NNT : 2024IPPAX116

Thèse de doctorat

Tropical Polynomial Systems and Game Theory

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'École polytechnique

École doctorale n°574 École doctorale de mathématiques Hadamard (EDMH)
Spécialité de doctorat : Mathématiques Appliquées

Thèse présentée et soutenue à Palaiseau, le 27 novembre 2024, par

ANTOINE BÉREAU

Composition du Jury :

Fabrice Rouillier Directeur de recherche, Sorbonne Université et Inria Paris	Président
Frédéric Bihan Maître de conférences, Université Savoie Mont Blanc	Rapporteur
Josephine Yu Professor, Georgia Institute of Technology	Rapporteuse
Alicia Dickenstein Profesora regular plenaria, Universidad de Buenos Aires	Examinatrice
Dmitry Grigoriev Directeur de recherche, Université de Lille	Examinateur
Pierre Lairez Chargé de recherche, Inria Saclay	Examinateur
Stéphane Gaubert Directeur de recherche, Inria Saclay et Ecole polytechnique (CMAP)	Directeur de thèse
Marianne Akian Directeur de recherche, Inria Saclay et Ecole polytechnique (CMAP)	Directrice de thèse

Tropical Polynomial Systems and Game Theory

Abstract

The main objects under study in this manuscript are systems of tropical polynomial equations and inequations. Given such a system, our aim is to be able to efficiently decide its solvability, with a Nullstellensatz-type result adapted into the tropical setting. A second, more elaborate question, consists in the effective computation of the solution set of such a system.

In 2018, Grigoriev and Podolskii established a tropical analogue of the effective Nullstellensatz, showing that the solvability of a system of tropical polynomial equations is equivalent to the solvability of a linearized system, obtained by truncating the associated tropical *Macaulay matrix* up to some degree bound, and provided an upper bound of the optimal truncation degree as a function the number of variables and the number of the tropical polynomials in the system, and their degrees. We establish an improved *tropical Nullstellensatz*, taking into consideration the possible *sparsity* of the tropical polynomials in a system. We rely on a polyhedral construction of Canny and Emiris from 1993, refined one year later by Sturmfels. On top of accounting for sparsity, our result closes the gap between the truncation degree obtained by Grigoriev and Podolskii and the classical *Macaulay degree bound*. Furthermore, we establish a more general *tropical Positivstellensatz* based on the very same construction, at the cost of an inflation of the truncation degree. This tropical Positivstellensatz allows one to decide the inclusion of tropical basic semialgebraic sets, thus reducing decision problems for tropical semialgebraic sets to the solution of systems of tropical linear equalities and inequalities. We combine these two results in a global *hybrid Positivstellensatz*.

Such tropical linear systems are known to be reducible to *mean payoff games*, which can be solved in practice, in a scalable way, by *value iteration* or *policy iteration* methods. In particular, we propose a speedup of the classical value iteration algorithm of Zwick and Paterson, which we then use in order to decide the solvability of a system of tropical polynomial equalities and inequalities. This speedup relies on two ingredients: the use of the *Krasnoselskii-Mann damping* in the iteration process, as well as the introduction of a *widening step*, allowing for a quicker exit in case of infeasibility. This value iteration algorithm with widening was implemented in Python.

We then develop a tropical analogue of *eigenvalue methods* in order to effectively compute the solution set of tropical polynomial systems. Relying on the connection between tropical linear systems and mean payoff games, we show that this solution set can be obtained by solving parametric mean-payoff games, arising from approriate linearizations of the tropical polynomial system using tropical Macaulay matrices. We present two approaches: a first one based on a dichotomic search, which simply allows one to certify the solvability of a tropical polynomial system, and a second, more elaborate approach, based on a *tropical homotopy* technique, allowing one to compute projections of the solution set onto any coordinate.

Finally, we present a generalization of the *Ishikawa fixed-point convergence theorem*, extending it so as to tackle the case of polyhedral fixed-point free maps. This provides a theoretical framework motivating the use of the Krasnoselskii-Mann damping in the construction of our accelerated value iteration-type algorithm.

Résumé en français

Les systèmes d'équations et inéquations polynomiales tropicales constituent le principal objet d'étude de cette thèse. Le but des travaux présentés ici est de décider efficacement la résolubilité de tels systèmes, en énonçant un résultat sous la forme d'un Nullstellensatz adapté au contexte tropical. Un second problème, plus élaboré, vient alors : celui de déterminer effectivement l'ensemble des solutions d'un système polynomial tropical.

En 2018, Grigoriev et Podolskii ont établi un analogue tropical du Nullstellensatz effectif, montrant ainsi que la résolubilité d'un système d'équations tropicales polynomiales était équivalente à la résolubilité d'un système linéarisé, obtenu en tronquant à un certain degré la *matrice de Macaulay* tropicale associée au système. Leur résultat donnait de plus une borne supérieure du degré de troncature optimal, en fonction du nombre de variables, du nombre de polynômes dans le système, ainsi que de leurs degrés. Nous établissons une version améliorée du *Nullstellensatz tropical*, prenant en compte la possible structure creuse des polynômes tropicaux du système. Notre résultat repose sur une construction polyédrale initialement due à Canny et Emiris en 1993, et qui fut raffinée un an plus tard par Sturmfels. En plus d'être adapté aux systèmes creux, notre résultat permet de combler l'écart entre le degré de troncature de Grigoriev et Podolskii, et la *borne de Macaulay* dans le cas classique. En outre, nous établissons grâce à la même construction un *Positivstellensatz tropical*, au prix d'une dilatation du degré de troncature de la matrice de Macaulay. Ce Positivstellensatz permet de résoudre les problèmes d'inclusion d'ensembles semi-algébriques tropicaux basiques, réduisant ainsi les problèmes de décision pour les ensembles semi-algébriques tropicaux à la résolution d'un système d'égalités et inégalités linéaires tropicales. Nous combinons ultimement ces deux résultats dans un *Positivstellensatz hybride*.

Il est connu que la résolution des systèmes linéaires tropicaux mentionnés ci-dessus se réduit à celle des jeux avec paiement moyen, qui peuvent en pratique être résolus sur de grandes instances par des méthodes d'*itération sur les valeurs* ou bien d'*itération sur les politiques*. En particulier, nous proposons une accélération de l'algorithme classique d'itération sur les valeurs de Zwick et Paterson, que l'on emploie ensuite pour déterminer la résolubilité d'un système d'équations et inéquations polynomiales tropicales. Cette accélération repose sur deux ingrédients : le recours à l'*amortissement de Krasnoselskii-Mann* dans le processus d'itération, ainsi que l'introduction d'une *étape d'élargissement*, qui fournit une condition de sortie rapide en cas d'infaisabilité. Cette itération sur les valeurs avec élargissement a été implémentée en Python.

Nous développons ensuite un analogue tropical des *méthodes de valeurs propres* afin de calculer de manière effective l'ensemble des solutions d'un système polynomial tropical. En nous reposant sur la correspondance entre systèmes linéaires tropicaux et *jeux avec paiement moyen*, nous montrons que cet ensemble de solutions peut être déterminé en résolvant des jeux paramétriques, provenant de linéarisations adéquates du système polynomial initial avec la matrice de Macaulay tropicale. Nous présentons deux approches : une première basée sur la recherche dichotomique, permettant de simplement certifier la résolubilité d'un système polynomial tropical, et une seconde, plus élaborée, basée sur le *suivi de chemin homotopique*, permettant de calculer des projections sur chaque coordonnée de l'ensemble des solutions.

Enfin, nous présentons une généralisation du *théorème de convergence d'Ishikawa* sur l'itération de Krasnoselskii-Mann en l'étendant au cas d'applications polyédrales sans point fixe. Cette généralisation fournit le contexte théorique nécessaire motivant le recours à l'amortissement de Krasnoselskii-Mann dans notre accélération de l'itération sur les valeurs.

Remerciements

Mes tous premiers remerciements s'adressent d'abord à Stéphane Gaubert et Marianne Akian, qui ont tous les deux dirigé cette thèse. Ce fut un plaisir de travailler ensemble pendant un petit peu plus de trois années, de mon stage de recherche de master jusqu'à la fin de cette thèse. L'enthousiasme et les nombreuses idées de Stéphane, combinés à l'exigence et à la rigeur de Marianne ont permis au présent travail d'atteindre sa forme la plus aboutie. Les nombreuses heures de travail et de discussions ensemble m'ont permis d'apprécier pleinement la beauté des mathématiques tropicales, ainsi que d'accéder à un niveau de compréhension des mathématiques — pas uniquement tropicales — bien au-delà de tout ce que j'aurais pu espérer en commençant cette thèse. Pour cela, merci encore une fois.

J'aimerais également remercier Frédéric Bihan et Joséphine Yu pour avoir accepté avec enthousiasme de rapporter cette thèse, ainsi qu'Alicia Dickenstein, Dima Grigoriev, Pierre Lairez et Fabrice Rouillier pour avoir accepté de prendre part à mon jury de thèse.

Je dois également remercier l'ensemble du CMAP et de l'Inria, et en particulier de l'équipe TROPICAL pour m'avoir fourni un environnement de travail extrêmement fertile. Un remerciement particulier à Matìas Bender ainsi qu'à Xavier Allamigeon avec qui j'ai pu avoir plusieurs discussions mathématiques très éclairantes. Je dois également remercier l'ensemble des équipes administratives du CMAP, en particulier Hanadi Dib, Alexandra Liot, Nassera Naar, Alexandra Noiret et Amandine Sainsard, pour leur aide précieuse qui a permis de mener à bien les différentes missions que j'ai pu effectuer au cours de ma thèse, ainsi qu'à Nathalie Rodrigues pour son aide dans l'organisation de la soutenance. Enfin, je remercie également pour leur bonne humeur l'ensemble des doctorants du CMAP, et en particulier mes co-chargés de TD ces différentes années : Alexandre, Diego, Doan, Dorian, Francesco, Nicolas et Yuanyang.

J'ai eu la chance d'évoluer dans des communautés bienveillantes, et je suis très reconnaissant pour toutes les rencontres mathématiques que j'ai pu faire durant ma thèse, à travers les pays voire les continents, d'une conférence à l'autre, qu'il s'agisse d'autres doctorants ou jeunes chercheurs, comme Christina, Denae, Dom et Hiroshi, ou bien de chercheurs plus experimentés, qui m'ont permis d'en apprendre toujours plus ne serait-ce que par une remarque ou une question pertinente lors d'un échange ou bien à la suite d'un exposé. Je remercie particulièrement Yue Ren pour les pistes enrichissantes qu'il a pu me suggérer à plusieurs reprises, ainsi qu'Elias Tsigaridas pour son intérêt et pour nos différents échanges sur ce projet de postdoc qui n'a malheureusement pas abouti.

Je dois enfin remercier tous mes proches, à commencer par toute ma famille, dont le soutien sans faille m'a permis d'en arriver jusqu'ici aujourd'hui. Je dois également remercier tous mes amis, qui ont vécu par procuration pour certains, et par eux-même également pour d'autres l'expérience de la thèse. Merci à tous les *STEM Gays*, Antoine, les deux Baptiste et Hugo, pour leurs conseils toujours avisés, sur tous les domaines adjacents à la thèse, des sessions code à l'écriture de lettres de motivation où à la direction artistique de mes présentations, ainsi que pour tout le reste. Merci à Thelma pour son baptême du feu en étant la première à soutenir sa thèse. Merci à tous mes proches du lycée et d'avant qui ont suivi mon parcours de près ou de loin depuis maintenant des années : Antoine, Camille, Dina, Estelle, Forgaill, Juliette, Marianne et Sophie. Merci à Corentin dont l'expérience a été précieuse dans la transition entre le monde de la recherche et celui de l'enseignement. Merci à Mikayil pour tous les moments de détente et les balades noctures quand la tension devenait trop intense. Merci à Moana d'avoir été le principal public des blagues mathématiques qui étaient trop obscures pour que je les envoie à qui que ce soit d'autre. Merci à Armand, Emmanuel, Nathan ainsi qu'à toutes celles et ceux qui ont pris la peine de venir m'encourager à ma soutenance de thèse.

Enfin, je ne pourrais jamais assez remercier Charlie, qui a été pendant ces trois années de thèse un soutien inébranlable, lors des périodes intenses, quand il y avait des articles à corriger à pas d'heure pour les soumettre à temps, des dossiers urgents à compléter, des documents à imprimer en urgence, mais également lors des meilleurs moments. Ton amour inconditionnel constitue une source d'énergie inépuisable qui m'a permis de traverser toutes les épreuves jusqu'ici et qui me permettra de traverser toutes celles qui se présenteront à l'avenir.

Contents

Introduction

1 Context and motivation of this work

The *tropical semifield* $\mathbb{T}$ refers to the semifield $\mathbb{R} \cup \{-\infty\}$, endowed with addition $\oplus = \max$ and multiplication $\odot = +$, and with zero element $\mathbb{0} = -\infty$ and unit $\mathbb{1} = 0$. It satisfies all the usual properties of a field, except for the existence of an opposite element for the tropical addition $\oplus$. A *tropical polynomial function* can be defined as the maximum of a finite number of affine functions with integer slopes. It is thus a convex and piecewise affine function. The *tropical hypersurface* associated to a tropical polynomial function is then defined as its non-differentiability locus. In other words, it consists in the points such that the maximum is simultaneously achieved by at least two distinct affine functions. These points are considered as the 'zeros' of these functions, in the tropical sense. Tropical objects naturally arise in valued field theory. Indeed, considering non-archimedian fields such as the field of Puiseux series, tropical hypersurfaces coincide with the image of algebraic hypersurfaces by the valuation map, as per the Kapranov theorem [EKL06]. More generally, one studies *non-archimedean amoebas*, which are images by a non-archimedean valuation of an algebraic set. A super-approximation of a non-archimedean amoeba can be obtained by translating the defining equations of the algebraic set in the tropical semifield, leading to a system of tropical polynomials equations. The solution set of such a tropical polynomial system is called a *tropical prevariety*. Moreover, this approximation becomes exact under appropriate genericity conditions when working over an algebraically closed non-archimedean field, in which case the non-archimedean amoeba is then refered to as a *tropical variety*, as per the 'fundamental theorem of tropical geometry' as stated in [MS15]. Similarly, the solution sets of systems of tropical polynomial inequalities provide upper approximations of images by a convex non-archimedean valuation of semi-algebraic sets over a real closed field, and these approximations are also exact under some genericity conditions [IV96, AGS20, JSY20]. Furthermore, the combinatorial nature of tropical polynomial systems makes them often easier to study than their classical analogues. All the previous ideas, which can be traced back to works of Viro [Vir89], Gelfand, Kapranov, Zelevinsky [GKZ94], Sturmfels [HS95], or Mikhalkin [Mik05], are at the very heart of tropical geometry. For even more background, the reader is advised to refer to [IMS09, MS15].

Tropical polynomial systems also arise in various specific applications, among which the computation of equilibrium points of the n-body problem [HM06], or, independently of the previous non-archimedean interpretation, in the computation of stationary behaviors of discrete event systems, such as Petri nets [CGQ98], see also [ABG15] for an application to performance evaluation of emergency call centers. Other motivations include auction theory [BK19], chemical reaction networks [DSR22], or max-out networks [MRZ22], just to list a few of them. Finally, tropical polynomial functions also provide a suitable language to describe some optimisation problems. In particular, the solvability of a tropical linear system is proven to be equivalent to the computation of the value of a certain type of zero-sum two-player games called *mean payoff games* [AGG12].

Given a collection of tropical polynomial functions, a fundamental question in the field of computational tropical geometry, consists in investigating the vacuity of the intersection of their associated tropical hypersurfaces. In other words, one wishes to determine whether there exists tropical 'zeros' common to all the tropical polynomial functions of the collection considered. Such questions have already been thouroughly explored — and they keep being explored — in the case of classical polynomial systems, where numerous tools have been developed to study the solvability of a classical polynomial system on the one hand, and on the other hand to effectively give its solutions. Notably, one can mention the computation of Gröbner bases in the context of symbolic computation (see *e.g.* [CLO15]), and homotopy continuation methods in the context of numerical algebraic geometry (see *e.g.* [AG90]).

A core result of the theory of classical polynomial system solving theory is that proving the existence of a solution to a polynomial system can be reduced to finding nontrivial elements in the kernel of some submatrices of the so-called *Macaulay matrix* (see *e.g.* [EM07]), whose entries are coefficients of the polynomials of the system. This process constitutes a *linearization* of the polynomial system. The size of the submatrices whose kernel constitute a certificate for the existence or nonexistence of a solution to a system of polynomial equations depends on a quantity N, called the *truncation degree*. The minimal value of the truncation degree such that the associated submatrix of the Macaulay matrix generically constitutes such a certificate is called the *Macaulay bound* (see [Laz81, Laz83, Giu84], and also [BCC$^+$05, CLO15] for background). The sparse case is also handled with a construction developed by Canny and Emiris in [CE93] and [Emi05], allowing one to construct minors of the Macaulay matrix from which one can obtain a Sylvester-type formula for the sparse resultant.

Besides tropical objects, one of the main focus of the present manuscript is a specific class of zero-sum two-player games consisting in the so-called *mean payoff games*. The broadest class of (stochastic, concurrent) mean payoff games was introduced by Gillette in [Gil57]. In the scope of this work, we shall only rely on *deterministic* mean payoff games. The question of finding optimal positional strategies for both players, as well as the complexity of computing the vector of values of a mean payoff game, has then been thoughrouly studied in [LL69, EM79, GKK90, ZP96] among others. (Deterministic) Mean payoff games are among the undecided problems in computational complexity — they belong to the complexity class $\mathsf{NP} \cap \mathsf{coNP}$ [KL93, ZP96], their membership to the complexity class P is a long standing open question.

The connection between mean payoff games and tropical geometry is the following: being able to solve a system of tropical linear weak inequalities has been proven to be equivalent to being able to solve a mean payoff game in [AGG12]. In particular, deciding the non-emptyness of a tropical linear prevariety reduces to a mean payoff game. Grigoriev and Podolskii refined this result, showing that deciding the non-emptyness of tropical linear prevarieties is actually polytime Turing-equivalent to mean-payoff games [GP13]. Earlier results relating tropical linear inequalities and mean-payoff games can be found in [CTGG99, MSS04, DG06, Kat07]. Each mean payoff game can be associated to a min-max operator, named the *Shapley operator*. More specifically, there exists an extensive theory of *nonlinear eigenvalues* for these game operators (see [BCOQ92, But10] for general references, see also [Koh80, GG98, GG04] on min-max operators), on which the results of [AGG12] rely.

Grigoriev and Podolskii established in [GP18] a *tropical Nullstellensatz*, stating that a system of tropical polynomial equations is solvable if and only if its linearization, obtained by truncating the Macaulay matrix up to an appropriate degree bound, is solvable. Their results also apply to polynomial inequations. Since systems of tropical linear equalities and weak inequalities reduce to mean payoff games [AGG12], and so do systems containing strict inequalities [AFG$^+$14], this provides both theoretical tools (strong duality theorems) and algorithms to effectively investigate the solvability of tropical polynomial systems. The proof of [GP18], confirming a conjecture made by Grigoriev [Gri12], is based on very ingenious geometric arguments. However, Grigoriev and Podolskii observed that their proof leads to an estimate of the truncation degree which may not be optimal, as it notably does not match the Macaulay bound for generic classical polynomial systems.

2 Contributions and organisation of the manuscript

The present manuscript is divided in five chapters whose order reflects the chronological progression of the work that was conducted during the three years of my PhD, whose starting point stemed from [GP18]. Whereas the first part of the thesis deals with the question of the solvability (deciding feasibility) of tropical polynomial systems by formulating tropical Null- and Positivstellensätze, the second part of the work focuses on the effective resolution of tropical polynomial systems, *i.e.* on the computation of solutions, relying on tools from game theory, and in particular on mean payoff games. The nonlinear eigentheory of Shapley operators of mean payoff games resides at the core of the proposed effective resolution of tropical polynomial systems, via the aforementioned tropical Positivstellensatz. Our resolution algorithms rely on classical *value iteration* and *policy iteration* methods, which are used in order to solve mean payoff games. On top of that, while exploring the theory of nonlinear eigenvalues of Shapley operators, and refining value iteration, we obtained a result of independent interest, an extension of the Ishikawa theorem [Ish76], dealing with the convergence of the Krasnoselskii-Mann iteration for non-expansive mappings, to the case of polyhedral fixed-point free mappings.

Below is a detailled summary of the content of each chapter of the present manuscript, highlighting in particular the main contributions it contains.

- ➢ **Chapter 1** is a detailed summary of all the theoretical framework of the present work, where the different notions used throughout this manuscript are introduced, and their main properties are recalled and commented. These preliminaries begin with very general and useful concepts from polyhedral geometry, which are necessary in order to accurately describe tropical objects. Then, we briefly recall some generalities of classical algebraic and semialgebraic sets, in order to later on define their tropical counterparts. We then present the main objects and ideas from tropical geometry: tropical polynomials, tropical prevarieties, tropical semialgebraic sets. We also recall some results from classical elimination theory, to serve as a point of reference later on in the manuscript when dealing with the tropical Nullstellensatz. Finally, we thoroughly summarize the language and objects from mean payoff games and nonlinear eigenvalue theory, upon which we shall rely when tackling the effective resolution of tropical polynomial systems.

➢ **Chapter 2** details the first contribution of this thesis. A fundamental problem in tropical geometry consists in deciding the solvability of a tropical polynomial system. By adaptating classical tools from elimination theory to the tropical setting, we state an improved tropical Nullstellensatz (Theorem 2.1.5), suited to *sparse* polynomial systems. Our result relies on the notion of *Canny-Emiris sets* [CE93, Emi05], involved in the construction of a submatrix of the Macaulay matrix. The columns of said submatrix are selected by considering the integer points of a generic perturbation of the Minkowski sum of the Newton polytopes of the polynomials for the system, and the selection of the row rely on an appropriate notion of 'row content'. This construction was generalized by Sturmfels in [Stu94, §3], by considering a generic collection of polyhedral concave functions. Here, we apply this tool in the present setting, and show that any Canny-Emiris set leads to a tropically valid linearization. More precisely, if a tropical system is unfeasible, then, the 'row contents' arising in the Canny-Emiris construction yield a minor of the Macaulay matrix which is 'nonsingular' in a tropical sense, therefore serving as an unsolvability certificate for our linearized tropical system. On top of the efficient handling of sparsity, our Nullstellensatz additionally provides an improved degree bound in the case of full polynomials, matching the classical *Macaulay bound* for systems of $n + 1$ equations in n unknowns closing the gap left in [GP18]. Additionally, at the cost of a slightly inflated degree bound, our construction leads to a *tropical Positivstellensatz* (Theorem 2.2.1), relying on additional ingredients, notably the *Shapley-Folkman lemma*. This tropical Positivstellensatz allows one to tackle the more delicate case of *two-sided* polynomial systems, combining a mixture of polynomial equalities, as well as weak and strict inequalities, and in particular to check the inclusion between two basic tropical semialgebraic subsets of $\mathbb{R}^n$. The provided inflated degree bound is however not known to be optimal. These two results are ultimately combined in the most general form in a *hybrid tropical Positivstellensatz* (Theorem 2.2.13). The contents of this chapter have been initially presented in the conference paper [ABG23a], and developed in [ABG23b]. For the sake of simplicity, the present tropical Null- and Positivstellensätze are stated and proven over the usual tropical semifield $\mathbb{R} \cup \{-\infty\}$. However, a tropical semifield can be constructed over any divisible (totally) ordered group, and this broader setting is suitable for the study 'higher rank' tropicalizations [Aro10a, Aro10b, AGT16, AI22, JS23]. The completeness of the first order theory of divisible ordered groups [Rob56, §4.3], fortunately allows all results provided to carry over to this case.

➢ **Chapter 3** presents our first algorithmic contribution, which consists in a speedup of the classical *value iteration* algorithm, which is one of the two main class of algorithms used in order to solve mean payoff games [ZP96, AGG12]. The two main ingredients of this speedup are the use of the Krasnoselskii-Mann damping, as well as the introduction of a 'widening step', leading to our accelerated *value iteration with widening* algorithm (Algorithm 1). We provide a proof of the termination and correction of this algorithm, as well as a theoretical complexity bound (Theorem 3.2.1 and subsequent corollaries), as well as a Python implementation in [Bé23]. This algorithm and its implementation have been presented at ISSAC 2023, and have also been the subject of [ABG23a].

➢ **Chapter 4** deals with the effective resolution of tropical polynomial systems. Indeed, tropical linear systems are known to reduce to mean payoff games, and therefore, as consequence of the linearization results from Chapter 2, this means that one can use classical mean payoff games algorithms in order to effectively implement an oracle deciding the solvability of a tropical polynomial system. One of the main advantages of this approach lies in the *scalability* of mean payoff games algorithms. Indeed, although the existence of a polynomial time algorithm to solve mean payoff games is a long standing open question, large sparse instances of mean payoff games can typically be solved efficiently in practice by *value iteration* type algorithms [ZP96, AGG12], or by *policy iteration* algorithms, see in particular [CTGG99, DG06, Cha09]. In this chapter, we lay the foundations to a general homotopy method to solve parametric mean payoff games, and then explain how this method can be used in order to solve effectively tropical polynomial systems. We present two methods. The first method is based on a *dichotomic search* (Algorithm 3), and provides a certificate for the feasibility of a tropical polynomial system. The second method is more elaborate and is based on a *tropical homotopy* technique, where one of the coordinates of the solution is passed as a parameter, and then the solvability of the system is examined in function of the value of the parameter, thus allowing one to compute projections of the solution set onto any coordinate. In order to introduce this method, we develop a general *singular homotopy method for parametric mean payoff games* (Algorithms 2a and 2b), whose interest goes beyond the question of tropical polynomial solving. This method relies on the theory of *nonlinear eigenvalues* of mean payoff game operators. The ideas of this chapter have partially been exposed at ICMS 2024, leading to [ABG24].

➢ **Chapter 5** is mainly independent from the previous chapters and presents a very simple convergence result on the Krasnoselskii-Mann iterates of polyhedral maps without fixed points (Theorem 5.1.2). This theorem constitutes a conveniant generalization of the well-known Ishikawa fixed-point convergence theorem ([Ish76, Corollary 2]), and allows one to handle the fixed-point free case. Then, instead of converging to a fixed point, the iteration asymptotically 'converges' to an invariant half-line. This new finding has not yet been published.

3 Related work

As mentioned above, the Nullstellensatz stated by Grigoriev and Podolskii [GP18] constituted the starting point of the present work. The main contribution here is the handling of sparsity, with a new proof, based on classical elimination theory tools, leading to an improved degree bound in the dense case, and moreover handling two-sided equalities and inequalities, allowing us to interpret these results in terms of a tropical Positivstellensatz. Grigoriev and Podolskii also considered solutions $x \in (\mathbb{R} \cup \{-\infty\})^n$, instead of $x \in \mathbb{R}^n$, showing that this leads to a blow up of the truncation degree bound by an exponential factor. In this manuscript, we limit our attention to the *toric* case only, looking only for solutions over the tropical torus $\mathbb{R}^n$. Indeed, whereas the tropical Nullstellensatz in the non-toric case is of intrinsic theoretical interest, from an algorithmic perspective, it is cheaper to reduce to the toric case by exhaustively looking for all the possible supports for a given solution. This leads to solving games with a number of states simply exponential in the input state, whereas the exponential character of the truncation degree in the non-toric case would generally lead to solving games with a doubly exponential number of states.

The standard approach to the computation of tropical prevarieties exploits the duality between an arrangement of tropical hypersurfaces, and the *dual subdivision* of the Minkowski sum of the Newton polytopes of the polynomials defining each tropical hypersurface. This dual subdivision is a mixed regular polyhedral subdivision that is obtained lifting each monomial in the Newton polytope of each polynomial by the value of the associated coefficient. Thanks to this duality, decision problems regarding tropical prevarieties can be reduced to the combinatorial problem of enumerating all the mixed cells that appear in the dual subdivision see [Jen16, Mal16]. A number of current works deal with the efficient computation of tropical varieties and prevariarieties (see [MR19, GRZ22] and the references therein) relying on these approaches, where *all* candidate solutions are typically constructed. In contrast, the application of the tropical Nullstellensatz or Positivstellensatz provided in this work, does not require the exhaustive enumeration of all the cells of a polyhedral complex (amounting to testing all the potential solutions). Instead, the feasibility or unfeasibility is directly decided by reduction to a mean payoff game. In the particular case of tropical linear inequalities, one can quantify the advantage of the latter approach over the former, as there are exponentially many tropically extreme solutions [AGK11a], hence exponentially many cells, rendering unfeasible the enumeration as soon as the dimension surpasses a few dozens. On the contrary, checking the feasibility reduces to solving a mean payoff game with a size linearly bounded in input size, which can be done for large systems with pseudo-polynomial complexity bounds. This advantage subsists for a significant class of higher degree instances, although the size of the game, that is the size of the certificate submatrix of the Macaulay matrix, now becomes exponential in the input size. The present approach is expected to be particularly useful in the Positivstellensatz case, where the exhaustive cell enumeration is especially heavy, while only a single feasible cell, or otherwise an unfeasibility certificate, is looked for.

Other approaches rely on the application of general-purpose SMT solving algorithms, see *e.g.* [Lü20]. Moreover, a tropical Nullstellensatz for *tropical ideals*, building upon [GP18], has been established in [MR18a]. Other tropical Nullstellensätze of varied natures have also been established, *e.g.* in [SI07, BE17, JM17, GP20].

As for the game theoretical aspect of this manuscript, the effective resolution of a tropical polynomial system relies on the computation of the value of the mean payoff game arising from its linearization. More precisely, links between solutions of tropical linear systems and the vector of value of a mean payoff game have been established in [AGG12], where the authors proved that there is a one-to-one correspondance between finite tropical linear systems and mean payoff games with a finite number of states, and that the maximal support (in the tropical sense, meaning the indices of the coefficients different to $-\infty$) of a solution of a tropical linear system are in correspondance with the initial winning moves for the first player of the associated mean payoff game. The determination of these initial winning moves relies on the computation of the vector of value of the game, and for mean payoff game, different approches have been proposed, notably the value iteration algorithm [ZP96] and the policy iteration algorithm [CTGG99, DG06, Cha09]. The question of solving parametric mean payoff games, which is at the core of the effective resolution of tropical polynomial systems, has already been studied in a more restrictive context in [GKS12], where the *spectral function*, that is the application mapping a value of the parameter onto the vector of values of the parametric mean payoff game at that particular parameter choice, is computed via a Newton-like algorithm.

4 Notation index

We summarize in the following list the notation and conventions that shall be used throughout this manuscript

- $\mathbb{N} = \{0, 1, 2, \ldots\}$ denote the set of natural numbers, with zero included. The set of strictly positive integers is denoted by $\mathbb{N}_{>0} = \{1, 2, 3, \ldots\}$.
- For any integer $n \in \mathbb{N}_{>0}$, one denotes by $[n] = \{1, \ldots, n\}$ the set of integers between 1 and n.
- $\mathbb{R}$ denotes the set of real number and $\mathbb{R}_{\geqslant 0}$ denotes the set of nonnegative real numbers.
- $\mathbb{T} = \mathbb{R} \cup \{-\infty\}$ denotes the tropical semifield, and $\mathbb{T}^* = \mathbb{R}$ denote the set of tropically nonzero elements.
- The tropical operations are $\oplus = \max$ and $\odot = +$. Moreover, we set $\mathbb{0} = -\infty$ and $\mathbb{1} = 0$.
- Given a n-variate tropical polynomial $f = \bigoplus_{\alpha \in \mathcal{A}} f_\alpha X^\alpha$ and a point $x \in \mathbb{R}^n$, we detone $f(x) \mathrel{\nabla} \mathbb{0}$ whenever the maximum in the expression $f(x) = \max_{\alpha \in \mathcal{A}} f_\alpha + \langle x, \alpha \rangle$ is achieved for at least two distinct values of $\alpha \in \mathcal{A}$.
- The cardinality of a set X is denoted by $|X|$.

For the rest of this section, fix $n \in \mathbb{N}_{>0}$ a strictly positive integer.

- The standard scalar product on $\mathbb{R}^n$ is denoted by $\langle \cdot, \cdot \rangle$, where we recall that $\langle u, v \rangle = \sum_{j=1}^n u_j v_j$ if $u = (u_j)_{j \in [n]} \in \mathbb{R}^n$ and $v = (v_j)_{j \in [n]} \in \mathbb{R}^n$.
- For $p \geqslant 1$, the p-norm on $\mathbb{R}^n$ is denoted by $\| \cdot \|_p$, and the sup-norm is denoted by $\| \cdot \|_\infty$. Recall that the p-norm is given for all $u = (u_j)_{j \in [n]} \in \mathbb{R}^n$ by $\|u\|_p = \left(\sum_{j=1}^n |u_j|^p \right)^{\frac{1}{p}}$ and the sup-norm by $\|u\|_\infty = \max_{j \in [n]}(|u_j|)$.
- The Hilbert seminorm on $\mathbb{R}^n$ is denoted by $\| \cdot \|_{\mathcal{H}}$. Recall that it is given for all $u = (u_j)_{j \in [n]} \in \mathbb{R}^n$ by $\|u\|_{\mathcal{H}} = \max_{1 \leqslant j \leqslant n}(u_j) - \min_{1 \leqslant j \leqslant n}(u_j)$.
- If $u = (u_j)_{j \in [n]}$ is a vector in $(\mathbb{R} \cup \{\pm\infty\})^n$, then $\lambda + u$ denotes the vector of $(\mathbb{R} \cup \{\pm\infty\})^n$ with entries $(\lambda + u_j)_{j \in [n]}$.
- For two vectors $u, v \in (\mathbb{R} \cup \{+\infty\})^n$, we write $v \leqslant u$ if for all $j \in [n]$, $v_j \leqslant u_j$, and $v \ll u$ if for all $j \in [n]$ such that $u_j < +\infty$, one has $v_j < u_j$.

Let X be a subset of $\mathbb{R}^n$. Then relatively to the standard topology of $\mathbb{R}^n$:

- $\mathrm{int}(X)$ denotes the interior of X;
- $\mathrm{relint}(X)$ denotes the relative interior of X;
- $\mathrm{cl}(X)$ denotes the closure of X.

Moreover, independently of the topology of $\mathbb{R}^n$,

- $\mathrm{vect}(X)$ denotes the vector subspace generated by X;
- $\mathrm{aff}(X)$ denotes the affine subspace generated by X, or affine hull of X;

- $\diamond$ $\mathrm{conv}(X)$ denotes the convex hull of X;
- $\diamond$ $\mathrm{recc}(X)$ denotes the recession cone of X.

Finally, given a Shapley operator $T : (\mathbb{R} \cup \{\pm\infty\})^n \to (\mathbb{R} \cup \{\pm\infty\})^n$,

- $\diamond$ $\mathrm{G}(T)$ denotes the graph associated to the mean payoff game of operator T;
- $\diamond$ $\mathrm{Eig}(T)$ denotes the nonlinear eigenspace of T, that is the set of vectors $u \in \mathbb{R}^n$ such that $T(u) = \lambda + u$ for $\lambda \in \mathbb{R}$.
- $\diamond$ $\mathrm{SAT}(T, u)$ denotes the saturation graph of the operator T evaluated at a nonlinear eigenvector $u \in \mathrm{Eig}(T)$.
- $\diamond$ $\mathbb{R}^n/\mathbb{R}\mathbf{1}$ is the vector space obtained by taking the quotient of the vector space $\mathbb{R}^n$ by the line directed by the vector $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^n$. In particular, two vectors $u, v \in \mathbb{R}^n$ are in the same quotient class if and only if there exists a constant $\lambda \in \mathbb{R}$ such that $v = \lambda + u$.

Chapter 1

Preliminaries

1.1 Notions of polyhedral geometry

In order to describe what we shall later call tropical prevarieties, that is intersection of tropical hypersurfaces, we first need to introduce a few notions from polyhedral geometry. For a more exhaustive exploration of polyhedral geometry, the interested reader can refer to [Zie95] on the general topic of polytopes, or to [DRS10] for the specific topic of polyhedral complexes, polyhedral subdivisions and triangulations.

Fix, for the remainder of this section, an integer $n \in \mathbb{N}_{>0}$.

1.1.1 Polyhedra, faces of a polyhedron

We start by describing polyhedra and giving some very general properties. First recall that a set C is called *convex* whenever any barycentric combination of two points in C remains within C.

Definition 1.1.1. A subset P of $\mathbb{R}^n$ of the form

$$P = \{x \in \mathbb{R}^n : \forall i \in [k],\ \langle a_i, x\rangle \leqslant b_i\} \tag{1.1}$$

where $k \in \mathbb{N}$, $a_i \in \mathbb{R}^n$ and $b_i \in \mathbb{R}$ for all $1 \leqslant i \leqslant k$ is called a *(closed convex) polyhedron*. In other words, a polyhedron is a subset of $\mathbb{R}^n$ defined by a set of affine inequalities. Moreover, if all these inequalities are in fact linear, meaning that $b_i = 0$ for all $1 \leqslant i \leqslant k$, then P is said to be a *polyhedral (closed convex) cone*.

Polyhedra are very structured subsets of $\mathbb{R}^n$. Some of their main properties are recalled in the following statement.

Property 1.1.2.

(a) A polyhedron $P \subseteq \mathbb{R}^n$ is a closed convex subset of $\mathbb{R}^n$.

(b) The empty set as well as all affine subspaces of $\mathbb{R}^n$ are polyhedra.

(c) If $P, Q \subseteq \mathbb{R}^n$ are two polyhedra, then so is their intersection $P \cap Q$.

(d) The projection of a polyhedron P onto any vector subspace of $\mathbb{R}^n$ forms a lower-dimensional polyhedron. *A fortiori*, the image of P by any linear map $\mathbb{R}^n \to \mathbb{R}^m$ with $m \in \mathbb{N}_{>0}$ is a polyhedron of $\mathbb{R}^m$.

Recall that the affine hull of a subset X of $\mathbb{R}^n$ is the intersection of all affine subspaces of $\mathbb{R}^n$ containing X. It is thus the smallest affine subspace of $\mathbb{R}^n$ which contains X. Equivalently, it is the set of all affine combinations of a finite number of points of X. This allows us to define the notion of dimension of a polyhedron.

Definition 1.1.3. The *dimension* $\dim(P)$ of a polyhedron $P \subseteq \mathbb{R}^n$ is the dimension of (the direction of) its affine hull. In particular, P is said to be *full-dimensional* whenever $\dim(P) = \dim(\mathbb{R}^n) = n$.

Let P be a convex subset of $\mathbb{R}^n$ in the following definition and proposition.

Definition 1.1.4. A closed convex subset F of P is called an *extremal face* of P whenever $tx + (1-t)y \in F$ for $x, y \in \mathbb{R}^n$ and for some $0 < t < 1$ implies that $x, y \in F$. Moreover, F is called an *exposed face* of P whenever there exists a linear form ϕ over $\mathbb{R}^n$ such that $F = \arg\max\{\phi(z) : z \in P\}$.

Recall that the Riesz representation theorem entails that for each linear form ϕ over $\mathbb{R}^n$ there exists a unique vector $c \in \mathbb{R}^n$ such that for all $z \in \mathbb{R}^n$, $\phi(z) = \langle c, z \rangle$. Therefore, the definition of exposed faces can be rewritten equivalently using scalar products.

Proposition 1.1.5. *All exposed faces of P are also extremal faces of P. Whenever P is a polyhedron, the converse also holds for nonempty extremal faces and the two notions coincide. In that case, they are simply refered to as faces of P — this definition includes the empty set — and the set of all faces of P is denoted by $\mathscr{F}(P)$. Moreover, all faces of a polyhedron are also polyhedra.*

Sketch of the proof. Let $F = \arg\max\{\langle c, z\rangle : z \in P\}$ with $c \in \mathbb{R}^n$ be an exposed face of P and let $x, y \in P$ and $0 < t < 1$ be such that $tx + (1-t)y \in F$. Then one has

$$\begin{aligned}\max_{z\in P}\langle c, z\rangle &= \langle c, tx + (1-t)y\rangle \\ &= t\langle c, x\rangle + (1-t)\langle c, y\rangle \\ &\leqslant t\left(\max_{z\in P}\langle c, z\rangle\right) + (1-t)\left(\max_{z\in P}\langle c, z\rangle\right) = \max_{z\in P}\langle c, z\rangle\end{aligned}$$

and the inequality is strict whenever $\langle c, x\rangle < \max_{z\in P}\langle c, z\rangle$ or $\langle c, y\rangle < \max_{z\in P}\langle c, z\rangle$, which entails a contradiction, hence $\langle c, x\rangle = \langle c, y\rangle = \max_{z\in P}\langle c, z\rangle$ and thus $x, y \in F$.

For the converse implication, the core of the proof resides in a classical result (see for instance [Roc70, §18]), stating the existence of a supporting hyperplane at any point x on the relative boundary of a nonempty polyhedron $P \subsetneq \mathbb{R}^n$, that is an hyperplane of the form $H := \{z \in \mathbb{R}^n : \langle a, z\rangle = b\}$ with $a \in \mathbb{R}^n$, such that $\langle a, x\rangle = b$ and $\langle a, z\rangle \leqslant b$ for all $z \in P$ but such that there exists $z \in P$ for which the inequality is strict, meaning that P is not contained in H. The idea of this direction of the proof is mainly contained in [Roc70, Theorem 19.1], see also [RG95] for a proof of the same result in the more general case of spectrahedra.

Finally, all faces of P are also polyhedra, as they are formed by intersecting P with a collection of affine hyperplanes. $\square$

We state some straight-forward properties of the faces of a polyhedron.

Property 1.1.6. Let $P \subseteq \mathbb{R}^n$ be a polyhedron. Then:

(a) the empty set as well as the polyhedron P are faces of P;

(b) the intersection of any two faces of P is a face of P;

(c) any face of a face of P is a face of P.

Recall the following vocabulary: a *proper face* of a polyhedron is a nonempty face which is also a proper subset, *i.e.* which is strictly included in the polyhedron. A maximal-dimensional proper face is called a *facet* and a one-dimensional face is called a *vertex*. If a polyhedron has at least one vertex, then it is called *pointed*. For a pointed polyhedral cone, the unique vertex is also refered to as the *apex*. Finally, the union of all the k-dimensional faces of a polyhedron P is called the *k-skeleton* of P.

Now consider a polyhedron $P \subseteq \mathbb{R}^n$ of the form of (1.1), and define

$$\mathcal{I} := \{i \in [k] : \exists x \in P,\ \langle a_i, x\rangle < b_i\} \subseteq [k] \ .$$

Then one can describe the relative interior of P as follows

$$\operatorname{relint}(P) = \left\{ x \in \mathbb{R}^n : \forall i \in [k], \left\{ \begin{array}{ll} \langle a_i, x\rangle < b_i & \text{if } i \in \mathcal{I} \\ \langle a_i, x\rangle = b_i & \text{otherwise} \end{array} \right. \right\} .$$

In particular, P is full-dimensional if and only if it has nonempty interior if and only if $\mathcal{I} = [k]$, in which case $\operatorname{int}(P) = \operatorname{relint}(P)$.

Similarly, one can describe the faces of P by constraining some of the inequalities in (1.1) to be equalities, and thus for all face F of P, there exists a unique set $\mathcal{I}_F \subseteq \mathcal{I}$ of indices such that

$$F = \left\{ x \in \mathbb{R}^n : \forall i \in [k], \left\{ \begin{array}{ll} \langle a_i, x\rangle \leqslant b_i & \text{if } i \in \mathcal{I}_F \\ \langle a_i, x\rangle = b_i & \text{otherwise} \end{array} \right. \right\} .$$

One furthermore obtains a description of $\operatorname{relint}(F)$ by replacing the weak inequalities in the above expression by strict inequalities. In particular, P is also a face of itself, and one has $\mathcal{I}_P = \mathcal{I}$. This leads to the following useful lemma.

Lemma 1.1.7. *Let $P \subseteq \mathbb{R}^n$ be a polyhedron and let x be a point of P. Then there exists a unique face F of P such that $x \in \operatorname{relint}(F)$.*

Sketch of the proof. Set $\mathcal{I}_x = \{i \in [k] : \langle a_i, x\rangle < b_i\}$. Then the face

$$F = \left\{ z \in \mathbb{R}^n : \forall i \in [k], \left\{ \begin{array}{ll} \langle a_i, z\rangle \leqslant b_i & \text{if } i \in \mathcal{I}_x \\ \langle a_i, z\rangle = b_i & \text{otherwise} \end{array} \right. \right\}$$

satisfies the claim. ☐

1.1.2 Polyhedral complexes, normal fan of a polyhedron

Definition 1.1.8. A finite collection $\mathscr{C}$ of polyhedra of $\mathbb{R}^n$ is called a *polyhedral complex* of $\mathbb{R}^n$ whenever it satisfies the two following conditions:

(i) every face of a polyhedron in $\mathscr{C}$ also belong in $\mathscr{C}$;

(ii) for all $C, D \in \mathscr{C}$, $C \cap D$ is a face of both C and D.

The polyhedra in a polyhedral complex are also refered to as *cells* of the complex. Moreover, if all the cells in $\mathscr{C}$ are polyhedral cones, then it is called a *fan*.

Remark 1.1.9. It follows readily from this definition that the empty set belongs to any polyhedral complex $\mathscr{C}$, and moreover that the relative interiors of polyhedra in $\mathscr{C}$ are pairwise distinct.

Definition 1.1.10. The *support* of a polyhedral complex $\mathscr{C}$ of $\mathbb{R}^n$ is the set $\operatorname{supp}(\mathscr{C})$ given by

$$\operatorname{supp}(\mathscr{C}) = \bigcup_{C \in \mathscr{C}} C \ .$$

In particular, the polyhedral complex $\mathscr{C}$ is said to be *complete* whenever $\operatorname{supp}(\mathscr{C}) = \mathbb{R}^n$.

Definition 1.1.11. Let $\mathscr{C}, \mathscr{D}$ be two polyhedral complexes of $\mathbb{R}^n$. Then:

(a) $\mathscr{D}$ is said to be a *subcomplex* of $\mathscr{C}$ whenever $\mathscr{D} \subseteq \mathscr{C}$;

(b) $\mathscr{D}$ is said to be a *refinement* of $\mathscr{C}$ whenever $\operatorname{supp}(\mathscr{C}) = \operatorname{supp}(\mathscr{D})$ and for all cell $D \in \mathscr{D}$, there exists a cell $C \in \mathscr{C}$ such that $D \subseteq C$. This is denoted as $\mathscr{D} \preccurlyeq \mathscr{C}$;

(c) the *common refinement* of $\mathscr{C}$ and $\mathscr{D}$ is the polyhedral complex $\mathscr{C} \wedge \mathscr{D}$ defined — even whenever $\operatorname{supp}(\mathscr{C}) \neq \operatorname{supp}(\mathscr{D})$ — by $\mathscr{C} \wedge \mathscr{D} = \{C \cap D : (C, D) \in \mathscr{C} \times \mathscr{D}\}$.

Remark 1.1.12. The relation $\preccurlyeq$ defines a partial order on the set of polyhedral complexes of $\mathbb{R}^n$ with a common given support. In particular, it forms a meet-semilattice, and the common refinement corresponds to the meet associated to this order.

We now proceed to describe some standard examples of polyhedral complexes and fans that we will be handling throughout this manuscript.

Proposition 1.1.13. *Let $P \subseteq \mathbb{R}^n$ be a polyhedron. Then the set $\mathscr{F}(P)$ of its faces forms a polyhedral complex.* A fortiori, *for all $0 \leqslant k \leqslant n$, the set of faces of dimension at most k of P also forms a polyhedral complex whose support corresponds to the k-skeleton of P.*

Proof. This statement is a straight-forward translation of Property 1.1.6. ☐

Definition 1.1.14. Let $P \subseteq \mathbb{R}^n$ be a polyhedron and let $x \in \mathbb{R}^n$. A vector $y \in \mathbb{R}^n$ is said to be *normal to P at point x* whenever $\langle y, z - x\rangle \leqslant 0$ for all $z \in P$. The *normal cone of P at the point x* is the cone $N_x(P)$ consisting of all vectors of $\mathbb{R}^n$ normal to P at point x, that is

$$N_x(P) := \{y \in \mathbb{R}^n : \forall z \in P,\ \langle y, z - x\rangle \leqslant 0\} \ .$$

Moreover, let F be a nonempty face of P. Then the *normal cone of P at the face F* is the set $N_F(P)$ defined by $N_F(P) = N_x(P)$ where x is an arbitrary point in the relative interior of F — this object is well-defined because it can easily be checked that two points in the relative interior of the same face of P yield the same normal cone.

Proposition 1.1.15. *Let $P \subseteq \mathbb{R}^n$ be a polyhedron. Then the collection*

$$\mathscr{N}(P) := \left\{ N_F(P) : F \in \mathscr{F}(P) \setminus \{\emptyset\} \right\}$$

of normal cones at every face of P forms a fan, called the normal fan of P.

Sketch of the proof. The idea of the proof resides in the following fact: the normal cone of a facet of P is the outer halfline orthogonal to (the hyperplane generated by) the facet, and for a general face of P, the normal cone of F is the cone generated all the halflines orthogonal to a facet of P containing F. It follows readily from this that a face of a normal cone $N_F(P)$ of a face F of P is the normal cone $N_G(P)$ of a face G of P containing F, and moreover that for two faces F and G of P, $N_F(P) \cap N_G(P)$ is the normal cone of the smallest face of P containing F and G. □

Remark 1.1.16. Note that the normal fan of a polyhedron $P \subseteq \mathbb{R}^n$ is the same as the normal fan of any translation $a + P$ of P. More precisely, if $a \in \mathbb{R}^n$, then the map $F \mapsto a + F$ yields a one-to-one correspondance between the faces of P and the faces of $a + P$, thus entailing $N_{a+F}(a + P) = N_F(P)$.

The set of faces $\mathscr{F}(P)$ of a polyhedron forms a lattice for the partial order given by the inclusion, and is therefore refered to as the *face lattice* of P. Moreover, this lattice is endowed with a grading given by the dimension. The same property holds for the normal fan $\mathscr{N}(P)$ of P, except that the order and the grading are reversed. In fact, these two lattices are related by the following duality result.

Proposition 1.1.17 (See [Tho06, pp. 62–63])**.** *Let P be a polyhedron in $\mathbb{R}^n$. Then the map*

$$\begin{array}{ccc} \mathscr{F}(P) & \longrightarrow & \mathscr{N}(P) \\ F & \longmapsto & N_F(P) \end{array}$$

is a poset anti-isomorphism — meaning an application between two posets which reverses inclusions — from the lattice of faces of P to the lattice of normal cones of P. In particular, $\dim(F) + \dim(N_F(P)) = n$ for all $F \in \mathscr{F}(P)$.

We deduce the following corollary from the previous proposition, which will help us deal with the case where the polyhedron we are working with is not full-dimensional.

Corollary 1.1.18. *Let P be a polyhedron in $\mathbb{R}^n$ and let W be the vector subspace of $\mathbb{R}^n$ directing the affine hull of P. Then for all $F \in \mathscr{F}(P)$,*

$$\dim(N_F(P) \cap W) = \dim(P) - \dim(F) \ .$$

Proof. Since translating a polyhedron does not affect its normal fan, we can assume without loss of generality that P is included in W. Then we notice that for any $x \in P$,

$$N_x(P) \cap W = \{y \in W : \forall z \in P, \ \langle y, z - x \rangle \leqslant 0\} \ .$$

This means that $N_x(P) \cap W$ corresponds to the normal cone at point x of the polyhedron P seen as a full dimensional polyhedron of W. We can then apply the previous proposition to obtain the equality

$$\dim(F) + \dim(N_F(P) \cap W) = \dim(W) = \dim(P) \ ,$$

from which the desired result follows. □

Finally, we define the geometric notions of vertical, upper and lower face using normal cones as follows.

Definition 1.1.19. Let $P \subseteq \mathbb{R}^{n+1} \simeq \mathbb{R}^n \times \mathbb{R}$ be a full-dimensional polyhedron and let F be a facet of P. Then, the facet F is said to be *vertical* whenever $N_F(P) \subseteq \mathbb{R}^n \times \{0\}$, and otherwise, F is said to be *non-vertical*. Moreover, if F is non-vertical, then it is called an *upper facet* of P (resp. *lower facet* of P) whenever there exists a vector $y = (y_i)_{i \in [n+1]} \in N_F(P)$ such that $y_{n+1} < 0$ (resp. $y_{n+1} > 0$). Finally, any face of an upper facet (resp. lower facet) is called an upper face (resp. lower face).

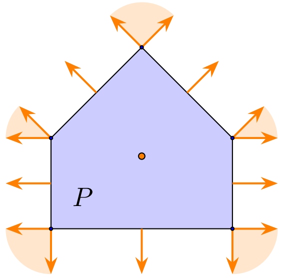

$$(x_1, x_2) \in P \iff \begin{cases} -x_1 & \leqslant & 0 \\ -x_2 & \leqslant & 0 \\ -x_1 + x_2 & \leqslant & 1 \\ x_1 + x_2 & \leqslant & 3 \\ x_2 & \leqslant & 2 \end{cases}$$

Figure 1.1: A polyhedron $P \subseteq \mathbb{R}^2$ represented with its normal fan.

1.1.3 Minkowski sum of two polyhedra

Now we describe the Minkowski sum, which constitutes a fundamental operation on convex sets, and *a fortiori* on polyhedra.

Definition 1.1.20. The *Minkowski sum* or *addition* of two subsets X and Y of $\mathbb{R}^n$ is the set denoted $X+Y$, defined as $X + Y = \{x + y : (x, y) \in X \times Y\}$. The sets X and Y are refered to as the *terms* or *summands* — or more rarely *factors* — of the Minkowski sum.

The following statement lists some of the main properties of the Minkowski sum.

Property 1.1.21.

(a) The Minkowski addition defines an internal law on the power set of $\mathbb{R}^n$, which is commutative, associative and admits a zero element given by the singleton $\{0\}$.

(b) Given a collection $X_1, \dots, X_k$ of k subsets of $\mathbb{R}^n$, one has

$$X_1 + \cdots + X_k = \{x_1 + \cdots + x_k : (x_1, \dots, x_k) \in X_1 \times \cdots \times X_k\} \ .$$

(c) The Minkoswki addition preserves closure, openness, boundedness, compactness as well as convexity. Moreover, the Minkowski sum of two polyhedra is a polyhedron and the Minkowski sum of two (polyhedral) cones is a (polyhedral) cone.

(d) If P and Q are two polyhedra, then $\dim(P + Q) \leqslant \dim(P) + \dim(Q)$.

The Minkowski sum also satisfies the following cancellation law for compact convex subsets of $\mathbb{R}^n$, and *a fortiori* for polyhedra.

Proposition 1.1.22. *Let X, Y, Z be three closed convex subsets of $\mathbb{R}^n$ such that Z is bounded and assume that $X + Z \subseteq Y + Z$. Then $X \subseteq Y$.*

Proof. Assume that $X \nsubseteq Y$ and consider $x \in X \setminus Y$. Let π be the projection onto the closed convex set Y. Then, by Hilbert projection theorem, one has $\langle x - \pi(x), y - \pi(x)\rangle \leqslant 0$ for all $y \in Y$, while $\langle x - \pi(x), x - \pi(x)\rangle = \|x - \pi(x)\|_2^2 > 0$ since $x \notin Y$. Moreover, since Z is compact, there exists $z^* \in Z$ such that $\langle x - \pi(x), z^*\rangle = \max_{z\in Z}\langle x - \pi(x), z\rangle$. Therefore, it follows that for all $(y, z) \in Y \times Z$, $\langle x - \pi(x), y + z\rangle < \langle x - \pi(x), x + z^*\rangle$, and thus $X + Z \nsubseteq Y + Z$. □

Recall that the convex hull $\operatorname{conv}(X)$ of a subset X of $\mathbb{R}^n$ is the intersection of all convex sets containing X, and thus it is the smallest convex set containing X in the sense of the inclusion. Equivalently, it coincides with the set of barycentric combinations of a finite number of points of X. Note that Minkowski addition commutes with the convex hull operation in the following sense.

Proposition 1.1.23. *Let X and Y be two subsets of $\mathbb{R}^n$, then* $\operatorname{conv}(X + Y) = \operatorname{conv}(X) + \operatorname{conv}(Y)$.

Proof. Since $\operatorname{conv}(X)+\operatorname{conv}(Y)$ is a convex set containing $X+Y$, it therefore contains $\operatorname{conv}(X+Y)$. Conversely, let

$$x = \sum_{i=1}^{k} \lambda_i x_i \in X \quad \text{with} \quad k \in \mathbb{N}_{>0},\ x_1, \dots, x_k \in X \text{ and } \lambda_1, \dots, \lambda_k \in \mathbb{R}_{\geqslant 0} \text{ such that } \sum_{i=1}^{k} \lambda_i = 1$$

and

$$y = \sum_{j=1}^{\ell} \mu_j y_j \in Y \quad \text{with} \quad \ell \in \mathbb{N}_{>0},\, y_1, \dots, y_\ell \in Y \text{ and } \mu_1, \dots, \mu_\ell \in \mathbb{R}_{\geqslant 0} \text{ such that } \sum_{j=1}^{\ell} \mu_j = 1 \ .$$

Then

$$x + y = \sum_{i=1}^{k} \lambda_i \underbrace{\left(\sum_{j=1}^{\ell} \mu_j (x_i + y_j) \right)}_{\cdot \in \operatorname{conv}(X+Y)} \in \operatorname{conv}(X + Y) \ ,$$

hence the converse inclusion. □

The following result of convex geometry lets us deal with Minkowski sums of convex sets with a large number of summands.

Lemma 1.1.24 (Shapley-Folkman, see [Sch13, Theorem 3.1.2]). *Let $X_1, \dots, X_k \subseteq \mathbb{R}^n$, and let*

$$x \in \sum_{i=1}^{k} \operatorname{conv}(X_i) \ .$$

Then there is an index set $I \subseteq [k]$ with $|I| \leqslant n$ such that

$$x \in \sum_{i \in I} \operatorname{conv}(X_i) + \sum_{i \in [k] \setminus I} X_i \ .$$

The Shapley-Folkman is mainly a consequence of the Carathéodory theorem, which states that for any set $X \subseteq \mathbb{R}^n$, if $x \in \operatorname{conv}(X)$, then x can be written as a barycentric combination of at most $n+1$ elements of X. However, in the same way that these $n+1$ elements depend on the point x, in the Shapley-Folkman lemma, the index set I cannot be prescribed and depends on the point x considered.

In the case where the Minkowski sum of the X_i is not full-dimensional, one has the following corollary.

Corollary 1.1.25. *If in Lemma 1.1.24 $\sum_{i=1}^{k} \operatorname{conv}(X_i)$ has (affine) dimension $d < n$, then the index set I can be choosen such that $|I| \leqslant d$.*

Proof. Let V be the vector space directing the affine hull of $\sum_{i=1}^{k} \operatorname{conv}(X_i)$, and let T be an isomorphism from $\mathbb{R}^d$ to V, and choose $x_i \in X_i$ for $i = 1, \dots, k$. We have $X_i = x_i + T(Y_i)$ for some subset Y_i of $\mathbb{R}^d$ containing 0, and $\operatorname{conv}(X_i) = x_i + T(\operatorname{conv}(Y_i))$. Applying Lemma 1.1.24 to the sets $Y_1, \dots, Y_k$ gives the result. □

The following proposition allows one to express the normal fan of a Minkowski sum of polyhedra in function of the normal fans of the terms of the Minkowski sum.

Proposition 1.1.26 (See [Zie95, Proposition 7.12]). *Let P, Q be two polyhedra of $\mathbb{R}^n$. Then the normal fan of the Minkowski sum of P and Q corresponds the common refinement of their normal fans,* i.e.

$$\mathscr{N}(P + Q) = \mathscr{N}(P) \wedge \mathscr{N}(Q) \ .$$

One also has the following weaker result for the face lattice of a Minkowski sum of polyhedra.

Proposition 1.1.27. *Let P, Q be two polyhedra of $\mathbb{R}^n$. Then any face of the Minkowski sum of P and Q can be written as a sum of a face of P and a face of Q,* i.e.

$$\mathscr{F}(P + Q) \subseteq \{F + G : (F, G) \in \mathscr{F}(P) \times \mathscr{F}(Q)\} \ .$$

However, this inclusion may be strict.

Sketch of the proof. Let z be a point in a face H of $P + Q$, and let $x \in P$ and $y \in Q$ be such that $z = x + y$. If c is a vector in the normal cone of H, then

$$\max_{w \in P+Q} \langle c, w \rangle = \langle c, z \rangle = \langle c, x \rangle + \langle c, y \rangle \leqslant \max_{u \in P} \langle c, u \rangle + \max_{v \in Q} \langle c, v \rangle = \max_{w \in P+Q} \langle c, w \rangle \ .$$

However if $\langle c, x \rangle < \max_{u \in P} \langle c, u \rangle$ or $\langle c, y \rangle < \max_{v \in Q} \langle c, v \rangle$, then the inequality would be strict, thus entailing a contradiction. Therefore $\langle c, x \rangle = \max_{u \in P} \langle c, u \rangle$ and $\langle c, y \rangle = \max_{v \in Q} \langle c, v \rangle$, and thus

$$z \in F + G \quad \text{with} \quad \left\{ \begin{array}{lcl} F & = & \arg\max\{\langle c, u \rangle : u \in P\} \in \mathscr{F}(P) \\ G & = & \arg\max\{\langle c, v \rangle : v \in Q\} \in \mathscr{F}(Q) \ . \end{array} \right.$$

Conversely, if $z \in F + G$ then the same reasoning shows that z must belong in the face H. □

1.1.4 Polytopes, subdivisions of a polytope

In this section, we focus on a particular class of polyhedra obtained by taking the convex hull of a finite collection of points.

Definition 1.1.28. A subset P of $\mathbb{R}^n$ of the form

$$P = \operatorname{conv}(\{x_1, \dots, x_k\}) \tag{1.2}$$

where $k \in \mathbb{N}_{>0}$ and $x_1, \dots, x_k$ is a finite collection of points of $\mathbb{R}^n$ is called a *polytope*.

A central example of polytope is given by the case where the polytope P in (1.2) is $k-1$-dimensional. In that case, P is called a *simplex*. Equivalently, a simplex is a polytope that is generated by an affinely independent family of points of $\mathbb{R}^n$.

As stated above, a polytope is a special case of polyhedron. More precisely, one has the following result.

Proposition 1.1.29. *A polytope is a polyhedron. Moreover, if P is a polyhedron, then the following are equivalent:*

(i) P is a polytope;

(ii) P is bounded;

(iii) the normal fan of P is complete;

(iv) P does not contain a halfline.

The previous proposition is in fact a direct consequence of the stronger Minkowski-Weyl theorem.

Theorem 1.1.30 (Minkowski-Weyl, see for instance [Cha21, Theorem 1.3.4])**.** *A subset P of $\mathbb{R}^n$ is a polyhedron if and only if there exist a polytope Q and a polyhedral cone C such that $P = Q + C$.*

The cone C of the previous theorem can be described precisely with the following results.

Definition 1.1.31. The *recession cone* of a polyhedron P of $\mathbb{R}^n$ (or more generally of any subset of $\mathbb{R}^n$) is the cone $\operatorname{recc}(P)$ defined by

$$\{y \in \mathbb{R}^n : \forall x \in \mathbb{R}^n, \, \forall t \geqslant 0, \, x + ty \in P\} \ .$$

Proposition 1.1.32 ([Cha21, Proposition 1.3.9])**.** *Let P be a polyhedron of $\mathbb{R}^n$, such that $P = Q + C$ where Q and C are respectively a polytope and a polyhedral cone C of $\mathbb{R}^n$. Then $C = \operatorname{recc}(P)$. In particular, P is a polytope if and only if $\operatorname{recc}(P) = \{0\}$.*

As a consequence of the Minkowski-Weyl theorem, there exist two equivalent ways of describing a polytope P of $\mathbb{R}^n$. It can be described by giving a set of points of $\mathbb{R}^n$ whose convex hull is equal to P. This is refered to as a V-*representation* or V-*description* of P — where the V stands for 'vertices'. Alternatively, it can be described by giving a list of affine inequalities whose solution set is equal to P. This is refered to as a H-*representation* or H-*description* of P — where the H stands for 'hyperplanes'. A key issue in computational geometry when dealing with polytopes resides in obtaining one of these descriptions from the other one, as depending on the algorithmic context, the V-representation or the H-representation of a polytope might be more suited. Different algorithms have been developed in order to compute one representation from the other, notably the double description method (see [MRTT53, FP96]). Having simultaneously access to both representations of a polytope is especially useful as it can reduce the cost of different manipulations such as computing a Minkowski sum of polytopes or computing a projection of a polytope. However, the double description method itself is inevitably quite costly, as it was shown not to be output-sensitive, and there is a class of polytopes for which the complexity of the double description method is superpolynomial [Bre99].

Another recurring problem, given either a V- or H-representation of a polytope P, is to decide whether it is minimal or not, in the sense that there is no redundant point in the generators of the convex hull in the former case, or no redundant inequality in the affine system in the latter case. In fact, computing an optimal V-representation of P is equivalent to computing its vertices, while computing an optimal H-representation of P is equivalent to computing its normal fan.

Lastly, we mention subdivisions and triangulations of polytopes.

Definition 1.1.33. A *polyhedral subdivision* or *polyhedral decomposition* $\mathscr{S}$ of a polytope P is a refinement of its face lattice $\mathscr{F}(P)$ seen as a polyhedral complex, *i.e.* it is any polyhedral complex $\mathscr{S}$ such that $\mathscr{S} \preccurlyeq \mathscr{F}(P)$. In particular, if $\mathscr{S}$ is a simplicial complex — meaning that all the polyhedra of $\mathscr{S}$ are simplexes — then it is said to be a *triangulation* of P.

Among these subdivisions, some have special properties that are of particular interest in the context of polynomial system solving and tropical geometry, being able to encode the geometry of an arrangment of tropical hypersurfaces. We shall take a particular interest in the following class of subdivisions.

Definition 1.1.34. Let $P \subseteq \mathbb{R}^n$ be a polytope and denote by $\pi : \mathbb{R}^{n+1} \simeq \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}^n$ the first projection. Then a subdivision $\mathscr{S}$ of P is said to be *coherent* or *regular* whenever there exists a polyhedron $Q \subseteq \mathbb{R}^{n+1}$ such that $P = \pi(Q)$ and every cell of $\mathscr{S}$ can be obtained as the image by π of an upper face of Q.

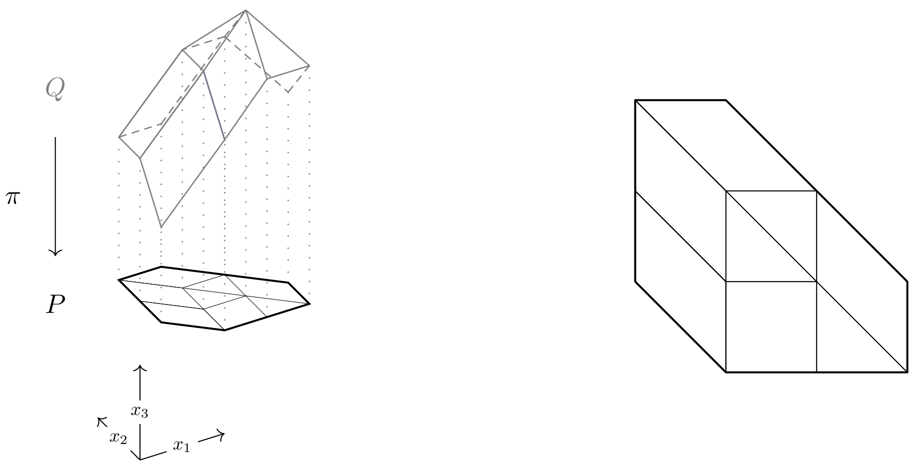

Figure 1.2: A coherent subdivision of a polytope $P \subseteq \mathbb{R}^2$.

Remark 1.1.35. Note that not all subdivisions of a polytope can be achieved by this process, or in other words: there exist subdivisions of a polytope that are not coherent. However, in the following chapter of this manuscript, we shall only be interested in coherent subdivision, as these subdivisions will play a central role in tropical geometry.

1.2 Algebraic and semialgebraic objects

In the previous section, we have discussed polyhedral sets, which are described by linear equalities and inequalities. We shall now describe a broder class of objects, known as algebraic or semialgebraic, consisting in sets defined by polynomial equations and inequations. For more thourough details, the curious reader is advised to check [CLO15].

For the remainder of this section, we fix $n \in \mathbb{N}_{>0}$ an integer, as well as $\mathbb{K}$ a field, which may be thought of as $\mathbb{C}$ for the purpose of our concerns.

1.2.1 Affine and projective algebraic varieties

In order to introduce algebraic varieties, we start by recalling the precise definition of a polynomial as well as the standard vocabulary to describe polynomials. Let $(R, +, \cdot, 0, 1)$ be a commutative ring in the following statements.

Definition 1.2.1. A *formal Laurent polynomial* f in n variables over R is a map

$$\begin{array}{ccc} \mathbb{Z}^n & \longrightarrow & R \\ \alpha & \longmapsto & f_\alpha \end{array}$$

such that $f_\alpha = 0$ for all $\alpha \in \mathbb{Z}^n$ but a finite number. Setting $\alpha = (\alpha_1, \ldots, \alpha_n)$, one uses the following notation:

$$f = \sum_{\alpha \in \mathbb{Z}^n} f_\alpha X^\alpha = \sum_{\alpha \in \mathbb{Z}^n} f_\alpha X_1^{\alpha_1} \cdots X_n^{\alpha_n} \ .$$

The set of formal n-variate Laurent polynomials over the ring R is denoted by $R[X_1^{\pm 1}, \ldots, X_n^{\pm 1}]$.

The *support* of a formal Laurent polynomial $f \in R[X_1^{\pm 1}, \ldots, X_n^{\pm 1}]$ is the set $\mathrm{supp}(f)$ given by

$$\mathrm{supp}(f) := \{\alpha \in \mathbb{Z}^n : f_\alpha \neq 0\} \ .$$

Following the standard vocabulary of symbolic computation, for $\alpha \in \mathrm{supp}(f)$, the *term of exponent* α designates the expression $f_\alpha X^\alpha$, in which the scalar f_α is refered to as the *coefficient of exponent* α, and the power product $X^\alpha = X_1^{\alpha_1} \cdots X_n^{\alpha_n}$ as the *monomial of exponent* α.

Finally, the *polynomial function* associated to a formal Laurent polynomial $f \in R[X_1^{\pm 1}, \ldots, X_n^{\pm 1}]$ is the function

$$\begin{array}{ccc} (R^\times)^n & \longrightarrow & R \\ x & \longmapsto & \sum_{\alpha \in \mathbb{Z}^n} f_\alpha x^\alpha \end{array} \ ,$$

where $R^\times$ denotes the multiplicative group of invertible elements of R, and $x^\alpha = x_1^{\alpha_1} \cdots x_n^{\alpha_n}$ for $x = (x_1, \ldots, x_n)$ and $\alpha = (\alpha_1, \ldots, \alpha_n)$. More generally, any function of the previous form is called a polynomial function.

In the case where the support of f is included in $\mathbb{N}^n$, it is simply refered to as a *formal polynomial*, and the set of formal n-variate polynomials over S is denoted by $R[X_1, \ldots, X_n]$. Moreover, in that case, the polynomial function is actually defined on all R^n rather than just on $(R^\times)^n$. In that case, recall that the *total degree* (or simply *degree*) of f is given by

$$\deg(f) := \max\{\|\alpha\|_1 = \alpha_1 + \cdots + \alpha_n : \alpha \in \mathrm{supp}(f)\} \ .$$

Finally, a n-variate polynomial f over R is called *homogeneous of degree* d whenever it only consists of terms of degree d, that is $\alpha_1 + \cdots + \alpha_n = d$ for all $\alpha = (\alpha_1, \ldots, \alpha_n) \in \mathrm{supp}(f)$.

Remark 1.2.2. Note that a formal polynomial f over R is homogeneous of degree d precisely whenever its associated polynomial function is d-homogeneous, *i.e.* whenever $f(\lambda x) = \lambda^d f(x)$ for all $\lambda \in R$ and for all $x \in R^n$.

We also recall that the addition and multiplication on R expand naturally on $R[X_1^{\pm}, \ldots, X_n^{\pm}]$ and $R[X_1, \ldots, X_n]$, endowing them of a ring structure, with zero and unit element respectively given by the zero polynomial 0 and the constant 1 polynomial. The addition for two polynomial over R is simply obtained by taking the sum coefficient-wise, and the multiplication is simply obtained by taking the Cauchy product of the coefficients.

Definition 1.2.3. Let $f \in R[X_1^{\pm}, \ldots, X_n^{\pm}]$ be a formal Laurent polynomial over R with support $\mathcal{A} \subseteq \mathbb{Z}^n$. An element $x = (x_1, \ldots, x_n)$ of $(R^\times)^n$ is called a *root* or *zero* of f whenever the associated polynomial function vanishes at x, that is whenever $f(x) = 0$.

Moreover, if the support $\mathcal{A}$ of f is included in $\mathbb{N}^n$, then a root x of f is instead defined to be an element R^n satisfying the same property, and in this case one defines the *support* of the root x as the set of indices $1 \leqslant j \leqslant n$ such that $x_j \neq 0$.

For the sake of readbility of the text, we shall drop the adjective *formal* in the remainder whenever the context is clear, and simply refer to formal (Laurent) polynomials as (Laurent) *polynomials*. This does not pose an issue in the usual case of polynomials over a 0-characteristic field, because in this case, there is a one-to-one correspondance between formal polynomials and polynomial functions, allowing us to identify both objects. However, this correspondance may be compromised in a more general case, for instance over rings with nonzero characteristic. We will also see later on that it does not hold for tropical polynomials, prompting us to proceed with a lot of caution.

From now on, we shall solely be concerned in the case of polynomials over a field $\mathbb{K}$ — algebraically closed if necessary — and denote by $\mathfrak{R} := \mathbb{K}[X_1, \ldots, X_n]$ the ring of n-variate polynomials over $\mathbb{K}$, and take interest in the zero locii of polynomials of $\mathfrak{R}$

Definition 1.2.4. The *affine hypersurface* (or simply *hypersurface*) *associated to a polynomial* $f \in \mathfrak{R}$ is the subset $\mathcal{V}_{\mathbb{K}}(f)$ of $\mathbb{K}^n$ defined by

$$\mathcal{V}_{\mathbb{K}}(f) := \{x \in \mathbb{K}^n : f(x) = 0\} \ .$$

In other words, it corresponds to the zero locus of the polynomial function associated to f. Any subset of $\mathbb{K}^n$ of this form is called an *hypersurface* of $\mathbb{K}^n$.

Likewise the *affine variety* (or simply *variety*) *associated to a collection* $f_1, \ldots, f_k \in \mathfrak{R}$ *of polynomials* is the subset $\mathcal{V}_{\mathbb{K}}(f_1, \ldots, f_k)$ of $\mathbb{K}^n$ defined by

$$\mathcal{V}_{\mathbb{K}}(f_1, \ldots, f_k) := \{x \in \mathbb{K}^n : \forall i \in [k],\ f_i(x) = 0\} \ .$$

It thus corresponds to the intersection of all the hypersurfaces associated to $f_1, \dots, f_k$. Any subset of $\mathbb{K}^n$ of this form is called a *variety* or an *algebraic subset* of $\mathbb{K}^n$.

Whenever the context is clear enough, we may sometime omit the field $\mathbb{K}$ in the notation of the previous variety and simply denote it as $\mathcal{V}(f_1, \dots, f_k)$.

More generally, given any subset S of $\mathfrak{R}$, we shall denote by $\mathcal{V}_{\mathbb{K}}(S)$ the subset of $\mathbb{K}^n$ defined by

$$\mathcal{V}_{\mathbb{K}}(S) := \{x \in \mathbb{K}^n : \forall f \in S,\ f(x) = 0\} \ .$$

Remark 1.2.5. One can also take interest in the intersection of an affine variety with the torus $(\mathbb{K}^*)^n$. This allows one more generally to describe the set of roots of a collection of Laurent polynomials.

Affine varieties are stable by union and intersection, as illustrated by the following proposition.

Proposition 1.2.6. *Let $k, \ell \in \mathbb{N}_{>0}$, and let $f_1, \dots, f_k$ and $g_1, \dots, g_\ell$ be two collections of polynomials in $\mathfrak{R}$. Then*

(a) $\mathcal{V}(f_1, \dots, f_k) \cap \mathcal{V}(g_1, \dots, g_\ell) = \mathcal{V}(f_1, \dots, f_k, g_1, \dots, g_\ell)$;

(b) $\mathcal{V}(f_1, \dots, f_k) \cup \mathcal{V}(g_1, \dots, g_\ell) = \mathcal{V}(f_i g_j : i \in [k],\ j \in [\ell])$.

Proof. The first equality follows readily from the definition of the variety associated to a collection of polynomials. For the second equality, the direct inclusion is straight-forward since for $x \in \mathbb{K}^n$ and $(i, j) \in [k] \times [\ell]$, if $f_i(x) = 0$ or $g_j(x) = 0$, then $f_i g_j(x) = 0$, and conversely, if $f_i g_j(x) = 0$ for all $(i, j) \in [k] \times [\ell]$, then if $x \notin \mathcal{V}(f_1, \dots, f_k)$, there exists $i_0 \in [k]$ such that $f_{i_0}(x) \neq 0$, but for all $j \in [\ell]$, the equality $f_{i_0} g_j(x) = 0$ implies that $g_j(x) = 0$, and thus $x \in \mathcal{V}(g_1, \dots, g_\ell)$. □

Varieties are very closely related to the notion of ideal, whose definition, as well as key properties are recalled below.

Definition 1.2.7. An *ideal* of $\mathfrak{R}$ is a subset $\mathfrak{a}$ of $\mathfrak{R}$ satisfying the following properties

(i) $\mathfrak{a}$ is an additive subgroup of $\mathfrak{R}$;

(ii) if $f \in \mathfrak{a}$, then $hf \in \mathfrak{a}$ for all $h \in \mathfrak{R}$.

If $\mathfrak{a}$ and $\mathfrak{b}$ are two ideals of $\mathfrak{R}$, then their *sum* is defined as

$$\mathfrak{a} + \mathfrak{b} = \{f + g : f \in \mathfrak{a},\ g \in \mathfrak{b}\}$$

and their *product* is defined as

$$\mathfrak{a}\mathfrak{b} = \left\{ \sum_{i=1}^{k} f_i g_i : k \in \mathbb{N}_{>0},\ \forall i \in [k],\ f_i \in \mathfrak{a},\ g_i \in \mathfrak{b} \right\} \ .$$

Property 1.2.8.

(a) $\{0\}$ and $\mathfrak{R}$ are ideals of $\mathfrak{R}$.

(b) If $\mathfrak{a}, \mathfrak{b}$ are two ideals of $\mathfrak{R}$, then so are $\mathfrak{a} \cap \mathfrak{b}$, $\mathfrak{a} + \mathfrak{b}$ and $\mathfrak{a}\mathfrak{b}$.

In particular, since the intersection operation preserves ideals, one can define the *ideal generated by a subset S of* $\mathfrak{R}$ as the intersection of all ideals of $\mathfrak{R}$ containing S. It is in particular the smallest ideal of $\mathfrak{R}$ containing S for the order given by the set inclusion, and it is denoted by $\langle S \rangle$. If $S = \{f_1, \dots, f_k\}$ is a finite subset of $\mathfrak{R}$, then it is simply denoted by $\langle f_1, \dots, f_k \rangle$. A finite generating set for an ideal is usually called a *basis* of that ideal. The following theorem states that one can in fact always find a basis of an ideal of $\mathfrak{R}$.

Theorem 1.2.9 (Hilbert basis theorem, see [CLO15, p. 77])**.** *The ring $\mathfrak{R}$ of n-variate polynomials over $\mathbb{K}$ is* Noetherian, *meaning that every ideal of $\mathfrak{R}$ is finitely generated.*

Proposition 1.2.10. *Let $\mathfrak{a} = \langle f_1, \dots, f_k \rangle$ be an ideal of $\mathfrak{R}$. Then* $\mathfrak{a} = \left\{ \sum_{i=1}^{k} h_i f_i : \forall i \in [k],\ h_i \in \mathfrak{R} \right\}$.

Proof. The reverse inclusion $\supseteq$ is immediate by property of ideals. For the direct inclusion $\subseteq$, the set on the righthandside of the equality contains $f_1, \dots, f_k$, and moreover it is straight-forward to check that this set is an ideal, thus this entails by minimality of the generated ideal that it contains $\mathfrak{a} = \langle f_1, \dots, f_k \rangle$. □

We deduce from the previous proposition the following corollary, stating that an affine variety associated to a collection of polynomials depends only on the ideal generated by these polynomials.

Corollary 1.2.11. *Let $\mathfrak{a}$ be an ideal of $\mathfrak{R}$, and let $(f_1, \ldots, f_k)$ be a basis of $\mathfrak{a}$. Then for all $x \in \mathbb{K}^n$,*

$$f_1(x) = \cdots = f_k(x) = 0 \quad \Longleftrightarrow \quad \forall f \in \mathfrak{a},\ f(x) = 0 \ .$$

In particular, one has $\mathcal{V}(\mathfrak{a}) = \mathcal{V}(f_1, \ldots, f_k)$.

Proof. The first equivalence follows readily from Proposition 1.2.10, and the equality $\mathcal{V}(\mathfrak{a}) = \mathcal{V}(f_1, \ldots, f_k)$ follows from this equivalence. □

The following properties of the variety of an ideal can be easily checked from the previous definitions and results.

Property 1.2.12. Let $\mathfrak{a}, \mathfrak{b}$ be two ideals of $\mathfrak{R}$. Then:

(a) $\mathcal{V}(\{0\}) = \mathbb{K}^n$ and $\mathcal{V}(\mathfrak{R}) = \emptyset$;

(b) if $\mathfrak{a} \subseteq \mathfrak{b}$, then $\mathcal{V}(\mathfrak{b}) \subseteq \mathcal{V}(\mathfrak{a})$;

(c) $\mathcal{V}(\mathfrak{a}) \cap \mathcal{V}(\mathfrak{b}) = \mathcal{V}(\mathfrak{a} + \mathfrak{b})$;

(d) $\mathcal{V}(\mathfrak{a}) \cup \mathcal{V}(\mathfrak{b}) = \mathcal{V}(\mathfrak{a} \cap \mathfrak{b}) = \mathcal{V}(\mathfrak{a}\mathfrak{b})$.

Definition 1.2.13. Let V be a subset of $\mathbb{K}^n$. The *vanishing ideal of V* is the ideal $\mathfrak{I}(V)$ of $\mathfrak{R}$ defined by

$$\mathfrak{I}(V) = \{f \in \mathfrak{R} : \forall x \in V,\ f(x) = 0\} \ .$$

The ideal of a subset of $\mathbb{K}^n$ verifies the immediate following properties.

Property 1.2.14. Let V, W be two subsets of $\mathbb{K}^n$. Then:

(a) $\mathfrak{I}(\emptyset) = \mathfrak{R}$ and $\mathfrak{I}(\mathbb{K}^n) = \{0\}$;

(b) if $V \subseteq W$, then $\mathfrak{I}(W) \subseteq \mathfrak{I}(V)$

(c) $\mathfrak{I}(V) \cap \mathfrak{I}(W) = \mathfrak{I}(V \cup W)$;

(d) $\mathfrak{I}(V) + \mathfrak{I}(W) \subseteq \mathfrak{I}(V \cap W)$.

Varieties and ideals of definitions are related by the following proposition.

Proposition 1.2.15.

(a) Let $\mathfrak{a}$ be an ideal of $\mathfrak{R}$. Then $\mathfrak{a} \subseteq \mathfrak{I}(\mathcal{V}(\mathfrak{a}))$.

(b) Let V be a subset of $\mathbb{K}^n$. Then $V \subseteq \mathcal{V}(\mathfrak{I}(V))$, with equality whenever V is a variety of $\mathbb{K}^n$.

Proof. For the first inclusion, let $(f_1, \ldots, f_k)$ be a set of generators of $\mathfrak{a}$. Then by definition, $f_i \in \mathfrak{I}(\mathcal{V}(\mathfrak{a}))$ for all $i \in [k]$, hence $\mathfrak{a} \subseteq \mathfrak{I}(\mathcal{V}(\mathfrak{a}))$ by minimality of the generated ideal.

For the second inclusion, let $x \in V$. Then by definition of the ideal of V, for all $f \in \mathfrak{I}(V)$, $f(x) = 0$, which translates precisely into the fact that $x \in \mathcal{V}(\mathfrak{I}(V))$. Moreover, if V is a variety, let $f_1, \ldots, f_k \in \mathfrak{R}$ be such that $V = \mathcal{V}(f_1, \ldots, f_k)$. Then by the inclusion proven above, one has $\langle f_1, \ldots, f_k \rangle \subseteq \mathfrak{I}(\mathcal{V}(f_1, \ldots, f_k))$, and thus it follows that $\mathcal{V}(\mathfrak{I}(\mathcal{V}(f_1, \ldots, f_k))) \subseteq \mathcal{V}(f_1, \ldots, f_k) = V$ from Property 1.2.12 (b), hence the result. □

1.2.2 Semialgebraic sets

After discussing algebraic objects in the previous section, we now move on to semialgebraic objects. The results in this section will be expressed in the general context of an arbitrary real closed field $\mathbf{R}$, which can simply be thought of as a generalization of the usual field $\mathbb{R}$ of real numbers. More precisely, following from quantifier elimination results Tarski [Tar48], the theory of real closed field is complete, meaning that any first-order proposition that is true over $\mathbb{R}$ will also hold over any real closed field $\mathbf{R}$. The results of Tarski have later been built upon by Denef [Den86] and Pas [Pas89b] in the case of valued field, ultimately leading to the completeness of the theory of real closed fields with valuation (see [AGS20, Theorem 10]). Most of the contents of this short section are detailed in [BCR98, §2].

Definition 1.2.16. An ordered field $(\mathbf{R}, \leqslant)$ is called a *real closed field* whenever it satisfies the following two conditions:

(i) every positive element of $\mathbf{R}$ can be written as a square;

(ii) every odd degree univariate polynomial over $\mathbf{R}$ has at least one root.

Remark 1.2.17. Equivalently, $\mathbf{R}$ is a real closed field whenever $\mathbf{R}$ is not algebraically closed, but the extension $\mathbf{R}[\sqrt{-1}]$ is.

We now fix a real closed field $(\mathbf{R}, \leqslant)$ for the remainder of this section.

Definition 1.2.18. Let $f_1, \dots, f_k \in \mathbf{R}[X_1, \dots, X_n]$ be a family of k polynomials over $\mathbf{R}$. The *basic open semialgebraic set* associated to $f_1, \dots, f_k$ is the subset $\mathcal{U}(f_1, \dots, f_k)$ of $\mathbf{R}^n$ defined by

$$\mathcal{U}(f_1, \dots, f_k) := \{x \in \mathbf{R}^n : \forall i \in [k],\ f_i(x) > 0\}$$

and the *basic closed semialgebraic set* associated to $f_1, \dots, f_k$ is the subset $\overline{\mathcal{U}}(f_1, \dots, f_k)$ of $\mathbf{R}^n$ defined by

$$\overline{\mathcal{U}}(f_1, \dots, f_k) := \{x \in \mathbf{R}^n : \forall i \in [k],\ f_i(x) \geqslant 0\}\ .$$

Moreover, any set of the previous form is respectively called a *basic open* and *basic closed semialgebraic subset* of $\mathbf{R}^n$.

Basic semialgebraic sets are the most elementary examples of semialgebraic sets, from which all semialgebraic sets can be constructed. Notice also that $\overline{\mathcal{U}}(f_1, \dots, f_k) = \mathrm{cl}(\mathcal{U}(f_1, \dots, f_k))$ in the previous definition, where the closure is taken with respect to the order topology.

Definition 1.2.19. Any subset of $\mathbf{R}^n$ which can be obtained as a finite boolean combination of basic closed semialgebraic subsets of $\mathbf{R}^n$ is called a *semialgebraic subset* of $\mathbf{R}^n$.

Equivalently, semialgebraic sets constitute the smallest class of subsets of $\mathbf{R}^n$ containing algebraic subsets as well as open basic semialgebraic subsets of $\mathbf{R}^n$, and that is closed under the union, the intersection and the complement. More precisely, one has the following proposition.

Proposition 1.2.20 (see [BCR98, Propositions 2.1.3 and 2.1.8])**.** *Every semialgebraic subset of $\mathbf{R}^n$ can be written as a finite union of sets of the form $\mathcal{V}(f) \cap \mathcal{U}(g_1, \dots, g_k)$ where $f, g_1, \dots, g_k \in \mathbf{R}[X_1, \dots, X_n]$, possibly with $k = 0$.*

In the closed case, the previous proposition entails the following theorem.

Theorem 1.2.21 (Finiteness theorem)**.** *Every closed semialgebraic set can be written as a finite union of basic closed semialgebraic sets.*

Finally, semialgebraic sets are preserved under projection, as per the following theorem.

Theorem 1.2.22 (Projection theorem, see [BCR98, Theorem 2.2.1])**.** *For $n, p \in \mathbb{N}_{>0}$, let S be a semialgebraic subset of $\mathbf{R}^{n+p}$ and let $\pi : \mathbf{R}^{n+p} \simeq \mathbf{R}^n \times \mathbf{R}^p \to \mathbf{R}^n$ denote the first projection. Then $\pi(S)$ is a semialgebraic subset of $\mathbf{R}^n$. In other words, semialgebraic subsets are stable by projection.*

This projection property is useful for several purposes, one of which is describing semi-algebraic sets with their cylindrical algebraic decomposition. We do not go into more details, but for more information, the reader can for instance refer to [ACM84] where it was introduced, or more generally, to [BCR98, §2.3].

1.3 Notions of tropical geometry

Armed with the knowledge of the previous sections, we are now in a position to introduce the main subject of this thesis, that is tropical geometry and tropical polynomial systems. The canonical reference on the domain of tropical geometry is the eponymous book by Maclagan and Sturmfels [MS15]. Additionally, an involved reader may also want to check the lectures notes of Chambert-Loir's master course on tropical geometry [Cha21].

1.3.1 The tropical semifield $\mathbb{T}$

The reference algebraic structure used for calculations in the context of tropical algebra is the semiring, which generalises the notion of ring, allowing for the addition to be non-invertible. More precisely, the definition of a semiring is the following.

Definition 1.3.1. A *semiring* $(S, +, \cdot, 0, 1)$ is a set S equipped with addition $+$ and multiplication $\cdot$ such that

(i) $(S, +)$ is a commutative monoid with identity element 0 called the *zero element*;

(ii) $(S, \cdot)$ is a monoid with identity element 1 called the *unit element*;

(iii) the multiplication is distributive over the addition;

(iv) 0 is the absorbing element for the multiplication.

If moreover, the multiplication is commutative, then S is called a *commutative semiring*, and if every nonzero element has an inverse for the multiplication, then S is called a *semifield*.

A semiring S satisfies all the properties of a ring except for one: the existence of an opposite element for the addition is not guaranteed. In particular, a ring is a semiring. All the operations of a semiring can be expanded to tuples and matrices in order to define tuple and matrix addition, as well as matrix multiplication, allowing one to perform linear algebra over a semifield.

There are many interesting examples of semirings besides rings, such as the set $\mathbb{N}$ of nonnegative integers, endowed with the usual arithmetic operations, the sets of ideals of a ring, with ideal addition and multiplication, or any Boolean algebra, with $\vee$ as addition and $\wedge$ as multiplication. However, in the context of this work, we will be specifically focusing on the following semiring.

Definition 1.3.2. The *tropical* (or *max-plus*) *semiring* is the semiring given by $(\mathbb{T}, \oplus, \odot, \mathbb{0}, \mathbb{1})$, where

◇ the underlying set is $\mathbb{T} := \mathbb{R} \cup \{-\infty\}$;

◇ the *tropical addition* is $\oplus := \max$;

◇ the *tropical multiplication* is $\odot := +$;

◇ the *tropical zero element* is $\mathbb{0} := -\infty$;

◇ the *tropical unit element* is $\mathbb{1} := 0$.

Remark 1.3.3. More generally, one can likewise define a semiring structure for any partially ordered abelian group $(\Gamma, +, \leqslant)$, with a join-semilattice structure, and enriched with a bottom element $\perp$ satisfying $\perp \leqslant v$ and $\perp + v = v + \perp = \perp$ for all $v \in \Gamma$. These semirings share two common properties. The first one is that they are actually semifields, since the addition in $(\Gamma, +)$ is invertible. The second one is that the semifield addition $\oplus$ is by definition *idempotent*, meaning that $v \oplus v = v$ for all $v \in \Gamma \cup \{\perp\}$. Conversely, any such idempotent semifield can be endowed with an order giving it a structure of join-semilattice. Moreover, to anticipate Section 1.3.3, the case where the group $(\Gamma, +, \leqslant)$ is totally ordered coincides with the case where the associated semifield arises from a nonarchimedian valuated field.

The tropical context corresponds to the case $\Gamma = \mathbb{R}$, and in that case, the set of invertible elements for the tropical multiplication is $\mathbb{T}^* := \mathbb{R}$.

Remark 1.3.4. Some authors rather adopt the *min-plus* convention, and define the tropical semiring with the minimum as tropical addition, and $+\infty$ as the zero element. The preference for one of these two conventions over the other depends mainly on the mathematician's field of research, as well as personal taste. One may also occasionally encounter *max-times* or *min-times* semirings, for which the multiplication is taken to be the usual multiplication, but the underlying set is replaced respectively by $\mathbb{R}_{\geqslant 0}$ and $\mathbb{R}_{>0} \cup \{+\infty\}$. All these different models for tropical algebra are all isomorphic as you can go from one to the other via multiplication by -1 and the exponential function.

Polynomials can be defined over a commutative semiring S the same way as they are defined over a ring R, simply by replacing the ring R by the semiring S in Definition 1.2.1. For our concerns, we shall only consider polynomials defined over the tropical semiring $\mathbb{T}$. Such a (Laurent) polynomial over S will be called a *tropical (Laurent) polynomial*, and its associated polynomial function a *tropical polynomial function*.

If f is a tropical (Laurent) polynomial in n variables with support $\mathcal{A} \subseteq \mathbb{Z}^n$, then the tropical polynomial function associated to f is the function

$$\begin{array}{ccc} (\mathbb{T}^*)^n & \longrightarrow & \mathbb{T} \\ x & \longmapsto & \max_{\alpha \in \mathcal{A}} (f_\alpha + \langle x, \alpha \rangle) \end{array} ,$$

where $\langle x, \alpha \rangle$ denotes the usual scalar product of $x = (x_1, \dots, x_n)$ and $\alpha = (\alpha_1, \dots, \alpha_n)$ in $\mathbb{R}^n$.

Remark 1.3.5. While in the case of an infinite field, there is a one-to-one correspondance between polynomial functions and formal polynomials, this is not the case in the tropical setting, where two distinct formal tropical polynomial can have the same tropical polynomial function. Therefore, we will always be explicit about the nature of the tropical objects we manipulate. This distinction notably matters when considering questions related to factorization of tropical polynomials. Despite this absence of correspondance, there exists a way to choose a canonical representant in the class of tropical polynomial sharing the same given tropical polynomial function, which is obtained as the coefficient-wise maximum of all tropical polynomials in said class.

Now that the definition of polynomials and polynomial function has been properly established, the notion of root can be tackled. Usually, a *root* of a polynomial f over a ring R is an element x of R such that evaluating the polynomial function of f at the point x outputs the zero element. This definition, however, is not suited for polynomials over semirings, and in particular for tropical polynomials, due to the absence of an opposite operation for the semiring addition. For the tropical semifield, the definition of a root of a polynomial is adapted the following way.

Definition 1.3.6. Let $f = \bigoplus_{\alpha \in \mathcal{A}} f_\alpha X^\alpha \in \mathbb{T}[X_1^\pm, \dots, X_n^\pm]$ be a tropical Laurent polynomial with support $\mathcal{A} \subseteq \mathbb{Z}^n$. An element $x = (x_1, \dots, x_n)$ of $(\mathbb{T}^*)^n$ is called a (*tropical*) *root* or *zero* of f whenever the maximum in the expression

$$\max_{\alpha \in \mathcal{A}} (f_\alpha + \langle x, \alpha \rangle)$$

is attained for at least two distinct values of $\alpha \in \mathcal{A}$ or equal to $-\infty$. This is denoted as $f(x) \nabla \mathbb{0}$.

Moreover, if f is a tropical polynomial with support $\mathcal{A} \subseteq \mathbb{N}^n$, then a root x of f is defined to be an element $\mathbb{T}^n$ — instead of $(\mathbb{T}^*)^n$ — satisfying the same property, and in this case one defines the *support* of the root x as the set of indices $1 \leqslant j \leqslant n$ such that $x_j \neq \mathbb{0}$.

Remark 1.3.7. The notation $f(x) \nabla \mathbb{0}$ is not only used to convey the fact that it is the tropical analogue of the equation $f(x) = 0$ in a ring, but it also has a deeper meaning as it comes from a natural binary relation that arises in signed extensions of the tropical semifield. These *signed* or *symmetrised* semifields are described in greater details in [BCOQ92, §3.4] or [AGG09, §4.1]. This extensions are also described with a different terminology in [IR10].

There is a tropical zero-product property, as per the following proposition.

Proposition 1.3.8 (Tropical zero-product property)**.** *Let f and g be two univariate tropical polynomials. Then for all $x \in \mathbb{T}$*

$$(fg)(x) \nabla \mathbb{0} \quad \iff \quad f(x) \nabla \mathbb{0} \text{ or } g(x) \nabla \mathbb{0} \ .$$

Sketch of the proof. One simply need to remark that for all $x \in \mathbb{T}^n$,

$$\arg\max_{\gamma \in \mathbb{N}} \left(\underbrace{\max_{\substack{(\alpha,\beta) \in \mathcal{A} \times \mathcal{B} \\ \alpha + \beta = \gamma}} (f_\alpha + g_\beta)}_{:= (fg)_\gamma} + \langle x, \gamma \rangle \right) = \arg\max_{\alpha \in \mathcal{A}} (f_\alpha + \langle x, \alpha \rangle) + \arg\max_{\beta \in \mathcal{B}} (g_\beta + \langle x, \beta \rangle) \ ,$$

which then entails the result as $(fg)(x) = \max_{\gamma \in \mathbb{N}} (fg)_\gamma + \langle x, \gamma \rangle$. □

This definition of a tropical root leads to the expected behaviour for an root of a polynomial. In particular, over an algebraically closed field, the process of factorizing a univariate polynomial is equivalent to finding all its roots. The following proposition is the tropical analogue of this result. Again, keep in mind that the lack of one-to-one correspondance between formal tropical polynomials and tropical polynomial functions forces one to use polynomial functions in order to achieve a working factorisation property.

Proposition 1.3.9. *Let f be a univariate tropical polynomial. Then $x_0 \in \mathbb{T}$ is a root of f if and only if there exists a univariate tropical polynomial g such that $f(x) = (x \oplus x_0) \odot g(x)$ for all $x \in \mathbb{T}$, where the equality is to be understood as an pointwise equality of two tropical polynomial functions.*

In fact, there is a stronger result on univariate tropical polynomials, stating that the tropical semifield $\mathbb{T}$ is algebraically closed in a tropical sense. More precisely, one has the following.

Definition 1.3.10. Let $f = \bigoplus_{\alpha \in \mathcal{A}} f_\alpha X^\alpha$ be a univariate tropical polynomial with support $\mathcal{A}$ and let x be a tropical root of f. Then the *multiplicity* of the root x is the number k defined by

$$k := \max \left\{ \left|\alpha^{(1)} - \alpha^{(2)}\right| : \alpha^{(1)}, \alpha^{(2)} \in \arg\max\{f_\alpha + \langle x, \alpha \rangle : \alpha \in \mathcal{A}\} \right\} \ .$$

Theorem 1.3.11 (Cuninghame-Green, see [CM80, Theorems 8 and 11])**.** *Let f be a nonconstant univariate tropical polynomial. Then the tropical polynomial function associated to f can be uniquely factored in the following way:*

$$f(x) = c \odot (x \oplus x_1)^{\odot \mu_1} \odot \cdots \odot (x \oplus x_k)^{\odot \mu_k} \quad \forall x \in \mathbb{T} \ ,$$

where $c \in \mathbb{T}^*$ denotes the leading coefficient of f and $x_1, \ldots, x_k \in \mathbb{T}$ denote the roots of f, with respective multiplicity $\mu_1, \ldots, \mu_k \in \mathbb{N}_{>0}$.

1.3.2 Tropical prevarieties from a combinatorial point of view

Elaborating on the notion of tropical root detailed in the previous sections, we now define tropical hypersurfaces and tropical varieties, and give their general properties. Recall that for a polynomial $f = \bigoplus_{\alpha \in \mathcal{A}} f_\alpha X^\alpha$ and a point $x \in (\mathbb{T}^*)^n$, the notation $f(x) \ \nabla \ \mathbb{0}$ means that the maximum in the evaluation $\max_{\alpha \in \mathcal{A}}(f_\alpha + \langle x, \alpha \rangle)$ of the tropical polynomial function of f at point x is achieved for at least two distinct values of α or equal to $-\infty$.

Definition 1.3.12. The *tropical hypersurface associated to a tropical polynomial* $f \in \mathbb{T}[X_1, \ldots, X_n]$ is the subset $\mathcal{V}_{\mathsf{trop}}(f)$ of $\mathbb{T}^n$ defined by

$$\mathcal{V}_{\mathsf{trop}}(f) := \{x \in \mathbb{T}^n : f(x) \ \nabla \ \mathbb{0}\} \ .$$

In other words, it corresponds to the set of tropical roots of the tropical polynomial f. Any subset of $\mathbb{T}^n$ of this form is called a *tropical hypersurface* of $\mathbb{T}^n$.

Likewise the *tropical prevariety associated to a collection* $f_1, \ldots, f_k \in \mathbb{T}[X_1, \ldots, X_n]$ *of tropical polynomials* is the subset $\mathcal{V}_{\mathsf{trop}}(f_1, \ldots, f_k)$ of $\mathbb{T}^n$ defined by

$$\mathcal{V}_{\mathsf{trop}}(f_1, \ldots, f_k) := \{x \in \mathbb{T}^n : \forall i \in [k], \ f_i(x) \ \nabla \ \mathbb{0}\} \ .$$

It thus corresponds to the intersection of all the tropical hypersurfaces associated to $f_1, \ldots, f_k$. Any subset of $\mathbb{T}^n$ of this form is called a *tropical prevariety* of $\mathbb{T}^n$.

Whenever the context is clear enough, we may sometimes omit the 'trop' in the notation of the previous variety and simply denote it as $\mathcal{V}(f_1, \ldots, f_k)$.

Remark 1.3.13. Similarly to the classical case, given a tropical prevariety, one can look at its intersection with the tropical torus $(\mathbb{T}^*)^n = \mathbb{R}^n$, allowing one again to more generally describe the set of roots of a collection of tropical Laurent polynomials.

We remark the following immediate result.

Proposition 1.3.14. *Let $f_1, \ldots, f_k \in \mathbb{T}[X_1, \ldots, X_n]$ be a collection of tropical polynomials. Then*

$$\bigcap_{i=1}^{k} \mathcal{V}_{\mathsf{trop}}(f_i) = \mathcal{V}_{\mathsf{trop}}(f_1, \ldots, f_k) \ .$$

Proof. This follows immediately from Proposition 1.3.8. □

We now focus on some elementary topological properties of tropical hypersurfaces. Most of the following results can be found proven in [PR04]. Fix $f = \bigoplus_{\alpha \in \mathcal{A}} f_\alpha X^\alpha \in \mathbb{T}[X_1, \ldots, X_n]$ a tropical polynomial, with support $\mathcal{A}$ in the following results.

Proposition 1.3.15. *The set $\mathcal{V}_{\mathsf{trop}}(f) \cap \mathbb{R}^n$ coincides with the nondifferentiability locus of the tropical polynomial function associated to f over $\mathbb{R}^n$.*

Sketch of the proof. Wherever the maximum in the tropical polynomial function associated to f is achieved by a single monomial of exponent $\alpha \in \mathcal{A}$, it locally coincides with the affine function $x \mapsto f_\alpha + \langle x, \alpha \rangle$, and is thus differentiable. Hovewer, if $x \in \mathbb{R}^n$ is such that the maximum is achieved by at least two distinct monomials $\alpha_1, \alpha_2 \in \mathcal{A}$, then there exist generic infinitesimal perturbations of x such that only the monomial α_1 dominates, and at these points, the differential of f is $\langle \cdot, \alpha_1 \rangle$, and likewise there exist generic infinitesimal perturbations of x such that only the monomial α_2 dominates, and at these points, the differential of f is $\langle \cdot, \alpha_2 \rangle \neq \langle \cdot, \alpha_1 \rangle$, which proves that f is non-differentiable at x. ☐

Remark 1.3.16. We gave an elementary proof of the previous statement in order to understand the intuition behind the link between the non-differentiability of a tropical polynomial function and the roots of the associated tropical polynomial. However, the subdifferential of a maximum of linear functions can easily be expressed, allowing for a more 'high-tech' proof of this result (see [Roc70, Theorem 16.5]).

Corollary 1.3.17. *The set $\mathcal{V}_{\mathsf{trop}}(f) \cap \mathbb{R}^n$ is the support of a $(n-1)$-dimensional polyhedral complex of $\mathbb{R}^n$, denoted by $\mathscr{T}_f$.*

Proof. Proposition 1.3.15 entails that $\mathcal{V}_{\mathsf{trop}}(f) \cap \mathbb{R}^n$ coincides with the projection onto $\mathbb{R}^n$ of the $(n-1)$-skeleton of the epigraph of the tropical polynomial function associated to f, and by Proposition 1.1.13, the set of faces of dimension $n-1$ or less is a polyhedral complex, whose image by the projection onto $\mathbb{R}^n$ corresponds to $\mathscr{T}_f$. ☐

Proposition 1.3.18. *Consider $\mathbb{R}^n$ endowed with its standard topology. Then $\mathcal{V}_{\mathsf{trop}}(f) \cap \mathbb{R}^n$ is a closed subset of $\mathbb{R}^n$, and moreover any connected component of the complement of $\mathcal{V}_{\mathsf{trop}}(f) \cap \mathbb{R}^n$ coincides with the interior of a full-dimensionnal polyhedron of $\mathbb{R}^n$ whose faces belong to $\mathscr{T}_f$.*

Sketch of the proof. $\mathcal{V}_{\mathsf{trop}}(f) \cap \mathbb{R}^n$ is closed as a union of finitely many closed polyhedra (the projections of the $(n-1)$-dimensional faces of the epigraph of the tropical polynomial function associated to f). Moreover, the connected components of the complement of $\mathcal{V}_{\mathsf{trop}}(f) \cap \mathbb{R}^n$ coincide with the projection of the relative interior of the facets of the epigraph, hence the result. ☐

Remark 1.3.19. There exists in general an explicit description of all the polyhedral complexes of $\mathbb{R}^n$ that arise from a tropical hypersurface, which is given by the *balancing condition*. We refer the reader to [MR18b, §2.4] for more details.

It follows in particular from the previous proposition that the polyhedral complex $\mathscr{T}_f$ induces a polyhedral subdivision $\mathscr{C}_f$ of $\mathbb{R}^n$ obtained by adding to $\mathscr{T}_f$ the closure of all connected components of the complement of $\mathcal{V}_{\mathsf{trop}}(f) \cap \mathbb{R}^n$ as the n-dimensional cells. We refer to the polyhedral subdivision $\mathscr{C}_f$ as the *(primal) subdivision* of $\mathbb{R}^n$ associated to f. This decomposition satisfies the following property.

Proposition 1.3.20 (See [PR04]). *Set for a point $x \in \mathbb{R}^n$, $\mathcal{A}_x := \{\alpha \in \mathcal{A} : f(x) = f_\alpha + \langle x, \alpha \rangle\}$. Then $\mathcal{A}_x$ depends only on the relative interior of the cell $C \in \mathscr{C}_f$ containing x.*

Proposition 1.3.20 allows us to pose the following definition.

Definition 1.3.21. Let C be a cell of $\mathscr{C}_f$. Then the *dual cell* of C is the subset $C^\diamond$ of $\mathbb{R}^n$ defined by $C^\diamond := \operatorname{conv}(\mathcal{A}_x)$ for an arbitrary point x in the relative interior of C.

The combinatorial properties of the tropical hypersurface associated to a tropical polynomial f can be related to a subdivision of a certain polytope constructed from the exponents which appear in the support of f. It is more precisely defined as follows.

Definition 1.3.22. Let $f \in S[X_1, \dots, X_n]$ be a polynomial over a semiring S. Then the *Newton polytope* of f is the polytope NP_f defined by $\mathrm{NP}_f := \operatorname{conv}(\operatorname{supp}(f))$.

In other words, the Newton polytope of f is obtained by taking the convex hull of the finitely-many nonzero exponents $\alpha \in \mathbb{Z}^n$ for which the associated coefficient of f is nonzero — in the sense of the semiring S.

The Newton polytope of a polynomial $f \in S[X_1, \dots, X_n]$ can be defined for any semiring S. However, in the tropical case, one can construct a special lifting of the Newton polytope of f, enriching it with information on the value of the coefficients of the polynomials.

Definition 1.3.23. Let $f = \bigoplus_{\alpha \in \mathcal{A}} f_\alpha X^\alpha \in \mathbb{T}[X_1, \dots, X_n]$ be a tropical polynomial with support $\mathcal{A}$. Then the *extended* or *lifted Newton polytope* of f is the polyhedron $\mathrm{NP}_f^{\mathsf{lift}}$ of $\mathbb{R}^{n+1} \simeq \mathbb{R}^n \times \mathbb{R}$ defined as

$$\mathrm{NP}_f^{\mathsf{lift}} := \mathrm{conv}\{(\alpha, t) \in \mathbb{R}^n \times \mathbb{R} : \alpha \in \mathcal{A},\, t \leqslant f_\alpha\} \ .$$

In other words, it corresponds to the lower hull of the set of all (α, f_α) where α ranges over the support of f.

The lifted Newton polytope $\mathrm{NP}_f^{\mathsf{lift}} \subseteq \mathbb{R}^n \times \mathbb{R}$ yields a coherent subdivision $\mathscr{D}_f$ of the Newton polytope $\mathrm{NP}_f \subseteq \mathbb{R}^n$ obtained by projecting the upper faces of $\mathrm{NP}_f^{\mathsf{lift}}$ onto $\mathbb{R}^n$. The subdivision $\mathscr{D}_f$ is often refered to as the *dual subdivision* (of NP_f) associated to f, and its combinatorics encode the geometric properties of the tropical hypersurface associated to f. The subdivisions $\mathscr{C}_f$ and $\mathscr{D}_f$ are dual in the following sense.

Theorem 1.3.24 (See [PR04])**.** *The cells of the dual subdivision $\mathscr{D}_f$ coincide precisely with the dual cells of the primal subdivision $\mathscr{C}_f$,* i.e. *$\mathscr{D}_f = \{C^\diamond : C \in \mathscr{C}_f\}$, and the map*

$$\begin{array}{ccc} \mathscr{C}_f & \longrightarrow & \mathscr{D}_f \\ C & \longmapsto & C^\diamond \end{array}$$

is a poset anti-isomorphism.

Moreover, for any cell C of $\mathscr{C}_f$, the following hold:

(a) $\dim(C) + \dim(C^\diamond) = n$*;*

(b) C and $C^\diamond$ span orthogonal affine subspaces of $\mathbb{R}^n$;

(c) C is unbounded if and only if $C^\diamond$ lies on the boundary of the Newton polytope NP_f.

Given a collection $f_1, \dots, f_k \in \mathbb{T}[X_1, \dots, X_n]$, we shall often consider $Q := \mathrm{NP}_{f_1} + \cdots + \mathrm{NP}_{f_k}$ the Minkowski sum of the Newton polytopes of the $f_1, \dots, f_k$. Notice that Proposition 1.3.14 entails the equality $Q = \mathrm{NP}_{f_1 \cdots f_k}$. In fact, similarly the hypersurface case above, one can construct a subdivision of the Minkowski sum Q, which will combinatorially encode the geometric properties of the arrangement of tropical hypersurfaces given by the tropical polynomials $f_1, \dots, f_k$. Namely, the lifted Minkowski sum $Q^{\mathsf{lift}} := \mathrm{NP}_{f_1}^{\mathsf{lift}} + \cdots + \mathrm{NP}_{f_k}^{\mathsf{lift}} \subseteq \mathbb{R}^n \times \mathbb{R}$ yields a coherent subdivision of Q, obtained likewise by projecting the upper faces of Q^{lift} onto $\mathbb{R}^n$. This subdivision is then once again dual to the subdivision of $\mathbb{R}^n$ induced by the arrangement of tropical hypersurfaces $\mathcal{V}_{\mathsf{trop}}(f_1), \dots, \mathcal{V}_{\mathsf{trop}}(f_k)$ in the sense of Theorem 1.3.24. This is illustrated in the following figures.

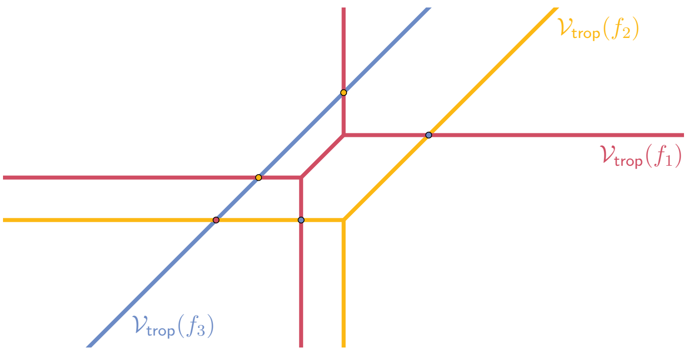

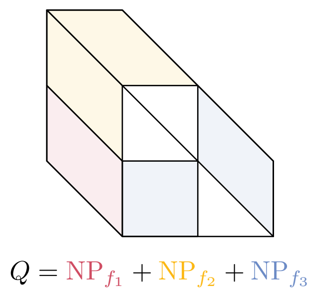

Figure 1.3: An arrangement of tropical hypersurfaces with the dual subdivision of the Minkowski sum Q of the associated Newton polytopes. The duality between the intersection points of the hypersurface arrangment and the mixed cells of the subdivision of Q is highlighted by the colouring.

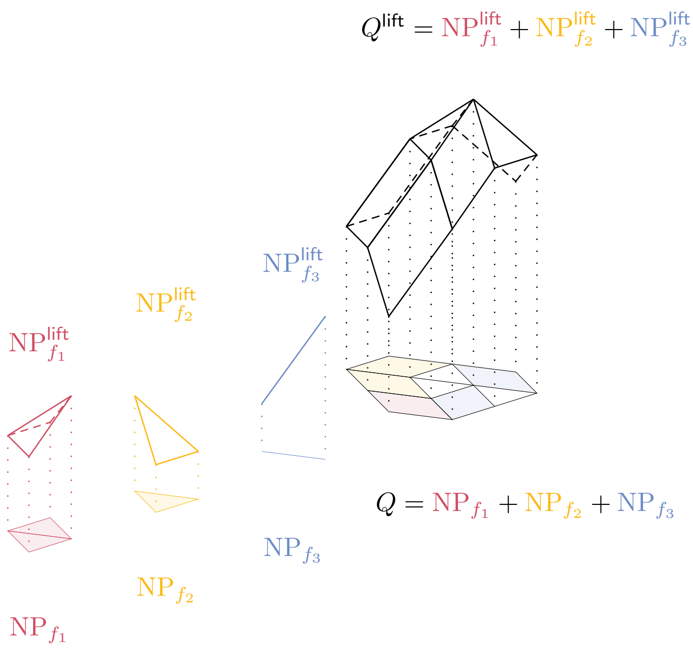

Figure 1.4: The dual subdivision of Q arises from the projection of the Minkowski sum Q^{lift} of the lifted Newton polytopes.

Similarly to the way tropical prevarieties have been defined, there exists a notion of tropical semialgebraic set. However, before properly defining tropical semialgebraic sets, we give a little bit of intuition through the following remark.

Remark 1.3.25. Let f be a real polynomial. Then by separating the positive and negative coefficients of f, one can write f as $f^+ - f^-$, where f^+ and f^- are two polynomials with only positive coefficients. This gives us a way to rewrite the equation $f(x) = 0$ without using any substraction, since writing $f(x) = 0$ is simply equivalent to writing the *two-sided* equality $f^+(x) = f^-(x)$. Similarly, one can rewrite the inequation $f(x) \geqslant 0$ as the two-sided inequation $f^+(x) \geqslant f^-(x)$, without ever using substractions, and likewise for strict inequalities.

While tropicalisation, via the image by a non-archimedian valuation as described above, forgets the sign of the coefficients, the previous rewriting of polynomial (in)equations into two-sided (in)equations give us an easy way to remember the information of the sign of each coefficient, by tropicalising the two-sided rewriting of the (in)equation instead.

Based on the previous remark, one can now define the notion of basic tropical semialgebraic set by analogy with the classical case.

Definition 1.3.26. A *basic tropical semialgebraic subset* of $\mathbb{R}^n$ is a subset of $\mathbb{R}^n$ that can be obtained as the set of solutions $x \in \mathbb{R}^n$ of a collection of two-sided polynomial (in)equations of the form

$$\begin{cases} f_1^+(x) & \rhd_1 & f_1^-(x) \\ & \vdots & \\ f_k^+(x) & \rhd_k & f_k^-(x) \end{cases}$$

where the $f_1^\pm, \dots, f_k^\pm$ are pairs of tropical polynomials, and $\rhd_1, \dots, \rhd_k \in \{=, \geqslant, >\}$, and a *tropical semialgebraic subset* of $\mathbb{R}^n$ is a subset of $\mathbb{R}^n$ that can be obtained as a finite boolean combination of basic tropical semialgebraic sets.

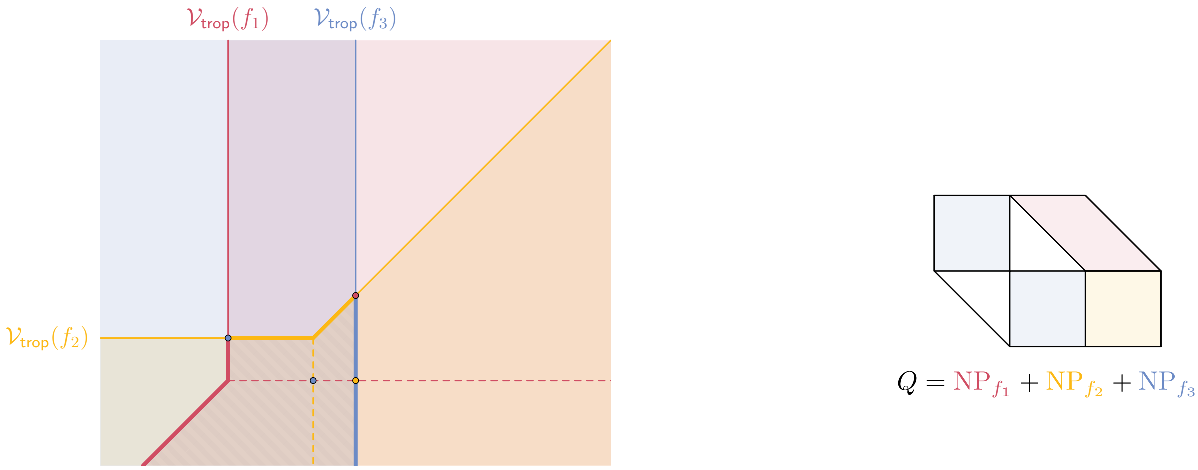

Figure 1.5: An arrangment of tropical semialgebraic sets given by three tropical polynomial weak inequalities of the form $f_i^+ \geqslant f_i^-$ represented next to the associated dual subdivision of the Minkowski sum of the Newton polytopes of the $f_i = f_i^+ \oplus f_i^-$. The boundary of the intersection of the three semi-algebraic sets is marked with hatching

Remark 1.3.27. For a single two-sided tropical polynomial inequality $f^+(x) \geqslant f^-(x)$, the set of solutions $x \in \mathbb{R}^n$ can simply be written as a union of maximal-dimensional cells of the subdivision $\mathscr{C}_f$ of $\mathbb{R}^n$ associated with the polynomial $f = f^+ \oplus f^-$. More precisely, the cells that appear in the union simply correspond by duality to the points $x \in \mathbb{R}^n$ such that the maximum in the expression $f(x) = \max_{\alpha \in \mathrm{supp}(f)} f_\alpha X^\alpha$ is achieved by a monomial of f^+

1.3.3 Valued fields

In this section, we briefly recall notions and vocabulary from the theory of non-archimedian valued fields, in order later on to explain how tropical objects naturally emerge from valued fields. For more details about the following results, or more generally about valued fields, one may refer to [EP05].

Throughout this section as well as the following one, we choose the convention to denote objects related to a valued field such as elements of a valued field or polynomials on a valued field in boldface, while keeping the regular weight font for their tropical counterpart.

Let $(\Gamma, +, \leqslant)$ be a totally ordered abelian group. Then the group law and ordering on Γ can be extended to the set $\Gamma \cup \{\bot\}$, where $\bot$ denotes the *bottom* element, by setting $\bot \leqslant v$ and $\bot + v = v + \bot = \bot$ for all $v \in \Gamma$. In $\Gamma \cup \{\bot\}$, the maximum function is defined as usual.

Definition 1.3.28. Let $\mathbf{K}$ be a field and Γ a totally ordered abelian group. A surjection val from $\mathbf{K}$ to $\Gamma \cup \{\bot\}$ is called a *valuation* if it satisfies the following properties for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbf{K}$:

(i) $\mathrm{val}(\boldsymbol{x}) = \bot$ if and only if $\boldsymbol{x} = \mathbf{0}$;

(ii) $\mathrm{val}(\boldsymbol{xy}) = \mathrm{val}(\boldsymbol{x}) + \mathrm{val}(\boldsymbol{y})$;

(iii) $\mathrm{val}(\boldsymbol{x} + \boldsymbol{y}) \leqslant \max(\mathrm{val}(\boldsymbol{x}), \mathrm{val}(\boldsymbol{y}))$.

A field endowed with a valuation is called a *valued field*.

Note that for number theorists working on valuation theory, it is usually more common to call valuation the opposite of the map val defined above. This would lead us to work which min-plus type, rather than max-plus type, tropical semifields. For our purposes, it is however more comfortable to work in the max-plus setting, hence the present choice of sign for the definition of the valuation.

It readily follows from items (i) and (ii) of the previous definition that val induces a group morphism from $(\mathbf{K}^*, \times)$ to $(\Gamma, +)$, where $\mathbf{K}^* := \mathbf{K} \setminus \{\mathbf{0}\}$ is the set of invertible elements of $\mathbf{K}$. This motivates the following definition.

Definition 1.3.29. The *value group* of a valued field $\mathbf{K}$ is the group $\Gamma := \mathrm{val}(\mathbf{K}^*)$.

Remark 1.3.30. In a lot of cases, we will be working with $\Gamma = \mathbb{R}$. In this case, we say that $\mathbf{K}$ is a field with *real valuation* or a *real valued* field and that val is a *real valuation*.

Example 1.3.31. A standard example of valued field is given by the field $\mathbb{C}\{\{t\}\}$ of complex univariate *Puiseux series*. This field is endowed with the valuation given by

$$\mathrm{val}(\boldsymbol{x}) := -\min\{q \in \mathbb{Q} : a_q \neq 0\} \quad \text{for} \quad \boldsymbol{x} = \sum_{q\in\mathbb{Q}} \boldsymbol{x}_q t^q \ ,$$

and in this case, the value group is $\Gamma = \mathbb{Q}$.

The former is in fact a subfield of the bigger field $\mathbb{C}[[t^{\mathbb{R}}]]$ of univariate complex *Hahn series*[1], that is the set of formal sums

$$\boldsymbol{x} = \sum_{r\in\mathbb{R}} \boldsymbol{x}_r t^r \quad \text{such that } \{r \in \mathbb{R} : \boldsymbol{x}_r \neq 0\} \text{ is a well ordered subset of } \mathbb{R} \ ,$$

in which case the quantity

$$\mathrm{val}(\boldsymbol{x}) := -\min\{r \in \mathbb{R} : \boldsymbol{x}_r \neq 0\}$$

is well-defined and yields a valuation of $\mathbb{C}[[t^{\mathbb{R}}]]$, which is surjective in this case, *i.e.* $\Gamma = \mathbb{R}$.

The following definition and proposition properly introduce the terminology of ‘non-archimedian field’, which is central in tropical geometry.

Definition 1.3.32. Let $\mathbb{K}$ be a real field and let $\mathbf{R}$ be a real closed field. A map $|\cdot| : \mathbf{K} \to \mathbf{R}$ is called an *absolute value* if it satisfies the following properties for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbf{K}$:

(i) $|\boldsymbol{x}| = 0$ if and only if $\boldsymbol{x} = \mathbf{0}$;

(ii) $|\boldsymbol{x}\boldsymbol{y}| = |\boldsymbol{x}|\,|\boldsymbol{y}|$;

(iii) $|\boldsymbol{x}+\boldsymbol{y}| \leqslant |\boldsymbol{x}| + |\boldsymbol{y}|$.

The inequality in (iii) is called the *triangular inequality*. Moreover, if the absolute value $|\cdot|$ satisfies the following stronger inequality for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbf{K}$

(iii’) $|\boldsymbol{x}+\boldsymbol{y}| \leqslant \max(|\boldsymbol{x}|, |\boldsymbol{y}|)$,

then it is said to be a *non-archimedian* absolute value, and the inequality in (iii’) is called the *ultrametric inequality*.

Remark 1.3.33. The term *non-archimedian* refers to the following property : if $\mathbf{K}$ is a field endowed with an absolute value, then it is non-archimedian if and only if the image of $\mathbb{Z} \cdot \mathbf{1} = \{n\mathbf{1} : n \in \mathbb{Z}\}$ by the absolute value is bounded. Otherwise, it is called *archimedian*.

The notions of absolute value and valuation are very closely related, as one can define a valuation from an absolute version and conversely as per the following proposition.

Proposition 1.3.34. *Let* $\mathbf{K}$ *be a field and for a map* $|\cdot| : \mathbf{K} \to \mathbb{R}$*, let* $\mathrm{val} : \mathbf{K} \to \mathbb{R} \cup \{-\infty\}$ *be the map defined by* $\mathrm{val}(\boldsymbol{x}) = \ln(|\boldsymbol{x}|)$*. Then* val *is a real valuation over* $\mathbf{K}$ *if and only if* $|\cdot|$ *is a non-archimedian absolute value.*

Proof. The proof of this result is straight-forward and rely on the fundamental property of the logarithm as well as the fact that it is an increasing function. □

Remark 1.3.35. The previous proposition was stated in the real case just for the sake of simplicity, but more generally, non-archimedian absolute values taking values in a real closed field are in correspondance with valuations whose value group is a totally ordered divisible group. Again, by completeness of the theory of real closed fields and divisible groups, the real case contains all the generality, since every first order proposition that is true over $\mathbb{R}$ also holds over any real closed field.

[1] Sometimes the notation $\mathbb{C}((\mathbb{R}))$ is also encountered, notably in [MS15].

We now fix $\mathbf{K}$ a valued field and recall some general knowledge. The *valuation ring* $\mathfrak{O}$ of $\mathbf{K}$ is the ring defined by

$$\mathfrak{O} := \{\boldsymbol{x} \in \mathbf{K} : \mathrm{val}(\boldsymbol{x}) \leqslant 0\} \ .$$

It is a local ring, with group of units

$$\mathfrak{O}^\times = \{\boldsymbol{x} \in \mathbf{K} : \mathrm{val}(\boldsymbol{x}) = 0\} \ ,$$

and its unique maximal ideal $\mathfrak{m}$ is given by

$$\mathfrak{m} = \{\boldsymbol{x} \in \mathbf{K} : \mathrm{val}(\boldsymbol{x}) < 0\} \ .$$

The residue field $\mathbf{k}$ of $\mathbf{K}$ is given by

$$\mathbf{k} := \mathfrak{O}/\mathfrak{m}$$

and we shall denote the projection from $\mathfrak{O}$ onto $\mathbf{k}$ by

$$\begin{array}{ccc} \mathfrak{O} & \longrightarrow & \mathbf{k} \\ \boldsymbol{x} & \longmapsto & \overline{\boldsymbol{x}} \end{array} \ .$$

Recall that a *splitting* of the surjection $\mathbf{K}^* \to \Gamma$ is a group morphism

$$\begin{array}{ccc} \Gamma & \longrightarrow & \mathbf{K}^* \\ v & \longmapsto & t^v \end{array} \ .$$

such that $\mathrm{val}(t^v) = v$ for all $v \in \Gamma$. Some authors also talk about *cross-sections*, which is the more standard terminology in the language of Denef-Pas, see [Pas89a].

The existence of such a splitting is guaranteed whenever $\mathbf{K}$ is algebraically closed, see [MS15, Lemma 2.1.15]. In that case, the projection $\mathfrak{O} \to \mathbf{k}$ induces a group morphism from $\mathfrak{O}^\times$ to $\mathbf{k}^*$ which can be extended into a group morphism π from $\mathbf{K}^*$ to $\mathbf{k}^*$ as follows

$$\pi : \left\{ \begin{array}{ccc} \mathbf{K}^* & \longrightarrow & \mathbf{k}^* \\ \boldsymbol{x} & \longmapsto & \overline{\boldsymbol{x}t^{-\mathrm{val}(\boldsymbol{x})}} \end{array} \right. ,$$

known as the *angular componant* in the language of Denef-Pas.

1.3.4 Tropical objects arising from valued fields

Let $\mathbf{K}$ be a field with valuation val. Recall that the hypersurface associated to a Laurent polynomial $\boldsymbol{f} = \sum_{\alpha \in \mathbb{Z}^n} \boldsymbol{f}_\alpha X^\alpha \in \mathbf{K}[X_1^{\pm 1}, \dots, X_n^{\pm 1}]$ is given the set

$$\mathcal{V}_{\mathbf{K}}(\boldsymbol{f}) := \{\boldsymbol{x} \in (\mathbf{K}^*)^n : \boldsymbol{f}(\boldsymbol{x}) = 0\} \ .$$

Similarly, the tropical hypersurface associated to a tropical Laurent polynomial $f \in \mathbb{T}[X_1^{\pm 1}, \dots, X_n^{\pm 1}]$ is given the set

$$\mathcal{V}_{\mathsf{trop}}(f) := \{x \in (\mathbb{T}^*)^n : f(x) \nabla \mathbb{0}\} \ ,$$

recalling that $\mathbb{T}^* = \mathbb{R}$. The link between the previous two objects is given by the following definition.

Definition 1.3.36. Let $\mathbf{K}$ be a field with valuation val, and consider a Laurent polynomial $\boldsymbol{f} = \sum_{\alpha \in \mathbb{Z}^n} \boldsymbol{f}_\alpha X^\alpha \in \mathbf{K}[X_1^{\pm 1}, \dots, X_n^{\pm 1}]$.

The *tropicalization* of $\boldsymbol{f}$ is the tropical polynomial function f defined by

$$f : \left\{ \begin{array}{ccc} \mathbb{R}^n & \longrightarrow & \mathbb{T} \\ x & \longmapsto & \max_{\alpha \in \mathbb{Z}^n}(\mathrm{val}(\boldsymbol{f}_\alpha) + \langle x, \alpha \rangle) \end{array} \right. .$$

We shall sometimes write $f = \mathrm{trop}(\boldsymbol{f})$ for short.

The *tropicalization* of $\mathcal{V}_{\mathbf{K}}(\boldsymbol{f})$ or the *tropical hypersurface* associated to the polynomial $\boldsymbol{f}$ is the set $\mathcal{T}_{\mathbf{K}}(\boldsymbol{f})$ defined by

$$\mathcal{T}_{\mathbf{K}}(\boldsymbol{f}) = \mathcal{V}_{\mathsf{trop}}(f) \ .$$

i.e. it corresponds to the set of points x in $\mathbb{R}^n$ such that the maximum in the expression $f(x)$ is achieved at least twice.

The previous definition can be expanded in order to define the tropicalization of any affine variety of $(\mathbf{K}^*)^n$, as follows.

Definition 1.3.37. Let $\mathfrak{a}$ be an ideal of $\mathbf{K}[X_1^{\pm 1}, \dots, X_n^{\pm 1}]$, and $\mathcal{V}_{\mathbf{K}}(\mathfrak{a})$ the variety it defines in $(\mathbf{K}^*)^n$. Then the *tropicalization* of the variety $\mathcal{V}_{\mathbf{K}}(\mathfrak{a})$, or the *tropical prevariety* associated to the ideal $\mathfrak{a}$ is the subset $\mathcal{T}_{\mathbf{K}}(\mathfrak{a})$ of $\mathbb{R}^n$ defined by

$$\mathcal{T}_{\mathbf{K}}(\mathfrak{a}) = \bigcap_{\boldsymbol{f} \in \mathfrak{a}} \mathcal{T}_{\mathbf{K}}(\boldsymbol{f}) \ . \tag{1.3}$$

More generally, we call *tropical variety* in $\mathbb{R}^n$ any subset of the previous form.

Remark 1.3.38. Note that the terminology of *tropical variety* being usually reserved to describe sets that can be achieved as the tropicalization of an affine variety over some non-archimedian field $\mathbf{K}$ is the reason why sets of the form $\mathcal{V}_{\mathsf{trop}}(f_1, \dots, f_k)$, where $f_1, \dots, f_k$ is a collection of tropical polynomials, are refered to as *tropical prevarieties* instead. These two notions do of course overlap, but not all tropical prevarieties can be obtained as the tropicalization of an affine variety. In particular, given any ideal $\mathfrak{a} = \langle \boldsymbol{f}_1, \dots, \boldsymbol{f}_k \rangle$, then one has

$$\mathcal{T}_{\mathbf{K}}(\mathfrak{a}) \subseteq \mathcal{V}_{\mathsf{trop}}(f_1, \dots, f_k)$$

where $f_i = \operatorname{trop}(\boldsymbol{f}_i)$ for all $i \in [k]$, but the latter inclusion may be strict, see for example [MS15, Example 2.6.7]. However, there is a notion of *tropical basis* in the following sense: there exists a finite collection of polynomials of $\mathfrak{a}$, called a *tropical basis* of $\mathfrak{a}$ such that the intersection on the righthandside of equality (1.3) can be restricted to this finite collection of polynomials (see [MS15, §2.6] for more details on tropical bases).

In this context, we can define initial forms and initial ideal, which will allow us to give another description of tropical hypersurfaces and varieties.

Definition 1.3.39. Let $\boldsymbol{f} = \sum_{\alpha \in \mathbb{Z}^n} \boldsymbol{f}_\alpha X^\alpha$ be a formal polynomial in $\mathbf{K}[X_1^{\pm 1}, \dots, X_n^{\pm 1}]$ and let $x \in \mathbb{R}^n$. Then the *initial form* of $\boldsymbol{f}$ with respect to x is the formal polynomial $\mathrm{in}_x(\boldsymbol{f}) \in \mathbf{k}[X_1^{\pm 1}, \dots, X_n^{\pm 1}]$ given by

$$\mathrm{in}_x(\boldsymbol{f}) = \sum_\beta \pi(\boldsymbol{f}_\beta) X^\beta \quad \text{where } \beta \text{ runs over } \operatorname*{arg\,max}_{\alpha \in \mathbb{Z}^n} \left(\mathrm{val}(\boldsymbol{f}_\alpha) + \langle x, \alpha \rangle \right) \ .$$

Moreover, if $\mathfrak{a}$ is an ideal of $\mathbf{K}[X_1^{\pm 1}, \dots, X_n^{\pm 1}]$, then the *initial ideal* of $\mathfrak{a}$ with respect to x is the ideal $\mathrm{in}_x(\mathfrak{a})$ generated by all the initial forms $\mathrm{in}_x(\boldsymbol{f})$ for $\boldsymbol{f}$ in $\mathfrak{a}$.

Remark 1.3.40. With the previous definition, the tropical variety $\mathcal{T}_{\mathbf{K}}(\boldsymbol{f})$ also corresponds to the set of points $x \in \mathbb{R}^n$ such that the initial form $\mathrm{in}_x(\boldsymbol{f})$ is not a monomial.

In the case where $\mathbf{K}$ has a nontrivial valuation, then there is an equivalent way to express $\mathcal{T}_{\mathbf{K}}(\mathfrak{a})$ which is given by the Fundamental theorem of Tropical Algebraic Geometry ([MS15, Theorem 3.2.3]) below, which generzlizes the Kapranov theorem from tropical hypersurfaces to tropical varieties altogether.

Theorem 1.3.41 (Fundamental theorem of Tropical Algebraic Geometry)**.** *Let $\mathbf{K}$ be an algebraically closed field endowed with a nontrivial valuation* val *and let $\mathfrak{a}$ be an ideal in $\mathbf{K}[X_1^{\pm 1}, \dots, X_n^{\pm 1}]$. Then the following three subsets of $\mathbb{R}^n$ coincide:*

(i) the tropical variety $\mathcal{T}_{\mathbf{K}}(\mathfrak{a})$;

(ii) the set of all vectors $x \in \mathbb{R}^n$ such that $\mathrm{in}_x(\mathfrak{a}) \neq \langle 1 \rangle$;

(iii) the closure of the set of coordinatewise valuations of points of $\mathcal{V}_{\mathbf{K}}(\mathfrak{a})$, i.e. of the set

$$\mathrm{val}(\mathcal{V}_{\mathbf{K}}(\mathfrak{a})) = \{ (\mathrm{val}(\boldsymbol{x}_1), \dots, \mathrm{val}(\boldsymbol{x}_n)) \in \mathbb{R}^n : (\boldsymbol{x}_1, \dots, \boldsymbol{x}_n) \in \mathcal{V}_{\mathbf{K}}(\mathfrak{a}) \} \ .$$

Remark 1.3.42. Note that in item (ii) of the previous theorem, the condition $\mathrm{in}_x(\mathfrak{a}) \neq \langle 1 \rangle$ can be replaced with the equivalent condition that the ideal $\mathrm{in}_x(\mathfrak{a})$ contains no monomial, since monomials are exactly the invertible elements of the ring of Laurent polynomials in n variables.

Remark 1.3.43. In particular, notice that in the case where the ideal $\mathfrak{a}$ is principal and generated by a polynomial $\boldsymbol{g} \in \mathbf{K}[X_1^{\pm 1}, \dots, X_n^{\pm 1}]$, then the set $\mathcal{T}_{\mathbf{K}}(\mathfrak{a})$ simply corresponds to the tropical hypersurface $\mathcal{T}_{\mathbf{K}}(\boldsymbol{g})$.

Theorem 1.3.41 motivates the present work. Indeed, for any finite family of polynomials $\boldsymbol{f}_1, \dots, \boldsymbol{f}_k$ in the ideal $\mathfrak{a}$, checking that $\bigcap_{j\in[k]} \mathcal{T}_{\mathbf{K}}(\boldsymbol{f}_k) = \emptyset$, which can be done by applying the tropical Nullstellensatz presented in Section 2.1, entails that $\operatorname{val}(\mathcal{V}_{\mathbf{K}}(\mathfrak{a})) = \emptyset$. Hence, the tropical Nullstellensatz provides a certificate of emptyness for an algebraic variety over a valued field. Moreover, thanks to the existence of tropical bases of $\mathcal{T}_{\mathbf{K}}(\mathfrak{a})$, the intersection (1.3) is achieved by considering only a subintersection over a finite set, which entails that the collection of certificates obtained in this way is complete.

1.4 Solving classical sparse polynomial systems

In this section, we briefly recall some of the notions and vocabulary from classical elimination theory, and in particular resultant theory. This language will in particular help to put in context the results of Chapter 2. We fix in the remainder of this section, $n, k \in \mathbb{N}_{>0}$ two strictly positive integers, as well as $\mathbb{K}$ an algebraically closed 0-characteristic field.

1.4.1 The Macaulay matrix and sparse resultant theory

Given a n-variate polynomial over $\mathbb{K}$, we have previously defined its support, as the subset of $\mathbb{N}^n$, or $\mathbb{Z}^n$ in the case of Laurent polynomials, of n-tuples corresponding to the exponents appearing in the polynomial. Our key concern regarding polynomials, is being able to solve polynomial systems, that is computing the set $\mathcal{V}_{\mathbb{K}}(f_1, \dots, f_k)$ for a collection $f_1, \dots, f_k$ of polynomials. In particular, whenever the supports $\mathcal{A}_1, \dots, \mathcal{A}_k$ of the polynomials $f_1, \dots, f_k$ are known in advance — and have a relatively 'small' cardinality in some sense — then one can hope to obtain some information on the behaviour of the variety $\mathcal{V}_{\mathbb{K}}(f_1, \dots, f_k)$. According to the usual terminology, we shall refer to such polynomial systems where the supports $\mathcal{A}_1, \dots, \mathcal{A}_k$ of the polynomials are prescribed, as *sparse polynomial systems*. More precisely, in the 'square' case, that is whenever $k = n + 1$, then $\mathcal{V}_{\mathbb{K}}(f_1, \dots, f_k)$ is generically a finite set, and in that case one can hope to obtain a bound on the cardinality of $\mathcal{V}_{\mathbb{K}}(f_1, \dots, f_k)$. A first bound, attributed by Bézout, is given by the product of the degrees of the $f_1, \dots, f_k$. However, depending on the supports $\mathcal{A}_1, \dots, \mathcal{A}_k$, one may sometimes find tighter bounds for structured (sparse) polynomial systems. Famously, the *BKK bound*, where the acronym BKK stands for Bernstein-Khovanskii-Kushnirenko, links the number of solutions of a sparse square polynomial system to some geometric property — the mixed volume — of the Minkowski sum of the Newton polytopes of the polynomials $f_1, \dots, f_{n+1}$. We do not give any further details in this manuscript, however, one can for instance refer to [Emi05] for a deeper understanding of this topic. Another central reference on the topic of resultant theory is [EM07].

We are now ready to introduce some standard vocabulary from classical elimination theory. For our concern, we will only take interest in the toric case, meaning that we are looking for solutions of polynomial systems over the torus $(\mathbb{K}^*)^n$.

Definition 1.4.1. Let $\mathcal{A} = (\mathcal{A}_1, \dots, \mathcal{A}_k)$ be a collection of finite subsets of $\mathbb{Z}^n$. Then the *incidence variety* associated to the collection of supports $\mathcal{A}$ (over the field $\mathbb{K}$) is the set $\mathcal{W}_{\mathbb{K}}(\mathcal{A})$ defined by

$$\mathcal{W}_{\mathbb{K}}(\mathcal{A}) := \left\{(f, x) \in \left((\mathbb{K}^*)^{\mathcal{A}_1} \times \cdots \times (\mathbb{K}^*)^{\mathcal{A}_k}\right) \times (\mathbb{K}^*)^n : \forall i \in [k],\, f_i(x) = 0\right\} \enspace .$$

Moreover, let $\mathcal{Z}_{\mathbb{K}}(\mathcal{A})$ be the projection of $\mathcal{W}_{\mathbb{K}}(\mathcal{A})$ onto the first factor of the cartesian product, that is

$$\mathcal{Z}_{\mathbb{K}}(\mathcal{A}) := \left\{f \in (\mathbb{K}^*)^{\mathcal{A}_1} \times \cdots \times (\mathbb{K}^*)^{\mathcal{A}_k} : \exists x \in (K^*)^n,\, \forall i \in [k],\, f_i(x) = 0\right\} \enspace .$$

Then, the *resultant variety* $\mathcal{RV}_{\mathbb{K}}(\mathcal{A})$ associated to $\mathcal{A}$ (over the field $\mathbb{K}$) is defined as the Zariski closure of $\mathcal{Z}_{\mathbb{K}}(\mathcal{A})$.

In the square case ($k = n + 1$) and under some combinatorial condition on the collection $\mathcal{A}$ — the supports $\mathcal{A}_1, \dots, \mathcal{A}_k$ must be *essential* in the sense of [Stu94] — then the resultant variety $\mathcal{RV}_{\mathbb{K}}(\mathcal{A})$ can be described as the hypersurface associated to a single irreducible polynomial $\mathcal{R} \in \mathbb{K}[f]$. This polynomial is called the *resultant polynomial* associated to $\mathcal{A}$, or simply *resultant* for short whenever the context is unambiguous.

The resultant polynomial plays a central role in polynomial system solving, as given a collection of polynomials $f_1, \dots, f_{n+1}$ of n-variate polynomials over $\mathbb{K}$, with respective support $\mathcal{A}_1, \dots, \mathcal{A}_{n+1}$, then generically, the choices of coefficients $f_{i,\alpha} \in \mathcal{A}_i$ for $i \in [n + 1]$ such that the polynomial system $f_1(x) = \cdots = f_k(x) = 0$ has a solution $x \in (\mathbb{K}^*)^n$ correspond to the points of the resultant variety. Therefore, being able to compute the resultant variety and the resultant polynomial is a crucial goal. In some cases, the resultant polynomial can be expressed as a quotient of minors of the following matrix, which constitutes a mutlivariate generalization of the Sylvester matrix of two univariate polynomials.

Definition 1.4.2. Let $f_1, \dots, f_k$ be a collection of n-variate polynomials over $\mathbb{K}$. Then the *Macaulay matrix* associated to $f_1, \dots, f_k$ is the infinite matrix $\mathcal{M}$ defined as such: the rows of $\mathcal{M}$ are indexed by pairs $(i, \alpha) \in [k] \times \mathbb{Z}^n$, the columns of $\mathcal{M}$ are indexed by exponents $\beta \in \mathbb{Z}^n$, and for given (i, α) and β, the entry $\mathcal{M}_{(i,\alpha),\beta}$ of $\mathcal{M}$ is set to be the coefficient of the monomial X^β in the polynomial $X^\alpha f_i(X)$ — or 0 if no such monomial exists.

Despite the matrix $\mathcal{M}$ having infinitely many rows and colums, one can define finite submatrices of $\mathcal{M}$ in the following way.

Definition 1.4.3. Let $\mathcal{E}$ be a *nonempty* finite subset of $\mathbb{Z}^n$, and let $\mathcal{A} = (\mathcal{A}_1, \dots, \mathcal{A}_k)$ be a collection of subsets of $\mathbb{Z}^n$. Then $\mathcal{M}_\mathcal{E}^\mathcal{A}$ denotes the submatrix of $\mathcal{M}$ obtained by keeping only the columns indexed by exponents $\beta \in \mathcal{E}$, and the rows indexed by pairs $(i, \alpha) \in [k] \times \mathbb{Z}^n$ such that $\alpha + \mathcal{A}_i \subseteq \mathcal{E}$. Additionally, whenever $\mathcal{A}_i = \operatorname{supp}(f_i)$ for all $i \in [k]$, then we simply write $\mathcal{M}_\mathcal{E}$ instead of $\mathcal{M}_\mathcal{E}^\mathcal{A}$.

Remark 1.4.4. Equivalently, $\mathcal{M}_\mathcal{E}$ is the submatrix of $\mathcal{M}$ obtaing by keeping only columns with indices $\beta \in \mathcal{E}$, and the rows that have all their finite entries in these columns. Moreover, if $\mathcal{E}$ is nonempty but too small, it might be possible that there are no such row of $\mathcal{M}$, and thus the set of rows of $\mathcal{M}_\mathcal{E}^\mathcal{A}$ might be empty, so we have to ensure that this does not happen by always choosing a suitable subset $\mathcal{E}$ of $\mathbb{Z}^n$. We shall always implicitly make this assumption unless stated otherwise.

Given a collection $f_1, \dots, f_{n+1}$ of $n+1$ polynomials in n variables, there exist different constructions of suitable subsets $\mathcal{E} \subseteq \mathbb{Z}^n$ such that the resultant of $f_1, \dots, f_{n+1}$ can be expressed as a quotient of two minors of the matrix $\mathcal{M}_\mathcal{E}$ (see [Emi05, EM07] for more details about this construction). In that case, we say that the resultant has a *Sylvester-type formula*. Are Sylvester-type formulae essential in order to explicitely be able to compute the resultant, and therefore proving the existence of such formulae is a crucial question in elimination theory.

We finally give a few straightforward properties of the Macaulay matrix, in order to get some intuition as to why this matrix naturally appears in resultant theory. Fix in the following a nonempty finite subset $\mathcal{E}$ of $\mathbb{Z}^n$.

Definition 1.4.5. The *Veronese embedding* from $\mathbb{K}^n$ to $\mathbb{K}^\mathcal{E}$ refers to the following map

$$\mathrm{ver} : \left\{ \begin{array}{rcl} \mathbb{K}^n & \longrightarrow & \mathbb{K}^\mathcal{E} \\ x & \longmapsto & (x^p)_{p \in \mathcal{E}} \end{array} \right. .$$

The following two propositions follow directly from the very definition of the Macaulay matrix.

Proposition 1.4.6. *If $x \in (\mathbb{K}^*)^n$, then $\mathcal{M}_\mathcal{E} \cdot \mathrm{ver}(x) = (x^\alpha f_i(x))_{(i,\alpha) \in [k] \times \mathcal{A}_i}$.*

Proposition 1.4.7. *For all $i \in [k]$, let $\mathcal{E}_i := \{\alpha \in \mathbb{Z}^n : \alpha + \mathcal{A}_i \subseteq \mathcal{E}\}$. Then the matrix of the map*

$$\left\{ \begin{array}{rcl} (\mathbb{K}^*)^{\mathcal{E}_1} \times \cdots \times (\mathbb{K}^*)^{\mathcal{E}_k} & \longrightarrow & (\mathbb{K}^*)^\mathcal{E} \\ (g_1, \dots, g_k) & \longmapsto & g_1 f_1 + \cdots + g_k f_k \end{array} \right. ,$$

where for all $i \in [k]$, $g_i \in (\mathbb{K}^*)^{\mathcal{E}_i}$ is assimilated to the polynomial $g_i = \sum_{\alpha \in \mathcal{E}_i} g_{i,\alpha} X^\alpha$ of support $\mathcal{E}_i$, is given by the transpose of $\mathcal{M}_\mathcal{E}$.

1.4.2 The effective Nullstellensatz and the Macaulay bound

A *Nullstellensatz*, standing in German for ‘zero locus theorem’ is a statement relating the existence of a root of a polynomial system to some algebraic properties of the ideal generated by the polynomials of the system. There are different ways to state such results. The standard formulation of the Nullstellensatz, stated in 1893 by Hilbert in [Hil93], relies on the notion of *radical* of an ideal, and state that the vanishing ideal of the variety associated to an ideal $\mathfrak{a}$ coincides with the radical of $\mathfrak{a}$. However, in the context of this work, we focus more specifically on *effective* formulations of the Nullstellensatz, that is formulations allowing one to effectively construct a certificate for the existence or non-existence of a root of a polynomial system (see *e.g.* [Bro87, Kol88]).

Theorem 1.4.8 (Weak effective Nullstellensatz)**.** *Let $f_1, \dots, f_k \in \mathbb{K}[X_1, \dots, X_n]$ be a collection of n-variate polynomials over $\mathbb{K}$. Then, one has*

$$\mathcal{V}(f_1, \dots, f_k) = \emptyset \quad \iff \quad \exists g_1, \dots, g_k \in \mathbb{K}[X_1, \dots, X_n] \text{ such that } g_1 f_1 + \cdots + g_k f_k = 1 \ .$$

In this effective version, the certificate is given by the polynomials $g_1, \dots, g_k \in \mathbb{K}[X_1, \dots, X_n]$. A natural question then arises: how 'small' can this certificate be. In other words, if the degrees of the polynomials $f_1, \dots, f_k$ are bounded, then one looks for the smallest *degree bound* δ such that one can find such certificates $g_1, \dots, g_k$, each with degree less than or equal to δ. A first step towards this direction is given by the following statement, which states that such a degree bound, depending only on the number n of variables, the number k of polynomials, and the maximal degree d of the $f_1, \dots, f_k$, but uniform in the coefficients of the polynomials, exists.

Theorem 1.4.9 (Effective Nullstellensatz)**.** *Let $f_1, \dots, f_k \in \mathbb{K}[X_1, \dots, X_n]$ be a collection of n-variate polynomials over $\mathbb{K}$, and let $d = \max\{\deg(f_i) : i \in [k]\}$. Then there exists a uniform degree bound $\delta(n, k, d)$ such that*

$$\mathcal{V}(f_1, \dots, f_k) = \emptyset \quad \iff \quad \exists g_1, \dots, g_k \in \mathbb{K}[X_1, \dots, X_n] \ \textit{such that} \ \begin{cases} g_1 f_1 + \cdots + g_k f_k = 1 \\ \forall i \in [k],\ \deg(g_i) \leqslant \delta \ . \end{cases}$$

The idea of the previous effective Nullstellensatz is therefore to reduce the problem of determining the solvability of a polynomial system, to simply computing the (left) kernel of a submatrix $\mathcal{M}_{\mathcal{E}}$ of the Macaulay matrix, obtain by taking $\mathcal{E}$ to be the set of monomials with degree less than or equal to the degree bound δ. This degree bound can be estimated in function of n, k and d. In particular, one has the following result, proven by Lazard in 1983.

Theorem 1.4.10 (Macaulay bound, see [Laz83])**.** *For a generic choice of the coefficients of the polynomials $f_1, \dots, f_k$, one can choose $\delta(n, k, d) = 1 + \sum_{i=1}^{k} (\deg(f_i) - 1)$ in the previous theorem.*

The previous value is refered to as the Macaulay bound. Note that the Macaulay bound only gives a suitable degree bound for generic systems, but can however fail for generic instances, as in the following example.

Example 1.4.11 (Masser-Phillipon example, see [Bro87, §1]). Let $f_1, \dots, f_n \in \mathbb{K}[X_1, \dots, X_n]$ given by

$$f_1 = X_1^d, f_2 = X_1 - X_2^d, \dots, f_{n-1} = X_{n-1} - X_n^d \quad \text{and} \quad f_n = 1 - X_{n-1} X_n^{d-1} \ .$$

Then the equation $f_1(x) = \cdots = f_n(x) = 0$ does not have a solution, as the first $n-1$ equations would otherwise imply that $x_1 = \cdots = x_n = 0$, hence contradicting the equality $f_n(x) = 0$. Therefore, by the previous Nullstellensatz, there exists polynomials $g_1, \dots, g_n \in \mathbb{K}[X_1, \dots, X_n]$ such that $g_1 f_1 + \cdots + g_n f_n = 1$.

However, by evaluating the previous equality at the point $x(t) = (t^{d^{n-1}(d-1)}, t^{d^{n-2}(d-1)}, \dots, t^{d-1}, \frac{1}{t})$ for $t \neq 0$ yields that $g_n(x(t)) = 1$, and therefore the degree of g_n in the variable x_n must be at least equal to $d^{n-1}(d-1)$ in order for the term $x_n(t) = \frac{1}{t}$ to cancel out all other terms in the expression $a_n(x(t))$. This bound being exponential in the number n of variables, overpasses the classical Macaulay bound, which would be otherwise equal to $1 + n(d-1)$.

1.5 Tropical linear systems and mean payoff games

This section presents some generalities on a particular class of zero-sum games, namely the so called mean payoff games, as well as a summary of the link between tropical linear systems and mean payoff games, through the theory of nonlinear eigen values. In particular, it describes how the solvability of tropical linear inequalities over $\mathbb{R}$ can be reduced to computing the value of a mean payoff game. This correspondance is explained in fuller details in [AGG12, Section 2], which the main results of this sections are drawn from.

1.5.1 Generalities on mean payoff games

Zero-sum games constitute a particular class of two-player games in which each gain for one of the players result in an equivalent loss for the second player. In a sense, no wealth is created in these kind of games, as at any time, the sum of gains (counted positively) and losses (counted negatively) of both players cancels out exactly. This specific characteristic gives zero-sum games some very interesting properties, among which the existence of *minimax theorems* or *Nash equilibria*.

In this section and more generally all throughout this manuscript, we will take a particular interest in a specific subclass of zero-sum games, called mean payoff games. We refer the reader to [EM79, ZP96] for more detailed information on mean payoff games.

Definition 1.5.1. Let G be a (finite) oriented weighted bipartite graph, given with its set of vertices $I \sqcup J$ and its set of arcs $E \subseteq (I \times J) \cup (J \times I)$. The vertices of G are refered to as the *states* or *positions* of the game, and the arcs of G are refered to as the *actions* or *moves*.

The *mean payoff game* associated to the graph G is the zero-sum two-player game defined as follows. The first player is called the *minimizer* and the second one the *maximizer*. Each turn, the minimizer, from a state $j_0 \in J$, chooses the next state $i_0 \in I$ along an arc (j_0, i_0) with weight $-a_{i_0 j_0}$, and receives in turn a payment of $a_{i_0 j_0}$ from the maximizer — or equivalently, in order for all the payments to always go from the minimizer to the maximizer, one may say that the maximizer receives a payment of $-a_{i_0 j_0}$ from the minimizer. Then the maximizer, from the current state $i_0 \in I$, chooses a state $j_1 \in J$ such that (i_0, j_1) is an arc of G with weight $b_{i_0 j_1}$, and receives a payment of $b_{i_0 j_1}$ from the minimizer. These steps repeat indefinitely, and the winner is the player who manages to ensure the highest average payment per turn.

For a mean payoff game as described above, we denote by $A = (a_{ij})_{(i,j) \in I \times J}$ and $B = (b_{ij})_{(i,j) \in I \times J}$ the *payment matrices* associated to the game, where we set $a_{ij} = -\infty$ whenever the arc (j, i) does not exist in the graph G and likewise $b_{ij} = -\infty$ whenever the arc (i, j) does not exist. Therefore, the payment matrices A and B have entries in $\mathbb{R} \cup \{-\infty\}$ and as such they can be considered as $I \times J$ tropical matrices over the semiring $\mathbb{T}$. Since the graph G is uniquely described by A and B, we shall denote this graph as $\mathrm{G}(A, B)$.

A variant of the above game, called the *finite horizon game*, consists in stopping the game after a fixed number $N \in \mathbb{N}$ of steps. In that case, the winning player is simply the one who got the bigger total payment. Since this is a zero-sum game, it means equivalently that the winner is the beneficiary player, with a positive net gain at the end of the game, while the loser is the deficit player, with a negative net gain.

Example 1.5.2. Figure 1.6 below displays graph of the mean payoff game whose corresponding sets of states are

$$I = \left\{\boxed{1}, \boxed{2}, \boxed{3}\right\} \quad \text{and} \quad J = \left\{1, 2\right\}$$

and whose corresponding payment matrices are

$$A = \begin{array}{cc} & \begin{array}{cc} 1 & 2 \end{array} \\ \begin{array}{c} \boxed{1} \\ \boxed{2} \\ \boxed{3} \end{array} & \begin{pmatrix} 2 & -\infty \\ 8 & -15 \\ -\infty & 0 \end{pmatrix} \end{array} \quad \text{and} \quad B = \begin{array}{cc} & \begin{array}{cc} 1 & 2 \end{array} \\ \begin{array}{c} \boxed{1} \\ \boxed{2} \\ \boxed{3} \end{array} & \begin{pmatrix} 1 & -\infty \\ 10 & -\infty \\ -\infty & 5 \end{pmatrix} \end{array}.$$

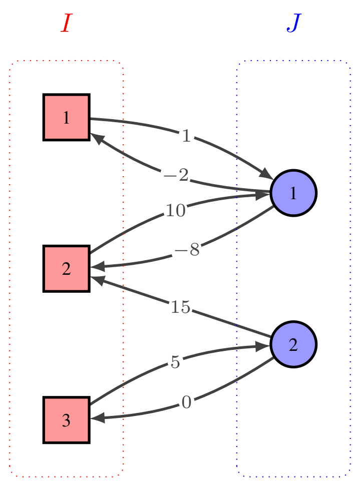

Figure 1.6: The graph $\mathrm{G}(A, B)$ of a mean payoff game.

We shall keep for the reminder of this section the notation of the previous definition. The mean payoff game described above belongs to the family of *turn-based games*, which consists in all the games where the two players alternate their moves, playing non-simultaneously, and each time with full knowledge of the last move played by the opponent. These games thus constitute a subclass of *perfect information games*. Such a game, where the initial state $j_0 \in J$ as well as the number of turns N are prescribed in advance and known by both players, admits a *Nash* *equilibrium*. More specifically, considering the finite-horizon game, this means that there exists a saddle point in the product of spaces of strategies for both the minimizer and the maximizer, *i.e.* there exists an optimal strategy τ^* for the minimizer player and an optimal strategy σ^* for the maximizer player such that after N turns, the profit of each player is optimized if they both follow their respective optimal strategy, in such a way that neither player has no interest whatsoever in deviating from these strategies. We formalize the previous intuition with the following definitions and results.

Definition 1.5.3. The *history* of the game at a given time is the set of all moves that have been chosen by both players so far. A *strategy* for either player is a map which takes the whole history of the game at the given time as an input, and returns a move to play. The set of strategies for the minimizer and the maximizer are denoted by Σ_{min} and Σ_{max} respectively.

Definition 1.5.4. A strategy which depends solely on the current state is called a *positional strategy*. In other words, a positional strategy for the minimizer player is a map $\sigma : J \to I$ such that $b_{\sigma(j)j} > -\infty$ for every state $j \in J$, and a positional strategy for the minimizer is a map $\tau : I \to J$ such that $b_{i\tau(i)} > -\infty$ for every state $i \in J$. The set of positional strategies for the minimizer and the maximizer are denoted by Π_{min} and Π_{max} respectively.

Definition 1.5.5. For $\sigma \in \Sigma_{\mathsf{min}}$, $\tau \in \Sigma_{\mathsf{min}}$, and $j \in J$, the respective profit of the minimizer and the maximizer players after N turns, where the minimizer has followed the strategy τ and the maximizer has followed the strategy σ, and where the initial state of the game was j, is denoted by $G^N_{j,\mathsf{min}}(\sigma,\tau)$ and $G^N_{j,\mathsf{max}}(\sigma,\tau)$.

Definition 1.5.6. A *Nash-equilibrium* of a mean payoff game consists in a pair $(\sigma^*,\tau^*) \in \Sigma_{\mathsf{min}} \times \Sigma_{\mathsf{max}}$ of strategies such that

$$\forall \sigma \in \Sigma_{\mathsf{min}}, G^N_{j,\mathsf{min}}(\sigma^*,\tau^*) \geqslant G^N_{j,\mathsf{min}}(\sigma,\tau^*) \tag{1.4a}$$

$$\forall \tau \in \Sigma_{\mathsf{max}}, G^N_{j,\mathsf{max}}(\sigma^*,\tau^*) \geqslant G^N_{j,\mathsf{max}}(\sigma^*,\tau) \ . \tag{1.4b}$$

The strategies σ^*, τ^* are refered to as *optimal strategies* for their respective player.

By Nash's existence theorem, the finite horizon game admits a Nash-equilibrium. In fact, mean payoff games in general do admit a Nash-equilibrium, and moreover the associated optimal strategies can be chosen positional (see [EM79]).

Notice that since a mean payoff game is in particular a zero-sum game, the following equality holds

$$G^N_{j,\mathsf{max}}(\sigma^*,\tau^*) = -G^N_{j,\mathsf{min}}(\sigma^*,\tau^*) \ .$$

This motivates the following definition of the value of a zero-sum game in general, and more precisely of a mean payoff game.

Definition 1.5.7. Let N be a nonnegative integer. The *value of the finite-horizon game, for the horizon N and the initial state j* corresponds to the quantity v^N_j defined by

$$v^N_j := G^N_{j,\mathsf{max}}(\sigma^*,\tau^*) = -G^N_{j,\mathsf{min}}(\sigma^*,\tau^*) \ .$$

The vector $v^N := (v^N_j)_{j\in[n]}$ is refered to as the *vector of values of the game in horizon N*.

The value of v^N_j corresponds to the minimal profit that the maximizer can guarantee after N turns, starting at the initial state j. It is also the maximal loss that the minimizer can ensure after N turns, for the same initial state. The sign of the value of the game (in horizon N) thus determines which player the game is most beneficial to, after N turns. We shall see in the next sections how the value of a mean payoff game is related to the theory of nonlinear eigenvalues of Shapley operators.

1.5.2 Shapley operator and value of a mean payoff game

We now describe the dynamic programming operator that arises from a mean payoff game. Given a payment matrix $B = (b_{ij})_{(i,j)\in I\times J}$, one can consider the max-plus linear operator defined with the convention that $(-\infty)+(+\infty) = (-\infty)$ as follows

$$\begin{array}{rcl} (\mathbb{R}\cup\{\pm\infty\})^J & \longrightarrow & (\mathbb{R}\cup\{\pm\infty\})^I \\ u & \longmapsto & Bu := (\max_{j\in J}(b_{ij}+u_j))_{i\in I} \ . \end{array}$$

This operator simply corresponds to the result of the tropical (max-plus) multiplication $B \odot u$ of the matrix B by the vector u.

Likewise, from a payment matrix $A = (a_{ij})_{(i,j)\in I\times J}$, one can consider the min-plus linear operator defined as follows, with the convention $(+\infty) + (-\infty) = (+\infty)$ this time.

$$\begin{array}{rcl} (\mathbb{R}\cup\{\pm\infty\})^I & \longrightarrow & (\mathbb{R}\cup\{\pm\infty\})^J \\ v & \longmapsto & A^\sharp v := (\min_{i\in I}(-a_{ij}+v_i))_{j\in J} \end{array} .$$

This operator corresponds this time to the min-plus multiplication of the transpose of the matrix $-A$ by the vector v. Given the max-plus linear operator $u \mapsto Au$ defined as before, the operator $A^\sharp$, is the so-called *residuated operator* of A, refering to the following property

$$Au \leqslant v \quad \Longleftrightarrow \quad u \leqslant A^\sharp v$$

where the inequalities are taken componentwise.

Finally, the *Shapley operator* or *dynamic programming operator* of the mean-payoff or finite horizon games associated to the graph $\mathrm{G}(A,B)$ as described in the previous section is simply the operator T defined by

$$\begin{array}{rcl} (\mathbb{R}\cup\{\pm\infty\})^J & \longrightarrow & (\mathbb{R}\cup\{\pm\infty\})^J \\ u & \longmapsto & T(u) := A^\sharp B u \end{array} .$$

In other words, one has

$$T(u) = \left(\min_{i\in I}(-a_{ij} + \max_{k\in J}(b_{ik}+u_k))\right)_{j\in J} \tag{1.5}$$

for all $u \in (\mathbb{R}\cup\{\pm\infty\})^J$. The operator T characterizes the mean payoff games associated to the graph $\mathrm{G}(A,B)$. For that reason, we shall denote this graph simply as $\mathrm{G}(T)$ for short. Note in particular that from the residuation property, one has that the tropical linear system $A\odot u \leqslant B\odot u$ is equivalent to the inequality $u \leqslant T(u)$.

Example 1.5.8. Consider the mean payoff game described in Example 1.5.2, with payment matrices

$$A = \begin{pmatrix} 2 & -\infty \\ 8 & -15 \\ -\infty & 0 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 1 & -\infty \\ 10 & -\infty \\ -\infty & 5 \end{pmatrix} .$$

Then, setting $T = A^\sharp B$ the Shapley operator of this game, the inequality $u \leqslant T(u)$ is equivalent to the following tropical linear system

$$\left\{ \begin{array}{rcl} 2+u_1 & \leqslant & 1+u_1 \\ \max(8+u_1, -15+u_2) & \leqslant & 10+u_1 \\ u_2 & \leqslant & 5+u_2 \end{array} \right. .$$

For given strategies $\sigma \in \Pi_{\mathsf{min}}$ and $\tau \in \Pi_{\mathsf{max}}$, we moreover define the associated *one-player dynamic programming maps* $T^{\sigma,\cdot}$ and $T^{\cdot,\tau}$ by

$$\begin{aligned} T_j^{\sigma,\cdot}(u) &:= -a_{\sigma(j)j} + \max_{k\in J}(b_{\sigma(j)k}+u_k) \\ T_j^{\cdot,\tau}(u) &:= \min_{i\in I}(-a_{ij} + b_{i\tau(i)} + u_{\tau(i)}) \ , \end{aligned}$$

as well as the *zero-player dynamic programming map* $T^{\sigma,\tau}$ by

$$T_j^{\sigma,\tau}(u) := -a_{\sigma(j)j} + (b_{\sigma(j)\tau(\sigma(j))} + u_{\tau(\sigma(j))}) \ .$$

Notice that the one-player dynamic programming operators simply correspond respectively to a max-plus and to a min-plus matrix multiplication.

Property 1.5.9. The map T of (1.5) satisfies the following properties:

(i) T is *order-preserving*: $\forall u,v \in (\mathbb{R}\cup\{\pm\infty\})^J$, $u\leqslant v \implies T(u)\leqslant T(v)$;

(ii) T is *additively homogeneous*: $\forall u \in (\mathbb{R}\cup\{\pm\infty\})^J$, $\forall\lambda\in\mathbb{R}$, $T(\lambda+u) = \lambda + T(u)$.

We shall sometimes require the following assumption in the remainder of this section.

Assumption 1.5.10. Let $T = A^\sharp B$ with $A = (a_{ij})_{(i,j)\in I\times J}$ and $B = (b_{ij})_{(i,j)\in I\times J}$. Then

(a) for all $j \in J$, there exists $i \in I$ such that $a_{ij} \neq -\infty$;

(b) for all $i \in I$, there exists $j \in J$ such that $b_{ij} \neq -\infty$.

In other words, the matrix A does not have a column identically equal to $-\infty$ and the matrix B does not have a row identically equal to $-\infty$, which translates into the fact that at every state, the set of possible actions for both players is never empty.

Remark 1.5.11. Under Assumption 1.5.10 (a) the map T preserves $(\mathbb{R}\cup\{-\infty\})^J$ and under Assumption 1.5.10 (b), T preserves $(\mathbb{R} \cup \{+\infty\})^J$. Moreover, whenever both assumptions are satisfied, T also preserves $\mathbb{R}^J$, in which case it follows from Property 1.5.9 that it is *sup-norm nonexpansive* over $\mathbb{R}^J$, that is

$$\forall u, v \in \mathbb{R}^J,\ \|T(u) - T(v)\|_\infty \leqslant \|u - v\|_\infty \ ,$$

where $\|\cdot\|_\infty$ denotes the sup-norm. *A fortiori*, the map T is continuous and piecewise affine over $\mathbb{R}^J$.

Remark 1.5.12. Recall that any order-preserving additively homogeneous operator T is the Shapley operator of a deterministic or stochastic zero-sum two player game with a possibly infinite number of actions for the minimizer player (see [AGG12, Section 2.2] for references). If T is also piecewise affine over $\mathbb{R}^J$, then it can also be seen as the Shapley operator of a zero-sum two player game with a finite number of actions, however the game may not be deterministic, that is it may not be of the form of (1.5).

The following crucial theorem from Kohlberg applies for maps that satisty Property 1.5.9 and Assumption 1.5.10 above.

Theorem 1.5.13 ([Koh80, Theorem 2.1])**.** *A self-map f of $\mathbb{R}^n$ that is nonexpansive in any norm and piecewise affine admits an invariant halfline, meaning that there exist two vectors $u, \eta \in \mathbb{R}^n$, with η unique, such that*

$$f(u + s\eta) = u + (s + 1)\eta \quad \textit{for all scalars } s \in \mathbb{R} \textit{ large enough.}$$

The vector u will be refered to as the *base point* of the invariant halfline, and the vector η as its *direction*. In particular, the Shapley operator T from (1.5) satisfies the conditions of the Kohlberg theorem, thus there exist $u, \eta \in \mathbb{R}^n$ such that $T(u + s\eta) = u + (s + 1)\eta$ for $s \in \mathbb{R}$ large enough. We then set $\chi(T) = (\chi_j(T))_{j\in J} := \eta$. For all $j \in J$, the coordinate $\chi_j(T)$ corresponds to the *value* of the mean payoff game described above, that is the average payment per turn of the minimizer to the maximizer when they both go on playing optimal strategies and the game starts at state j. The vector $\chi(T)$ is called the *vector of values* of the mean payoff game. For more information on mean-payoff games and their values, the reader can for instance refer to the introductive section 2.2 of [AGQS23].

For the particular case of Shapley operators, the additive homogeneousness entails the following corollary of the Kohlberg theorem.

Corollary 1.5.14. *Let T be the Shapley operator given in* (1.5) *and let $u = (u_j)_{j\in J} \in \mathbb{R}^J$ be such that for $s \in \mathbb{R}$ large enough, $T(u + s\chi(T)) = u + (s + 1)\chi(T)$. Moreover, let $\lambda = \max_{j\in J} \chi_j(T)$ and let $v = (v_j)_{j\in J} \in (\mathbb{R} \cup \{-\infty\})^J$ where $v_j = u_j$ whenever $\chi_j(T) = \lambda$ and $v_j = -\infty$ otherwise. Then $T(v) = \lambda + v$.*

Proof. First, notice that by additive homogeneousness of T, for all $s \in \mathbb{R}$ large enough, the equality $T(u + s(\chi(T) - \lambda)) = u + s(\chi(T) - \lambda) + \chi(T)$ holds. Moreover, by construction, $\lim_{s\to+\infty} u + s(\chi(T) - \lambda) = v$, and $\chi(T) + v = \lambda + v$. Therefore, taking the limit as s goes to $+\infty$ in the previous equality entails that $T(v) = \lambda + v$. $\square$

The vector of values of a mean payoff game can also be thought of as the limit of the value of a family of games in horizon N, as N tends to infinity. The vector of values $v^N = (v_j^N)_{j\in J}$ of the game in horizon Ncan be computed recursively using the Shapley operator T of (1.5) via the following dynamic programming relation

$$\begin{cases} v^0 = (0, \dots, 0) \in \mathbb{R}^J \\ v^{N+1} = T(v^N) \quad \text{for all } N \geqslant 0 \ . \end{cases}$$

The link between the value of the mean payoff game and the value of the finite horizon game is given by the following relation

$$\chi(T) = \lim_{N\to+\infty} \frac{v^N}{N} = \lim_{N\to+\infty} \frac{T^N(0)}{N} \ ,$$

which follows from Theorem 1.5.13 and is also a special case of the existence of a uniform value (see [EM79, MN81]) whenever Assumption 1.5.10 holds, *i.e.* whenever T preserves $\mathbb{R}^n$.

As a consequence of [LL69, Theorem 1], there is a pair of positional strategies $(\sigma^*, \tau^*) \in \Pi_{\mathsf{min}} \times \Pi_{\mathsf{max}}$ such that for all initial state $j \in J$, the minimizer player can ensure by following strategy σ^* that the average gain per turn of the maximizer does not exceed $\chi_j(T)$ and conversely. In fact, it is whown in [EM79, Theorem 2] that $\chi_j(T)$ corresponds to the average weight of a cycle in the graph $\mathrm{G}(T)$ of the mean payoff game that can be reached from the initial position $j \in J$. More precisely, for positional strategies $\sigma \in \Pi_{\mathsf{min}}$ of the minimizer and $\tau \in \Pi_{\mathsf{max}}$ of the maximizer, and for an initial position $j \in J$, a cycle $j_0, i_0, \dots, j_{\ell-1}, i_{\ell-1}, j_0$ is eventually reached in $\mathrm{G}(T)$, and its mean weight is denoted by

$$\Phi_j^{\sigma,\tau}(T) = \frac{1}{\ell} \sum_{t=0}^{\ell-1} -a_{i_t j_t} + b_{i_t j_{t+1}} \enspace .$$

Then one has the following result.

Theorem 1.5.15 (see [LL69, EM79]). *For every initial state $j \in J$, one has*

$$\chi_j(T) = \min_{\sigma \in \Pi_{\mathsf{min}}} \max_{\tau \in \Pi_{\mathsf{max}}} \Phi_j^{\sigma,\tau}(T) = \max_{\tau \in \Pi_{\mathsf{max}}} \min_{\sigma \in \Pi_{\mathsf{min}}} \Phi_j^{\sigma,\tau}(T) \enspace .$$

Moreover, there is a Nash equilibrium in the space of positional strategies i.e. *there exists a pair $(\sigma^*, \tau^*) \in \Pi_{\mathsf{min}} \times \Pi_{\mathsf{max}}$ of positional strategies such that $\Phi_j^{\sigma^*,\tau}(T) \leqslant \Phi_j^{\sigma^*,\tau^*}(T) \leqslant \Phi_j^{\sigma,\tau^*}(T)$ for all positional strategies $(\sigma, \tau) \in \Pi_{\mathsf{min}} \times \Pi_{\mathsf{max}}$.*

Remark 1.5.16. The length ℓ of a cycle in the graph $\mathrm{G}(T)$ of the mean payoff game is bounded above by $2\min(|I|, |J|)$. Therefore, if all the weights of the edges of $\mathrm{G}(T)$ are integer, then it follows from Theorem 1.5.15 that $\chi_j(T)$ is a rational number with a denominator bounded above by $2\min(|I|, |J|)$.

The strategies σ^* and τ^* of the previous theorem are refered to as *optimal strategies*, in the sense that both players have no interest in deviating from these strategies, as per the above inequality.

In a similar fashion to the quantity $v_j^N = T_j^N(0)$, the quantity $(T^{\sigma,\cdot})_j^N(0)$ represents the maximal amount the maximizer player can guarantee to win after N turns, if the minimizer follows the positional strategy $\sigma \in \Pi_{\mathsf{min}}$, given the initial state $j \in [n]$ and the quantity $-(T^{\cdot,\tau})_j^N(0)$ has an analogous interpretation for a positional strategy $\tau \in \Pi_{\mathsf{max}}$ of the maximizer player.

Moreover, the following equality holds by construction

$$\forall u \in (\mathbb{R} \cup \{\pm\infty\})^J,\, T(u) = \min_{\sigma \in \Pi_{\mathsf{min}}} T^{\sigma,\cdot}(u) = \max_{\tau \in \Pi_{\mathsf{max}}} T^{\cdot,\tau}(u) \enspace , \tag{1.6}$$

and the maximum as well as the minimum are both attained by at least one positional strategy each since the space of strategies for each player is finite.

In the context of mean payoff games, the quantity of interest is precisely the limit as the number N of turns tends to $+\infty$ of the average payment received by the maximizer from the minimizer each turn. From now on, this quantity will be simply refered to as the *average gain per turn for the maximizer*, or as the *average loss per turn for the minimizer*. If there exists a strategy for the maximizer which ensures that this quantity exists and is nonnegative, then the game can be considered as favourable to the maximizer. Thus, if $\chi_j(T) \geqslant 0$, we will say that $j \in J$ is a *winning initial state for the maximizer*. Likewise, we define the *winning initial states for the minimizer* as the set of states $j \in J$ such that $\chi_j(T) \leqslant 0$. This is illustrated by the following theorem, which can be regarded as a particular case of the *duality conjecture*, proven in [GG98], see also [AGG12, §2.4].

Theorem 1.5.17. *Make Assumption 1.5.10. Then one has*

$$\chi(T) = \min_{\sigma \in \Pi_{\mathsf{min}}} \chi(T^{\sigma,\cdot}) = \max_{\tau \in \Pi_{\mathsf{max}}} \chi(T^{\cdot,\tau}) \enspace . \tag{1.7}$$

In particular, the maximizer can choose a positional strategy such that he can guarantee an average gain per turn greater than or equal to $\chi_j(T)$, given the initial state j, whatever strategy the minimizer chooses, and likewise, the minimizer can choose a positional strategy such that he can guarantee an average loss per turn of no more than $\chi_j(T)$, given the initial state j, whatever strategy the maximizer chooses.

Proof. By monotonicity of the limit, the map $T \mapsto \chi(T)$ is order preserving. Moreover, it follows from the previous remark that $T^{\cdot\sigma} \leqslant T$ for all strategy σ, hence

$$\chi(T^{\cdot\sigma}) \leqslant \chi(T), \quad \forall \sigma \in \Pi_{\mathsf{max}}.$$

Let then (u, η) be an invariant half-line of T, such that $T(u + s\eta) = u + (s + 1)\eta$ for s large enough, and fix $i \in I$. Noticing that the maximum of a finite family of affine functions coincides with one of these functions on a neighbourhood of $+\infty$, we deduce that for s large enough, there exists $\sigma(i) \in I$ such that

$$\max_{k \in J}(b_{ik} + u_k + s\eta_k) = b_{i\sigma(i)} + u_{\sigma(i)} + s\eta_{\sigma(i)} \ .$$

Now adding $-a_{ij}$ for $j \in J$, then taking the infimum on $i \in I$, we then deduce that for t large enough, $T_j(u+s\eta) = T_j^{\cdot\sigma}(u + s\eta)$. Since this is true for any $j \in J$, we have indeed for s large enough,

$$u + (s + 1)\eta = T(u + s\eta) = T^{\cdot\sigma}(u + s\eta) \ ,$$

therefore (u, η) is also an invariant half-line of $T^{\cdot\sigma}$, and thus by Kohlberg's theorem, it follows that $\chi(T) = \chi(T^{\cdot\sigma}) = \eta$, and in particular σ attains the maximum in (1.7). The proof of the second equality is analogous. □

1.5.3 Nonlinear eigenvalue theory and solvability of tropical linear systems

Given a Shapley operator $T : (\mathbb{R} \cup \{\pm\infty\})^J \to (\mathbb{R} \cup \{\pm\infty\})^J$, we shall take interest in the following problem:

$$\text{find} \quad (\lambda, u) \in \mathbb{R} \times \mathbb{R}^J \quad \text{such that} \quad T(u) = \lambda + u \ . \tag{1.8}$$

This problem has been thouroughly studied in the particular case of one-player operators, where it simply reduces to finding nonlinear eigenvalues and eigenvectors of tropical matrices. Therefore, we shall use this same terminology of nonlinear eigenvalue theory to describe solutions of (1.8) in the two-player case. Moreover, many results about two-player Shapley operators rely on the one-player case. For an overview of the nonlinear eigenvalue theory for one-player Shapley operators, the reader is advised to refer to [BCOQ92, But10], which set the foundation for many of the following results.

Definition 1.5.18. The above equation shall be refered to as the *ergodic equation* associated to the operator T, and if $(\lambda, u) \in \mathbb{R} \times \mathbb{R}^J$ is a solution of (1.8), then λ is called a *nonlinear eigenvalue* of T and u a *nonlinear eigenvector* or *bias vector* of T. The *nonlinear eigenspace* of T is the set $\operatorname{Eig}(T)$ of all nonlinear eigenvectors associated to operator T.

Remark 1.5.19. Note that if $u \in \mathbb{R}^J$ is a nonlinear eigenvector of T, then adding a constant to all coordinates of u also yields another eigenvector of T by additive homogeneity, thus $\operatorname{Eig}(T)$ can be written as a union of lines in $\mathbb{R}^J$ directed by the vector $\mathbf{1} = (1, \ldots, 1) \in \mathbb{R}^J$. Therefore, the set $\operatorname{Eig}(T)$ naturally induces a subset $\operatorname{Eig}(T)/\mathbb{R}\mathbf{1}$ of the tropical projective space $\mathbb{R}^J/\mathbb{R}\mathbf{1}$. In particular, if $\operatorname{Eig}(T)$ is reduced to a single such line, then the bias vector u solution of (1.8) is said to be unique in the projective sense.

The ergodic equation can be rewritten into the following form

$$\min_{i \in I} \quad -a_{ij} + v_i \quad = \quad \lambda + u_j \quad \forall j \in J \tag{1.9a}$$

$$\max_{j \in J} \quad b_{ij} + u_j \quad = \quad v_i \qquad \forall i \in I \ , \tag{1.9b}$$

refered to as the *ergodic problem*, where the unknowns are $\lambda \in \mathbb{R}$, $u \in \mathbb{R}^J$ and $v \in \mathbb{R}^I$. The sets of *active constraints* (or *saturated constraints*) of the ergodic problem are given by

$$I_j(u) := \operatorname*{arg\,min}_{i \in I} -a_{ij} + v_i \quad \text{for all } j \in J$$

and

$$J_i(u) := \operatorname*{arg\,min}_{j \in J} b_{ij} + u_j \quad \text{for all } i \in I \ .$$

These active constraints generate a subgraph of the graph of the game associated to the operator T according to the following definition.

Definition 1.5.20. The *saturation graph* associated to a bias vector u of the operator T is the subgraph $\mathrm{SAT}(T,u)$ of the graph $\mathrm{G}(T)$ of the mean payoff game associated to the operator T obtained by only keeping the arcs (j,i) of G such that $i \in I_j(u)$ as well as the arcs (i,j) such that $j \in J_i(u)$.

Now, we state a few results regarding the existence and uniqueness of solutions to the ergodic equation, starting with the following particular case of Corollary 1.5.14 of the Kohlberg theorem.

Corollary 1.5.21. *The ergodic equation* (1.8) *has a solution* $(\lambda, u) \in \mathbb{R} \times \mathbb{R}^J$ *if and only if the value vector* $\chi(T)$ *is constant, in which case* $\chi(T) \equiv \lambda$*, and thus the nonlinear eigenvalue is always unique whenever it exists.*

Moreover, equality (1.6) entails that if $u \in \mathbb{R}^J$ is a nonlinear eigenvector of the operator T, then there exists a positional strategy $\sigma \in \Pi_{\mathsf{min}}$ of the minimizer (resp. $\tau \in \Pi_{\mathsf{max}}$ of the maximizer) such that u is also a nonlinear eigenvector of $T^{\sigma,\cdot}$ (resp. $T^{\cdot,\tau}$) with the same associated nonlinear eigenvalue. In particular, this entails the following inclusions

$$\mathrm{Eig}(T) \subseteq \bigcup_{\sigma \in \Pi_{\mathsf{min}}} \mathrm{Eig}(T^{\sigma,\cdot}) \quad \text{and} \quad \mathrm{Eig}(T) \subseteq \bigcup_{\tau \in \Pi_{\mathsf{max}}} \mathrm{Eig}(T^{\cdot,\tau}) \ . \tag{1.10}$$

Now fix a positional strategy $\sigma \in \Pi_{\mathsf{min}}$ of the minimzer and let $\lambda^{\sigma,\cdot}$ be the maximal average weight of all cycles in the graph $\mathrm{G}(T^{\sigma,\cdot})$ of the mean payoff game associated to the operator $T^{\sigma,\cdot}$. A cycle of maximal weight in $\mathrm{G}(T^{\sigma,\cdot})$ is refered to as a critical cycle and the subgraph of $\mathrm{G}(T^{\sigma,\cdot})$ consisting in all vertices and edges belonging to a *critical cycle* is refered to as the *critical graph*. For the one-player operator $T^{\sigma,\cdot}$, one can characterize the uniqueness of the nonlinear eigenvector in the projective sense — that is up to an additive constant — with the following classical result found for instance in [But10, BCOQ92].

Theorem 1.5.22 (Corollary of [BCOQ92, Theorem 3.101])**.** *The one-player ergodic equation* $T^{\sigma,\cdot}(u) = \lambda + u$ *has a unique solution* $(\lambda, u) \in \mathbb{R} \times \mathbb{R}^J$ *if and only if the critical graph has a unique connected component, where the uniqueness of the bias vector* u *is taken in the projective sense.*

Remark 1.5.23. *A fortiori*, the previous theorem entails that if there is a unique critical cycle in the graph $\mathrm{G}(T^{\sigma,\cdot})$, then the solution to the one-player ergodic equation is unique.

Example 1.5.24. Again, consider the mean payoff game from Example 1.5.2, with payment matrices

$$A = \begin{pmatrix} 2 & -\infty \\ 8 & -15 \\ -\infty & 0 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 1 & -\infty \\ 10 & -\infty \\ -\infty & 5 \end{pmatrix} ,$$

and Shapley operator $T = A^\sharp B$. Then the associated nonlinear eigenproblem is the following problem

$$\begin{aligned} \min(-2 + v_1, -8 + v_2) &= \lambda + u_1 \\ \min(15 + v_2, v_3) &= \lambda + u_2 \\[1em] 1 + u_1 &= v_1 \\ 10 + u_1 &= v_2 \\ 5 + u_2 &= v_3 \end{aligned}$$

As per Remark 1.5.19, one can look for a bias vector $u = (u_1, u_2)$ such that $u_1 = 0$, and the system can thus be solved by hand, to find the solution $\lambda = -1$, $u = (0, 26)$ and $v = (1, 10, 31)$. The vector of value os this game is thus $\chi(T) = (-1, -1)$ meaning that both starting positions are winning initial states for the minimizer. Moreover, given the previous solution, one can compute the sets of active constraints, and the resulting saturation graph is represented on Figure 1.7 below. In particular, there is a unique critical cycle given by $1 \xrightarrow{-2} \boxed{1} \xrightarrow{1} 1$, whose average weight is indeed equal to $\frac{-2+1}{2} = 1$.

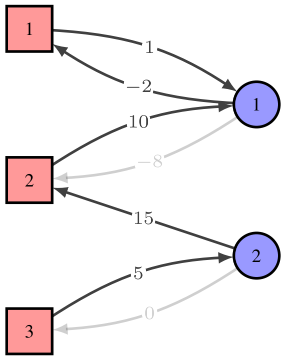

Figure 1.7: The saturation graph $\mathrm{SAT}(T, u)$ consists in the subgraph of $\mathrm{G}(T)$ obtained by removing all the transparent edges.

One finally provides the following result, which states that this generically holds.

Proposition 1.5.25. *Consider for $\delta \in \mathbb{R}^{I\times J}$ the perturbed operator T_δ defined by $T_\delta = A^\sharp(B + \delta)$. Then for a generic choice of δ — meaning for all $\delta \in \mathbb{R}^{I\times J}$ outside of a finite union of hyperplanes — there is a unique critical cycle in the graph $\mathrm{G}(T_\delta^{\sigma,\cdot})$ of the mean payoff game associated to the perturbed one-player operator for all $\sigma \in \Pi_{\mathsf{min}}$.*

Proof. Assume that there are two distinct critical cycles $j_0, i_0, \dots, j_{\ell-1}, j_{\ell-1}$ and $j'_0, i'_0, \dots, j'_{\ell'-1'}, i'_{\ell'-1}$ in $\mathrm{G}(T_\delta^{\sigma,\cdot})$. Then by definition of the one-player operator, one has $i_k = \sigma(j_k)$ for all $0 \leqslant k \leqslant \ell - 1$ and likewise $i'_{k'} = \sigma(j'_{k'})$ for all $0 \leqslant k' \leqslant \ell' - 1$, and their weights are both equal to the maximal possible weight, hence

$$\frac{1}{\ell}\sum_{k=0}^{\ell} a_{\sigma(j_k)j_k} + b_{\sigma(j_k)j_{k+1}} + \delta_{\sigma(j_k)j_{k+1}} = \frac{1}{\ell'}\sum_{k'=0}^{\ell'} a_{\sigma(j'_{k'})j'_{k'}} + b_{\sigma(j'_{k'})j'_{k'+1}} + \delta_{\sigma(j'_{k'})j'_{k'+1}} \ ,$$

which immediately entails a nontrivial linear relationship on δ, which thus does not hold generically. $\square$

One can in fact show the following stronger result, to be compared with [AGH18, Theorem 3.2].

Theorem 1.5.26. *Consider the perturbed operator T_δ as defined for all $\delta \in \mathbb{R}^{I\times J}$ in the previous proposition., and fix a positional strategy $\sigma \in \Pi_{\mathsf{min}}$ of the minimizer player. Then there exists a full dimensional polyhedral complex $\mathscr{C}^{\sigma,\cdot}$ of $\mathbb{R}^{I\times J}$ such that for all δ in the interior of a maximal dimensional cell, $\mathrm{G}(T_\delta^{\sigma,\cdot})$ has a unique critical cycle, which is moreover independent of the choice of δ.*

Proof. Fix a positional strategy $\sigma \in \Pi_{\mathsf{min}}$ of the minimizer and let $\lambda^{\sigma,\cdot}(\delta)$ be the maximal weight of all cycles in the graph $\mathrm{G}(T_\delta^{\sigma,\cdot})$. Since the payments are affine in $\delta \in \mathbb{R}^{I\times J}$, then so is $\lambda^{\sigma,\cdot}$. Set $\mathscr{C}^{\sigma,\cdot}$ the linearity complex of $\lambda^{\sigma,\cdot}$, that is the subdivision of $\mathbb{R}^{I\times J}$ such that $\lambda^{\sigma,\cdot}$ is affine in the interior of all maximal-dimensional cells of $\mathscr{C}^{\sigma,\cdot}$. In particular, since $\lambda^{\sigma,\cdot}$ is a maximum of finitely many affine functions, for every cell C of $\mathscr{C}^{\sigma,\cdot}$, the set of cycles of maximal weight in $\mathrm{G}(T_\delta^{\sigma,\cdot})$ do not depend on the choice of $\delta \in \mathrm{relint}(C)$. In particular, in the interior of the maximal-dimensional cells, Proposition 1.5.25 prevents the existence of two cycles of maximal weight by genericity, hence the result. $\square$

The vector of values of the mean payoff game also has an interpretation in terms of the nonlinear eigenvalues of the Shapley operator T. More precisely, we set the following quantities for any order-preserving additively homogeneous operator T:

- the *Collatz-Wielandt number* $\mathrm{cw}(T)$ of T defined by

$$\mathrm{cw}(T) := \inf\{\lambda \in \mathbb{R} : \exists u \in \mathbb{R}^J, T(u) \leqslant \lambda + u\} \tag{1.11a}$$

- the *symmetrical Collatz-Wielandt number* $\mathrm{cw}'(T)$ of T defined by

$$\mathrm{cw}'(T) := \sup\{\lambda \in \mathbb{R}\cup\{-\infty\} : \exists u \in (\mathbb{R}\cup\{-\infty\})^J, u \not\equiv -\infty, T(u) \geqslant \lambda + u\} \tag{1.11b}$$

- the *nonlinear spectral radius* $\rho(T)$ of T defined by

$$\rho(T) := \sup\{\lambda \in \mathbb{R} \cup \{-\infty\} : \exists u \in (\mathbb{R} \cup \{-\infty\})^J, u \not\equiv -\infty, T(u) = \lambda + u\} \ . \tag{1.11c}$$

Then, the following result can be applied to any operator T of the form (1.5) satisfying Assumption 1.5.10.

Theorem 1.5.27 (Collatz-Wielandt property [AGG12, Lemma 2.8] and [AGQS23, Theorem 1])**.** *Let T be an order-preserving additively homogeneous self-map of $(\mathbb{R} \cup \{\pm\infty\})^J$. Assume that T preserves $(\mathbb{R} \cup \{-\infty\})^J$ and set $\overline{\chi}(T) := \max\{\chi_j(T) : j \in J\}$. Then*

$$\overline{\chi}(T) = \mathrm{cw}(T) = \mathrm{cw}'(T) = \rho(T) \ ,$$

and the suprema in (1.11c) *and* (1.11b) *are attained. If in addition T is piecewise affine or $-\infty$ over $\mathbb{R}^J$, and $\overline{\chi}(T) \neq -\infty$, then the infimum in* (1.11a) *is attained.*

Remark 1.5.28. If T is an order-preserving additively homogeneous self-map of $(\mathbb{R} \cup \{\pm\infty\})^J$ and if T preserves $\mathbb{R}^J$, then *a fortiori* T preserves $(\mathbb{R} \cup \{-\infty\})^J$ and thus $\mathrm{cw}'(T) \neq -\infty$, so all the conclusions of the above result hold if T is piecewise affine.

All the quantities in (1.11) can be dualized, so we obtain the following similar result for the dual quantities.

Corollary 1.5.29. *Let T be an order-preserving additively homogeneous self-map of $(\mathbb{R} \cup \{\pm\infty\})^J$ and assume that T preserves $(\mathbb{R} \cup \{+\infty\})^J$. Then, the following quantities coincide and they are all equal to $\underline{\chi}(T) := \min\{\chi_j(T) : j \in J\}$*

$$\sup\{\lambda \in \mathbb{R} : \exists u \in \mathbb{R}^J, T(u) \geqslant \lambda + u\} \tag{1.12a}$$

$$\inf\{\lambda \in \mathbb{R} \cup \{+\infty\} : \exists u \in (\mathbb{R} \cup \{+\infty\})^J, u \not\equiv +\infty, T(u) \leqslant \lambda + u\} \tag{1.12b}$$

$$\inf\{\lambda \in \mathbb{R} \cup \{+\infty\} : \exists u \in (\mathbb{R} \cup \{+\infty\})^J, u \not\equiv +\infty, T(u) = \lambda + u\} \ . \tag{1.12c}$$

Moreover, if T preserves $\mathbb{R}^J$ and is piecewise affine, then the infima and supremum in (1.12) *are attained.*

At last, one can state the correspondance between mean payoff games and tropical linear systems, in the form of the next theorem, which is a reformulation of the main result of [AGG12].

Theorem 1.5.30 ([AGG12, Theorem 3.2])**.** *With the same notation and conditions as above, then for all $j \in J$, there exists $u \in \mathbb{T}^J$ with $u_j \neq \mathbb{0}$ such that $A \odot u \leqslant B \odot u$ if and only if $\chi_j(T) \geqslant 0$.*

In other words, the support of the solutions $u \in \mathbb{T}^J$ of the tropical linear inequality $A \odot u \leqslant B \odot u$ corresponds exactly to the set of initial positions $j \in J$ that are winning for the maximizer. In particular, the next corollary allows one to check the existence of a solution in $\mathbb{R}^J$ to a tropical linear system.

Corollary 1.5.31 ([AGG12, Corollary 3.4])**.** *The tropical linear system $A \odot u \leqslant B \odot u$ has a solution $u \in \mathbb{R}^J$ if and only if all the initial states of the associated game have a nonnegative value,* i.e. $\chi(T) \geqslant 0$.

Example 1.5.32. The tropical linear system

$$\left\{ \begin{array}{rcl} 2 + u_1 & \leqslant & 1 + u_1 \\ \max(8 + u_1, -15 + u_2) & \leqslant & 10 + u_1 \\ u_2 & \leqslant & 5 + u_2 \end{array} \right.$$

arising from Example 1.5.2 does not have a solution $u \in \mathbb{R}^2$, as it has already been established that $\chi(T) = (-1, -1) \not\geqslant 0$.

Remark 1.5.33. Thanks to the above tropical Positivstellensatz, checking for the existence of a solution in $\mathbb{R}^n$ of a n-variate tropical polynomial system of inequalities, reduces to checking the solvability of a linearized system of the form $A \odot u \leqslant B \odot u$, by solving the associated mean payoff game, that is computing its vector of values. In fact, this method generalizes to all tropical polynomial systems consisting of a mixture of tropical linear equalities, weak and strict inequalities as well as ∇ relations. Indeed, strict inequalities and equalities both reduce to weak inequalities as for all $A, B \in \mathbb{T}^{I \times J}$ and for all $u \in \mathbb{T}^J$,

$$A \odot u = B \odot u \iff \left\{ \begin{array}{l} A \odot u \geqslant B \odot u \\ A \odot u \leqslant B \odot u \end{array} \right.$$

and, using the fact that the value of a mean-payoff game is a short rational,

$$A \odot u > B \odot u \iff A \odot u \geqslant \lambda \odot B \odot u$$

for a short rational $\lambda > \mathbb{1} = 0$, see [AFG+14] for the systems mixing strict and weak tropical inequalities. Moreover, Theorem 4.7 from [AGG12] also shows that tropical linear systems of the form $A \odot u \;\nabla\; \mathbb{0}$ reduce to tropical linear systems of the form $\tilde{A} \odot u \leqslant \tilde{B} \odot u$, and gives the construction of the tropical matrices $\tilde{A}$ and $\tilde{B}$ in function of A.

Following the previous remark, for simplicity, we will restrict ourselves to tropical linear systems of the form $A \odot u \leqslant B \odot u$ in the remainder of this paper.

1.5.4 Classical algorithms for solving mean payoff games

In this section, we recall two algorithms for solving mean payoff games: the value iteration algorithm and the policy iteration algorithm.

Mean payoff games oracles

To check the solvability of a tropical linear system using Corollary 1.5.31, it suffices to compute the value vector of a mean payoff game. Actually, a weaker information will be enough for some of our results.

We call *weak mean payoff game oracle* a procedure which takes as input two tropical matrices A, B, and decides whether $\underline{\chi}(T) \geqslant 0$ with $T = A^\sharp B$. We denote by $\mathsf{w\text{-}MPG}(|I|, |J|, r^\infty)$ the number of arithmetic operations of a mean payoff oracle taking as input $|I| \times |J|$ matrices A, B whose entries are either relative integers of absolute values bounded by r^∞ or $-\infty$. We observe that $\mathsf{w\text{-}MPG}(|I|, |J|, r^\infty) \geqslant |I||J|$ since the input size is $\Omega(|I \times J|)$.

We shall also use the notion of *strong mean payoff game oracle*, which not only decides whether $\underline{\chi}(T) \geqslant 0$, as a weak oracle does, but also decides whether $\chi(T)$ is a constant vector, and if this the case, returns a nonlinear eigenvector vector $u \in \mathbb{R}^J$ associated to the nonlinear eigenvalue λ, with $\chi(T) \equiv \lambda$ as per Corollary 1.5.21. The bias vector u serves as an optimality certificate, allowing one to identify optimal policies. It may be non unique, even up to an additive constant. In fact, the set of possible biases belongs to a particular class of polyhedral complexes, called an 'ambitropical polyhedra', which are characterized in [AGV23]. We denote by $\mathsf{MPG}(|I|, |J|, r^\infty)$ the number of arithmetic operations of a strong mean payoff oracle.

The value iteration algorithm

A classical algorithm to solve mean payoff games is the value iteration, analyzed in [ZP96]. It consists in computing the sequence $T^N(0)$ and inferring the limit $\lim_{N\to+\infty} T^N(0)/N$ by specializing N to an explicit sufficiently large value, exploiting the fact that the value of a mean payoff game is a rational number with a small denominator, so that the exact value can be obtained from an approximate value by a rounding argument.

Theorem 1.5.34 (Corollary of [ZP96, Theorem 2.4])**.** *The value iteration algorithm provides a weak mean payoff oracle requiring $\mathcal{O}(|J|^2 r^\infty)$ evaluations of the Shapley operator T, entailing*

$$\mathsf{w\text{-}MPG}(|I|, |J|, r^\infty) = \mathcal{O}(|I||J|^3 r^\infty) \enspace .$$

The number of iterations of the method of [ZP96] is always in $\Omega(|J|^2 r^\infty)$, which is unpracticable in our application, as J will be exponentially large in the input size. We shall however present in Chapter 3 a refinement of the value iteration, first introduced in [ABG23a], exploiting the ideas of Krasnoselskii-Mann damping with an acceleration or *widening* step. This accelerated version will allow in practice for a much quicker check of feasibility. We coin the term "widening" by analogy with the field of static analysis of program by abstract interpretation, in which various accelerations of Kleene's fixed point iteration, of a different nature, are commonly used [CC77].

We also have the following result concerning strong mean payoff game oracles.

Theorem 1.5.35. *A strong mean payoff oracle can be implemented by making $\mathcal{O}(|J|^3 r^\infty)$ evaluations of the Shapley operator T, leading to*

$$\mathsf{MPG}(|I|, |J|, r^\infty) = \mathcal{O}(|I||J|^4 r^\infty) \enspace .$$

Proof. We first compute $\chi(T)$ by means of [ZP96, Theorem 2.3]. Moreover, when $\chi(T) \equiv \lambda \in \mathbb{R}$ is a constant vector, we first perform the iteration $u^{k+1} = (-\lambda + T)(u^k) \wedge u^k$, starting from $u^0 = 0$, and show it converges to a vector u such that $u \leqslant (-\lambda + T)(u)$, in $\mathcal{O}(|J|^3 r^\infty)$ iterations, then we perform the iteration $v^{k+1} = (-\lambda + T)(v^k)$, starting from $v^0 = u$, and show it converges to a bias vector v, satisfying $v = (-\lambda + T)(v)$, again in $\mathcal{O}(|J|^3 r^\infty)$ iterations, leading to Theorem 1.5.35. $\square$

The policy iteration algorithm

Another approach to solve mean payoff games is the policy iteration algorithm. The idea of policy iteration was introduced by Howard in the one-player case, and extended by Hoffman and Karp [HK66] to a class of two-player stochastic games satisfying an ergodicity condition. This condition is generally not satisfied in the deterministic case, so we rely on the extension of this algorithm developed in [CTGG99, GG98, DG06], in which degenerate steps induced by the absence of ergodicity are dealt with using a spectral projector technique.

Recall that the operator $T = A^\sharp B$ is associated to the mean payoff game described by the weighted bipartite graph $G = (I \sqcup J, E)$ such that for all $(i,j) \in J \times I$, a_{ij} is the weight of the arc (j,i) if it exists, or $-\infty$ otherwise, and b_{ij} is the weight of the arc (i,j) if it exists, or $-\infty$ otherwise. If we fix a positional strategy $\sigma \in \Pi_{\mathsf{min}}$ of the minimizer, then one can consider the one-player operator $T^{\sigma,\cdot}$ defined by $T^\sigma(u)_j = -a_{\sigma(j)j} + \max_{k\in J}(b_{\sigma(j)k} + u_k$ for all $j \in J$ and for all $u \in (\mathbb{R} \cup \{\pm\infty\})^J$. Then, the idea of the policy iteration algorithm such as described *e.g.* in [DG06, Algorithm 2] is to compute an invariant half-line for a family $(\sigma_k)_{k\geqslant 0}$ of stationary policies of the minimizer, that will eventually converge (in finite time) towards the optimal policy σ^*, which will then satisfy $\chi(T^{\sigma^*,\cdot}) = \chi(T)$ as per Theorem 1.5.17. A counterexample of Friedmann [Fri11] entails that policy iteration methods can take an exponential time in the worst case, although they appear to be remarkably efficient in practice. In particular, an experimental study [Cha09] suggests that the algorithm of [DG06] (used in our experiments) is among the fastest ones.

Remark 1.5.36. A different policy iteration approach would rely on the approximation of the mean-payoff problem by a discounted problem, leading to a pseudo-polynomial bound. Indeed, Zwick and Paterson showed that solving the discounted game with a discount factor of $\alpha = 1 - 1/(4|J|^3 r^\infty)$ allows one to solve the mean-payoff game. Ye showed that policy iteration with a fixed discount factor is strongly polynomial in [Ye05], and the same is true for games [HMZ11]. Applying the refined complexity bound from [AG13, Theorem 5], one can show that the discounted game (and so the mean-payoff game) can be solved in total of $O(s^2|I||J|^7(r^\infty)^2 \log(|J|r^\infty)^2)$ policy iterations where s is the maximal number of finite entries in every row of the matrices A and B. Note that in this bound, the term r^∞ (maximal absolute value of an instantaneous payment) is squared, whereas r^∞ only appears linearly in the complexity bound of value iteration.

Chapter 2

The tropical Nullstellensatz and Positivstellensatz

In the present chapter, we describe the construction of a Nullstellensatz and a Positivstellensatz adatped to sparse tropical polynomial systems, as first introduced in [ABG23a, ABG23b]. A first tropical analogue of the effective Nullstellensatz was established in [GP18], showing that a system of n-variate tropical polynomial equations is solvable over $(\mathbb{T}^*)^n$ if and only if a linearized system obtained from a truncated Macaulay matrix is solvable over $(\mathbb{T}^*)^N$ for some truncation degree N. Grigoriev and Podolskii provided an upper bound of the minimal admissible truncation degree, as a function of the degrees of the tropical polynomials. Our approach is inspired by a construction of Canny-Emiris (1993), refined by Sturmfels (1994), and leads to an improved bound of the truncation degree, which coincides precisely with the classical Macaulay degree in the case of $n + 1$ equations in n unknowns. It also leads to a more efficient result when the polynomial system under consideration is sparse. Moreover, we also establish a tropical Positivstellensatz based on the same construction, allowing one to decide the inclusion of tropical basic semialgebraic sets. In particular, this reduces decision problems for tropical semi-algebraic sets to the solution of systems of tropical linear equalities and inequalities.

2.1 The Sparse Tropical Nullstellensatz

2.1.1 Statement of the theorem

The idea of the main theorem in this section is to reduce the problem of the existence of a solution to a system of polynomials equations to the existence of a solution to a system of tropical linear equations arising from a certain matrix called the Macaulay matrix, which can be constructed using the coefficients of the polynomials $f_1, \ldots, f_k$. From now on, we denote by f the collection $(f_1, \ldots, f_k)$ of polynomials, and by $f \ \nabla \ \mathbb{0}$ the system

$$\bigoplus_{\alpha \in \mathcal{A}_i} f_{i,\alpha} \odot x^{\odot \alpha} \ \nabla \ \mathbb{0} \quad \text{for all} \quad 1 \leqslant i \leqslant k$$

of tropical polynomial equations with unknown $x \in \mathbb{R}^n$.

We start by giving a proper setting to talk about tropical linear equations. We call *tropical matrix* a matrix with coefficients in $\mathbb{T}$. For two integers $p, q \in \mathbb{N}_{>0}$, the set of $p \times q$ tropical matrices is denoted by $\mathbb{T}^{p\times q}$. We can define tropical addition $\oplus$ and multiplication $\odot$ on tropical matrices by replacing the usual operations by their tropical version in the definition of the usual matrix operations. This notably gives a semiring structure to the set $\mathbb{T}^{p\times p}$.

Particularly, for $A = (a_{ij})_{(i,j)\in[p]\times[q]} \in \mathbb{T}^{p\times q}$ and $y = (y_j)_{j\in[q]} \in \mathbb{T}^q$, one has

$$A \odot y = \left(\max_{1\leqslant j\leqslant q} a_{ij} + y_j \right)_{i\in[p]} . \tag{2.1}$$

Definition 2.1.1. Let A be a $p \times q$ tropical matrix and let $y \in \mathbb{T}^q$. Then we write that $A \odot y \ \nabla \ \mathbb{0}$ whenever the maximum is attained twice for every coordinate in the righthandside of (2.1). The set of vectors $y \in \mathbb{T}^q$ such that $A \odot y \ \nabla \ \mathbb{0}$ is called the *tropical right null space* or *kernel* of the matrix A. Moreover, we set by convention that all vectors $y \in \mathbb{T}^q$ are in the tropical kernel of a $0 \times m$ matrix.

Remark 2.1.2. Note that as in the usual case, the tropical matrix equation $A \odot y \nabla \mathbb{0}$ can be written as the following q-variate tropical polynomial — linear in fact — system

$$\forall i \in [p], \quad \bigoplus_{j=1}^{q} a_{ij} \odot y_j \nabla \mathbb{0} \ .$$

Now, we define the tropical Macaulay matrix associated to the system f, which plays a crucial role in the determination of the solvability of a polynomial system, with no restriction on the number of equations, nor on the dimension of the resultant variety, as we will see in the next theorem.

Definition 2.1.3. Given a collection of tropical polynomials $f = (f_1, \dots, f_k)$, we define the *tropical Macaulay matrix* $\mathcal{M}$ of the system as such: the rows of $\mathcal{M}$ are indexed by pairs (i, α) where $1 \leqslant i \leqslant k$ and $\alpha \in \mathbb{Z}^n$, the columns of $\mathcal{M}$ are indexed by integer vectors $\beta \in \mathbb{Z}^n$, and for given (i, α) and β, we set the entry $\mathcal{M}_{(i,\alpha),\beta}$ of $\mathcal{M}$ equal to the coefficient of the monomial X^β in the polynomial $X^\alpha f_i(X)$, or $-\infty$ if no such monomial exists.

Given a Macaulay matrix $\mathcal{M}$ as above, a *nonempty* finite subset $\mathcal{E}$ of $\mathbb{Z}^n$, and a collection $\mathcal{A} = (\mathcal{A}_1, \dots, \mathcal{A}_k)$ of subsets of $\mathbb{Z}^n$, we denote by $\mathcal{M}_{\mathcal{E}}^{\mathcal{A}}$ the submatrix of $\mathcal{M}$ consisting only of the columns with indices $\beta \in \mathcal{E}$, and the rows indexed by pairs (i, α) where $1 \leqslant i \leqslant k$ and $\alpha \in \mathbb{Z}^n$ such that $\alpha + \mathcal{A}_i \subseteq \mathcal{E}$. When the polynomials f_i have their support equal to $\mathcal{A}_i$, and $\mathcal{M}$ is associated to $f = (f_1, \dots, f_k)$, we simply write $\mathcal{M}_{\mathcal{E}}$ instead of $\mathcal{M}_{\mathcal{E}}^{\mathcal{A}}$.

Remark 2.1.4. Note that, equivalently, $\mathcal{M}_{\mathcal{E}}$ is the submatrix of $\mathcal{M}$ consisting of the columns with indices $\beta \in \mathcal{E}$, and the rows that have all their finite entries in these columns. Moreover, if $\mathcal{E}$ is nonempty but too small, it might be possible that there are no such row of $\mathcal{M}$, and thus the set of rows of $\mathcal{M}_{\mathcal{E}}^{\mathcal{A}}$ might be empty, in which case by the convention of the previous definition, all vectors are considered to be in the tropical kernel of $\mathcal{M}_{\mathcal{E}}^{\mathcal{A}}$.

Now, let us denote as previously by $\mathcal{A} = (\mathcal{A}_1, \dots, \mathcal{A}_k)$ a collection of subsets of $\mathbb{Z}^n$, and for all $1 \leqslant i \leqslant k$, let Q_i be the convex hull of $\mathcal{A}_i$ and set $Q := Q_1 + \dots + Q_k$. Take a generic vector $\delta \in V + \mathbb{Z}^n$ where $V \subseteq \mathbb{R}^n$ is the vector space directing the affine hull of Q, and consider the set

$$\mathcal{E} := (Q + \delta) \cap \mathbb{Z}^n \ .$$

We will refer to sets of this form as *Canny-Emiris subsets* of $\mathbb{Z}^n$ associated to the collection $\mathcal{A}$. Note that for δ small enough, we always have the inclusion

$$\mathrm{relint}(Q) \cap \mathbb{Z}^n \subseteq \mathcal{E} \subseteq Q \cap \mathbb{Z}^n \ ,$$

where relint denotes the relative interior.

Now, for a collection $f = (f_1, \dots, f_k)$ of tropical polynomials, we shall consider in particular the collection $\mathcal{A} = (\mathcal{A}_1, \dots, \mathcal{A}_k)$ where $\mathcal{A}_i$ is the support of f_i for all $i \in [k]$. In that case, the set Q_i corresponds to the Newton polytope NP_{f_i} of f_i, and we shall also refer to Q as the *Newton polytope* of f. Also the Canny-Emiris subsets $\mathcal{E}$ associated to the collection $\mathcal{A}$ of supports are refered to as the *Canny-Emiris sets associated to* f.

The tropical linear system $\mathcal{M}_{\mathcal{E}} \odot y \nabla \mathbb{0}$ will be of interest to us in the resolution of the previous problem. More precisely, we have the following result, which will be proven in Section 2.1.3.

Theorem 2.1.5 (Sparse tropical Nullstellensatz)**.** *There exists a common root $x \in \mathbb{R}^n$ to the system $f(x) \nabla \mathbb{0}$ if and only if there exists a vector $y \in \mathbb{R}^{\mathcal{E}'}$ in the tropical kernel of the submatrix $\mathcal{M}_{\mathcal{E}'}$ of the Macaulay matrix $\mathcal{M}$ associated to the collection f —* i.e. *$\mathcal{M}_{\mathcal{E}'} \odot y \nabla \mathbb{0}$ — where $\mathcal{E}'$ is any subset of $\mathbb{Z}^n$ containing a nonempty Canny-Emiris subset $\mathcal{E}$ of $\mathbb{Z}^n$ associated to f.*

Moreover, if $\mathcal{E}' = \mathcal{E}$, these conditions are equivalent to the existence of a vector $y \in \mathbb{T}^{\mathcal{E}} \setminus \{\mathbb{0}\}$ such that $\mathcal{M}_{\mathcal{E}} \odot y \nabla \mathbb{0}$.

When we have no particular information on the supports of the polynomials f_i besides that they are ordinary tropical polynomials with respective degrees d_i, we denote by $\mathcal{M}_N$ the submatrix $\mathcal{M}_{\overline{\mathcal{E}}}^{\mathcal{A}}$ of $\mathcal{M}$ with $\overline{\mathcal{E}} = N\Delta \cap \mathbb{N}^n$, and $\mathcal{A}_i = d_i \Delta \cap \mathbb{N}^n$, where

$$\Delta := \{\alpha \in \mathbb{R}^n_{\geqslant 0} : |\alpha| = \alpha_1 + \dots + \alpha_n \leqslant 1\}$$

denotes the unit simplex. In this case, the integer N is called the *truncation degree* of the Macaulay submatrix $\mathcal{M}_N$.

More generally, if it is only known, for all $i = 1, \dots, k$, that the support of f_i is included in $\mathcal{A}_i$ (which plays now the role of an *a priori* support), then one shall consider a bigger set

$$\overline{\mathcal{E}} := Q \cap \mathbb{Z}^n \ ,$$

where $Q = \mathrm{conv}(\mathcal{A}_1) + \cdots + \mathrm{conv}(\mathcal{A}_k)$ is defined using the collection $\mathcal{A} = (\mathcal{A}_1, \ldots, \mathcal{A}_k)$. Recall that for δ small enough $\mathcal{E} \subseteq \overline{\mathcal{E}}$. In that case, one has the following theorem, in which we consider the matrices $\mathcal{M}^{\mathcal{A}}_{\mathcal{E}'}$ with $\mathcal{E}' \supseteq \overline{\mathcal{E}}$.

Theorem 2.1.6 (Nullstellensatz for sparse Tropical Polynomial Systems with *a priori* supports)**.** *There exists a common root $x \in \mathbb{R}^n$ to the system $f(x)\ \nabla\ \mathbb{0}$ if and only if there exists a vector $y \in \mathbb{R}^{\mathcal{E}'}$ in the tropical kernel of the submatrix $\mathcal{M}^{\mathcal{A}}_{\mathcal{E}'}$ of the Macaulay matrix $\mathcal{M}$ associated to the collection f, where $\mathcal{E}'$ is any subset of $\mathbb{Z}^n$ containing $\overline{\mathcal{E}}$ and $\mathcal{A}$ is a collection of a priori supports of the f_i.*

Moreover, when the Newton polytope of f has the same dimension as Q, one can replace $\overline{\mathcal{E}}$ by any nonempty Canny-Emiris set $\mathcal{E}$ associated to $\mathcal{A}$.

Proof. The inclusions $\mathcal{E} \subseteq \mathcal{E}'$ and $\mathrm{supp}(f_i) \subseteq \mathcal{A}_i$ for all $1 \leqslant i \leqslant k$ imply that the matrix $\mathcal{M}_{\mathcal{E}}$ is a submatrix of the matrix $\mathcal{M}^{\mathcal{A}}_{\mathcal{E}'}$. More precisely, if (i, α) is the index of a row of $\mathcal{M}_{\mathcal{E}}$, then this indicates that the support of polynomial $x^\alpha f_i$ is included in $\mathcal{E}$, and therefore

$$\alpha \in \mathrm{conv}\left(\sum_{1 \leqslant j \neq i \leqslant n} \mathrm{supp}(f_j)\right) \subseteq \mathrm{conv}\left(\sum_{1 \leqslant j \neq i \leqslant n} \mathcal{A}_j\right) ,$$

which shows that (i, α) is also the index of a row of the matrix $\mathcal{M}^{\mathcal{A}}_{\mathcal{E}'}$. The theorem then follows directly from Theorem 2.1.5 and from the latter remark. $\square$

When $f = (f_1, \ldots, f_k)$ is a collection of ordinary tropical polynomials f_i with respective degree d_i and the matrix $\mathcal{M}_N$ is defined as above, applying Theorem 2.1.6 with $\mathcal{A}_i = d_i\Delta \cap \mathbb{N}^n$, we deduce the following result.

Corollary 2.1.7. *There exists a common root $x \in \mathbb{R}^n$ to the system $f(x)\ \nabla\ \mathbb{0}$ if and only if there exists a vector $y \in \mathbb{R}^{\mathcal{E}'}$ such that $\mathcal{M}_N \odot y\ \nabla\ \mathbb{0}$, where $\mathcal{E}' = N\Delta \cap \mathbb{N}^n$ and $N \geqslant d_1 + \cdots + d_k$. Moreover, when the Newton polytope of f has full dimension, one can replace the above lower bound on N by $N \geqslant d_1 + \cdots + d_k - n$.*

Proof. This result is obtained by applying Theorem 2.1.6 with $\mathcal{A}_i = d_i\Delta \cap \mathbb{N}^n$, and $\mathcal{E}' = N\Delta \cap \mathbb{N}^n$, since $\overline{\mathcal{E}} = (d_1 + \cdots + d_k)\Delta \cap \mathbb{N}^n \subset \mathcal{E}'$. Moreover, for the full-dimensional case, one can perturb the simplex $(d_1 + \cdots + d_k)\Delta$ by a perturbation $(\varepsilon, \ldots, \varepsilon)$ for $\varepsilon > 0$ small enough and the resulting Canny-Emiris set is $(d_1 + \cdots + d_k - n)\Delta \cap \mathbb{N}^n$. $\square$

Example 2.1.8. Let us illustrate Theorem 2.1.5 with some explicit examples. Consider the following two systems:

$$(\mathcal{S}_1) : \begin{cases} f_1 &= 1 \oplus 2x_1 \oplus 1x_2 \oplus 1x_1x_2 \\ f_2 &= 0 \oplus 0x_1 \oplus 1x_2 \\ f_3 &= 2x_1 \oplus 0x_2 \end{cases} \quad \text{and} \quad (\mathcal{S}_2) : \begin{cases} f_1 &= 1 \oplus 4x_1 \oplus 1x_2 \oplus 3x_1x_2 \\ f_2 &= 0 \oplus 0x_1 \oplus 1x_2 \\ f_3 &= 2x_1 \oplus 0x_2 \end{cases}$$

Both systems have the same supports, and thus yield the same polytope Q.

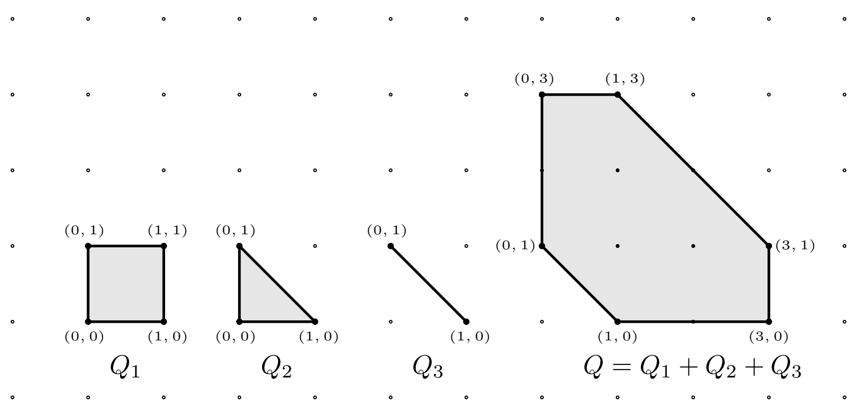

Figure 2.1: The Newton polytopes associated to f_1, f_2, f_3 and their Minkowski sum.

For this collection of supports, one can take $\delta = (-1+\varepsilon, -1+\varepsilon)$ with $\varepsilon > 0$ sufficiently small, for instance $\varepsilon = 0.1$, which gives us the Canny-Emiris set

$$\mathcal{E} := (Q+\delta) \cap \mathbb{Z}^n = \{(0,0),(1,0),(0,1),(2,0),(1,1),(0,2)\} \ ,$$

corresponding to the set of monomials

$$\{1, x_1, x_2, x_1^2, x_1x_2, x_2^2\} \ .$$

Note that we could simply have chosen δ to be $(\varepsilon,\varepsilon)$ instead, but we chose to also translate Q by $(-1,-1)$ so that $Q+\delta$ contains the origin, thus ensuring that $\mathcal{E}$ contains smaller degree monomials. We can therefore write the respective submatrices $\mathcal{M}_{\mathcal{E}}^{(1)}$ and $\mathcal{M}_{\mathcal{E}}^{(2)}$ of the Macaulay matrix associated to the set $\mathcal{E}$ and we obtain the following 7×6 matrices

$$\mathcal{M}_{\mathcal{E}}^{(1)} = \begin{array}{c} \\ f_1 \\ f_2 \\ x_1f_2 \\ x_2f_2 \\ f_3 \\ x_1f_3 \\ x_2f_3 \end{array} \begin{array}{c} \begin{array}{cccccc} 1 & x_1 & x_2 & x_1^2 & x_1x_2 & x_2^2 \end{array} \\ \begin{pmatrix} 1 & 2 & 1 & & 1 & \\ 0 & 0 & 1 & & & \\ & 0 & & 0 & 1 & \\ & & 0 & & 0 & 1 \\ & 2 & 0 & & & \\ & & & 2 & 0 & \\ & & & & 2 & 0 \end{pmatrix} \end{array} \quad \text{and} \quad \mathcal{M}_{\mathcal{E}}^{(2)} = \begin{array}{c} \\ f_1 \\ f_2 \\ x_1f_2 \\ x_2f_2 \\ f_3 \\ x_1f_3 \\ x_2f_3 \end{array} \begin{array}{c} \begin{array}{cccccc} 1 & x_1 & x_2 & x_1^2 & x_1x_2 & x_2^2 \end{array} \\ \begin{pmatrix} 1 & 4 & 1 & & 3 & \\ 0 & 0 & 1 & & & \\ & 0 & & 0 & 1 & \\ & & 0 & & 0 & 1 \\ & 2 & 0 & & & \\ & & & 2 & 0 & \\ & & & & 2 & 0 \end{pmatrix} \end{array} .$$

One can check that the system $(\mathcal{S}_1)$ does not have a common root, as the different intersection points of the tropical hypersurfaces associated to f_1, f_2 and f_3 are listed on Figure 2.2 (and we know that two tropical lines only have at most one intersection point, and a line and a quadric have at most two intersection points).

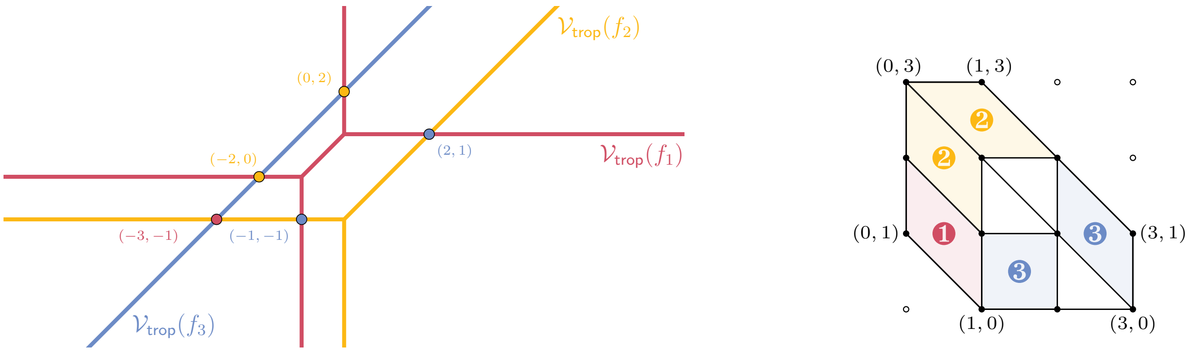

Figure 2.2: The arrangement of tropical hypersurfaces of the polynomials from the system $(\mathcal{S}_1)$ and the associated subdivision of Q.

With a similar argument to the first system, one can check that $(-3,-1)$ is the only common root of the system $(\mathcal{S}_2)$ (see Figure 2.3), and indeed by choosing

$$y = \operatorname{ver}(-3,-1) = \begin{pmatrix} 0 \\ -3 \\ -1 \\ -6 \\ -4 \\ -2 \end{pmatrix} \ ,$$

we observe that

$$\mathcal{M}_{\mathcal{E}}^{(2)} \odot y \,\nabla\, \mathbb{0} \ .$$

Moreover, note that the set of solutions $y \in \mathbb{R}^6$ to the tropical linear system $\mathcal{M}_{\mathcal{E}}^{(2)} \odot y \,\nabla\, \mathbb{0}$ consists precisely in the set $\{\lambda + \operatorname{ver}(-3,-1) : \lambda \in \mathbb{R}\}$ of tropical multiples of the Veronese embedding of the point $(-3,-1)$, which indeed attests to the uniqueness of the solution $(-3,-1)$, as two distinct solutions would have two non-collinear Veronese embeddings, in the tropical sense.

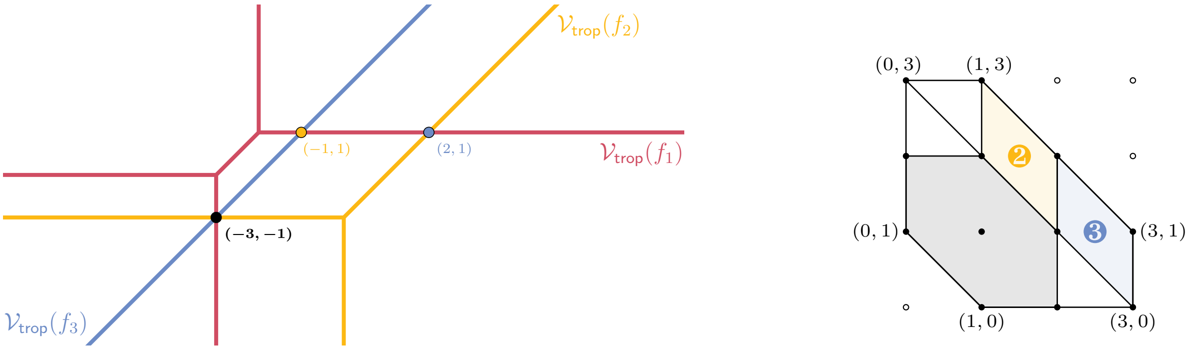

Figure 2.3: The arrangement of tropical varieties of the polynomials from the system $(\mathcal{S}_2)$ and the associated subdivision of Q.

Remark 2.1.9. This improves on Grigoriev and Podolskii's Tropical Dual Nullstellensatz from [GP18, Theorem 3.3 (i)], which requires $N \geqslant (n+2)(d_1 + \cdots + d_k)$. Moreover, under the condition that the Newton polytope of f is full-dimensional, and when $k = n+1$, we retrieve the classical Macaulay bound $N \geqslant d_1 + \cdots + d_{n+1} - n$ (see [Laz81, Laz83, Giu84]).

In [GP18, §4.6], the authors provide for all degree $d \geqslant 2$ and all number $n \geqslant 2$ of variables the following family of $n+1$ polynomials of degree at most d

$$\begin{array}{rcl} f_1 &=& 0 \oplus 0x_1 \\ f_i &=& 0x_{i-1}^d \oplus 0x_i, \quad 2 \leqslant i \leqslant n \\ f_{n+1} &=& 0 \oplus 1x_n \end{array}$$

and show that the linearized system with truncation degree $N = (n-1)(d-1)$

$$\mathcal{M}_{(n-1)(d-1)} \odot y \mathrel{\nabla} \mathbb{0}$$

has a solution in $\mathbb{R}^{\binom{N+n}{n}}$ while the polynomial system does not have a solution in $\mathbb{R}^n$, showing that our improved bound is tight, as in this example, our bound yields

$$d_1 + \cdots + d_{n+1} - n = 1 + (n-1)d + 1 - n = (n-1)(d-1) + 1 \ .$$

The previous system is inspired by the Masser-Philippon example (÷Crefexpl:masser-philippon, see also [GV01]) of a system of n degree d polynomials in n variables for which the minimal truncation degree needed in Hilbert's Nullstellensatz is bounded below by $(d-1)d^{n-1}$. Hence, in the classical case, the truncation degree appearing in the Nullstellensatz may be exponential in n, whereas in the tropical case, it is bounded in terms of the *sum* of degrees of the polynomials.

However, in non-square cases, the lower bound in Corollary 2.1.7 is not necessarily optimal.

For instance, in the case of $k > n+1$ degree one polynomials, for all $1 \leqslant i \leqslant k$, the tropical polynomial function associated to f_i is simply a tropical affine function

$$(x_1, \ldots, x_n) \mapsto f_{i0} \oplus f_{i1}x_1 \oplus \cdots \oplus f_{in}x_n \ ,$$

and thus $x = (x_1, \ldots, x_n) \in \mathbb{R}^n$ is a common root of $f_1, \ldots, f_k$ if and only if

$$\begin{pmatrix} f_{10} & f_{11} & \cdots & f_{1n} \\ \vdots & \vdots & & \vdots \\ f_{k0} & f_{k1} & \cdots & f_{kn} \end{pmatrix} \odot \begin{pmatrix} 1 \\ x_1 \\ \vdots \\ x_n \end{pmatrix} \mathrel{\nabla} \mathbb{0} \ ,$$

and thus the collection of polynomials f has a common root if and only if the matrix $(f_{ij})_{\substack{1 \leqslant i \leqslant k \\ 0 \leqslant j \leqslant n}}$ has an element $(y_0, \ldots, y_n) \in \mathbb{R}^{n+1}$ in its right null space, in which case $(y_1 - y_0, \ldots, y_n - y_0) \in \mathbb{R}^n$ is a common root of f. But this matrix corresponds to the submatrix of the Macaulay matrix $\mathcal{M}$ obtained by taking $N = 1$ as the truncation degree for the Macaulay matrix, while the previous bound gives $N = d_1 + \cdots + d_k - n = k - n$.

Remark 2.1.10. Although Theorem 2.1.5 only deals with the toric case, *i.e.* only accounts for solutions in $x \in \mathbb{R}^n$, one can still use it to deal with the non-toric case and find solutions in $x \in \mathbb{T}^n$: for any subset I of $\{1, \dots, n\}$, if $x \in \mathbb{T}^n$ is a solution of a tropical polynomial system such that

$$\begin{cases} x_i \neq \mathbb{0} & \text{for all } i \in I \\ x_i = \mathbb{0} & \text{for all } i \in J := \{1, \dots, n\} \setminus I \end{cases},$$

then $x_I := (x_i)_{i \in I} \in \mathbb{R}^I$ is a root of the tropical polynomial system obtained by removing all the monomials in which the variables X_j for $j \in J$ appear.

Note however that in [GP18, Theorems 3.3 (ii) and 4.20], it is shown that the linearization remains valid with $-\infty$ but at the price of an exponential blow up of the truncation degree, which becomes $N = 2(n+2)^2 k(4d)^{\min(n,k)+2}$, and thus for practical applications, enumerating the 2^n possible supports of a solution leads to a faster method.

Remark 2.1.11. The assumption that the considered Canny-Emiris set $\mathcal{E}$ is nonempty is needed because it is possible to find systems both with and without a common root for which the empty set is a Canny-Emiris subset associated to the system. For instance for $n = 3$ and $k = 2$, consider the system

$$(S_1) : \begin{cases} f_1 &= 0 \oplus 0x_1 \oplus 0x_2 \oplus 0x_3 \\ f_2 &= 0 \end{cases}.$$

In this case, Q is simply the tetrahedron with vertices $(0,0,0)$, $(0,0,1)$, $(0,1,0)$ and $(1,0,0)$. Now if we take $\delta = (\varepsilon, \varepsilon, \varepsilon)$ for $\varepsilon > 0$ small enough, we obtain a Canny-Emiris set $\mathcal{E} := (Q + \delta) \cap \mathbb{Z}^n$ which is empty, as illustrated in Figure 2.4. Of course, the system (S_1) does not have a common root since f_2 is a constant.

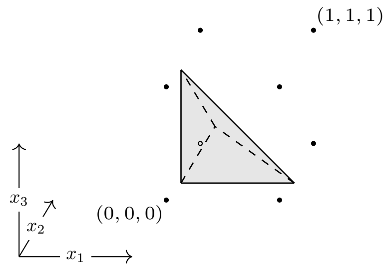

Figure 2.4: The polytope $Q + \delta$ with $\varepsilon = 0.1$ for the system (S_1).

Now, still for $n = 3$ and $k = 2$, consider the system

$$(S_2) : \begin{cases} f_1 &= 0 \oplus 0x_1 \oplus 0x_3 \\ f_2 &= 0 \oplus 0x_2 \end{cases}.$$

Now, Q is the triangular prism with vertices $(0,0,0)$, $(1,0,0)$, $(0,0,1)$, $(0,1,0)$, $(1,1,0)$ and $(0,1,1)$. Once again, if we take $\delta = (\varepsilon, \varepsilon, \varepsilon)$ for $\varepsilon > 0$ small enough, we obtain a Canny-Emiris set which is empty, as illustrated in Figure 2.5. This time however, the system (S_2) has a common root, namely $(0,0,0)$.

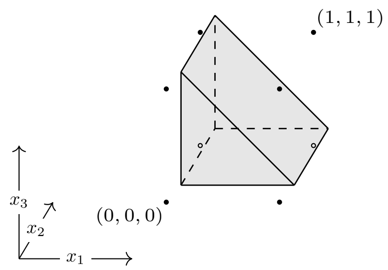

Figure 2.5: The polytope $Q + \delta$ with $\varepsilon = 0.1$ for the system (S_2).

Therefore, it is *a priori* not possible to conclude if the considered Canny-Emiris set is taken empty, although excluding the case where some of the polynomials of the system are monomials is rather degenerate. Excluding this particular case, one can wonder if it is then possible to reach a conclusion in the case of systems where there exists a empty associated Canny-Emiris set.

2.1.2 Preliminary results

In this section, we state and prove a number of lemmas which will be used in order to prove Theorem 2.1.5.

Nonsingular and diagonally dominant tropical matrices

We first recall the definition of nonsingularity and diagonal dominance for tropical matrices. Set two integers $p, q \in \mathbb{N}_{>0}$.

Definition 2.1.12. Let $A = (a_{ij})_{(i,j)\in[p]\times[q]}$ be a $p \times q$ tropical matrix. Then the matrix A is said to be *tropically nonsingular* whenever the only solution to the equation $A \odot y \nabla \mathbb{0}$ of unknown $y \in \mathbb{T}^q$ is $y = \mathbb{0}$.

Remark 2.1.13. In the case where $A \in \mathbb{T}^{p\times p}$ is a square matrix, one can consider its *tropical determinant* $\operatorname{tdet}(A)$ which is given by

$$\operatorname{tdet}(A) = \max_{\sigma\in\mathfrak{S}_p} a_{1\sigma(1)} + \cdots + a_{p\sigma(p)} \enspace .$$

If the maximum in the previous expression is attained exactly once, hence if $\operatorname{tdet}(A) \not\nabla \mathbb{0}$, then as a direct consequence of Corollary 6.12 of [AGG09], the only possible solution to the equation $A \odot y \nabla \mathbb{0}$ is $y = \mathbb{0}$, *i.e.* the matrix A is tropically nonsingular.

Definition 2.1.14. A matrix $A = (a_{ij})_{(i,j)\in[p]\times[p]}\mathbb{T}^{p\times p}$ is said to be *weakly diagonally dominant* in the tropical sense whenever we have the inequalities

$$a_{ii} \geqslant a_{ij} \quad \text{for all} \quad 1 \leqslant i, j \leqslant p \text{ such that } i \neq j \enspace ,$$

and we say that it is *(strictly) diagonally dominant* if these inequalities are strict.

Remark 2.1.15. This notion of tropical diagonal dominance just corresponds to the tropical version of classical diagonal dominance, as the inequality of the above definition is equivalent to

$$a_{ii} \geqslant \bigoplus_{1\leqslant j\neq i\leqslant p} a_{ij} = \max_{1\leqslant j\neq i\leqslant p} a_{ij} \quad \text{for all} \quad 1 \leqslant i \leqslant p \enspace .$$

A notable fact about diagonally dominant tropical matrices, which will play a crucial role in the proof of Theorem 2.1.5, is the fact that similarly to classical diagonally dominant matrices, these matrices are non-singular. More precisely, we have the following result.

Lemma 2.1.16. *Let $A \in \mathbb{T}^{p\times p} = (a_{ij})_{(i,j)\in[p]\times[p]}$ be a diagonally dominant tropical matrix. Then A is tropically nonsingular.*

Proof. Let $y = (y_1, \dots, y_p) \in \mathbb{T}^p$ be such that $A \odot y \nabla \mathbb{0}$, and consider $1 \leqslant i \leqslant p$ such that $y_i = \max_{1\leqslant j\leqslant p} y_j$. Then from the relation $A \odot y \nabla \mathbb{0}$, it follows in particular that the maximum in the expression

$$\max_{1\leqslant j\leqslant p} (a_{ij} + y_j)$$

is attained twice, but since for all $1 \leqslant j \leqslant p$, we have

$$a_{ii} > a_{ij} \quad \text{and} \quad y_i \geqslant y_j \enspace ,$$

the only possible way such that the maximum in the previous expression is attained twice is that

$$y_i = -\infty$$

and thus

$$y = \mathbb{0} \enspace . \qquad \square$$

Remark 2.1.17. Alternatively, one can retrieve the previous lemma with the following argument: since A is diagonally dominant, it means that the maximum in the expression

$$\operatorname{tdet}(A) = \max_{\sigma\in\mathfrak{S}_d} a_{1\sigma(1)} + \cdots + a_{d\sigma(d)}$$

is attained exactly once, hence $\operatorname{tdet}(A) \not\nabla \mathbb{0}$, and thus by the previous remark, A is tropically nonsingular.

Finally we will also make use of the following two lemmas.

Lemma 2.1.18. *Let $A = (a_{ij})_{(i,j)\in[p]\times[q]}$ be a $p \times q$ tropical matrix. Fix for $1 \leqslant j \leqslant q$, $\varepsilon_j \in \mathbb{R}$, and set $\widetilde{A} = (\widetilde{a}_{ij})_{(i,j)\in[p]\times[q]} \in \mathbb{T}^{p\times q}$ with $\widetilde{a}_{ij} = a_{ij} + \varepsilon_j$ for all $1 \leqslant i \leqslant p$ and $1 \leqslant j \leqslant m$. Then A is tropically nonsingular if and only if $\widetilde{A}$ is tropically nonsingular.*

Proof. Assume that A is nonsingular, and let $\widetilde{y} = (\widetilde{y}_j)_{1\leqslant j\leqslant q} \in \mathbb{T}^q$ be such that $\widetilde{A} \odot \widetilde{y} \mathrel{\nabla} \mathbb{0}$. Then, this means that for all $1 \leqslant i \leqslant p$, the maximum in the expression

$$\max_{1\leqslant j\leqslant q} (\widetilde{a}_{ij} + \widetilde{y}_j) = \max_{1\leqslant j\leqslant q} (a_{ij} + (\widetilde{y}_j + \varepsilon_j))$$

is attained twice. In other words, setting $y = (\widetilde{y}_j + \varepsilon_j)_{1\leqslant j\leqslant q}$, we obtain that $A \odot y \mathrel{\nabla} \mathbb{0}$. Therefore, by nonsingularity of A, we must have $y = \mathbb{0}$, and since ε_j is finite for all $1 \leqslant j \leqslant q$, this implies that $\widetilde{y} = \mathbb{0}$, hence $\widetilde{A}$ is tropically nonsingular.

As for the converse implication, it is obtained by swapping A and $\widetilde{A}$, and by changing $(\varepsilon_j)_{1\leqslant j\leqslant q}$ to $(-\varepsilon_j)_{1\leqslant j\leqslant m}$. ☐

Lemma 2.1.19. *Let A be a $p \times q$ tropical matrix, and assume that A can be written by block as a lower-triangular matrix*

$$A = \begin{pmatrix} A^{(m)} & \mathbb{0} \\ * & * \end{pmatrix}$$

with $A^{(m)}$ a $m \times m$ square submatrix with $0 < m \leqslant p, q$. Moreover, assume that $A^{(m)}$ is tropically nonsingular. Then the equation $A \odot y \mathrel{\nabla} \mathbb{0}$ of unknown y has no solution in $\mathbb{R}^q$.

Proof. Let $y = (y_j)_{1\leqslant j\leqslant q} \in \mathbb{T}^q$ be such that $A \odot y \mathrel{\nabla} \mathbb{0}$. Then by setting $y^{(m)} = (y_j)_{1\leqslant j\leqslant m} \in \mathbb{T}^m$, we obtain in particular, by looking at the first m rows of the product $A \odot y$, that $A^{(m)} \odot y^{(m)} \mathrel{\nabla} \mathbb{0}$, which implies by nonsingularity of $A^{(m)}$ that $y^{(m)} = \mathbb{0}$ and thus y does not belong to $\mathbb{R}^q$. ☐

Generalities on sup-convolution and Minkowski sums

In order to write the proof of Theorem 2.1.5, we also need to introduce the following definition, which simply corresponds to the tropical equivalent of the convolution product. The *support* of a function $h : \mathbb{R}^n \to \mathbb{R} \cup \{\pm\infty\}$ is defined by $\operatorname{supp}(h) = \operatorname{cl}(\{x \in \mathbb{R}^n : h(x) > -\infty\})$. The *hypograph* of h is the set $\operatorname{hypo}(h) = \{(x,t) \in \mathbb{R}^n \times \mathbb{R} : t \leqslant h(x)\}$.

Definition 2.1.20. The *sup-convolution* is the binary operator $\square$ defined for all functions $f, g : \mathbb{R}^n \to \mathbb{R} \cup \{\pm\infty\}$ by

$$f \mathbin{\square} g(x) = \sup_{y+z=x} f(y) + g(z) \ ,$$

with the convention $(-\infty) + (+\infty) = -\infty$.

In particular, if f and g are upper semicontinuous, take values in $\mathbb{R} \cup \{-\infty\}$, and have compact support, then, the supremum in the expression of $f \mathbin{\square} g(x)$ is achieved and $f \mathbin{\square} g$ also has compact support. The operations of sup-convolution and Minkowski sum are commutative and associative, and that we have the following immediate properties:

Property 2.1.21. Let $E_1, \dots, E_\ell$ be a collection of subsets of $\mathbb{R}^n$ and let $h_1, \dots, h_\ell$ be a family of upper semicontinuous functions with compact support from $\mathbb{R}^n$ to $\mathbb{R} \cup \{\pm\infty\}$. Let $E = E_1 + \cdots + E_\ell$ and $h = h_1 \mathbin{\square} \cdots \mathbin{\square} h_\ell$. Then,

(a) For all $q \in \mathbb{R}^n$, $h(q) = \max_{q_1+\cdots+q_k=q} h_1(q_1) + \cdots + h_\ell(q_\ell)$;

(b) $\operatorname{hypo}(h) = \operatorname{hypo}(h_1) + \cdots + \operatorname{hypo}(h_\ell)$ and $\operatorname{supp}(h) = \operatorname{supp}(h_1) + \cdots + \operatorname{supp}(h_\ell)$;

(c) For $1 \leqslant i \leqslant \ell$, let $\widehat{h}_i := h_1 \mathbin{\square} \cdots \mathbin{\square} h_{i-1} \mathbin{\square} h_{i+1} \mathbin{\square} \cdots \mathbin{\square} h_\ell$. Then we have $h_i \mathbin{\square} \widehat{h}_i = h$.

Definition 2.1.22. A concave function $\mathbb{R}^n \to \mathbb{R} \cup \{-\infty\}$ is said to be *polyhedral* if its hypograph is a (closed) polyhedron.

Example 2.1.23. The main example of concave polyhedral functions that will be of interest in this paper are the functions obtained by taking the concavification of the coefficient map $\omega : \mathbb{R}^n \to \mathbb{R}\cup\{-\infty\}$ — that is the infimum of all concave functions greater than or equal to ω — defined by

$$\omega(\alpha) = \begin{cases} f_\alpha & \text{if} \quad \alpha \in \operatorname{supp}(f), \\ \mathbb{0} & \text{otherwise.} \end{cases}$$

of a tropical polynomial f. These functions satisfy in particular the property that the projection of the singularities of their graph onto $\mathbb{R}^n \times \{0\}$ is a rational polyhedral complex. Moreover, the projection of vertices of their hypograph are elements of $\operatorname{supp}(f)$.

For all functions $h : \mathbb{R}^n \to \mathbb{R} \cup \{-\infty\}$, and for all $x \in \mathbb{R}^n$, we set

$$\mathcal{C}(x, h) := \operatorname*{arg\,max}_{q\in\mathbb{R}^n}(\langle q, x\rangle + h(q)) \quad \text{and} \quad \mathcal{F}(x, h) := \{(q, h(q)) : q \in \mathcal{C}(x, h)\} \ .$$

We make the following crucial observations.

Observation 2.1.24.

(a) If h is a concave polyhedral function nonidentically $-\infty$, then $\mathcal{F}(x, h)$ is the face of the hypograph of h, which is obtained as the intersection of this hypograph with a supporting hyperplane of outer normal vector $(x, 1)$. In particular, this face is non-vertical, and hence it is a proper face of $\operatorname{hypo}(h)$, *i.e.* $\mathcal{F}(x, h) \subsetneq \operatorname{hypo}(h)$.

(b) If h is the concavification of the coefficient map ω as in Example 2.1.23, then $\mathcal{C}(x, h)$ is the convex hull of

$$\mathcal{C}(x, \omega) = \operatorname*{arg\,max}_{\alpha\in\mathbb{Z}^n}(f_\alpha + \langle x, \alpha\rangle) \ ,$$

which coincides with the intersection of $\mathcal{C}(x, h)$ with the elements $\alpha \in \operatorname{supp}(f)$ such that $h(\alpha) = f_\alpha$. Then, $\mathcal{F}(x, h)$ is the convex hull of $\mathcal{F}(x, \omega)$. We also have that $\mathcal{F}(x, h)$ is the convex hull of its intersection with the set of vertices of $\operatorname{hypo}(h)$.

(c) Moreover, when h is the concavification of the coefficient map ω, if F is a non-vertical face of $\operatorname{hypo}(h)$, then $F = \mathcal{F}(x, h)$ if and only if $(x, 1)$ is in the relative interior of $N_F(\operatorname{hypo}(h))$. In particular, if F is a facet of $\operatorname{hypo}(h)$, then there exists a unique vector $x \in \mathbb{R}^n$ such that $F = \mathcal{F}(x, h)$ and x is in the vector space directing the affine hull of $Q := \operatorname{supp}(h)$. Indeed, from Corollary 1.1.18, we have $\dim(N_F(\operatorname{hypo}(h)) \cap W) = 1$ where $W = V \times \mathbb{R}$ is the vector space directing the affine hull of $\operatorname{hypo}(h)$, and thus if $(x, 1)$ and $(x', 1)$ are both in the half-line $N_F(\operatorname{hypo}(h)) \cap W$, then it follows that $x' = x$.

We now state a useful lemma on convex polyhedra.

Lemma 2.1.25. *Let $P_1, \dots, P_\ell$ be a finite collection of convex polyhedra, and denote by P their Minkowski sum $P_1 + \cdots + P_\ell$. Let F be a face of P and let y be in the relative interior of $N_F(P)$. If $p = p_1 + \cdots + p_\ell \in F$ with $p_i \in P_i$ for all $1 \leqslant i \leqslant \ell$, then for all $1 \leqslant i \leqslant \ell$, $y \in N_{p_i}(P_i)$, i.e. $p_i \in \operatorname{arg\,max}_{p'_i\in P_i}\langle p'_i, y\rangle$.*

Proof. Saying that y is in the relative interior of $N_F(P)$ is equivalent to saying that $F = \operatorname{arg\,max}_{p'\in P}\langle p', y\rangle$. Therefore, you have

$$\begin{aligned} \langle p_1, y\rangle + \cdots + \langle p_\ell, y\rangle &= \langle p, y\rangle \\ &= \max\nolimits_{p'\in P}\langle p', y\rangle \\ &= \max\nolimits_{p'_1\in P_1,\dots,p'_\ell\in P_\ell}(\langle p'_1, y\rangle + \cdots + \langle p'_\ell, y\rangle) \\ &= \max\nolimits_{p'_1\in P_1}\langle p'_1, y\rangle + \cdots + \max\nolimits_{p'_\ell\in P_\ell}\langle p'_\ell, y\rangle \ , \end{aligned}$$

therefore $\langle p_i, y\rangle = \max_{p'_i\in P_i}\langle p'_i, y\rangle$ for all $1 \leqslant i \leqslant \ell$, *i.e.* $y \in N_{p_i}(P_i)$ for all $1 \leqslant i \leqslant \ell$. ☐

Remark 2.1.26. In particular, from the previous lemma, by setting $F_i := \operatorname{arg\,max}_{p'_i\in P_i}\langle p'_i, y\rangle$, we obtain that the decomposition of F as a sum of faces of the P_i is precisely $F = F_1 + \cdots + F_\ell$.

The following result is a direct corollary of the previous lemma.

Corollary 2.1.27. *Let $h_1, \dots, h_\ell$ denote concave non-identically $-\infty$ polyhedral functions from $\mathbb{R}^n$ to $\mathbb{R} \cup \{-\infty\}$. Consider the sup-convolution*

$$h = h_1 \,\square \cdots \square\, h_\ell \enspace .$$

Let F be a non-vertical face of $\operatorname{hypo}(h)$, and let $x \in \mathbb{R}^n$ be such that $(x, 1)$ is in the relative interior of the normal cone of $\operatorname{hypo}(h)$ at this face. Consider a point $(q, \lambda) \in F$ such that

$$(q, \lambda) = (q_1, \lambda_1) + \cdots + (q_\ell, \lambda_\ell) \quad \textit{with} \quad (q_i, \lambda_i) \in \operatorname{hypo}(h_i) \textit{ for all } 1 \leqslant i \leqslant \ell \enspace .$$

Then for all $1 \leqslant i \leqslant \ell$, $(q_i, \lambda_i) \in \mathcal{F}(x, h_i)$.

In particular,

$$F = \mathcal{F}(x, h) = \mathcal{F}(x, h_1) + \cdots + \mathcal{F}(x, h_\ell)$$

and moreover, for all $1 \leqslant i \leqslant \ell$, $(x, 1)$ is in the normal cone of $\operatorname{hypo}(h_i)$ at point (q_i, λ_i).

Proof. We apply the previous lemma for $P_i = \operatorname{hypo}(h_i)$ for all $1 \leqslant i \leqslant \ell$ and $P = \operatorname{hypo}(h)$. Since $(x, 1)$ is in the relative interior of $N_F(P)$, then we have $F = \mathcal{F}(x, h)$ by definition, and moreover from the previous lemma, we have $(x, 1) \in N_{(q_i, \lambda_i)}(P_i)$, or equivalently $(q_i, \lambda_i) \in \mathcal{F}(x, h_i)$ for all $1 \leqslant i \leqslant \ell$. The remainder follows immediately. $\square$

2.1.3 Proving the Tropical Nullstellensatz

The Canny-Emiris construction

In the following, we assume given the polynomials $f_1, \dots, f_k$, a nonempty Canny-Emiris subset $\mathcal{E}$ associated to the system f, and a subset $\mathcal{E}'$ of $\mathbb{Z}^n$ containing $\mathcal{E}$. We can now describe the construction of Canny and Emiris from [CE93] and [Emi05], which was generalized by Sturmfels in part 3 of [Stu94], which we will apply to the system $f = (f_1, \dots, f_k)$ in the particular case where the polynomials f_i do not share a common root, in order to prove Theorem 2.1.5. In particular, contrary to the previous constructions, we do not assume that the coefficients satisfy any genericity condition. This construction is illustrated on an example in Example 2.1.32 below.

First of all, let us settle some notations for the rest of this section. For all $1 \leqslant i \leqslant k$, let $\mathcal{A}_i$ denote the support of f_i, Q_i the convex hull of $\mathcal{A}_i$, $Q := Q_1 + \cdots + Q_k$. The Canny-Emiris subset $\mathcal{E}$ of $\mathbb{Z}^n$ is such that $\mathcal{E} := \mathbb{Z}^n \cap (Q + \delta)$, for some generic vector $\delta \in V + \mathbb{Z}^n$, where $V \subseteq \mathbb{R}^n$ is the vector space directing the affine hull of Q. Let $\omega = (\omega_i)_{1 \leqslant i \leqslant k}$ be the collection of coefficient maps of the f_i, that is defined by

$$\omega_i(\alpha) = \begin{cases} f_{i,\alpha} & \text{if} \quad \alpha \in \mathcal{A}_i \\ \mathbb{0} & \text{else,} \end{cases}$$

and let $h_i : \mathbb{R}^n \to \mathbb{R} \cup \{-\infty\}$ denote the concave hull of ω_i. Then we can consider the following liftings of the polytopes $Q, Q_1, \dots, Q_k$. Let $Q_i^{\mathsf{lift}} := \operatorname{hypo}(h_i)$ (this is refered to by Grigoriev and Podolskii in [GP18] as the *extended Newton polytope* of f_i) and $Q^{\mathsf{lift}} := Q_1^{\mathsf{lift}} + \cdots + Q_k^{\mathsf{lift}} = \operatorname{hypo}(h)$ with $h := h_1 \,\square \cdots \square\, h_k$.

Observation 2.1.28. Note that for all $1 \leqslant i \leqslant k$, by construction of the maps h_i, we have

$$h_i(\alpha) \geqslant f_{i,\alpha} \quad \forall \alpha \in \mathcal{A}_i \enspace ,$$

with equality whenever the monomial $f_{i,\alpha} X^\alpha$ is *essential* in f_i, that is whenever there exists a point $x \in \mathbb{R}^n$ at which the monomial of exponent α is the only one achieving the maximum over all the monomials of f_i. Equivalently, the monomial $f_{i,\alpha} X^\alpha$ of f_i is essential if, and only if, the point $(\alpha, f_{i,\alpha})$ is a vertex of $\operatorname{hypo}(h_i)$. Note also that if $(q, h(q))$ is an extreme point of Q^{lift}, then $h(q)$ corresponds to the coefficient of X^q in the product $f_1 \cdots f_k$.

Now let us apply the Canny-Emiris construction, to the collection ω of maps — note that as opposed to what is done in [Stu94], ω is given and might not be generic. The projection of Q^{lift} onto Q induces a mixed coherent subdivision Δ_ω of Q, given by the points of non-differentiability of h. The following observation follows readily from the genericity of δ.

Observation 2.1.29. Given $p \in \mathcal{E}$, the couple $(p - \delta, h(p - \delta))$ lies in the relative interior of a unique non-vertical facet F of Q^{lift}, or equivalently, $p - \delta$ lies in the relative interior of the associated cell C of Q, obtained as the image of F by the projection mapping. Moreover, there is a unique $x \in V$ such that $F = \mathcal{F}(x, h)$ and $C = \mathcal{C}(x, h)$.

We construct a matrix $\mathcal{M}_{\mathcal{E}\mathcal{E}'}$, whose rows are indexed by $\mathcal{E}$ and whose columns are indexed by $\mathcal{E}'$. We proceed as follows. First, for $p \in \mathcal{E}$, consider the associated $x \in V \subseteq \mathbb{R}^n$ and F as above, and for all $1 \leqslant i \leqslant k$, let F_i denote the face $\mathcal{F}(x, h_i)$ of Q_i^{lift}, and C_i the cell $\mathcal{C}(x, h_i)$ of the associated subdivision of Q_i. Then by Corollary 2.1.27, we have

$$F = F_1 + \cdots + F_k \quad \text{and} \quad C = C_1 + \cdots + C_k \ .$$

Moreover, under the assumption that the tropical polynomials $f_1, \ldots, f_k$ do not share a common root, we know that at least one of the C_i is a singleton, and we let j be the maximal index such that $C_j = \{a_j\}$ is a singleton. We call the couple (j, a_j) the *row content* of p. Note that this row content only depends on the cell C containing $p - \delta$ in its interior. Then since $p - \delta \in C$, thus it can be written as

$$p - \delta = q_1 + \cdots + q_j + \cdots + q_k \quad \text{with} \quad \left\{ \begin{array}{lll} q_i & \in & C_i \quad \text{for all} \quad 1 \leqslant i \neq j \leqslant k \\ q_j & = & a_j \ . \end{array} \right. \tag{2.2}$$

We then construct the matrix $\mathcal{M}_{\mathcal{E}\mathcal{E}'}$ as follows : for every $p \in \mathcal{E}$, we associate to p its row content (j, a_j), and then we put the row $(j, p - a_j)$ of the Canny-Emiris submatrix $\mathcal{M}_{\mathcal{E}'}$ in $\mathcal{M}_{\mathcal{E}\mathcal{E}'}$, whose coefficients are given by the coefficients of the polynomial $X^{p-a_j} f_j(X)$. Note that all exponents appearing in the support of $X^{p-a_j} f_j(X)$ belong to $\mathcal{E}$ and thus to $\mathcal{E}'$. Indeed, if $a'_j \in \mathcal{A}_j$, then

$$\begin{array}{rcl} p - a_j + a'_j & = & \delta + q_1 + \cdots + q_j + \cdots + q_k - a_j + a'_j \\ & = & \delta + q_1 + \cdots + a'_j + \cdots + q_k \in (Q + \delta) \cap \mathbb{Z}^n = \mathcal{E} \subseteq \mathcal{E}' \ . \end{array}$$

Thus, we end up with a matrix $\mathcal{M}_{\mathcal{E}\mathcal{E}'} = (m_{pp'})_{(p,p') \in \mathcal{E} \times \mathcal{E}'}$ indexed by $\mathcal{E} \times \mathcal{E}'$. We shall show in the proof of Theorem 2.1.5 that this matrix $\mathcal{M}_{\mathcal{E}\mathcal{E}'}$ is actually a submatrix of $\mathcal{M}_{\mathcal{E}'}$ — it might not be the case because the rows selected by the row content are *a priori* not necessarily distinct. Also notice that by grouping together the columns of $\mathcal{M}_{\mathcal{E}\mathcal{E}'}$ indexed by $\mathcal{E}$ and by $\mathcal{E}' \setminus \mathcal{E}$, the matrix $\mathcal{M}_{\mathcal{E}\mathcal{E}'}$ can be written as a block matrix

$$\mathcal{M}_{\mathcal{E}\mathcal{E}'} = \begin{pmatrix} \mathcal{M}_{\mathcal{E}\mathcal{E}} & \mathbb{0} \end{pmatrix}$$

where $\mathcal{M}_{\mathcal{E}\mathcal{E}} = (m_{pp'})_{(p,p') \in \mathcal{E} \times \mathcal{E}}$ is a square matrix indexed by $\mathcal{E} \times \mathcal{E}$, which we shall show is a submatrix of $\mathcal{M}_{\mathcal{E}}$.

The proof of Theorem 2.1.5

In order to prove Theorem 2.1.5, we will make use of the following lemma.

Lemma 2.1.30. *Consider h and $h_1, \ldots, h_k$ as defined as in Section 2.1.3, and moreover for all $1 \leqslant j \leqslant k$, set $\widehat{h}_j = h_1 \,\square \cdots \square\, h_{j-1} \,\square\, h_{j+1} \,\square \cdots \square\, h_k$. Let $p \in \mathcal{E}$ and let (j, a_j) be its row content. Then for all $p' \in \mathcal{E}$ and $a'_j \in \mathbb{Z}^n$ such that $p' = p - a_j + a'_j$, we have*

$$h(p' - \delta) \geqslant h_j(a'_j) + \widehat{h}_j(p - \delta - a_j), \tag{2.3}$$

with equality if and only if $p' = p$ and $a'_j = a_j$.

Proof. Since $p' - \delta = a'_j + (p - \delta - a_j)$, the inequality (2.3) follows from the definition of sup-convolution, noting that $h = h_j \,\square\, \widehat{h}_j$.

We next show that the equality holds if $p' = p$, which entails that $a'_j = a_j$. Indeed, since

$$(p - \delta, h(p - \delta)) \in Q^{\mathsf{lift}} = \operatorname{hypo}(h_j) + \operatorname{hypo}(\widehat{h}_j) \ ,$$

we can write

$$(p - \delta, h(p - \delta)) = (q_j, \lambda_j) + (\widehat{q}_j, \widehat{\lambda}_j) \tag{2.4}$$

with $(q_j, \lambda_j) \in \operatorname{hypo}(h_j)$ and $(\widehat{q}_j, \widehat{\lambda}_j) \in \operatorname{hypo}(\widehat{h}_j)$. Then by Corollary 2.1.27, defining x as in Observation 2.1.29, it follows that $(q_j, \lambda_j) \in \mathcal{F}(x, h_j)$. In particular, since (j, a_j) is the row content of p, we have $\mathcal{F}(x, h_j) = \{(a_j, h_j(a_j))\}$, hence,

$$(q_j, \lambda_j) = (a_j, h_j(a_j)) \ .$$

Observation 2.1.29 also entails that $(\widehat{q}_j, \widehat{\lambda}_j) \in \mathcal{F}(x, \widehat{h}_j)$ and therefore that $(\widehat{q}_j, \widehat{\lambda}_j) = (\widehat{q}_j, \widehat{h}_j(\widehat{q}_j))$. Moreover, since $\widehat{q}_j = p - \delta - q_j = p - \delta - a_j$, it follows that

$$(\widehat{q}_j, \widehat{\lambda}_j) = (p - \delta - a_j, \widehat{h}_j(p - \delta - a_j)) \ .$$

Hence, we deduce from (2.4) that

$$h(p-\delta) = h_j(a_j) + \widehat{h}_j(p-\delta-a_j) \ . \tag{2.5}$$

We now show that (2.3) is strict whenever $p \neq p'$. Indeed, assume that the equality is achieved. Then this implies that

$$(p'-\delta, h(p'-\delta)) = (a'_j, h_j(a'_j)) + (p-\delta-a_j, \widehat{h}_j(p-\delta-a_j)) \ .$$

Now consider $x' \in \mathbb{R}^n$ such that $F' = \mathcal{F}(x', h)$ is the facet in the interior of which $(p'-\delta, h(p'-\delta))$ lies. Then from Corollary 2.1.27, we have

$$(a'_j, h_j(a'_j)) \in \mathcal{F}(x', h_j) \quad \text{and} \quad (p-\delta-a_j, \widehat{h}_j(p-\delta-a_j)) \in \mathcal{F}(x', \widehat{h}_j) \ .$$

However, we also know from equality (2.5) and Corollary 2.1.27 that

$$(p-\delta-a_j, \widehat{h}_j(p-\delta-a_j)) \in \mathcal{F}(x, \widehat{h}_j) \ ,$$

and since $\mathcal{F}(x, h_j) = \{(a_j, h_j(a_j))\}$ is a singleton, then $\mathcal{F}(x, \widehat{h}_j)$ is simply a translation of $\mathcal{F}(x, h)$, and moreover since $(p-\delta, h(p-\delta))$ is in the relative interior of $\mathcal{F}(x, h)$, then it means that $(p-\delta-a_j, \widehat{h}_j(p-\delta-a_j))$ is in the relative interior of $\mathcal{F}(x, \widehat{h}_j)$, and that $\mathcal{F}(x, \widehat{h}_j)$ is a facet of $\widehat{Q}_j^{\mathsf{lift}} := \operatorname{hypo}(\widehat{h}_j)$. Therefore, since $(p-\delta-a_j, \widehat{h}_j(p-\delta-a_j))$ is in both $\mathcal{F}(x, \widehat{h}_j)$ and $\mathcal{F}(x', \widehat{h}_j)$ and since it is in particular in the relative interior of the first face, we deduce that

$$\mathcal{F}(x, \widehat{h}_j) \subseteq \mathcal{F}(x', \widehat{h}_j) \ .$$

Since by Observation 2.1.24 (a), $\mathcal{F}(x', \widehat{h}_j)$ is a proper face of $\widehat{Q}_j^{\mathsf{lift}}$, then it implies that it is also a facet of $\widehat{Q}_j^{\mathsf{lift}}$, and thus that

$$\mathcal{F}(x, \widehat{h}_j) = \mathcal{F}(x', \widehat{h}_j) \ .$$

Hence, it follows from Observation 2.1.24 (b) that $x = x'$, and therefore $(a'_j, h_j(a'_j)) \in \mathcal{F}(x, h_j) = \{(a_j, h_j(a_j))\}$, hence $a'_j = a_j$ and thus $p' = p$. ☐

Using Lemma 2.1.30, we are now all set to give a proof of Theorem 2.1.5.

Proof of Theorem 2.1.5. The 'only if' implication is straightforward, for if x is a common root of the system f, then its image by the Veronese embedding

$$\operatorname{ver} : \left\{ \begin{array}{ccc} \mathbb{R}^n & \longrightarrow & \mathbb{R}^{\mathcal{E}'} \\ x & \longmapsto & \operatorname{ver}(x) = (x^\nu)_{\nu \in \mathcal{E}'} \end{array} \right.$$

yields a finite vector in the right null space of the matrix $\mathcal{M}_{\mathcal{E}'}$. Note that by convention, this also holds when $\mathcal{M}_{\mathcal{E}'}$ is empty.

For the converse implication, we rather show the contrapositive. Assume that the tropical polynomial system f does not have a solution, and consider the matrix $\mathcal{M}_{\mathcal{E}\mathcal{E}'} = (m_{pp'})_{(p,p') \in \mathcal{E} \times \mathcal{E}'}$ obtained from the Canny-Emiris construction described in Section 2.1.3. Note that this construction implies that the matrix $\mathcal{M}_{\mathcal{E}\mathcal{E}}$ and *a fortiori* the matrix $\mathcal{M}_{\mathcal{E}}$ are nonempty, leading to a contradiction when $\mathcal{M}_{\mathcal{E}}$ is empty, thus proving the 'if' implication in that case. We now assume that $\mathcal{M}_{\mathcal{E}}$ and thus $\mathcal{M}_{\mathcal{E}'}$ are nonempty.

Let

$$\widetilde{\mathcal{M}}_{\mathcal{E}\mathcal{E}'} = (\widetilde{m}_{pp'})_{(p,p') \in \mathcal{E} \times \mathcal{E}'} \quad \text{with} \quad \widetilde{m}_{pp'} = m_{pp'} - h(p'-\delta) \ .$$

Similarly to $\mathcal{M}_{\mathcal{E}\mathcal{E}'}$, this matrix can also be written as a block matrix in the following way:

$$\widetilde{\mathcal{M}}_{\mathcal{E}\mathcal{E}'} = \begin{pmatrix} \widetilde{\mathcal{M}}_{\mathcal{E}\mathcal{E}} & \mathbb{0} \end{pmatrix}$$

where $\widetilde{\mathcal{M}}_{\mathcal{E}\mathcal{E}} = (\widetilde{m}_{pp'})_{(p,p') \in \mathcal{E} \times \mathcal{E}}$ is a square matrix indexed by $\mathcal{E} \times \mathcal{E}$. We show that the tropical matrix $\widetilde{\mathcal{M}}_{\mathcal{E}\mathcal{E}}$ is diagonally dominant. Indeed, for all $p, p' \in \mathcal{E}$, and let (j, a_j) be the row content of p. Then, we have for $a'_j \in \mathbb{Z}^n$ such that $p' = p - a_j + a'_j$

$$m_{pp'} = \left\{ \begin{array}{ll} f_{j,a'_j} & \text{if} \quad a'_j \in \mathcal{A}_j \\ \mathbb{0} & \text{otherwise,} \end{array} \right.$$

or in other words, $m_{pp'} = \omega_j(a'_j)$. In particular, for $p' = p$, we have

$$\widetilde{m}_{pp} = f_{j,a_j} - h(p - \delta) \ .$$

Moreover, by assumption, the monomial $f_{j,a_j} X^{a_j}$ of f_j is essential. By Observation 2.1.28, we have therefore $h_j(a_j) = f_{j,a_j}$, hence we obtain from the equality case of Lemma 2.1.30 that

$$\widetilde{m}_{pp} = -\widehat{h}_j(p - \delta - a_j) \ .$$

Now for the case of $p' \neq p$, again since $h_j(a'_j) \geqslant f_{j,a'_j}$, then by rearranging the inequality of Lemma 2.1.30, we obtain that

$$\widetilde{m}_{pp} = -\widehat{h}_j(p - \delta - a_j) \geqslant f_{j,a'_j} - h(p' - \delta) = \widetilde{m}_{pp'} \ ,$$

with equality if and only if $p' = p$. Thus, the matrix $\widetilde{\mathcal{M}}_{\mathcal{E}\mathcal{E}}$ is tropically diagonally dominant. In particular, it follows that any two rows of $\widetilde{\mathcal{M}}_{\mathcal{E}\mathcal{E}'}$ are distinct, and thus it is also true for $\mathcal{M}_{\mathcal{E}\mathcal{E}'}$ because $\widetilde{\mathcal{M}}_{\mathcal{E}\mathcal{E}'}$ was obtained from $\mathcal{M}_{\mathcal{E}\mathcal{E}'}$ by adding the same vector to all of its rows. Therefore, it means that the matrix $\mathcal{M}_{\mathcal{E}\mathcal{E}'}$ is indeed a submatrix of $\mathcal{M}_{\mathcal{E}'}$ — and likewise the rows of $\mathcal{M}_{\mathcal{E}\mathcal{E}}$ are pairwise distinct, thus $\mathcal{M}_{\mathcal{E}\mathcal{E}}$ is a (square) submatrix of $\mathcal{M}_{\mathcal{E}}$.

Finally, since $\widetilde{\mathcal{M}}_{\mathcal{E}\mathcal{E}}$ is tropically diagonally dominant, by Lemma 2.1.16, it also is nonsingular, and thus applying Lemma 2.1.18 to $\mathcal{M}_{\mathcal{E}\mathcal{E}}$, we obtain that $\mathcal{M}_{\mathcal{E}\mathcal{E}}$ is also nonsingular. Therefore, the matrix $\mathcal{M}_{\mathcal{E}'}$ has the following form

$$\mathcal{M}_{\mathcal{E}'} = \begin{pmatrix} \mathcal{M}_{\mathcal{E}\mathcal{E}} & \mathbb{0} \\ * & * \end{pmatrix}$$

and thus Lemma 2.1.19 entails that the equation $\mathcal{M}_{\mathcal{E}'} \odot y \ \nabla \ \mathbb{0}$ has no solution y in $\mathbb{R}^{\mathcal{E}'}$, *i.e.* there cannot exist a finite vector in the tropical right null space of $\mathcal{M}_{\mathcal{E}'}$. □

Remark 2.1.31. Notice that in the case where $\mathcal{E}' = \mathcal{E}$, following the same proof, you obtain a weaker condition on the right null space of $\mathcal{M}_{\mathcal{E}}$ in order to find a finite solution of the system f. More precisely, you can show that there exists a solution $x \in \mathbb{R}^n$ to the system f if and only if there exists a vector $y \in \mathbb{T}^{\mathcal{E}} \setminus \{\mathbb{0}\}$ such that $\mathcal{M}_{\mathcal{E}} \odot y \ \nabla \ \mathbb{0}$. In other words, in the case where $\mathcal{E}' = \mathcal{E}$, the existence of a nonzero vector in the right null space of $\mathcal{M}_{\mathcal{E}}$, even possibly with some coordinates equal to $\mathbb{0}$, is enough to guarantee the existence of a finite solution of the system f.

Example 2.1.32. We illustrate the use of the Canny-Emiris construction in our proof by applying it to system

$$(\mathcal{S}_1) : \left\{ \begin{array}{lcl} f_1 & = & 1 \oplus 2x_1 \oplus 1x_2 \oplus 1x_1x_2 \\ f_2 & = & 0 \oplus 0x_1 \oplus 1x_2 \\ f_3 & = & 2x_1 \oplus 0x_2 \ , \end{array} \right.$$

which was shown in Example 2.1.8 not to have any solution. First, we obtain a subdivision of Q by projecting the non-vertical faces of the Minkowski sum of the hypographs of h_1, h_2 and h_3 onto the horizontal hyperplane $\mathbb{R}^n \times \{0\}$ as shown on Figure 2.6.

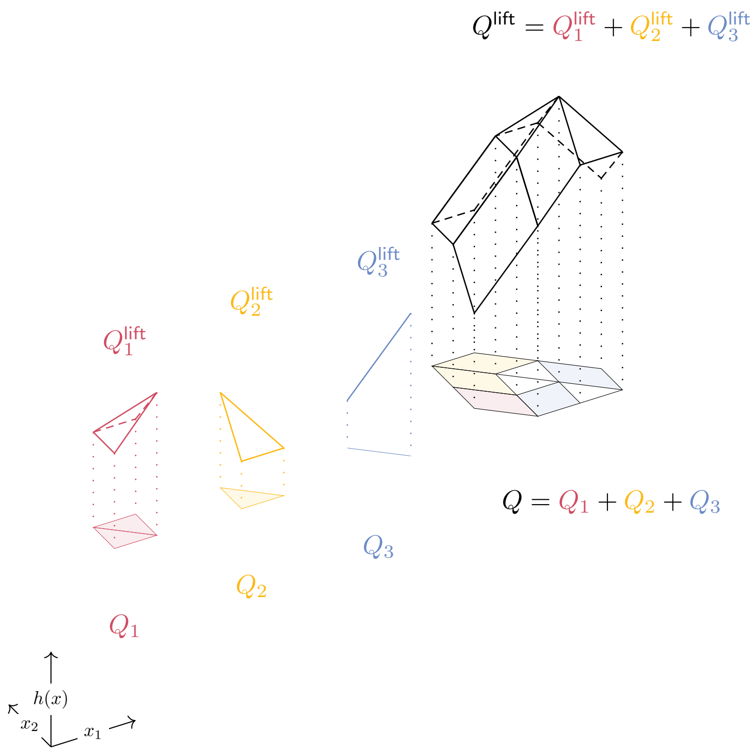

Figure 2.6: The subdivision of Q arises from the projection of the Minkowski sum of the hypographs of the h_i.

With this subdivision, we associate to every point $p \in \mathcal{E}$ its row content i, a_i, which is univocally determined by the maximal-dimensional cell of the decomposition of Q to which $p - \delta$ belongs. This process is illustrated in Figure 2.7.

More precisely, in the following table, for each point p of $\mathcal{E}$ in the first row:

- the second row displays the monomial x^p which corresponds to a column of the Macaulay matrix,
- the third row displays the row content i, a_i of p,
- the fourth row displays the polynomial $x^{p-a_i} f_i$ which corresponds to a row of the Macaulay matrix,
- and finally the last row displays the tropical scaling factor $h(p - \delta)$ which must be substracted (in the usual sense) to the column p of the matrix $\mathcal{M}^{(1)}_{\mathcal{E}\mathcal{E}}$ in order to obtain the matrix $\widetilde{\mathcal{M}}^{(1)}_{\mathcal{E}\mathcal{E}}$.

$p \in \mathcal{E}$	$(0,0)$	$(1,0)$	$(0,1)$	$(2,0)$	$(1,1)$	$(0,2)$
x^p	1	x_1	x_2	x_1^2	x_1x_2	x_2^2
i, a_i	$1, (0,0)$	$3, (1,0)$	$2, (0,1)$	$3, (1,0)$	$3, (1,0)$	$2, (0,1)$
$x^{p-a_i} f_i$	f_1	f_3	f_2	$x_1 f_3$	$x_2 f_3$	$x_2 f_2$
$h(p-\delta)$	3.6	4.8	3.8	3.3	4.1	2.2

With the information from the previous table, we obtain the following 6×6 square submatrix $\mathcal{M}^{(1)}_{\mathcal{E}\mathcal{E}}$ of $\mathcal{M}^{(1)}_{\mathcal{E}}$

$$\mathcal{M}^{(1)}_{\mathcal{E}\mathcal{E}} = \begin{array}{cc} & \begin{array}{cccccc} 1 & x_1 & x_2 & x_1^2 & x_1x_2 & x_2^2 \end{array} \\ \begin{array}{c} f_1 \\ f_3 \\ f_2 \\ x_1 f_3 \\ x_2 f_3 \\ x_2 f_2 \end{array} & \begin{pmatrix} 1 & 2 & 1 & & 1 & \\ & 2 & 0 & & & \\ 0 & 0 & 1 & & & \\ & & & 2 & 0 & \\ & & & & 2 & 0 \\ & & 0 & & 0 & 1 \end{pmatrix} \end{array},$$

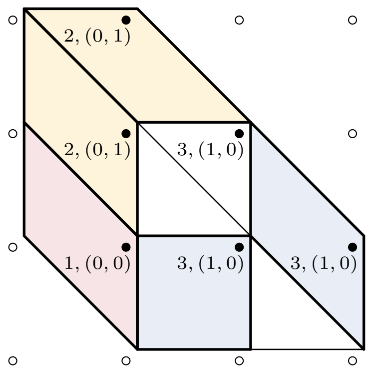

Figure 2.7: The polytope $Q + \delta$, with the integer points inside the maximal dimensional cells of the decomposition of $Q + \delta$ labelled by the row content the cell they belong to.

and after applying the tropical scaling of factor $h(p - \delta)$ to the column p of the previous matrix for all $p \in \mathcal{E}$, we obtain the following matrix, in which we highlighted the diagonal coefficients by writing them in bold

$$\widetilde{\mathcal{M}}^{(1)}_{\mathcal{E}\mathcal{E}} = \begin{array}{c} \\ f_1 \\ f_3 \\ f_2 \\ x_1 f_3 \\ x_2 f_3 \\ x_2 f_2 \end{array} \begin{array}{c} \begin{array}{cccccc} 1 & x_1 & x_2 & x_1^2 & x_1 x_2 & x_2^2 \end{array} \\ \begin{pmatrix} \mathbf{-2.6} & -2.8 & -2.8 & & -3.1 & \\ & \mathbf{-2.8} & -3.8 & & & \\ -3.6 & -4.8 & \mathbf{-2.8} & & & \\ & & & \mathbf{-1.3} & -4.1 & \\ & & & & \mathbf{-2.1} & -2.2 \\ & & -3.8 & & -4.1 & \mathbf{-1.2} \end{pmatrix} \end{array} .$$

We finally observe that the matrix $\widetilde{\mathcal{M}}^{(1)}_{\mathcal{E}\mathcal{E}}$ is tropically diagonally dominant, and therefore it follows from Lemmas 2.1.16 and 2.1.18 that $\mathcal{M}^{(1)}_{\mathcal{E}\mathcal{E}}$ is also nonsingular[1], thus there is no solution $y \in \mathbb{R}^{\mathcal{E}}$ to the equation $\mathcal{M}^{(1)}_{\mathcal{E}} \odot y \nabla \mathbb{0}$.

Besides, note that by choosing a different ordering of the polynomials, in which f_3 comes before f_2, then the row content of $(1, 1)$ would be $2, (1, 0)$ and thus the fifth row of $\mathcal{M}^{(1)}_{\mathcal{E}\mathcal{E}}$ would change to be

$$\begin{array}{c} \\ x_1 f_2 \end{array} \begin{array}{c} \begin{array}{cccccc} 1 & x_1 & x_2 & x_1^2 & x_1 x_2 & x_2^2 \end{array} \\ \begin{pmatrix} & 0 & & 0 & 1 & \end{pmatrix} \end{array}$$

and we can still check likewise that the right null space of this matrix is simply equal to $\{\mathbb{0}\}$.

2.2 The Tropical Positivstellensatz for two-sided systems

2.2.1 Statement of the theorem

Now we formulate a result similar to Theorem 2.1.5, for systems mixing equalities and weak and strict inequalities between tropical polynomial functions (also called 'two sided' systems). More precisely, consider $f^+ = (f_1^+, \dots, f_k^+)$ and $f^- = (f_1^-, \dots, f_k^-)$ two collections of k tropical polynomials. For $1 \leqslant i \leqslant k$, we denote by $\mathcal{A}_i^+$ and $\mathcal{A}_i^-$ respectively the supports of f_i^+ and f_i^-, and set $\mathcal{A}_i = \mathcal{A}_i^+ \cup \mathcal{A}_i^-$, and $\mathcal{A} = (\mathcal{A}_1, \dots, \mathcal{A}_k)$. Equivalently, $\mathcal{A}$ is the collection of supports of the polynomials $f_i := f_i^+ \oplus f_i^-$.

[1]Alternatively, one can verify that $\mathrm{tdet}(\mathcal{M}^{(1)}_{\mathcal{E}\mathcal{E}}) \not\nabla \mathbb{0}$ if tdet denotes the tropical determinant defined in Remark 2.1.13.

Consider also a collection $\rhd = (\rhd_1, \dots, \rhd_k)$ of relations $\rhd_i \in \{\geqslant, =, >\}$, for $1 \leqslant i \leqslant k$. We denote by $f^+(x) \rhd f^-(x)$ the system

$$\max_{\alpha \in \mathcal{A}_i^+} \left(f^+_{i,\alpha} + \langle \alpha, x\rangle\right) \rhd_i \max_{\alpha \in \mathcal{A}_i^-} \left(f^-_{i,\alpha} + \langle \alpha, x\rangle\right) \text{ for all } 1 \leqslant i \leqslant k, \tag{2.6}$$

of unknown $x \in \mathbb{T}^n$. Note that this is equivalent to the expressions $f_i^+(x) \rhd_i f_i^-(x)$, $i = 1, \dots, k$. Finally, we denote by $\mathcal{M}^+$ and $\mathcal{M}^-$ the Macaulay matrices associated to f^+ and f^- respectively, so with entries $\mathcal{M}^{\pm}_{(i,\alpha),\beta} = f^{\pm}_{i,\beta-\alpha}$. Then, for any subset $\mathcal{E}$ of $\mathbb{Z}^n$, we denote by $\mathcal{M}^+_{\mathcal{E}}$ and $\mathcal{M}^-_{\mathcal{E}}$ the submatrices associated to $\mathcal{E}$ and the collection $\mathcal{A}$ defined above, that is $\mathcal{M}^+_{\mathcal{E}} = (\mathcal{M}^+)^{\mathcal{A}}_{\mathcal{E}}$ and $\mathcal{M}^-_{\mathcal{E}} = (\mathcal{M}^-)^{\mathcal{A}}_{\mathcal{E}}$, and we likewise set for $N \in \mathbb{N}$, $\mathcal{M}^+_N = (\mathcal{M}^+)^{\mathcal{A}}_N$ and $\mathcal{M}^-_N = (\mathcal{M}^-)^{\mathcal{A}}_N$.

Finally, we set for $1 \leqslant i \leqslant k$:

$$r_i = \begin{cases} \dim(\mathrm{aff}(\mathcal{A}_i^-)) + 1 & \text{if } \rhd_i \in \{\geqslant, >\} \\ \max\left(\dim(\mathrm{aff}(\mathcal{A}_i^-)), \dim(\mathrm{aff}(\mathcal{A}_i^+))\right) + 1 & \text{if } \rhd_i \in \{=\} \end{cases}.$$

Notice that in particular, this definition implies the following inequality

$$r_i \leqslant \begin{cases} \min(|\mathcal{A}_i^-|, n+1) & \text{if } \rhd_i \in \{\geqslant, >\} \\ \min(\max(|\mathcal{A}_i^-|, |\mathcal{A}_i^+|), n+1) & \text{if } \rhd_i \in \{=\} \end{cases}.$$

We now call *Canny-Emiris subsets* of $\mathbb{Z}^n$ associated to the system $f^+ \rhd f^-$ any set $\mathcal{E}$ of the form

$$\mathcal{E} := \left(\widetilde{Q} + \delta\right) \cap \mathbb{Z}^n \quad \text{with} \quad \widetilde{Q} = r_1 Q_1 + \cdots + r_k Q_k \ ,$$

where $Q_i = \mathrm{conv}(\mathcal{A}_i)$ for $1 \leqslant i \leqslant k$, and δ is a generic vector in $V + \mathbb{Z}^n$, with V the direction of the affine hull of $\widetilde{Q}$ (note that this is the same as for $Q_1 + \cdots + Q_k$). Finally, for any subset $\mathcal{E}'$ of $\mathbb{Z}^n$ containing a Canny-Emiris subset $\mathcal{E}$ associated to $f^+ \rhd f^-$ and $y \in \mathbb{R}^{\mathcal{E}'}$, we denote by $\mathcal{M}^+_{\mathcal{E}'} \odot y \rhd \mathcal{M}^-_{\mathcal{E}'} \odot y$ the following system of tropical linear equalities:

$$\max_{\beta \in \mathcal{E}'} \left(m^+_{(i,\alpha),\beta} + y_\beta\right) \rhd_i \max_{\beta \in \mathcal{E}'} \left(m^-_{(i,\alpha),\beta} + y_\beta\right) \quad \text{for all} \quad 1 \leqslant i \leqslant k \quad \text{and} \quad \alpha \in \mathcal{A}_i \ .$$

We now state a Positivstellensatz for the case of tropical polynomial systems allowing both weak and strict inequalities, and equalities.

Theorem 2.2.1 (Sparse tropical Positivstellensatz). *There exists a solution $x \in \mathbb{R}^n$ to the system $f^+ \rhd f^-$ if and only if there exists a vector $y \in \mathbb{R}^{\mathcal{E}'}$ satisfying $\mathcal{M}^+_{\mathcal{E}'} \odot y \rhd \mathcal{M}^-_{\mathcal{E}'} \odot y$, where $\mathcal{E}'$ is any subset of $\mathbb{Z}^n$ containing a nonempty Canny-Emiris subset $\mathcal{E}$ of $\mathbb{Z}^n$ associated to the system $f^+ \rhd f^-$.*

Remark 2.2.2. The following consideration justifies the name of *tropical Positivstellensatz* given to the previous theorem. By analogy with classical semialgebraic sets, let us define a *tropical basic semialgebraic subset* of $\mathbb{R}^n$ to be the set of solutions of a collection of inequalities of the form $f_i^+ \rhd_i f_i^-$, $i \in [k]$ where $f_1^{\pm}, \dots, f_k^{\pm}$ are pairs of tropical polynomials, and $\rhd_1, \dots, \rhd_k \in \{\geqslant, >\}$. Then, Theorem 2.2.1 provides an effective way to check the inclusion of two tropical basic semialgebraic sets. Let us illustrate this by the following typical special case: given $f_1^{\pm}, \dots, f_{k+1}^{\pm}$ a collection of pairs of tropical polynomials, check whether the following implication holds:

$$\left\{ \begin{array}{ccc} f_1^+(x) & \geqslant & f_1^-(x) \\ & \vdots & \\ f_k^+(x) & \geqslant & f_k^-(x) \end{array} \right. \implies f_{k+1}^+(x) \geqslant f_{k+1}^-(x) \quad \forall x \in \mathbb{R}^n \ . \tag{$\mathscr{P}$}$$

Checking ($\mathscr{P}$) is equivalent to showing that the following system

$$(\mathcal{S}) : \left\{ \begin{array}{ccc} f_1^+(x) & \geqslant & f_1^-(x) \\ & \vdots & \\ f_k^+(x) & \geqslant & f_k^-(x) \\ f_{k+1}^+(x) & < & f_{k+1}^-(x) \end{array} \right.$$

has no solution $x \in \mathbb{R}^n$, which can be done using Theorem 2.2.1, by reduction to the problem of the unsolvability a system of linear tropical (in) equations, allowing both strict and weak inequalities. Recall that the latter system reduces to a mean payoff game [AFG+14, Theorem 4.7], see also [AGK11b, Theorem 18]. Therefore, the certificate of unfeasibility shall be given in the form of a strategy for the minimizer player in the associated mean payoff game.

Example 2.2.3. Now let us illustrate Theorem 2.1.5 with an example. Consider the following problem

$$\left\{\begin{array}{lll} 0x_1 \oplus 0x_1x_2 & \geqslant & 0x_2 \\ 2 \oplus 0x_1 & \geqslant & 1x_2 \\ 3 & \geqslant & 0x_1 \end{array}\right. \implies \quad 1x_1 \ \geqslant \ 0x_2 \oplus (-3)x_1^2 \ . \tag{$\mathcal{P}$}$$

We want to show that the implication in ($\mathcal{P}$) holds. This is the case if and only if the system

$$\left\{\begin{array}{rcl} 0x_1 \oplus 0x_1x_2 & \geqslant & 0x_2 \\ 2 \oplus 0x_1 & \geqslant & 1x_2 \\ 3 & \geqslant & 0x_1 \\ 0x_2 \oplus (-3)x_1^2 & > & 1x_1 \end{array}\right. \tag{$\mathcal{S}$}$$

does not have a solution $x \in \mathbb{R}^2$. We can turn the strict inequality into a weak inequality by noticing that the latter system does not have a solution on $\mathbb{R}^2$ if and only if for all $\lambda > \mathbb{1} = 0$, the system

$$\left\{\begin{array}{rcl} 0x_1 \oplus 0x_1x_2 & \geqslant & 0x_2 \\ 2 \oplus 0x_1 & \geqslant & 1x_2 \\ 3 & \geqslant & 0x_1 \\ 0x_2 \oplus (-3)x_1^2 & \geqslant & \lambda \odot 1x_1 \end{array}\right. \tag{$\mathcal{S}'_\lambda$}$$

does not have a solution $x \in \mathbb{R}^2$.

We rewrite the system ($\mathcal{P}$) as $\forall i \in \{1, 2, 3\}$, $f_i^+(x) \geqslant f_i^-(x) \implies f_4^+(x) \geqslant f_4^-(x)$ by setting

$$\begin{array}{lll} f_1^+ & = & 0x_1 \oplus 0x_1x_2 \\ f_2^+ & = & 2 \oplus 0x_1 \\ f_3^+ & = & 3 \\ f_4^+ & = & 1x_1 \end{array} \qquad \begin{array}{lll} f_1^- & = & 0x_2 \\ f_2^- & = & 1x_2 \\ f_3^- & = & 0x_1 \\ f_4^- & = & 0x_2 \oplus (-3)x_1^2 \end{array} \qquad \begin{array}{lll} f_1 & = & 0x_1 \oplus 0x_2 \oplus 0x_1x_2 \\ f_2 & = & 2 \oplus 0x_1 \oplus 1x_2 \\ f_3 & = & 3 \oplus 0x_1 \\ f_4 & = & 1x_1 \oplus 0x_2 \oplus (-3)x_1^2 \end{array}$$

On Figure 2.8, we show the Newton polytopes associated to the polynomials f_1, f_2, f_3 and f_4 as well as their Minkowski sum. Moreover, on Figures 2.9 and 2.10, we show on the left the configuration of hypersurfaces associated respectively to ($\mathcal{P}$) and ($\mathcal{S}'_\lambda$). In this configuration, we colored every intersection point by the colors of the curves that do not take part in the intersection, and we reported the coloring on the matching cell of the decomposition of Q on the right, as to illustrate the duality between the hypersurface configuration and the subdivision of Q.

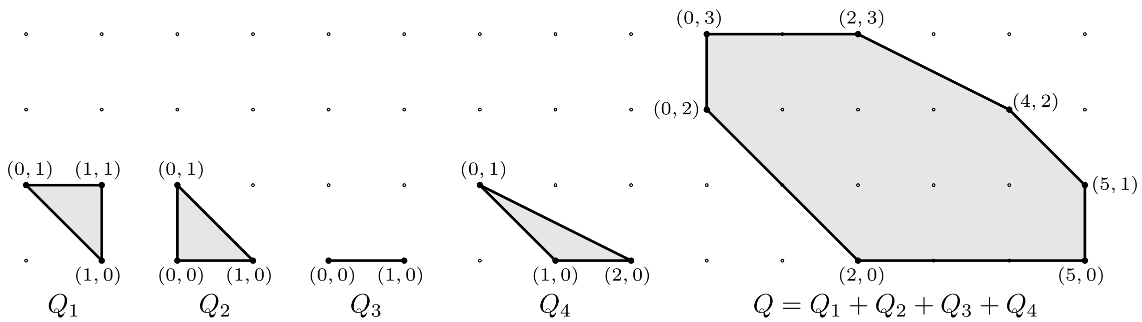

Figure 2.8: The Newton polytopes associated f_1, f_2, f_3, f_4 and their Minkowski sum.

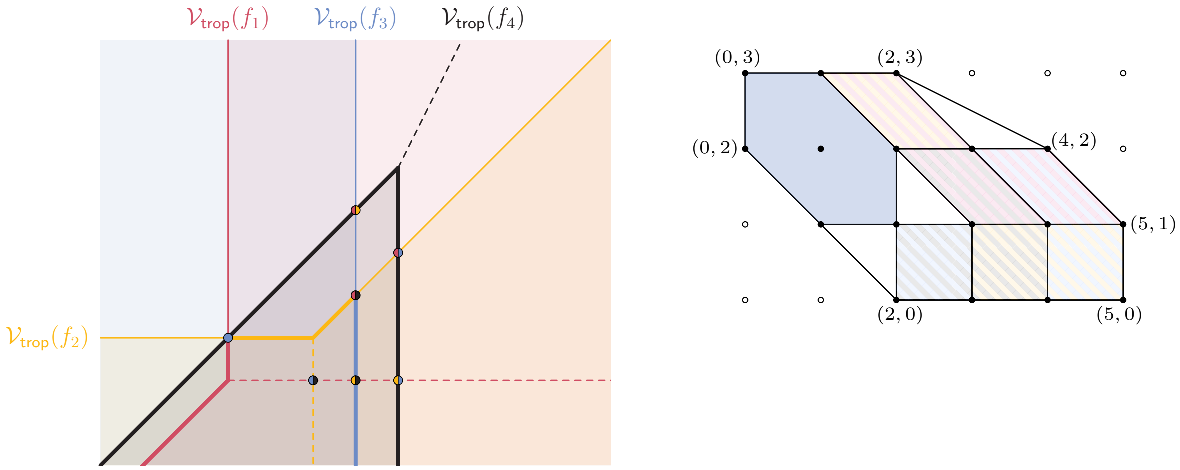

Figure 2.9: Problem $(\mathcal{P})$ is illustrated here as the tropical basic semialgebraic set $\{f_1^+ \geqslant f_1^-\} \cap \{f_2^+ \geqslant f_2^-\} \cap \{f_3^+ \geqslant f_3^-\}$ is indeed included in the set $\{f_4^+ \geqslant f_4^-\}$. The subdivision of Q associated to this system is displayed on the right.

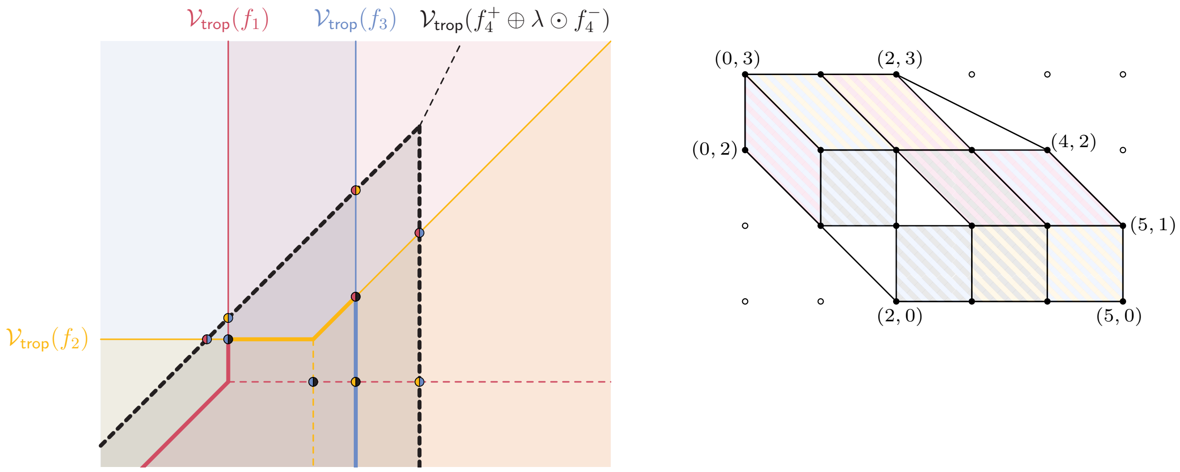

Figure 2.10: For all $\lambda > 0$, the tropical basic semialgebraic set $\{f_1^+ \geqslant f_1^-\} \cap \{f_2^+ \geqslant f_2^-\} \cap \{f_3^+ \geqslant f_3^-\}$ is indeed included in the set $\{\lambda \odot f_4^+ > f_4^-\}$, showing that system $(\mathcal{S}'_\lambda)$ does not have a solution in $\mathbb{R}^2$. The subdivision of Q associated to this system is displayed on the right.

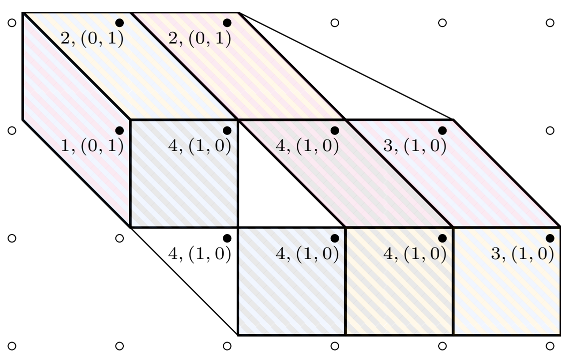

Figure 2.11: The polytope $Q+\delta$, with the integer points inside the maximal dimensional cells of the decomposition of $Q+\delta$ labelled by the row content the cell they belong to.

For this collection of supports, one can once again take $\delta = (-1+\varepsilon, -1+\varepsilon)$ with $\varepsilon = \frac{1}{10}$, which gives us the Canny-Emiris set

$$\mathcal{E} = \{(1,0),(0,1),(2,0),(1,1),(0,2),(3,0),(2,1),(1,2),(4,0),(3,1)\}$$

corresponding to the set of monomials

$$\{x_1, x_2, x_1^2, x_1x_2, x_2^2, x_1^3, x_1^2x_2, x_1x_2^2, x_1^4, x_1^3x_2\} \ .$$

We thus obtain the 21×10 submatrices $\mathcal{M}^+_{\mathcal{E}}$ and $\mathcal{M}^-_{\mathcal{E}}$ of the Macaulay matrix associated to the set $\mathcal{E}$, which we combined in the single following matrix for the sake of readability and to save some space, with the normal font weight coeffficients corresponding to the coefficients of $\mathcal{M}^+_{\mathcal{E}}$, and the bold ones corresponding to $\mathcal{M}^-_{\mathcal{E}}$

$$\mathcal{M}^{\pm}_{\mathcal{E}} = \begin{array}{cc}
 & \begin{array}{cccccccccc} x_1 & x_2 & x_1^2 & x_1x_2 & x_2^2 & x_1^3 & x_1^2x_2 & x_1x_2^2 & x_1^4 & x_1^3x_2 \end{array} \\
\begin{array}{c} f_1 \\ x_1f_1 \\ x_2f_1 \\ x_1^2f_1 \\ x_1f_2 \\ x_2f_2 \\ x_1^2f_2 \\ x_1x_2f_2 \\ x_1^3f_2 \\ x_1f_3 \\ x_2f_3 \\ x_1^2f_3 \\ x_1x_2f_3 \\ x_2^2f_3 \\ x_1^3f_3 \\ x_1^2x_2f_3 \\ f_4 \\ x_1f_4 \\ x_2f_4 \\ x_1^2f_4 \\ x_1x_2f_4 \end{array} &
\begin{pmatrix}
0 & \mathbf{0} & & 0 & & & & & & \\
 & & 0 & \mathbf{0} & & & 0 & & & \\
 & & & 0 & \mathbf{0} & & & 0 & & \\
 & & & & & 0 & \mathbf{0} & & & 0 \\
2 & & 0 & \mathbf{1} & & & & & & \\
 & 2 & & 0 & \mathbf{1} & & & & & \\
 & & 2 & & & 0 & \mathbf{1} & & & \\
 & & & 2 & & & 0 & \mathbf{1} & & \\
 & & & & & 2 & & & 0 & \mathbf{1} \\
3 & & \mathbf{0} & & & & & & & \\
 & 3 & & \mathbf{0} & & & & & & \\
 & & 3 & & & \mathbf{0} & & & & \\
 & & & 3 & & & \mathbf{0} & & & \\
 & & & & 3 & & & \mathbf{0} & & \\
 & & & & & 3 & & & \mathbf{0} & \\
 & & & & & & 3 & & & \mathbf{0} \\
\boldsymbol{1+\lambda} & 0 & -3 & & & & & & & \\
 & & \boldsymbol{1+\lambda} & 0 & & -3 & & & & \\
 & & & \boldsymbol{1+\lambda} & 0 & & -3 & & & \\
 & & & & & \boldsymbol{1+\lambda} & 0 & & -3 & \\
 & & & & & & \boldsymbol{1+\lambda} & 0 & & -3
\end{pmatrix}
\end{array} .$$

Applying the Canny-Emiris construction to the collection of polynomials in the system $(\mathcal{S}'_\lambda)$, we obtain the subdivision of Q which allows us to associate to every point $p \in \mathcal{E}$ its row content i, a_i, which we recall is univocally determined by the maximal-dimensional cell of the decomposition of Q to which $p - \delta$ belongs, as illustrated in Figure 2.11.

More precisely, in the following table, for each point p of $\mathcal{E}$ in the first row:

- the second row displays the monomial x^p which corresponds to a column of the Macaulay matrix,
- the third row displays the row content i, a_i of p,
- the fourth row displays the polynomial $x^{p-a_i}f_i$ which corresponds to a row of the Macaulay matrix,
- and finally the last row displays the tropical scaling factor $h(p-\delta)$ which must be substracted (in the usual sense) to the column p of the matrices $\mathcal{M}^{\pm}_{\mathcal{E}\mathcal{E}}$ in order to obtain the matrix $\widetilde{\mathcal{M}}^{\pm}_{\mathcal{E}\mathcal{E}}$.

$p \in \mathcal{E}$	$(1,0)$	$(0,1)$	$(2,0)$	$(1,1)$	$(0,2)$	$(3,0)$	$(2,1)$	$(1,2)$	$(4,0)$	$(3,1)$
x^p	x_1	x_2	x_1^2	x_1x_2	x_2^2	x_1^3	$x_1^2x_2$	$x_1x_2^2$	x_1^4	$x_1^3x_2$
i, a_i	$4,(1,0)$	$1,(0,1)$	$4,(1,0)$	$4,(1,0)$	$2,(0,1)$	$4,(1,0)$	$4,(1,0)$	$2,(0,1)$	$3,(1,0)$	$3,(1,0)$
$x^{p-a_i}f_i$	f_4	f_1	x_1f_4	x_2f_4	x_2f_2	$x_1^2f_4$	$x_1x_2f_4$	$x_1x_2f_2$	$x_1^3f_3$	$x_1^2x_2f_3$
$h(p-\delta)$	$6+\lambda$	$\frac{51+9\lambda}{10}$	$\frac{42}{10}+\lambda$	$\frac{51}{10}+\lambda$	$\frac{41+\lambda}{10}$	$\frac{13}{10}+\lambda$	$\frac{25}{10}+\lambda$	$\frac{17+\lambda}{10}$	$\frac{-26+\lambda}{10}$	$\frac{-13+2\lambda}{10}$

With the information from the previous table, we obtain the following 10×10 pair of square submatrices $\mathcal{M}^{\pm}_{\mathcal{E}\mathcal{E}}$ of $\mathcal{M}^{\pm}_{\mathcal{E}}$

$$
\mathcal{M}^{\pm}_{\mathcal{E}\mathcal{E}} = \begin{array}{c} \\ f_4 \\ f_1 \\ x_1 f_4 \\ x_2 f_4 \\ x_2 f_2 \\ x_1^2 f_4 \\ x_1 x_2 f_4 \\ x_1 x_2 f_2 \\ x_1^3 f_3 \\ x_1^2 x_2 f_3 \end{array}
\begin{array}{c}
\begin{array}{cccccccccc} x_1 & x_2 & x_1^2 & x_1x_2 & x_2^2 & x_1^3 & x_1^2x_2 & x_1x_2^2 & x_1^4 & x_1^3x_2 \end{array} \\
\begin{pmatrix}
\mathbf{1+\lambda} & 0 & -3 & & & & & & & \\
0 & \mathbf{0} & & 0 & & & & & & \\
 & & \mathbf{1+\lambda} & 0 & & -3 & & & & \\
 & & & \mathbf{1+\lambda} & 0 & & -3 & & & \\
 & 2 & & 0 & \mathbf{1} & & & & & \\
 & & & & & \mathbf{1+\lambda} & 0 & & -3 & \\
 & & & & & & \mathbf{1+\lambda} & 0 & & -3 \\
 & & & 2 & & & 0 & \mathbf{1} & & \\
 & & & & & 3 & & & \mathbf{0} & \\
 & & & & & & 3 & & & \mathbf{0}
\end{pmatrix}
\end{array} .
$$

Finally, after applying the tropical scaling of factor $h(p-\delta)$ to the column p of the previous matrices for all $p \in \mathcal{E}$, we obtain the following pair of matrices

$$
\widetilde{\mathcal{M}}^{\pm}_{\mathcal{E}\mathcal{E}} = -\frac{1}{10}
\begin{pmatrix}
\mathbf{50} & 51+9\lambda & 72+10\lambda & & & & & & & \\
60+10\lambda & \mathbf{51+9\lambda} & & 51+10\lambda & & & & & & \\
 & & \mathbf{32} & 51+10\lambda & & 43+10\lambda & & & & \\
 & & & \mathbf{41} & 41+\lambda & & 55+10\lambda & & & \\
 & 31+9\lambda & & 51+10\lambda & \mathbf{31+\lambda} & & & & & \\
 & & & & & \mathbf{3} & 25+10\lambda & & 4+\lambda & \\
 & & & & & & \mathbf{15} & 17+\lambda & & 17+2\lambda \\
 & & & 31+10\lambda & & & 25+10\lambda & \mathbf{7+\lambda} & & \\
 & & & & & -17+10\lambda & & & \mathbf{-26+\lambda} & \\
 & & & & & & -5+10\lambda & & & \mathbf{-13+2\lambda}
\end{pmatrix}
$$

and we indeed observe for every $\lambda > 0$ that for each row, the diagonal coefficient of $\widetilde{\mathcal{M}}^{-}_{\mathcal{E}\mathcal{E}}$ is strictly greater than all the coefficients of $\widetilde{\mathcal{M}}^{+}_{\mathcal{E}\mathcal{E}}$ in the same row, or in other words that $\widetilde{\mathcal{M}}^{+}_{\mathcal{E}\mathcal{E}}$ is diagonally dominated by $\widetilde{\mathcal{M}}^{-}_{\mathcal{E}\mathcal{E}}$, and thus from Lemmas 2.2.7 and 2.2.9, the only solution to the equation $\mathcal{M}^{+}_{\mathcal{E}} \odot y \geqslant \mathcal{M}^{-}_{\mathcal{E}} \odot y$ in $\mathbb{R}^{10}$ is $y = \mathbb{0}$.

Remark 2.2.4. We may define a tropical semialgebraic subset of $\mathbb{R}^n$ as a finite union of tropical basic semialgebraic subsets. This leads to a more general class of sets than the one arising by considering the images by the valuation of semialgebraic sets over a real closed non-archimedean field. Indeed, it is shown in [AGS20, Theorem 3.1] that the image by a non-trivial and convex valuation of a semialgebraic set over a non-archimedean field is a closed semilinear set. In particular, when the value group is $\mathbb{R}$, this image is a finite union of *closed* tropical basic semialgebraic subsets, whereas our definition allows more generally tropical basic semialgebraic subsets not be closed, owing to the presence of strict inequalities.

Remark 2.2.5. Once again, the proof of the implication that if there is a solution $x \in \mathbb{R}^n$ to the system $f^+ \rhd f^-$, then there exists a vector $y \in \mathbb{R}^{\mathcal{E}'}$ such that $\mathcal{M}^+ \odot y \rhd \mathcal{M}^- \odot y$ is immediate by choosing y to be the Veronese embedding $\mathrm{ver}(x)$ of x, which we recall is equal to the vector $(x^\nu)_{\nu \in \mathcal{E}'}$. Therefore, we will only focus on the converse implication in what follows, or rather on its contrapositive.

2.2.2 Further preliminary results

In this section, we adapt the results from Section 2.1.2 to the two-sided case in order to prove Theorem 2.2.1.

Diagonal dominance for a pair of matrices

In order to prove this theorem, we introduce a notion of diagonal dominance adapted to systems of inequalities and equalities.

Definition 2.2.6. Let $A = (a_{ij})_{(i,j) \in [p] \times [p]}$ and $B = (b_{ij})_{(i,j) \in [p] \times [p]}$ be a pair of matrices in $\mathbb{T}^{p \times p}$. One says that B *diagonally dominates* A, or that A is *diagonally dominated* by B whenever

$$b_{ii} > a_{ij} \quad \text{for all} \quad 1 \leqslant i, j \leqslant p \ .$$

Lemma 2.2.7. *Let A and B be a pair of matrices in $\mathbb{T}^{p \times p}$ such that A is diagonally dominated by B. Then the only solution to the inequation $A \odot y \geqslant B \odot y$ of unknown $y \in \mathbb{T}^p$ is $y = \mathbb{0}$.*

Proof. Let $y = (y_1, \dots, y_p) \in \mathbb{T}^p$ be such that $A \odot y \geqslant B \odot y$ and consider $1 \leqslant i \leqslant p$ such that $y_i = \max_{1\leqslant j\leqslant p} y_j$. Assume that $y \neq \mathbb{0}$, and thus that $y_i > -\infty$. Then from the relation $A \odot y \geqslant B \odot y$, it follows in particular that

$$\max_{1\leqslant j\leqslant p} (a_{ij} + y_j) \geqslant \max_{1\leqslant j\leqslant p} (b_{ij} + y_j) \geqslant b_{ii} + y_i \ ,$$

but by choice of i and since A is diagonally dominated by B, we have

$$b_{ii} + y_i > a_{ij} + y_j \quad \text{for all} \quad 1 \leqslant j \leqslant p \ ,$$

hence we get a contradiction. ☐

Corollary 2.2.8. *Let A and B be a pair of matrices in $\mathbb{T}^{p\times p}$ such that for all $1 \leqslant i \leqslant p$, we have either*

$$b_{ii} > \max_{1\leqslant j\leqslant p} a_{ij} \quad \textit{or} \quad a_{ii} > \max_{1\leqslant j\leqslant p} b_{ij} \ . \tag{2.7}$$

Then the only solution to the inequation $A \odot y = B \odot y$ of unknown $y \in \mathbb{T}^p$ is $y = \mathbb{0}$.

Proof. The idea of the proof is simply that by swapping some of the rows of A and B, we can obtain a pair of matrices $\widetilde{A}$ and $\widetilde{B}$ such that $\widetilde{A}$ is diagonally dominated by $\widetilde{B}$. More precisely, set

$$I = \{1 \leqslant i \leqslant p : b_{ii} > a_{ij} \text{ for all } 1 \leqslant j \leqslant p\} \quad \text{and} \quad J = \{1, \dots, p\} \setminus I \ ,$$

and let

$$\widetilde{A} = (\widetilde{a}_{ij})_{(i,j)\in[p]\times[p]} \quad \text{with} \quad \widetilde{a}_{ij} = \begin{cases} a_{ij} & \text{if} \quad i \in I \\ b_{ij} & \text{if} \quad i \in J \end{cases}$$

and

$$\widetilde{B} = (\widetilde{b}_{ij})_{(i,j)\in[p]\times[p]} \quad \text{with} \quad \widetilde{b}_{ij} = \begin{cases} b_{ij} & \text{if} \quad i \in I \\ a_{ij} & \text{if} \quad i \in J \ . \end{cases}$$

By construction, notice that for $y \in \mathbb{R}^p$, we have

$$\widetilde{A} \odot y = \widetilde{B} \odot y \iff A \odot y = B \odot y \ ,$$

and moreover $\widetilde{A}$ is diagonally dominated by $\widetilde{B}$, thus by the previous lemma, the only solution to the equality $\widetilde{A} \odot y = \widetilde{B} \odot y$ is $y = \mathbb{0}$, and thus $\mathbb{0}$ is also the only solution to the equality $A \odot y = B \odot y$. ☐

Finally we will also make use of the following two lemmas, which are immediate adaptations of Lemma 2.1.18 and Lemma 2.1.19 to the two-sided case.

Lemma 2.2.9. *Let $\rhd \in \{\geqslant, =\}$. Let $A = (a_{ij})_{(i,j)\in[p]\times[q]}$ and $B = (b_{ij})_{(i,j)\in[p]\times[q]} \in \mathbb{T}^{p\times q}$ be two $p \times q$ tropical matrices. Fix for $1 \leqslant j \leqslant q$, $\varepsilon_j \in \mathbb{R}$, and set $\widetilde{A} = (\widetilde{a}_{ij})_{(i,j)\in[p]\times[q]} \in \mathbb{T}^{p\times q}$ and $\widetilde{B} = (\widetilde{b}_{ij})_{(i,j)\in[p]\times[q]} \in \mathbb{T}^{p\times q}$ with $\widetilde{a}_{ij} = a_{ij} + \varepsilon_j$ and $\widetilde{b}_{ij} = b_{ij} + \varepsilon_j$ for all $1 \leqslant i \leqslant p$ and $1 \leqslant j \leqslant q$. Then the (in)equality $A \odot y \rhd B \odot y$ of unknown $y \in \mathbb{T}^q$ has no nonzero solution if and only if the (in)equality $\widetilde{A} \odot \widetilde{y} \rhd \widetilde{B} \odot \widetilde{y}$ of unknown $\widetilde{y} \in \mathbb{T}^q$ has no nonzero solution.*

Lemma 2.2.10. *Let $\rhd \in \{\geqslant, =\}$. Let A and B be two $p \times q$ tropical matrices, and assume that A and B can both be written by block as lower-triangular matrices*

$$A = \begin{pmatrix} A^{(m)} & \mathbb{0} \\ * & * \end{pmatrix} \quad \textit{and} \quad B = \begin{pmatrix} B^{(m)} & \mathbb{0} \\ * & * \end{pmatrix}$$

with $A^{(m)}$ and $B^{(m)}$ two $m \times m$ square submatrices, with $0 < m \leqslant p, q$. Moreover, assume that the only solution to the equation $A^{(m)} \odot y^{(m)} \rhd B^{(m)} \odot y^{(m)}$ of unknown $y^{(m)} \in \mathbb{T}^m$ is $y^{(m)} = \mathbb{0}$. Then the equation $A \odot y \rhd B \odot y$ of unknown y has no solution in $\mathbb{R}^q$.

2.2.3 Proving the Tropical Positivstellensatz

Before detailing the proof of Theorem 2.2.1, notice that we can simply consider the case where all relations in system (2.6) are of the form $\rhd_i \in \{\geqslant, =\}$. Indeed, suppose the system $f^+(x) \rhd f^-(x)$ comprises relations of the form $f_i^+(x) > f_i^-(x)$ for $i \in I \subseteq \{1, \dots, k\}$. Let $\lambda > \mathbb{1} = 0$, and let us consider the transformed system $\mathcal{S}(\lambda)$, in which every relation

$$f_i^+(x) > f_i^-(x)$$

is replaced by

$$f_i^+(x) \geqslant \lambda \odot f_i^-(x) \enspace .$$

Let $\mathcal{M}^-_{\mathcal{E}'}(\lambda)$ denote the negative Macaulay matrix associated to $\mathcal{S}(\lambda)$, and let us take $\rhd'_j$ to be equal to $\geqslant$ for $j \in I$, and to $\rhd_j$ otherwise.

Then, the relation $\mathcal{M}^+_{\mathcal{E}'} \odot y \rhd \mathcal{M}^-_{\mathcal{E}'} \odot y$ holds iff $\mathcal{M}^+_{\mathcal{E}'} \odot y \rhd' \mathcal{M}^-_{\mathcal{E}'}(\lambda) \odot y$ holds for some $\lambda > \mathbb{1}$. Thus, if the theorem is proven for systems without strict inequalities, it follows that the latter relation is equivalent to the solvability of $\mathcal{S}(\lambda)$, which proves the result for systems which include strict inequalities.

In the following, we will therefore simply focus on systems such that $\rhd_i \in \{\geqslant, =\}$ for all $1 \leqslant i \leqslant k$.

A row content adapted to two-sided relations

We consider the collection $f = (f_1, \dots, f_k)$ of tropical polynomials, their associated Newton polytopes Q_i as in Section 2.2.1 and the constants r_i defined in Section 2.2.1. We also consider the maps $h_1, \dots, h_k$ defined as in Section 2.1.3. Moreover, we set

$$\widetilde{h} = h_1^{\square r_1} \,\square \cdots \square\, h_k^{\square r_k} \enspace ,$$

and note that $\widetilde{Q} = \operatorname{supp}(\widetilde{h})$.

Recall that now $\mathcal{E}$ is a Canny-Emiris set associated to the system $f^+ \rhd f^-$, that is $\mathcal{E} = (\widetilde{Q} + \delta) \cap \mathbb{Z}^n$, where δ is a generic vector in $V + \mathbb{Z}^n$ and V is the direction of the affine hull of $\widetilde{Q}$. Assume that there is no solution $x \in \mathbb{R}^n$ to the system $f^+(x) \rhd f^-(x)$, *i.e.* that for all $x \in \mathbb{R}^n$, there exists $1 \leqslant j \leqslant k$ such that

$$f_j^+(x) \not\rhd_j f_j^-(x) \enspace . \tag{2.8}$$

More precisely, (2.8) implies that either

$$f_j^-(x) > f_j^+(x) \enspace , \tag{2.9a}$$

or

$$\rhd_j \text{ is an equality} \quad \text{and} \quad f_j^-(x) < f_j^+(x) \enspace . \tag{2.9b}$$

Then for $p \in \mathcal{E}$, $(p - \delta, \widetilde{h}(p - \delta))$ is in the relative interior of a facet $\widetilde{F}$ of $\operatorname{hypo}(\widetilde{h}) = \sum_{i=1}^k r_i \operatorname{hypo}(h_i)$. This facet satisfies $\widetilde{F} = \mathcal{F}(x, \widetilde{h})$ for some $x \in V$, and then

$$\widetilde{F} = r_1 F_1 + \cdots + r_k F_k \enspace ,$$

with $F_i = \mathcal{F}(x, h_i)$. This means that we can write $p - \delta = r_1 q_1 + \cdots + r_k q_k$ with $(q_i, h_i(q_i)) \in F_i$ for all $1 \leqslant i \leqslant k$. Set j to be the maximal index such that (2.8) is satisfied. We have

$$r_j(q_j, h_j(q_j)) \in r_j F_j = F_j + \cdots + F_j$$

where the sum has r_j terms. Moreover, F_j is isomorphic to its projection $C_j := \mathcal{C}(x, h_j)$ on $\mathbb{R}^n$, and the set of extremal points of C_j is included in the finite set $\mathcal{C}(x, \omega_j)$, where ω_j is the coefficient map of the tropical polynomial f_j (see Section 2.1.3 and Observation 2.1.24(b)). Moreover, by assumption j and x satisfy (2.8), hence, in case (2.9a), one must have $\mathcal{C}(x, \omega_j) \subset \mathcal{A}_j^-$, because the maximum in the expression

$$\max_{\alpha \in \mathcal{A}_j} f_{j,\alpha} + \langle x, \alpha \rangle$$

cannot be attained by a monomial of f_j^+, and likewise in case (2.9b), one must have $\mathcal{C}(x, \omega_j) \subset \mathcal{A}_j^+$. Therefore, this means by definition of r_j that the cell C_j, as well as the facet F_j have dimension at most $r_j - 1$, and so does the facet $r_j F_j = F_j + \cdots + F_j$. Thus, applying Corollary 1.1.25 (of the Shapley-Folkman lemma), there exists $(q'_j, h_j(q'_j)) \in F_j$ and a_j an extremal point of C_j such that $r_j(q_j, h_j(q_j)) = (r_j - 1)(q'_j, h_j(q'_j)) + (a_j, h_j(a_j))$.

Moreover, since a_j is an extremal point of C_j, we know that in case (2.9a), one has $a_j \in \mathcal{A}_j^-$, while in case (2.9b), one has $a_j \in \mathcal{A}_j^+$.

We then define the *row content* of p in this context to be equal to a couple (j, a_j) satisfying the above properties. Note that the element $a_j \in \mathcal{A}_j$ satisfying the above condition need not be unique, but we will see in the proof that this does not matter. Since a_j is an extremal point of $\mathcal{C}(x, h_j)$, we have that $(a_j, h_j(a_j))$ is a vertex of $\operatorname{hypo}(h_j)$, which implies that $h_j(a_j) = f_{j,a_j}$, by Observation 2.1.28.

Now for any subset $\mathcal{E}'$ of $\mathbb{Z}^n$ containing $\mathcal{E}$, we can construct the matrices $\mathcal{M}^{\pm}_{\mathcal{E}\mathcal{E}'} = (m^{\pm}_{pp'})_{(p,p')\in\mathcal{E}\times\mathcal{E}'}$ similarly to Section 2.1.3, by setting the row p of $\mathcal{M}^{\pm}_{\mathcal{E}\mathcal{E}'}$ to be the row $(j, p - a_j)$ of $\mathcal{M}^{\pm}_{\mathcal{E}'}$, if (j, a_j) is the row content of p.

The proof of Theorem 2.2.1

In order to prove Theorem 2.2.1, we will make use of the following lemma.

Lemma 2.2.11. *Consider $h_1, \dots, h_k$ and $\widetilde{h}$ as defined in Section 2.1.3, and let $p \in \mathcal{E}$ and let (j, a_j) be its row content. Then for all $p' \in \mathcal{E}$ and $a'_j \in \mathbb{Z}^n$ such that $p' = p - a_j + a'_j$, we have*

$$\widetilde{h}(p'-\delta) \overset{(\dagger)}{\geqslant} \widetilde{h}(p-\delta) - h_j(a_j) + h_j(a'_j) \overset{(\ddagger)}{\geqslant} \widetilde{h}(p-\delta) - f^-_{j,a_j} + f^+_{j,a'_j} \ , \tag{2.10}$$

and moreover at least one of the two inequalities is strict.

Proof. By setting

$$\widehat{h}_j = \underbrace{h_1 \square \cdots \square h_1}_{r_1 \text{ terms}} \square \cdots \square \underbrace{h_j \square \cdots \square h_j}_{r_j - 1 \text{ terms}} \square \cdots \square \underbrace{h_k \square \cdots \square h_k}_{r_k \text{ terms}} \ ,$$

so that $\widetilde{h} = h_j \square \widehat{h}_j$, the exact same reasoning as in the proof of Lemma 2.1.30 shows that

$$\widetilde{h}(p'-\delta) \geqslant h_j(a'_j) + \widehat{h}_j(p - \delta - a_j) \ ,$$

and that if $p' = p$, which entails that $a'_j = a_j$, then this inequality is an equality. This implies inequality (†).

Inequality (‡) simply follows from the fact that $f^-_{j,a_j} = h_j(a_j)$ and $f^+_{j,a'_j} \leqslant h_j(a'_j)$.

We now show that either (†) or (‡) is strict. If $p' = p$, then inequality (‡) reduces to $f^-_{j,a_j} \geqslant f^+_{j,a_j}$ which is known to be strict. Now assume that $p' \neq p$, and suppose that the equality is achieved in (†). Then this implies that

$$(p'-\delta, \widetilde{h}(p'-\delta)) = (a'_j, h_j(a'_j)) + (p - \delta - a_j, \widehat{h}_j(p-\delta-a_j)) \ . \tag{2.11}$$

Now consider $x, x' \in \mathbb{R}^n$ such that $F = \mathcal{F}(x, \widetilde{h})$ and $F' = \mathcal{F}(x', \widetilde{h})$ are the facet in the interior of which $(p-\delta, \widetilde{h}(p-\delta))$ and $(p'-\delta, \widetilde{h}(p'-\delta))$ respectively lie. Then from Corollary 2.1.27, we have

$$(a'_j, h_j(a'_j)) \in \mathcal{F}(x', h_j) \quad \text{and} \quad (p-\delta-a_j, \widehat{h}_j(p-\delta-a_j)) \in \mathcal{F}(x', \widehat{h}_j) \ .$$

However, we also know from equality (2.11) and Corollary 2.1.27 that

$$(p-\delta-a_j, \widehat{h}_j(p-\delta-a_j)) \in \mathcal{F}(x, \widehat{h}_j) \ .$$

Moreover, by the Shapley-Folkman lemma (Lemma 1.1.24), we have that

$$r_j \mathcal{F}(x, h_j) = \mathcal{S} + (r_j - 1)\mathcal{F}(x, h_j)$$

with $\mathcal{S} := \{(\alpha, h_j(\alpha)) \in \mathbb{Z}^n \times \mathbb{R} : (\alpha, h_j(\alpha)) \text{ is a vertex of } \mathcal{F}(x, h_j)\}$ which is in particular a finite set, and thus

$$\mathcal{F}(x, \widetilde{h}) = \mathcal{S} + \mathcal{F}(x, \widehat{h}_j) \ .$$

Since $(p-\delta, h(p-\delta))$ is in the relative interior of $\mathcal{F}(x, \widetilde{h})$, and since δ was taken generic, then it follows that $(p-\delta-a_j, \widehat{h}_j(p-\delta-a_j))$ is in the relative interior of $\mathcal{F}(x, \widehat{h}_j)$, and that $\mathcal{F}(x, \widehat{h}_j)$ is a facet of $\operatorname{hypo}(\widehat{h}_j)$. Therefore, since $(p-\delta-a_j, \widehat{h}_j(p-\delta-a_j))$ is in both $\mathcal{F}(x, \widehat{h}_j)$ and $\mathcal{F}(x', \widehat{h}_j)$ and since it is in particular in the relative interior of the first face, we deduce that

$$\mathcal{F}(x, \widehat{h}_j) \subseteq \mathcal{F}(x', \widehat{h}_j) \ .$$

Since by Observation 2.1.24 (a), $\mathcal{F}(x',\widehat{h}_j)$ is a proper face of $\mathrm{hypo}(\widehat{h}_j)$, then it implies that it is also a facet of $\mathrm{hypo}(\widehat{h}_j)$, and thus that

$$\mathcal{F}(x,\widehat{h}_j) = \mathcal{F}(x',\widehat{h}_j) \enspace .$$

Hence, it follows from Observation 2.1.24 (b) that $x = x'$, and therefore $(a'_j, h_j(a'_j)) \in \mathcal{F}(x, h_j)$. Therefore, we have

$$h_j(a'_j) + \langle x, a'_j\rangle = \max_{q_j\in Q_j} h_j(q_j) + \langle x, q_j\rangle = \max_{\alpha\in\mathcal{A}_j} f_{j,\alpha} + \langle x,\alpha\rangle = f(x) \enspace ,$$

and since $f^-(x) > f^+(x)$, this implies that $a'_j \in A_j^-$, and thus

$$f^-_{j,a'_j} + \langle x, a'_j\rangle = f^-(x) > f^+(x) \geqslant f^+_{j,a'_j} + \langle x, a'_j\rangle \enspace ,$$

hence $f^+_{j,a'_j} < f^-_{j,a'_j} = h_j(a'_j)$, thus inequality (‡) is strict. □

Proof of Theorem 2.2.1. We keep the notation of Section 2.2.3. We first look at the case of systems in the form $f^+(x) \geqslant f^-(x)$, as the result for the general case will follow from this case. For any subset $\mathcal{E}'$ from $\mathbb{Z}^n$ containing $\mathcal{E}$, we can construct the submatrices $\mathcal{M}^+_{\mathcal{E}\mathcal{E}'} = (m^+_{pp'})_{(p,p')\in\mathcal{E}\times\mathcal{E}'}$ and $\mathcal{M}^-_{\mathcal{E}\mathcal{E}'} = (m^-_{pp'})_{(p,p')\in\mathcal{E}\times\mathcal{E}'}$ of respectively $\mathcal{M}^+_{\mathcal{E}'}$ and $\mathcal{M}^-_{\mathcal{E}'}$ similarly to Section 2.1.3, by setting $m^+_{pp'}$ to be the coefficient in p' of $X^{p-a_j}f_j^+$ where (j,a_j) is the row content of p, *i.e.*

$$m^\pm_{pp'} = \begin{cases} f^\pm_{j,a'_j} & \text{if} \quad a'_j \in \mathcal{A}_j^\pm \\ \mathbb{0} & \text{otherwise,}\end{cases}$$

where $a'_j \in \mathbb{Z}^n$ is such that $p' = p - a_j + a'_j$, as well as the matrices $\widetilde{\mathcal{M}}^\pm_{\mathcal{E}\mathcal{E}} = (\widetilde{m}^\pm_{pp'})_{(p,p')\in\mathcal{E}\times\mathcal{E}}$ defined by

$$\widetilde{m}^\pm_{pp'} = m^\pm_{pp'} - \widetilde{h}(p'-\delta)$$

for all $p, p' \in \mathcal{E}$. We now show that the matrix $\widetilde{\mathcal{M}}^+_{\mathcal{E}\mathcal{E}}$ is diagonally dominated by $\widetilde{\mathcal{M}}^-_{\mathcal{E}\mathcal{E}}$, which will lead to the desired result. Let $p, p' \in \mathcal{E}$. Then:

- ◇ if $p' \neq p$, then $a'_j \neq a_j$ and thus the inequality $\widetilde{m}^-_{pp} > \widetilde{m}^+_{pp'}$ is equivalent to the inequality

$$\widetilde{h}(p'-\delta) > \widetilde{h}(p-\delta) - f^-_{j,a_j} + f^+_{j,a'_j} \enspace , \tag{2.12}$$

  which follows directly from Lemma 2.2.11;

- ◇ if $p' = p$, then $a'_j = a_j$ and the inequality $\widetilde{m}^-_{pp} > \widetilde{m}^+_{pp'}$ is equivalent to the inequality

$$f^-_{j,a_j} > f^+_{j,a_j},$$

  which is satisfied since $f^-_{j,a_j} + \langle x, a_j\rangle = f^-(x) > f^+(x) \geqslant f^+_{j,a_j} + \langle x, a_j\rangle$.

Thus, the matrix $\widetilde{\mathcal{M}}^+_{\mathcal{E}\mathcal{E}}$ is diagonally dominated by $\widetilde{\mathcal{M}}^-_{\mathcal{E}\mathcal{E}}$, and therefore thanks to Lemma 2.2.7, this shows that the only solution to the inequation

$$\widetilde{\mathcal{M}}^+_{\mathcal{E}\mathcal{E}} \odot \widetilde{z} \geqslant \widetilde{\mathcal{M}}^-_{\mathcal{E}\mathcal{E}} \odot \widetilde{z}$$

of unknown $\widetilde{z} \in \mathbb{T}^{\mathcal{E}}$ is $\widetilde{z} = \mathbb{0}$.

Finally, using Lemma 2.2.9 and Lemma 2.2.10, we find that this implies that there cannot exist a vector $y \in \mathbb{R}^{\mathcal{E}'}$ such that $\mathcal{M}^+_{\mathcal{E}'} \odot y \geqslant \mathcal{M}^-_{\mathcal{E}'} \odot y$.

Now for the general case, assume that there is no solution in $\mathbb{R}^n$ to the system $f^+(x) \rhd f^-(x)$. Then, for all $x \in \mathbb{R}^n$, there exists $1 \leqslant j \leqslant k$ such that

$$f_j^+(x) \ntrianglerighteq_j f_j^-(x) \enspace .$$

We reiterate the construction of the matrices $\mathcal{M}^\pm_{\mathcal{E}\mathcal{E}}$ as above, as well as the renormalized matrices $\widetilde{\mathcal{M}}^\pm_{\mathcal{E}\mathcal{E}}$, with the only difference that for $p \in \mathcal{E}$ with row content (j, a_j), one might have either $a_j \in \mathcal{A}_j^+$ or $a_j \in \mathcal{A}_j^-$. In fact, applying (2.12) gives us this time that

$$\widetilde{m}^-_{pp} > \max_{p'\in\mathcal{E}} \widetilde{m}^+_{pp'} \iff a_j \in \mathcal{A}_j^- \quad \text{and} \quad \widetilde{m}^+_{pp} > \max_{p'\in\mathcal{E}} \widetilde{m}^-_{pp'} \iff a_j \in \mathcal{A}_j^+ \enspace .$$

In particular, this shows that the matrices $\widetilde{\mathcal{M}}^+_{\mathcal{E}\mathcal{E}}$ and $\widetilde{\mathcal{M}}^-_{\mathcal{E}\mathcal{E}}$ satisfy the condition (2.7) of Corollary 2.2.8. It thus follows from Lemma 2.2.11 that the matrices $\widetilde{\mathcal{M}}^+_{\mathcal{E}\mathcal{E}}$ and $\widetilde{\mathcal{M}}^-_{\mathcal{E}\mathcal{E}}$ satisfy the condition (2.7) of Corollary 2.2.8, hence the only solution to the system $\widetilde{\mathcal{M}}^+_{\mathcal{E}\mathcal{E}} \odot \widetilde{z} \rhd \widetilde{\mathcal{M}}^-_{\mathcal{E}\mathcal{E}} \odot \widetilde{z}$ of unknown $\widetilde{z} \in \mathbb{T}^{\mathcal{E}}$ is $\widetilde{z} = \mathbb{0}$, and using Lemma 2.2.9 and Lemma 2.2.10, we obtain there is no solution $y \in \mathbb{R}^{\mathcal{E}'}$ to the equation $\mathcal{M}^+_{\mathcal{E}'} \odot y \rhd \mathcal{M}^-_{\mathcal{E}'} \odot y$. □

Remark 2.2.12. Note that in the previous construction, the matrices $\mathcal{M}^{\pm}_{\mathcal{E}\mathcal{E}}$ need not necessarily be submatrices of the matrices $\mathcal{M}^{\pm}_{\mathcal{E}}$, as the construction of the row content in this case could allow for one row of the Macaulay matrix to appear multiple times in $\mathcal{M}^{\pm}_{\mathcal{E}\mathcal{E}}$. This however does not have any impact on the outcome of the proof.

2.2.4 The case of hybrid systems

In fact, Theorem 2.2.1 can be generalized even further so as to also include relations of the form $f_i(x)\ \nabla\ \mathbb{0}$. More precisely, let $f_1^{\pm}, \dots, f_k^{\pm}$ be a collection of pairs of tropical polynomials and let $f_{k+1}, \dots, f_\ell$ be a second collection of tropical polynomials. Keeping the notation of Section 2.2.1, the system we consider is the following

$$\begin{cases} f_i^+(x) \rhd_i f_i^-(x) & \text{for } 1 \leqslant i \leqslant k \\ f_i(x)\ \nabla\ \mathbb{0} & \text{for } k+1 \leqslant i \leqslant \ell \end{cases} , \qquad (\mathcal{S}_{\text{pol}})$$

with $\rhd_i \in \{\geqslant, =, >\}$ for all $1 \leqslant i \leqslant k$. Recall also that

$$r_i = \begin{cases} \dim(\operatorname{aff}(\mathcal{A}_i^-)) + 1 & \text{if } \rhd_i \in \{\geqslant, >\} \\ \max\big(\dim(\operatorname{aff}(\mathcal{A}_i^-)), \dim(\operatorname{aff}(\mathcal{A}_i^+))\big) + 1 & \text{if } \rhd_i \in \{=\} \end{cases} .$$

We define Canny-Emiris subsets $\mathcal{E}$ of $\mathbb{Z}^n$ associated to system $(\mathcal{S}_{\text{pol}})$ to be sets of the form

$$\mathcal{E} := \left(\widetilde{Q} + \delta\right) \cap \mathbb{Z}^n \quad \text{with} \quad \widetilde{Q} = r_1 Q_1 + \cdots + r_k Q_k + Q_{k+1} + \cdots + Q_\ell ,$$

where δ is a generic vector in $V + \mathbb{Z}^n$, with V the direction of the affine hull of $\widetilde{Q}$. Then we have the following result.

Theorem 2.2.13 (Sparse tropical Positivstellensatz for hybrid systems)**.** *The system $(\mathcal{S}_{\text{pol}})$ has a solution $x \in \mathbb{R}^n$ if and only if the linear tropical system*

$$\begin{cases} \mathcal{M}^+_{\mathcal{E}'} \odot y \rhd \mathcal{M}^-_{\mathcal{E}'} \odot y \\ \mathcal{M}_{\mathcal{E}'} \odot y\ \nabla\ \mathbb{0} \end{cases} \qquad (\mathcal{S}_{\text{lin}})$$

has a solution $y \in \mathbb{R}^{\mathcal{E}'}$, where $\mathcal{M}^+, \mathcal{M}^-$ is the pair of Macaulay matrices associated to system $f_i^+(x) \rhd_i f_i^-(x)$ for $1 \leqslant i \leqslant k$, $\mathcal{M}$ is the Macaulay matrix associated to system $f_i(x)\ \nabla\ \mathbb{0}$ for $k+1 \leqslant i \leqslant \ell$, and $\mathcal{E}'$ is any subset of $\mathbb{Z}^n$ containing a nonempty Canny-Emiris subset $\mathcal{E}$ of $\mathbb{Z}^n$ associated to the system $(\mathcal{S}_{\text{pol}})$.

Proof. This proof is mainly combination of the proofs of Theorems 2.1.5 and 2.2.1, so we will just give the outline of the proof and skip the details to avoid repetition. Moreover, as previously, the existence of a root $x \in \mathbb{R}^n$ to the polynomial easily implies the existence of an element $y \in \mathbb{R}^{\mathcal{E}'}$ solution of the system $(\mathcal{S}_{\text{lin}})$, given by the Veronese embedding.

Assume that system $(\mathcal{S}_{\text{pol}})$ does not have a root in $\mathbb{R}^n$. Then this means that for all $x \in \mathbb{R}^n$, either $f_j(x) \not\nabla \mathbb{0}$ for some $k+1 \leqslant j \leqslant \ell$, or $f_j^+(x) \not\rhd_j f_j^-(x)$ for some $1 \leqslant i \leqslant k$ (or possibly both cases can happen at the same time). This allows us to construct a notion of row content for this case in the following way.

First, we define the maps $h_1, \dots, h_\ell$ as in the proofs of Theorems 2.1.5 and 2.2.1, and set

$$\widetilde{h} := h_1^{\square r_1} \square \cdots \square h_k^{\square r_k} \square h_{k+1} \square \cdots \square h_\ell .$$

Then for all $p \in (\widetilde{Q} + \delta) \cap \mathbb{Z}^n$, the point $(p - \delta, \widetilde{h}(p - \delta))$ lies in the relative interior of a facet $\widetilde{F}$ of $\operatorname{hypo}(\widetilde{h})$, and this facet can be written as

$$\widetilde{F} = r_1 F_1 + \cdots + r_k F_k + F_{k+1} + \cdots + F_\ell ,$$

where for all $1 \leqslant j \leqslant \ell$, F_j is a facet of $\operatorname{hypo}(h_j)$, and there exists an element x of V such that $\widetilde{F} = \mathcal{F}(\widetilde{h}, x)$ and for all $1 \leqslant j \leqslant \ell$, $F_j = \mathcal{F}(h_j, x)$. Then we have the following two possible cases *(i)* if there exists $k+1 \leqslant j \leqslant \ell$ such that $f_j(x) \not\nabla \mathbb{0}$, then F_j is a singleton $\{a_j\}$, and we set the row content of p to be the pair (j, a_j), where such an index j is taken to be maximal;

(ii) if $f_j(x) \nabla \mathbb{0}$ for all $k+1 \leqslant j \leqslant \ell$, then since x is not a root of ($\mathcal{S}_{\text{pol}}$), this implies that there exists $1 \leqslant j \leqslant k$ such that $f_j^+(x) \not\triangleright_j f_j^-(x)$, and this time we can construct the row content of x using the Shapley-Folkman lemma as in the proof of Theorem 2.2.1.

Now for $p \in \mathcal{E}$ with row content (j, a_j), if $1 \leqslant j \leqslant k$, then we apply Lemma 2.2.11 to obtain that for all $p' \in \mathcal{E}$ and $a'_j \in \mathbb{Z}^n$ such that $p' = p - a_j + a'_j$,

$$\widetilde{h}(p' - \delta) > \widetilde{h}(p - \delta) - f^-_{j,a_j} + f^+_{j,a'_j} \enspace ,$$

and if $k+1 \leqslant j \leqslant \ell$, then we apply Lemma 2.1.30 to obtain that for all $p' \in \mathcal{E}$ and $a'_j \in \mathbb{Z}^n$ such that $p' = p - a_j + a'_j$,

$$\widetilde{h}(p' - \delta) > \widetilde{h}(p - \delta) - f_{j,a_j} + f_{j,a'_j} \enspace ,$$

with equality if and only if $p' = p$. We denote $\mathcal{E}_1$ the set of $p \in \mathcal{E}$ satisfying the first condition, and $\mathcal{E}_2$ the set of $p \in \mathcal{E}$ satisfying the second condition, such that $\mathcal{E} = \mathcal{E}_1 \sqcup \mathcal{E}_2$.

Thus, we can construct as previously matrices $\mathcal{M}^+_{\mathcal{E}_1\mathcal{E}}$, $\mathcal{M}^-_{\mathcal{E}_1\mathcal{E}}$ and $\mathcal{M}_{\mathcal{E}_2\mathcal{E}}$ such that $\mathcal{M}^-_{\mathcal{E}_1\mathcal{E}}$ and $\mathcal{M}^+_{\mathcal{E}_1\mathcal{E}}$ satisfy the condition (2.7) of Corollary 2.2.8, and $\mathcal{M}_{\mathcal{E}_2\mathcal{E}}$ is diagonally dominant in the tropical sense.

From this point, it can easily be shown using a combination of Lemma 2.1.16 and 2.2.7 that the only solution $z \in \mathbb{T}^{\mathcal{E}}$ to the system

$$\begin{cases} \mathcal{M}^+_{\mathcal{E}_1\mathcal{E}} \odot z \rhd \mathcal{M}^-_{\mathcal{E}_1\mathcal{E}} \odot z \\ \mathcal{M}_{\mathcal{E}_2\mathcal{E}} \odot z \,\nabla\, \mathbb{0} \enspace , \end{cases}$$

is $z = \mathbb{0}$, and thus it follows that there is no solution $y \in \mathbb{R}^{\mathcal{E}'}$ to system ($\mathcal{S}_{\text{lin}}$). □

Chapter 3

A speedup of the value iteration to detect the solvability of tropical polynomial systems

In this chapter, we lay the groundwork to tackle the explicit resolution of tropical polynomial systems by proposing with Algorithm 1 an acceleration of the classical value iteration algorithm introduced by Zwick and Paterson in [ZP96]. This algorithm determines whether all coordinates of the vector of values of a given mean payoff game are greater than or equal to 0, and thus by virtue of Corollary 1.5.31, one can use this algorithm in order to check whether a given tropical linear system has a solution over the tropical torus. In particular, following from the tropical Nullstellensatz and Positivstellensatz presented in Chapter 2, this means that this algorithm can more broadly be used in order to determine the solvability of any tropical polynomial system. In the remainder of this manuscript, we shall specifically focus on systems of weak tropical polynomial inequalities. Besides readability considerations, this is sufficient because all the other types of tropical polynomial equations or inequations can be reduced to systems with only weak inequalities. This reduction is trivial for standard tropical polynomial equations $f(x) \nabla \mathbb{0}$ as well as two-sided inequations $f^+(x) = f^-(x)$, and the reduction from two-sided strict inequalities $f^+(x) > f^-(x)$ to two-sided weak inequalities $f^+(x) \geqslant f^-(x)$ uses arguments of short rationals (see [AFG$^+$14] for more details).

3.1 Preliminary results on Shapley operators

In this first section, we present some preliminary vocabulary and results relating to Shapley operators and mean payoff games, that will play a central role in the proper statement as well as in the proof of Algorithm 1 below.

Definition 3.1.1. Let $T : (\mathbb{R} \cup \{\pm\infty\})^J \to (\mathbb{R} \cup \{\pm\infty\})^J$ be an order-preserving additively homogeneous operator. Then one denotes by $\underline{T} : (\mathbb{R} \cup \{\pm\infty\})^J \to (\mathbb{R} \cup \{\pm\infty\})^J$ the operator defined for all $u \in (\mathbb{R} \cup \{\pm\infty\})^J$ by

$$\underline{T}(u) = u \wedge T(u) \enspace ,$$

where $\wedge$ denotes the infimum for the partial order, or equivalently the coordinatewise minimum of vectors.

Remark 3.1.2. The operator $\underline{T}$ is also order-preserving and additively homogeneous. Moreover, if $T = A^\sharp B$, then $\underline{T}$ can be expressed as a Shapley operator of the form $\underline{T} = \underline{A}^\sharp \underline{B}$ with $\underline{A}$ and $\underline{B}$ given by block as

$$\underline{A} = \begin{pmatrix} A \\ \boldsymbol{I} \end{pmatrix} \quad \text{and} \quad \underline{B} = \begin{pmatrix} B \\ \boldsymbol{I} \end{pmatrix} \enspace ,$$

where $\boldsymbol{I}$ denotes the $|J| \times |J|$ tropical identity matrix, with diagonal entries equal to 0 and all other entries set to $-\infty$. In particular, this implies that $\underline{T}$ satisfies the same properties as T, and thus all results that apply to T can also be applied to $\underline{T}$.

Moreover, notice that it follows from the Collatz-Wielandt property that $\underline{\chi}(\underline{T}) = \underline{\chi}(T) \wedge 0$, and thus $\chi(T) \geqslant 0$ if and only if $\chi(\underline{T}) \geqslant 0$ if and only if $\chi(\underline{T}) \equiv 0$.

The operator $\underline{T}$ will be of interest because whenever T satisfies Assumption 1.5.10 (a), *i.e.* preserves $(\mathbb{R} \cup \{+\infty\})^J$, then $\underline{T}$ satisfies both Assumption 1.5.10 (a) and (b), and thus preserves $\mathbb{R}^J$. Therefore, all the results of the previous section will be applicable to $\underline{T}$. We shall also take interest in the following operator.

Definition 3.1.3. Let $T : (\mathbb{R} \cup \{\pm\infty\})^J \to (\mathbb{R} \cup \{\pm\infty\})^J$ be an order-preserving additively homogeneous operator, and assume moreover that T preserves $(\mathbb{R} \cup \{+\infty\})^J$. Then the *damped Krasnoselskii-Mann operator* is the operator $T_{\mathsf{KM}} : (\mathbb{R} \cup \{+\infty\})^J \to (\mathbb{R} \cup \{+\infty\})^J$ defined for all $u \in (\mathbb{R} \cup \{+\infty\})^J$ by

$$T_{\mathsf{KM}}(u) = \frac{1}{2}(u + T(u)) \enspace .$$

Remark 3.1.4. The operator T_{KM} is also order-preserving and additively homogeneous. Moreover, whenever T preserves $\mathbb{R}^J$, then $\chi(T_{\mathsf{KM}}) = \frac{\chi(T)}{2}$. Indeed, let $u, \eta \in \mathbb{R}^J$ be such that $T(u + s\eta) = u + (s+1)\eta$ for all $s \in \mathbb{R}$ big enough, say $s \geqslant s^*$, then for all $s \geqslant 2s^*$,

$$T_{\mathsf{KM}}\left(u + s\frac{\eta}{2}\right) = \frac{1}{2}\Big(u + s\frac{\eta}{2} + \underbrace{T\left(u + s\frac{\eta}{2}\right)}_{=u+\left(\frac{s}{2}+1\right)\eta}\Big) = u + (s+1)\frac{\eta}{2} \enspace ,$$

hence the result by Theorem 1.5.13.

In particular, if T preserves $(\mathbb{R} \cup \{+\infty\})^J$, then $\underline{T}$ preserves $\mathbb{R}^J$, and thus $\chi(\underline{T}_{\mathsf{KM}}) = \frac{\chi(\underline{T})}{2}$, and thus from the previous remark, $\underline{\chi}(\underline{T}_{\mathsf{KM}}) = \frac{\underline{\chi}(T)}{2} \wedge 0$.

Definition 3.1.5. Let $T : (\mathbb{R} \cup \{\pm\infty\})^J \to (\mathbb{R} \cup \{\pm\infty\})^J$ be an order-preserving and additively homogeneous operator. Then, for a subset $\mathscr{D}$ of J, one defines the *reduced operator* $T_{\mathscr{D}} : (\mathbb{R} \cup \{\pm\infty\})^{\mathscr{D}} \to (\mathbb{R} \cup \{\pm\infty\})^{\mathscr{D}}$ by

$$T_{\mathscr{D}}(u) = (\pi_{\mathscr{D}} \circ T \circ \iota_{\mathscr{D}})(u) \quad \text{for all } u \in (\mathbb{R} \cup \{\pm\infty\})^{\mathscr{D}}$$

with $\iota_{\mathscr{D}} : (\mathbb{R} \cup \{\pm\infty\})^{\mathscr{D}} \to (\mathbb{R} \cup \{\pm\infty\})^J$ given by

$$\iota_{\mathscr{D}}(v) = w \text{ with } w_j = \begin{cases} v_j & \text{if } j \in \mathscr{D} \\ +\infty & \text{otherwise} \end{cases} \qquad \text{for all } v = (v_j)_{j \in \mathscr{D}} \in (\mathbb{R} \cup \{\pm\infty\})^{\mathscr{D}} \enspace ,$$

and $\pi_{\mathscr{D}} : (\mathbb{R} \cup \{\pm\infty\})^J \to (\mathbb{R} \cup \{\pm\infty\})^{\mathscr{D}}$ given by

$$\pi_{\mathscr{D}}(w) = (w_j)_{j \in \mathscr{D}} \quad \text{for all } w = (w_j)_{j \in J} \in (\mathbb{R} \cup \{\pm\infty\})^J \enspace ;$$

Definition 3.1.6. A *dominion* for the minimizer is a nonempty subset $\mathscr{D}$ of J such that the associated reduced operator $T_{\mathscr{D}}$ preserves $\mathbb{R}^{\mathscr{D}}$, *i.e.* such that $T_{\mathscr{D}}(\mathbb{R}^{\mathscr{D}}) \subseteq \mathbb{R}^{\mathscr{D}}$.

Remark 3.1.7. The above definition of dominions for the minimizer is dual to [AGKS22, Definition 14] which defines dominions for the maximizer. The following game theoretical interpretation of dominions is also given: a subset $\mathscr{D} \subseteq J$ of states is a dominion for one of the players if whenever the game starts at a state $j \in \mathscr{D}$, then this player can force the game to remain within the set $\mathscr{D}$, regardless of the strategy of the opposing player. They thus coincide with the dominions considered in [AGH20]. Note that in [JPZ08] states of dominions are also required to be winning for the player.

In particular, if $\mathscr{D}$ is a dominion for the minimizer, then one can define the *subgame induced by $\mathscr{D}$ for the minimizer* as follows. The set of states of the minimizer is the subset $\mathscr{D} \subseteq J$ and the set of actions from any state (in $\mathscr{D}$) consists in all the actions of the initial game leading to a state from which only states in $\mathscr{D}$ can be attained; the set of states of the maximizer is the subset $\mathscr{C} \subseteq I$ of states of the game that are accessible by one of the actions of the minimizer describes above, and the set of actions from any state in $\mathscr{C}$ consists in all the possible actions of the initial game (that will by construction lead to a state in $\mathscr{D}$). The payoffs are kept the same as in the initial game. It can then be shown that if $\mathscr{D}$ is a dominion for the maximizer of a mean payoff game with Shapley operator T, then $T_{\mathscr{D}}$ is the Shapley operator of the subgame induced by $\mathscr{D}$ for the minimizer.

Recall that if T preserves $\mathbb{R}^n$, $\chi(T)$ denotes the vector of values of the mean payoff game associated to the operator T. We take interest in the set of losing initial states, *i.e.* of initial states $j \in J$ such that $\chi_j(T) < 0$.

Lemma 3.1.8. *Assume that $\underline{\chi}(T) < 0$ and let $\mathscr{D} = \{j \in J : \chi_j < 0\}$. Then $\mathscr{D}$ is a dominion for the minimizer player. Moreover, $\overline{\chi}(T_{\mathscr{D}}) < 0$.*

Proof. If the game starts at a state $j_0 \in \mathscr{D}$, then by definition, the value $\chi_{j_0}(T)$ associated to this position is strictly negative, which means that by playing optimaly, the maximizer can only ensure a strictly negative mean payoff. This means that if the minimizer plays optimally, then for all states $j \in J$ that are visited during the game, one must also have $\chi_j(T) < 0$, *i.e.* $j \in \mathscr{D}$, otherwise the maximizer could ensure a nonnegative mean payoff. Thus $\mathscr{D}$ is indeed a dominion for the minimizer.

The proof of the inequality $\overline{\chi}(T_{\mathscr{D}}) < 0$ follows readily from the game theoretical interpretation of the operator $T_{\mathscr{D}}$ in the above remark, as if there was a state $j_0 \in \mathscr{D}$ for which $\chi_{j_0}(T_{\mathscr{D}})$ was nonnegative, then this would mean that the maximizer player, starting from this state, would have a strategy ensuring a nonnegative mean payoff per turn in the reduced games, and thus *a fortiori* in the initial game, as all the states $J \setminus \mathscr{D}$ have a nonnegative value, so if the minimizer allows the maximizer to ever reach a state in $j \in J \setminus \mathscr{D}$, then the maximizer player will also be able to guarantee a nonnegative mean payoff from state j. □

Finally, we state the following technical lemma, which gives a bound on the norm of the subharmonic and superharmonic vectors associated to extremal nonlinear eigenvalues of a Shapley operator $T = A^\sharp B$, still under Assumption 1.5.10. By the residuation property, it entails a 'short solution property' (to be compared with the 'small model property' from [BNR08, Lemma 2], as well as [AGK11b, Proposition 10]) of the tropical linear system $A \odot y \leqslant B \odot y$, on which the construction of the mean payoff game oracle deciding the solvability of a tropical linear system will rely.

Lemma 3.1.9. *Let $T = A^\sharp B$ be an order-preserving additively homogeneous and piecewise affine self-map of $(\mathbb{R} \cup \{\pm\infty\})^J$, with $A, B \in \mathbb{T}^{I\times J}$. Assume that T preserves $\mathbb{R}^J$. Let $\lambda = \overline{\chi}(T)$. Then there exists $u \in \mathbb{R}^J$ such that $T(u) \leqslant \lambda + u$ with*

$$\|u\|_{\mathcal{H}} \leqslant (|J| - 1)W \ ,$$

where $\|\cdot\|_{\mathcal{H}}$ denotes the Hilbert seminorm defined by $\|u\|_{\mathcal{H}} = \max_{j\in J}(u_j) - \min_{j\in J}(u_j)$, and

$$W := \max_{\substack{(i,j,k)\in I\times J\times J \\ a_{ij}, b_{ik} \neq -\infty}} (-a_{ij} + b_{ik}) - \min_{\substack{(i,j,k)\in I\times J\times J \\ a_{ij}, b_{ik} \neq -\infty}} (-a_{ij} + b_{ik}) \leqslant 4r^\infty \ .$$

Similarly, let $\mu = \underline{\chi}(T)$. Then there exists $u \in \mathbb{R}^J$ such that $T(u) \geqslant \mu + u$ with the same bound on the Hilbert seminorm of u.

Proof. The existence of $u \in \mathbb{R}^J$ such that $T(u) \leqslant \lambda + u$ is ensured by definition of the Collatz-Wielandt number as well as by Theorem 1.5.27 and Remark 1.5.28. In particular, for such a $u \in \mathbb{R}^J$, the previous inequality is equivalent to

$$\forall j \in J, \quad \min_{i\in I}(-a_{ij} + \max_{k\in J}(b_{ik} + u_k)) \leqslant \lambda + u_j \ .$$

Now, from [AGG12, Theorem 2.13], we know that there exists an optimal positional strategy $\sigma^* : J \to I$ for the minimizer, and moreover that for this strategy, the previous inequality becomes

$$\forall j \in J, \quad -a_{\sigma^*(j)j} + \max_{k\in J}(b_{\sigma^*(j)k} + u_k) \leqslant \lambda + u_j \ .$$

Thus, letting $C = (c_{jk})_{(j,k)\in J^2}$ with $c_{jk} = -a_{\sigma^*(j)j} + b_{\sigma^*(j)k}$ for all $j, k \in J$, the previous inequality can be rewritten as the tropical linear inequality

$$\left(\forall j \in J, \quad \max_{k\in J}(c_{jk} + u_k) \leqslant \lambda + u_j\right) \iff C \odot u \leqslant \lambda \odot u \ .$$

Finally, let $((-\lambda)\odot C)^* = \bigoplus_{\ell=0}^{|J|-1}((-\lambda)\odot C)^{\odot \ell}$ be the metric closure or *Kleene star* of the matrix $(-\lambda)\odot C$, then a standard result of tropical spectral theory (see for instance [But10, Theorem 1.6.18]) entails that the vector

$$u = ((-\lambda)\odot C)^* \odot \begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix}$$

satisfies the inequality $C \odot u \leqslant \lambda \odot u$. Hence u satisfies $T(u) \leqslant \lambda + u$. Since the vector u is nonnegative, one has

$$\|u\|_{\mathcal{H}} \leqslant \|u\|_\infty \leqslant (|J| - 1)\left(-\lambda + \max_{j,k\in J} c_{jk}\right) \ .$$

Moreover, since $C \odot u \leqslant \lambda \odot u$, we obtain $\max_{k\in J} c_{jk} + \min_{k\in J} u_k \leqslant \lambda + u_j$ for all $j \in J$. Taking the infimum over j and using that $u \in \mathbb{R}^J$, we deduce that $\min_{j\in J}\max_{k\in J} c_{jk} \leqslant \lambda$. Therefore, using that for all $j \in J$, $i = \sigma^*(j)$ is such that $a_{ij} \neq -\infty$, one obtains the inequality

$$\begin{aligned}
\|u\|_{\mathcal{H}} &\leqslant (|J|-1)\left(-\min_{j\in J}\max_{k\in J} c_{jk} + \max_{j,k\in J} c_{jk}\right) \\
&\leqslant (|J|-1)\left(-\min_{\substack{(j,k)\in J\times J\\ c_{jk}\neq -\infty}} c_{jk} + \max_{j,k\in J} c_{jk}\right) \\
&\leqslant (|J|-1)\left(-\min_{\substack{(i,j,k)\in I\times J\times J\\ b_{jk}\neq -\infty}} (-a_{ij}+b_{ik}) + \max_{\substack{(i,j,k)\in I\times J\times J\\ a_{ij}\neq -\infty}} (-a_{ij}+b_{ik})\right) \\
&\leqslant (|J|-1)W \enspace .
\end{aligned}$$

The proof of the second assertion of the lemma relies on Corollary 1.5.29 and is dual, though not identical, so we will just roughly give the sketch. As above, the Collatz-Wielandt property entails the existence of $u \in \mathbb{R}^J$ such that $T(u) \geqslant \mu + u$, *i.e.*

$$\forall j \in J, \quad \min_{i\in I}(-a_{ij} + \max_{k\in J}(b_{ik}+u_k)) \geqslant \mu + u_j \enspace .$$

Now we can similarly consider an optimal positional strategy $\tau^* : I \to J$ for the maximizer, to transfrom the previous inequality into

$$\forall j \in J, \quad \min_{i\in I}(-a_{ij} + b_{i\tau^*(i)} + u_{\tau^*(i)}) \geqslant \mu + u_j \enspace .$$

However, one can rewrite the lefthandside of the above inequality as

$$-\max_{i\in I}(a_{ij} - b_{i\tau^*(i)} - u_{\tau^*(i)}) = -\max_{k\in\tau^*(I)} \max_{\substack{i\in I\\ \tau^*(i)=k}} (a_{ij} - b_{ik} - u_k) \enspace .$$

Therefore, letting $C = (c_{jk})_{(j,k)\in J^2}$ with

$$c_{jk} = \begin{cases} \max\limits_{\substack{i\in I\\ \tau^*(i)=k}} (a_{ij} - b_{ik}) & \text{if } k \in \tau^*(I) \\ -\infty & \text{otherwise,} \end{cases}$$

for all $j, k \in J$, the previous inequality can be rewritten as the tropical linear inequality

$$\left(\forall j \in J, \quad \max_{k\in J}(c_{jk} - u_k) \leqslant -\mu - u_j\right) \iff C \odot (-u) \leqslant (-\mu) \odot (-u) \enspace ,$$

and finally, one obtains as above

$$\|u\|_{\mathcal{H}} \leqslant (|J|-1)\left(-\min_{\substack{j,k\in J\\ c_{jk}\neq -\infty}} c_{jk} + \max_{j,k\in J} c_{jk}\right) \leqslant (|J|-1)W \enspace . \qquad \square$$

Remark 3.1.10. In [AGKS22, Lemma 20] a similar bound is obtained for stochastic zero-sum games, but under the condition that $\overline{\chi}(T) = \underline{\chi}(T)$. Under this restriction, applying the bound of [AGKS22] to deterministic games would lead to an additional factor 4 in the bound of $\|u\|_{\mathcal{H}}$ with respect to the one in Lemma 3.1.9.

Corollary 3.1.11. *Let $\lambda = \overline{\chi}(T)$ and $\mu = \underline{\chi}(T)$. Then, for all $N \in \mathbb{N}$, one has*

$$-(|J|-1)W + N\mu \leqslant T^N(0) \leqslant (|J|-1)W + N\lambda \enspace .$$

Proof. From Lemma 3.1.9, we know that there exists $u \in \mathbb{R}^J$ such that $T(u) \leqslant u + \lambda$ with $\|u\|_{\mathcal{H}} \leqslant (|J|-1)W$.

Thus one has for all $N \in \mathbb{N}$,

$$\begin{aligned} T^N(0) = T^N(-u+u) &\leqslant T^N\left(\max_{j\in J}(-u_j) + u\right) \\ &= \max_{j\in J}(-u_j) + T^N(u) \\ &\leqslant \max_{j\in J}(-u_j) + u + N\lambda \\ &\leqslant \underbrace{\max_{j\in J}(-u_j) + \max_{j\in J}(u_j)}_{=\|u\|_{\mathcal{H}}} + N\lambda \leqslant (|J|-1)W + N\lambda \ , \end{aligned}$$

where we used that T is order-preserving to get the first inequality, and that it is additively homogeneous to obtain the equality that follows.

The proof of the second inequality is dual. □

3.2 The value iteration algorithm with widening

In this section, we now present how to construct by means of mean payoff games an oracle capable of deciding the solvability of a tropical linear system, and thus of a tropical polynomial system. The resulting oracle relies on Algorithm 1 below, which consists in an acceleration of the classical value iteration algorithm. Two main ideas reside together at the core of this speed up : the first one is the idea of applying Algorithm 1 to the Krasnoselskii-Mann operator T_{KM} to ensure the termination of the algorithm. More general results on the convergence of Krasnoselskii-Mann iterates are detailed in Chapter 5, sheding a light on the use of the Krasnoselskii-Mann damping. The second idea is the introduction of a 'widening step', which consists in an easy check allowing one to stop the algorithm in a quicker way in unfeasible cases.

As mentioned in Section 1.5.4, the standard value iteration algorithm of Zwick and Paterson takes a time of $\mathcal{O}(|J|^2 r^{\infty})$, but by the nature of the algorithm, this 'worst-case' bound is in fact always achieved on every instance. This entails that this algorithm is inadapted to the large scale instances arising from the linearization of tropical polynomial systems. The present refinement of the vanilla value iteration algorithm, exploiting the idea of 'widening' together with Krasnolselskii-Mann damping, accelerates the termination of the algorithm on average, allowing one to construct a more efficient mean payoff game oracleThis improved algorithm was first presented in [ABG23a]. Moreover, the precise mathematical intuition behind the use of the Krasnoselskii-Mann damping will be detailed in Chapter 5.

Consider a system of tropical linear inequalities $A \odot y \leqslant B \odot y$, where $A, B \in \mathbb{T}^{I\times J}$ are two tropical matrices. In the previous subsection, we recalled in Corollary 1.5.31 the main result of [AGG12], which states that this system has a solution $y \in \mathbb{R}^J$ if and only if all the initial positions of a mean-payoff game associated with the Shapley operator $T = A^\sharp B$ are winning.

We propose Algorithm 1 for deciding the solvability of a tropical linear system, where we recall that for $u, v \in (\mathbb{R}\cup\{+\infty\})^J$, $v \ll u$ means that for all $j \in J$ such that $u_j < +\infty$, one has $v_j < u_j$. This algorithm exploits the classical idea of value iteration for mean-payoff games (see [ZP96, AGKS22]), with the added introduction of a *widening* step, consisting in the construction of the vector $\hat{u}$ at line 11, which together with the Krasnoselskii-Mann damping, allows for a quicker test of unfeasibility.

The appeal of the proposed algorithm indeed resides in the fact that even though the worst-case complexity is the same as for the classical algorithm proposed by Zwick and Paterson, the widening step gives an exit case for the algorithm that is usually reached quickly, lowering significantly the average complexity of the algorithm, as opposed to the usual value iteration algorithm whose average complexity matches its worst-case complexity.

We aim to apply Algorithm 1 to the Krasnoselskii-Mann operator T_{KM}. We thus consider the sequence $(u^N)_{N\in\mathbb{N}} \in (\mathbb{R}^J)^{\mathbb{N}}$ defined by

$$\begin{cases} u^0 = (0, \dots, 0) \in \mathbb{R}^J \\ u^{N+1} = \underline{T_{\mathsf{KM}}}(u^N) \quad \text{for all } N \geqslant 0 \ . \end{cases}$$

The aims of iterating over the operator $\underline{T_{\mathsf{KM}}}$ instead of the operator T are twofold. First of all, the use of the Krasnoselskii-Mann operator empirically accelerates the detection of the unfeasibility cases. In fact, it can be proven that $u^{N+1} - u^N \underset{N\to+\infty}{\longrightarrow} \frac{\chi(T)}{2}$. Moreover, taking the operator $\underline{T_{\mathsf{KM}}}$ instead of just T_{KM} forces the sequence $(u^N)_{N\in\mathbb{N}}$ to be decreasing, which guarantees the existence of a timeout before which the algorithm will terminate. Using the results from the previous subsection, we now prove the termination and correctness of our algorithm.

Algorithm 1: Value iteration algorithm with widening.

input: T a Shapley operator from $(\mathbb{R} \cup \{+\infty\})^J$ to $(\mathbb{R} \cup \{+\infty\})^J$
$\varepsilon > 0$ the approximation error for comparisons
N^* a timeout on the number of iterations which guarantees the existence of a solution whenever reached

output: Decides the feasibility of the system $u \leqslant T(u)$ in $\mathbb{R}^J$

`/* Initialization */`

$u := 0 \in \mathbb{R}^J$
$v := 0 \in \mathbb{R}^J$
$N := 0$
repeat
  `/* Value iteration step */`
  $u := v$
  $v := u \wedge T(u)$
  $N := N + 1$
  `/* Widening step */`
  $I := \{i : v_i \geqslant -\varepsilon + u_i\}$
  $\hat{u} := (\hat{u}_i) \in (\mathbb{R} \cup \{+\infty\})^m$ with $\begin{cases} \hat{u}_i = +\infty & \text{if } i \in I \\ \hat{u}_i = u_i & \text{otherwise} \end{cases}$
  $\hat{v} := T(\hat{u})$
until $v \geqslant -\varepsilon + u$ *or* $v \lll -\varepsilon + u$ *or* $\hat{v} \lll -\varepsilon + \hat{u}$ *or* $\min_{i \in J}(u_i) < -(|J| - 1)W$ *or* $N \geqslant N^*$
if $v \lll -\varepsilon + u$ *or* $\hat{v} \lll -\varepsilon + \hat{u}$ *or* $\min_{i \in J}(u_i) < -(|J| - 1)W$ **then**
  return "Unfeasible"
else
  return "Feasible"

Theorem 3.2.1 (Value iteration for tropical linear systems with integer coefficients)**.** *Algorithm 1 applied to the Krasnoselskii-Mann damped Shapley operator T_{KM} correctly decides (in exact arithmetic) the feasibility of a tropical linear system with integer coefficients in at most $N^* = 4\min(|I|, |J|)(|J| - 1)W + 1 = \mathcal{O}(|J|^2 r^\infty)$ iterations for $\varepsilon < \frac{1}{\min(|I|,|J|)}$, where W is as given in Lemma 3.1.9.*

Proof. The algorithm terminates by design in at most $N^* = 4\min(|I|, |J|)(|J| - 1)W + 1$ iterations, which proves the termination as well as the worst case complexity.

For the correction, start by noticing that $\underline{T_{\mathsf{KM}}} = \underline{T}_{\mathsf{KM}}$, so the two operators can be used interchangeably. Now, we see that the algorithm terminates after iteration N if one of the four following conditions is satisfied:

(i) $u^{N+1} = \underline{T_{\mathsf{KM}}}(u^N) \geqslant -\varepsilon + u^N$.

This entails by Corollary 1.5.29 that $\underline{\chi}(\underline{T_{\mathsf{KM}}}) \geqslant -\varepsilon$, and thus from Remark 3.1.4 that $\underline{\chi}(T) \geqslant -2\varepsilon$ which proves the correctness for ε small enough, depending on the weights of the game.

However, in the case where the coefficients of A and B are all integers, then by Remark 1.5.16, all coordinates of $\chi(T)$ are rational numbers with a denominateur less than or equal to $2\min(|I|, |J|)$, and thus for $\varepsilon < \frac{1}{\min(|I|,|J|)}$, the inequality $\chi(T) \geqslant -2\varepsilon$ implies that $\chi(T) \geqslant 0$, proving by Corollary 1.5.31 that the system $A \odot y \leqslant B \odot y$ is feasible.

(ii) $u^{N+1} = \underline{T_{\mathsf{KM}}}(u^N) \lll -\varepsilon + u^N$.

This entails by Corollary 1.5.29 that $\underline{\chi}(\underline{T_{\mathsf{KM}}}) \leqslant -\varepsilon < 0$, which means by Remark 3.1.4 that $\underline{\chi}(T) < 0$, and thus $\chi(T) \ngeqslant 0$ and thus Corollary 1.5.31 implies this time that the system $A \odot y \leqslant B \odot y$ does not have a solution in $\mathbb{R}^J$.

(iii) $T_{\mathsf{KM}}(\hat{u}^N) \lll -\varepsilon + \hat{u}^N$ where $\hat{u}^N$ is the vector constructed in the widening step of Algorithm 1.

Similarly, this entails by Corollary 1.5.29 that $\underline{\chi}(T_{\mathsf{KM}}) \leqslant -\varepsilon < 0$ and we can conclude as above, using Remark 3.1.4 and Corollary 1.5.31.

(iv) $\min_{i\in J}(u_i^N) < -(|J|-1)W$

Assume that $\underline{\chi}(T) \geqslant 0$. Then from Corollary 3.1.11, one has

$$u^N \geqslant -(|J|-1)W$$

for all $N \geqslant 0$. Therefore, if $\min_{i\in J}(u_i^N) < -(|J|-1)W$, then this entails that $\underline{\chi}(T) < 0$.

(v) $N > N^* = 4\min(|I|,|J|)(|J|-1)W$.

In the case where $\underline{\chi}(T) < 0$, then with the notation of Lemma 3.1.8, one has $\overline{\chi}(T_{\mathscr{D}}) < 0$. Moreover, applying Remark 1.5.16 to the operator $T_{\mathscr{D}}$ entails that $\overline{\chi}(T_{\mathscr{D}}) < \frac{1}{\min(|I|,|J|)}$, and thus one obtains that for all $N \geqslant 0$,

$$\pi_{\mathscr{D}}(u^N) \leqslant (T_{\mathscr{D}})^N_{\mathsf{KM}}(0) \leqslant -\frac{N}{2\min(|I|,|J|)} + (|J|-1)W \ ,$$

where the first inequality comes from the fact that T is order-preserving, and the second one follows from Corollary 3.1.11 applied to the reduced operator. Thus, since

$$-\frac{N}{2\min(|I|,|J|)} + (|J|-1)W < -(|J|-1)W$$

as soon as $N > 4\min(|I|,|J|)(|J|-1)W$, if the condition (iv) has not been reached at before $4\min(|I|,|J|)(|J|-1)W$, this entails that $\underline{\chi}(T) \geqslant 0$. □

Remark 3.2.2. Theorem 3.2.1 shows in fact that the value iteration algorithm is pseudo-polynomial, as the cost of each iteration is in $\mathcal{O}(|I||J|)$, which is just the cost of an evaluation of the Shapley operator.

The previous theorem also generalizes to system of tropical linear systems with potentially non-integer coefficients.

Theorem 3.2.3 (Value iteration for general tropical linear systems)**.** *Algorithm 1 applied to the Krasnoselskii-Mann damped Shapley operator T_{KM} correctly decides the feasibility of a tropical linear system (with potentially non-integer coefficients) in at most $N^* = \frac{|J|(|J|-1)W}{\varepsilon} + 1$ iterations for small enough values of ε.*

Proof. The proof of this theorem is mainly identical to the previous one, with the only difference residing in the choice of the timeout $N^* = \frac{|J|(|J|-1)W}{\varepsilon} + 1$. If the number of iterations performed by the algorithm exceeds $\frac{|J|(|J|-1)W}{\varepsilon}$, then it implies that $\min_{i\in J}(u_i^N) < -(|J|-1)W$, and thus the termination condition (iv) of the previous proof is met and the system does not have a solution. Indeed, since the operator $\underline{T}_{\mathsf{KM}}$ is order-preserving, it follows that the sequence $(u^N)_{N\in\mathbb{N}}$ is decreasing. In the worst case scenario, if after N iterations, condition (i) was not met, then it means that at each iteration, there exists $j \in J$ such that $u_j^{N+1} < -\varepsilon + u_j^N$,or in other words, at least one coordinate has decreased by at least ε. However, this can only happen for at most $\frac{|J|(|J|-1)W}{\varepsilon}$ iterations, because after the $\frac{|J|(|J|-1)W}{\varepsilon} + 1$-th iteration, this entails that $u_j^{N+1} < -(|J|-1)W$ for some $j \in J$, and thus condition (iv) is met, thus concluding the proof of the correction of the algorithm. □

One can also perform Algorithm 1 in approximate arithmetics, thanks to the next theorem.

Theorem 3.2.4 (Approximate value iteration for tropical linear systems with integer coefficients)**.** *For $\eta > 0$, let $\tilde{T}$ be an η-approximation of a Shapley operator T, meaning that*

$$\forall u \in \mathbb{R}^J, \quad \|\tilde{T}(u) - T(u)\|_\infty \leqslant \eta \ .$$

Then, keeping the notation of Theorem 3.2.1, Algorithm 1 applied to the operator $\tilde{T}_{\mathsf{KM}}$ correctly decides the feasibility of a tropical linear system with integer coefficients in at most $N^* = 4\min(|I|,|J|)(|J|-1)W + 1$ iterations for sufficiently small values of ε and sufficiently small approximation errors η. In particular, for $\varepsilon = \frac{1}{2\min(|I|,|J|)}$ it is sufficient to take $\eta < \frac{1}{(8|J|^2W+3)|J|}$.

Proof. Fix $\eta > 0$. First, notice that if $\tilde{T}$ is an η-approximation of T, then $\underline{\tilde{T}_{\mathsf{KM}}}$ is an $\frac{\eta}{2}$-approximation of T_{KM}. Let $(u^N)_{N\in\mathbb{N}} \in (\mathbb{R}^J)^{\mathbb{N}}$ be the sequence as defined above Theorem 3.2.1, and similarly, define $(\tilde{u}^N)_{N\in\mathbb{N}} \in (\mathbb{R}^J)^{\mathbb{N}}$ by

$$\begin{cases} \tilde{u}^0 = (0,\dots,0) \in \mathbb{R}^J \\ \tilde{u}^{N+1} = \underline{\tilde{T}_{\mathsf{KM}}}(\tilde{u}^N) \quad \text{for all } N \geqslant 0 \ . \end{cases}$$

The triangular inequality and non-expansivity of the operator $\underline{T_{\mathsf{KM}}}$ give the inequality

$$\|\tilde{u}^{N+1} - u^{N+1}\|_\infty \leqslant \|\underline{\tilde{T}_{\mathsf{KM}}}(\tilde{u}^N) - \underline{T_{\mathsf{KM}}}(\tilde{u}^N)\|_\infty + \|\underline{T_{\mathsf{KM}}}(\tilde{u}^N) - \underline{T_{\mathsf{KM}}}(u^N)\|_\infty \leqslant \frac{\eta}{2} + \|\tilde{u}^N - u^N\|_\infty$$

for all $N \in \mathbb{N}$, and it follows by induction that

$$\forall N \in \mathbb{N}, \quad \|\tilde{u}^N - u^N\|_\infty \leqslant \frac{N\eta}{2} \ ,$$

and thus, again by triangular inequality,

$$\forall N \in \mathbb{N}, \quad \tilde{u}^{N+1} - \tilde{u}^N - \frac{(2N+1)\eta}{2} \leqslant u^{N+1} - u^N \leqslant \tilde{u}^{N+1} - \tilde{u}^N + \frac{(2N+1)\eta}{2} \ ,$$

and similarly for the widening step,

$$\forall N \in \mathbb{N}, \quad \tilde{T}_{\mathsf{KM}}(\hat{\tilde{u}}^N) - \hat{\tilde{u}}^N - \frac{(2N+1)\eta}{2} \leqslant T_{\mathsf{KM}}(\hat{u}^{N+1}) - \hat{u}^N \leqslant \tilde{T}_{\mathsf{KM}}(\hat{\tilde{u}}^N) - \hat{\tilde{u}}^N + \frac{(2N+1)\eta}{2} \ .$$

These inequalities show the correction of the algorithm for sufficiently small values of ε and sufficiently small approximation errors η in general.

In particluar, if we take $\varepsilon = \frac{1}{2\min(|I|,|J|)}$, then for the algorithm to be correct, one needs to have $\frac{(2N+1)\eta}{2} < \frac{1}{2\min(|I|,|J|)}$, *i.e.* $\eta < \frac{1}{(2N+1)\min(|I|,|J|)}$ for the conclusion of (i) to still hold in the previous proof, and since in that case the convergence happen in a number of iterations $N \leqslant 4\min(|I|,|J|)(|J|-1)W + 1$, one thus has

$$\frac{1}{(2N+1)\min(|I|,|J|)} \geqslant \frac{1}{(8|J|^2W+3)|J|} \ ,$$

hence $\eta \leqslant \frac{1}{(8|J|^2W+3)|J|}$ suffices. □

Theorem 3.2.5 (Approximate value iteration for general tropical linear systems)**.** *For $\eta > 0$, let $\tilde{T}$ be an η-approximation of a Shapley operator T. Then, keeping the notation of Theorem 3.2.1, Algorithm 1 applied to the operator $\tilde{T}_{\mathsf{KM}}$ correctly decides the feasibility of a tropical linear system with integer coefficients in at most $N^* = \frac{|J|(|J|-1)W}{\varepsilon} + 1$ iterations for sufficiently small values of ε and sufficiently small approximation errors η.*

Proof. Up to the change in the value of N^*, the proof of this theorem is identical to the first part of the proof of Theorem 3.2.4. □

One can thus check if the system of tropical polynomial equations $f^+(x) \geqslant f^-(x)$ has a finite solution, by applying the above algorithm to the Shapley operators $T = (\mathcal{M}^-_{\mathcal{E}'})^\sharp \mathcal{M}^+_{\mathcal{E}'}$ and T_{KM}, with $\mathcal{E}'$ any superset of nonempty a Canny-Emiris subset associated to the system, as described in Theorem 2.2.1: indeed, by construction, no row of the Macaulay matrices $\mathcal{M}^\pm_{\mathcal{E}'}$ can be identically equal to $-\infty$, which means that the operator T satisfies Assumption 1.5.10 (a), and thus satisfies the requirement that it must preserve $\mathbb{R} \cup \{+\infty\}$. In particular, in the case where the coefficients of the polynomials are integers, then if the system is overconstrained, then the matrices $\mathcal{M}^\pm_{\mathcal{E}}$ will have more rows than columns, and their number of columns will be exactly $|\mathcal{E}|$, hence one can choose $\varepsilon < \frac{1}{|\mathcal{E}|}$ to ensure the correction of the algorithm.

Chapter 4

Nonlinear eigenvalue methods for tropical polynomial systems

In this chapter, we develop a nonlinear eigenvalue method, based on the solution of parametric mean payoff games, in order to compute the solution set of a tropical polynomial system, whenever it is finite, or otherwise in order to describe the projection of the solution set on each coordinate otherwise.

4.1 Solving parametric mean payoff games: a path-following method

In this section, we focus on the resolution of a class of parametric mean payoff games. The ability to solve these parametric mean payoff games will be crucial in the resolution of tropical polynomial systems. It is also a tool of interest in the context of mean payoff games, as it allows one to perform general game homotopy, which the following class of games is a particular case of. The payment matrices of the parametric games under consideration are of the form $A_\zeta = (a_{ij}(\zeta))_{(i,j)\in I\times J}$ and $B_\zeta = (b_{ij}(\zeta))_{(i,j)\in I\times J}$, where the entries are either continuous piecewise affine functions of the real parameter $\zeta \in \mathbb{R}$ or identically equal to $-\infty$. Moreover, denote by r_ζ^∞ the maximal absolute value of all the finite coefficients of A_ζ and B_ζ. We take interest in the parametric Shapley operator $T_\zeta := A_\zeta^\sharp B_\zeta$, as well as the parametric system of homogeneous tropical linear inequalities of the form $A_\zeta \odot y \leqslant B_\zeta \odot y$ of unknown $y \in \mathbb{R}^J$.

4.1.1 The spectral function

We introduce the key notion of spectral function, on which the whole path-following method relies.

Definition 4.1.1. The *spectral function* of the operator $T_\zeta = A_\zeta^\sharp B_\zeta$ is the map ϕ defined by

$$\phi : \left\{ \begin{array}{ccc} \mathbb{R} & \longrightarrow & \mathbb{R} \\ \zeta & \longmapsto & \underline{\chi}(T_\zeta) \end{array} \right. . \tag{4.1}$$

By definition of the spectral function, the super-level set $\{\zeta \in \mathbb{R} : \phi(\zeta) \geqslant 0\}$ corresponds to the set of parameters $\zeta \in \mathbb{R}$ for which the maximizer player is winning. Therefore, being able to compute it efficiently allows one to solve the parametric mean payoff game of Shapley operator $T_\zeta = A_\zeta^\sharp B_\zeta$, and thus determine the values of the parameter $\zeta \in \mathbb{R}$ for which the parametric tropical linear system $A_\zeta \odot y \leqslant B_\zeta \odot y$ of unknown $y \in \mathbb{R}^J$ has a solution.

The following theorem gives some of the main properties of the spectral function of piecewise affine parametric mean payoff games, generalizing the results of [GKS12, §3.2] to a more general class of parametrizations.

Theorem 4.1.2. *The spectral function ϕ is continuous and piecewise affine over $\mathbb{R}$. Moreover, it is Lipschitz of constant equal to the biggest Lipschitz constant of the a_{ij} and b_{ij} for all $(i, j) \in I \times J$.*

Proof. The proof of this result uses arguments very similar to the one used in the proof of [GKS12, Theorem 8]. We provide the proof for the sake of completeness. as per Theorem 1.5.15, for all $j \in J$, the function $\zeta \mapsto \chi_j(T_\zeta)$ is obtained as a min-max function of the mean weight of the cycles of the graph $\mathrm{G}(T_\zeta)$ that can be reached from the initial state j. Since the weights $-a_{ij}(\zeta), b_{ij}(\zeta)$ are piecewise affine and Lipschitz in ζ for all $(i,j) \in I \times J$, then so is $\chi_j(T_\zeta)$ for all $j \in J$, and therefore so is $\phi(\zeta) = \underline{\chi}(T_\zeta)$. □

The spectral function can be interpreted in the following way: whenever $\phi(\zeta) \leqslant 0$, then it indicates how close the system at parameter ζ is to be feasible, and whenever $\phi(\zeta) \geqslant 0$, it morally measures the size of the solution set. This measure can be related to the radius of the largest Hilbert ball contained in the solution set, possibly up to a condition number. This is to be compared with the duality results from [AGQS23].

We also state the following remark on genericity, on which further results shall rely.

Remark 4.1.3. Assume that A and B are chosen generic in the sense of Theorem 1.5.26. Then the average weights of all cycles of the graph $\mathrm{G}(T_\zeta)$ are generic linear functions of $\zeta \in \mathbb{R}$, and thus there exists only a single cycle of maximal average weight except for finitely many values of ζ for which there are exactly two cycles of maximal average weight. Anything more than two simultaneous maximal average weight cycles would entail a nongeneric condition on the coefficients of A and B.

We next show how to compute the spectral function over a given segment $\mathcal{I}$ of $\mathbb{R}$, using a path-following method. First, we construct the matrix $\bar{A}_\zeta$, obtained by replacing every $-\infty$ entry of A by a number $-M$, with M large enough, in order for the vector of values of the parametric game to be constant at every point $\zeta \in \mathbb{R}$, as per the following lemma. The intuition behind this 'big M' trick is that at the cost of a large penalty M, it allows the minimizer player to teleport from any state $j \in J$, to any desired state $i \in I$. This ensures in particular that asymptitocally, the behaviour of the game, if both players go on playing optimally, does not depend on whatever the starting position $j_0 \in J$ was, hence rendering the vector of values of the mean payoff game constant. This is detailled in the following lemma.

Lemma 4.1.4. *If M is chosen larger than $4|J|\max_{\zeta\in\mathcal{I}}(r^\infty_\zeta)$, then for all $\zeta \in \mathcal{I}$, $\chi(\bar{A}^\sharp_\zeta B_\zeta)$ is a constant vector whose entries coincide with $\underline{\chi}(A^\sharp_\zeta B_\zeta)$. Moreover, the eigenproblem $\bar{A}^\sharp_\zeta B_\zeta u = \lambda + u$, with $\lambda \in \mathbb{R}$ and $u \in \mathbb{R}^J$, is solvable, and λ is unique and satisfies $\lambda = \phi(\zeta) = \underline{\chi}(A^\sharp_\zeta B_\zeta)$.*

Proof. Let $\zeta \in \mathcal{I}$. Replacing all the $-\infty$ entries of the matrices A_ζ by a finite entry $-M$ with M large means that from any state $j \in J$, at the cost of a payment of M to the maximizer player, the minimizer can move to every state $i \in I$. Therefore, the set of cycles eventually reached in the game does not depend on the initial position $j_0 \in J$ of the game, hence $\chi(\bar{A}^\sharp_\zeta B_\zeta)$ is a constant vector. The same reasoning holds for any additive perturbation of the transition payments $-a_{ij}, b_{ij}$, and thus [AGH15, Theorem 3.1] entails the existence of a solution $(\lambda, u) \in \mathbb{R} \times \mathbb{R}^J$ to the eignevalue problem $\bar{A}^\sharp_\zeta B_\zeta u = \lambda + u$, and moreover, λ is unique as per the Collatz-Wielandt property (Theorem 1.5.27 and Corollary 1.5.29) since $\underline{\chi}(\bar{A}^\sharp_\zeta B_\zeta) = \overline{\chi}(\bar{A}^\sharp_\zeta B_\zeta)$.

Moreover, if M is chosen large enough, then in the modified game, the mean weight of any cycle that contains an arc with weight M — that is of any newly formed cycle — will be greater than the weight of any cycle present in the initial game, encouraging the minimizer player to only play finitely many times along an arc of weight M in order to minimize the mean weight of the eventually reached cycle. Since the length of a cycle can not exceed $2\min(|I|,|J|) \leqslant 2|J|$ and since the weights are bounded above by r^∞_ζ for all $\zeta \in \mathcal{I}$ it does indeed follow that for all $\zeta \in \mathcal{I}$, one has $\chi(\bar{A}^\sharp B) \equiv \underline{\chi}(A^\sharp B)$ with $M \geqslant 2|J|\max_{\zeta\in\mathcal{I}}(r^\infty_\zeta)$. □

For the remainder of this section, we assume that A_ζ has already only finite real entries, and shall apply the following algorithm to $\bar{A}_\zeta$ instead of A_ζ. This hypothesis, as per the previous lemma, is crucial, in order for the parametric ergodic equation defined below to always have a solution, allowing us to apply properly our path-following method. Note however that we rely on this 'big M' trick because it makes the following results easier to state. In fact, this trick and the increase in complexity it entails can be avoided altogether by using more complex lexicographic methods.

4.1.2 Outline of the path-following method

In order to state Algorithm 2a below, we first need to introduce the ergodic equation and its associated derivated equation, which will play a central role in this method.

The *parametric ergodic equation* refers to the non-linear eigenvalue problem $T_\zeta(u) = \lambda + u$ of unknown $(\lambda, u) \in \mathbb{R} \times \mathbb{R}^J$. From the equality $T_\zeta = A^\sharp_\zeta B_\zeta$, it can be rewritten as the following problem

$$\min_{i\in I} \quad -a_{ij}(\zeta) + v_i \quad = \quad \lambda + u_j \quad \forall j \in J \tag{4.2a}$$

$$\max_{j\in J} \quad b_{ij}(\zeta) + u_j \quad = \quad v_i \quad \forall i \in I \ , \tag{4.2b}$$

refered to as the *parametric ergodic problem*, where the unknowns are $\lambda \in \mathbb{R}$, $u \in \mathbb{R}^J$ and $v \in \mathbb{R}^I$. Given a solution (λ, u, v) of (4.2), we denote the sets of active constraints by

$$I_{j,\zeta}(u) = \operatorname*{arg\,min}_{i \in I} -a_{ij}(\zeta) + v_i \quad \text{and} \quad J_{i,\zeta}(u) = \operatorname*{arg\,max}_{j \in J} b_{ij}(\zeta) + u_j \ .$$

These sets only depend on the choice of ζ and u since v is determined by $v = B_\zeta u$. We also recall that by definition, the sets $I_{j,\zeta}(u)$ and $J_{i,\zeta}(u)$ describe all the arcs in the saturation graph $\mathrm{SAT}(T_\zeta, u)$ of the operator T_ζ associated to the bias vector u (see Definition 1.5.20).

Let a'_{ij} and b'_{ij} denote the right derivatives of the piecewise affine functions a_{ij} and b_{ij}, with the convention that the derivative of the constant $-\infty$ function is the zero function. Notice that the piecewise affine assumption for the a_{ij} and b_{ij} entails that the functions a'_{ij} and b'_{ij} are constant on a right-neighbourhood of any point $\zeta_0 \in \mathbb{R}$, and for any ζ_0, we denote by $\mathcal{I}^{\mathsf{lin}}(\zeta_0) \subseteq \mathbb{R}$ the greatest right interval of the form $[\zeta_0, \zeta[$ over which the coefficients of the matrices A_ζ and B_ζ are affine. This means that these matrices can be written in the form $A_\zeta = A_{\zeta_0} + (\zeta - \zeta_0)A'_{\zeta_0}$ and $B_\zeta = B_{\zeta_0} + (\zeta - \zeta_0)B'_{\zeta_0}$ for all $\zeta \in \mathcal{I}^{\mathsf{lin}}(\zeta_0)$, where $A'_\zeta = (a'_{ij}(\zeta))_{(i,j)\in I\times J}$, $B'_\zeta = (b'_{ij}(\zeta))_{(i,j)\in I\times J} \in \mathbb{R}^{I\times J}$ for all $\zeta \in \mathbb{R}$.

If there exists a choice of solution $(\lambda(\zeta), u(\zeta)) \in \mathbb{R} \times \mathbb{R}^J$ such that $\lambda(\zeta)$ and $u(\zeta)$ can be written as right differentiable functions of ζ, then, for all $j \in J$, the function $\zeta \mapsto T_\zeta(u(\zeta))_j$ is also right differentiable and moreover its right derivative satisfies

$$\frac{\mathrm{d}}{\mathrm{d}\zeta}\Big(T_\zeta(u(\zeta))_j\Big) = \min_{i \in I_{j,\zeta}(u(\zeta))} \Big(-a'_{ij}(\zeta) + \max_{k \in J_{i,\zeta}(u(\zeta))} b'_{ik}(\zeta) + u'_k(\zeta)\Big) \quad ,$$

where $u'_k(\zeta)$ denotes the right-derivative of the function $\zeta \mapsto u_k(\zeta)$ at point ζ. This motivates the following definition. For $\zeta \in \mathbb{R}$ and $u \in \mathbb{R}^J$, denote by $A'_{\zeta,u}$ the matrix with entries $a'_{ij}(\zeta)$ for $i \in I_{j,\zeta}(u)$ and $-\infty$ otherwise. Similarly, denote by $B'_{\zeta,u}$ the matrix with entries $b'_{ij}(\zeta)$ for $j \in J_{i,\zeta}(u)$ and $-\infty$ otherwise. Then, the *derivated Shapley operator* at point ζ, u refers to the operator $T'_{\zeta,u}$ defined by $T'_{\zeta,u} := (A'_{\zeta,u})^\sharp B'_{\zeta,u}$, and the *derivated ergodic equation* refers to the non-linear eigenvalue problem $T'_{\zeta,u}(u') = \lambda' + u'$ of unknown $(\lambda', u') \in \mathbb{R} \times \mathbb{R}^J$. This equation can similarly be rewritten as the following problem

$$\begin{aligned} \min_{i \in I_{j,\zeta}(u)} \quad -a'_{ij}(\zeta) + v'_i \quad &= \quad \lambda' + u'_j \quad \forall j \in J && \text{(4.3a)} \\ \max_{j \in J_{i,\zeta}(u)} \quad b'_{ij} + u'_j \quad &= \quad v'_i \qquad \forall i \in I \ , && \text{(4.3b)} \end{aligned}$$

refered to as the *derivated ergodic problem*, where the unknowns are $\lambda' \in \mathbb{R}$, $u' \in \mathbb{R}^J$ and $v' \in \mathbb{R}^I$. Similarly to above, given a solution (λ', u', v') of (4.3), we denote the sets of active constraints by

$$I'_{j,\zeta}(u') = \operatorname*{arg\,min}_{i \in I_{j,\zeta}(u)} -a'_{ij}(\zeta) + v'_i \quad \text{and} \quad J'_{i,\zeta}(u') = \operatorname*{arg\,max}_{j \in J_{i,\zeta}(u)} b'_{ij}(\zeta) + u'_j \ ,$$

which depend only on the choice of ζ and u'.

Finally, for a fixed ζ, the map T_ζ, being piecewise affine, admits a directional derivative at every point, and the derivative $\partial_h T_\zeta(u)$ of T_ζ at point $u \in \mathbb{R}^J$ in the direction $h \in \mathbb{R}^J$ satisfies

$$\partial_h T_\zeta(u) = \Big(\min_{i \in I_{j,\zeta}(u)} \max_{k \in J_{i,\zeta}(u)} h_k\Big)_{j \in J} \ .$$

Using this directional derivative, one can always write the first-order expansion of T_ζ at point u and in the direction h, which is locally exact since again T_{ζ_0} is piecewise affine, *i.e.*

$$T_\zeta(u + th) = T_\zeta(u) + t\partial_h T_\zeta(u) \quad \text{for } t > 0 \text{ small enough.} \tag{4.4}$$

The idea of Algorithm 2a is to construct a continuous piecewise affine function $\lambda : \mathcal{I} \to \mathbb{R}$, as well as two *possibly noncontinuous* piecewise affine functions $u : \mathcal{I} \to \mathbb{R}^J$ and $v : \mathcal{I} \to \mathbb{R}^I$ satisfying the parametric ergodic problem (4.2). In particular, ϕ will coincide with λ, hence the continuity of λ. However, u and v may not be continuous, but we shall construct them as right continuous functions. We denote by $u(\zeta_0^-)$ and $v(\zeta_0^-)$ the left limit of u and v in ζ_0 respectively.

Now suppose that λ, u and v have been evaluated at a point ζ_0. Then, we look for a solution of the parametric ergodic problem (4.2) defined on a small right neighborhood of ζ_0, and satisfying the following *Ansatz*

$$\begin{cases} \lambda(\zeta) = \lambda(\zeta_0) + (\zeta - \zeta_0)\lambda'(\zeta_0) \\ u(\zeta) = u(\zeta_0) + (\zeta - \zeta_0)u'(\zeta_0) \\ v(\zeta) = v(\zeta_0) + (\zeta - \zeta_0)v'(\zeta_0) \end{cases} \tag{4.5}$$

where $\lambda'(\zeta_0) \in \mathbb{R}$, $u'(\zeta_0) \in \mathbb{R}^J$ and $v'(\zeta_0) \in \mathbb{R}^I$, representing respectively the right derivative of λ, u and v at ζ_0, will be computed as solutions of the derivated ergodic problem (4.3). The latter problem reduces to solving a mean payoff game in which the payments matrices are $A'_{\zeta_0,u(\zeta_0)}$ and $B'_{\zeta_0,u(\zeta_0)}$. The points $\zeta_0 \in \mathbb{R}$ such that there exists a solution of the parametric ergodic problem (4.2) satisfying the *Ansatz* (4.5) on a small right neighborhood of ζ_0 are refered to as *regular points*, and any other point is refered to as a *singular point*. The next lemma gives a condition on the solvability of this problem.

Lemma 4.1.5 (Affine right-continuation of the solution)**.** *The derivated ergodic problem* (4.3) *has a solution* $\lambda'(\zeta_0), u'(\zeta_0), v'(\zeta_0)$ *if and only if there exists a solution* $\lambda(\zeta), u(\zeta), v(\zeta)$ *of the parametric ergodic problem* (4.2) *which is affine in* ζ *on a right neighbourhood of* ζ_0*, and satisfies the Ansatz* (4.5)*. In that case, for any* ζ *in such a right neighbourhood of* ζ_0*, with* $\zeta > \zeta_0$*, the sets of active constraints satisfy* $I_{j,\zeta}(u(\zeta)) = I'_{j,\zeta_0}(u'(\zeta_0))$ *and* $J_{i,\zeta}(u(\zeta)) = J'_{i,\zeta_0}(u'(\zeta_0))$*.*

Proof. Assume that $\lambda'(\zeta_0), u'(\zeta_0), v'(\zeta_0)$ is solution of (4.3) and consider $\lambda(\zeta)$, $u(\zeta), v(\zeta)$ as defined in (4.5). For ζ in a right neighbourhood of ζ_0, one thus has

$$\begin{aligned} -a_{ij}(\zeta) + v_i(\zeta) - \lambda(\zeta) - u_j(\zeta) = &\; -a_{ij}(\zeta_0) + v_i(\zeta_0) - \lambda(\zeta_0) - u_j(\zeta_0) \\ &+ (\zeta - \zeta_0)\big(-a'_{ij}(\zeta_0) + v'_i(\zeta_0) - \lambda'(\zeta_0) - u'_j(\zeta_0)\big) \end{aligned}$$

and

$$\begin{aligned} b_{ij}(\zeta) + u_j(\zeta) - v_i(\zeta) = &\; b_{ij}(\zeta_0) + u_j(\zeta_0) - v_i(\zeta_0) \\ &+ (\zeta - \zeta_0)\big(b'_{ij}(\zeta_0) + u'_j(\zeta_0) - v'_i(\zeta_0)\big) \enspace . \end{aligned}$$

One has from the previous equations that for all $(i,j) \in I \times J$,

$$-a_{ij}(\zeta_0) + v_i(\zeta_0) - \lambda(\zeta_0) - u_j(\zeta_0) \geqslant 0 \quad \text{and} \quad b_{ij}(\zeta_0) + u_j(\zeta_0) - v_i(\zeta_0) \leqslant 0 \enspace ,$$

with equality exactly whenever $i \in I_{j,\zeta_0}(u(\zeta_0))$ for the former inequality, and whenever $j \in J_{i,\zeta_0}(u(\zeta_0))$ for the latter. Since $-a_{ij}(\zeta_0) + v_i(\zeta_0) - \lambda(\zeta_0) - u_j(\zeta_0) = 0$ for $i \in I_{j,\zeta_0}(u(\zeta_0))$, by taking the minimum over $I_{j,\zeta_0}(u(\zeta_0))$, one obtains from (4.3) that

$$\min_{i \in I_{j,\zeta_0}(u(\zeta_0))} -a_{ij}(\zeta) + v_i(\zeta) - \lambda(\zeta) - u_j(\zeta) = (\zeta - \zeta_0) \min_{i \in I_{j,\zeta_0}(u(\zeta_0))} -a'_{ij}(\zeta_0) + v'_i(\zeta_0) - \lambda'(\zeta_0) - u'_j(\zeta_0) = 0 \enspace ,$$

and for $i \in I \setminus I_{j,\zeta_0}(u(\zeta_0))$, one has $-a_{ij}(\zeta_0) + v_i(\zeta_0) - \lambda(\zeta_0) - u_j(\zeta_0) > 0$, and thus it follows that $-a_{ij}(\zeta) + v_i(\zeta) - \lambda(\zeta) - u_j(\zeta) > 0$ on a right neighbourhood of ζ_0. Likewise, for $j \in J_i(\zeta_0)$, one has $b_{ij}(\zeta_0) + u_j(\zeta_0) - v_i(\zeta_0) = 0$, and this time taking the maximum over $J_{i,\zeta_0}(u(\zeta_0))$ yields

$$\max_{j \in J_{i,\zeta_0}(u(\zeta_0))} b_{ij}(\zeta) + u_j(\zeta) - v_i(\zeta) = (\zeta - \zeta_0) \max_{j \in J_{i,\zeta_0}(u(\zeta_0))} b'_{ij}(\zeta_0) + u'_j(\zeta_0) - v'_i(\zeta_0) = 0 \enspace ,$$

and for $j \in J \setminus J_{i,\zeta_0}(u(\zeta_0))$, one has $b_{ij}(\zeta_0) + u_j(\zeta_0) - v_i(\zeta_0) < 0$, thus $b_{ij}(\zeta) + u_j(\zeta) - v_i(\zeta) < 0$ on a right neighbourhood of ζ_0. This entails that $\lambda(\zeta)$, $u(\zeta), v(\zeta)$ is indeed a solution of the parametric ergodic problem (4.2) on a right neighbourhood of ζ_0 with the same set of active constraints as $\lambda'(\zeta_0), u'(\zeta_0), v'(\zeta_0)$.

The converse implication follows readily from the above calculations, which also prove that if there exists an affine right-continuation $\lambda(\zeta)$, $u(\zeta), v(\zeta)$ of a solution $\lambda(\zeta_0)$, $u(\zeta_0), v(\zeta_0)$ of (4.2) such that $\lambda(\zeta)$, $u(\zeta), v(\zeta)$ is also a solution of the parametric ergodic problem (4.2) for ζ in a right-neighbourhood of ζ_0, then its gradient must be a solution of the derivated ergodic problem (4.3). □

Corollary 4.1.6. *There exists a solution* $\lambda(\zeta)$*,* $u(\zeta), v(\zeta)$ *of the parametric ergodic problem* (4.2) *satisfying the Ansatz* (4.5) *on a right neighbourhood of* ζ_0 *if and only if the vector of values* $\chi(T'_{\zeta_0,u(\zeta_0)})$ *of the derivated Shapley operator at the point* $\zeta_0, u(\zeta_0)$ *is a constant vector.*

Proof. As per Corollary 1.5.21, the equation $T(u) = \lambda + u$ has a solution $(\lambda, u) \in \mathbb{R} \times \mathbb{R}^J$ if and only if the vector of values $\chi(T)$ of T is constant. Therefore, the derivated ergodic equation at point $\zeta_0, u(\zeta_0)$ has a solution if and only if the vector of values $\chi(T'_{\zeta_0,u(\zeta_0)})$ of the derivated Shapley operator at the point $\zeta_0, u(\zeta_0)$ is a constant vector, and thus Lemma 4.1.5 entails the desired result. □

We then propose Algorithms 2a and 2b to compute the spectral function over the interval $\mathcal{I}$. We shall describe more in depth how the above algorithms works, as well as prove their correction and termination in the following sections.

4.1.3 The uniqueness and linearity complexes

Whereas the eigenvalue $\lambda(\zeta)$ of the operator $T_\zeta = A^\sharp_\zeta B_\zeta$ is unique for all values of the parameter $\zeta \in \mathbb{R}$, the nonlinear eigenvector of this operator may not be unique in the projective sense, meaning that the parametric ergodic problem (4.2) may have several solutions $u(\zeta), v(\zeta)$ in the projective sense. However, we show that if A_0 or B_0 has *generic entries*, in the sense of not belonging to an explicit finite union of hyperplanes, then the eigenvector map $\zeta \mapsto u(\zeta)$ becomes uniquely defined and then the number of singular points must be finite. The proof of the generic uniqueness of the parametric eigenvector relies on ideas similar to the ones used in the proof of [AGH18, Theorem 3.2], but for a stronger type of genericity.

Lemma 4.1.7 (Uniqueness complex)**.** *Consider the parameterized operator $T_{\delta,\zeta}$ defined for all $(\delta,\zeta) \in \mathbb{R}^{I\times J} \times \mathbb{R}$ by $T_{\delta,\zeta} = A^\sharp_\zeta(B_\zeta + \delta)$ and assume that for every value of the parameter (δ,ζ), the perturbed ergodic equation $T_{\delta,\zeta}(u) = \lambda + u$ has a solution $(\lambda, u) \in \mathbb{R} \times \mathbb{R}^J$. Then, there exists a polyhedral subdivision $\mathscr{C}^{\mathsf{uniq}}$ of $\mathbb{R}^{I\times J} \times \mathbb{R}$ such that for all parameter $(\delta,\zeta) \in \mathbb{R}^{I\times J} \times \mathbb{R}$ in the interior of a maximal-dimensional cell, the solution of the perturbed ergodic equation is unique, where the uniqueness of the bias vector u is taken in the projective sense.*

Proof. Recall that for the nonlinear eigenvalue λ, the uniqueness is guaranteed by the Kohlberg theorem, via Corollary 1.5.21.

Fix a positional strategy $\sigma \in \Pi_{\mathsf{min}}$ of the minimizer and let $\lambda^{\sigma,\cdot}(\delta,\zeta)$ be the maximal weight of all cycles in the graph $\mathrm{G}(T^{\sigma,\cdot}_{\delta,\zeta})$. Since the payments are piecewise affine in $(\delta,\zeta) \in \mathbb{R}^{I\times J} \times \mathbb{R}$, then so is $\lambda^{\sigma,\cdot}$. Set $\mathscr{C}^{\sigma,\cdot}$ the linearity complex of $\lambda^{\sigma,\cdot}$, that is the subdivision of $\mathbb{R}^{I\times J} \times \mathbb{R}$ such that $\lambda^{\sigma,\cdot}$ is affine in the interior of all maximal-dimensional cells of $\mathscr{C}^{\sigma,\cdot}$. In particular, fixing one value of ζ_0 such that (δ_0,ζ_0) is in the interior of a maximal-dimensional cell, we get that δ_0 is in the interior of a maximal-dimensional cell C of the linearity complex of $\delta \mapsto \lambda^{\sigma,\cdot}(\delta,\zeta_0)$. Then, as per Theorem 1.5.26, $\mathrm{G}(T^{\sigma,\cdot}_{\delta,\zeta_0})$ has a unique critical cycle for every value δ in the interior of C, and this unique critical cycle is independent of the choice of such a δ. This shows that $\mathrm{G}(T^{\sigma,\cdot}_{\delta,\zeta})$ has a unique critical cycle for every value (δ,ζ) of the parameter in the interior of every maximal-dimensional cell of $\mathscr{C}^{\sigma,\cdot}$. By the same arguments as for the proof of Theorem 1.5.26, one can show that, for a given maximal-dimensional cell C of $\mathscr{C}^{\sigma,\cdot}$, this unique critical cycle is independent of the choice of (δ,ζ) in the interior of C. Moreover, as per Theorem 1.5.22, the uniqueness of the critical cycle for (δ,ζ) implies the uniqueness of the solution $(\lambda, u) \in \mathbb{R}\times\mathbb{R}^J$, in the projective sense for the bias u, of the one-player ergodic equation $T^{\sigma,\cdot}_{\delta,\zeta}(u) = \lambda + u$.

Now go back to the operator $T_{\delta,\zeta}$ and let $(\lambda, u) \in \mathbb{R}\times\mathbb{R}^J$ be a solution to the ergodic equation $T_{\delta,\zeta}(u) = \lambda + u$. Then by (1.6), there exists a positional strategy $\sigma \in \Pi_{\mathsf{min}}$ of the minimizer such that $T_{\delta,\zeta}(u) = T^{\sigma,\cdot}_{\delta,\zeta}(u)$, and thus λ is the unique eigenvalue of $T^{\sigma,\cdot}_{\delta,\zeta}$ and u is a nonlinear eigenvector associated to the operator $T^{\sigma,\cdot}_{\delta,\zeta}$.

If (δ,ζ) lies in the interior of a maximal-dimensional cell of $\mathscr{C}^{\sigma,\cdot}$, then the set of nonlinear eigenvectors u associated to the operator $T^{\sigma,\cdot}_{\delta,\zeta}$ consists in a line directed by the vector $\mathbf{1} = (1,\dots,1) \in \mathbb{R}^J$. By finiteness of the set Π_{min} of positional strategies for the minimizer, this means that the bias vector u must belong to a finite union of lines parallel to $\mathbb{R}\mathbf{1}$. However, [AGH18, Theorem 3.10] shows that the set of eigenvectors associated to a Shapley operator satisfying Assumption 1.5.10 is arcwise connected, and thus it must coincide with a single such line, hence the uniqueness of the nonlinear eigenvector in the projective sense.

Thus, taking $\mathscr{C}^{\mathsf{uniq}}$ to be the smallest common refinement $\bigwedge_{\sigma\in\Pi_{\mathsf{min}}} \mathscr{C}^{\sigma,\cdot}$ of all the previous complexes $\mathscr{C}^{\sigma,\cdot}$ yields the result. □

The polyhedral complex from the previous theorem is refered to as the *uniqueness complex* for eigenvectors of the parametric Shapley operator $T_{\delta,\zeta}$. However, the proof of the finiteness of the path-following method described earlier requires a finer polyhedral complex, called the *linearity complex*, which is given by the following corollary of Lemma 4.1.7.

Algorithm 2a: Path-following method: ζ-pivoting

input: $T_\zeta = A^\sharp_\zeta B_\zeta$ a parametric Shapley operator from $(\mathbb{R} \cup \{+\infty\})^J$ to $(\mathbb{R} \cup \{+\infty\})^J$ such that $\chi(T_\zeta)$ is constant for all $\zeta \in \mathbb{R}$
$\mathcal{I}$ an interval on which the spectral function will be computed

output: Computes the spectral function $\phi : \zeta \mapsto \underline{\chi}(T_\zeta)$ over the interval $\mathcal{I}$

/* Initialization */

$\zeta_0 := \inf \mathcal{I}$

compute a solution (λ_0, u_0) of the ergodic equation $T_{\zeta_0}(u_0) = \lambda_0 + u_0$

$\lambda(\zeta_0) := \lambda_0$

$u(\zeta_0) := u_0$

repeat

  /* Enforcing the affine right continuability of the solution */

  if *the derivated ergodic equation* $T'_{\zeta_0, u(\zeta_0)}(u'_0) = \lambda'_0 + u'_0$ *has no solution* **then**

    compute a new solution $(\lambda_0, \tilde{u}_0)$ of the ergodic equation satisfying the affine right-continuation lemma via ε-pivoting (**Algorithm 2b**)

    $u(\zeta_0) := \tilde{u}_0$

  /* ζ-pivoting step */

  compute a solution (λ'_0, u'_0) of the derivated ergodic equation

  $\lambda(\zeta) := \lambda_0 + (\zeta - \zeta_0)\lambda'_0$

  $u(\zeta) := u_0 + (\zeta - \zeta_0)u'_0$

  compute $\zeta_1 = \sup\{\zeta \in \mathcal{I}^{\mathsf{lin}}(\zeta_0) : T_\zeta(u(\zeta)) = \lambda(\zeta) + u(\zeta)\}$

  $\zeta_0 := \zeta_1$

until $\zeta_0 \geqslant \sup \mathcal{I}$

return the list of pairs $(\zeta, \lambda(\zeta))$ for all pivoting points ζ encountered during the execution

Algorithm 2b: Path-following method: ε-pivoting at singular points

input: $T_\zeta = A^\sharp_\zeta B_\zeta$ a parametric Shapley operator from $(\mathbb{R} \cup \{+\infty\})^J$ to $(\mathbb{R} \cup \{+\infty\})^J$ such that $\chi(T_\zeta)$ is constant for all $\zeta \in \mathbb{R}$
$\zeta_0 \in \mathbb{R}$ a singular point
$(\lambda_0, u_0) \in \mathbb{R} \times \mathbb{R}^J$ a solution of the ergodic equation $T_{\zeta_0}(u_0) = \lambda_0 + u_0$ which admits an affine left continuation, but not an affine right continuation

output: Returns a solution $(\lambda_0, \tilde{u}_0) \in \mathbb{R} \times \mathbb{R}^J$ of the above ergodic equation which admits an affine right continuation

/* Initialization */

$\tilde{u}_0 := u_0$

compute the unique solution $h_{\mathsf{r}} \in \{0,1\}^J$ of the fixed-point equation $\partial_h T_{\zeta_0}(\tilde{u}_0) = h$ such that $h_{\mathsf{r}} \neq \mathbf{0} \mod \mathbb{R}\mathbf{1}$

repeat

  /* ε-pivoting step */

  compute $t^* = \max\{t \geqslant 0 : T_{\zeta_0}(\tilde{u}_0 + t h_{\mathsf{r}}) = \lambda_0 + (\tilde{u}_0 + t h_{\mathsf{r}})\}$

  $\tilde{u}_0 := \tilde{u}_0 + t^* h_{\mathsf{r}}$

  $h_{\mathsf{l}} := h_{\mathsf{r}}$

  compute the unique solution $h_{\mathsf{r}} \in \{0,1\}^J$ of the equation $\partial_h T_{\zeta_0}(\tilde{u}_0) = h$ such that $h_{\mathsf{r}} \neq -h_{\mathsf{l}} \mod \mathbb{R}\mathbf{1}$

until $h_{\mathsf{r}} = \mathbf{0} \mod \mathbb{R}\mathbf{1}$

return $(\lambda_0, \tilde{u}_0)$

Lemma 4.1.8 (Linearity complex)**.** *There exists a refinement $\mathscr{C}^{\mathsf{lin}}$ of the uniqueness complex $\mathscr{C}^{\mathsf{uniq}}$ on $\mathbb{R}^{I\times J}\times\mathbb{R}$ such that the solution $(\lambda,u)\in\mathbb{R}\times\mathbb{R}^J$ of $\lambda+u=T_{\delta,\zeta}(u)$ exists, is unique (up to an additive constant for the bias vector u) and affine with respect to (δ,ζ) on the interior of all cells with maximal dimension.*

Proof. One needs to prove that the unique bias vector u from Lemma 4.1.7 is piecewise affine in (δ,ζ). Take (δ,ζ) in the interior of a maximal-dimensional cell of $\mathscr{C}^{\mathsf{uniq}}$. Then, by the construction in the proof of Lemma 4.1.7, there exists a positional strategy $\sigma\in\Pi_{\mathsf{min}}$ depending only on the cell of the minimizer player such that the eigenpair (λ,u) of $T_{\delta,\zeta}$ is an eigenpair of the tropically linear operator $T^{\sigma,\cdot}_{\delta,\zeta}$, *i.e.* for all $j\in J$,

$$\max_{k\in J}\Big(a_{\sigma(j)j}(\zeta)+b_{\sigma(j)k}(\zeta)+\delta_{\sigma(j)k}+u_k\Big)=\lambda+u_j\ . \tag{4.6}$$

Moreover, the tropically linear operator $T^{\sigma,\cdot}_{\delta,\zeta}$ has a unique critical cycle, so that equation (4.6) has a unique solution (λ,u) up to the translation of u by a constant, and this critical cycle depends only on the cell. Denoting by C, for brevity, the tropical matrix representing the operator $T^{\sigma,\cdot}_{\delta,\zeta}$, so that $C_{j,k}=a_{\sigma(j)j}(\zeta)+b_{\sigma(j)k}(\zeta)+\delta_{\sigma(j)k}$, and selecting a fixed index j_0 in the critical cycle of C, we know for instance from [But10, Theorem 1.6.18] that the unique eigenvector of C coincides with the j_0-th column $(((-\lambda)\odot C)^*)_{\cdot,j_0}$ of the Kleene star

$$((-\lambda)\odot C)^*=\max_{0\leqslant\ell\leqslant|J|-1}((-\lambda)\odot C)^{\odot\ell}\ .$$

This is a finite maximum of terms which are piecewise-linear in (ζ,δ). This shows that the unique eigenvector is a piecewise-linear function of (ζ,δ) on the interior of each cell of maximal dimension of the uniqueness complex. Then, we can take for the linearity complex any complex refining the piecewise-linearity regions obtained in this way.

□

Corollary 4.1.9. *Fix a generic $\delta\in\mathbb{R}^{I\times J}$. Then when ζ runs over $\mathbb{R}$, there exists a selection $(\lambda(\zeta),u(\zeta))\in\mathbb{R}\times\mathbb{R}^J$ of nonlinear eigenpairs such that the map $\zeta\mapsto\lambda(\zeta)$ is continuous and piecewise affine, and the map $\zeta\mapsto u(\zeta)$ is piecewise affine, right-continuous with left limit and with finitely many discontinuity points on every bounded interval.*

Proof. This is a very straight-forward application of the previous lemma, as for a fixed generic $\delta\in\mathbb{R}^{I\times J}$, the line $\{(\delta,\zeta):\zeta\in\mathbb{R}\}$ will only cross lower-dimensional cells of the linearity complex $\mathscr{C}^{\mathsf{lin}}$ at isolated points, from which the result follows. □

Remark 4.1.10. Even though there exists a polyhedral complex $\mathscr{C}^{\mathsf{lin}}$ such that the bias vector is linear on the interior of each maximal-dimensional cell, Sturmfels and Tran noted – in the one-player case – that the collection of linearity domains of the eigenvector itself is not necessarly a polyhedral complex, see [ST11, §3]. Moreover, the eigenvector is a piecewise affine but generally *discontinuous* function of the matrix entries (see Example 4.1.15). This pathology originates from the nonuniqueness of the eigenvector on lower dimensional cells, allowing a 'jump' of the eigenvector when crossing the boundary of a maximal dimensional cell. This difficulty is central in the development of a tropical homotopy method, it will be handled below by the treatment of the *singular* 'ε-pivoting' steps (Section 4.1.5).

Remark 4.1.11. A related linearity complex (without the ζ parameter) has been used recently in [LS24] for other purposes.

In the results from the next sections, the expressions 'for generic instances' is to be understood in the following way: a statement on T_ζ holds *for generic instances* if it holds for all $T_{\delta,\zeta}$ except for a zero-measure set of values of the parameter δ. As we shall see in the proofs of the following results, this notion of genericity will be necessary in order to control the saturation graph of T_ζ at different points, as it will play a key role in the different pivoting steps described below.

4.1.4 The ζ-pivoting step at regular points

At regular points, that is whenever the condition of Lemma 4.1.5 is satisfied, we perform what shall be refered to as a *ζ-pivoting step*, similar in its principle to a pivoting in the simplex algorithm, in order to find the supremum ζ_1 of all $\zeta\in\mathcal{I}^{\mathsf{lin}}(\zeta_0)$ such that $\lambda(\zeta),u(\zeta),v(\zeta)$ as given by the Ansatz (4.5) remains a solution of the parametric ergodic problem (4.2).

Lemma 4.1.12 (ζ-pivoting step)**.** *Assume that the condition of Lemma 4.1.5 is satisfied. Then the constant $\zeta_1 \in \mathbb{R}$ defined by $\zeta_1 := \sup\{\zeta \in \mathcal{I}^{\mathsf{lin}}(\zeta_0) : T_\zeta(u(\zeta)) = \lambda(\zeta) + u(\zeta)\}$ satisfies*

$$\zeta_1 = \zeta_0 + \min\left\{\delta_{ij}^{(a)} : j \in J,\, i \in I\right\} \cup \left\{\delta_{ij}^{(b)} : i \in I,\, j \in J\right\} \cup \left\{\ell(\mathcal{I}^{\mathsf{lin}}(\zeta_0))\right\} \ , \tag{4.7}$$

where $\ell(\mathcal{I}^{\mathsf{lin}}(\zeta_0))$ denote the length of the interval $\mathcal{I}^{\mathsf{lin}}(\zeta_0)$ and

$$\delta_{ij}^{(a)} := \frac{-a_{ij}(\zeta_0) + v_i(\zeta_0) - \lambda(\zeta_0) - u_j(\zeta_0)}{a'_{ij}(\zeta_0) - v'_i(\zeta_0) + \lambda'(\zeta_0) + u'_j(\zeta_0)} \qquad \text{if} \quad -a'_{ij}(\zeta_0) + v'_i(\zeta_0) - \lambda'(\zeta_0) - u'_j(\zeta_0) < 0 \tag{4.8a}$$

$$\delta_{ij}^{(b)} := \frac{-b_{ij}(\zeta_0) - u_j(\zeta_0) + v_i(\zeta_0)}{b'_{ij}(\zeta_0) + u'_j(\zeta_0) - v'_i(\zeta_0)} \qquad \text{if} \quad b'_{ij}(\zeta_0) + u'_j(\zeta_0) - v'_i(\zeta_0) > 0 \tag{4.8b}$$

with $\delta_{ij}^{(a)} = +\infty$ and $\delta_{ij}^{(b)} = +\infty$ when the corresponding conditions are not satisfied.

Proof. Whenever the condition of Lemma 4.1.5 is satisfied, then $\zeta \in \mathcal{I}^{\mathsf{lin}}(\zeta_0)$ is such that $\lambda(\zeta), u(\zeta), v(\zeta)$ as given by the Ansatz (4.5) does no longer satisfy the parametric ergodic problem (4.2) if either for some $j \in J$,

$$\min_{i \in I \setminus I_{j,\zeta_0}(u(\zeta_0))} -a_{ij}(\zeta_0) + v_i(\zeta_0) - \lambda(\zeta_0) - u_j(\zeta_0) + (\zeta - \zeta_0)\big(-a'_{ij}(\zeta_0) + v'_i(\zeta_0) - \lambda'(\zeta_0) - u'_j(\zeta_0)\big) < 0 \ ,$$

or either for some $i \in I$,

$$\max_{j \in J \setminus J_{i,\zeta_0}(u(\zeta_0))} b_{ij}(\zeta) + u_j(\zeta) - v_i(\zeta) + (\zeta - \zeta_0)\big(b'_{ij}(\zeta_0) + u'_j(\zeta_0) - v'_i(\zeta_0)\big) > 0 \ .$$

The linear inequality

$$-a_{ij}(\zeta_0) + v_i(\zeta_0) - \lambda(\zeta_0) - u_j(\zeta_0) + (\zeta - \zeta_0)\big(-a'_{ij}(\zeta_0) + v'_i(\zeta_0) - \lambda'(\zeta_0) - u'_j(\zeta_0)\big) < 0$$

has a solution $\zeta > \zeta_0$ whenever $i \in \{i \in I : -a'_{ij}(\zeta_0) + v'_i(\zeta_0) - \lambda'(\zeta_0) - u'_j(\zeta_0) < 0\} \subseteq I \setminus I_{j,\zeta_0}(u(\zeta_0))$. In that case, the infimum of the solution set of the inequation is given by

$$\zeta_0 + \frac{-a_{ij}(\zeta_0) + v_i(\zeta_0) - \lambda(\zeta_0) - u_j(\zeta_0)}{a'_{ij}(\zeta_0) - v'_i(\zeta_0) + \lambda'(\zeta_0) + u'_j(\zeta_0)} \ ,$$

and otherwise, it is equal to $+\infty$, hence (4.8a).

Similarly, the linear inequality

$$b_{ij}(\zeta) + u_j(\zeta) - v_i(\zeta) + (\zeta - \zeta_0)\big(b'_{ij}(\zeta_0) + u'_j(\zeta_0) - v'_i(\zeta_0)\big) > 0$$

has a solution $\zeta > \zeta_0$ whenever $j \in \{j \in J : b'_{ij}(\zeta_0) + u'_j(\zeta_0) - v'_i(\zeta_0) > 0\} \subseteq J \setminus J_{i,\zeta_0}(u(\zeta_0))$. Likewise, in that case, the infimum of the solution set of the inequation is given by

$$\zeta_0 + \frac{-b_{ij}(\zeta_0) - u_j(\zeta_0) + v_i(\zeta_0)}{b'_{ij}(\zeta_0) + u'_j(\zeta_0) - v'_i(\zeta_0)} \ ,$$

and it is equal to $+\infty$ otherwise, hence (4.8b).

All the above finally entails that the supremum ζ_1 of all $\zeta \in \mathcal{I}^{\mathsf{lin}}(\zeta_0)$ for which the Ansatz (4.5) remains a solution of the parametric ergodic problem (4.2) is indeed given by (4.7). □

We then reevaluate the sets of active constraints in ζ_1, that is we compute the sets of active constraints $I_{j,\zeta_1}(u(\zeta_1))$ and $J_{i,\zeta_1}(u(\zeta_1))$ for the new bias $u(\zeta_1)$ and $v(\zeta_1)$. Note that the new sets $I_{j,\zeta_1}(u(\zeta_1))$ and $J_{i,\zeta_1}(u(\zeta_1))$ of active constraints are easily shown to satisfy

$$I_{j,\zeta_1}(u(\zeta_1)) = I'_{j,\zeta_0}(u(\zeta_0)) \cup \left\{i \in I : \delta_{ij}^{(a)} = \zeta_1 - \zeta_0\right\} \tag{4.9a}$$

$$\text{and} \quad J_{i,\zeta_1}(u(\zeta_1)) = J'_{i,\zeta_0}(u(\zeta_0)) \cup \left\{j \in J : \delta_{ij}^{(b)} = \zeta_1 - \zeta_0\right\} \ . \tag{4.9b}$$

We add the following more precise result on the set of active constraints for the values of ζ outside of pivoting points.

Lemma 4.1.13. *For generic instances, the saturation graph* $\mathrm{SAT}(T_\zeta, u(\zeta))$ *is a successor graph — meaning that each vertex has a unique successor,* i.e. *has outdegree exactly one — for all* $\zeta_0 < \zeta < \zeta_1$*, and is moreover independent of* ζ*. In other words, the set of active constraints* $I_{j,\zeta}(u(\zeta))$ *and* $J_{i,\zeta}(u(\zeta))$ *do not depend of* ζ *and have both cardinality* 1.

Proof. Without loss of generality, assume that there exists distinct indices $i_1, i_2 \in I_{j,\zeta}(u(\zeta))$. Then the equality $a_{i_1 j}(\zeta) + v_{i_1}(\zeta) = a_{i_2 j}(\zeta) + v_{i_2}(\zeta)$ entails that ζ is solution of an affine equation. By finiteness of the sets I and J, there exists thus only a finite number of values of ζ for which one of the sets of active constraints $I_{j,\zeta}(u(\zeta))$ or $J_{i,\zeta}(u(\zeta))$ is not reduced to a single point. Therefore the set of active constraints are all of cardinality one on an open interval $\mathcal{I}$ of lower bound ζ_0. Moreover, let ζ^* be the upper bound of $\mathcal{I}$. Then again without loss of generality, at ζ^*, there exists two distincts constraints $i_1, i_2 \in I_{j,\zeta^*}(u(\zeta^*))$ such that $i_2 \notin I_{j,\zeta_0}(u(\zeta_0))$, and the equality $a_{i_2 j}(\zeta^*) + v_{i_2}(\zeta^*) - \lambda(\zeta^*) - u_j(\zeta^*) = 0$ implies that $\zeta^* = \zeta_0 + \delta^{(a)}_{i_2 j} \geqslant \zeta_1$, as per (4.7). It follows that $\mathcal{I} =]\zeta_0, \zeta_1[$. Moreover, the saturation graph is independent of $\zeta \in \mathcal{I}$, because any change in the saturation graph happens at a pivoting point, hence the result. □

Lemma 4.1.14. *Assume that the solution* $u(\zeta_1)$ *satisfies the affine right continuation lemma Lemma 4.1.5 and consider two points* ζ^- *and* ζ^+ *respectively in arbitrarily small left-neighbourhood and right-neighbourhood of* ζ_1*. Then for generic instances,* $\mathrm{SAT}(T_{\zeta_1}, u(\zeta_1))$ *differs from* $\mathrm{SAT}(T_{\zeta^-}, u(\zeta^-))$*, as well as from* $\mathrm{SAT}(T_{\zeta^+}, u(\zeta^+))$ *by a unique additional arc, in the following way: there exists a unique vertex* v *of* $\mathrm{SAT}(T_{\zeta_1}, u(\zeta_1))$ *with exactly two successors* w^- *and* w^+*, such that the arc* (v, w^-) *is an arc of* $\mathrm{SAT}(T_{\zeta^-}, u(\zeta^-))$ *but not of* $\mathrm{SAT}(T_{\zeta^+}, u(\zeta^+))$*, and conversely the arc* (v, w^+) *is an arc of* $\mathrm{SAT}(T_{\zeta^+}, u(\zeta^+))$ *but not of* $\mathrm{SAT}(T_{\zeta^-}, u(\zeta^-))$

Proof. As per Lemma 4.1.5, one has the equalities $I_{j,\zeta^-}(u(\zeta^-)) = I'_{j,\zeta_0}(u(\zeta_0))$ and $J_{i,\zeta^-}(u(\zeta^-)) = J'_{i,\zeta_0}(u(\zeta_0))$ for all $i \in I$ and $j \in J$. Moreover, the minimum in (4.7) is generically achieved for a single i or j, and therefore (4.9) implies that there exists only a single $I_{j,\zeta_1}(u(\zeta_1))$ or $J_{i,\zeta_1}(u(\zeta_1))$ which contains an additional point, hence $\mathrm{SAT}(T_{\zeta_1}, u(\zeta_1))$ differs indeed from $\mathrm{SAT}(T_{\zeta^-}, u(\zeta^-))$ by a unique additional arc. Symetrically, one shows the same result for $\mathrm{SAT}(T_{\zeta_1}, u(\zeta_1))$ and $\mathrm{SAT}(T_{\zeta^+}, u(\zeta^+))$. □

4.1.5 The ε-pivoting at singular points

The condition of Lemma 4.1.5 is however not always met, and some solutions of the ergodic problem do not necessarily admit an affine right-continuation which remains a solution. This may be the case in particular, for the solution $u(\zeta_1)$ and $v(\zeta_1)$ obtained at the end of the previous regular step.

Example 4.1.15. Consider the parametric mean payoff game described for all $\zeta \in \mathbb{R}$ by the payment matrices

$$A_\zeta = \begin{pmatrix} 1 & 0 \\ 0 & 1-\zeta \end{pmatrix} \quad \text{and} \quad B_\zeta = \begin{pmatrix} 1 & -\infty \\ -\infty & 1 \end{pmatrix} \ .$$

Then, the graph of this mean payoff game as well as its Shapley operator $T_\zeta = A^\sharp_\zeta B_\zeta$ are as follow.

$$T_\zeta \begin{pmatrix} u_1 \\ u_2 \end{pmatrix} = \begin{pmatrix} \min(u_1, u_2 + 1) \\ \min(u_1 + 1, u_2 + \zeta) \end{pmatrix}$$

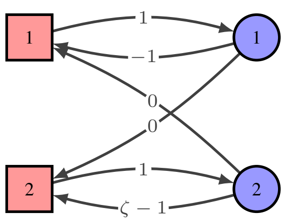

Figure 4.1: The graph $\mathrm{G}(T_\zeta)$.

Then one can easily prove that for all $\zeta \in \mathbb{R}$, the operator T_ζ has a nonlinear eigenvalue $\lambda(\zeta)$ given by

$$\lambda(\zeta) = \min(\zeta, 0) \ ,$$

and moreover, whenever $\zeta \neq 0$, T_ζ has a unique (in the projective sense) nonlinear eigenvector $u(\zeta)$, which is given by

$$u(\zeta) = \begin{cases} \begin{pmatrix} \zeta \\ 0 \end{pmatrix} & \text{if } \zeta < 0 \\ \\ \begin{pmatrix} 0 \\ 1 \end{pmatrix} & \text{if } \zeta > 0. \end{cases}$$

This prevents the selection $\zeta \mapsto u(\zeta)$ to be continuous at 0, and in fact, for $\zeta = 0$, the nonlinear eigenvector is not unique, as for all $0 \leqslant t \leqslant 1$, one has

$$T_0 \begin{pmatrix} 0 \\ t \end{pmatrix} = \begin{pmatrix} 0 \\ t \end{pmatrix} \quad .$$

In order to handle that case, we perform the following *singular pivoting step*. Assume for the rest of this section that ζ_0 is a singular point and fix a solution $\lambda(\zeta_0)$, $u(\zeta_0^-)$, $v(\zeta_0^-)$ of the parametric ergodic problem (4.2), for which the condition of Lemma 4.1.5 is not satisfied. Then, from this solution, we shall construct another solution $\lambda(\zeta_0)$, $u(\zeta_0^+)$, $v(\zeta_0^+)$ of the parametric ergodic problem (4.2) which does indeed admit an affine right continuation. The idea to find such a solution will be, given the fixed value of the parameter $\zeta_0 \in \mathbb{R}$, to perform another type of pivoting in the nonlinear eigenspace $\mathrm{Eig}(T_{\zeta_0})$ of the operator T_{ζ_0}, using the directional derivative of T_{ζ_0}, until such a solution is found. In accordance with the language of the theory of automata[1], this 'stationary' pivoting, for which the value ζ_0 of the parameter is 'frozen', shall be refered to as an *ε-pivoting*. This pivoting relies on the fact that, for generic instances the nonlinear eigenspace $\mathrm{Eig}(T_{\zeta_0})$ is a *simple polygonal chain*, that is a connected union of finitely many segments which intersect only at their consecutive endpoints, as stated by Theorem 4.1.19, whose proof relies on the following lemmas. We first start to prove that the nonlinear eigenspace is a connected metric graph — that is a realization of an abstract graph as a one dimensional polyhedral complex — which is moreover bounded and does not contain any cycles, making it a finite tree.

Lemma 4.1.16. *For generic instances, the nonlinear eigenspace $\mathrm{Eig}(T_{\zeta_0})$ at the singular point ζ_0 thought of as a subset of $\mathbb{R}^J/\mathbb{R}\mathbf{1}$ is a metric connected finite acyclic graph (tree).*

Proof. As per Theorem 3.10 and Remark 3.11 of [AGH18], the nonlinear eigenspace $\mathrm{Eig}(T_{\zeta_0})$ can be obtained as a deformation retract of $\mathbb{R}^J/\mathbb{R}\mathbf{1}$ by a piecewise-affine nonexpansive map, and therefore it is a bounded, simply connected, finite union of polyhedra. Moreover, recall from (1.10) that $\mathrm{Eig}(T_{\zeta_0}) \subseteq \bigcup_{\sigma\in\Pi_{\mathsf{min}}} \mathrm{Eig}(T_{\zeta_0}^{\sigma,\cdot})$. By Remark 4.1.3, there cannot be more than two cycles of maximal average weight in the graph associated to $T_{\zeta_0}^{\sigma,\cdot}$. Therefore, for all strategy $\sigma \in \Pi_{\mathsf{min}}$, the critical graph associated to $\mathrm{Eig}(T_{\zeta_0}^{\sigma,\cdot})$ has at most two connected components, and therefore, [BCOQ92, Theorem 3.101] entails that it is at most one dimensional, seen again as a subset of $\mathbb{R}^J/\mathbb{R}\mathbf{1}$. Hence, $\mathrm{Eig}(T_{\zeta_0})$ can be written as a simply connected union of a finite number of segments, which entails the result. ☐

Lemma 4.1.17. *Let $v \in \mathrm{Eig}(T_{\zeta_0}) \subset \mathbb{R}^J/\mathbb{R}\mathbf{1}$ be a bias vector and denote for all $\rho > 0$ the euclidian ball of center v and radius ρ by $B(v,\rho)$. Then*

$$\mathrm{Eig}(T_{\zeta_0}) \cap B(v,\rho) = \{v + h : h = \partial_h T_{\zeta_0}(v),\ \|h\| \leqslant \rho\} \quad \textit{for } \rho > 0 \textit{ small enough.}$$

Proof. Let $\rho > 0$ be arbitrarily small. Then as per (4.4), one has for all $h \in \mathbb{R}^J$ such that $\|h\| \leqslant \rho$

$$T_{\zeta_0}(v+h) = T_{\zeta_0}(v) + \partial_h T_{\zeta_0}(v) = \lambda + v + \partial_h T_{\zeta_0}(v) \quad .$$

Thus, $v + h \in \mathrm{Eig}(T_{\zeta_0})$ if and only if $h = \partial_h T_{\zeta_0}(v)$, hence the result. ☐

Lemma 4.1.18. *Let v be a node of $\mathrm{Eig}(T_{\zeta_0})$ (seen as a tree). Then for generic instances, the solution set in $\mathbb{R}^J/\mathbb{R}\mathbf{1}$ of the equation $h = \partial_h T_{\zeta_0}(v)$ consists in a single half-line if v is a leaf, or in two half-lines meeting at their common origin otherwise.*

In particular, for a fixed state $j \in J$, there exists exactly two solutions non-constant $h \in \{0,1\}^J$ of the previous equation if v is an internal node, and otherwise, if v is a leaf, then only one such solution exists.

[1] In automata theory, the term 'ε-transition' refers to transitions that take a null time, without consuming any input symbol. By analogy, we call ε-pivoting the present method. We may think of ζ as a physical time, which must be frozen at every singular point. Then several ε-pivotings are needed for the time ζ to be able to progress again.

Proof. To begin, recall that the equation $h = \partial_h T_{\zeta_0}(v)$ can be rewritten as

$$h_j = \min_{i \in I_{j,\zeta_0}(v)} \max_{k \in J_{i,\zeta_0}(v)} h_k \quad \text{for all} \quad j \in J \ . \tag{4.10}$$

Lemmas 4.1.13 and 4.1.14 entail that $\mathrm{SAT}(T_{\zeta_0}, v)$ consists in a successor graph to which a single additional arc has been added. Therefore, all vertices of the saturation graph but one have only one successor, which translates in the fact that $I_{j,\zeta_0}(v)$ and $J_{i,\zeta_0}(v)$ are singletons for these vertices $i \in I$ and $j \in J$, and there is a single vertex $k^* \in I \sqcup J$ of the saturation graph with exactly two successors.

For all $j_0 \in J$ such that j_0 has a unique successor i_0, which itself has a unique successor j_1 in $\mathrm{SAT}(T_{\zeta_0}, v)$, the equality (4.10) entails that $h_{j_0} = h_{j_1}$. Moreover, there are two possibilities for the vertex with two successor:

(i) if $k^* = j_0 \in J$, let $i_0, i'_0 \in I$ be its two successors, then i_0 and i'_0 both have a unique successor, which we denote respectively as j_1 and j'_1, and (4.10) entails that $h_{j_0} = \min(h_{j_1}, h_{j'_1})$;

(ii) if however $k^* = i_0 \in I$, denote by $j_1, j'_1 \in J$ its two successors, then for any predecessor j_0 of i_0, (4.10) entails that $h_{j_0} = \max(h_{j_1}, h_{j'_1})$.

As per Remark 4.1.3, we know that the graph $\mathrm{SAT}(T_{\zeta_0}, v)$ has only two critical cycles. We can thus partition the set J into three classes of vertices: J_1 consisting of vertices which have only access to the first cycle, J_2 consisting in vertices which have only access to the second cycle, and $J_{1,2}$ consinting in vertices which have access to both cycles.

From all the above, we deduce that every solution h of the equation (4.10) is of the form

$$h = \begin{pmatrix} \alpha \\ \vdots \\ \alpha \\ \diamond(\alpha,\beta) \\ \vdots \\ \diamond(\alpha,\beta) \\ \beta \\ \vdots \\ \beta \end{pmatrix} \begin{array}{l} \left.\vphantom{\begin{matrix}\alpha\\ \vdots\\ \alpha\end{matrix}}\right\} J_1 \\ \left.\vphantom{\begin{matrix}\alpha\\ \vdots\\ \alpha\end{matrix}}\right\} J_{1,2} \\ \left.\vphantom{\begin{matrix}\alpha\\ \vdots\\ \alpha\end{matrix}}\right\} J_2 \end{array} \quad \text{with} \quad \alpha, \beta \in \mathbb{R} \text{ and } \diamond = \begin{cases} \min \text{ in case (i)} \\ \max \text{ in case (ii)} \end{cases} .$$

In particular, since we are looking at solutions in $\mathbb{R}/\mathbb{R}\mathbf{1}$, this means that β can be set equal to 0.

Now if all three classes J_1, J_2 and $J_{1,2}$ are nonempty, then the solution set of equation (4.10) in $\mathbb{R}/\mathbb{R}\mathbf{1}$ consists in two halflines, one for the positive values of α and one for the negative values of α, meeting at their common origin when $\alpha = 0$.

If the vertex k^* only has access to one of the cycles of the saturation graph, then the class $J_{1,2}$ is empty, but the same conclusion as above holds, except that the two halflines are in opposite direction, *i.e.* the solution set consists in a line, meaning that the direction h of the pivoting does not change at point v.

Finally, if the vertex k^* belongs in one of the two cycles of the saturation graph and has access to the other cycle, then this means that every node accessing the former cycle has also access to the latter cycle, meaning that one of the classes J_1 or J_2 is empty. Then, the solution set consists this time in a single halfline of $\mathbb{R}/\mathbb{R}\mathbf{1}$. Indeed, assume without loss of generality that $J_2 = \emptyset$ and that we are in case (i) — all other cases work similarly — then for $\alpha < 0$, one has $h = \alpha\mathbf{1}$ which is just the zero vector modulo $\mathbb{R}\mathbf{1}$, and for $\alpha \geqslant 0$, h describes a halfline of $\mathbb{R}/\mathbb{R}\mathbf{1}$ with origin 0. In particular, this means that v is a leaf of $\mathrm{Eig}(T_{\zeta_0})$.

This concludes the proof of the lemma, as the second part of the statement is a direct consequence of the first one by setting $\{\alpha, \beta\} = \{0, 1\}$ in the proper order. □

Theorem 4.1.19. *For generic instances, the nonlinear eigenspace* $\mathrm{Eig}(T_{\zeta_0})$ *at the singular point* ζ_0 *seen as a subset of* $\mathbb{R}^J/\mathbb{R}\mathbf{1}$ *is a simple polygonal chain.*

Proof. This follows directly from Lemmas 4.1.16 and 4.1.18. Indeed, we know from the first lemma that $\mathrm{Eig}(T_{\zeta_0})$ is a connected acyclic graph, and moreover, the second lemma implies that every node besides the leaves has degree 2, and therefore $\mathrm{Eig}(T_{\zeta_0})$ must be a simple polygonal chain. □

Remark 4.1.20. The nonlinear eigenspace is in fact a tropical geodesic, as it can be obtained as the image of a geodesic passing through the two endpoints by a deformation retract that is nonexpansive in Hilbert seminorm (again see [AGH18, Theorem 3.10]).

The ε-pivoting step works in the following way. Given a starting solution $\lambda, u^{(0)}$ of the ergodic equation (1.8) at point ζ_0 such that $u^{(0)}$ is an endpoint of $\mathrm{Eig}(T_{\zeta_0})$, we compute the unique solution $h^{(0)} \in \{0,1\}^J$ of the equation $\partial_h T_{\zeta_0}(u^{(0)}) = h$ of unknown h such that $h^{(0)} \neq 0 \mod \mathbb{R}\mathbf{1}$. Then, the first-order expansion (4.4) of T_{ζ_0} at point $u^{(0)}$ in the direction $h^{(0)}$ entails the following equality for $t > 0$ small enough

$$T_{\zeta_0}(u^{(0)} + th^{(0)}) = T_{\zeta_0}(u^{(0)}) + t\partial_{h^{(0)}} T_{\zeta_0}(u^{(0)}) = \lambda + u^{(0)} + th^{(0)} \enspace .$$

We then set $u^{(1)} = u^{(0)} + t^* h^{(0)}$, where t^* denotes the maximal value of $t > 0$ such that $\lambda, u^{(0)} + th^{(0)}$ is still solution of the ergodic equation (1.8) at point ζ_0. We repeat the previous step replacing $u^{(0)}$ with $u^{(1)}$ above, and by computing the only solution $h^{(1)} \in \{0,1\}^J$ of the equation $\partial_h T_{\zeta_0}(u^{(0)}) = h$ such that $h^{(1)} \neq -h^{(0)} \mod \mathbb{R}\mathbf{1}$, in order to find the next vertex $u^{(2)}$ of $\mathrm{Eig}(T_{\zeta_0})$, and so on until the second endpoint $u^{(N)}$ of $\mathrm{Eig}(T_{\zeta_0})$ is reached, that is whenever the only solution $h^{(N)} \in \{0,1\}^J$ of the equation $\partial_h T_{\zeta_0}(u^{(N)}) = h$ such that $h^{(N)} \neq -h^{(N-1)}$ $\mod \mathbb{R}\mathbf{1}$ is given by $h^{(N)} = 0 \mod \mathbb{R}\mathbf{1}$.

Lemma 4.1.21 (ε-pivoting step)**.** *With the notation of the previous paragraph, the constant $t^* > 0$ defined by $t^* := \sup\{t > 0 : T_{\zeta_0}(u^{(0)} + th^{(0)}) = \lambda + u^{(0)} + th^{(0)}\}$ satisfies*

$$t^* = \min \left\{ t^{(a)}_{ij} : j \in J,\, i \in I \right\} \cup \left\{ t^{(b)}_{ij} : i \in I,\, j \in J \right\} \tag{4.11}$$

with

$$t^{(a)}_{ij} := \frac{a_{ij}(\zeta_0) - v^{(0)}_i + \lambda + u^{(0)}_j}{g^{(0)}_i - h^{(0)}_j} \tag{4.12a}$$

$$t^{(b)}_{ij} := \frac{b_{ij}(\zeta_0) + u^{(0)}_j - v^{(0)}_i}{g^{(0)}_i - h^{(0)}_j} \tag{4.12b}$$

if $g^{(0)}_i - h^{(0)}_j > 0$, where $g^{(0)}_i = \max_{k \in J_{i,\zeta_0}(u^{(0)})} h^{(0)}_k$ for all $i \in I$, and $t^{(a)}_{ij} = t^{(b)}_{ij} = +\infty$ otherwise.

Proof. The proof of the ε-pivoting step is very similar to the proof of the ζ-pivoting step. For all $t > 0$, the ergodic equation $T_{\zeta_0}(u^{(0)} + th^{(0)}) = \lambda + u^{(0)} + th^{(0)}$ can be rewritten as the problem

$$\begin{aligned} \min_{i \in I} \quad -a_{ij}(\zeta_0) + v^{(0)}_i + tg^{(0)}_i &= \lambda + u^{(0)}_j + th^{(0)}_j & \forall j \in J \\ \max_{j \in J} \quad b_{ij}(\zeta_0) + u^{(0)}_j + th^{(0)}_j &= v^{(0)}_i + tg^{(0)}_i & \forall i \in I \enspace . \end{aligned}$$

of unkonwns $\lambda, u^{(0)}, v^{(0)}$, which is satisfied for $t > 0$ small enough as per Lemma 4.1.17.

For all $i \in I_{j,\zeta_0}(u^{(0)})$, $-a_{ij}(\zeta_0) + v^{(0)}_i - \lambda - u^{(0)}_j = 0$, hence

$$\min_{i \in I_{j,\zeta_0}(u^{(0)})} -a_{ij}(\zeta_0) + v^{(0)}_i - \lambda - u^{(0)}_j + t(g_i - h_j) = t \min_{i \in I_{j,\zeta_0}(u^{(0)})} g_i - h_j = 0$$

and similarly, for all $j \in J_{i,\zeta_0}(u^{(0)})$, $b_{ij}(\zeta_0) + u^{(0)}_j - v^{(0)}_i = 0$, hence

$$\max_{j \in J_{i,\zeta_0}(u^{(0)})} b_{ij}(\zeta_0) + u^{(0)}_j - v^{(0)}_i + t(h_j - g_i) = t \max_{j \in J_{i,\zeta_0}(u^{(0)})} h_j - g_i = 0 \enspace .$$

Therefore, we look for the infimum of all values of $t > 0$ such that either

$$\min_{i \in I \setminus I_{j,\zeta_0}(u^{(0)})} -a_{ij}(\zeta_0) + v^{(0)}_i - \lambda - u^{(0)}_j + t(g_i - h_j) < 0$$

or

$$\max_{j \in J \setminus J_{i,\zeta_0}(u^{(0)})} b_{ij}(\zeta_0) + u^{(0)}_j - v^{(0)}_i + t(h_j - g_i) > 0 \enspace .$$

It is given in the first case by $t = \frac{a_{ij}(\zeta_0) - v^{(0)}_i + \lambda + u^{(0)}_j}{g^{(0)}_i - h^{(0)}_j}$, hence (4.12a), and by $\frac{b_{ij}(\zeta_0) + u^{(0)}_j - v^{(0)}_i}{g^{(0)}_i - h^{(0)}_j}$ in the second case, hence (4.12b). $\square$

4.1.6 Termination of the path-following method

Theorem 4.1.22 (Termination of the ε-pivoting)**.** *For generic instances, the solution $u(\zeta_0^-)$ corresponds to an endpoint of the polygonal chain $\mathrm{Eig}(T_{\zeta_0})$, and moreover the second endpoint $u(\zeta_0^+)$ of $\mathrm{Eig}(T_{\zeta_0})$ satisfies the affine right continuation lemma, hence the ε-pivoting starting at the point $u^{(0)} = u(\zeta_0^-)$ terminates and returns a solution $u(\zeta_0^+)$ from which the ζ-pivoting can be resumed.*

Proof. Lemma 4.1.18 and Theorem 4.1.19 entail that the ε-pivoting runs over the nodes of the polygonal chain $\mathrm{Eig}(T_{\zeta_0})$ without ever going back, and therefore by finiteness of the number of nodes in the chain, it must end at one of the endpoints. By symmetry, if the first endpoint admits an affine left continuation then the second one must admit an affine right continuation, as the previous path-following could be replicated in the opposite direction. □

Theorem 4.1.23 (Termination of the path-following method)**.** *The path-following method as described in Algorithm 2a finishes in a finite number of steps for generic instances.*

Proof. Since the ε-pivoting terminates as per the previous lemma, one simply needs to prove that only a finite number of ζ-pivoting points are encountered during the path-following. However, by construction, ζ-pivoting points correspond precisely to the points where the saturation graph $\mathrm{SAT}(T_\zeta, u(\zeta))$ changes. Fix a generic $\delta \in \mathbb{R}^{I\times J}$. Then by construction, the saturation graph $\mathrm{SAT}(T_{\delta,\zeta}, u(\zeta))$ changes whenever the point $(\delta,\zeta) \in \mathbb{R}^{I\times J} \times \mathbb{R}$ crosses a lower-dimensional cell of the linearity complex $\mathscr{C}^{\mathrm{lin}}$. However, a generic segment can only meet lower-dimensional cells of $\mathscr{C}^{\mathrm{lin}}$ finitely many times, and thus there are only finitely many ζ-pivoting points, thus proving the termination of the path-following method for generic instances. □

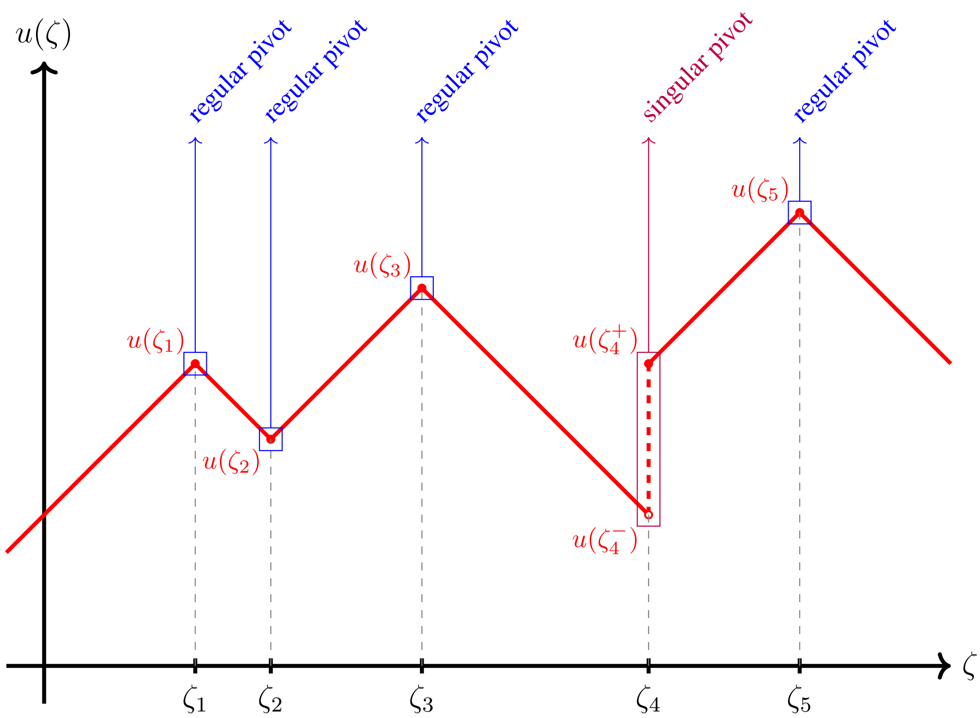

Figure 4.2: The ζ-pivoting from Algorithm 2a consists in computing piece by piece the red curve illustrated above, in order to compute the spectral function. At every non-pivot as well as at every regular pivot, the solution of the eigenproblem is unique in $\mathbb{R}^J/\mathbb{R}\mathbf{1}$. However, at singular pivots, the existence of more than one eigenvector causes a dincontinuity of the selection of the eigenvector $u(\zeta)$.

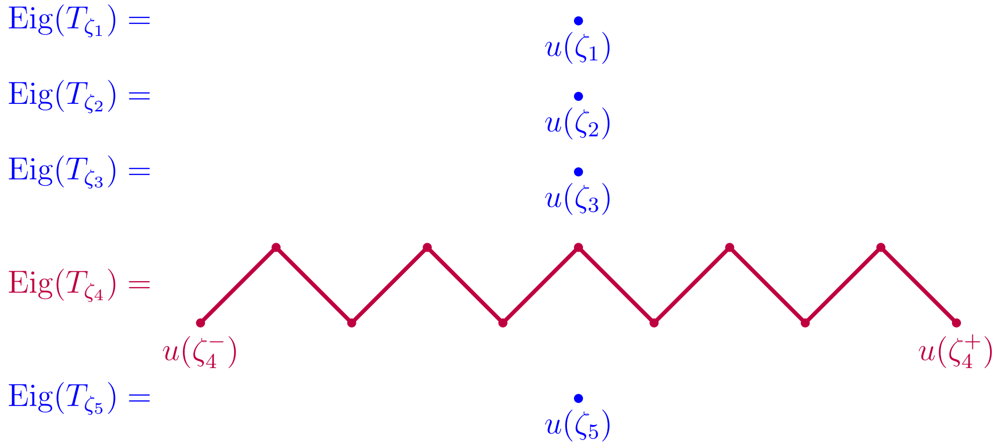

Figure 4.3: At regular pivots (as well as at non-pivots), the associated eigenspace seen as a subset of $\mathbb{R}^J/\mathbb{R}\mathbf{1}$ is reduced to a single point. However, at singular pivots, the eigenspace is a simple polygonal chain, whose endpoints correspond to eigenvectors for which the selection $u(\zeta)$ can be prolongated on the left or on the right. In the singular case, the ε-pivoting from Algorithm 2b explores the eigenspace, starting at one end-point, and finds the second end-point, from which the ζ-pivoting can the be resumed.

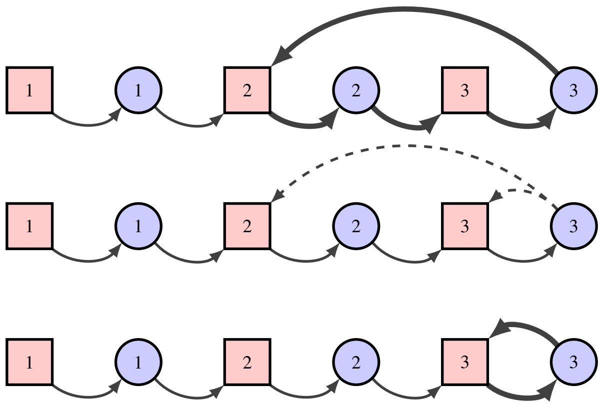

Figure 4.4: A flip of the saturation graph at a regular pivot. Before and after the pivoting, both saturation graphs have a unique critical cycle, marked with bolder edges.

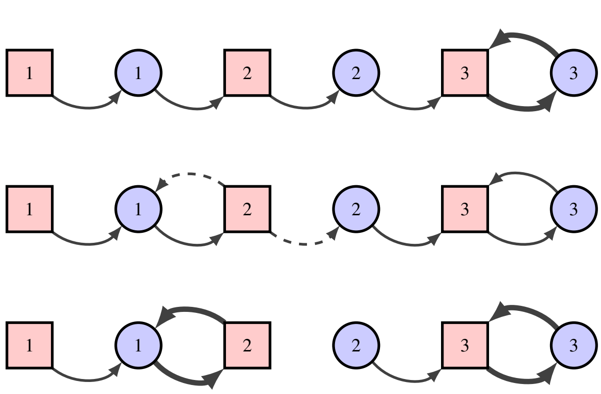

Figure 4.5: A flip of the saturation graph at a singular pivot. Before the pivoting, the saturation graph has a unique critical cycle, but after the pivoting, a second critical cycle appears.

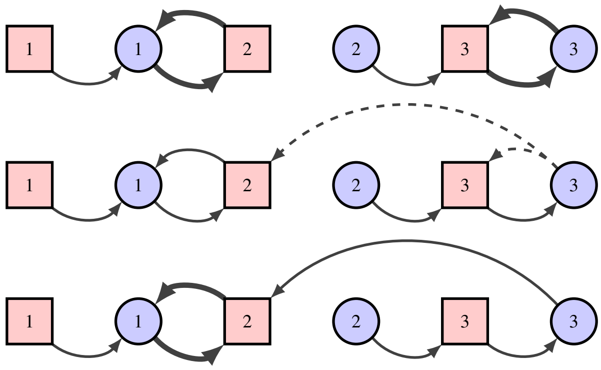

Figure 4.6: A flip of the saturation graph at the last step of the ε-pivoting. Before the pivoting, the saturation graph has two critical cycles, but after the pivoting, one cycle is broken, leaving only a single critical cycle remaining.

4.2 Solving tropical polynomial systems by means of parametric mean payoff games

In this section, we explain how parametric mean payoff games can be used in order to effectively compute solutions of tropical polynomial systems. We propose two approaches. The first approach, based on a dichotomic search, serves mainly to give an easy way to certify the solvability of a tropical polynomial system by exhibiting a particular solution whenever the solution set is nonempty. The second approach consists in the application of the results of the previous section on homotopy-path following for parametric mean payoff games, in order to compute projections on the solution set onto each coordinate.

As in Chapter 3, we shall solely focus on systems of weak tropical polynomial inequalities, as the generalization to other type of tropical polynomial equations and inequations is straight-forward.

4.2.1 Short solutions of tropical polynomial systems

In this first brief section, we state a second short solution property, this time for general tropical polynomial systems. This result will allow us to choose a suitable initialization point for the nonlinear eigenvalue algorithms described in the following sections.

Theorem 4.2.1. *Let $f^{\pm} = (f_1^{\pm}, \dots, f_k^{\pm})$ be two collections of tropical polynomials and let*

$$d = \max_{1 \leqslant i \leqslant k} \deg(f_i^{\pm}) \quad \textit{and} \quad W' = \max_{\substack{1 \leqslant i \leqslant k \\ (\alpha,\beta) \in \mathcal{A}_i^+ \times \mathcal{A}_i^-}} \left| f_{i,\beta}^- - f_{i,\alpha}^+ \right|,$$

and for $\epsilon \in \{\pm 1\}^n$, denote by $\epsilon\mathbb{R}^n_{\geqslant 0}$ the orthant $\{x \in \mathbb{R}^n : \forall j \in [n],\ \epsilon_j x_j \geqslant 0\}$. Then:

(i) the vertices of every polyhedral complex $\{x \in \mathbb{R}^n : \forall i \in [k],\ f_i^-(x) \leqslant f_i^+(x)\} \cap \epsilon\mathbb{R}^n_{\geqslant 0}$ are included in a $\| \cdot \|_\infty$-ball of radius $n(2d)^{n-1}W'$ centered at point 0;

(ii) if moreover all the coefficients of the polynomials $f_i^{\pm}$ are integer, these vertices have coordinates that are rational numbers with a denominator bounded above by $(2d)^n$.

Proof. Starting with the proof of (i), let $V = \{x \in \mathbb{R}^n : \forall i \in [k],\ f_i^-(x) \leqslant f_i^+(x)\}$. Then if V is nonempty, then there exists an orthant $\epsilon\mathbb{R}^n_{\geqslant 0}$, with $\epsilon \in \{\pm 1\}^n$, such that $V_\epsilon := V \cap \epsilon\mathbb{R}^n_{\geqslant 0}$ is nonempty. Thus V_ϵ is a polyhedral complex which is included in an orthant (which is a pointed polyhedron), and therefore V_ϵ admits a 0-dimensional cell, which we denote by x^*. By construction of V_ϵ, the point x^* lies in an intersection of exactly n hyperplanes of the form

$$\{x \in \mathbb{R}^n : f_{i,\alpha}^+ + \langle x, \alpha \rangle = f_{i,\beta}^- + \langle x, \beta \rangle\} \quad \text{and} \quad \{x \in \mathbb{R}^n : x_j = 0\} \ ,$$

where e_j denotes the j-th vector of the standard basis of $\mathbb{R}^n$. Therefore, x^* is solution to a system of equations of the form

$$\left\{ \begin{array}{rcl} \langle \alpha_\ell - \beta_\ell, x \rangle & = & f_{i_\ell,\beta_\ell}^- - f_{i_\ell,\alpha_\ell}^+ \\ \langle e_{j_\ell}, x \rangle & = & 0 \end{array} \right. \quad \text{with} \quad \begin{array}{l} i_\ell \in \{1, \dots, k\}, (\alpha_\ell, \beta_\ell) \in \mathcal{A}_i^+ \times \mathcal{A}_i^- \text{ for } 1 \leqslant \ell \leqslant r \\ j_\ell \in \{1, \dots, n\} \text{ for } r+1 \leqslant \ell \leqslant n \end{array}$$

where $\mathcal{A}_i^+$ and $\mathcal{A}_i^-$ denote respectively the support of f_i^+ and f_i^-. We can thus write this system in the following matrix form

$$Mx = b \quad \text{with} \quad M = \begin{pmatrix} \alpha_1 - \beta_1 \\ \hline \vdots \\ \hline \alpha_r - \beta_r \\ \hline e_{j_{r+1}} \\ \hline \vdots \\ \hline e_{j_n} \end{pmatrix} \quad \text{and} \quad b = \begin{pmatrix} f^+_{i_1,\beta_1} - f^-_{i_1,\alpha_1} \\ \vdots \\ f^+_{i_r,\beta_r} - f^-_{i_r,\alpha_r} \\ 0 \\ \vdots \\ 0 \end{pmatrix} .$$

Writing the matrix M in columns as $M = \big(M_1|\cdots|M_j|\cdots|M_n\big)$, one obtains from the Cramer formula that the solution x^* to the equation $Mx = b$ above is given by

$$x_j^* = \frac{D_j}{D} \quad \text{with} \quad D = \det(M) \text{ and } D_j = \det\big(M_1|\cdots|\underset{\substack{\uparrow \\ j\text{-th column}}}{b}|\cdots|M_n\big) .$$

Notice that M has integer coefficients, which means that $D \in \mathbb{Z}$, and since $D \neq 0$, then $|D| \geqslant 1$, therefore $\left|x_j^*\right| \leqslant |D_j|$. Moreover, expanding the determinant D_j along the j-th column then applying the triangular inequality yields

$$|D_j| \leqslant \sum_{\ell=1}^{r} \left| f^+_{i_\ell,\beta_\ell} - f^-_{i_\ell,\alpha_\ell} \right| |\Delta_{\ell,j}| \quad ,$$

where $\Delta_{\ell,j}$ denotes the minor (ℓ, j) of the matrix M. Applying Hadamard's inequality on the rows of the determinant $\Delta_{\ell,j}$, we obtain

$$|\Delta_{\ell,j}| \leqslant \prod_{1 \leqslant i \neq \ell \leqslant r} \|\alpha_i - \beta_i\|_2 .$$

Moreover, one can bound all the factors of the previous product, as

$$\|\alpha_i - \beta_i\|_2 \leqslant \|\alpha_i - \beta_i\|_1 \leqslant \|\alpha_i\|_1 + \|\beta_i\|_1 \leqslant d + d \leqslant 2d \quad ,$$

and thus $|\Delta_{\ell,j}| \leqslant (2d)^{r-1} \leqslant (2d)^{n-1}$. Finally, each term of the sum $\sum_{\ell=1}^{r} \left| f^+_{i_\ell,\beta_\ell} - f^-_{i_\ell,\alpha_\ell} \right| |\Delta_{\ell,j}|$ can be bounded above by $(2d)^{n-1}W'$, and this sum has at most n terms, hence $\left|x_j^*\right| \leqslant |D_j| \leqslant n(2d)^{n-1}W'$. Since this upper bound works for all $1 \leqslant j \leqslant n$, this shows that x^* belongs in the $\|\cdot\|_\infty$-ball of radius $n(2d)^{n-1}W'$.

Now for the proof of (ii), assume moreover that all the coefficients of the $f_i^\pm$ are integer. Then for all $1 \leqslant j \leqslant n$, $D_j \in \mathbb{Z}$, and $D \in \mathbb{Z}^*$, thus $x^* \in \mathbb{Q}$. Moreover, the Hadamard inequality applied on the rows of M entails the inequality

$$D \leqslant \prod_{i=1}^{r} \|\alpha_i - \beta_i\|_2 \leqslant (2d)^r \leqslant (2d)^n \quad ,$$

giving us the required bound on the denominator of x^*. □

Remark 4.2.2. Theorem 4.2.1 (i) can be seen as a generalization of a weaker version of Lemma 3.1.9 to all systems of polynomial inequalities instead of just linear ones. Indeed, applying the above result with $d = 1$ results in a bound worse than the one from Lemma 3.1.9 by a factor $2^{|J|}$.

4.2.2 First eigenvalue method: dichotomy

We present here a dichotomic method allowing one to certify the solvability of a system of tropical polynomial inequalities. The tropical Positivstellensatz that was established in Chapter 2 entails that checking whether a system of weak tropical polynomial inequalities

$$(S) : \begin{cases} f_1^+(x) & \geqslant & f_1^-(x) \\ & \vdots & \\ f_k^+(x) & \geqslant & f_k^-(x) \end{cases}$$

admits a solution $x \in \mathbb{R}^n$ reduces to solving a mean payoff game obtained by linearization. We enrich the previous system, by adding extra inequalities of the form $a \leqslant x_1 \leqslant b$. In this way, one can decide whether there is a solution such that $x_1 \in [a, b]$.

The 'short solution property' of Theorem 4.2.1 above provides an *a priori* bound for a solution, in the sense that it allows one to reduce the search space to a sup-norm box centered at the origin, of radius equal to the bound $n(2d)^{n-1}W'$ of the theorem, and to rational numbers with a denominator bounded above by $(2d)^n$.

By dichotomic search, a rational number x_1^* which belongs to the projection of the solution set on the first variable can be obtained. Then, we substitute x_1 by the fixed value x_1^* in the polynomial system, and perform again a dichotomic search, now on the variable x_2, leading to a rational value x_2^* such that (x_1^*, x_2^*) belongs to the projection of the solution set on the first two variables. We pursue this procedure by fixing gradually the variables $x_1, \dots, x_n$, so as to eventually retrieve a solution $(x_1^*, \dots, x_n^*)$ of the system (S). Observe that the dichotomic search stops at the first step whenever the solution set is empty. We summarize the search procedure of the first coordinate of a solution with Algorithm 3 below. This algorithm can then be applied recursively in order to compute all the coordinates of a solution of the system (S). We thus arrive at the following complexity result.

Algorithm 3: Dichotomy search.

input: $(S) : \forall i \in [k],\ f_i^+(x) \geqslant f_i^+(x)$ a system of tropical polynomial inequalities
$R > 0$ the radius of a $\| \cdot \|_\infty$-ball to initialise the dichotomy search
$\varepsilon > 0$ approximation error on the solution of the system
output: The first coordinate x_1 of a solution $x \in \mathbb{R}^n$ to the system (S) in $\mathbb{R}^n$ if it is solvable
`/* Initialization */`
$x = 0 \in \mathbb{R}^n$
if *the system (S) does not have a solution such that $-R \leqslant x_1 \leqslant R$* **then**
  `/* No solution to the system is found within the given window */`
  return "No solution within the provided bound"
$a := -R$
$b := R$
`/* Search for the first coordinate` x_1 `of a solution of` (S) `by dichotomy */`
repeat
  $c = \frac{a+b}{2}$
  if *the system (S) does not have a solution such that $a \leqslant x_1 \leqslant c$* **then**
    $b = c$
  else
    $a = c$
until $b - a < \varepsilon$
`/* The mid-point of the final window gives an approximation of` x_1 `*/`
return $\frac{a+b}{2}$

Theorem 4.2.3. *Consider a system of weak polynomial inequalities, as in Theorem 4.2.1. Then, the dichotomic search method returns a rational solution of this system (or decides that there is none) in $\mathcal{O}\big(n \log(n(2d)^{2n-1}W')\big)$ calls to a weak mean payoff oracle.*

Proof. From Theorem 4.2.1 (ii), it is enough to look for solutions with a denominator bounded above by $(2d)^n$, thus the algorithm can be stopped once one reaches a window width lower than $\frac{1}{(2d)^n}$. Since, following Theorem 4.2.1 (i) the initial window width is set to $2n(2d)^{n-1}W'$, the computation of the first coordinate of the solution is hence achieved in $\lfloor \log_2(2n(2d)^{2n-1}W') \rfloor + 1 = \mathcal{O}\big(\log(n(2d)^{2n-1}W')\big)$ steps. We repeat this procedure n times in total for each coordinate, hence a complexity bound in $\mathcal{O}\big(n \log(n(2d)^{2n-1}W')\big)$ ☐

Remark 4.2.4. Notice that the choice of the tiebreaker in the dichotomic search affects the solution returned by the algorithm. In the present case, since Algorithm 3 always favours the left direction, this means that the solution obtained by recursively applying this dichotomic search will be the smallest solution — among all rational solutions with a denominator bounded by $(2d)^n$ — for the lexicographic order of the centered $\| \cdot \|_\infty$-ball of radius R.

Algorithm 3 can also be adapted, in the case of a finite solution set, in order to compute the projections of all the onto the first coordinate. Indeed, if an initial search over the interval $[-R, R]$ returns the value x_1^*, then as per the previous remark, x_1^* is the smallest value of the interval $[-R, R]$ corresponding to the projection of a solution onto the first coordinate, and therefore one can perform a second search over the interval $\left[x_1^* + \frac{1}{2(2d)^n}, R\right]$ to find the second smallest value, and so on until the whole projection of the solution set is computed.

This gives us a method in order to compute the entirety of the solution set whenever it is finite. Assume that the system (S) has a finite number p of solutions. Then one can coordinate the projection of the solution set onto each coordinate, giving us each time at most p values. Thus, we obtain at most p^n candidates for solutions of the system, and then one just need to check for each of these candidates whether they actually form a solution or not. The complexity of this algorithm will then be of order $p^n \times \mathcal{O}\big(n \log(n(2d)^{2n-1}W')\big)$ multiplied by the complexity of checking whether an element $x \in \mathbb{R}^n$ is a solution of the the system (S). This algorithm is in particular output sensitive, since its complexity heavily rely on the number of solutions of the system (S), which can always be bounded above by the BKK bound.

4.2.3 Second eigenvalue method: path-following

We now present a second method based on the path-following of the spectral function of a parametric mean payoff game. The idea for this method is similar to the idea of the dichotomy method: again as per the tropical Positivstellensatz from Chapter 2, the solvability of the system

$$(S) : \begin{cases} f_1^+(x) & \geqslant & f_1^-(x) \\ & \vdots & \\ f_k^+(x) & \geqslant & f_k^-(x) \end{cases}$$

over $\mathbb{R}^n$ can be determined by computing the value of a Mean payoff game. Now in order to obtain the solution set of (S), we proceed as follow: setting ζ a real parameter, we partially evaluate the polynomials of (S) at $x_1 = \zeta$, which yields the following system

$$(S_\zeta) : \begin{cases} f_1^+(\zeta, x_2, \dots, x_n) & \geqslant & f_1^-(\zeta, x_2, \dots, x_n) \\ & \vdots & \\ f_k^+(\zeta, x_2, \dots, x_n) & \geqslant & f_k^-(\zeta, x_2, \dots, x_n) \end{cases} .$$

The system (S_ζ) consists of k $(n-1)$-variate polynomials inequalities, of unkown $(x_2, \dots, x_n) \in \mathbb{R}^{n-1}$, and the coefficients of the polynomials $f_1(\zeta, \cdot), \dots, f_k(\zeta, \cdot) \in \mathbb{T}[X_2, \dots, X_n]$ can be expressed as continuous piecewise affine functions of ζ. The system (S_ζ) can then be linearized into a system of homogeneous tropical linear inequalities of the form $A_\zeta \odot y \leqslant B_\zeta \odot y$ of unknown $y \in \mathbb{R}^J$, in which the entries of the tropical matrices $A_\zeta = (a_{ij}(\zeta))_{(i,j)\in I\times J}$ and $B_\zeta = (b_{ij}(\zeta))_{(i,j)\in I\times J}$ are piecewise affine functions of the scalar ζ. We define the parametric Shapley operator $T_\zeta := A_\zeta^\sharp B_\zeta$. By construction, the Macaulay matrices arising from the linearization will always satisfy Assumption 1.5.10 (b). The 'big M' trick ensures that assumption Assumption 1.5.10 (a) is satisfied as well. However, even without this trick, one could force the latter assumption to be satisfied by replacing T_ζ with the operator $\underline{T_\zeta} : u \mapsto u \wedge T_\zeta(u)$, which sends $\mathbb{R}^J$ to $\mathbb{R}^J$. This change does not affect the zero-locus of the spectral function although it loses the information of the value of the spectral function whenever it is strictly positive.

We recall that the *spectral function* of the operator T_ζ is the map

$$\phi : \begin{cases} \mathbb{R} & \longrightarrow & \mathbb{R} \\ \zeta & \longmapsto & \underline{\chi}(T_\zeta) \end{cases} .$$

We then have the following result.

Theorem 4.2.5. *The projection on the first coordinate of the solution set of (S) coincides with the super-level set $\{\zeta \in \mathbb{R} : \phi(\zeta) \geqslant 0\}$ of the spectral function.*

Proof. As per Theorem 2.2.1 or 2.2.13, the solvability of the system (S_ζ) over $\mathbb{R}^{n-1}$ is equivalent to the solvability of its linearization $A_\zeta \odot y \leqslant B_\zeta \odot y$, and as recalled by Corollary 1.5.31, proving the solvability of this system reduces to showing that $\chi(T_\zeta) \geqslant 0$, and thus that $\phi(\zeta) = \underline{\chi}(T_\zeta) \geqslant 0$. □

The previous result means that we can apply the methods from Section 4.1 in order to compute the spectral function $\zeta \mapsto \underline{\chi}(T_\zeta)$, as from this computation one can retrieve the projection of the solution set onto the first coordinate.

Example 4.2.6. Consider the following system.

$$\begin{cases} 0 \oplus 0x^2y & \geqslant & 2x \oplus 2xy \\ 2xy & \geqslant & 1x \oplus -1y \\ 0 & \geqslant & -3x \oplus -1y \end{cases} \tag{4.13}$$

Observing the zero set of the two associated spectral functions displayed on Figure 4.8, one finds that the solutions of (4.13) are included in $([-3,-2] \sqcup [2,3]) \times [-1,1]$, which is indeed confirmed by the representation of the solution set in Figure 4.7. Note that on Figure 4.8, the operator T_ζ was replaced by the operator $\underline{T_\zeta}$, whose value is $\chi(T) \wedge 0$, hence the feasible set corresponds exactly to the zero-locus of the spectral function $\overline{\phi}$ instead of the nonnegative locus.

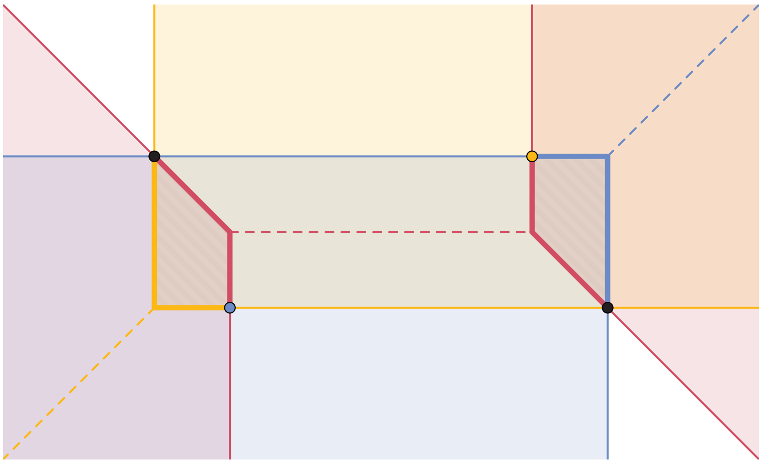

Figure 4.7: The collection of tropical semialgebraic sets arising from system (4.13)

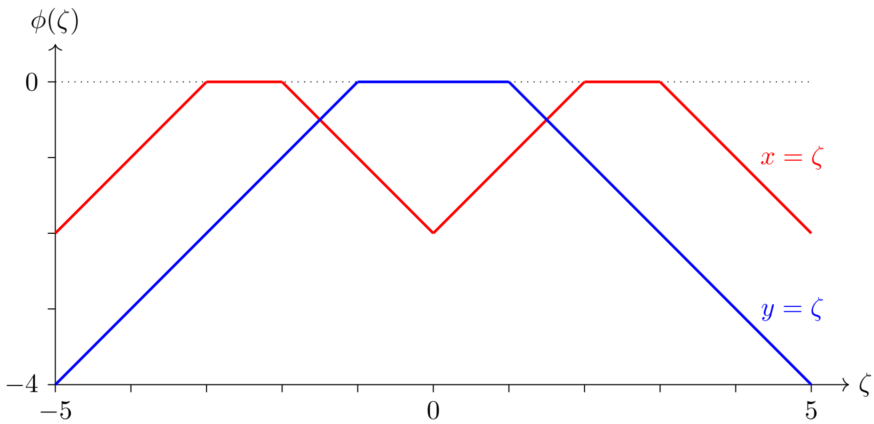

Figure 4.8: The spectral functions obtained when specializing each variable in system (4.13).

We resort to an unpleasant technicality for non-generic systems: we perturb explicitly the input to make it generic, at the cost of a dilation of W by a possibly large factor of at most $(2|\mathcal{E}|+1)^{|\mathcal{E}|^2}$ to ensure that the input remains integer. We denote by $W' \leqslant (2|\mathcal{E}|+1)^{|\mathcal{E}|^2} W$ this dilated cost. Finally, as per Theorem 4.2.1, we can reduce the computation of the spectral function ϕ to the interval $[-R, R]$ with $R = 2n(2d)^{n-1}W'$.

Theorem 4.2.7. *The projection of the solution set of* (S_ζ) *onto the first coordinate can be calculated in finite-time by computing the spectral function* ϕ *of a linearization of* (S_ζ) *with Algorithms 2a and 2b.*

Proof. The tropical Positivstellensatz Theorem 2.2.13 shows that the solvability of system (S_ζ) reduces to the solvability of a linearization obtained by truncation of the Macaulay matrix, and Theorem 4.1.23 shows that the spectral function associated to this parametric linearization can indeed be computed in finite-time hence the result. ☐

4.2.4 Lazy linearization method

Building upon the previous eigenvalue methods, we now describe an incremental method to solve tropical polynomial systems. The main idea is that when solving a tropical polynomial system of the form

$$(S): \left\{ \begin{array}{ccc} f_1^+(x) & \geqslant & f_1^-(x) \\ & \vdots & \\ f_k^+(x) & \geqslant & f_k^-(x) \end{array} \right.$$

with associated Canny-Emiris set $\mathcal{E}$, in a lot of cases, submatrices of $\mathcal{M}_{\mathcal{E}}^{\pm}$ will actually be enough to characterize the solvability of the system over $\mathbb{R}^n$. Indeed, if any linearized system constructed with a subset of $\mathcal{E}' \subseteq \mathcal{E}$ does not have a solution over $\mathbb{R}^{\mathcal{E}'}$, then the tropical polynomial system does not have a solution, or otherwise, the Veronese embedding of this solution would have given a solution of the linearized system. However, if the linearized system does have a solution, then one can try to apply one of the previous eigenvalue method to retrieve a candidate solution to the initial tropical polynomial system. If this candidate solution does indeed satisfy the inequalities of system (S), then this confirms the solvability of the polynomial system. Otherwise, one can expand the set $\mathcal{E}'$ into a bigger subset of $\mathcal{E}$, and repeat the above set until the reaching one of the two above conclusions. This process will necessarily stop, since we know from the results of Chapter 2 that for $\mathcal{E}' = \mathcal{E}$, the solvability of the linearized system over $\mathbb{R}^{\mathcal{E}'}$ is known to be equivalent to the solvability on $\mathbb{R}^n$ of the polynomial system (S).

In particular, in the full case, one can implement the lazy linearization by gradually increasing the truncation degree of the Macaulay matrix in the linearization, and then this process is garanteed to terminate at most at the truncation degree $N = (n+1)(d_1 + \cdots + d_k)$ in the worst case, where for all $i \in [k]$, $d_i = \deg(f_i^+ \oplus f_i^-)$, as per Theorem 2.2.1 or Theorem 2.2.13. One can also generalize this incremental construction to sparse polynomials systems by incrementally constructing a Canny-Emiris set. In a similar manner.

Chapter 5

The Krasnoselskii-Mann iteration for fixed-point free polyhedral nonexpansive mappings

In this brief chapter, we prove a theoretical result on the Krasnoselskii-Mann iteration in the broader setting of polyhedral nonexpansive maps, which in particular includes the case of Shapley operators arising from mean payoff games. This theoretical result lays the groundwork for the algorithmic implementation of the value iteration of Chapter 3, of which a crucial aspect consisted in the use of the Krasnoselskii-Mann damping in the iteration process, in order to ensure the termination.

5.1 Statement of the main convergence theorem

The Krasnoselskii-Mann iteration was originally introduced in [Man53] and [Kra55]. Given a nonexpansive self-map T of a Banach space X, it constructs the sequence

$$v^{N+1} = \frac{1}{2}\left(v^N + T(v^N)\right), \quad v^0 \text{ given.} \tag{5.1}$$

Introducing the operator $T_{\mathsf{KM}}(v) = \frac{1}{2}(v + T(v))$, which is non-expansive, we get $v^N = T_{\mathsf{KM}}^N(v^0)$.

Ishikawa established the following fundamental convergence result, see [Ish76, Corollary 2],

Theorem 5.1.1 (Ishikawa convergence theorem)**.** *Let D be a closed convex subset of a Banach space X and let T be a nonexpansive mapping from D to a compact subset of D. Then T has a fixed point in D, and for any $v^0 \in D$, the sequence $(v^N)_{N\in\mathbb{N}}$ constructed by the Krasnoselskii-Mann iteration converges to a fixed point of T.*

In what follows, we are interested in the degenerate situation in which the operator T has no fixed point. Then, the Krasnoselskii-Mann iteration cannot converge. Moreover, when $X = D = \mathbb{R}^n$, equipped with an arbitrary norm, the sequence $\|v^N\|$ necessarily tends to infinity. Indeed, a general result of Nussbaum (see proof of [Nus88, Theorem 4.1]) entails that if a nonexpansive self-map of $\mathbb{R}^n$ admits a bounded orbit, then, this map has a fixed-point. Our goal is to analyze the asymptotic behaviour of v^N as $N \to \infty$ in such degenerate cases.

To do so, we shall assume that T is *polyhedral*, meaning that $\mathbb{R}^n$ can be covered by finitely many polyhedra on which T has an affine restriction. Then, we recall the result of Kohlberg (Theorem 1.5.13 above) which guarantees the existence of an *invariant halfline* for T, that is a pair (u, η) of vectors of $\mathbb{R}^n$ such that $T(u+s\eta) = u+(s+1)\eta$ for all $s \geqslant 0$ big enough. The vectors u and η are respectively refered to as the *base point* and *direction* of the invariant halfline of T. Moreover, the direction η is unique, we denote it by $\chi(T)$. Using the nonexpansiveness of T, we readily check that

$$\chi(T) = \lim_{N\to\infty} \frac{1}{N} T^N(x),$$

independently of the choice of $x \in \mathbb{R}^n$. Finally, as per Remark 3.1.4, if $(u, \eta)$is an invariant half-line of T, then $(u, \frac{\eta}{2})$ is an invariant halfline of the Krasnoselskii-Mann operator T_{KM}.

The main result of this chapter shows that when T is polyhedral and nonexpansive, the sequence v^N constructed by the Krasnoselskii-Mann iteration heads off to infinity with a drift of $\frac{\chi(T)}{2}$, and moreover, that the deviation between v^N and $N\frac{\chi(T)}{2}$ does converge to a finite quantity.

Theorem 5.1.2 (Fixed-point free Krasnoselskii-Mann iteration)**.** *Let T be a polyhedral selfmap of $\mathbb{R}^n$ that is nonexpansive in an arbitrary norm, and let $(v^N)_{N\in\mathbb{N}}$ be the sequence defined by the Krasnoselskii-Mann iteration* (5.1)*. Then, for any choice of $v^0 \in \mathbb{R}$, the quantity $v^N - N\frac{\chi(T)}{2}$ converges in $\mathbb{R}^n$.*

This result is at the origin of the idea used in Algorithm 1, to work with the Krasnoselskii-Mann operator T_{KM} instead of the original Shapley operator T. Indeed, the convergence of the deviation $v^N - N\frac{\chi(T)}{2}$ is precisely what allows Algorithm 1 to terminate.

Before proving this theorem, we give two examples to show the tightness of its statement. The following first example shows a phenomenon of 'cyclic oscillation' which justifies the use of the Krasnoselskii-Mann operator T_{KM} instead of T in Chapter 3.

Example 5.1.3. Let $T : \mathbb{R}^2 \to \mathbb{R}^2$ be the Shapley operator defined by $T(x, y) = (\max(x, y), x + 1)$ for all $(x, y) \in \mathbb{R}^2$ or, in terms of mean payoff games, by $T = A^\sharp B$ with

$$A = \begin{pmatrix} 0 & -\infty \\ -\infty & 0 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 0 & 0 \\ 1 & -\infty \end{pmatrix} \ .$$

Then one can easily compute for all $N \geqslant 1$ that

$$T^N(0,0) = \left(\left\lfloor \frac{N}{2} \right\rfloor, \left\lfloor \frac{N+1}{2} \right\rfloor \right) \ ,$$

and thus

$$\chi(T) = \lim_{N\to+\infty} \frac{T^N(0,0)}{N} = \left(\frac{1}{2}, \frac{1}{2}\right) \ .$$

However, for all $N \in \mathbb{N}_{>0}$, one has

$$T^N(0,0) - N\chi(T) = \left(\left\lfloor \frac{N}{2} \right\rfloor - \frac{N}{2}, \left\lfloor \frac{N+1}{2} \right\rfloor - \frac{N}{2} \right) = \begin{cases} (0,0) & \text{if } N \text{ is even} \\ (-\frac{1}{2}, \frac{1}{2}) & \text{if } N \text{ is odd} \end{cases} ,$$

and therefore the quantity $T^N(0,0) - N\chi(T)$ keeps alternating without ever converging as N goes to $+\infty$. This first example thus illustrates how applying the value iteration with widening as presented in Algorithm 1 of Chapter 3 directly to the Shapley operator T could lead to a periodic sequence, preventing the termination of the algorithm, and thus this is why the Krasnoselskii-Mann damping is required in order to avoid such cases. Indeed, for the Krasnoselskii-Mann operator, one has $T_{\mathsf{KM}}(x, y) = \frac{1}{2}(\max(2x, x + y), x + y + 1)$ for all $(x, y) \in \mathbb{R}^2$, and then it is straightforward to verify for all $N \geqslant 1$ that

$$T_{\mathsf{KM}}^N(0,0) = \left(\frac{N-1}{4}, \frac{N+1}{4} \right) \ ,$$

and thus

$$T_{\mathsf{KM}}^N(0,0) - N\frac{\chi(T)}{2} = \left(-\frac{1}{4}, \frac{1}{4} \right)$$

which is constant in N and therefore does converge as N goes to $+\infty$.

Moreover, the polyhedral aspect of T in Theorem 5.1.2 is also crucial, as per the following counterexamples.

Example 5.1.4. If T is not polyhedral, then the asymptotic expansion of $T^N(0)$ may contain terms of intermediate order between N and 1. For instance, let $T : \mathbb{R}^2 \to \mathbb{R}^2$ be the sup-norm non-expansive operator defined by $T(x, y) = (f(y), y + 1)$, where

$$f(y) = \begin{cases} 1 & \text{if } y < 1 \\ \sqrt{y} & \text{if } y \geqslant 1 \end{cases} .$$

Then one can easily compute for all $N \geqslant 2$ that

$$T^N(0,0) = \left(\sqrt{N-1}, N \right) \ ,$$

and therefore
$$\frac{T^N(0,0)}{N} \underset{N\to+\infty}{\longrightarrow} (0,1) \ ,$$
and yet for any choice of $\eta = (\eta_1, \eta_2) \in \mathbb{R}^n$, the quantity
$$T^N(0,0) - N\eta = \left(\sqrt{N-1} - N\eta_1, N(1-\eta_2)\right) \ , \tag{5.2}$$
does not converge as N tends to $+\infty$ because of the term in $\sqrt{N-1}$.

In fact, it can get even worse. Kohlberg provided in [Koh80, Remark 2.4] an example of a nonexpansive mapping T, polyhedral on every compact set, with slopes oscillating between two values $0 < \alpha < \beta < 1$, and for which the quantity $\frac{T^N(0)}{N}$ does not even converge, as the value keeps oscillating. Such oscillations do not occur if T is definable in a o-minimal structure [BGV14], but o-minimality is not enough to prevent milder deviations like the one arising in (5.2) – the map T arising there is actually definable in the o-minimal structure consisting of semi-algebraic sets. This is why we make the *polyhedral* assumption in Theorem 5.1.2.

5.2 Preliminary results

Fix for the remainder of this section an integer $n \in \mathbb{N}_{>0}$ and T a polyhedral nonexpansive selfmap of $\mathbb{R}^n$ for some norm $\|\cdot\|$, as well as an invariant halfline of T of base point $u \in \mathbb{R}^n$ and of direction $\eta \in \mathbb{R}^n$. We construct the following operator which will play a crucial role in the proof of Theorem 5.1.2.

Definition 5.2.1. The *deviation operator* $\check{T}$ associated to T is the selfmap of $\mathbb{R}^n$ defined by
$$\check{T}(v) = \lim_{s\to+\infty} T(v+s\eta) - (s+1)\eta \quad \text{for all } v \in \mathbb{R}^n \ .$$

Remark 5.2.2. The deviation operator $\check{T}$ as it has been defined in the previous definition corresponds to an additive version of the usual recession operator $\hat{T}$ associated to T, which is given for all $v \in \mathbb{R}^n$ by $\hat{T}(v) = \lim_{s\to+\infty} \frac{T(sv)}{s}$ (see for instance [GG04] for more details). In particular, for a polyhedral map, the recession operator is always defined on all $\mathbb{R}^n$ as for all $v \in \mathbb{R}^n$, the map $s \mapsto T(sv)$ is eventually linear in s for $s \geqslant 0$ large enough, similarly to what happens for the deviation operator in the proofs of the following lemmas. In particular, this allows us to define a more general deviation operator $\check{T}_h$ for any direction $h \in \mathbb{R}^n$, by setting $\check{T}_h(v) = \lim_{s\to+\infty} T(v + s\eta) - (s+1)\hat{T}(h)$ for all $v \in \mathbb{R}^n$. In particular, since the direction η of the invariant halfline of T always satisfies $\hat{T}(\eta) = \eta$, this entails that for $h = \eta$, one has $\check{T}_\eta = \check{T}$.

We give a few straightforward properties of the deviation operator.

Proposition 5.2.3. *The map $\check{T}$ is defined on all $\mathbb{R}^n$, is nonexpansive, and moreover, for all $v \in \mathbb{R}^n$, there exists $s_v^* \in \mathbb{R}$ such that for all $s \geqslant s_v^*$, $\check{T}(v) = T(v+s\eta) - (s+1)\eta$.*

Proof. Let $v \in \mathbb{R}^n$. Since the map T is polyhedral, there is a polyhedral decomposition $\mathscr{C}$ of $\mathbb{R}^n$ into a finite union of closed polyhedral cells on which the map T is affine. The halfline $\{v + s\eta \in \mathbb{R}^n : s \geqslant 0\}$ only meets a finite number of these cells, and thus must eventually remain in a single minimal-dimensional cell C_v. Let ϕ be a linear selfmap of $\mathbb{R}^n$ and let $c \in \mathbb{R}^n$ be such that
$$T(x) = \phi(x) + c \quad \text{for all } x \in C_v \ .$$
Then, for all $x \in C_v$, and for all $s > 0$ large enough, one has
$$\begin{aligned} s\|\phi(\eta) - \eta\| - \|\phi(x) + c - u - \eta\| &\leqslant \|s\phi(\eta) - s\eta + \phi(x) + c - u - \eta\| \\ &= \|\phi(x + s\eta) + c - u - (s+1)\eta\| \\ &= \|T(x+s\eta) - T(u+s\eta)\| \\ &\leqslant \|x - u\| \end{aligned}$$
where the last inequality is by nonexpansivity of T, and since the inequality holds for s arbitrarily large, one must necessarily have $\|\phi(\eta) - \eta\| = 0$, *i.e.* $\phi(\eta) = \eta$.

Now, let $s_v^* = \min\{s \geqslant 0 : v + s\eta \in C_v\}$. Then for all $s \geqslant s_v^*$, one has from the previous result $T(v + s\eta) = \phi(v) + s\eta + c$, and thus the quantity $T(v+s\eta) - (s+1)\eta = \phi(v) + c - \eta$ does no longer depend on $s \geqslant s_v^*$, hence the equality $\check{T}(v) = T(v + s\eta) - (s + 1)\eta = \phi(v) + c - \eta$, showing that the deviation operator is indeed defined on all $\mathbb{R}^n$ and satisfies the desired equality.

Finally, for the nonexpansivity, take $v, w \in \mathbb{R}^n$. Then for $s \geqslant \max(s_v^*, s_w^*)$, one has simultaneously the equalities $\check{T}(v) = T(v + s\eta) - (s + 1)\eta$ and $\check{T}(w) = T(w + s\eta) - (s + 1)\eta$, and thus the nonexpansivity of T entails that $\|\check{T}(v) - \check{T}(w)\| = \|T(v + s\eta) - T(w + s\eta)\| \leqslant \|v - w\|$. □

Corollary 5.2.4. *The point u is a fixed point of the deviation operator $\check{T}$.*

Proof. For $s \geqslant 0$ large enough, one has $T(u + s\eta) = u + (s + 1)\eta$, and thus $T(u + s\eta) - (s + 1)\eta = u$, entailing $\check{T}(u) = u$. □

Proposition 5.2.5. *For all $v \in \mathbb{R}^n$ and $t \in \mathbb{R}$, $\check{T}(v + t\eta) = \check{T}(v) + t\eta$.*

Proof. Let $v \in \mathbb{R}^n$ and $t \in \mathbb{R}$. Then for $s \geqslant \max(s^*_{v+t\eta}, s^*_v - t)$, one has

$$\begin{aligned}\check{T}(v + t\eta) &= T((v + t\eta) + s\eta) - (s + 1)\eta \\ &= T(v + (s + t)\eta) - (s + t + 1)\eta + t\eta \\ &= \check{T}(v) + t\eta \ .\end{aligned}$$

□

Corollary 5.2.6. *The set of fixed points of the deviation operator $\check{T}$ coincides precisely with the invariant halfline of T.*

Proof. It follows directly from Corollary 5.2.4 and Proposition 5.2.5 that every point in the invariant halfline of T is in fact a fixed point of the deviation operator. Conversely, any fixed point v of $\check{T}$ satisfies for $s \geqslant 0$ large enough that $\check{T}(v) = T(v + s\eta) - (s + 1)\eta = v$, and thus v must belong to the invariant halfline of T. □

The previous properties of the deviation operator lead to the following two lemma, which is at the center of the proof of the main convergence theorem.

Lemma 5.2.7. *Let V be a bounded subset of $\mathbb{R}^n$. Then there exists a uniform $s^* \geqslant 0$ such that for all $v \in V$ and all $s \geqslant s^*$, $\check{T}(v) = T(v + s\eta) - (s + 1)\eta$.*

Proof. Consider the map $f : v \mapsto s_v^* = \min\{s \geqslant 0 : v + s\eta \in C_v\}$ where C_v is the minimal dimensional cell eventually — that is for $s \geqslant 0$ large enough — containing $v + s\eta$, as constructed in the proof of the previous proposition. In order to obtain the result, it suffices to show that the map f is locally bounded. Begin by considering the H-representation of C_v, and write

$$C_v = \{x \in \mathbb{R}^n : \forall i \in [k],\, \langle a_i, x\rangle \leqslant b_i\}$$

with $k \in \mathbb{N}$, $a_i \in \mathbb{R}^n$ and $b_i \in \mathbb{R}$ for all $i \in [k]$. Then by definition, $v + s\eta$ remains eventually in C implies that η belongs to the recession cone $\operatorname{recc}(C_v)$ of C_v, or equivalently, that $\langle a_i, \eta\rangle \leqslant 0$ for all $i \in [k]$. Therefore, the H-representation of C_v entails that

$$\begin{aligned} v + s\eta \in C_v \iff & s\langle a_i, \eta\rangle \leqslant b_i - \langle a_i, v\rangle \\ \iff & \begin{cases} s \geqslant \frac{b_i - \langle a_i, v\rangle}{\langle a_i, \eta\rangle} & \text{for all } i \in [k] \text{ such that } \langle a_i, \eta\rangle < 0 \\ 0 \leqslant b_i - \langle a_i, v\rangle & \text{for all } i \in [k] \text{ such that } \langle a_i, \eta\rangle = 0 \end{cases} ,\end{aligned}$$

and it follows readily that

$$s_v^* \leqslant \max_{\substack{i \in [k] \\ \langle a_i, \eta\rangle < 0}} \frac{b_i - \langle a_i, v\rangle}{\langle a_i, \eta\rangle} \ .$$

In fact, for every cell C of the complex $\mathscr{C}$ of linearity domains of the operator T such that $\eta \in \operatorname{recc}(C)$, one can consider the collection of inequalities of the form $\langle a, x\rangle \leqslant b$ defining C, and then, by taking the maximum of $\frac{b - \langle a, v\rangle}{\langle a, \eta\rangle}$ such that $\langle a, \eta\rangle < 0$ over all the coefficients $a \in \mathbb{R}^n$ and $b \in \mathbb{R}$ that appear in the H-representations of all the finitely-many cells C whose recession cone contains η, we obtain an upper-bound on s_v^* which does not depend on the cell C_v in which $v + s\eta$ eventually ends up, and this upper-bound is moreover continuous in $v \in \mathbb{R}^n$, which implies that the map $f : v \mapsto s_v^*$ is indeed locally bounded, hence the existence of a uniform bound over any bounded subset of $\mathbb{R}^n$. □

Remark 5.2.8. More precisely, it is in fact possible to prove that the map $f : v \mapsto s_v^*$ is upper-semicontinuous, which does *a fortiori* imply Lemma 5.2.7.

The deviation operator of the Krasnoselskii-Mann operator T_{KM} is related to the deviation operator of T via the relation $(\check{T})_{\mathsf{KM}} = (\check{T_{\mathsf{KM}}})$, and thus we can write the previous quantity as $\check{T}_{\mathsf{KM}}$ without ambiguity. Indeed, since $\chi(T_{\mathsf{KM}}) = \frac{\eta}{2}$, it follows for all $v \in \mathbb{R}^n$ that

$$
\begin{aligned}
(\check{T_{\mathsf{KM}}})(v) &= \lim_{s\to+\infty} T_{\mathsf{KM}}\left(v + s\frac{\eta}{2}\right) - (s+1)\frac{\eta}{2} \\
&= \lim_{s\to+\infty} \frac{1}{2}\left(v + s\frac{\eta}{2} + T\left(v + s\frac{\eta}{2}\right)\right) - (s+1)\frac{\eta}{2} \\
&= \lim_{s\to+\infty} \frac{1}{2}\left(v + T\left(v + s\frac{\eta}{2}\right)\right) - \left(\frac{s}{2}+1\right)\frac{\eta}{2} \\
&= \frac{1}{2}\Big(v + \underbrace{\lim_{s\to+\infty} \frac{1}{2}\left(v + T\left(v + \frac{s}{2}\eta\right)\right) - \left(\frac{s}{2}+1\right)\eta}_{=\check{T}(v)}\Big) = (\check{T})_{\mathsf{KM}}(v) \ .
\end{aligned}
$$

The existence of a fixed point of the deviation operator, paired with its nonexpansivity allows us use the Ishikawa convergence theorem to the deviation operator, leading to this second lemma on which the proof of the main convergence theorem relies.

Lemma 5.2.9. *For all $\check{v}^0 \in \mathbb{R}^n$, the sequence $(\check{v}^N)_{n\in\mathbb{N}}$ defined for all $N \in \mathbb{N}$ by $\check{v}^{N+1} = \check{T}_{\mathsf{KM}}(\check{v}^N)$ converges to a fixed point $\check{v}$ of $\check{T}$.*

Proof. This in an immediate application of Ishikawa's result. One only need to check that the deviation operator satisfies the conditions of the Ishikawa convergence theorem: for any $\check{v}^0 \in \mathbb{R}$, there exists a closed ball D centered at the base point u of the invariant halfline of T, and of radius large enough so that it contains $\check{v}^0$. The nonexpansivity of the deviation operator $\check{T}$ and the fact that u is a fixed point imply that the (convex) ball D is mapped into a compact subset of itself, and therefore Theorem 5.1.1 can be applied. ☐

5.3 The proof of the convergence theorem

Relying on the results of the previous section, one can finally give the proof of our main convergence theorem.

Proof of Theorem 5.1.2. We begin by proving that the sequence $(v^N - N\frac{\eta}{2})_{N\in\mathbb{N}}$ is bounded. Let $(\check{v}^N)_{n\in\mathbb{N}}$ be the sequence defined for all $N \in \mathbb{N}$ by $\check{v}^{N+1} = \check{T}_{\mathsf{KM}}(\check{v}^N)$ and $\check{v}^0 = v^0 \in \mathbb{R}^n$. Since the sequence $(\check{v}^N)_{N\in\mathbb{N}}$ converges, it is bounded, and thus applying Lemma 5.2.7, there exists $s^* \geqslant 0$ uniform in N such that for all $s \geqslant s^*$ and for all $N \geqslant 0$,

$$
\check{v}^{N+1} = \check{T}_{\mathsf{KM}}(\check{v}^N) = T_{\mathsf{KM}}\left(\check{v}^N + s\frac{\eta}{2}\right) - (s+1)\frac{\eta}{2} \ ,
$$

i.e.

$$
T_{\mathsf{KM}}\left(\check{v}^N + s\frac{\eta}{2}\right) = \check{v}^{N+1} + (s+1)\frac{\eta}{2} \ .
$$

Applying the above equality consecutively yields

$$
\begin{aligned}
T_{\mathsf{KM}}^N\left(\check{v}^0 + s^*\frac{\eta}{2}\right) &= T_{\mathsf{KM}}^{N-1}\Big(\underbrace{T_{\mathsf{KM}}\left(\check{v}^0 + s^*\frac{\eta}{2}\right)}_{=\check{v}^1+(s^*+1)\frac{\eta}{2}}\Big) \\
&= T_{\mathsf{KM}}^{N-2}\Big(\underbrace{T_{\mathsf{KM}}\left(\check{v}^1 + (s^*+1)\frac{\eta}{2}\right)}_{=\check{v}^2+(s^*+2)\frac{\eta}{2}}\Big) \\
&\ \ \vdots \\
&= T_{\mathsf{KM}}\left(\check{v}^{N-1} + (s^* + (N-1))\frac{\eta}{2}\right) \\
&= \check{v}^N + (s^*+N)\frac{\eta}{2} \ ,
\end{aligned}
$$

hence for all $N \geqslant 0$,

$$T_{\mathsf{KM}}^N\left(\check{v}^0 + s^*\frac{\eta}{2}\right) - N\frac{\eta}{2} = \left(\check{v}^N + s^*\frac{\eta}{2}\right) \quad . \tag{5.3}$$

Therefore, it follows that

$$\begin{aligned}
\|v^N - N\frac{\eta}{2}\| &= \|T_{\mathsf{KM}}^N(\check{v}^0) - N\frac{\eta}{2}\| \\
&\leqslant \|T_{\mathsf{KM}}^N(\check{v}^0) - T_{\mathsf{KM}}^N\left(\check{v}^0 + s^*\frac{\eta}{2}\right)\| + \|T_{\mathsf{KM}}^N\left(\check{v}^0 + s^*\frac{\eta}{2}\right) - N\frac{\eta}{2}\| \\
&\leqslant s^*\frac{\|\eta\|}{2} + \|\check{v}^N + s^*\frac{\eta}{2}\| \ ,
\end{aligned}$$

where the second inequality comes from the nonexpansivity of T_{KM} and from equality (5.3), and since the sequence $(\check{v}^N)_{N\in\mathbb{N}}$ converges, the quantity $s^*\frac{\|\eta\|}{2} + \|\check{v}^N + s^*\frac{\eta}{2}\|$ is bounded above by some constant $R > 0$.

By applying Lemma 5.2.7 again to the closed ball of radius R centered at the origin, we thus obtain the existence of another $s^{**} \geqslant 0$ uniform such that for all $s \geqslant s^{**}$ and for all $N \geqslant 0$,

$$\check{T}_{\mathsf{KM}}\left(v^N - N\frac{\eta}{2}\right) = T_{\mathsf{KM}}\left(v^N + (s-N)\frac{\eta}{2}\right) - (s+1)\frac{\eta}{2} \ .$$

and recalling from Proposition 5.2.5 that $\check{T}_{\mathsf{KM}}(v^N - N\frac{\eta}{2}) = \check{T}_{\mathsf{KM}}(v^N) - N\frac{\eta}{2}$ and substracting $\frac{\eta}{2}$, it follows that

$$\check{T}_{\mathsf{KM}}(v^N) - \frac{\eta}{2} = T_{\mathsf{KM}}\left(v^N + (s-N)\frac{\eta}{2}\right) - (s-N)\frac{\eta}{2} \ .$$

In particular, for all $N \geqslant s^{**}$, one can take $s = N$ to obtain for all $N \geqslant s^{**}$

$$v^{N+1} = T_{\mathsf{KM}}(v^N) = \check{T}_{\mathsf{KM}}(v^N) + \frac{\eta}{2} = \check{T}_{\mathsf{KM}}\left(v^N + \frac{\eta}{2}\right) \ .$$

Finally, one can again apply the previous equality recursively and obtain for all $N \geqslant N_0 := \lceil s^{**} \rceil$ that

$$v^N - N\frac{\eta}{2} = \check{T}_{\mathsf{KM}}^{N-N_0}(v^{N_0}) \ .$$

Finally, as per Lemma 5.2.9, the quantity $\check{T}_{\mathsf{KM}}^{N-N_0}(v^{N_0})$ is known to converge to some fixed point of the deviation operator, hence $v^N - N\frac{\eta}{2}$ also converges to that same fixed point as N goes to $+\infty$. □

Conclusion and perspectives

The value iteration with widening (Algorithm 1) as well as the dichotomy method (Algorithm 3) have both been implemented in Python. The implementation is available in [Bé23]. Related benchmarks can be viewed in [ABG23a, ABG24].

Following the differents result presented in this manuscript, a certain number of unanswered questions arise, constituting a possible research direction for future works.

- In **Chapter 2**, the main open question is to determine whether the inflation of the degree bound for the tropical Positivstellensatz is necessary or not. The proof we provided does not show that this degree bound is tight, and one could imagine that this inflation is an artefact, resulting specifically from the use of the Shapley-Folkman lemma in our construction. Experimental results on randomly chosen instances seem to indicate that even for systems of tropical polynomial inequalities, applying the tropical Positivstellensatz with the non-inflated degree bound is sufficient to decide the feasibility of the system. Another related question of interest concerns the sparse tropical resultant. Jensen and Yu [JY13] showed that the sparse tropical resultant variety can be obtained as the tropicalization of the classical sparse resultant variety, and described its fan structure. Under some conditions on the supports of the polynomials involved, the sparse resultant variety can be described by a polynomial, called the sparse resultant polynomial, which is then usually computed, via Sylvester-type formulae, as a quotient of two minors of the Macaulay matrix. The existence of such formulae in the tropical setting would complete the present Null- and Positivstellensätze. However, such formulae cannot be obtained 'trivially' by simply replacing the minors of the Macaulay matrix by the same minors of the tropical Macaulay matrix in the formula. Indeed, by comparing the determinant and the permanent of submatrices of the Macaulay matrix, we exhibited numerical examples of polynomial systems such that there are cancellations in the computation of the resultant as a quotient of minors of the Macaulay matrix, but such cancellations cannot happen in the tropical setting.

- In **Chapter 4**, the very first question of interest is to provide actual complexity bounds for the path-following method for parametric mean payoff games. In order to do so, one needs to be able to estimate the number of ζ- and ε-pivotings. One should expectedly be able to bound the number of ζ-pivotings by the number of cells appearing in the linearity complex $\mathscr{C}^{\mathsf{lin}}$. However, providing an upper-bound for the number of ε-pivotings seems trickier, as it apparently only seems to rely on geometrical properties of tropical geodesics. Once this question is addressed, another direction would be to state the previous theorem without the costly 'big M' trick, relying on lexicographic methods, in order to improve the complexity bounds. Also for the specific question of solving tropical polynomial systems, the lazy linearization method proposed could be improved, in particular by adapting it to the sparse case.

Bibliography

[ABG15] X. Allamigeon, V. Bœuf, and S. Gaubert. Performance evaluation of an emergency call center: Tropical polynomial systems applied to timed Petri nets. In S. Sankaranarayanan and E. Vicario, editors, *Proceedings of FORMATS*, volume 9268 of *LNCS*, pages 10–26. Springer, 2015.

[ABG23a] Marianne Akian, Antoine Béreau, and Stéphane Gaubert. The tropical Nullstellensatz and Positivstellensatz for sparse polynomial systems. In *Proceedings of the 2023 International Symposium on Symbolic and Algebraic Computation*, ISSAC '23, page 43–52, New York, NY, USA, 2023. Association for Computing Machinery.

[ABG23b] Marianne Akian, Antoine Béreau, and Stéphane Gaubert. The Nullstellensatz and Positivstellensatz for sparse tropical polynomial systems, 2023. *arXiv:2312.05859*.

[ABG24] Marianne Akian, Antoine Béreau, and Stéphane Gaubert. Eigenvalue methods for sparse tropical polynomial systems. In *Mathematical Software – ICMS 2024*, page 299–312. Springer Nature Switzerland, 2024.

[ACM84] Dennis S. Arnon, George E. Collins, and Scott McCallum. Cylindrical algebraic decomposition i: The basic algorithm. *SIAM Journal on Computing*, 13(4):865–877, 1984.

[AFG$^+$14] X. Allamigeon, U. Fahrenberg, S. Gaubert, A. Legay, and R. Katz. Tropical Fourier-Motzkin elimination, with an application to real-time verification. *International Journal of Algebra and Computation*, 24(5):569–607, 2014.

[AG90] Eugene L. Allgower and Kurt Georg. *Numerical Continuation Methods*. Springer Berlin Heidelberg, 1990.

[AG13] Marianne Akian and Stéphane Gaubert. Policy iteration for perfect information stochastic mean payoff games with bounded first return times is strongly polynomial, 2013.

[AGG09] M. Akian, S. Gaubert, and A. Guterman. Linear independence over tropical semirings and beyond. In G.L. Litvinov and S.N. Sergeev, editors, *Proceedings of the International Conference on Tropical and Idempotent Mathematics*, volume 495 of *Contemporary Mathematics*, pages 1–38. AMS, 2009.

[AGG12] M. Akian, S. Gaubert, and A. Guterman. Tropical polyhedra are equivalent to mean payoff games. *International Journal of Algebra and Computation*, 22(1):125001 (43 pages), 2012.

[AGH15] Marianne Akian, Stéphane Gaubert, and Antoine Hochart. Ergodicity conditions for zero-sum games. *Discrete & Continuous Dynamical Systems - A*, 35(9):3901–3931, 2015.

[AGH18] Marianne Akian, Stéphane Gaubert, and Antoine Hochart. Generic uniqueness of the bias vector of finite stochastic games with perfect information. *Journal of Mathematical Analysis and Applications*, 457(2):1038–1064, 2018.

[AGH20] Marianne Akian, Stéphane Gaubert, and Antoine Hochart. A game theory approach to the existence and uniqueness of nonlinear perron-frobenius eigenvectors. *Discrete and Continuous Dynamical Systems*, 40(1):207–231, 2020.

[AGK11a] X. Allamigeon, S. Gaubert, and R. Katz. The number of extreme points of tropical polyhedra. *J. Comb. Theory Series A*, 118(1):162–189, 2011.

[AGK11b] X. Allamigeon, S. Gaubert, and R. Katz. Tropical polar cones, hypergraph transversals, and mean payoff games. *Linear Algebra and its Applications*, 435(7):1549–1574, 2011.

[AGKS22] X. Allamigeon, S. Gaubert, R. D. Katz, and M. Skomra. Universal Complexity Bounds Based on Value Iteration and Application to Entropy Games. In Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff, editors, *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*, volume 229 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 110:1–110:20, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[AGQS23] Marianne Akian, Stéphane Gaubert, Yang Qi, and Omar Saadi. Tropical linear regression and mean payoff games: Or, how to measure the distance to equilibria. *SIAM J. Disc. Math.*, 37(2):632–674, May 2023.

[AGS20] X. Allamigeon, S. Gaubert, and M. Skomra. Tropical spectrahedra. *Discrete and Computational Geometry*, 63:507–548, February 2020.

[AGT16] Fuensanta Aroca, Cristhian Garay, and Zeinab Toghani. The fundamental theorem of tropical differential algebraic geometry. *Pacific Journal of Mathematics*, 283(2):257–270, June 2016.

[AGV23] Marianne Akian, Stephane Gaubert, and Sara Vannucci. Ambitropical geometry, hyperconvexity and zero-sum games, 2023.

[AI22] Omid Amini and Hernan Iriarte. Geometry of higher rank valuations, 2022.

[Aro10a] Fuensanta Aroca. Krull-tropical Hypersurfaces. *Annales de la Faculté des sciences de Toulouse : Mathématiques*, Ser. 6, 19(3-4):525–538, 2010.

[Aro10b] Fuensanta Aroca. Tropical geometry for fields with a Krull valuation: First definitions and a small result. *Boletin De La Sociedad Matematica Mexicana*, 16:9–14, 2010.

[BCC$^+$05] M. Bronstein, A. M. Cohen, H. Cohen, D. Eisenbud, B. Sturmfels, A. Dickenstein, and I. Z. Emiris, editors. *Solving Polynomial Equations*. Springer-Verlag, 2005.

[BCOQ92] François Louis Baccelli, Guy Cohen, Geert Jan Olsder, and Jean-Pierre Quadrat. *Synchronization and linearity. An algebra for discrete event systems*. Chichester: Wiley, 1992.

[BCR98] Jacek Bochnak, Michel Coste, and Marie-Françoise Roy. *Real Algebraic Geometry*. Springer Berlin Heidelberg, 1998.

[BE17] A. Bertram and R. Easton. The tropical Nullstellensatz for congruences. *Advances in Mathematics*, 308:36–82, February 2017.

[BGV14] Jérôme Bolte, Stéphane Gaubert, and Guillaume Vigeral. Definable zero-sum stochastic games. *Mathematics of Operations Research*, 40(1):171–191, 2014.

[BK19] E. Baldwin and P. Klemperer. Understanding preferences: “demand types”, and the existence of equilibrium with indivisibilities. *Econometrica*, 87(3):867–932, 2019.

[BNR08] Marc Bezem, Robert Nieuwenhuis, and Enric Rodríguez-Carbonell. *The Max-Atom Problem and Its Relevance*, page 47–61. Springer Berlin Heidelberg, 2008.

[Bre99] D. Bremner. Incremental convex hull algorithms are not output sensitive. *Discrete & Computational Geometry*, 21(1):57–68, January 1999.

[Bro87] W. Dale Brownawell. Bounds for the degrees in the Nullstellensatz. *Annals of Mathematics*, 126(3):577–591, 1987.

[But10] Peter Butkovič. *Max-linear Systems: Theory and Algorithms*. Springer London, 2010.

[Bé23] Antoine Béreau. Tp2s: a solver for systems of tropical polynomial equalities and inequalities, 2023. Python software available from `https://gitlab.inria.fr/abereau/tropical-polynomial-system-solving`.

[CC77] Patrick Cousot and Radhia Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the 4th ACM SIGACT-SIGPLAN symposium on Principles of programming languages - POPL '77*, POPL'77, page 238–252. ACM Press, 1977.

[CE93] J. Canny and I. Z. Emiris. An efficient algorithm for the sparse mixed resultant. In *Applied algebra, algebraic algorithms and error-correcting codes. Proceedings of 10th international symposium, AAECC-10, San Juan de Puerto Rico, May 10-14, 1993*, pages 89–104. Springer-Verlag, Berlin, 1993.

[CGQ98] G. Cohen, S. Gaubert, and J.P. Quadrat. Algebraic system analysis of timed Petri nets. In J. Gunawardena, editor, *Idempotency*, Publications of the Isaac Newton Institute, pages 145–170. Cambridge University Press, 1998.

[Cha09] J. Chaloupka. Parallel algorithms for mean-payoff games: an experimental evaluation. In *Algorithms—ESA 2009*, volume 5757 of *Lecture Notes in Comput. Sci.*, pages 599–610. Springer, Berlin, 2009.

[Cha21] Antoine Chambert-Loir. Topics in tropical geometry, 2021. Lecture notes for the 2020-2021 course Géométrie Tropicale at Sorbonne Université, Paris, see: `https://webusers.imj-prg.fr/~antoine.chambert-loir/enseignement/2020-21/gt/toptrop.pdf`.

[CLO15] D. A. Cox, J. Little, and D. O'Shea. *Ideals, Varieties, and Algorithms*. Springer International Publishing, 2015.

[CM80] R.A. Cuninghame-Green and P.F.J. Meijer. An algebra for piecewise-linear minimax problems. *Discrete Applied Mathematics*, 2(4):267–294, 1980.

[CTGG99] J. Cochet-Terrasson, S. Gaubert, and J. Gunawardena. A constructive fixed point theorem for min-max functions. *Dynamics and Stability of Systems*, 14(4):407–433, 1999.

[Den86] Jan Denef. p-adic semi-algebraic sets and cell decomposition. *Journal für die reine und angewandte Mathematik*, 369:154–166, 1986.

[DG06] Vishesh Dhingra and Stéphane Gaubert. How to solve large scale deterministic games with mean payoff by policy iteration. In *Valuetools '06: Proceedings of the 1st international conference on Performance evaluation methodologies and tools*, page 12, New York, NY, USA, 2006. ACM Press.

[DRS10] Jesús A. De Loera, Jörg Rambau, and Francisco Santos. *Triangulations*. Springer Berlin Heidelberg, 2010.

[DSR22] Aurélien Desoeuvres, Peter Szmolyan, and Ovidiu Radulescu. Qualitative dynamics of chemical reaction networks: An investigation using partial tropical equilibrations. In Ion Petre and Andrei Păun, editors, *Computational Methods in Systems Biology*, pages 61–85, Cham, 2022. Springer International Publishing.

[EKL06] M. Einsiedler, M. Kapranov, and D. Lind. Non-archimedean amoebas and tropical varieties. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 2006(601), January 2006.

[EM79] A. Ehrenfeucht and J. Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8(2):109–113, June 1979.

[EM07] Mohamed Elkadi and Bernard Mourrain. *Introduction à la résolution des systèmes polynomiaux*. Springer, 2007.

[Emi05] I. Z. Emiris. Toric resultants and applications to geometric modelling. In [BCC$^+$05], pages 269–300. 2005.

[EP05] A.J. Engler and A. Prestel. *Valued Fields*. Springer Monographs in Mathematics. Springer Berlin Heidelberg, 2005.

[FP96] Komei Fukuda and Alain Prodon. *Double description method revisited*, page 91–111. Springer Berlin Heidelberg, 1996.

[Fri11] Oliver Friedmann. An exponential lower bound for the latest deterministic strategy iteration algorithms. *Logical Methods in Computer Science*, Volume 7, Issue 3, October 2011.

[GG98] Séphane Gaubert and Jeremy Gunawardena. The duality theorem for min-max functions. *Comptes Rendus de l'Académie des Sciences - Series I - Mathematics*, 326(1):43–48, 1998.

[GG04] Stéphane Gaubert and Jeremy Gunawardena. The perron-frobenius theorem for homogeneous, monotone functions. *Transactions of the American Mathematical Society*, 356(12):4931–4950, March 2004.

[Gil57] Dean Gillette. Stochastic games with zero stop probabilities. *Contributions to the Theory of Games*, 3(39):179–187, 1957.

[Giu84] M. Giusti. Some effectivity problems in polynomial ideal theory. In John Fitch, editor, *EUROSAM 84*, pages 159–171, Berlin, Heidelberg, 1984. Springer Berlin Heidelberg.

[GKK90] V. A. Gurvich, A. V. Karzanov, and L. G. Khachivan. Cyclic games and an algorithm to find minimax cycle means in directed graphs. *USSR Comput. Math. Math. Phys.*, 28(5):85–91, April 1990.

[GKS12] Stéphane Gaubert, Ricardo D. Katz, and Sergei Sergeev. Tropical linear-fractional programming and parametric mean payoff games. *Journal of symbolic computation*, 47(12):1447–1478, December 2012.

[GKZ94] I. M. Gelfand, M. M. Kapranov, and A. V. Zelevinsky. *Discriminants, Resultants, and Multidimensional Determinants*. Birkhäuser, 1994.

[GP13] D. Grigoriev and V. V. Podolskii. Complexity of tropical and min-plus linear prevarieties. *Computational Complexity*, pages 1–34, 2013.

[GP18] D. Grigoriev and V. Podolskii. Tropical effective primary and dual Nullstellensätze. *Discrete & Computational Geometry*, 59:507–552, 04 2018.

[GP20] Dima Grigoriev and Vladimir V. Podolskii. Tropical combinatorial Nullstellensatz and sparse polynomials. *Found. Comput. Math.*, 20(4):753–781, aug 2020.

[Gri12] Dima Grigoriev. On a tropical dual Nullstellensatz. *Advances in Applied Mathematics*, 48(2):457–464, February 2012.

[GRZ22] P. Görlach, Y. Ren, and L. Zhang. Computing zero-dimensional tropical varieties via projections. *Computational complexity*, 31(1), May 2022.

[GV01] Dima Grigoriev and Nicolai Vorobjov. Complexity of null- and positivstellensatz proofs. *Annals of Pure and Applied Logic*, 113(1-3):153–160, December 2001.

[Hil93] David Hilbert. Ueber die vollen invariantensysteme. *Mathematische Annalen*, 42(3):313–373, September 1893.

[HK66] A. J. Hoffman and R. M. Karp. On nonterminating stochastic games. *Management Science. Journal of the Institute of Management Science. Application and Theory Series*, 12:359–370, 1966.

[HM06] Marshall Hampton and Richard Moeckel. Finiteness of relative equilibria of the four-body problem. *Invent. Math.*, 163(2):289–312, February 2006.

[HMZ11] T.D. Hansen, P.B. Miltersen, and U. Zwick. Strategy iteration is strongly polynomial for 2-player turn-based stochastic games with a constant discount factor. In *Innovations in Computer Science 2011*, pages 253–263. Tsinghua University Press, 2011.

[HS95] B. Huber and B. Sturmfels. A polyhedral method for solving sparse polynomial systems. *Mathematics of Computation*, 64(212):1541–1555, 1995.

[IMS09] I. Itenberg, G. Mikhalkin, and E. Shustin. *Tropical algebraic geometry*, volume 35 of *Oberwolfach Semin.* Birkhäuser, Basel, second edition, 2009.

[IR10] Zur Izhakian and Louis Rowen. Supertropical algebra. *Advances in Mathematics*, 225(4):2222–2286, November 2010.

[Ish76] Shiro Ishikawa. Fixed points and iteration of a nonexpansive mapping in a banach space. *Proceedings of the American Mathematical Society*, 59(1):65–71, 1976.

[IV96] I. Itenberg and O. Viro. Patchworking algebraic curves disproves the Ragsdale conjecture. *Math. Intelligencer*, 18(4):19–28, 1996.

[Jen16] A. N. Jensen. Tropical homotopy continuation, 2016. arXiv:1601.02818.

[JM17] Dániel Joó and Kalina Mincheva. Prime congruences of additively idempotent semirings and a nullstellensatz for tropical polynomials. *Selecta Mathematica*, 24(3):2207–2233, March 2017.

[JPZ08] M. Jurdziński, M. Paterson, and U. Zwick. A deterministic subexponential algorithm for solving parity games. *SIAM J. Comput.*, 38(4):1519–1532, 2008.

[JS23] Michael Joswig and Ben Smith. Convergent Hahn series and tropical geometry of higher rank. *Journal of the London Mathematical Society*, 107(4):1450–1481, 2023.

[JSY20] Ph. Jell, C. Scheiderer, and J. Yu. Real tropicalization and analytification of semialgeaic sets. *International Mathematics Research Notices*, 2022(2):928–958, May 2020.

[JY13] Anders Jensen and Josephine Yu. Computing tropical resultants. *Journal of Algebra*, 387:287–319, 2013.

[Kat07] Ricardo David Katz. Max-plus (a, b)-invariant spaces and control of timed discrete-event systems. *IEEE Trans. Autom. Control.*, 52(2):229–241, 2007.

[KL93] A.V. Karzanov and V.N. Lebedev. Cyclical games with prohibitions. *Math. Programming*, 60:277–293, 1993.

[Koh80] Elon Kohlberg. Invariant half-lines of nonexpansive piecewise-linear transformations. *Mathematics of Operations Research*, 5(3):366–372, 1980.

[Kol88] János Kollár. Sharp effective nullstellensatz. *Journal of the American Mathematical Society*, 1(4):963–975, 1988.

[Kra55] M. A. Krasnosel'skiĭ. Two remarks on the method of successive approximations. *Uspekhi Matematicheskikh Nauk*, 10:123–127, 1955.

[Laz81] Daniel Lazard. Résolution des systèmes d'équations algébriques. *Theoretical Computer Science*, 15(1):77–110, 1981.

[Laz83] D. Lazard. Gröbner bases, Gaussian elimination and resolution of systems of algebraic equations. In J. A. van Hulzen, editor, *Computer Algebra*, pages 146–156. Springer Berlin Heidelberg, 1983.

[LL69] T. M. Liggett and S. A. Lippman. Stochastic games with perfect information and time average payoff. *SIAM Rev.*, 11:604–607, 1969.

[LS24] Bruno Loff and Mateusz Skomra. Smoothed analysis of deterministic discounted and mean-payoff games. In *51st International Colloquium on Automata, Languages, and Programming (ICALP 2024)*, 2024.

[Lü20] Ch. Lüders. Computing tropical prevarieties with satisfiability modulo theories (smt) solvers. In P. Fontaine, K. Korovin, I. S. Kotsireas, Ph. Rümmer, and S. Tourret, editors, *Proceedings of SC2'20: Fifth International Workshop on Satisfiability Checking and Symbolic Computation, July 05, 2020, Paris, France*. CEUR Workshop Proceedings (CEUR-WS.org), 2020.

[Mal16] G. Malajovich. Computing mixed volume and all mixed cells in quermassintegral time. *Foundations of Computational Mathematics*, 17(5):1293–1334, May 2016.

[Man53] W. R. Mann. Mean value methods in iteration. *Proceedings of the American Mathematical Society*, 4:506–510, 1953.

[Mik05] G. Mikhalkin. Enumerative tropical algebraic geometry in $\mathbb{R}^2$. *J. Amer. Math. Soc.*, 18:313–377, 2005.

[MN81] J.-F. Mertens and A. Neyman. Stochastic games. *Internat. J. Game Theory*, 10(2):53–66, 1981.

[MR18a] D. Maclagan and F. Rincón. Tropical ideals. *Compositio Mathematica*, 154(3):640–670, 2018.

[MR18b] Grigory Mikhalkin and Johannes Rau. Tropical geometry, 2018. Textbook in preparation, see: `https://webusers.imj-prg.fr/~antoine.chambert-loir/enseignement/2020-21/gt/toptrop.pdf`.

[MR19] Th. Markwig and Y. Ren. Computing tropical varieties over fields with valuation. *Foundations of Computational Mathematics*, 20(4):783–800, August 2019.

[MRTT53] T. S. Motzkin, H. Raiffa, G. L. Thompson, and R. M. Thrall. *3. The Double Description Method*, page 51–74. Princeton University Press, December 1953.

[MRZ22] Guido Montúfar, Yue Ren, and Leon Zhang. Sharp bounds for the number of regions of maxout networks and vertices of minkowski sums. *SIAM Journal on Applied Algebra and Geometry*, 6(4):618–649, 2022.

[MS15] D. Maclagan and B. Sturmfels. *Introduction to Tropical Geometry*. Graduate Studies in Mathematics. American Mathematical Society, 2015.

[MSS04] Rolf H. Möhring, Martin Skutella, and Frederik Stork. Scheduling with and/or precedence constraints. *SIAM Journal on Computing*, 33(2):393–415, 2004.

[Nus88] R. D. Nussbaum. Hilbert's projective metric and iterated nonlinear maps. *Memoirs of the AMS*, 75(391), 1988.

[Pas89a] J. Pas. Uniform p-adic cell decomposition and local zeta functions. *J. Reine Angew. Math.*, 399:137–172, 1989.

[Pas89b] Johan Pas. Uniform p-adic cell decomposition and local zeta functions. *Journal für die reine und angewandte Mathematik*, 399:137–172, 1989.

[PR04] Mikael Passare and Hans Rullgård. Amoebas, monge-ampère measures, and triangulations of the newton polytope. *Duke Mathematical Journal*, 121(3), February 2004.

[RG95] Motakuri Ramana and A. J. Goldman. Some geometric results in semidefinite programming. *Journal of Global Optimization*, 7(1):33–50, July 1995.

[Rob56] A. Robinson. *Complete Theories*. Studies in logic and the foundations of mathematics. North-Holland Publishing Company, 1956.

[Roc70] Ralph Tyrell Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, 1970.

[Sch13] R. Schneider. *Convex Bodies: The Brunn–Minkowski Theory*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2 edition, 2013.

[SI07] E. Shustin and Z. Izhakian. A tropical Nullstellensatz. *Proc. Amer. Math. Soc.*, 135:815–3821, 2007.

[ST11] Bernd Sturmfels and Ngoc Tran. Combinatorial types of tropical eigenvectors. *Bulletin of the London Mathematical Society*, 45, 05 2011.

[Stu94] B. Sturmfels. On the Newton Polytope of the Resultant. *Journal of Algebraic Combinatorics*, 3(2):207–236, 1994.

[Tar48] A. Tarski. *A Decision Method for Elementary Algebra and Geometry*. Rand Corporation, 1948.

[Tho06] R. R. Thomas. *Lectures in Geometric Combinatorics*. IAS/Park city mathematical subseries. American Mathematical Society, 2006.

[Vir89] O. Ya. Viro. Real plane algebraic curves: constructions with controlled topology. *Algebra i Analiz*, 1(5):1–73, 1989.

[Ye05] Yinyu Ye. A new complexity result on solving the Markov decision problem. *Mathematics of Operations Research*, 30(3):733–749, 2005.

[Zie95] Günter M. Ziegler. *Lectures on Polytopes*. Springer New York, 1995.

[ZP96] U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoret. Comput. Sci.*, 158(1–2):343–359, 1996.

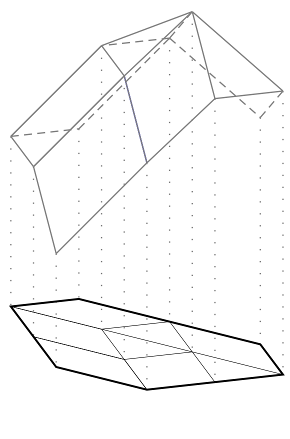

ECOLE
DOCTORALE
DE MATHEMATIQUES
HADAMARD

Titre : Systèmes polynomiaux tropicaux et théorie des jeux

Mots clés : Géométrie tropicale ; Jeux à somme nulle ; Systèmes polynomiaux ; Polyèdres ; Valeurs propres non-linéaires

Résumé : Le but des travaux présentés dans ce manuscrit est de décider efficacement la résolubilité de systèmes polynomiaux tropicaux, puis d'en déterminer effectivement l'ensemble des solutions.
En 2018, Grigoriev et Podolskii ont établi un analogue tropical du Nullstellensatz effectif, montrant ainsi que la résolubilité d'un système d'équations tropicales polynomiales était équivalente à la résolubilité d'un système linéarisé, obtenu en tronquant la *matrice de Macaulay*. Nous établissons une version améliorée du *Nullstellensatz tropical*, prenant en compte la possible structure creuse des polynômes tropicaux du système. Notre résultat repose sur une construction due à Canny et Emiris en 1993. Notre résultat permet de combler l'écart entre le degré de troncature de Grigoriev et Podolskii, et la *borne de Macaulay* dans le cas classique. En outre, nous établissons grâce à la même construction un *Positivstellensatz tropical*, permettant de résoudre les problèmes d'inclusion d'ensembles semi-algébriques tropicaux basiques.
La résolution de tels systèmes linéaires tropicaux se réduit à celle des jeux avec paiement moyen. En particulier, nous proposons une accélération de l'algorithme classique d'*itération sur les valeurs* de Zwick et Paterson, que l'on emploie ensuite pour déterminer la résolubilité d'un système d'équations et inéquations polynomiales tropicales. Cette itération sur les valeurs avec élargissement a été implémentée en Python.
Nous développons ensuite un analogue tropical des *méthodes de valeurs propres* afin de calculer de manière effective l'ensemble des solutions d'un système polynomial tropical. Nous montrons que cet ensemble de solutions peut être déterminé en résolvant des jeux paramétriques, provenant de linéarisations adéquates du système polynomial initial. Nous présentons deux approches : une première basée sur la recherche dichotomique, et une seconde, plus élaborée, basée sur le *suivi de chemin homotopique*.
Enfin, nous présentons une généralisation du *théorème de convergence d'Ishikawa* sur l'itération de Krasnoselskii-Mann, en l'étendant au cas sans point fixe.

Title : Tropical Polynomial Systems and Game Theory

Keywords : Tropical geometry ; Zero-sum games ; Polynomial systems ; Polyhedra ; Nonlinear eigenvalues

Abstract : Given a tropical polynomial system, the aim of the present work is to be able to efficiently decide its solvability, and then effectively compute the solution set.
In 2018, Grigoriev and Podolskii established a tropical analogue of the effective Nullstellensatz, showing that the solvability of a system of tropical polynomial equations is equivalent to the solvability of a linearized system. We establish an improved *tropical Nullstellensatz*, taking into consideration the possible *sparsity* of the tropical polynomials in a system. We rely on a construction of Canny and Emiris from 1993. Our result closes the gap between the truncation degree obtained by Grigoriev and Podolskii and the classical *Macaulay degree bound*. Furthermore, we establish a more general *tropical Positivstellensatz* based on the very same construction, allowing one to decide the inclusion of tropical basic semialgebraic sets.
Such tropical linear systems are known to be reducible to *mean payoff games*. In particular, we propose a speedup of the classical *value iteration* algorithm of Zwick and Paterson, which we then use in order to decide the solvability of a system of tropical polynomial equalities and inequalities. This value iteration algorithm with widening was implemented in Python.
We then develop a tropical analogue of *eigenvalue methods* in order to effectively compute the solution set of tropical polynomial systems. We show that this solution set can be obtained by solving parametric mean-payoff games, arising from approriate linearizations of the tropical polynomial system. We present two approaches : a first one based on a dichotomic search, and a second, more elaborate approach, based on a *tropical homotopy* technique.
Finally, we present a generalization of the *Ishikawa fixed-point convergence theorem*, expanding it to the fixed-point free case.

Institut Polytechnique de Paris
91120 Palaiseau, France