_{k+2}(v') \neq \mathbf{P}_k(v')$ (the fixpoint is not reached), otherwise $p \in \mathbf{P}_{k+2}(\rho_j)$ for all j and ρ is also a $(p, k + 2)$ -labeled play. Therefore the Remove operation has removed some gain profile p' from one $\mathbf{P}_k(v')$ and the Adjust operation has possibly removed p' from some other $\mathbf{P}_k(u)$. If $p' \neq p$, then clearly p still belongs to each $\mathbf{P}_{k+2}(\rho_j)$ and ρ is again a $(p, k + 2)$ -labeled play. If $p' = p$, then $v' \neq v$ since $p \in \mathbf{P}_{k+2}(v)$ by hypothesis. Moreover, by the Adjust operation, this means that there exists a $(p, k + 1)$ -labeled play $\pi = \pi_0\pi_1 \dots$ with gain profile p from v which never visits v' . Let us show that π is also a $(p, k + 2)$ -labeled play, that is, $p \in \mathbf{P}_{k+2}(\pi_j)$ for all j . Each suffix $\pi_j\pi_{j+1} \dots$ of π is a $(p, k + 1)$ -labeled play with gain profile p thanks to prefix-independence of Gain. By the Adjust operation, it follows that $\mathbf{P}_{k+2}(\pi_j) = \mathbf{P}_{k+1}(\pi_j)$ for all j . This concludes the proof. \square

The proof of Theorem A.2.1 uses the concept of (good) symbolic witness defined hereafter but we begin with some intuition about it.

A symbolic witness \mathcal{P} is a *compact* representation of some finite-memory strategy profile σ in (\mathcal{G}, v_0) . It is a finite set of lassoes that represent some subgame outcomes of σ : the lasso $\rho^{(0, v_0)}$ of \mathcal{P} represents the outcome $\langle \sigma \rangle_{v_0}$, and each other lasso $\rho^{(i, v')}$ represents the subgame outcome $\langle \sigma \rangle_{hv'}$ for some particular histories $hv' \in \text{Hist}(v_0)$. The index i records that Player i can move from v (the last vertex of h) to v' (with the convention that $i = 0$ for the outcome $\langle \sigma \rangle_{v_0}$). When σ is a weak SPE, the related symbolic witness \mathcal{P} is good, that is, its lassoes avoid profitable one-shot deviations between them.

Let us now define it properly.

Definition A.2.3 (Symbolic witness). Let (\mathcal{G}, v_0) be an initialized Boolean game with prefix-independent gain functions. Let $I \subseteq (\Pi \cup \{0\}) \times V$ be the set

$$\mathcal{I} = \{(0, v_0)\} \cup \{(i, v') \mid \text{there exists } (v, v') \in E \text{ such that } v, v' \in \text{Succ}^*(v_0) \text{ and } v \in V_i\}.$$

A *symbolic witness* is a set $\mathcal{P} = \{\rho_{i,v} \mid (i, v) \in \mathcal{I}\}$ such that each $\rho^{(i,v)} \in \mathcal{P}$ is a lasso in A with $\text{First}(\rho^{(i,v)}) = v$ and with length bounded by $2 \cdot |V|^2$.

A symbolic witness has thus at most $|V| \cdot |\Pi| + 1$ lassoes (by definition of \mathcal{I}) with polynomial length.

Definition A.2.4 (Good symbolic witness). A symbolic witness \mathcal{P} is *good* if for all $\rho^{(j,u)}, \rho^{(i,v')} \in \mathcal{P}$, for all vertices $v \in \rho^{(j,u)}$ such that $v \in V_i$ and $v' \in \text{Succ}(v)$, we have $\text{Gain}_i(\rho^{(j,u)}) \geq \text{Gain}_i(\rho^{(i,v')})$.

The condition of Definition A.2.4 is depicted in Figure A.1.




Figure A.1: The condition of Definition A.2.4

Example A.2.5. A good symbolic witness for the weak SPE of Example 2.4.19 depicted in Figure 7.2 is already given in Example 7.4.4.

In Proposition A.2.6 below, we are going to prove that there exists a weak SPE if and only if there exists a good symbolic witness representing a finite-memory weak SPE with the same gain profile, and that the existence of this witness is equivalent to the non-emptiness of the fixpoint $\mathbf{P}_{k^*}(v)$, $v \in V$. In this way, we will prove Theorem A.2.1. We will see that the lassoes $\rho_{i,v}$ of a good symbolic witness can be constructed from (p, k^*) -labeled plays for well-chosen gain profiles $p \in \mathbf{P}_{k^*}(v)$.

Notice that in the second assertion of Proposition A.2.6, we not only ask that $p_0 \in \mathbf{P}_{k^*}(v_0)$ (as in Theorem A.2.1) but also that $\mathbf{P}_{k^*}(v) \neq \emptyset$ for all $v \in \text{Succ}^*(v_0)$. We will come back to this observation when we will prove Theorem A.2.1.

Proposition A.2.6. *Let (\mathcal{G}, v_0) be an initialized Boolean game with prefix-independent gain functions. The following assertions are equivalent:*

1. *There exists a weak SPE with gain profile p_0 in (\mathcal{G}, v_0) ;*
2. *$\mathbf{P}_{k^*}(v) \neq \emptyset$ for all $v \in \text{Succ}^*(v_0)$ and $p_0 \in \mathbf{P}_{k^*}(v_0)$;*
3. *There exists a good symbolic witness \mathcal{P} that contains a lasso $\rho^{(0, v_0)}$ with gain profile p_0 ;*
4. *There exists a finite-memory weak SPE σ with gain profile p_0 in (\mathcal{G}, v_0) such that the size of each strategy σ_i is in $\mathcal{O}(|V|^3 \cdot |\Pi|)$.*

Proof. We prove that $1 \Rightarrow 2 \Rightarrow 3 \Rightarrow 4 \Rightarrow 1$.

(1 \Rightarrow 2) Suppose that there exists a weak SPE σ with gain profile p_0 in (\mathcal{G}, v_0) . To show that $\mathbf{P}_{k^*}(v) \neq \emptyset$ for all $v \in \text{Succ}^*(v_0)$, let us prove by induction on k that

$$\text{Gain}(\langle \sigma_{\uparrow h} \rangle_v) \in \mathbf{P}_k(v) \text{ for all } hv \in \text{Hist}(v_0). \quad (\text{A.1})$$

For case $k = 0$, this is true by definition of $\mathbf{P}_0(v)$. Suppose that this assertion is satisfied for an even k . Let us prove that it remains true for $k + 2$ by showing that gain profile $p = \text{Gain}(\langle \sigma_{\uparrow h} \rangle_v) \in \mathbf{P}_k(v)$ can be removed neither from $\mathbf{P}_k(v)$ at step $k + 1$, nor from $\mathbf{P}_{k+1}(v)$ at step $k + 2$.

- Gain profile p cannot be removed from $\mathbf{P}_k(v)$ by the Remove operation at step $k + 1$. Otherwise, if $v \in V_i$, this means that there exists $v' \in \text{Succ}(v)$ such that

$$\forall p' \in \mathbf{P}_k(v'), p_i < p'_i. \quad (\text{A.2})$$

By induction hypothesis,

$$\text{Gain}(\langle \sigma_{\uparrow hv} \rangle_{v'}) \in \mathbf{P}_k(v'). \quad (\text{A.3})$$

To get a contradiction, we prove that in the subgame $(\mathcal{G}_{\uparrow h}, v)$ there exists a one-shot deviating strategy σ'_i from $\sigma_{i \uparrow h}$ that is a profitable

deviation for Player i . We define σ'_i that only differs from $\sigma_{i|h}$ on v : $\sigma'_i(v) = v'$. Therefore we get $\text{Gain}(h\langle\sigma'_i, \sigma_{-i|h}\rangle_v) = \text{Gain}(hv\langle\sigma_{|hv}\rangle_{v'})$. It follows by (A.2), (A.3), and prefix-independence of Gain_i that $\text{Gain}_i(h\langle\sigma_{|h}\rangle_v) = p_i < p'_i = \text{Gain}_i(hv\langle\sigma_{|hv}\rangle_{v'}) = \text{Gain}(h\langle\sigma'_i, \sigma_{-i|h}\rangle_v)$. This is impossible since σ is a weak SPE.

- Gain profile p cannot be removed from $\mathbf{P}_{k+1}(v)$ by the Adjust operation at step $k + 2$. Otherwise, this means that there exists u such that $\mathbf{P}_{k+1}(u) = \mathbf{P}_k(u) \setminus \{p\}$ (by the Remove operation at step $k + 1$) and there is no $(p, k + 1)$ -labeled play with gain profile p from v . However by Lemma A.2.2, as $p \in \mathbf{P}_k(v)$, there exists a (p, k) -labeled play π with gain profile p from v . This means that π visit u . Let $h'u \in \text{Hist}(v)$ such that $h'u < \pi$. Then we get a contradiction with σ being a weak SPE if we repeat the argument done before in the previous item for u and the subgame $(\mathcal{G}_{|hh'}, u)$ (instead of v and $(\mathcal{G}_{|h}, v)$).

Now that we know that $\mathbf{P}_{k^*}(v) \neq \emptyset$ for all $v \in \text{Succ}^*(v_0)$, it remains to prove that $p_0 \in \mathbf{P}_{k^*}(v_0)$. By (A.1), we have $p_0 = \text{Gain}(\langle\sigma\rangle_{v_0}) \in \mathbf{P}_{k^*}(v_0)$.

(2 \Rightarrow 3) Let us show how to build a good symbolic witness \mathcal{P} from the non-empty fixpoint $\mathbf{P}_{k^*}(v)$, $v \in V$. First recall that if $p \in \mathbf{P}_{k^*}(v)$, then by Lemma A.2.2 there exists a (p, k^*) -labeled play with gain profile p from v . Notice that such a play can be supposed to be a lasso with length at most $2 \cdot |V|^2$. Indeed it is proved in [BBMU15, Proposition 3.1] that given a play ρ , one can construct a lasso ρ' of length bounded by $2 \cdot |V|^2$ such that $\text{First}(\rho) = \text{First}(\rho')$, $\text{Occ}(\rho) = \text{Occ}(\rho')$, and $\text{Inf}(\rho) = \text{Inf}(\rho')$ (this construction eliminates some cycles of ρ in a clever way). Therefore, if ρ is a (p, k^*) -labeled play with gain profile p from v , the lasso ρ' is also a (p, k^*) -labeled play with gain profile p from v . The required set \mathcal{P} will be composed of some of these lassoes.

We start with $\mathcal{P} = \emptyset$. As $p_0 \in \mathbf{P}_{k^*}(v_0)$, then there exists a (p_0, k^*) -labeled lasso $\rho^{(0, v_0)}$ with gain profile p_0 from v_0 that we add to \mathcal{P} . For all $v, v' \in \text{Succ}^*(v_0)$ such that $v \in V_i$ and $v' \in \text{Succ}(v)$, let p' be a gain profile in $\mathbf{P}_{k^*}(v')$ such that

$$p'_i = \text{Min}\{q_i \mid q \in \mathbf{P}_{k^*}(v')\}. \quad (\text{A.4})$$

This gain profile exists since $\mathbf{P}_{k^*}(v') \neq \emptyset$ by hypothesis. Then there exists a (p', k^*) -labeled lasso $\rho^{(i, v')}$ with gain profile p' from v' that we add to \mathcal{P} .

This set \mathcal{P} is a symbolic witness by construction. It remains to prove that it is good. Let $v \in \rho^{(j, u)}$ such that $v \in V_i$ and $\rho^{(j, u)} \in \mathcal{P}$. As $\rho^{(j, u)}$ is a (p, k^*) -labeled lasso for some gain profile p , we have $p \in \mathbf{P}_{k^*}(v)$. Furthermore, as $\mathbf{P}_{k^*}(v) = \mathbf{P}_{k^*+1}(v)$ (by the fixpoint), this means that p was not removed from $\mathbf{P}_{k^*}(v)$ by the Remove operation at step k^* . In particular, by definition of the gain profile p' of $\rho_{i, v'}$ (see (A.4)), we have $p_i \geq p'_i$, that is $\text{Gain}_i(\rho^{(j, u)}) \geq \text{Gain}_i(\rho^{(i, v')})$. This shows that \mathcal{P} is a good symbolic witness.

(3 \Rightarrow 4) Let $\mathcal{P} = \{\rho^{(i, v)} \mid (i, v) \in \mathcal{I}\}$ be a good symbolic witness that contains a lasso $\rho^{(0, v_0)}$ with gain profile p_0 . We define a strategy profile σ step by step by induction on the subgames of (\mathcal{G}, v_0) . We first partially build σ such that $\langle \sigma \rangle_{v_0} = \rho_{0, v_0}$. Consider next $hvv' \in \text{Hist}(v_0)$ with $v \in V_i$ such that $\langle \sigma_{\upharpoonright h} \rangle_v$ is already built but not $\langle \sigma_{\upharpoonright hv} \rangle_{v'}$. Then we extend the definition of σ such that

$$\langle \sigma_{\upharpoonright hv} \rangle_{v'} = \rho^{(i, v')}. \quad (\text{A.5})$$

Notice that $\langle \sigma_{\upharpoonright h} \rangle_v$ being already built means that there exists $h' \leq h$ and $(j, u) \in \mathcal{I}$ such that

$$h' \langle \sigma_{\upharpoonright h'} \rangle_u = h' \rho^{(j, u)} = h \langle \sigma_{\upharpoonright h} \rangle_v. \quad (\text{A.6})$$

Let us prove that σ is a very weak SPE (and so a weak SPE by Proposition 2.4.18). Consider the subgame $(\mathcal{G}_{\upharpoonright h}, v)$ and the one-shot deviating strategy σ'_i from $\sigma_{i \upharpoonright h}$ such that $\sigma'_i(v) = v'$. We have to prove that

$$\text{Gain}_i(h \langle \sigma_{\upharpoonright h} \rangle_v) \geq \text{Gain}_i(hv \langle \sigma_{\upharpoonright hv} \rangle_{v'}). \quad (\text{A.7})$$

By (A.5), (A.6), and prefix independence of Gain_i , we have

$$\begin{aligned} \text{Gain}_i(hv \langle \sigma_{\upharpoonright hv} \rangle_{v'}) &= \text{Gain}_i(\rho^{(i, v')}), \\ \text{Gain}_i(h \langle \sigma_{\upharpoonright h} \rangle_v) &= \text{Gain}_i(\rho^{(j, u)}). \end{aligned}$$

Inequality (A.7) follows from these equalities and the fact that \mathcal{P} is a good symbolic witness (see Figure 7.1).

Notice that σ has gain profile p_0 by construction ($\rho^{(0,v_0)}$ has gain profile p_0). It remains to show that σ is finite-memory. Having (j, u) in memory (the last deviating player j and the vertex u where he moved), the Moore machines \mathcal{M}_i , $i \in \Pi$, representing each strategy σ_i , have to produce together the lasso $\rho^{(j,u)}$ of length bounded by $2 \cdot |V|^2$. As $|\mathcal{P}| = |\mathcal{I}| \leq |V| \cdot |\Pi| + 1$, the size of each σ_i is in $\mathcal{O}(|\Pi| \cdot |V|^3)$.

(4 \Rightarrow 1) This implication is obvious. \square

We are now able to prove Theorem A.2.6.

Proof of Theorem A.2.1. If there exists a weak SPE with gain profile p_0 in (\mathcal{G}, v_0) , then $p_0 \in \mathbf{P}_{k^*}(v_0)$ by Proposition A.2.6. It is proved in [BRPR17] that each initialized game with prefix-independent gain functions always has a weak SPE. Therefore $\mathbf{P}_{k^*}(v) \neq \emptyset$ for each $v \in \text{Succ}^*(v_0)$ by Proposition A.2.6. The latter assertion together with the assertion that $p_0 \in \mathbf{P}_{k^*}(v_0)$ implies the existence of a weak SPE with gain profile p_0 in (\mathcal{G}, v_0) . \square

APPENDIX B

TECHNICAL PROOFS

In this Appendix, we provide some of the proofs of Part III and Part IV.

B.1 Proofs of Chapter 11

B.1.1 Proofs of Section 11.1.2

Proof of Proposition 11.1.7

To prove Proposition 11.1.7, we have to prove that the sequences $(\lambda^k(v))_{k \in \mathbb{N}}$, with $v \in V^X$, are non increasing.

Lemma B.1.1. *For all $v \in V^X$, the sequences $(\lambda^k(v))_{k \in \mathbb{N}}$ and $(\Lambda^k(v))_{k \in \mathbb{N}}$ are non increasing.*

Proof. Let us prove by induction on k that for all $v \in V^X$

$$\lambda^{k+1}(v) \leq \lambda^k(v). \tag{B.1}$$

We will get that $\Lambda^{k+1}(v) \subseteq \Lambda^k(v)$.

First, recall that $\lambda^k(v) = 0$ if and only if $i \in I(v)$, where i is the player owning

v . So, in this case, $\lambda^{k+1}(v) = \lambda^k(v) = 0$ and assertion (B.1) is proved. It remains to prove this assertion when $i \notin I(v)$.

For $k = 0$, let $v \in V_i^X$ such that $i \notin I(v)$, then $\lambda^0(v) = +\infty$ and obviously $\lambda^1(v) \leq \lambda^0(v)$.

Suppose that assertion (B.1) is true for k and let us prove it for $k + 1$. We know by induction hypothesis that for all $v' \in V^X$, $\lambda^{k+1}(v') \leq \lambda^k(v')$, and thus also

$$\Lambda^{k+1}(v') \subseteq \Lambda^k(v'). \quad (\text{B.2})$$

Let us prove that for all $v \in V_i^X$ such that $i \notin I(v)$, $\lambda^{k+2}(v) \leq \lambda^{k+1}(v)$. If $\lambda^{k+2}(v) = \lambda^{k+1}(v)$ or $\lambda^{k+1}(v) = +\infty$, then the assertion is proved. Otherwise,

$$\lambda^{k+1}(v) = 1 + \min_{(v,v') \in E^X} \sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v')\}.$$

By (B.2), it follows that

$$\lambda^{k+1}(v) \geq 1 + \min_{(v,v') \in E^X} \sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^{k+1}(v')\} = \lambda^{k+2}(v).$$

And so, the assertion again holds. \square

We can now prove Proposition 11.1.7.

Proof of Proposition 11.1.7. The base case is easily proved. Indeed, as V^{J_N} is a bottom region, we have $\lambda_1(v) = \lambda_0(v)$ for all $v \in V^{J_N}$, and thus the local fixpoint is immediately reached on V^{J_N} . Hence with $k_N^* = 0$, for all $m \in \mathbb{N}$ and all $v \in V^{J_N}$, $\lambda^{k_N^*+m}(v) = \lambda^{k_N^*}(v)$.

Let J_n be an element of \mathcal{I} , with $n \in \{1, \dots, N-1\}$. Suppose that a fixpoint has been reached in the arena $X^{\geq J_{n+1}}$ and that the labeling function λ^k is updated on the arena $X^{\geq J_n}$ as described in Definition 11.1.5. Recall (as already summarized in Lemma 11.1.8) that the previously computed values of λ^k do no longer change on $X^{\geq J_{n+1}}$ (a local fixpoint is reached) and they do not change outside of $V^{\geq J_n}$ (by construction). However they can be modified on X^{J_n} . In this region X^{J_n} , there are $|V|$ sequences $(\lambda^k(v))_{k \in \mathbb{N}}$, $v \in V^{J_n}$, since the vertices v are of the form $v = (u, J_n)$ where J_n is fixed. These sequences are non increasing by Lemma B.1.1. As the component-wise ordering over

$(\mathbb{N} \cup \{+\infty\})^{|V|}$ is a well-quasi-ordering, there exists a natural number k_n^* that we choose as small as possible such that for all $v \in V^{J_n}$, $\lambda^{k_n^*+1}(v) = \lambda^{k_n^*}(v)$. This equality also holds for all $v \in V^{\geq J_n}$ (and not only for V^{J_n}), and it follows that for all $m \in \mathbb{N}$ and all $v \in V^{\geq J_n}$, $\lambda^{k_n^*+m}(v) = \lambda^{k_n^*}(v)$.

Notice that when n is decremented in Algorithm 4, k is incremented at least once showing that the sequence $0 < k_N^* < k_{N-1}^* < \dots < k_1^* = k^*$ is strictly increasing.

Finally, the last arena processed by the algorithm is $X^{\geq J_1} = X$. So with $k^* = k_1^*$, we have that for all $v \in V$ and all $m \in \mathbb{N}$, $\lambda^{k^*+m}(v) = \lambda^{k^*}(v)$. \square

Proof of Proposition 11.1.10

Proof of Proposition 11.1.10. We assume that $\sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v)\} = +\infty$, that is, for all $n \in \mathbb{N}$, there exists $\rho^n \in \Lambda^k(v)$ such that $\text{Cost}_i(\rho^n) > n$. By König's lemma, there exist $(\rho^{n_\ell})_{\ell \in \mathbb{N}}$ a subsequence of $(\rho^n)_{n \in \mathbb{N}}$ and $\rho \in \text{Plays}_X(v)$ such that $\rho = \lim_{\ell \rightarrow +\infty} \rho^{n_\ell}$. Moreover, $\text{Cost}_i(\rho) = +\infty$ as Cost_i is a continuous function. Let us prove that

$$\rho \in \Lambda^k(v).$$

This will establish Proposition 11.1.10.

We prove by induction on t with $0 \leq t \leq k$ that $\rho \in \Lambda^t(v)$. If $t = 0$, then $\Lambda^0(v) = \text{Plays}_X(v)$ by Lemma 11.1.4. As $\rho = \lim_{\ell \rightarrow +\infty} \rho^{n_\ell} \in \text{Plays}_X(v)$, it follows that $\rho \in \Lambda^0(v)$.

Let $t > 0$ and assume that the assertion is true for $t - 1 < k$. Suppose by contradiction that $\rho \notin \Lambda^t(v)$. For all $\ell \in \mathbb{N}$, as $\rho^{n_\ell} \in \Lambda^k(v)$, we have $\rho^{n_\ell} \in \Lambda^t(v)$ by Lemma B.1.1. Moreover by induction hypothesis we have $\rho \in \Lambda^{t-1}(v)$. From $\rho \in \Lambda^{t-1}(v) \setminus \Lambda^t(v)$, it follows that there exists $m \in \mathbb{N}$ and $j \in \Pi$ with $\rho_m \in V_j$ such that

$$\text{Cost}_j(\rho_{\geq m}) > \lambda^t(\rho_m) \quad \text{and} \quad \text{Cost}_j(\rho_{\geq m}) \leq \lambda^{t-1}(\rho_m).$$

In particular, $\lambda^t(\rho_m) < \lambda^{t-1}(\rho_m)$ and thus $\lambda^t(\rho_m) < +\infty$, and player j does not reach his target set along $\pi = \rho_0 \dots \rho_m \dots \rho_{m+\lambda^t(\rho_m)}$. We choose n_ℓ large

enough such that ρ and $\rho^{n\ell}$ share a common prefix of length at least $|\pi|$. As $\text{Cost}_j(\rho_{\geq m}) > \lambda^t(\rho_m)$, it follows that $\text{Cost}_j(\rho_{\geq m}^{\ell n}) > \lambda^t(\rho_m)$. We can conclude that $\rho^{n\ell} \notin \Lambda^t(v)$ which leads to a contradiction. \square

Proof of Theorem 11.1.9

The two implications of Theorem 11.1.9 are proved in the following Propositions B.1.2 and B.1.3. Notice that in Proposition B.1.2, we derive the additional property that $\Lambda^*(v) \neq \emptyset$, for all $v \in \text{Succ}^*(x_0)$. This is necessary to prove Proposition B.1.3.

Proposition B.1.2. *If σ is an SPE in (\mathcal{X}, x_0) then for all $v \in \text{Succ}^*(x_0)$, $\Lambda^*(v) \neq \emptyset$ and $\langle \sigma \rangle_{x_0} \in \Lambda^*(x_0)$.*

Proof. Suppose that σ is an SPE in (\mathcal{X}, x_0) and let us prove by induction on k that for all $k \in \mathbb{N}$ and all $hv \in \text{Hist}_X(x_0)$,

$$\langle \sigma \upharpoonright h \rangle_v \in \Lambda^k(v).$$

For case $k = 0$, this is true by definition of $\Lambda^0(v)$ and Lemma 11.1.4.

Suppose that this assertion is satisfied for $k \geq 0$ and by contradiction, assume that there exists $hv \in \text{Hist}_X(x_0)$ such that $\langle \sigma \upharpoonright h \rangle_v \notin \Lambda^{k+1}(v)$. Let $\rho = \langle \sigma \upharpoonright h \rangle_v$, by Definition 5.1.2, it means that there exist $n \in \mathbb{N}$ and $i \in \Pi$ such that $\rho_n \in V_i$ and

$$\text{Cost}_i(\rho_{\geq n}) > \lambda^{k+1}(\rho_n). \quad (\text{B.3})$$

But by induction hypothesis, we know that

$$\text{Cost}_i(\rho_{\geq n}) \leq \lambda^k(\rho_n). \quad (\text{B.4})$$

It follows by (B.3) and (B.4) that $\lambda^{k+1}(\rho_n) < \lambda^k(\rho_n)$ and that

$$\lambda^{k+1}(\rho_n) < +\infty, \lambda^k(\rho_n) \neq 0 \text{ and } \lambda^{k+1}(\rho_n) \neq 0.$$

In regards of Definition 11.1.5, these three relations allow us to conclude that

$$\lambda^{k+1}(\rho_n) = 1 + \min_{(\rho_n, v') \in E^X} \sup\{\text{Cost}_i(\rho') \mid \rho' \in \Lambda^k(v')\}.$$

Let $v' \in V$ be such that

$$\lambda^{k+1}(\rho_n) - 1 = \sup\{\text{Cost}_i(\rho') \mid \rho' \in \Lambda^k(v')\}. \quad (\text{B.5})$$

Let $h' = h\rho_0 \dots \rho_{n-1} \in \text{Hist}_X(x_0)$, and let us define the one-shot deviating strategy τ_i from $\sigma_{i|h'}$ such that $\tau_i(\rho_n) = v'$. Let us prove that τ_i is a one-shot profitable deviation for player i in $(\mathcal{X}_{|h'}, \rho_n)$.

$$\begin{aligned} & \text{Cost}_i(h' \langle \tau_i, \sigma_{-i|h'} \rangle_{\rho_n}) \\ &= \text{Cost}_i(h' \rho_n \langle \sigma_{|h' \rho_n} \rangle_{v'}) \quad (\text{as } \tau_i \text{ is a one-shot deviating strategy}) \\ &= |h' \rho_n v'| + \text{Cost}_i(\langle \sigma_{|h' \rho_n} \rangle_{v'}) \quad (\text{as } \lambda^{k+1}(\rho_n) \neq 0) \\ &\leq |h' \rho_n v'| + \lambda^{k+1}(\rho_n) - 1 \quad (\text{by (B.5) and as } \langle \sigma_{|h' \rho_n} \rangle_{v'} \in \Lambda^k(v') \text{ by IH}) \\ &< |h' \rho_n| + \text{Cost}_i(\rho_{\geq n}) \quad (\text{by (B.3)}) \\ &= \text{Cost}_i(h' \langle \sigma_{|h'} \rangle_{\rho_n}) \end{aligned}$$

This proves that σ is not a very weak SPE and so not an SPE by Corollary 2.4.23, which is a contradiction. This concludes the proof. \square

Proposition B.1.3. *Let $\rho^0 \in \Lambda^*(x_0)$, then ρ^0 is the outcome of an SPE in (\mathcal{X}, x_0) .*

Proof. By Theorem 2.4.9, there exists an SPE in (\mathcal{X}, x_0) and thus, by Proposition B.1.2:

$$\Lambda^*(v) \neq \emptyset \text{ for all } v \in \text{Succ}^*(x_0). \quad (\text{B.6})$$

Let $\rho^0 \in \Lambda^*(x_0)$ and let us show how to construct a very weak SPE σ (and so an SPE by Corollary 2.4.23) with outcome ρ^0 in (\mathcal{X}, x_0) . We define σ step by step by induction on the subgames of (\mathcal{X}, x_0) . We first partially build σ such as it produces ρ^0 , i.e., $\langle \sigma \rangle_{x_0} = \rho^0$. Now, we define a set of plays which is useful to define σ in the subgames. For all (i, v') such that $(v, v') \in E^X$ with $v, v' \in \text{Succ}^*(x_0)$ and $v \in V_i^X$, we take $\rho_{i,v'} \in \Lambda^*(v')$ such that:

$$\text{Cost}_i(\rho_{i,v'}) = \max\{\text{Cost}_i(\rho') \mid \rho' \in \Lambda^*(v')\} \quad (\text{B.7})$$

Notice that such a play exists by (B.6) and Corollary 11.1.11. The construction of σ is done by induction as follows. Consider $hvv' \in \text{Hist}_X(x_0)$ with $v \in V_i^X$ such that $\langle \sigma_{\uparrow h} \rangle_v$ is already defined but not yet $\langle \sigma_{\uparrow hv} \rangle_{v'}$. We extend the definition of σ as follows:

$$\langle \sigma_{\uparrow hv} \rangle_{v'} = \rho_{i,v'}$$

Let us prove that σ is a very weak SPE. Consider the subgame $(\mathcal{X}_{\uparrow h}, v)$ for a given $hv \in \text{Hist}_X(x_0)$ with $v \in V_i^X$ and the one-shot deviating strategy σ'_i from $\sigma_{i|h}$ such that $\sigma'_i(v) = v'$. By construction, there exists $\rho_{j,u}$ and $\rho_{i,v'}$ as defined previously and $g \in \text{Hist}_X(x_0)$ such that $h\langle \sigma_{\uparrow h} \rangle_v = g\rho_{j,u}$ and $\langle \sigma_{\uparrow hv} \rangle_{v'} = \rho_{i,v'}$.

We prove that σ'_i is not a one-shot profitable deviation by proving that $\text{Cost}_i(h\langle \sigma_{\uparrow h} \rangle_v) \leq \text{Cost}_i(hv\langle \sigma_{\uparrow hv} \rangle_{v'})$. If $i \in I(v)$, then obviously $\text{Cost}_i(h\langle \sigma_{\uparrow h} \rangle_v) = \text{Cost}_i(hv\langle \sigma_{\uparrow hv} \rangle_{v'})$. Otherwise $i \notin I(v)$. We have that $\rho = \langle \sigma_{\uparrow h} \rangle_v$ is suffix of $\rho_{j,u}$, that is, $\rho = \rho_{j,u,\geq n}$ for some $n \in \mathbb{N}$ and that $\rho_{j,u} \in \Lambda^*(u) = \Lambda^{k^*}(u) = \Lambda^{k^*+1}(u)$ (by (B.7) and the fixpoint). It follows that:

$$\begin{aligned} \text{Cost}_i(\rho) &= \text{Cost}_i(\rho_{j,u,\geq n}) \\ &\leq \lambda^{k^*+1}(v) \\ &= 1 + \min_{(v,w) \in E^X} \sup\{\text{Cost}_i(\rho') \mid \rho' \in \Lambda^{k^*}(w)\} \\ &\leq 1 + \sup\{\text{Cost}_i(\rho') \mid \rho' \in \Lambda^{k^*}(v')\} \\ &= 1 + \text{Cost}_i(\rho_{i,v'}). \end{aligned} \tag{B.8}$$

Thus we have that

$$\begin{aligned} \text{Cost}_i(h\langle \sigma_{\uparrow h} \rangle_v) &= |hv| + \text{Cost}_i(\rho) \\ &\leq |hv| + 1 + \text{Cost}_i(\rho_{i,v'}) \quad (\text{by (B.8)}) \\ &= \text{Cost}_i(hv\rho_{i,v'}) \\ &= \text{Cost}_i(hv\langle \sigma_{\uparrow hv} \rangle_{v'}). \end{aligned}$$

This concludes the proof. \square

B.1.2 Proofs of Section 11.2

Proofs of Lemma 11.2.5 and Lemma 11.2.6

Proof of Lemma 11.2.5. Let $v \in V^X$ and ρ be a play in \mathcal{X} such that $\rho \in \Lambda(v)$ and $\rho_0 = v$. We build the corresponding infinite path π in $\mathbb{C}(\lambda)$ iteratively as follows: Let $\pi_0 := v^C$. Let $n \in \mathbb{N}$, $n \geq 1$. Suppose that $\pi_{<n}$ has been already constructed, we show how to choose π_n . Suppose $\pi_{n-1} = (v', (c'_i)_{i \in \Pi})$. Then $\pi_n := (v'', (c''_i)_{i \in \Pi})$, where:

- $v'' = \rho_n$,
- for every $i \in \Pi$:

$$c''_i = \begin{cases} 0 & \text{if } i \in I(v''), \\ c'_i - 1 & \text{if } i \notin I(v'') \text{ and } v'' \notin V_i^X, \\ \min(c'_i - 1, \lambda(v'')) & \text{if } i \notin I(v'') \text{ and } v'' \in V_i^X \end{cases}$$

Let us show that $((v', (c'_i)_{i \in \Pi}), (v'', (c''_i)_{i \in \Pi})) \in E^C$. As ρ is a play in \mathcal{X} , we clearly have $(v', v'') \in E^X$. Assume now, towards contradiction, that $((v', (c'_i)_{i \in \Pi}), (v'', (c''_i)_{i \in \Pi})) \notin E^C$, that is, there is no outgoing edge from vertex $(v', (c'_i)_{i \in \Pi})$ in the counter graph. By Definition 11.2.2, this means that there exists a player i such that $c'_i = 1$ and $i \notin I(v'')$.

As $c'_i = 1$, by Definition 11.2.2 and construction of $\pi_0 \dots \pi_{n-1}$, there must exist a largest index $m \leq n - 1$ such that $\rho_m \in V_i^X$, $\lambda(\rho_m) = d$ with $d > 0$ and such that the counter value for player i decreases by exactly 1 at each step from vertex π_m until reaching value $c'_i = 1$ at vertex $\pi_{n-1} = v'$. Since ρ is λ -consistent, player i visits his target set along ρ in at most d steps, thus at most at vertex v'' , and thus $i \in I(v'')$, which is a contradiction. \square

Proof of Lemma 11.2.6. Let $v^C = (v, (c_i)_{i \in \Pi})$ be a starting vertex in $\text{SV}(\lambda)$. Let π be an infinite path in $\mathbb{C}(\lambda)$ such that $\pi_0 = v^C$. Let ρ be the projection of π on V^X , that is, $\rho_n = v'$ with $\pi_n = (v', (c'_i)_{i \in \Pi})$, for every $n \in \mathbb{N}$.

- Clearly, ρ is a play in \mathcal{X} : by construction of $\mathbb{C}(\lambda)$, there exists an edge

between two vertices $(v, (c_i)_{i \in \Pi})$ and $(v', (c'_i)_{i \in \Pi})$ in $\mathbb{C}(\lambda)$ only if (v, v') is an edge in E^X .

- Furthermore, the play ρ is λ -consistent: Assume, towards contradiction, that it is not. Thus, there exists $n \in \mathbb{N}$ and $i \in \Pi$ such that $\rho_n \in V_i^X$ and

$$\text{Cost}_i(\rho_{\geq n}) > \lambda(\rho_n). \quad (\text{B.9})$$

Consider now π_n and the value c_i of the counter for player i at this vertex. Since $\rho_n \in V_i^X$, we know that $c_i \leq \lambda(\rho_n)$. From this vertex, the counter value for player i decreases at least by 1 at each step along $\pi_{\geq n}$ and hits the value 0 before $\lambda(\rho_n)$ steps. However, this means that in $\rho_{\geq n}$, player i visits his target set sooner than expected, which is a contradiction with (B.9). \square

Proof of Proposition 11.2.7

Before proving Proposition 11.2.7, we need the following technical lemma:

Lemma B.1.4. *Let v^C be a starting vertex in $\text{SV}(\lambda)$ associated with $v \in V^X$ such that $I(v) = J_\ell$. Let π be a finite prefix of a valid path in $\mathbb{C}(\lambda)$ such that:*


- $\pi_0 = v^C$,
- π does not contain any cycle.

Then,

$$|\pi| \leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \sum_{r=|J_\ell|+1}^{|\Pi|} |V| + 2 \cdot \max_{\substack{J_j > J_\ell \\ |J_j|=r}} \text{mR}(\lambda_j)$$

Proof sketch. The proof of this lemma is quite technical, so we give here the main ideas and refer to the full proof hereafter for more details.

Let π be a finite prefix of a valid path in $\mathbb{C}(\lambda)$ as in the statement. Let

Figure B.1: Region decomposition of π

$\pi[\ell] \dots \pi[m]$ be its region decomposition according to Lemma 11.1.1, graphically represented in Figure B.1. Let ρ be the corresponding path in \mathcal{X} and $\rho[\ell] \dots \rho[m]$ be its region decomposition. Let us consider a fixed non-empty section $\pi[n]$.

Suppose first that the counter values at $\pi[n]_0$ are either 0 or $+\infty$. Let us prove that along $\pi[n]$, there can be at most $|V|$ steps before reaching a vertex with a finite positive value of λ :

- assume there is a cycle in the corresponding section $\rho[n]$ in \mathcal{X} such that from $\rho[n]_0$ and along the cycle, all the values of λ are either 0 or $+\infty$,
- by construction of $\mathbb{C}(\lambda)$, the counter values in the corresponding prefix of $\pi[n]$ remain fixed for each vertex of this prefix: as no value of λ is positive and finite, no counter value can be decremented,
- thus, the cycle in $\rho[n]$ is also a cycle in $\pi[n]$ which is impossible by hypothesis,
- thus there is no such cycle in $\rho[n]$, and as there are at most $|V|$ vertices in region X^{J_n} , $\rho[n]$ can have a prefix of length at most $|V|$ with only values 0 or $+\infty$ for λ , implying that this is also the case for $\pi[n]$.

Therefore, we can decompose $\pi[n]$ into a (possibly empty) prefix of length at most $|V|$, and a (possibly empty) suffix where at least one counter value c'_i , for some i , is a positive finite value in its first vertex v' . This frontier between prefix and suffix of $\pi[n]$ is represented by a vertical double bar \parallel with caption $\infty \rightarrow c$ in Figure B.1. This value c'_i is bounded by $\text{mR}(\lambda_n)$, the maximal finite range of λ_n . From there, as the corresponding ρ is λ -consistent, player i

reaches his target set in at most c'_i steps, and ρ enters a new region, which means that the section $\pi[n]$ is over. So, in that case, the length of $\pi[n]$ can be bounded by $|V| + \text{mR}(\lambda_n)$.

Suppose now that at vertex $\pi[n]_0$, there exists a counter value c_i for some player i that is neither 0 nor $+\infty$. This means that there was a constraint for player i initialized in a previous section $\pi[n']$, with $n' < n$, that has carried over to $\pi[n]_0$, via decrements of at least 1 per step. We know that the initial finite counter value is bounded by $\text{mR}(\lambda_{n'})$, and appeared before the end of section $\pi[n']$. Thus the length from the end of section $\pi[n']$ to the end of section $\pi[n]$ is bounded by $\text{mR}(\lambda_{n'})$, as again, once the counter value attains 0 for player i , the path π has entered the next section.

Therefore, considering the possible cases for each section, we can bound the total length of π as follows:

$$|\pi| \leq \sum_{j=\ell}^m |V| + 2 \cdot \text{mR}(\lambda_j).$$

Finally, remark that by I -monotonicity, it is actually the case that only (and at most) $|\Pi|$ different non-empty sections can appear in the decomposition of π . Furthermore, for each $n \in \{\ell + 1, \dots, N\}$, we have

$$\text{mR}(\lambda_n) \leq \max_{\substack{J_j > J_\ell \\ |J_j| = |J_n|}} \text{mR}(\lambda_j)$$

by Definition 11.2.1. Thus, we have the following bound:

$$|\pi| \leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \sum_{r=|J_\ell|+1}^{|\Pi|} |V| + 2 \cdot \max_{\substack{J_j > J_\ell \\ |J_j|=r}} \text{mR}(\lambda_j)$$

which is the bound stated in Lemma B.1.4. \square

Proof of Lemma B.1.4. Let v^C be a starting vertex in $\text{SV}(\lambda)$. Let π be a finite prefix of a valid path in $\mathbb{C}(\lambda)$ such that:

- $\pi_0 = v^C$,

- π does not contain any cycle.

The following proof is quite technical, thus we alleviate some of the difficulties by proving a first upper bound on the length of finite paths without cycles in $\mathbb{C}(\lambda)$, then by showing how to obtain the desired bound. The main idea is to bound iteratively the length of prefixes of π , adding at each step of the reasoning the section for the next region traversed by π .

Let π' be a valid path in $\mathbb{C}(\lambda)$ such that π is a prefix of π' . By Lemma 11.1.1, we know that there exist two natural numbers $\ell, m' \leq N$ and $m' - \ell$ (possibly empty) paths $\pi^\ell, \dots, \pi^{m'}$ in $\mathbb{C}(\lambda)$ such that:

- $\pi' = \pi^\ell \dots \pi^{m'}$,
- for each $\ell \leq j \leq m'$, each vertex in π^j is of the form $(w, (c_i)_{i \in \Pi})$ with $I(w) = J_j$.

The finite path π is a finite prefix of π' . Thus there exists a natural number $m \leq m'$ and $m - \ell$ (possibly empty) paths $\pi[\ell], \dots, \pi[m]$ in $\mathbb{C}(\lambda)$ such that:

- $\pi = \pi[\ell] \dots \pi[m]$,
- for each $\ell \leq j < m$, $\pi[j] = \pi'[j]$,
- $\pi[m]$ if a finite prefix of $\pi'[m]$.

By Lemma 11.2.6, there exist a corresponding λ -consistent play ρ' in \mathcal{X} and a corresponding history ρ . Furthermore, since π contains no cycle, the sections $\pi[n]$ do not either.

We first treat the case where $m < m'$.

We first bound the length of the first section $\pi[\ell]$. Recall that since $\pi[\ell]_0 = v^C \in \text{SV}(\lambda)$, we have that $v^C = (v, (c_i)_{i \in \Pi})$ with $c_i = 0$ if $i \in J_\ell$, $c_i = \lambda(v)$ if $v \in V_i$ and $c_i = +\infty$ otherwise.

Suppose $0 < \lambda(v) < +\infty$. In that case, we know that along ρ' , which is λ -consistent, player i reaches his target after at most $\lambda(v)$ steps. Thus, there exists $n \leq m'$, $n \neq \ell$ such that $I(\rho_{\lambda(v)}) = J_n$, $J_\ell \subsetneq J_n$, and $\pi_{\lambda(v)}$ belongs to section $\pi[n]$. This means that section $\pi[\ell]$ is shorter than $\lambda(v)$. Since $I(v) = J_\ell$, we have that $\lambda(v) \leq \text{mR}(\lambda_\ell)$, and thus $|\pi[\ell]| \leq \text{mR}(\lambda_\ell)$.

Suppose now that $\lambda(v) = 0$ or $\lambda(v) = +\infty$. The counter values at v^C are thus either 0 or $+\infty$. Along π' , they can stay stable for at most $|V|$ steps (see Proof sketch of Lemma B.1.4. Otherwise, the cycle induced in ρ' is also a cycle in π' as the counter values are fixed. If $\pi'_{|V|+1}$ is in section $\pi'[n]$ with $n > \ell$, then we immediately get that $|\pi[\ell]| \leq |V|$. If $\pi'_{|V|+1}$ is in section $\pi'[\ell]$, then it means that a counter value for some player i has become finite along the first $|V| + 1$ vertices of $\pi'[\ell]$. Let t be the first index it does so along $\pi'[\ell]$. By the same argument as in the previous case, we know that from $\pi'[\ell]_t$, there is at most $\text{mR}(\lambda_\ell)$ vertices before entering the next section of π' . Thus, we obtain that:

$$\begin{aligned} |\pi[\ell]| &\leq |\pi'[\ell]| \\ &\leq |V| + \text{mR}(\lambda_\ell) \\ &\leq |V| + 2 \cdot \text{mR}(\lambda_\ell). \end{aligned}$$

If $\ell = m$, we can already conclude that:

$$\begin{aligned} |\pi| &\leq |\pi[\ell]| \\ &\leq |V| + 2 \cdot \text{mR}(\lambda_\ell) \\ &\leq \sum_{j=\ell}^m |V| + 2 \cdot \text{mR}(\lambda_j). \end{aligned}$$

Suppose now that $\ell < m$. We show that for each $n > \ell$, we have:

$$|\pi[\ell] \dots \pi[n]| \leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \left(\sum_{j=\ell+1}^{n-1} |V| + 2 \cdot \text{mR}(\lambda_j) \right) + |V| + \text{mR}(\lambda_n).$$

Let $n = \ell + 1$. We assume $\pi[n]$ is not empty, otherwise its length is 0. Consider now $\pi[n]_0$.

Suppose there exists a player i such that his counter value c_i at $\pi[n]_0 = (w, J_n, (c_i)_{i \in \Pi})$ is a finite non-zero value. If $w \notin V_i$, it means that the counter value has decreased since a vertex $(w', (c'_i)_{i \in \Pi})$ in the previous section $\pi[\ell]$

such that $w' \in V_i$ and $c'_i = \lambda(w)$. Thus, the length of the path from this vertex $(w', (c'_i)_{i \in \Pi})$ to the next section $\pi[n+1]$ is smaller or equal than $\lambda(w')$. In particular, the whole section $\pi[n]$ is “covered” by this path. Therefore, we can conclude that $|\pi[n]| \leq \lambda(w') \leq \text{mR}(\lambda_\ell)$. Since we already know that $|\pi[\ell]| \leq |V| + \text{mR}(\lambda_\ell)$, we obtain that:

$$\begin{aligned} |\pi[\ell]\pi[n]| &\leq |V| + \text{mR}(\lambda_\ell) + \text{mR}(\lambda_\ell) \\ &\leq |V| + 2 \cdot \text{mR}(\lambda_\ell). \end{aligned}$$

If $w \in V_i$, it means that either the counter value c_i is equal to $\lambda(w)$ or has decreased since a vertex $(w', (c'_i)_{i \in \Pi})$ in the previous section $\pi[\ell]$ such that $w' \in V_i$ and $c'_i = \lambda(w')$. Thus, we have that $c_i \leq \text{mR}(\lambda_n)$ or $c_i \leq \text{mR}(\lambda_\ell)$, and, in turn, $|\pi[n]| \leq \text{mR}(\lambda_n)$ or $|\pi[n]| \leq \text{mR}(\lambda_n)$. Therefore, we can conclude that:

$$\begin{aligned} |\pi[\ell]\pi[n]| &\leq |V| + \text{mR}(\lambda_\ell) + \text{mR}(\lambda_\ell) + \text{mR}(\lambda_n) \\ &\leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \text{mR}(\lambda_n). \end{aligned}$$

Suppose now that for every player i , his counter value c_i at $\pi[n]_0 = (w, (c_i)_{i \in \Pi})$ is either 0 or $+\infty$. In that case, we are in a similar case than for section $\pi[\ell]$, thus we can conclude that $|\pi[n]| \leq |v| + \text{mR}(\lambda_n)$. Thus, we have indeed:

$$|\pi[\ell]\pi[n]| \leq |V| + \text{mR}(\lambda_\ell) + |V| + \text{mR}(\lambda_n),$$

and also

$$|\pi[\ell]\pi[n]| \leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + |V| + \text{mR}(\lambda_n).$$

If $m = \ell + 1$, we are done.

Suppose now that $m > \ell + 1$. Let n be such that $\ell + 1 < n \leq m$. Assume that for each $n' < n$, it holds that

$$|\pi[\ell] \dots \pi[n']| \leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \left(\sum_{j=\ell+1}^{n'-1} |V| + 2 \cdot \text{mR}(\lambda_j) \right) + |V| + \text{mR}(\lambda_{n'}).$$

We assume $\pi[n]$ is not empty, otherwise its length is 0. Consider now $\pi[n]_0$.

Suppose there exists a player i such that his counter value c_i at $\pi[n]_0 = (w, (c_i)_{i \in \Pi})$ is a finite non-zero value. If $w \notin V_i$, it means that the counter value has decreased since a vertex $(w', (c'_i)_{i \in \Pi})$ in a previous section $\pi[n']$ such that $w' \in V_i$ and $c'_i = \lambda(w')$. Thus, the length of the path from this vertex $(w', (c'_i)_{i \in \Pi})$ to the next section $\pi[n+1]$ is smaller or equal than $\lambda(w')$. In particular, the whole sections from $\pi[n'+1]$ to $\pi[n]$ are “covered” by this path. Therefore, we can conclude that $|\pi[n'+1] \dots \pi[n]| \leq \lambda(w') \leq \text{mR}(\lambda_{n'})$. Since we already know that:

$$|\pi[\ell] \dots \pi[n']| \leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \left(\sum_{j=\ell+1}^{n'-1} |V| + 2 \cdot \text{mR}(\lambda_j) \right) + |V| + \text{mR}(\lambda_{n'}),$$

we obtain that:

$$\begin{aligned} |\pi[\ell] \dots \pi[n]| &\leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \left(\sum_{j=\ell+1}^{n'-1} |V| + 2 \cdot \text{mR}(\lambda_j) \right) \\ &\quad + |V| + \text{mR}(\lambda_{n'}) + \text{mR}(\lambda_{n'}) \\ &\leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \left(\sum_{j=\ell+1}^{n'-1} |V| + 2 \cdot \text{mR}(\lambda_j) \right) \\ &\quad + |V| + 2 \cdot \text{mR}(\lambda_{n'}) \\ &\leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \left(\sum_{j=\ell+1}^{n-1} |V| + 2 \cdot \text{mR}(\lambda_j) \right) \\ &\quad + |V| + \text{mR}(\lambda_n). \end{aligned}$$

If $w \in V_i$, it means that either the counter value c_i is equal to $\lambda(w)$ or has decreased since a vertex $(w', (c'_i)_{i \in \Pi})$ in a previous section $\pi[n']$ such that $w' \in V_i$ and $c'_i = \lambda(w')$. Thus, we have that $c_i \leq \text{mR}(\lambda_n)$ or $c_i \leq \text{mR}(\lambda_{n'})$, and, in turn, $|\pi[n]| \leq \text{mR}(\lambda_n)$ or $|\pi[n]| \leq \text{mR}(\lambda_{n'})$.

If $|\pi[n]| \leq \text{mR}(\lambda_n)$, we can conclude that:

$$\begin{aligned} |\pi[\ell] \dots \pi[n]| &\leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \left(\sum_{j=\ell+1}^{n-1} |V| + 2 \cdot \text{mR}(\lambda_j) \right) + \text{mR}(\lambda_n) \\ &\leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \left(\sum_{j=\ell+1}^{n-1} |V| + 2 \cdot \text{mR}(\lambda_j) \right) + |V| + \text{mR}(\lambda_n). \end{aligned}$$

If $|\pi[n]| \leq \text{mR}(\lambda_{n'})$, we can conclude that:

$$\begin{aligned} |\pi[\ell] \dots \pi[n]| &\leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \left(\sum_{j=\ell+1}^{n'-1} |V| + 2 \cdot \text{mR}(\lambda_j) \right) \\ &\quad + |V| + \text{mR}(\lambda_{n'}) + \text{mR}(\lambda_{n'}) \\ &\leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \left(\sum_{j=\ell+1}^{n-1} |V| + 2 \cdot \text{mR}(\lambda_j) \right) + |V| + \text{mR}(\lambda_n). \end{aligned}$$

Suppose now that for every player i , his counter value c_i at $\pi[n]_0 = (w, (c_i)_{i \in \Pi})$ is either 0 or $+\infty$. In that case, we are in a similar case than for section $\pi[\ell]$, thus we can conclude that $|\pi[n]| \leq |v| + \text{mR}(\lambda_n)$. Thus, we have indeed:

$$|\pi[\ell] \dots \pi[n]| \leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \left(\sum_{j=\ell+1}^{n-1} |V| + 2 \cdot \text{mR}(\lambda_j) \right) + |V| + \text{mR}(\lambda_n)$$

and thus finally:

$$|\pi| \leq |\pi[\ell] \dots \pi[m]| \leq \sum_{j=\ell}^m |V| + 2 \cdot \text{mR}(\lambda_j).$$

Assume now that $m = m'$. In that case, we cannot rely on the section π'^m to be finite, as the play ρ' never reaches another region. However, in this situation, it is guaranteed that the counter values in π'^m are fixed and are equal to either 0 or $+\infty$: indeed, a finite counter value would imply that some

player i such that $i \notin J_m$ reaches his target in ρ' exactly when his counter value becomes 0 in π' . But π'^m is the last section of ρ' , thus no new region is reached after J_m and no new player can visit his target set than the players $i \in J_m$. Therefore, finite prefix of π'^m that contains no cycle has its length bounded by $|V|$. Thus, we can conclude:

$$\begin{aligned} |\pi| &= |\pi[\ell] \dots \pi[m]| \\ &\leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \left(\sum_{j=\ell+1}^{m-1} |V| + 2 \cdot \text{mR}(\lambda_j) \right) + |V| \\ &\leq \sum_{j=\ell}^m |V| + 2 \cdot \text{mR}(\lambda_j). \end{aligned}$$

In fact, the bound given above can be slightly changed to give the desired bound. Indeed, the bound above relies on the fact that $m \leq N$ and covers the case where a path traverses every region from J_ℓ to J_N . In all generality, the number N of different regions can be exponential in the number $|\Pi|$ of players. However, by the I -monotonicity property, we know that a path can actually traverse at most $|\Pi|$ regions. Thus, in the region decomposition of a path, only (and at most) $|\Pi|$ sections are relevant and of length greater than 0. Therefore, we can define a subsequence of indices $(n_r)_{r \leq |\pi|}$, with $n_1 = \ell$, such that in fact $\pi = \pi[n_1] \dots \pi[n_{|\Pi|}]$. Hence, we obtain the following bound on the length t of π :

$$t \leq \sum_{r=1}^{|\Pi|} |V| + 2 \cdot \text{mR}(\lambda_{n_r})$$

Finally, as for every $r \leq |\pi|$, $r > 1$, we have $\text{mR}(\lambda_{n_r}) \leq \max\{\text{mR}(\lambda_j) \mid J_j > J_\ell, |J_j| = |J_{n_r}|\}$, we obtain the desired bound:

$$t \leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \sum_{r=|J_\ell|+1}^{|\Pi|} |V| + 2 \cdot \max_{\substack{J_j > J_\ell \\ |J_j|=r}} \text{mR}(\lambda_j)$$

□

We are now ready to prove Proposition 11.2.7.

Proof of Proposition 11.2.7. Let $v \in V^X$ with $I(v) = J_\ell$ and $i \in \Pi$. Let $c \in \mathbb{N} \cup \{+\infty\}$ be such that $\sup \{\text{Cost}_i(\rho) \mid \rho \in \Lambda(v)\} = c$. Consider $\rho \in \Lambda^k(v)$ such that $\text{Cost}_i(\rho) = c$. Notice that such a play always exists by Corollary 11.1.11. Consider also π the valid path in $\mathbb{C}(\lambda)$ that starts in v^C and corresponds to ρ .

Suppose first that $c < +\infty$ and let us prove that

$$c \leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \sum_{r=|J_\ell|+1}^{|\Pi|} |V| + 2 \cdot \max_{\substack{J_j > J_\ell \\ |J_j|=r}} \text{mR}(\lambda_j). \quad (\text{B.10})$$

If $i \in J_\ell$, then $c = 0$, as every play starting in v has a cost 0 for player i . Hence (B.10) trivially holds.

Suppose now that $i \notin J_\ell$. As player i eventually reaches his target set along ρ , this means that ρ eventually leaves region J_ℓ and eventually reaches another region J_n such that $i \in J_n$. Consider the prefix $\rho_{\leq c}$ of ρ of length c . Let $\pi_{\leq c} := \pi_{\leq c}[\ell] \dots \pi_{\leq c}[m]$ be the prefix of π and its region decomposition such that $\pi_{\leq c}$ is associated with $\rho_{\leq c}$. Notice that $\pi_{\leq c}[m]$ consists only of one vertex corresponding to ρ_c , and that for every $n < m$, we have $i \notin J_n$.

Suppose that $\pi_{\leq c}$ contains a cycle. By construction of $\mathbb{C}(\lambda)$, this cycle is included in one single section $\pi_{\leq c}[n]$, where $n < m$ (as $\pi_{\leq c}[m]$ contains only one vertex), and thus $i \notin J_n$. Consider the infinite path π' in $\mathbb{C}(\lambda)$ that follows $\pi_{\leq c}$ until the cycle and then repeats the cycle forever. By Lemma 11.2.6, there exists a λ -consistent play ρ' in \mathcal{X} corresponding to π' . We have $\text{Cost}_i(\rho') = +\infty$ for player i , as ρ' never reaches a region where player i visits his target set. This is a contradiction with the fact that $\sup \{\text{Cost}_i(\rho) \mid \rho \in \Lambda(v)\} = c$ is finite.

Therefore $\pi_{\leq c}$ contains no cycle, and by Lemma B.1.4,

$$|\pi_{\leq c}| \leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \sum_{r=|J_\ell|+1}^{|\Pi|} |V| + 2 \cdot \max_{\substack{J_j > J_\ell \\ |J_j|=r}} \text{mR}(\lambda_j).$$

Since $|\pi_{\leq c}| = c = \text{Cost}_i(\rho)$, we obtain Inequality (B.10).

Let us now prove the second part of Proposition 11.2.7 for both cases $c < +\infty$ and $c = +\infty$. Given the valid path π , we consider the first occurrence of a cycle in π . We then construct the infinite path π' in $\mathbb{C}(\lambda)$ that follows π until this cycle and then repeats it forever. Then π' is a lasso hg^ω with the length $|hg|$ bounded by $2 \cdot |\mathbb{C}(\lambda)|$. Clearly if $c = +\infty$, then the corresponding play ρ' in \mathcal{X} belongs to $\Lambda(v)$ and has a cost $\text{Cost}_i(\rho')$ equal to $\text{Cost}_i(\rho) = +\infty$. If $c < +\infty$, we know by the first part of the proof that $\pi_{\leq c}$ contains no cycle and thus is prefix of h . Therefore we also have that the corresponding play ρ' has cost $\text{Cost}_i(\rho') = \text{Cost}_i(\rho) = c$. \square

Remark B.1.5. In the proof of Proposition 11.2.7 and Lemma B.1.4, we consider paths π in the counter graph $\mathbb{C}(\lambda)$ that starts in a vertex v such that $I(v) = J_\ell$. This means that such paths only visit vertices of regions V^{J_j} such that $j \in \{\ell, \ell + 1, \dots, N\}$. There are therefore paths in the counter graph restricted $V^{\geq J_\ell}$ that we denote by $\mathbb{C}(\lambda_{\geq \ell})$.

Proof of Theorem 11.2.8

Proof of Theorem 11.2.8. The proof is done by a *double* induction: First, we exploit the fact that Algorithm 4 treats every region one after the other, following the total order on regions in reverse. That is, to compute the values of the fixpoint function λ^* over V^X , Algorithm 4 computes first the values of λ^* on region V^{J_N} , then on region $V^{J_{N-1}}$ etc. . . until finally on region V^{J_1} . Thus, we follow this order to prove the local bounds on $\text{mR}(\lambda_\ell^{k_\ell^*})$, starting by $\text{mR}(\lambda_N^{k_N^*})$ and making our way up to $\text{mR}(\lambda_1^{k_1^*})$, assuming for each region V^{J_ℓ} such that $\ell < N$ that the bound is true for every region already treated by Algorithm 4. Second, given a non-bottom region V^{J_ℓ} and assuming the bound is true for every already treated region, we proceed to show the local bound for V^{J_ℓ} by induction on the number k of steps in the computation, which corresponds to the values of function λ^k in the sequence of functions leading to the fixpoint, up to step k_ℓ^* , where the values stabilize on region V^ℓ .

Let us now detail the proof. It is structured into several parts.

Part 1. We begin with some notations and basic properties. We introduce a useful notation: for each $\ell < N$ and each $k \in \mathbb{N}$, we define $\alpha(\lambda^k, \ell)$ as follows^a:

$$\alpha(\lambda^k, \ell) := \sum_{r=|J_\ell|+1}^{|\Pi|} |V| + 2 \cdot \max_{\substack{J_j > J_\ell \\ |J_j|=r}} \text{mR}(\lambda_j^k)$$

Let $\ell \leq N$ and $v \in V^{J_\ell}$. By Lemma 11.1.8, we know the following:

1. the values of $\lambda_\ell^{k+1}(v)$ and $\lambda_\ell^k(v)$ may differ only when $k_{\ell+1}^* < k < k_\ell^*$;
2. for $k \leq k_{\ell+1}^*$, we have $\lambda_\ell^k(v) = \lambda_\ell^{k+1}(v) = \lambda_\ell^0(v)$;
3. for $k \geq k_\ell^*$, we have $\lambda_\ell^k(v) = \lambda_\ell^{k+1}(v) = \lambda_\ell^{k_\ell^*}(v)$

By 3., we have that for each $\ell < N$, and for $k \geq k_{\ell+1}^*$, we have $\text{mR}(\lambda_{\ell+1}^k) = \text{mR}(\lambda_{\ell+1}^{k_{\ell+1}^*})$. Thus, for $k \geq k_{\ell+1}^*$, we also have

$$\alpha(\lambda^k, \ell) = \alpha(\lambda^{k_{\ell+1}^*}, \ell). \quad (\text{B.11})$$

Part 2. Let $\ell < N$ and assume J_ℓ is not a bottom region. We start by proving, this time by induction on the number k of algorithm steps, a bound on $\text{mR}(\lambda_\ell^k)$. Note, by Lemma 11.1.8 recalled above, that the only relevant steps for region J_ℓ are the steps k , where $k_{\ell+1}^* \leq k \leq k_\ell^*$. Let us show, for each such k , that:

$$\text{mR}(\lambda_\ell^k) \leq \left(\sum_{i=0}^{|\text{Dom}(\lambda_\ell^k)|} 2^i \right) \cdot (1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell)) \quad (\text{B.12})$$

where $\text{Dom}(\lambda_\ell^k) = \{v \in V^{J_\ell} \mid \lambda_\ell^k(v) \neq 0, \lambda_\ell^k(v) \neq +\infty\}$.

- *Base case:* Assume $k = k_{\ell+1}^*$. By Lemma 11.1.8, we know that $\lambda_\ell^k = \lambda_\ell^0$. Thus, $\text{mR}(\lambda_\ell^k) = \text{mR}(\lambda_\ell^0) = 0$. Furthermore, $\text{Dom}(\lambda_\ell^k) = \text{Dom}(\lambda_\ell^0) = \emptyset$, thus $\sum_{i=0}^{|\text{Dom}(\lambda_\ell^k)|} 2^i = 1$. Clearly, Equation (B.12) is satisfied.

- *General case:* Let now k be such that $k_{\ell+1}^* \leq k < k_\ell^*$. Assume that Equation (B.12) holds for k . Let us show that

$$\text{mR}(\lambda_\ell^{k+1}) \leq \left(\sum_{i=0}^{|\text{Dom}(\lambda_\ell^{k+1})|} 2^i \right) \cdot (1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell)) \quad (\text{B.13})$$

We distinguish the two following subcases: (a) when $\text{Dom}(\lambda_\ell^{k+1}) = \text{Dom}(\lambda_\ell^k)$, and (b) when $\text{Dom}(\lambda_\ell^{k+1}) \neq \text{Dom}(\lambda_\ell^k)$.

- (a) Assume $\text{Dom}(\lambda_\ell^{k+1}) = \text{Dom}(\lambda_\ell^k)$. We know by Lemma B.1.1 that for each $v \in V^{J_\ell}$, $\lambda_\ell^{k+1}(v) \leq \lambda_\ell^k(v)$. Thus in the considered case we have:

$$\begin{aligned} \text{mR}(\lambda_\ell^{k+1}) &\leq \text{mR}(\lambda_\ell^k) \\ &\leq \left(\sum_{i=0}^{|\text{Dom}(\lambda_\ell^k)|} 2^i \right) \cdot (1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell)) \quad \text{by IH (B.12)} \\ &\leq \left(\sum_{i=0}^{|\text{Dom}(\lambda_\ell^{k+1})|} 2^i \right) \cdot (1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell)) \quad \text{by (a)} \end{aligned}$$

That is, Equation (B.13) holds.

- (b) Assume $\text{Dom}(\lambda_\ell^{k+1}) \neq \text{Dom}(\lambda_\ell^k)$. Either we have $\text{mR}(\lambda_\ell^{k+1}) \leq \text{mR}(\lambda_\ell^k)$, and we can proceed as in subcase (a), or we know that there exists $v \in V^{J_\ell}$ such that $\lambda_\ell^k(v) = +\infty$ and $\lambda_\ell^{k+1}(v) < +\infty$. Recall that by Definition 11.1.5 we have that

$$\lambda_\ell^{k+1}(v) = 1 + \min_{(v, v') \in E^X} \sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v')\} \quad (\text{B.14})$$

Since $\lambda_\ell^{k+1}(v) \neq +\infty$, we know that $\sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v')\}$ is finite, for at least one successor v' of v . Thus, by Proposition 11.2.7, we obtain:

$$\begin{aligned}
\lambda_\ell^{k+1}(v) &\leq 1 + |V| + 2 \cdot \text{mR}(\lambda_\ell^k) + \alpha(\lambda^k, \ell) \\
&\leq 1 + |V| + 2 \cdot \text{mR}(\lambda_\ell^k) + \alpha(\lambda^{k_{\ell+1}^*}, \ell) && \text{by (B.11)} \\
&\leq 1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell) \\
&\quad + 2 \cdot \left(\sum_{i=0}^{|\text{Dom}(\lambda_\ell^k)|} 2^i \right) \cdot (1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell)) \text{ by IH (B.12)} \\
&\leq \left[2 \cdot \left(\sum_{i=0}^{|\text{Dom}(\lambda_\ell^k)|} 2^i \right) + 1 \right] \cdot (1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell))
\end{aligned}$$

Since $\text{Dom}(\lambda_\ell^{k+1}) \neq \text{Dom}(\lambda_\ell^k)$ by (b) (thus indeed $|\text{Dom}(\lambda_\ell^{k+1})| > |\text{Dom}(\lambda_\ell^k)|$), we have

$$\left[2 \cdot \left(\sum_{i=0}^{|\text{Dom}(\lambda_\ell^k)|} 2^i \right) + 1 \right] \leq \left(\sum_{i=0}^{|\text{Dom}(\lambda_\ell^{k+1})|} 2^i \right)$$

Hence, we have

$$\lambda_\ell^{k+1}(v) \leq \left(\sum_{i=0}^{|\text{Dom}(\lambda_\ell^{k+1})|} 2^i \right) \cdot (1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell))$$

Finally, as this holds for any such $v \in V^{J_\ell}$, we obtain

$$\text{mR}(\lambda_\ell^{k+1}) \leq \left(\sum_{i=0}^{|\text{Dom}(\lambda_\ell^{k+1})|} 2^i \right) \cdot (1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell))$$

That is, Equation (B.13) holds for region J_ℓ .

We proved that Equation (B.12) holds for each non-bottom region J_ℓ and each step k such that $k_{\ell+1}^* \leq k \leq k_\ell^*$.

Part 3. We can now come back to the induction on the regions J_ℓ , following the order provided by Algorithm 4. Let us show that each $\ell \leq N$, we

have

$$\text{mR}(\lambda_\ell^{k_\ell^*}) \leq 2^{(|V|+1)} \cdot (|\Pi||V| + |V| + 1) \cdot \sum_{i=0}^{|\Pi \setminus J_\ell|} 2^{(|V|+1)} \cdot 2^{|\Pi|} \quad (\text{B.15})$$

- *Base case:* If J_ℓ is a bottom region, we have that $\text{mR}(\lambda_\ell^{k_\ell^*}) = 0$ showing that (B.15) holds in this case and thus in particular when $\ell = N$.
- *General case:* Let $\ell < N$ and suppose J_ℓ is not a bottom region. Assume that, for each $j > \ell$, Inequality (B.15) holds. By (B.12), we know that

$$\text{mR}(\lambda_\ell^k) \leq \left(\sum_{i=0}^{|\text{Dom}(\lambda_\ell^k)|} 2^i \right) \cdot (1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell)).$$

Furthermore, as $|\text{Dom}(\lambda_\ell^{k_\ell^*})| \leq |V|$, we have

$$\text{mR}(\lambda_\ell^k) \leq 2^{(|V|+1)} \cdot (1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell)) \quad (\text{B.16})$$

We turn now our attention towards the term $\alpha(\lambda^{k_{\ell+1}^*}, \ell)$:

$$\begin{aligned} \alpha(\lambda^{k_{\ell+1}^*}, \ell) &= \sum_{r=|J_\ell|+1}^{|\Pi|} |V| + 2 \cdot \max_{\substack{J_j > J_\ell \\ |J_j|=r}} \text{mR}(\lambda_j^{k_{\ell+1}^*}) && \text{by Definition} \\ &= \sum_{r=|J_\ell|+1}^{|\Pi|} |V| + 2 \cdot \max_{\substack{J_j > J_\ell \\ |J_j|=r}} \text{mR}(\lambda_j^{k_j^*}) && \text{by Lemma 11.1.8} \end{aligned}$$

By the induction hypothesis (B.15), we can bound each term $\text{mR}(\lambda_j^{k_j^*})$ appearing in the above sum, as $|J_j| > |J_\ell|$:

$$\begin{aligned} \text{mR}(\lambda_j^{k_j^*}) &\leq 2^{(|V|+1)} \cdot (|\Pi||V| + |V| + 1) \cdot \sum_{i=0}^{|\Pi \setminus J_j|} 2^{(|V|+1)} \cdot 2^{|\Pi|} \\ &\leq 2^{(|V|+1)} \cdot (|\Pi||V| + |V| + 1) \cdot \sum_{i=0}^{|\Pi \setminus J_\ell|-1} 2^{(|V|+1)} \cdot 2^{|\Pi|}. \end{aligned}$$

Coming back to $\alpha(\lambda^{k_{\ell+1}^*}, \ell)$ whose the number of terms in the sum can be bounded by $|\Pi|$, we get

$$\alpha(\lambda^{k_{\ell+1}^*}, \ell) \tag{B.17}$$

$$\begin{aligned} &= \sum_{r=|J_{\ell}|+1}^{|\Pi|} |V| + 2 \cdot \max_{\substack{J_j > J_{\ell} \\ |J_j|=r}} \text{mR}(\lambda_j^{k_{\ell+1}^*}) \\ &\leq |\Pi||V| + 2|\Pi| \left(2^{(|V|+1)} \cdot (|\Pi||V| + |V| + 1) \cdot \sum_{i=0}^{|\Pi \setminus J_{\ell}|-1} 2^{(|V|+1)} \cdot 2^{|\Pi|} \right). \end{aligned} \tag{B.18}$$

Hence we have, by combining (B.18) and (B.16):

$$\begin{aligned} &\text{mR}(\lambda_{\ell}^{k_{\ell}^*}) \\ &\leq 2^{(|V|+1)} \cdot (1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell)) \\ &\leq 2^{(|V|+1)} \cdot (1 + |V| + |\Pi||V|) \\ &\quad + 2^{(|V|+1)} \cdot 2|\Pi| \left(2^{(|V|+1)} \cdot (|\Pi||V| + |V| + 1) \cdot \sum_{i=0}^{|\Pi \setminus J_{\ell}|-1} 2^{(|V|+1)} \cdot 2^{|\Pi|} \right) \\ &\leq 2^{(|V|+1)} \cdot (1 + |V| + |\Pi||V|) \left[1 + 2^{(|V|+1)} \cdot 2|\Pi| \left(\sum_{i=0}^{|\Pi \setminus J_{\ell}|-1} 2^{(|V|+1)} \cdot 2^{|\Pi|} \right) \right] \\ &\leq 2^{(|V|+1)} \cdot (1 + |V| + |\Pi||V|) \cdot \left(\sum_{i=0}^{|\Pi \setminus J_{\ell}|} 2^{(|V|+1)} \cdot 2^{|\Pi|} \right). \end{aligned}$$

Hence, Equation (B.15) holds for region J_{ℓ} .

Part 4. We can now prove the three statements of Theorem 11.2.8. We obtain the first one from Inequality (B.15) by recalling that $|\Pi| \leq |V|$ and $|V| \geq 2$:

$$\begin{aligned}
\text{mR}(\lambda_\ell^{k_\ell^*}) &\leq 2^{(|V|+1)} \cdot (1 + |V| + |\Pi||V|) \cdot \left(\sum_{i=0}^{|\Pi \setminus J_\ell|} 2^{(|V|+1)} \cdot 2^{|\Pi|} \right) \\
&\leq |V|^{(|V|+1)} \cdot (1 + |V| + |V|^2) \cdot \left(\sum_{i=0}^{|\Pi \setminus J_\ell|} |V|^{(|V|+3)} \right) \\
&\leq |V|^{(|V|+1)} \cdot (1 + |V| + |V|^2) \cdot (|V|^{(|V|+3)})^{(|\Pi \setminus J_\ell|+1)} \\
&\leq \mathcal{O}(|V|^{(|V|+3)}) \cdot (|V|^{(|V|+3)})^{(|\Pi \setminus J_\ell|+1)} \\
&\leq \mathcal{O}(|V|^{(|V|+3)(|\Pi \setminus J_\ell|+2)}) \tag{B.19}
\end{aligned}$$

This proves the first statement of Theorem 11.2.8.

For the second one, remark that $\text{mR}(\lambda_{\geq \ell}^{k_\ell^*}) = \max\{\text{mR}(\lambda_j^{k_j^*}) \mid J_j \geq J_\ell\}$ (since only regions already treated are considered). For each such $j \geq \ell$, we have $\text{mR}(\lambda_j^{k_j^*}) \leq \mathcal{O}(|V|^{(|V|+3)(|\Pi \setminus J_j|+2)})$ by (B.19) and thus $\text{mR}(\lambda_j^{k_j^*}) \leq \mathcal{O}(|V|^{(|V|+3)(|\Pi|+2)})$. It follows that $\text{mR}(\lambda_{\geq \ell}^{k_\ell^*}) \leq \mathcal{O}(|V|^{(|V|+3)(|\Pi|+2)})$.

For the last statement, recall that $\text{mR}(\lambda^*) = \text{mR}(\lambda_{\geq 1}^{k_1^*})$. Hence in particular we obtain $\text{mR}(\lambda^*) \leq \mathcal{O}(|V|^{(|V|+3)(|\Pi|+2)})$. \square

^aThis sum appears in the statement of Proposition 11.2.7.

Proofs of Corollary 11.2.10 and Corollary 11.2.11

Proof of Corollary 11.2.10. Assume, without loss of generality, that J_ℓ is not a bottom region (otherwise we immediately have $\text{mR}(\lambda_\ell^k) = 0$ for every $k \in \mathbb{N}$). Let $k \in \mathbb{N}$. Again, by Lemma 11.1.8, if $k \leq k_{\ell+1}^*$, we already have $\text{mR}(\lambda_\ell^k) = 0$, and if $k \geq k_\ell^*$, we have $\lambda_\ell^k(v) = \lambda_\ell^{k_\ell^*}(v)$. Thus we assume that $k_{\ell+1}^* < k \leq k_\ell^*$. Using the terminology of the proof of Theorem 11.2.8, we know, by (B.12) in its proof, that

$$\text{mR}(\lambda_\ell^k) \leq \left(\sum_{i=0}^{|\text{Dom}(\lambda_\ell^k)|} 2^i \right) \cdot (1 + |V| + \alpha(\lambda_\ell^{k_{\ell+1}^*}, \ell)) \tag{B.20}$$

As $|\text{Dom}(\lambda_\ell^k)| \leq |V|$, we have

$$\text{mR}(\lambda_\ell^k) \leq 2^{|V|+1} \cdot (1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell)) \quad (\text{B.21})$$

From there, following the same steps (from Equation (B.16) onwards) as in the proof of Theorem 11.2.8 leads to the desired bound: $\text{mR}(\lambda_\ell^k) \leq \mathcal{O}(|V|^{(|V|+3)(|\Pi|+2)})$.

Similarly, as $\text{mR}(\lambda_{\geq \ell}^k) = \max\{\text{mR}(\lambda_j^k) \mid j \geq \ell\}$, and as we just showed that $\text{mR}(\lambda_j^k) \leq \mathcal{O}(|V|^{(|V|+3)(|\Pi|+2)})$ for every j and k , we immediately get $\text{mR}(\lambda_{\geq \ell}^k) \leq \mathcal{O}(|V|^{(|V|+3)(|\Pi|+2)})$. \square

Proof of Corollary 11.2.11. Let $v \in V^X$ with $I(v) = J_\ell$ and $k \in \mathbb{N}$. Suppose there exists $c \in \mathbb{N}$ such that $\sup \{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v)\} = c$. If J_ℓ is a bottom region, then $c = 0$, thus we assume from now on that J_ℓ is not a bottom region. Then, by Proposition 11.2.7, we know that:

$$c \leq |V| + 2 \cdot \text{mR}(\lambda_\ell^k) + \sum_{r=|J_\ell|+1}^{|\Pi|} |V| + 2 \cdot \max_{\substack{J_j > J_\ell \\ |J_j|=r}} \text{mR}(\lambda_j^{k^*}).$$

Using the terminology of the proof of Theorem 11.2.8, we have:

$$\begin{aligned} c &\leq |V| + 2 \cdot \text{mR}(\lambda_\ell^k) + \alpha(\lambda^{k_{\ell+1}^*}, \ell) \\ &\leq 1 + |V| + 2 \cdot \text{mR}(\lambda_\ell^k) + \alpha(\lambda^{k_{\ell+1}^*}, \ell) \\ &\leq 1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell) + 2 \cdot \left(\sum_{i=0}^{|\text{Dom}(\lambda_\ell^k)|} 2^i \right) \cdot (1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell)) \text{ by (B.12)} \\ &\leq \left(\sum_{i=0}^{|\text{Dom}(\lambda_\ell^k)|+1} 2^i \right) \cdot (1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell)) \end{aligned}$$

As $|\text{Dom}(\lambda_\ell^k)| \leq |V|$, we obtain:

$$c \leq 2 \cdot \left[2^{(|V|+1)} \cdot (1 + |V| + \alpha(\lambda^{k_{\ell+1}^*}, \ell)) \right].$$

From there, following the same steps (from Equation (B.16) onwards) as in the proof of Theorem 11.2.8 leads to the desired bound:

$$c \leq 2 \cdot \mathcal{O}\left(|V|^{(|V|+3)(|\Pi|+2)}\right) = \mathcal{O}\left(|V|^{(|V|+3)(|\Pi|+2)}\right).$$

□

B.2 Proofs of Chapter 12

B.2.1 Proof of Section 12.2

Proof of Theorem 12.2.1

To prove Theorem 12.2.1, we begin with a preliminary lemma and the proof of Theorem 12.2.1 follows.

Lemma B.2.1. *Let \mathcal{G} be a multiplayer quantitative Reachability game. Then for all $v_0 \in V$ for which some target set F_j , $j \in \Pi$, is reachable from v_0 , there exists an SPE in (\mathcal{G}, v_0) whose outcome ρ visits at least one target set F_i , $i \in \Pi$, that is, $|\text{Visit}(\rho)| \geq 1$.*

Proof. By Theorem 2.4.9, there exists an SPE in (\mathcal{G}, v_0) for each initial vertex $v_0 \in V$. Consider the set $U \subseteq V$ of vertices u for which some F_j is reachable from u , and the set $U' \subseteq U$ of those vertices u for which there is an SPE in (\mathcal{G}, u) that visits at least one target set. We have to prove that $U = U'$.

Let us assume that $v_0 \in U \setminus U'$. We claim that there exists an edge (u, u') such that $u \in U \setminus U'$ and $u' \in U'$. Indeed as $v_0 \in U$, there exists a history $h = v_0 v_1 \dots v_k$ with $v_k \in F_j$ for some j . Hence $v_k \in U'$ since the outcome of all SPEs in (\mathcal{G}, v_k) immediately visits F_j . As along h we begin with $v_0 \in U \setminus U'$ and we end with $v_k \in U'$, there must exist an edge $(v_\ell, v_{\ell+1}) = (u, u')$ with $u \in U \setminus U'$ and $u' \in U'$.

Let σ^u (resp. $\sigma^{u'}$) be an SPE in (\mathcal{G}, u) (resp. in (\mathcal{G}, u')). As $u' \in U'$, we can suppose that the outcome of $\sigma^{u'}$ visits some target set F_j . From σ^u and $\sigma^{u'}$,

we are going to construct another SPE τ in (\mathcal{G}, u) whose outcome will now visit this set F_j . This will lead to a contradiction with $u \in U \setminus U'$. We define such a strategy profile τ equal to σ^u except that it is replaced by $\sigma^{u'}$ for all histories with prefix uu' . More precisely,

- for the particular history u , if $u \in V_i$, then $\tau_i(u) = u'$,
- for each history $uu'h \in \text{Hist}_i$, $i \in \Pi$, we define $\tau_i(uu'h) = \sigma_i^{u'}(u'h)$,
- for each history $uv'h \in \text{Hist}_i$, $i \in \Pi$, with $v' \neq u'$, we define $\tau_i(uv'h) = \sigma_i^u(uv'h)$.

Clearly the outcome of τ is equal to $u \langle \sigma^{u'} \rangle_{u'}$ and thus visits F_j . It remains to show that τ is an SPE, *i.e.*, that $\tau_{|h}$ is an NE in the subgame $(\mathcal{G}_{|h}, v)$ for all $hv \in \text{Hist}_i(u)$, $i \in \Pi$.

- For all histories hv that begin with uv' with $v' \neq u'$, clearly $\tau_{|h}$ is an NE in $(\mathcal{G}_{|h}, v)$ because $\tau_{|h} = \sigma_{|h}^u$ and σ^u is an SPE.
- Take any history hv that begin with uu' , and let $h = uh'$. Let τ'_i be a deviating strategy for Player i in $(\mathcal{G}_{|h}, v)$. By definition of τ we have

$$\begin{aligned} \langle \tau_{|h} \rangle_v &= \langle \sigma_{|h'}^{u'} \rangle_v \\ \langle (\tau'_i, \tau_{|h, -i}) \rangle_v &= \langle (\tau'_i, \sigma_{|h', -i}^{u'}) \rangle_v \end{aligned}$$

Moreover, as u belongs to no target set, we have $\text{Cost}_i(u\rho) = 1 + \text{Cost}_i(\rho)$ for all plays $\rho \in \text{Plays}(u')$. It follows that if τ'_i is a profitable deviation for Player i with respect to $\tau_{|h}$, it is also a profitable deviation with respect to $\sigma_{|h'}^{u'}$. The latter case never holds because $\sigma^{u'}$ is an SPE (and in particular $\sigma_{|h'}^{u'}$ is an NE). Therefore $\tau_{|h}$ is an NE in $(\mathcal{G}_{|h}, v)$.

- It remains to consider the history u and to prove that τ is an NE in (\mathcal{G}, u) . From what has been gathered so far, only Player i such that $u \in V_i$ might have a profitable deviation by deviating at the initial vertex u with a strategy τ'_i such that $\tau'_i(u) = v' \neq u' = \tau_i(u)$. Notice that since $u \in U \setminus U'$, we have $\text{Cost}_i(\langle \sigma^u \rangle_u) = +\infty$ and since σ^u is an SPE (and in particular an NE), we have $\text{Cost}_i(\langle \tau'_i, \sigma_{-i}^u \rangle_u) = +\infty$. Moreover

as $\tau'_i(u) = v' \neq u'$ and by definition of τ , we have $\text{Cost}_i(\langle \tau'_i, \sigma_{-i}^u \rangle_u) = \text{Cost}_i(\langle \tau'_i, \tau_{-i} \rangle_u) = +\infty$. It follows that τ'_i is not a profitable deviation for Player i with respect to τ , and then τ is an NE in (\mathcal{G}, u) . \square

Proof of Theorem 12.2.1. Let (\mathcal{G}, v_0) , with $\mathcal{G} = (\mathcal{A}, (\text{Cost}_i)_{i \in \Pi}, (F_i)_{i \in \Pi})$, be an initialized quantitative Reachability game such that its arena is strongly connected. Assume by contradiction that there exists no SPE in (\mathcal{G}, v_0) whose outcome visits all target sets F_i , $i \in \Pi$, that are non-empty. By Theorem 2.4.9, there exists an SPE σ in (\mathcal{G}, v_0) , and we take such an SPE σ whose outcome $\rho = \langle \sigma \rangle_{v_0}$ visits a maximum number of target sets, say $F_{i_1}, F_{i_2}, \dots, F_{i_k}$. Thus by assumption there exists at least one $F_j \neq \emptyset$ with $j \notin \{i_1, \dots, i_k\}$ that is not visited by ρ . Thanks to Lemma B.2.1, we are going to define from σ another SPE τ in (\mathcal{G}, v_0) whose outcome visits all F_{i_1}, \dots, F_{i_k} as well as an additional target set. This will lead to a contradiction.

Consider a prefix $\rho_0 \rho_1 \dots \rho_\ell$ of ρ that visits all F_{i_1}, \dots, F_{i_k} . We denote it by gu with $u = \rho_\ell$. From \mathcal{G} we define the quantitative Reachability game $\mathcal{G}' = (\mathcal{A}, (\text{Cost}'_i)_{i \in \Pi}, (F'_i)_{i \in \Pi})$ with the same arena A and such that $F'_i = \emptyset$ if $i \in \{i_1, \dots, i_k\}$ and $F'_i = F_i$ otherwise ($(\text{Cost}'_i)_{i \in \Pi}$ is defined with respect to $(F'_i)_{i \in \Pi}$). Notice that $F'_j = F_j$ is not empty and it is reachable from u since A is strongly connected. Therefore by Lemma B.2.1, there exists an SPE σ' in (\mathcal{G}', u) that visits at least one target set $F'_{j'}$. From σ and σ' , we define a strategy profile τ in (\mathcal{G}, v_0) as follows: let $h \in \text{Hist}_i(v_0)$,

- if $h = guh'$ for some h' , then $\tau_i(h) = \sigma'_i(uh')$,
- otherwise $\tau_i(h) = \sigma_i(h)$.

Thus, τ acts as σ , except that after a history beginning with gu , it acts as σ' . Clearly the outcome of τ is equal to $g\langle \sigma' \rangle_u$ and thus visits $F'_{j'} = F_{j'}$ in addition to F_{i_1}, \dots, F_{i_k} . It remains to show that τ is an SPE. Consider $hv \in \text{Hist}_i(v_0)$, $i \in \Pi$, and let us show that $\tau_{\upharpoonright h}$ is an NE in $(\mathcal{G}_{\upharpoonright h}, v)$.

- If neither hv is a prefix of gu nor gu is a prefix of hv , then $\tau_{\upharpoonright h} = \sigma_{\upharpoonright h}$ by definition of τ , and $\tau_{\upharpoonright h}$ is an NE in $(\mathcal{G}_{\upharpoonright h}, v)$ because σ is an SPE in (\mathcal{G}, v_0) .

- If gu is a prefix of hv , let h' such that $gh' = h$. Suppose first that hv visits F_i , then Player i has clearly no incentive to deviate in $(\mathcal{G}_{\upharpoonright h}, v)$. Suppose now that hv does not visit F_i , then $i \notin \{i_1, \dots, i_k\}$ and $F'_i = F_i$ by definition of \mathcal{G}' . Hence for all plays π in $(\mathcal{G}_{\upharpoonright h}, v)$ that start in v , $h'\pi$ is a play in (\mathcal{G}', u) that starts in u , and we have $\text{Cost}_i(h\pi) = |gu| + \text{Cost}'_i(h'\pi)$. Hence by definition of τ , a profitable deviation for Player i with respect to $\tau_{\upharpoonright h}$ in $(\mathcal{G}_{\upharpoonright h}, v)$ would be a profitable deviation with respect to $\sigma'_{\upharpoonright h'}$ in $(\mathcal{G}'_{\upharpoonright h'}, v)$. The latter case cannot happen as σ' is an SPE in (\mathcal{G}', u) and it follows that $\tau_{\upharpoonright h}$ is an NE in $(\mathcal{G}_{\upharpoonright h}, v)$.
- Consider the last case where hv is a prefix of gu with $hv \neq gu$, and let $hh' = g$. Consider τ'_i a deviating strategy for Player i with respect to $\tau_{\upharpoonright h}$ in the subgame $(\mathcal{G}_{\upharpoonright h}, v)$, and let $\rho' = \langle (\tau'_i, \tau_{\upharpoonright h, -i}) \rangle_v$. Without loss of generality, we can suppose that $h'u$ is not a prefix of ρ' since this case was treated at the previous item. Notice that if $i \in \{i_1, \dots, i_k\}$, then $\text{Cost}_i(\langle \tau_{\upharpoonright h} \rangle_v) = \text{Cost}_i(\langle \sigma_{\upharpoonright h} \rangle_v)$, otherwise $\text{Cost}_i(\langle \tau_{\upharpoonright h} \rangle_v) \leq +\infty = \text{Cost}_i(\langle \sigma_{\upharpoonright h} \rangle_v)$. In both cases, as $h'u$ is a prefix of both $\langle \tau_{\upharpoonright h} \rangle_v$ and $\langle \sigma_{\upharpoonright h} \rangle_v$, but not a prefix of ρ' , if τ'_i was a profitable deviation for Player i with respect to $\tau_{\upharpoonright h}$, it would also be a profitable deviation with respect to $\sigma_{\upharpoonright h}$ which is impossible since σ is an SPE. \square

B.2.2 Proofs of Section 12.3

Proof of Proposition 12.3.9

Proof of Proposition 12.3.9. Let (\mathcal{G}, v_0) be an initialized multiplayer quantitative Reachability game and let (\mathcal{X}, x_0) be its associated extended game. Let $c \in (\mathbb{N} \cup \{+\infty\})^{|\Pi|}$ be a cost profile and let $M = \max_{i \in \Pi} \{c_i \mid c_i < +\infty\}$ if this maximum exists and $M = 0$ otherwise.

(1 \Rightarrow 2) : Let us assume that there exists an SPE σ in (\mathcal{X}, x_0) such that $\text{Cost}(\langle \sigma \rangle_{x_0}) = c$. Then by Corollary 7.5.2 and Lemma 7.3.2 (applied on the extended game), we have that $\langle \sigma \rangle_{x_0} \in \Lambda^*(x_0)$ and for all $x \in \text{Succ}^*(x_0)$, $\Lambda^*(x) \neq \emptyset$. In the proof of Theorem 7.3.1, we use these sets of λ^* -consistent

plays in order to build a good symbolic witness with $\rho^{(0,x_0)} = \langle \sigma \rangle_{x_0}$.

We now show how to obtain a finite good symbolic witness \mathcal{P} . We begin with $\mathcal{P} = \emptyset$.

Let $\rho = \langle \sigma \rangle_{x_0}$. We apply (P2) on ρ to obtain a lasso $\rho^{(0,x_0)}$ such that the length of this lasso is bounded by $M + |V|$ and $\text{Cost}(\rho^{(0,x_0)}) = c$. Moreover, as ρ is λ^* -consistent, $\rho^{(0,x_0)}$ is also λ^* -consistent (by Lemma 12.3.5)^a. We add $\rho^{(0,x_0)}$ to \mathcal{P} .

For each $(i, x) \in \mathcal{I}$, let ρ be such that $\text{Cost}_i(\rho) = \max\{\text{Cost}_i(\rho') \mid \rho' \in \Lambda^*(x)\}$. We obtain $\rho^{(i,x)}$ by copying ρ until Player i has visited his target set, then by removing the unnecessary cycles and applying (P2). If Player i does not visit his target set along ρ , we remove all the unnecessary cycles (by applying iteratively (P1)) and then we apply (P2). By the same kind of arguments used in Lemma 12.3.3 and Lemma 12.3.5, we obtain that: *i*) $\rho^{(i,x)}$ is λ^* -consistent, *ii*) $\text{Cost}_i(\rho^{(i,x)}) = \text{Cost}_i(\rho)$ and *iii*) the length of the lasso $\rho^{(i,x)}$ is bounded by a value in $\mathcal{O}(|V|^{(|\Pi|+2) \cdot (|V|+3)} + |\Pi| \cdot |V|)$ (by Corollary 11.2.11^b). We add $\rho^{(i,x)}$ to \mathcal{P} .

By construction \mathcal{P} is a finite good symbolic witness.

(2 \Rightarrow 3) : Let us assume that there exists a finite good symbolic witness which satisfies the properties given in Proposition 12.3.9. We have to prove that there exists an SPE τ in (\mathcal{X}, x_0) such that $\text{Cost}(\langle \tau \rangle_{x_0}) = c$ and τ is a finite-memory strategy profile.

Since the lengths of the lassoes are not bounded by the same value L , we do not directly apply Corollary 7.2.6. We have a lasso $\rho^{(0,x_0)}$ of length bounded by $M + |V|$ and at most $|\Pi| \cdot |V| \cdot 2^{|\Pi|}$ lassoes of lengths bounded by a value in $\mathcal{O}(|V|^{(|\Pi|+2) \cdot (|V|+3)} + |\Pi| \cdot |V|)$. Thus if we look at the proof of Corollary 7.2.6, we obtain from these lassoes a weak SPE τ in (\mathcal{X}, x_0) such that $\text{Cost}(\langle \tau \rangle_{x_0}) = c$ and its memory size is in $\mathcal{O}(M + |\Pi| \cdot 2^{|\Pi|} \cdot |V|^{(|\Pi|+2) \cdot (|V|+3)+1})$ (we assume without loss of generality that $|\Pi| \leq |V|$).

(3 \Rightarrow 1) : Obvious. □

^aLet us recall that being λ^* -consistent or Visit λ^* -consistent is equivalent in the extended game (Lemma 7.4.15)

^bEven if this result is based on the labeling functions λ^k computed region by region, we can rely on it since it is used on the fixpoint $\Lambda^*(x)$ which is equivalent for both computations of λ^*

Proof of Proposition 12.3.6 and Proposition 12.3.7

We are now able to prove Propositions 12.3.6 and 12.3.7. We begin by the first one. Recall that it states that if there exists an NE σ (resp. SPE) in a reachability game, then one can construct another one, τ , such that its outcome is a lasso of polynomial length and τ is composed of finite-memory strategies with polynomial (resp. exponential) size. Moreover, if σ is a solution to Problem 3 (resp. Problem 4), it is also the case for τ .

Proof of Proposition 12.3.6. • **For NEs:** Let (\mathcal{G}, v_0) be an multiplayer quantitative Reachability game and σ be an NE in (\mathcal{G}, v_0) . Let $\rho = \langle \sigma \rangle_{v_0}$. We apply procedure (P1) on ρ until there is no longer any unnecessary cycle and then we apply (P2). In this way, we obtain a lasso $\rho' = h\ell^\omega \in \text{Plays}(v_0)$. By Lemma 12.3.3, $|h\ell| \leq (|\Pi| + 1) \cdot |V|$ and $\text{Cost}_i(h\ell^\omega) \leq \min\{\text{Cost}_i(\langle \sigma \rangle_{v_0}), |\Pi| \cdot |V|\}$ if $i \in \text{Visit}(\langle \sigma \rangle_{v_0})$ and $\text{Cost}_i(h\ell^\omega) = +\infty$ otherwise.

By hypothesis and thanks to Theorem 6.2.3, we know that ρ is Visit Val*-consistent. Thus, by Lemma 12.3.5, ρ' is Visit Val*-consistent. Thanks to Corollary 6.2.4, there exists an NE τ such that $\langle \tau \rangle_{v_0} = \rho' = h\ell^\omega$ with memory $\mathcal{O}(|\Pi| \cdot |V|)$ (we can assume without loss of generality that $|V|, |\Pi| \geq 1$).

Let $y \in (\mathbb{N} \cup \{+\infty\})^{|\Pi|}$, let us assume that $(\text{Cost}_i(\langle \sigma \rangle_{v_0}))_{i \in \Pi} \leq y$, by Lemma 12.3.3 we have that $(\text{Cost}_i(\langle \tau \rangle_{v_0})) \leq y$.

Let $k \in \{0, \dots, |\Pi|\}$ and $c \in \mathbb{N} \cup \{+\infty\}$, let us assume that $\text{SW}(\langle \sigma \rangle_{v_0}) \succeq (k, c)$. If $\text{SW}(\langle \sigma \rangle_{v_0}) = (k_1, c_1)$ and $\text{SW}(\langle \tau \rangle_{v_0}) = (k_2, c_2)$, we have $k_1 = k_2$ and $c_2 \leq c_1 \leq c$ thanks to Lemma 12.3.3. Thus, we have that $\text{SW}(\langle \tau \rangle_{v_0}) \succeq (k, c)$.

• **For SPEs:**

Let (\mathcal{G}, v_0) be a quantitative Reachability game and (\mathcal{X}, x_0) be its ex-

tended game. Let σ be an SPE in (\mathcal{X}, x_0) . Let $\rho = \langle \sigma \rangle_{x_0}$. We apply procedure (P1) on ρ until there is no longer any unnecessary cycle and then we apply (P2). In this way, we obtain a lasso $\rho' = h\ell^\omega \in \text{Plays}(x_0)$. By Lemma 12.3.3, $|h\ell| \leq (|\Pi| + 1) \cdot |V|$ and $\text{Cost}_i(h\ell^\omega) \leq \min\{\text{Cost}_i(\langle \sigma \rangle_{x_0}), |\Pi| \cdot |V|\}$ if $i \in \text{Visit}(\langle \sigma \rangle_{x_0})$ and $\text{Cost}_i(h\ell^\omega) = +\infty$ otherwise.

By hypothesis and by Corollary 7.5.2 and Proposition 7.4.15, we know that ρ is Visit λ^* -consistent. Thus, by Lemma 12.3.5, ρ' is Visit λ^* -consistent. Finally, thanks to Proposition 7.4.15, Corollary 7.5.2 and Proposition 12.3.9, there exists an SPE τ such that $(\text{Cost}_i(\langle \tau \rangle_{x_0}))_{i \in \Pi} = (\text{Cost}_i(\rho'))_{i \in \Pi}$ with memory $\mathcal{O}(2^{|\Pi|} \cdot |\Pi| \cdot |V|^{(|\Pi|+2) \cdot (|V|+3)+1})$.

Let $y \in (\mathbb{N} \cup \{+\infty\})^{|\Pi|}$, let us assume that $(\text{Cost}_i(\langle \sigma \rangle_{x_0}))_{i \in \Pi} \leq y$, by Lemma 12.3.3 we have that $(\text{Cost}_i(\langle \tau \rangle_{x_0}))_{i \in \Pi} = (\text{Cost}_i(\rho'))_{i \in \Pi} \leq y$.

Let $k \in \{0, \dots, |\Pi|\}$ and $c \in \mathbb{N} \cup \{+\infty\}$, let us assume that $\text{SW}(\langle \sigma \rangle_{x_0}) \succeq (k, c)$. If $\text{SW}(\langle \sigma \rangle_{x_0}) = (k_1, c_1)$ and $\text{SW}(\langle \tau \rangle_{x_0}) = (k_2, c_2)$, we have $k_1 = k_2$ and $c_2 \leq c_1 \leq c$ thanks to Lemma 12.3.3. Thus, we have that $\text{SW}(\langle \tau \rangle_{x_0}) \succeq (k, c)$. □

Finally, we prove Proposition 12.3.7. Recall that this proposition states that if there exists an NE σ (resp. SPE) whose outcome is Pareto optimal, then one can construct another one, τ , such that its outcome is a lasso of polynomial length, has the same cost as σ (thus is also Pareto optimal), and τ uses finite-memory strategies with polynomial (resp. exponential) size.

Proof of Proposition 12.3.7. The second and third items are a direct consequence of the first one. Thus, let us prove the first item.

Let σ be an NE such that its cost profile is Pareto optimal in $\text{Plays}(v_0)$. To get a contradiction, assume that there exists $i \in \text{Visit}(\langle \sigma \rangle_{v_0})$ such that $\text{Cost}_i(\langle \sigma \rangle_{v_0}) > |V| \cdot |\Pi|$. It means that there exists an unnecessary cycle before Player i reaches his target set. By removing this cycle (applying (P1)), we obtain a new play ρ' such that $\text{Cost}_i(\rho') < \text{Cost}_i(\langle \sigma \rangle_{v_0})$ and for Player j ($j \neq i$), $\text{Cost}_j(\rho') \leq \text{Cost}_j(\langle \sigma \rangle_{v_0})$ (by Lemma 12.3.3). It leads

to a contradiction with the fact that $(\text{Cost}_i(\langle\sigma\rangle_{v_0}))_{i \in \Pi}$ is Pareto optimal in $\text{Plays}(v_0)$.

The same proof holds for SPE. \square

B.3 Proofs of Chapter 15

This section is devoted to the proofs of Chapter 15. For the sake of clarity, we denote $\text{Plays}_{\bar{A}}$, $\text{Hist}_{\bar{A}}$ and $\text{Hist}_{i, \bar{A}}$ by $\widetilde{\text{Plays}}$, $\widetilde{\text{Hist}}$ and $\widetilde{\text{Hist}}_i$ respectively.

B.3.1 Proof of Proposition 15.3.3

In this section, when we consider a history $h = h_0 \dots h_n$ for some $n \in \mathbb{N}$, the length of h , denoted by $|h|$, is its number of vertices.

This section is devoted to the proof of Proposition 15.3.3. Let $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \Pi})$ be a game and \sim be a bisimulation equivalence on (\mathcal{G}, v_0) which respects the partition and such that for each ρ and ρ' in Plays , if $\rho \sim \rho'$ then $\text{Gain}(\rho) = \text{Gain}(\rho')$.

If there exists an SPE τ in (\mathcal{G}, v_0) which is uniform and such that $\text{Gain}(\langle\tau\rangle_{v_0}) = p$, clearly there exists an SPE σ in (\mathcal{G}, v_0) such that $\text{Gain}(\langle\sigma\rangle_{v_0}) = p$.

The difficult part is the other implication: if there exists an SPE σ in (\mathcal{G}, v_0) such that $\text{Gain}(\langle\sigma\rangle_{v_0}) = p$, then there exists an SPE τ in (\mathcal{G}, v_0) which is uniform and such that $\text{Gain}(\langle\tau\rangle_{v_0}) = p$. Let us prove it.

Let σ be an SPE in (\mathcal{G}, v_0) such that $\text{Gain}(\langle\sigma\rangle_{v_0}) = p$. In order to build τ , we need some additional material and notations that we explain below.

- for each $h \in \text{Hist}(v_0)$: $[h] = \{h' \in \text{Hist}(v_0) \mid h \sim h'\}$;
- $\mathcal{C}^n = \{[h] \mid h \in \text{Hist}(v_0) \wedge |h| = n\}$;
- $\mathcal{R} : \bigcup_{n \in \mathbb{N}} \mathcal{C}^n \rightarrow \text{Hist}(v_0) \cup \{\perp\}$ which allow us to indentify a witness for each class;
- $P : \text{Hist}(v_0) \rightarrow \{0, 1\}$

Inductive construction of \mathcal{R} and P

The first step, is to choose in a proper way a witness to each each class $[h]$. We proceed by induction on the length of histories. Moreover, we claim that the following properties are satisfied all along the inductive construction.

Invariant 1: For each $hv \in \text{Hist}(v_0)$ such that $\mathcal{R}([hv]) \neq \perp$,

$$hv \sim \mathcal{R}([hv]).$$

Invariant 2: For each $hv \in \text{Hist}(v_0)$ such that $\mathcal{R}([hv]) \neq \perp$ and $|hv| > 1$,

$$\mathcal{R}([hv]) = \mathcal{R}([h]) \text{Last}(\mathcal{R}([hv])).^1$$

Invariant 3: For each $hv \in \text{Hist}(v_0)$ such that $\mathcal{R}([hv]) \neq \perp$,

$$h'v' \leq \mathcal{R}([hv]) < h'\langle \sigma_{\uparrow h'} \rangle_{v'}$$

for some $h'v'$ such that $P(h'v') = 1$.

Before beginning the induction, we initialize P and \mathcal{R} in the following way: for all $C \in \bigcup_{n \in \mathbb{N}} \mathcal{C}^n$, $\mathcal{R}(C) = \perp$ and for all $h \in \text{Hist}(v_0)$, $P(h) = 0$.

- For $n = 1$: $\mathcal{C}^1 = \{[v_0]\}$, we define $P(v_0) = 1$. Then , for each h such that $v_0 \leq h < \langle \sigma \rangle_{v_0}$, we define $\mathcal{R}([h]) = h$. Thus, Invariant 3 is satisfied with $h'v' = v_0$ and for each $v_0 < hv < \langle \sigma \rangle_{v_0}$, $\mathcal{R}([hv])$ is defined in this step and satisfies Invariant 2. Since $\mathcal{R}([h]) = h$ for each witness defined in this step, Invariant 1 is satisfied too.
- Let us assume that these two invariant are satisfied after step k , and let us prove it remains true after step $k + 1$.
- In this step, we first define \mathcal{R} for each $C \in \mathcal{C}^{k+1}$ such that $\mathcal{R}(C) = \perp$. We know that for all $h_1v_1, h_2v_2 \in C$, $h_1 \sim h_2$ and $\mathcal{R}([h_1]) = \mathcal{R}([h_2])$ are already defined (*i.e.*, $\neq \perp$). Moreover, by Invariant 1, $h_1 \sim \mathcal{R}([h_1])$, let $h = \mathcal{R}([h_1])$, by bisimulation \sim , there exists $v \in V$ such that $h_1v_1 \sim hv$. We define $P(hv) = 1$ and $\mathcal{R}(C) = hv$. Then, for all $h_2v_2 \in C$,

¹If $h = h_0 \dots h_n$ for some $n \in \mathbb{N}$, $\text{Last}(h) = h_n$.

$h_2v_2 \sim h_1v_1 \sim hv$ this implies that $h_2v_2 \sim \mathcal{R}([h_2v_2])$ (Inv 1 ok). For all $h_2v_2 \in C$, $\mathcal{R}([h_2v_2]) = \mathcal{R}(C) = hv = \mathcal{R}([h_1])v = \mathcal{R}([h_2])v = \mathcal{R}([h_2]) \text{Last}(\mathcal{R}([h_2v_2]))$ (Inv 2 ok). Moreover, for all $h_2v_2 \in C$, $\mathcal{R}([h_2v_2]) = hv$ and $P(hv) = 1$, since $hv \leq hv < h\langle\sigma_{\uparrow h}\rangle_v$ (Inv 3 ok).

Now, we extend the construction of \mathcal{R} and P from hv in the following way:

$\forall h'v' \in \text{Hist}(v_0)$ st. $hv < h'v' < h\langle\sigma_{\uparrow h}\rangle_v$ we define $\mathcal{R}([h'v']) = h'v'$. \star

Now, we have to prove that the invariants remains satisfied for all these new defined classes.

- $\forall \bar{h}\bar{v} \in [h'v']$: $\bar{h}\bar{v} \sim h'v' = \mathcal{R}([h'v']) = \mathcal{R}([\bar{h}\bar{v}])$ (Inv 1 ok);
- $\forall \bar{h}\bar{v} \in [h'v']$, we have that $\bar{h}\bar{v} \sim h'v'$ thus: $\mathcal{R}([\bar{h}\bar{v}]) = \mathcal{R}([h'v']) = \mathcal{R}([h']) \text{Last}(\mathcal{R}([h'v']))$ (by construction \star). Thus, since $\bar{h} \sim h'$: $\mathcal{R}([h']) = \mathcal{R}([\bar{h}])$ (Inv 2 ok).
- $\forall \bar{h}\bar{v} \in [h'v']$, we have by construction \star that $P(hv) = 1$ and $hv < h'v' < h\langle\sigma_{\uparrow h}\rangle_v$. Since $h'v' = \mathcal{R}([h'v'])$ and $\mathcal{R}([h'v']) = \mathcal{R}([\bar{h}\bar{v}])$ ($h'v' \sim \bar{h}\bar{v}$), we are done (Inv 3 ok).

Construction of τ

To build the uniform strategy profile τ , we proceed as follows: for all $n \in \mathbb{N}$, for all $C \in \mathcal{C}^n$, for all $h \in C$, by assuming that $\text{Last}(h) \in V_i$:

- If $\mathcal{R}([h]) = h$ (h is a witness, thus we want to follow σ): $\tau_i(h) = \sigma_i(h)$;
- If $\mathcal{R}([h]) \neq h$ (we simulate σ): we know by Invariant 1 that $h \sim \mathcal{R}([h])$, thus in particular $\text{Last}(h) \sim \text{Last}(\mathcal{R}([h]))$, by bisimulation \sim , there exists $x \in V$ such that $\text{Last}(h)x \sim \text{Last}(\mathcal{R}([h]))\sigma_i(\mathcal{R}([h]))$. Thus, we define $\tau_i(h) = x$.

We state now, some properties about τ and σ . First, we define $\mathcal{P}^2 = \{h \in \text{Hist}(v_0) \mid \exists C \in \bigcup_{n \in \mathbb{N}} \mathcal{C}^n \text{ st. } \mathcal{R}(C) = h\}$.

²Notice that it is not the same \mathcal{P} as the one used to depict a symbolic witness in the previous parts of this document.

Lemma B.3.1. For all $h \in \text{Hist}(v_0)$ such that $h \in \mathcal{P}$ and $\text{Last}(h) \in V_i$: $\tau_i(h) = \sigma_i(h)$.

Proof. This assertion is true due to the construction of τ . \square

Lemma B.3.2. For all $h, h' \in \text{Hist}(v_0)$ such that $h \sim h'$: $\tau_i(h) \sim \tau_i(h')$ by assuming that $\text{Last}(h) \in V_i$.

Notice that, since \sim respects the partition, if $\text{Last}(h) \in V_i$ then $\text{Last}(h') \in V_i$, and vice versa.

Proof. Let $h, h' \in \text{Hist}(v_0)$ such that $h \sim h'$ and $\text{Last}(h) \in V_i$ for some $i \in \Pi$ then $\text{Last}(h') \in V_i$. We have that $\mathcal{R}([h]) = \mathcal{R}([h'])$. By construction, $\tau_i(h) \sim \sigma_i(\mathcal{R}([h]))$ and $\tau_i(h') \sim \sigma_i(\mathcal{R}([h']))$, by transitivity, we have: $\tau_i(h) \sim \tau_i(h')$. \square

Lemma B.3.3. For all $h \in \mathcal{P}$, $h\tau_i(h) \in \mathcal{P}$ (by assuming that $\text{Last}(h) \in V_i$ for some $i \in \Pi$).

Proof. Let $h \in \mathcal{P}$, such that $\text{Last}(h) \in V_i$ for some $i \in \Pi$. Since $h \in \mathcal{P}$, by Invariant 3, there exists $h'v' \in \text{Hist}(v_0)$ such that:

$$h'v' \leq h < h'\langle\sigma_{\uparrow h'}\rangle_{v'}.$$

Thus, we have that

$$h'v' \leq h\sigma_i(h) < h'\langle\sigma_{\uparrow h'}\rangle_{v'}.$$

It follows by construction of \mathcal{R} , that $h\sigma_i(h) \in \mathcal{P}$. Moreover, $h \in \mathcal{P}$ implies that $\tau_i(h) = \sigma_i(h)$ (by Lemma B.3.1). Thus, $h\tau_i(h) \in \mathcal{P}$. \square

Lemma B.3.4. For all $hv \in \mathcal{P}$, $\langle \sigma_{|h} \rangle_v = \langle \tau_{|h} \rangle_v$.

Proof. Let $hv \in \mathcal{P}$, let $\rho = \langle \sigma_{|h} \rangle_v$ and let $\bar{\rho} = \langle \tau_{|h} \rangle_v$. Let us prove by induction that for all $n \in \mathbb{N}$:

1. $\rho_n = \bar{\rho}_n$;
2. $h\rho_0 \dots \rho_n \in \mathcal{P}$.

For $n = 0$, $\rho_0 = v$ and $\bar{\rho}_0 = v$. And by hypothesis, $hv \in \mathcal{P}$. Let us assume that both assertions are satisfied for all n such that $n \leq k$. Let us prove that it remains true for $n = k + 1$. By assuming that $\bar{\rho}_k \in V_i$,

1.

$$\begin{aligned} \bar{\rho}_{k+1} &= \tau_i(h\bar{\rho}_0 \dots \bar{\rho}_k) \\ &= \tau_i(h\rho_0 \dots \rho_k) && \text{By IH, } \bar{\rho}_0 \dots \bar{\rho}_k = \rho_0 \dots \rho_k \\ &= \sigma_i(h\rho_0 \dots \rho_k) && \text{By IH, } h\rho_0 \dots \rho_k \in \mathcal{P} \text{ and by Lemma B.3.1} \\ &= \rho_{k+1}. \end{aligned}$$

2. By IH, $h\rho_0 \dots \rho_k \in \mathcal{P}$, moreover we have that, by Lemma B.3.1:

$$h\rho_0 \dots \rho_k \rho_{k+1} = h\rho_0 \dots \rho_k \sigma_i(h\rho_0 \dots \rho_k) = h\rho_0 \dots \rho_k \tau_i(h\rho_0 \dots \rho_k)$$

And by Lemma B.3.3, we can conclude that $h\rho_0 \dots \rho_k \tau_i(h\rho_0 \dots \rho_k) \in \mathcal{P}$. \square

Proof that τ is an uniform SPE with gain profile p

There is still to prove that τ is an uniform SPE in (\mathcal{G}, v_0) such that $\text{Gain}(\langle \tau \rangle_{v_0}) = p$. By Lemma B.3.2, τ is uniform, let us prove this is an SPE with the gain profile p .

Proof. First, since $v_0 \in \mathcal{P}$ and by Lemma B.3.4, we have that $\langle \sigma \rangle_{v_0} = \langle \tau \rangle_{v_0}$. Thus, in particular, $\text{Gain}(\langle \tau \rangle_{v_0}) = \text{Gain}(\langle \sigma \rangle_{v_0}) = p$.

By absurdum, let us assume that τ is not an SPE in (\mathcal{G}, v_0) . It means that there exist $hv \in \text{Hist}(v_0)$, $i \in \Pi$ and a strategy τ'_i of Player i in $(\mathcal{G}_{|h}, v)$ such that τ'_i is a profitable deviation of $\tau_{i|h}$, *i.e.*,

$$\text{Gain}_i(h\langle \tau_{|h} \rangle_v) < \text{Gain}_i(h\langle \tau'_i, \tau_{-i|h} \rangle_v). \quad (\text{B.22})$$

Let $h'v' = \mathcal{R}([hv]) = \mathcal{R}([h]) \text{Last}(\mathcal{R}([hv]))$ by Invariant 2.

First step: Let $\rho = \langle \tau_{|h} \rangle_v$ and $\rho' = \langle \tau_{|h'} \rangle_{v'}$, let us prove by induction that for all $n \in \mathbb{N}$:

1. $\rho_n \sim \rho'_n$;
2. $h'\rho'_0 \dots \rho'_n \in \mathcal{P}$.

For $n = 0$, we have that $\rho_0 = v$ and $\rho'_0 = v'$, thus $v \sim v'$ since $hv \sim h'v'$. Moreover, $h'v' \in \mathcal{P}$ by hypothesis. Let us assume that these two properties are satisfied for all n such that $n \leq k$, let us prove they remain true for $n = k + 1$. Let us assume that $\rho_k \in V_j$ for some $j \in \Pi$, since \sim respects the partition and due to the fact that $\rho_k \sim \rho'_k$ by IH, we have that $\rho'_k \in V_j$.

1.

$$\begin{aligned} \rho_{k+1} &= \tau_j(h\rho_0 \dots \rho_k) \\ &\sim \tau_j(h'\rho'_0 \dots \rho'_k) \quad \text{By IH, } h\rho_0 \dots \rho_k \sim h'\rho'_0 \dots \rho'_k \text{ and Lem. B.3.2} \\ &= \rho'_{k+1} \end{aligned}$$

2. $h'\rho'_0 \dots \rho'_k \rho'_{k+1} = h'\rho'_0 \dots \rho'_k \tau_i(h'\rho'_0 \dots \rho'_k)$, $h'\rho'_0 \dots \rho'_k \in \mathcal{P}$ by IH, thus by Lemma B.3.3: $h'\rho'_0 \dots \rho'_k \tau_i(h'\rho'_0 \dots \rho'_k) \in \mathcal{P}$.

It allows us to state by (1) that $h\rho \sim h'\rho'$, thus by hypothesis on \sim , we have that

$$\text{Gain}(h\langle \tau_{|h} \rangle_v) = \text{Gain}(h'\langle \tau_{|h'} \rangle_{v'}). \quad (\text{B.23})$$

By (2) and Lemma B.3.1, we have that $\langle \tau_{|h'} \rangle_{v'} = \langle \sigma_{|h'} \rangle_{v'}$ and thus:

$$\text{Gain}(h' \langle \tau_{|h'} \rangle_{v'}) = \text{Gain}(h' \langle \sigma_{|h'} \rangle_{v'}). \quad (\text{B.24})$$

Second step: Let $\rho = \langle \tau'_i, \tau_{-i|h} \rangle_v$, we will build a strategy $\tilde{\tau}_i$ in $(\mathcal{G}_{|h'}, v')$ such that $\rho \sim \langle \tilde{\tau}_i, \tau_{-i|h'} \rangle_{v'}$. Let $\tilde{p} \in \text{Hist}_i(v')$ and let us assume that $\tilde{p} = \tilde{p}_0 \dots \tilde{p}_m$ for some $m \in \mathbb{N}$.

- If $\tilde{p} \sim \rho_0 \dots \rho_m$, then $\rho_m \in V_i$ (\sim respects the partition) and $\rho_{m+1} = \tau'_i(\rho_0 \dots \rho_m)$. Thus, by \sim there exists $x \in V$ such that $\tilde{p}x \sim \rho_0 \dots \rho_m \rho_{m+1}$. We define $\tilde{\tau}_i(\tilde{p}) = x$. Thus, $\tilde{\tau}_i(\tilde{p}_0 \dots \tilde{p}_m) \sim \tau'_i(\rho_0 \dots \rho_m)$.
- Otherwise, we define $\tilde{\tau}_i(\tilde{p}) = x$ for some $x \in \text{Succ}(\tilde{p}_m)$.

Let $\tilde{\rho} = \langle \tilde{\tau}_i, \tau_{-i|h'} \rangle_{v'}$, let us prove that for all $n \in \mathbb{N}$, $\tilde{\rho}_n \sim \rho_n$. For $n = 0$, $\tilde{\rho}_0 = v'$ and $\rho_0 = v$, since $h'v' \sim hv$, $v' \sim v$. Let us assume that this property is true for all $n \leq k$ and let us prove it remains true for $n = k + 1$.

- If $\tilde{\rho}_k \in V_i$, then since $\tilde{\rho}_k \sim \rho_k$ by IH, $\rho_k \in V_i$ (\sim respects the partition). It follows:

$$\begin{aligned} \tilde{\rho}_{k+1} &= \tilde{\tau}_i(\tilde{\rho}_0 \dots \tilde{\rho}_k) \\ &\sim \tau'_i(\rho_0 \dots \rho_k) \quad \text{By IH, } \tilde{\rho}_0 \dots \tilde{\rho}_k \sim \rho_0 \dots \rho_k \text{ and by constr. of } \tilde{\tau}_i \\ &= \rho_{k+1}. \end{aligned}$$

- If $\tilde{\rho}_k \in V_j$ ($j \neq i$), as previously $\rho_k \in V_j$. Thus:

$$\begin{aligned} \tilde{\rho}_{k+1} &= \tau_{j|h'}(\tilde{\rho}_0 \dots \tilde{\rho}_k) = \tau_j(h' \tilde{\rho}_0 \dots \tilde{\rho}_k) \\ &\sim \tau_j(h \rho_0 \dots \rho_k) \quad \text{By IH, } h' \tilde{\rho}_0 \dots \tilde{\rho}_k \sim h \rho_0 \dots \rho_k \text{ and Lem. B.3.2} \\ &= \rho_{k+1}. \end{aligned}$$

From this we have that $h\rho \sim h'\tilde{\rho}$ and in particular:

$$\text{Gain}_i(h \langle \tau'_i, \tau_{-i|h} \rangle_v) = \text{Gain}_i(h' \langle \tilde{\tau}_i, \tau_{-i|h'} \rangle_{v'}). \quad (\text{B.25})$$

Third step: From $\tilde{\tau}_i$, we build $\tilde{\sigma}_i$ in $(\mathcal{G}_{|h'}, v')$ which is a profitable deviation of $\sigma_{i|h'}$. Let $p \in \text{Hist}_i(v')$

- If $h'p \in \mathcal{P}$, we consider $\mathcal{R}([h'p\tilde{\tau}_i(p)]) = \mathcal{R}([h'p]) \text{Last}(\mathcal{R}([h'p\tilde{\tau}_i(p)]))$ by Invariant 2. Let $x = \text{Last}(\mathcal{R}([h'p\tilde{\tau}_i(p)]))$. We define $\tilde{\sigma}_i(p) = x$, in particular, we have that $\mathcal{R}([h'p\tilde{\tau}_i(p)]) = \mathcal{R}([h'p])\tilde{\sigma}_i(p)$ and thus $\mathcal{R}([h'p])\tilde{\sigma}_i(p) \in \mathcal{P}$.
- If $h'p \notin \mathcal{P}$, we define $\tilde{\sigma}_i(p) = \tilde{\tau}_i(p)$.

Let $\pi = \langle \tilde{\sigma}_i, \sigma_{-i|h'} \rangle_{v'}$ and $\pi' = \langle \tilde{\sigma}_i, \tau_{-i|h'} \rangle_{v'}$. Let us prove that for all $n \in \mathbb{N}$:

1. $\pi_n = \pi'_n$;
2. $h'\pi_0 \dots \pi_n \in \mathcal{P}$

For $n = 0$, we have that $\pi_0 = v' = \pi'_0$. Moreover, $h'v' \in \mathcal{P}$ by hypothesis. Let us assume that these two properties are true for all $n \leq k$ and let us prove that they remain true for $n = k + 1$.

- If $\pi_k \in V_i$, then by IH, $\pi_k = \pi'_k \in V_i$.

1.

$$\begin{aligned} \pi_{k+1} &= \tilde{\sigma}_i(\pi_0 \dots \pi_k) \\ &= \tilde{\sigma}_i(\pi'_0 \dots \pi'_k) && \text{By IH, } \pi_0 \dots \pi_k = \pi'_0 \dots \pi'_k. \\ &= \pi'_{k+1} \end{aligned}$$

2.

$$\begin{aligned} h'\pi_0 \dots \pi_k \pi_{k+1} &= h'\pi_0 \dots \pi_k \tilde{\sigma}_i(\pi_0 \dots \pi_k) \\ &= \mathcal{R}([h'\pi_0 \dots \pi_k])\tilde{\sigma}_i(\pi_0 \dots \pi_k) && \text{By IH, } h'\pi_0 \dots \pi_k \in \mathcal{P} \\ &\in \mathcal{P} && \text{By constr. of } \tilde{\sigma}_i. \end{aligned}$$

- If $\pi_k \in V_j$ ($j \neq i$), then by IH, $\pi_k = \pi'_k \in V_j$.

1.

$$\begin{aligned}
\pi_{k+1} &= \sigma_j(h'\pi_0 \dots \pi_k) \\
&= \tau_j(h'\pi_0 \dots \pi_k) && \text{By IH, } h'\pi_0 \dots \pi_k \in \mathcal{P} \text{ and by Lem. B.3.1} \\
&= \tau_j(h'\pi'_0 \dots \pi'_k) && \text{By IH} \\
&= \pi'_{k+1}.
\end{aligned}$$

2.

$$\begin{aligned}
h'\pi_0 \dots \pi_k \pi_{k+1} &= h'\pi_0 \dots \pi_k \sigma_j(h'\pi_0 \dots \pi_k) \\
&= h'\pi_0 \dots \pi_k \tau_j(h'\pi_0 \dots \pi_k) && \text{(By IH, } h'\pi_0 \dots \pi_k \in \mathcal{P} \\
&\in \mathcal{P} && \text{and by Lemma B.3.1)} \\
&&& \text{(By Lemma B.3.3)}
\end{aligned}$$

Thus, we can conclude that:

$$\text{Gain}_i(h'\langle \tilde{\sigma}_i, \sigma_{-i|h'} \rangle_{v'}) = \text{Gain}_i(h'\langle \tilde{\sigma}_i, \tau_{-i|h'} \rangle_{v'}). \quad (\text{B.26})$$

Now, we want to prove that $\pi' = \langle \tilde{\sigma}_i, \tau_{-i|h'} \rangle_{v'} \sim \tilde{\rho} = \langle \tilde{\tau}_i, \tau_{-i|h'} \rangle_{v'}$. Let us recall, that from the second step, we know that $\tilde{\rho} \sim \rho = \langle \tau'_i, \tau_{-i|h} \rangle_v$. Let us prove that for all $n \in \mathbb{N}$: $\pi'_n \sim \tilde{\rho}_n$.

For $n = 0$: $\pi'_0 = v' = \tilde{\rho}_0$. Let us assume that this property is true for all $n \leq k$ and let us prove that it remains true for $n = k + 1$.

- If $\pi'_k \in V_i$ then, by IH we have that $\pi'_k \sim \tilde{\rho}_k$ and so $\tilde{\rho}_k \in V_i$.

$$\begin{aligned}
\pi'_{k+1} &= \tilde{\sigma}_i(\pi'_0 \dots \pi'_k) \\
&= \text{Last}(\mathcal{R}([h'\pi'_0 \dots \pi'_k \tilde{\tau}_i(\pi'_0 \dots \pi'_k)])) && (h'\pi'_0 \dots \pi'_k \in \mathcal{P}) \\
&\sim \tilde{\tau}_i(\pi'_0 \dots \pi'_k)
\end{aligned}$$

By IH, we know that $\pi'_0 \dots \pi'_k \sim \tilde{\rho}_0 \dots \tilde{\rho}_k$ and by hypothesis, we have that $\tilde{\rho}_0 \dots \tilde{\rho}_k \sim \rho_0 \dots \rho_k$. It follows from the construction of $\tilde{\tau}_i$ that $\tilde{\tau}_i(\pi'_0 \dots \pi'_k) \sim \tau'_i(\rho_0 \dots \rho_k)$ and $\tilde{\tau}_i(\tilde{\rho}_0 \dots \tilde{\rho}_k) \sim \tau'_i(\rho_0 \dots \rho_k)$. Thus, by transitivity, $\pi'_{k+1} \sim \tilde{\tau}_i(\tilde{\rho}_0 \dots \tilde{\rho}_k) = \tilde{\rho}_{k+1}$.

- If $\pi'_k \in V_j$ ($j \neq i$) then as previously $\tilde{\rho}_k \in V_j$.

$$\begin{aligned}
\pi'_{k+1} &= \tau_j(h'\pi'_0 \dots \pi'_k) \\
&\sim \tau_j(h'\tilde{\rho}_0 \dots \tilde{\rho}_k) && \text{(By IH, } h'\pi'_0 \dots \pi'_k \sim h'\tilde{\rho}_0 \dots \tilde{\rho}_k \\
& && \text{and by Lemma B.3.2)} \\
&= \tilde{\rho}_{k+1}.
\end{aligned}$$

Thus $h'\pi' \sim h'\tilde{\rho}$ and it follows that:

$$\text{Gain}_i(h'\langle \tilde{\sigma}_i, \tau_{-i|h'} \rangle_{v'}) = \text{Gain}_i(h'\langle \tilde{\tau}_i, \tau_{-i|h'} \rangle_{v'}). \quad (\text{B.27})$$

Fourth step: putting all together

By (B.22), (B.23), (B.24), (B.25), (B.26) and (B.27), we can conclude that

$$\text{Gain}_i(h'\langle \sigma_{|h'} \rangle_{v'}) < \text{Gain}_i(h'\langle \tilde{\sigma}_i, \sigma_{-i|h'} \rangle_{v'}).$$

Thus, there exists a profitable deviation of $\sigma_{i|h'}$ for Player i in $(\mathcal{G}_{|h'}, v')$. This is impossible, since σ is an SPE in (\mathcal{G}, v_0) . □

B.3.2 Proof of Theorem 15.3.1

In this section we prove Theorem 15.3.1. In order to do so, we prove the two implications of the equivalence in two different propositions.

Proposition B.3.5. *Let $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \pi})$ be a game and $(\tilde{\mathcal{G}}, [v_0]) = (\tilde{A}, (\tilde{\text{Gain}}_i)_{i \in \Pi})$ its associated quotient game where \sim is a bisimulation equivalence on (\mathcal{G}, v_0) . If \sim respects the partition and the gain functions, we have that: if there exists an SPE σ in (\mathcal{G}, v_0) such that $\text{Gain}(\langle \sigma \rangle_{v_0}) = p$ for some $p \in \{0, 1\}^{|\Pi|}$ then there exists an SPE τ in $(\tilde{\mathcal{G}}, [v_0])$ such that $\tilde{\text{Gain}}(\langle \tau \rangle_{[v_0]}) = p$.*

Proof. Let $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \pi})$ be a game and $(\tilde{\mathcal{G}}, [v_0]) = (\tilde{A}, (\tilde{\text{Gain}}_i)_{i \in \Pi})$ its associated quotient game where \sim is a bisimulation equivalence on (\mathcal{G}, v_0) which respects the partition and the gain functions. We assume that there exists an SPE σ in (\mathcal{G}, v_0) such that $\text{Gain}(\langle \sigma \rangle_{v_0}) = p$ for some $p \in \{0, 1\}$.

Without loss of generality, we can assume thanks to Proposition 15.3.3 that σ is uniform, *i.e.*, for all histories $h, h' \in \text{Hist}(v_0)$ such that $\text{Last}(h) \in V_i \Leftrightarrow \text{Last}(h') \in V_i$, $\sigma_i(h) \sim \sigma_i(h')$.

Let $\tilde{h} \in \widetilde{\text{Hist}}([v_0])$ be a history in the quotient game, by bisimulation $\sim_q \subseteq V \times \tilde{V}$, there exists $h = h_0 \dots h_n \in \text{Hist}(v_0)$ such that $h \sim_q \tilde{h} = [h_0] \dots [h_n]$. Let $v \in V$ be the vertex such that $\sigma_i(h) = v$, by assuming that $\text{Last}(h) \in V_i$. We have that $v \sim_q [v]$ and, we define $\tau_i(\tilde{h}) = [v]$.

CLAIM 1: $\forall \tilde{h} \in \widetilde{\text{Hist}}([v_0]), \forall h \in \text{Hist}(v_0)$ such that $h \sim_q \tilde{h}$, if $\text{Last}(h) \in V_i$, $\sigma_i(h) \sim_q \tau_i(\tilde{h})$.

Proof: Let $\tilde{h} \in \widetilde{\text{Hist}}([v_0])$ and $h \in \text{Hist}(v_0)$ such that $\text{Last}(h) \in V_i$ for some $i \in \Pi$ and $h \sim_q \tilde{h}$. By construction of τ , there exists $h' \in \text{Hist}_i(v_0)$ such that $h' \sim_q \tilde{h}$ and $\tau_i(\tilde{h}) = [\sigma_i(h')]$. If $h \sim_q \tilde{h}$ and $h' \sim_q \tilde{h}$, we have that $h \sim h'$. Thus, by uniformity of σ , $\sigma_i(h) \sim \sigma_i(h')$ and in particular $[\sigma_i(h)] = [\sigma_i(h')]$. In conclusion, $\sigma_i(h) \sim_q [\sigma_i(h)] = [\sigma_i(h')] = \tau_i(\tilde{h})$. \square

Let $\rho = \langle \sigma \rangle_{v_0}$ and $\tilde{\rho} = \langle \tau \rangle_{[v_0]}$. Let us prove that: $\forall n \in \mathbb{N} \rho_n \sim_q \tilde{\rho}_n$. Thus, $\rho \sim_q \tilde{\rho}$ and since \sim respects the gain functions, $\text{Gain}(\rho) = \tilde{\text{Gain}}(\tilde{\rho}) = p$.

For $n = 0$: $\rho_0 = v_0 \sim_q [v_0] = \tilde{\rho}_0$. We assume that this is true for all $n \leq k$ and we prove it remains true for $n = k + 1$. By induction hypothesis, we have that $\rho_0 \dots \rho_k \sim_q \tilde{\rho}_0 \dots \tilde{\rho}_k$. By \sim_q , $\rho_k \in V_i$ if and only if $\tilde{\rho}_k \in \tilde{V}_i$. Thus, $\rho_{k+1} = \sigma_i(\rho_0 \dots \rho_k) \sim_q \tau_i(\tilde{\rho}_0 \dots \tilde{\rho}_k) = \tilde{\rho}_{k+1}$ by Claim 1.

To conclude, we have to prove that τ is an SPE in $(\tilde{\mathcal{G}}, [v_0])$. Ad absurdum, we assume that there exists $\tilde{h}\tilde{v} \in \widetilde{\text{Hist}}([v_0])$ such that there exists a player

$i \in \Pi$ and a profitable deviation τ'_i of $\tau_{i|\tilde{h}}$ in $(\tilde{\mathcal{G}}_{|\tilde{h}}, \tilde{v})$, *i.e.*,

$$\tilde{\text{Gain}}_i(\tilde{h}\langle\tau_{|\tilde{h}}\rangle_{\tilde{v}}) < \tilde{\text{Gain}}_i(\tilde{h}\langle\tau'_i, \tau_{-i|\tilde{h}}\rangle_{\tilde{v}}). \quad (\text{B.28})$$

By bisimulation \sim_q , there exists $hv \in \text{Hist}(v_0)$ such that $hv \sim_q \tilde{h}\tilde{v}$. We prove that Player i has a profitable deviation of $\sigma_{i|h}$ in $(\mathcal{G}_{|h}, v)$. From which a contradiction follows since σ has to be an SPE in (\mathcal{G}, v_0) .

We build the profitable deviation σ'_i . Let $p \in \text{Hist}_i(v)$ a history such that $p = vp_1p_2 \dots p_m$ for some $m \in \mathbb{N}$. By bisimulation \sim_q , there exists a unique \tilde{p} such that $p = [v][p_1] \dots [p_m]$ and thus $p \sim_q \tilde{p}$. Let $r \in \tilde{V}$ be such that $\tau'_i(\tilde{p}) = r$. By bisimulation \sim_q , there exists $x \in V$ such that $\rho_m x \sim_q [\rho_m]r$. We define $\sigma'_i(p) = x$. In particular, $\sigma'_i(p) \sim_q \tau'_i(\tilde{p})$.

Let $\rho = \langle\sigma'_i, \sigma_{-i|h}\rangle_v = v\rho_1\rho_2 \dots$ and $\tilde{\rho} = \langle\tau'_i, \tau_{-i|\tilde{h}}\rangle_{\tilde{v}} = \tilde{v}\tilde{\rho}_1\tilde{\rho}_2 \dots$. Let us show by induction that for all n , $\rho_n \sim_q \tilde{\rho}_n$. It means that $h\rho \sim_q \tilde{h}\tilde{\rho}$ and since \sim_q respects the gain functions,

$$\text{Gain}_i(h\langle\sigma'_i, \sigma_{-i|h}\rangle_v) = \text{Gain}_i(h\rho) = \tilde{\text{Gain}}_i(\tilde{h}\tilde{\rho}) = \tilde{\text{Gain}}_i(\tilde{h}\langle\tau'_i, \tau_{-i|\tilde{h}}\rangle_{\tilde{v}}). \quad (\text{B.29})$$

For $n = 0$: $\rho_0 = v \sim_q \tilde{v} = \tilde{\rho}_0$. Assume that this property is true for all $n \leq k$ and let us prove it remains true for $n = k + 1$.

- **First case:** if $\rho_k \in V_i$, by IH $\rho_k \sim_q \tilde{\rho}_k$ and thus $\tilde{\rho}_k \in \tilde{V}_i$. It follows that:

$$\begin{aligned} \rho_{k+1} &= \sigma'_i(\rho_0 \dots \rho_k) \\ &\sim_q \tau'_i(\tilde{\rho}_0 \dots \tilde{\rho}_k) && \text{(By construction of } \sigma_i \\ & && \text{and } \rho_0 \dots \rho_k \sim_q \tilde{\rho}_0 \dots \tilde{\rho}_k \text{ (IH))} \\ &= \tilde{\rho}_{k+1} \end{aligned}$$

- **Second case:** if $\rho_k \in V_j$ with $(j \neq i)$ then as previously $\tilde{\rho}_k \in \tilde{V}_j$

and we have:

$$\begin{aligned}\rho_{k+1} &= \sigma_{j|h}(\rho_0 \dots \rho_k) \\ &\sim_q \tau_{j|\tilde{h}}(\tilde{\rho}_0 \dots \tilde{\rho}_k) \quad (\rho_0 \dots \rho_k \sim_q \tilde{\rho}_0 \dots \tilde{\rho}_k \text{ (HI) and Claim 1.}) \\ &= \tilde{\rho}_{k+1}\end{aligned}$$

There is still to prove that

$$\tilde{\text{Gain}}_i(\tilde{h}\langle\tau_{|\tilde{h}}\rangle_{\tilde{v}}) = \text{Gain}_i(h\langle\sigma_{|h}\rangle_v). \quad (\text{B.30})$$

By Claim 1, we have that $\langle\sigma_{|h}\rangle_v \sim_q \langle\tau_{|\tilde{h}}\rangle_{\tilde{v}}$ thus $h\langle\sigma_{|h}\rangle_v \sim_q \tilde{h}\langle\tau_{|\tilde{h}}\rangle_{\tilde{v}}$. The fact that \sim_q respects the gain functions concludes the reasoning.

By (B.28), (B.29) and (B.30), we conclude that σ'_i is a profitable deviation in $(\mathcal{G}_{|h}, v)$. □

Proposition B.3.6. *Let $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \pi})$ be a game and $(\tilde{\mathcal{G}}, [v_0]) = (\tilde{A}, (\tilde{\text{Gain}}_i)_{i \in \Pi})$ its associated quotient game where \sim is a bisimulation equivalence on (\mathcal{G}, v_0) . If \sim respects the partition and the gain functions, we have that: if there exists an SPE τ in $(\tilde{\mathcal{G}}, [v_0])$ such that $\tilde{\text{Gain}}(\langle\tau\rangle_{[v_0]}) = p$ for some $p \in \{0, 1\}^{|\Pi|}$ then there exists an SPE σ in (\mathcal{G}, v_0) such that $\text{Gain}(\langle\sigma\rangle_{v_0}) = p$.*

Proof. Let $(\mathcal{G}, v_0) = (A, (\text{Gain}_i)_{i \in \pi})$ be a game and $(\tilde{\mathcal{G}}, [v_0]) = (\tilde{A}, (\tilde{\text{Gain}}_i)_{i \in \Pi})$ its associated quotient game where \sim is a bisimulation equivalence on (\mathcal{G}, v_0) which respects the partition and the gain functions. We assume that there exists an SPE τ in $(\tilde{\mathcal{G}}, [v_0])$ such that $\tilde{\text{Gain}}(\langle\tau\rangle_{[v_0]}) = p$ for some $p \in \{0, 1\}$.

Let $h \in \text{Hist}(v_0)$ such that $\text{Last}(h) \in V_i$ for some $i \in \Pi$. Thanks to bisimulation \sim_q , there exists a unique $\tilde{h} \in \tilde{\text{Hist}}_i([v_0])$ such that $h \sim_q \tilde{h}$ (\star). We have that $\tau_i(\tilde{h}) = \tilde{v}$ for some $\tilde{v} \in \tilde{V}$, thus by \sim_q there exists $v \in V$ such that

$hv \sim_q \tilde{h}\tilde{v}$. We define $\sigma_i(h) = v$.

CLAIM 2:

1. $\forall h, h' \in \text{Hist}(v_0)$ such that $h \sim h'$: $\sigma_i(h) \sim \sigma_i(h')$ (if $\text{Last}(h) \in V_i$).
2. $\forall h \in \text{Hist}(v_0), \forall \tilde{h} \in \widetilde{\text{Hist}}([v_0])$ such that $h \sim_q \tilde{h}$: $\sigma_i(h) \sim_q \tau_i(\tilde{h})$ (if $\text{Last}(h) \in V_i$).

Proof:

1. By (\star) , we have that for all $h \sim h' \in \text{Hist}_i(v_0)$ there exists a unique \tilde{h} such that $h \sim_q \tilde{h}$ and $h' \sim_q \tilde{h}$. It follows by construction of σ that $\sigma_i(h) \sim_q \tau_i(\tilde{h})$ and $\sigma_i(h') \sim_q \tau_i(\tilde{h})$ and thus $\sigma_i(h) \sim \sigma_i(h')$. It means that σ is uniform.
2. let $h \in \text{Hist}(v_0)$ and $\tilde{h} \in \widetilde{\text{Hist}}([v_0])$ be two histories such that $\text{Last}(h) \in V_i$ iff $\text{Last}(\tilde{h}) \in \tilde{V}_i$ for some $i \in \Pi$ and such that $h \sim_q \tilde{h}$. By construction of σ , there exists $\tilde{g} \in \widetilde{\text{Hist}}([v_0])$ such that $h \sim_q \tilde{g}$ and $\sigma_i(h) \sim_q \tau_i(\tilde{g})$. But by \sim_q if $h \sim_q \tilde{g}$ and $h \sim_q \tilde{h}$, then $\tilde{g} = \tilde{h}$. It concludes the proof. □

By (2) in Claim 2, we have that $\langle \sigma \rangle_{v_0} \sim_q \langle \tau \rangle_{[v_0]}$. It follows, due to the fact that \sim respects the gain functions, that $\text{Gain}(\langle \sigma \rangle_{v_0}) = \tilde{\text{Gain}}(\langle \tau \rangle_{[v_0]}) = p$.

Now, we prove that σ is an SPE. Ad absurdum, we assume that there exists $hv \in \text{Hist}(v_0)$, there exists $i \in \Pi$ and there exists σ'_i a profitable deviation of $\sigma_{i|h}$ for Player i in $(\mathcal{G}_{|h}, v)$, i.e.,

$$\text{Gain}_i(h\langle \sigma_{|h} \rangle_v) < \text{Gain}_i(h\langle \sigma'_i, \sigma_{-i|h} \rangle_v) \quad (\text{B.31})$$

Let $\tilde{h}\tilde{v} = [h_0][h_1] \dots [v]$ with $[h_0] = [v_0]$ we have that $hv \sim_q \tilde{h}\tilde{v}$. By (2) in Claim 2, we have that $h\langle \sigma_{|h} \rangle_v \sim_q \tilde{h}\langle \tau_{|\tilde{h}} \rangle_{\tilde{v}}$, since \sim respects the gain functions,

it follows:

$$\text{Gain}_i(h\langle\sigma_{|h}\rangle_v) = \widetilde{\text{Gain}}_i(\tilde{h}\langle\tau_{|\tilde{h}}\rangle_{\tilde{v}}) \quad (\text{B.32})$$

To obtain the contradiction, we build τ'_i a profitable deviation of $\tau_{i|\tilde{h}}$ for Player i in $(\tilde{\mathcal{G}}_{|\tilde{h}}, \tilde{v})$.

Let $\rho = \langle\sigma'_i, \sigma_{i|h}\rangle_v$, let $\tilde{p} \in \widetilde{\text{Hist}}_i(\tilde{v})$, we define $\tau_i(\tilde{p})$ as follows:

$$\tau'_i(\tilde{p}) = \begin{cases} [\rho_{n+1}] & \text{if } \tilde{p} < [\rho_0][\rho_1] \dots \text{ and } \text{Last}(\tilde{p}) = [\rho_n] \\ \text{some } r \in \text{Succ}(\text{Last}(\tilde{p})) & \text{otherwise} \end{cases}$$

Let $\tilde{\rho} = \langle\tau'_i, \tau_{-i|\tilde{h}}\rangle_{\tilde{v}}$ and let us prove that $\rho \sim_q \tilde{\rho}$, *i.e.*, $\forall n \in \mathbb{N} \rho_n \sim_q \tilde{\rho}_n$. We proceed by induction on n .

For $n = 0$, $\rho_0 = v \sim_q [v] = \tilde{v} = \tilde{\rho}_0$. Let us assume that this property is true for all $n \leq k$ and let us prove it remains true for $n = k + 1$.

- **First case:** If $\rho_k \in V_i$, then since $\rho_k \sim_q \tilde{\rho}_k$ by IH, $\tilde{\rho}_k \in \tilde{V}_i$. It follows that:

$$\begin{aligned} \rho_{k+1} &\sim_q [\rho_{k+1}] && (\text{By definition of } \sim_q) \\ &= \tau'_i(\tilde{\rho}_0 \dots \tilde{\rho}_k) && (\text{By IH, } \rho_0 \dots \rho_k \sim_q \tilde{\rho}_0 \dots \tilde{\rho}_k = [\rho_0] \dots [\rho_k]) \\ &= \tilde{\rho}_{k+1} \end{aligned}$$

- **Second case:** If $\rho_k \in V_j$ (with $j \neq i$), then as previously $\tilde{\rho}_k \in \tilde{V}_j$. It follows that:

$$\begin{aligned} \rho_{k+1} &= \sigma_{j|h}(\rho_0 \dots \rho_k) = \sigma_j(h\rho_0 \dots \rho_k) \\ &\sim_q \tau_j(\tilde{h}\tilde{\rho}_0 \dots \tilde{\rho}_k) && (\text{By IH, } h\rho_0 \dots \rho_k \sim_q \tilde{h}\tilde{\rho}_0 \dots \tilde{\rho}_k \\ & && \text{and by (2) in Claim 2)} \\ &= \tau_{j|\tilde{h}}(\tilde{\rho}_0 \dots \tilde{\rho}_k) = \tilde{\rho}_{k+1} \end{aligned}$$

Thus, $\rho \sim_q \tilde{\rho}$ and so $h\rho \sim_q \tilde{h}\tilde{\rho}$. Since, \sim respects the gain functions, we can conclude that:

$$\text{Gain}_i(h\langle\sigma'_i, \sigma_{-i|h}\rangle_v) = \widetilde{\text{Gain}}_i(\tilde{h}\langle\tau'_i, \tau_{-i|\tilde{h}}\rangle_{\tilde{v}}). \quad (\text{B.33})$$

By (B.31), (B.32) and (B.33), we can state that:

$$\tilde{\text{Gain}}_i(\tilde{h}\langle\tau'_i, \tau_{i|\tilde{h}}\rangle\tilde{v}) = \text{Gain}_i(h\langle\sigma'_i, \sigma_{-i|h}\rangle v) > \text{Gain}_i(h\langle\sigma|h\rangle v) = \tilde{\text{Gain}}_i(\tilde{h}\langle\tau_{|\tilde{h}}\rangle\tilde{v})$$

□