

Tropical linear programming and parametric mean payoff games

May 7, 2010 Stéphane Gaubert*

Stephane.Gaubert@inria.fr, Ricardo D. Katz
rkatz@fceia.unr.edu.ar and Sergeĭ Sergeev†sergeevs@maths.bham.ac.uk

¹ INRIA and CMAP

École Polytechnique, France

² CONICET

Universidad de Rosario, Argentina

³ University of Birmingham

Birmingham, U.K.

Abstract

Tropical polyhedra have been recently used to represent disjunctive invariants in static analysis. To handle larger instances, the tropical analogues of classical linear programming results need to be developed. This motivation leads us to study a general tropical linear programming problem. We construct an associated parametric mean payoff game problem, and show that the optimality of a given point, or the unboundedness of the problem, can be certified by exhibiting a strategy for one of the players having certain infinitesimal properties (involving the value of the game and its derivative) that we characterize combinatorially. In other words, strategies play in tropical linear programming the role of Lagrange multipliers in classical linear programming. We use this idea to design a Newton-like algorithm to solve tropical linear programming problems, by reduction to a sequence of auxiliary mean payoff game problems.

1 Introduction

Motivation from static analysis Tropical algebra is the structure in which the set of real numbers, completed with $-\infty$, is equipped with the “additive” law “ $a+b$ ” := $\max(a, b)$ and the “multiplicative” law “ ab ” := $a+b$. The max-plus or tropical analogues of convex sets have been studied by a number of authors [Zim77, CG79, GP97, LMS01, CGQ04, DS04, BH04, Jos05], under various names (idempotent spaces, semimodules, \mathbb{B} -convexity, extremal convexity), with different degrees of generality, and various motivations.

In a recent work [AGG08], Allamigeon, Gaubert and Goubault have used tropical polyhedra to compute disjunctive invariants in static analysis. A general (affine) tropical polyhedron can be represented as

$$P := \{x \in (\mathbb{R} \cup \{-\infty\})^n \mid \max_{j \in [n]} (a_{ij} + x_j), b_i \leq \max_{j \in [n]} (c_{ij} + x_j), d_i \text{ , } \forall i \in [m]\} . \quad (1)$$

Here, we use the notation $[n] := \{1, \dots, n\}$, and the parameters a_{ij} , c_{ij} , b_i , d_i are given, with values in $\mathbb{R} \cup \{-\infty\}$. The analogy with classical polyhedra becomes clearer with the

*Partially supported by the Arpege programme of the French National Agency of Research (ANR), project “ASOPT”, number ANR-08-SEGI-005 and by the Digiteo project DIM08 “PASO” number 3389.

†Supported by the EPSRC grant RRAH12809 and the RFBR grant 08-01-00601.

tropical notation, which allows us to write the constraints as “ $Ax + b \leq Cx + d$ ”, to be compared with classical systems of linear inequalities, $Ax \leq d$ (in the tropical setting, we need to put affine maps on both sides of the inequality due to the absence of opposite law of the addition). The previous representation of P is the analogue of the *external* representation of polyhedra, as the intersection of half-spaces. As in the classical case, tropical polyhedra have a dual (internal) representation, which involves extreme points and extreme rays. The tropical analogue of Motzkin double description method allows one to pass from one representation to the other [AGG10b].

Disjunctive invariants arise naturally when analyzing sorting algorithms or in the verification of string manipulation programs. The well known `memcpy` function of C is discussed in [AGG08] as a simple illustration: when copying the first `n` characters of a string buffer `src` to a string buffer `dst`, the length `len_dst` of the latter buffer may differ from the length `len_src` of the former, for if `n` is smaller than `len_src`, the null terminal character of the buffer `src` is not copied. However, the relation $\min(\text{len_src}, n) = \min(\text{len_dst}, n)$ is valid. This can be expressed geometrically by saying that the vector $(-\text{len_src}, -\text{len_dst}, -n)$ belongs to a tropical polyhedron. Examples of this nature motivated the development in [AGG08, All09] of an abstract interpretation method in which tropical polyhedra are efficiently handled by the double description algorithm. However, as in the case of usual polyhedra, the scalability of the approach is inherently limited by the exponential blow up of the size of representations of polyhedra [AGK10a, AGK10b].

The complexity of earlier polyhedral approaches led Sankaranarayanan, Colon, Sipma and Manna to introduce the method of *templates* [SCSM06]. In a nutshell, a template consists of a finite set $\mathcal{T} = \{h_1, \dots, h_m\}$ of linear forms on \mathbb{R}^n . The latter define a parametric family of polyhedra with precisely m degrees of freedom $\alpha_1, \dots, \alpha_m \in \mathbb{R} \cup \{+\infty\}$,

$$P_\alpha = \{x \in \mathbb{R}^n \mid h_k(x) \leq \alpha_k, \quad k \in [m]\} .$$

The classical domains of boxes or the domain of *zones* (potential constraints) [Min04] are recovered by incorporating in the template the linear forms $h(x) = \pm x_i$ or $h(x) = x_i - x_j$, respectively. Fixing the template, or changing it dynamically while keeping m bounded, avoids the exponential blow up.

The method of [SCSM06] relies critically on linear programming, which allows one to evaluate quickly the fixed point functional of abstract interpretation. However, the precision of the invariants remains limited by the linear nature of templates, and it is natural to ask whether the machinery of templates carries over to the non-linear case. A framework based on quadratic templates has been introduced by Assale, Gaubert and Goubault [AGG10a]: then, linear programming is replaced by semidefinite programming, thanks to Shor’s relaxation. More generally, every tractable subclass of optimization problems yields a tractable template. Thus, in order to compute disjunctive invariants, it is natural to ask for a different generalization, relying on tropical polyhedra. To this end, the relevant results of linear programming must be tropicalized first, which is what we do here. In particular, comparing the expressions of P and P_α , we see that the linear forms of [SCSM06] must now be replaced by

$$h(x) = \max \left(\max_{j \in [n]} (p_j + x_j), r \right) - \max \left(\max_{j \in [n]} (q_j + x_j), s \right) \quad (2)$$

where p_j, r, q_j, s are given parameters with values in $\mathbb{R} \cup \{-\infty\}$.

The problem In this paper, we study the following general tropical linear programming problem:

$$\text{minimize } h(x) \quad \text{subject to } x \in P, \quad (3)$$

where P is given by (1) and h is given by (2). This is the tropical analogue of the classical problem

$$\text{minimize } (px + r)/(qx + s) \quad \text{subject to } Ax + b \leq Cx + d, \quad x \geq 0,$$

where p, q are nonnegative vectors, r, s are nonnegative scalars, and A, C, b, d are matrices and vectors. The constraint “ $x \geq 0$ ” is implicit in (3), since any number is “positive” (i.e. $\geq -\infty$) in the tropical world.

Problem (3) includes as special cases

$$h(x) = \max\left(\max_{j \in [n]}(p_j + x_j), r\right) \quad \text{and} \quad h(x) = -\max\left(\max_{j \in [n]}(q_j + x_j), s\right) \quad (4)$$

(take $q_j \equiv -\infty$ and $s = 0$, or $p_j \equiv -\infty$ and $r = 0$). Minimizing the ratio of affine forms is not so usual in linear programming, however, in the tropical context, this turns out to be needed both in applications and for theoretical purposes. In particular, in static analysis, we need typically to compute the tightest inequality of the form $x_i \leq K + x_j$ satisfied by the elements of P . This fits in the general form (2), but not in the special cases (4).

Contribution A basic question in linear programming is to certify the optimality of a given point. This is classically done by exhibiting a feasible solution (i.e. a vector of Lagrange multipliers) of the dual problem with the same value. There is no such a simple result in the tropical setting, because as remarked in [GK09], there are (tropically linear) inequalities which can be logically deduced from some finite system of inequalities but which cannot be obtained by taking (positive) linear combinations of the inequalities of this system. In other words, the usual statement of Farkas lemma is not valid in the tropical setting. However, in a recent work, building on [AGG09], Allamigeon, Gaubert, and Katz [AGK10b] have established a tropical analogue of Farkas lemma, in which Lagrange multipliers are replaced by strategies of an associated mean payoff game. We use the same idea here, and show (Theorem 8 below) that the optimality of a solution can be (concisely) certified by exhibiting a strategy of a game, having certain combinatorial properties. Similarly, whether the value of the linear programming problem is unbounded can also be certified in terms of strategies.

The second ingredient is to think of the tropical linear programming problem as a parametric mean payoff game problem. Then, the tropical linear programming problem reduces to the computation of the minimal parameter for which the value of the game is nonnegative.

The main contribution of this paper is a Newton-like method, in which at each iteration we select a strategy (whose existence is implied by the fact that the current feasible point is not optimal). This defines a *one player* parametric game problem, and we show that the smallest value of the parameter making the value zero can be computed in polynomial time for this subgame. The master algorithm requires solving at each step one or two auxiliary mean payoff games (which can be done either by value iteration, which is pseudo-polynomial, or by policy iteration, for which exponential time instances have been recently constructed by Friedmann [Fri09], although the algorithm is fast on typical examples). The number of Newton type iterations of the present algorithm has a trivial exponential bound (the number of strategies). Although the algorithm seems to behave well on typical examples, some further work would be needed to assess its worst case complexity (its behavior is likely to be similar to the one of policy iteration).

Related work Butkovič and Aminu [BA08] studied the special cases (4). At each step, they solve a feasibility problem (whether a tropical polyhedron is non-empty), which is equivalent to checking whether a mean payoff game is winning. However, their algorithm does not involve a Newton-like iteration, but rather a dichotomy argument. This leads to a worst case complexity bound which is incomparable with the one of the present algorithm: the number of calls to a mean payoff oracle in [BA08] depends on the size of the integers in the input, whereas the number of calls in the present algorithm can be bounded independently of these, just in terms of strategies.

Further motivation Tropical polyhedra have been used in [Kat07] to determine invariants of discrete event systems. Systems of constraints equivalent to the ones defining tropical polyhedra have also appeared in the analysis of delays in digital circuits, and in the study of scheduling problems with both “and” and “or” constraints [MSS04]. Such systems have been studied by Bezem, Nieuwenhuis, and Rodríguez-Carbonell [BNgC08, BNgC10], under the name of “max-atom problem”. The latter is motivated by SAT Modulo theory solving, since conjunctions of max-atoms determine a remarkable fragment of linear arithmetic. Tropical polyhedra also turn out to be interesting mathematical objects in their own right [DS04, Jos05]. A final motivation arises from mean payoff games, the complexity of which is a well known open problem: a series of works show that a number of problems which can be expressed in terms of tropical polyhedra are polynomial time equivalent to mean payoff games problems [MSS04, DG06, AGG09, BNgC10, AGK10b]. The results of the present paper rely on [AGG09, AGK10b]. Most proofs are omitted here due to the lack of space.

Acknowledgments The authors thank Peter Butkovič for many useful discussions concerning tropical linear programming and tropical linear algebra. The first author thanks Xavier Allamigeon and Éric Goubault for having shared with him their insights on disjunctive invariants and static analysis.

2 Preliminaries

2.1 Mean-payoff games and min-max functions

Consider a two-player deterministic game, where the players “Max” and “Min” make alternate moves of a pawn on a weighted bipartite digraph \mathcal{G} . The set of nodes of \mathcal{G} is the disjoint union of the nodes $[m] := \{1, \dots, m\}$ where Max is active, and the nodes $[n] := \{1, \dots, n\}$ where Min is active. When the pawn is in node $k \in [m]$ of Max, he must choose an arc in \mathcal{G} connecting node k with some node $l \in [n]$ of Min, and while moving the pawn along this arc, he receives payment b_{kl} from Min, which is the weight of the selected arc. When the pawn is in node $j \in [n]$ of Min, she must choose an arc in \mathcal{G} connecting node j with some node $i \in [m]$ of Max, and pays $-a_{ij}$ to Max, where $-a_{ij}$ is the weight of the selected arc. We assume that $b_{kl}, a_{ij} \in \mathbb{R}$. Moreover, certain moves may be prohibited, meaning that the corresponding arcs are not present in \mathcal{G} . Then, we set $b_{kl} = -\infty$ and $a_{ij} = -\infty$. Thus, the whole game is equivalently defined by two $m \times n$ matrices $A = (a_{ij})$ and $B = (b_{kl})$ with entries in $\mathbb{R} \cup \{-\infty\}$. We make the following assumptions, which mean that both players have at least one move allowed in each node.

Assumption 1. For all $k \in [m]$ there exists $l \in [n]$ such that $b_{kl} \neq -\infty$.

Assumption 2. For all $j \in [n]$ there exists $i \in [m]$ such that $a_{ij} \neq -\infty$.

Consider the infinite horizon game, which starts in a certain node j of Min, and assume the players are interested in the value of average payment of Min per turn:

$$\limsup_{N \rightarrow \infty} \left(\sum_{t=1}^N -a_{i_t j_t} + b_{i_t j_{t+1}} \right) / N, \quad j_1 = j. \quad (5)$$

Here, $j_1 \in [n]$, $i_1 \in [m]$, $j_2 \in [n]$, $i_2 \in [m]$, \dots is the infinite sequence of positions of the pawn. The ambition of Max is to maximize (5) while Min is seeking the opposite. In other words, we are considering a *mean payoff game* in which the payment of an infinite trajectory is given by (5). We next recall why this game has a value (meaning that the payment thought of as a function of the strategies of both players has a saddle point).

To compute this value, we define the dynamic operator $f : \mathbb{R}^n \mapsto \mathbb{R}^n$,

$$f_j(x) = \min_{k \in [m]} (-a_{kj} + \max_{l \in [n]} (b_{kl} + x_l)). \quad (6)$$

This is known as a *min-max function* [CTGG99]. Min-max functions are isotone ($x \leq y \Rightarrow f(x) \leq f(y)$) and additively homogeneous ($f(\lambda + x) = \lambda + f(x)$). Hence, they are nonexpansive in the sup-norm. Moreover, they are piecewise affine (\mathbb{R}^n can be covered by a finite number of polyhedra on which f is affine). We are interested in the following limit (*cycle-time vector*):

$$\chi(f) = \lim_{k \rightarrow \infty} f^k(x)/k. \quad (7)$$

The j th entry of the vector $\chi(f)$ can be interpreted as the limit of a game with finite horizon k with initial state j , as k tends to infinity. The existence of $\chi(f)$ follows from a theorem of Kohlberg.

Theorem 1 (Kohlberg [Koh80]). *Let $f : \mathbb{R}^n \mapsto \mathbb{R}^n$ be a nonexpansive and piecewise affine map. Then, there exist $v \in \mathbb{R}^n$ and $\chi \in \mathbb{R}^n$ such that*

$$f(v + t\chi) = v + (t + 1)\chi, \quad \forall t \geq T, \quad (8)$$

where T is a large enough real number.

The map $t \mapsto v + t\chi$ is known as an *invariant half-line*. Using the nonexpansiveness of f , one deduces that the limit (7) exists, is the same for all $x \in \mathbb{R}^n$ and is equal to the growth rate χ of any invariant half-line.

A *positional (stationary) strategy* of Min is a map $\tau : [n] \rightarrow [m]$ such that $a_{\tau(j)j}$ is finite for all $j \in [n]$ (in other words, it is a selection rule, requiring to move the pawn to node $\tau(j)$ when in node j). Similarly, a *positional (stationary) strategy* of Max is a map $\sigma : [m] \rightarrow [n]$ such that $b_{i\sigma(i)}$ is finite, for all $i \in [m]$. The corresponding dynamic operators are max-only and min-only:

$$f_j^\tau(x) = -a_{\tau(j)j} + \max_{l \in [n]} b_{\tau(j)l} + x_l, \quad f_j^\sigma(x) = \min_{i \in [m]} (-a_{ij} + b_{i\sigma(i)} + x_{\sigma(i)}). \quad (9)$$

The following result can be derived as a standard corollary of Kohlberg's theorem. Indeed, we define a positional strategy τ of Min and a positional strategy σ of Max by the condition that $f(v + t\chi) = f^\sigma(v + t\chi) = f_\tau(v + t\chi)$ for t large enough, where $t \mapsto v + t\chi$ is an invariant half-line. These strategies are easily seen to be optimal for the mean payoff game.

Theorem 2 (Coro. of [Koh80]). *For $f(x)$ given by (6), the j th coordinate of $\chi(f)$ is the value of the mean payoff game which starts in j .*

We shall use repeatedly the following duality result, proved in [GG98], as a corollary of the termination of the policy iteration algorithm of [CTGG99, GG98]. It can also be derived from Kohlberg’s theorem. (See [DG06] for a more recent presentation). A closely related result is [LL69].

Theorem 3 (Coro. of [Koh80], see also [LL69] and [GG98]). *Let $A, B \in (\mathbb{R} \cup \{-\infty\})^{m \times n}$ satisfy Assumptions 1 and 2, and let S and T be the sets of all positional strategies of Max and Min, respectively. Then,*

$$\max_{\sigma \in S} \chi(f^\sigma) = \chi(f) = \min_{\tau \in T} \chi(f^\tau) . \quad (10)$$

2.2 Tropical linear systems and mean payoff games

Max-only and min-only maps of the form (9) belong to tropical linear algebra. Max-only maps are linear in the *max-plus semiring* \mathbb{R}_{\max} , which is the set $\mathbb{R} \cup \{-\infty\}$ equipped with the operations of “addition” “ $a + b$ ” := $a \vee b = \max(a, b)$ and “multiplication” “ ab ” := $a + b$. For min-only maps, we use the *min-plus semiring* \mathbb{R}_{\min} := $\mathbb{R} \cup \{+\infty\}$ equipped with the operations of “addition” “ $a + b$ ” := $a \wedge b = \min(a, b)$ and the same “additive” multiplication. The setting in which both structures are considered simultaneously has been called minimax algebra by Cuninghame-Green [CG79]. Then, we need to allow the scalars to belong to the enlarged set $\overline{\mathbb{R}} := \mathbb{R} \cup \{-\infty\} \cup \{+\infty\}$. Note that in $\overline{\mathbb{R}}$, $(-\infty) + (+\infty) = -\infty$ if the max-plus convention is understood, and $(-\infty) + (+\infty) = +\infty$ if the min-plus convention is understood. The tropical operations are extended to matrices and vectors in the usual way. In particular,

$$(Bx)_i = \bigvee_j b_{ij} + x_j \quad (\text{max-plus}). \quad (11)$$

Here, and in the sequel, the action of matrices on vectors (by tropical multiplication) is denoted by concatenation, whereas we shall use the notation \vee and $+$ with its standard sense for scalars, the law $+$ having priority over \vee . In tropical algebra, there is no obvious subtraction. However, for any $A \in \overline{\mathbb{R}}^{n \times n}$ we can define the *Kleene star*

$$A^* := “(I - A)^{-1}” = I \vee A \vee A^2 \vee \dots \quad (\text{max-plus}), \quad (12)$$

and analogously with \wedge in the min-plus case. Here I is the max-plus identity matrix with 0 entries on the main diagonal and $-\infty$ off the diagonal. The powers are understood in the tropical (max-plus) sense. Due to the order completeness of $\overline{\mathbb{R}}$, series (12) is well-defined for all matrices. We shall need the following result.

Proposition 4 (See e.g. [BCOQ92, Th. 3.17]). *Let $A \in \overline{\mathbb{R}}^{n \times n}$ and $b \in \overline{\mathbb{R}}^n$. Then, A^*b is the least solution of $z \geq Az \vee b$.*

The maximal (minimal) circuit mean is another important object of tropical algebra. For $A = (a_{ij}) \in \mathbb{R}_{\max}^{n \times n}$ ($A = (a_{ij}) \in \mathbb{R}_{\min}^{n \times n}$), it is defined as

$$\begin{aligned} \mu^{\max}(A) &= \bigvee_{p=1}^n \bigvee_{i_1, \dots, i_p} \frac{a_{i_1 i_2} + \dots + a_{i_p i_1}}{p} \quad (\text{max-plus}), \\ \mu^{\min}(A) &= \bigwedge_{p=1}^n \bigwedge_{i_1, \dots, i_p} \frac{a_{i_1 i_2} + \dots + a_{i_p i_1}}{p} \quad (\text{min-plus}). \end{aligned} \quad (13)$$

For $A \in \overline{\mathbb{R}}^{n \times n}$ define the associated digraph $\mathcal{D}(A) = \{[n], E\}$ so that $(i, j) \in E$ whenever a_{ij} is finite. We will say that i *accesses* j if there exists a path from i to j in $\mathcal{D}(A)$. Denote by $\mu_i^{\max}(A)$ ($\mu_i^{\min}(A)$) the maximal (minimal) circuit mean of the component of $\mathcal{D}(A)$ to which

i belongs. These numbers are given by the same expressions as in (13), but with i_1, \dots, i_p restricted to that component.

Using $\mu_i^{\max}(A)$ ($\mu_i^{\min}(A)$), we can write explicit expressions for the cycle-time vector of a max-plus linear (min-plus linear) map $x \mapsto Ax$:

$$\begin{aligned}\chi_i^{\max}(A) &= \max\{\mu_j^{\max}(A), i \text{ accesses } j\} \text{ (max-plus),} \\ \chi_i^{\min}(A) &= \min\{\mu_j^{\min}(A), i \text{ accesses } j\} \text{ (min-plus).}\end{aligned}\tag{14}$$

See [CTCG⁺98] or [HOvdW05] for a proof. Note that these cycle-time vectors of max-plus and min-plus maps appear in (10). Computing A^*b is equivalent to solving a single destination shortest path problem, which can be done in $O(n^3)$ time (for instance by the Bellman-Ford algorithm). Moreover, χ^{\max} and χ^{\min} can also be computed in $O(n^3)$ time (decompose first the digraph $\mathcal{D}(A)$ in strongly connected components, and apply Karp's algorithm to compute the maximal or minimal circuit mean of each component).

Max-plus and min-plus linear maps are mutually adjoint, or *residuated*. Recall that for a max-plus linear map A from $\overline{\mathbb{R}}^n$ to $\overline{\mathbb{R}}^m$, the *residuated operator* A^\sharp from $\overline{\mathbb{R}}^m$ to $\overline{\mathbb{R}}^n$ is defined by

$$(A^\sharp y)_j := \bigwedge_{i=1}^m (-a_{ij} + y_i) ,\tag{15}$$

with the convention $(-\infty) + (+\infty) = +\infty$. Note that this operator, also known as *Cuninghame-Green inverse*, sends \mathbb{R}_{\max}^m to \mathbb{R}_{\max}^n whenever A does not have columns identically equal to $-\infty$. The term “residuated” refers to the property

$$Ax \leq y \Leftrightarrow x \leq A^\sharp y ,\tag{16}$$

where \leq is the partial order on \mathbb{R}_{\max}^m or \mathbb{R}_{\max}^n . This residuated operator is crucial for max-plus two-sided systems of inequalities, since

$$Ax \leq Bx \Leftrightarrow x \leq A^\sharp Bx .\tag{17}$$

Writing the last inequality explicitly, we have

$$x_j \leq \bigwedge_{k \in [m]} (-a_{kj} + \bigvee_{l \in [n]} (b_{kl} + x_l)) , \quad \forall j \in [n] .\tag{18}$$

Thus we obtain the same min-max function as in (6).

Moreover, positional strategies $\sigma: [m] \mapsto [n]$ and $\tau: [n] \mapsto [m]$ correspond to affine mappings B^σ and A_τ defined by

$$(A_\tau)_{ij} = \begin{cases} a_{ij} & \text{if } i = \tau(j), \\ -\infty & \text{otherwise,} \end{cases} \quad (B^\sigma)_{ij} = \begin{cases} b_{ij} & \text{if } j = \sigma(i), \\ -\infty & \text{otherwise.} \end{cases}\tag{19}$$

Recasting (10) in max(min)-plus algebra, we obtain

$$\max_{\sigma \in S} \chi(A^\sharp B^\sigma) = \chi(A^\sharp B) = \min_{\tau \in T} \chi(A_\tau B) .\tag{20}$$

The following result obtained by Akian, Gaubert and Guterman relates the solutions of $Ax \leq Bx$ and the nonnegative coordinates of $\chi(A^\sharp B)$. These coordinates correspond to *winning nodes* of the game: if the game starts in these nodes, then Max can ensure nonnegative profit with any positional strategy of Min.

Theorem 5 ([AGG09, Th. 3.2]). *Let $A, B \in \mathbb{R}_{\max}^{m \times n}$. Then, $\chi_j(A^\sharp B) \geq 0$ if and only if there exists $x \in \mathbb{R}_{\max}^n$ such that $Ax \leq Bx$ and $x_j \neq -\infty$.*

This is derived in [AGG09] from Kohlberg’s theorem. The vector x is constructed by taking an invariant half-line, $t \mapsto v + t\chi$, setting $x_i = v_i + t\chi_i$ if $\chi_i \geq 0$, and $x_i = -\infty$ otherwise, and t large enough.

Theorem 5 shows that to decide whether $Ax \leq Bx$ can be satisfied by a vector x such that $x_i \neq -\infty$, we can call a *mean-payoff oracle*. This oracle will decide whether i is a winning node of the associated mean payoff game and give a winning strategy of Max. This oracle can be implemented either by using the value iteration method, which is pseudo-polynomial [ZP96], by the approach of Puri (solving an associated discounted game for a discount factor close enough to 1 by policy iteration [Pur95]), by using the policy iteration algorithm for mean payoff games of [CTGG99, GG98, DG06] or the one of [BV07].

3 Tropical linear programming

3.1 The spectral function

The tropical linear programming problem (3) can be recast as:

$$\text{minimize } \lambda \quad \text{subject to} \quad px \vee r \leq \lambda + (qx \vee s), \quad Ax \vee c \leq Bx \vee d, \quad x \in \mathbb{R}_{\max}^n, \quad (21)$$

where $r, s \in \mathbb{R}_{\max}$, $p, q \in \mathbb{R}_{\max}^n$, $c, d \in \mathbb{R}_{\max}^m$, and $A, B \in \mathbb{R}_{\max}^{m \times n}$. Every problem concerning affine polyhedra has a “homogeneous” analogue concerning cones, which is equivalent to it. The homogeneous version of (3) reads:

$$\text{minimize } \lambda \quad \text{subject to} \quad uy \leq \lambda + vy, \quad Cy \leq Dy, \quad y_{n+1} \neq -\infty, \quad y \in \mathbb{R}_{\max}^{n+1}, \quad (22)$$

where $u, v \in \mathbb{R}_{\max}^{n+1}$ and $C, D \in \mathbb{R}_{\max}^{m \times (n+1)}$. Problem (22) is equivalent to (21) if we set $u = [p, r]$, $v = [q, s]$, $C = [A, c]$ and $D = [B, d]$. We may drop the constraint $y_{n+1} \neq -\infty$ in (22), but then it may be not equivalent to (21).

We can still reformulate (22) in a more compact way:

$$\min\{\lambda \mid Uy \leq V(\lambda)y, \quad y_{n+1} \neq -\infty \text{ is solvable}\}, \quad (23)$$

where

$$U = \begin{pmatrix} C \\ u \end{pmatrix} \quad \text{and} \quad V(\lambda) = \begin{pmatrix} D \\ \lambda + v \end{pmatrix} \quad (24)$$

have dimensions $(m+1) \times (n+1)$. Thus, we consider the original programming problem as a family of two-sided tropical linear systems where the right hand side depends on λ .

Using (16), we write

$$Uy \leq V(\lambda)y \Leftrightarrow y \leq U^\sharp V(\lambda)y, \quad (25)$$

where $U^\sharp V(\lambda)$ is a min-max function from \mathbb{R}_{\max}^n to \mathbb{R}_{\max}^n for each λ , if U does not have identically $-\infty$ columns. We next reformulate the tropical linear programming problem in terms of a spectral function, which gives the value of a parametric mean payoff game: the payments are given by the matrices U and $V(\lambda)$, and the initial state is $n+1$.

Proposition-Definition 6 (Spectral function). *The tropical linear programming problem (22) is equivalent to*

$$\min\{\lambda \mid \phi(\lambda) \geq 0\} \quad (26)$$

where the spectral function ϕ is given by $\phi(\lambda) := \chi_{n+1}(U^\sharp V(\lambda))$.

Proof. By Theorem 5, $Uy \leq V(\lambda)y$ is solvable with finite y_{n+1} if, and only if, $\chi_{n+1}(U^\sharp V(\lambda)) \geq 0$. \square

Given a strategy $\sigma \in S$ of Max and a strategy $\tau \in T$ of Min, we define the min-plus matrix $U^\sharp V^\sigma(\lambda)$ and the max-plus matrix $U_\tau^\sharp V(\lambda)$, respectively, see (19) and (20). We introduce *partial spectral functions*, $\phi^\sigma(\lambda) := \chi_{n+1}(U^\sharp V^\sigma(\lambda))$ and $\phi_\tau(\lambda) := \chi_{n+1}(U_\tau^\sharp V(\lambda))$. With this notation, (20) reads

$$\phi(\lambda) = \max_{\sigma \in S} \phi^\sigma(\lambda) = \min_{\tau \in T} \phi_\tau(\lambda). \quad (27)$$

All spectral functions introduced above are non-decreasing, due to the monotonicity of $V(\lambda)$ and of the map $f \mapsto \chi_{n+1}(f)$.

We have a finite number of strategies σ and τ , and hence a finite number of functions $\phi^\sigma(\lambda)$ and $\phi_\tau(\lambda)$ in (27). Further, $U_\tau^\sharp V(\lambda)$ is a max-plus linear operator. It is expressed by a max-plus matrix, where some entries are of the form $\alpha + \lambda$ and the other entries are constant. Using Expression (14) for the components of the cycle-time vector of a max-plus matrix, we obtain that $\phi_\tau(\lambda)$ is a piecewise-linear convex function. It is also 1-Lipschitz since the modulus of any slope does not exceed 1. It consists of no more than $O(n^2)$ linear pieces, since each possible slope cannot appear more than once. Indeed, each linear piece must be of the form $(\alpha + \beta\lambda)/k$, where $\alpha + \beta\lambda$ can be interpreted as the weight of an elementary circuit in the digraph of $U_\tau^\sharp V(\lambda)$ and k is the length of this circuit. It is easily seen that $0 \leq \beta \leq n + 1$ and $1 \leq k \leq n + 1$, which gives $O(n^2)$ possible values for the slope. Similarly, it can be shown that $\phi^\sigma(\lambda)$ is a piecewise-linear concave 1-Lipschitz function, this time with only $O(n)$ linear pieces (because now $\beta \in \{0, 1\}$). Then, (27) implies that $\phi(\lambda)$ is piecewise-linear and 1-Lipschitz.

3.2 Strategies as certificates

In the classical simplex method, the optimality of a feasible solution is certified by the sign of Lagrange multipliers. In our case, the certificate is of a different nature: it is a strategy. We shall also use such strategies to guide the next iteration of the algorithm, when the current feasible solution is not optimal.

Definition 7 (Left and right optimal strategies). *A strategy σ^- of Max (resp. τ^- of Min) is left-optimal at point $\lambda \in \mathbb{R}$ if there exists $\epsilon > 0$ such that*

$$\phi(\mu) = \phi^{\sigma^-}(\mu) = \phi_{\tau^-}(\mu) \quad \forall \mu \in [\lambda - \epsilon, \lambda].$$

Right optimal strategies σ^+ and τ^+ are defined in a similar way, by replacing $[\lambda - \epsilon, \lambda]$ by $[\lambda, \lambda + \epsilon]$.

The existence of left and right optimal strategies at each point follows readily from (27), together with the finiteness of the number of strategies and the piecewise affine character of each map $\phi^\sigma(\lambda)$ and $\phi_\tau(\lambda)$.

Let \mathcal{G}_λ be the bipartite digraph of the mean payoff game associated with the matrices U and $V(\lambda)$. Observe that in this digraph, only the weights of the arcs connecting node $m + 1$ of Max with some nodes $l \in [n + 1]$ of Min depend on λ . Given any strategy σ of Max (resp. τ of Min), denote by $\mathcal{G}_\lambda^\sigma$ (resp. \mathcal{G}_λ^τ) the bipartite subgraph of \mathcal{G}_λ obtained by deleting from \mathcal{G}_λ all the arcs (i, j) such that $i \in [m + 1]$ and $j \neq \sigma(i)$ (resp. the arcs (j, i) such that $j \in [n + 1]$ and $i \neq \tau(j)$). Then, it can be checked that circuits in $\mathcal{D}(U^\sharp V^\sigma(\lambda))$ (resp. $\mathcal{D}(U_\tau^\sharp V(\lambda))$) correspond to circuits in $\mathcal{G}_\lambda^\sigma$ (resp. \mathcal{G}_λ^τ).

A left-optimal strategy τ of Min at point λ^* is such that $\phi(\lambda) = \phi_\tau(\lambda)$ for all $\lambda \leq \lambda^*$ close enough to λ^* . Hence, by (26), we can certify that λ^* is the optimal value of the tropical linear

program by checking that $\phi_\tau(\lambda^*) = 0$ and that the left derivative of the function ϕ_τ at point λ^* is positive. This can be reformulated in purely combinatorial terms, leading to the following certificate.

Theorem 8. *The tropical linear programming problem (22) has the optimal value $\lambda^* \in \mathbb{R}$ if, and only if, $\phi(\lambda^*) \geq 0$ and there exists a strategy τ of Min such that the digraph $\mathcal{G}_{\lambda^*}^\tau$ satisfies the following conditions: (i) all circuits accessible from node $n + 1$ of Min have nonpositive weight, (ii) there is a circuit accessible from node $n + 1$ of Min with zero weight, (iii) each circuit of zero weight accessible from node $n + 1$ of Min passes through node $m + 1$ of Max. Moreover, these conditions are always satisfied when τ is left-optimal at point λ^* .*

The proof is inspired by the one of the tropical Farkas lemma in [AGK10b]. Pictures illustrating the game can be found there. In the same way, we can certify when the tropical linear programming problem (22) is unbounded.

Theorem 9. *The tropical linear programming problem (22) is unbounded if, and only if, there exists a strategy σ of Max such that all circuits in the digraph \mathcal{G}_0^σ accessed from node $n + 1$ of Min do not contain node $m + 1$ of Max and have nonnegative weight.*

Remark 1. If the strategies σ or τ and the scalar λ^* are fixed (considered as inputs) the conditions in Theorems 8 and 9, i.e., the validity of the certificates, can be checked in polynomial time. We leave details to the reader.

3.3 Newton iterations

In (26), we need to find the least λ such that $\phi(\lambda) \geq 0$, where $\phi(\lambda)$ is non-decreasing and Lipschitz continuous. Thus, we can consider certain classical methods for finding zeroes of a “good enough” function of one variable. In particular, the bisection method for $\phi(\lambda)$ corresponds to the approach of Butkovič and Aminu [BA08]. For this method it is not important to know the actual value of $\phi(\lambda)$, but just whether $\phi(\lambda) \geq 0$, i.e., whether $Uy \leq V(\lambda)y$ is solvable with $y_{n+1} \neq -\infty$. Here, we adopt a different approach: the concept of (left-, right-) optimal strategy yields an analogue of (left-, right-) derivative, and leads to an analogue of Newton iteration.

Algorithm 1 (Positive Newton iteration). **Start.** A point λ_0 at which $\phi(\lambda_0) \geq 0$.

Iteration k . Find a left-optimal strategy $\sigma \in S$ at λ_{k-1} and compute $\lambda_k = \min\{\lambda : \phi^\sigma(\lambda) \geq 0\}$.

Stop. Verify $\lambda_k = \lambda_{k-1}$ or $\lambda_k = -\infty$.

Proposition 10. *The number of iterations of Algorithm 1 does not exceed the number $|S|$ of strategies of Player Max. \square*

We next explain how each step of the algorithm can be implemented.

3.4 Newton iterations by means of Kleene star

We shall see that each Newton iteration in Algorithm 1 reduces to solving a shortest path problem. We start with a technical observation.

Lemma 11. *The minimization problems*

$$\min\{\lambda \mid Uy \leq V(\lambda)y, y_1 \neq -\infty \text{ is solvable}\} \quad (28)$$

$$\min\{\lambda \mid Uy \leq V(\lambda)y, y_2 \neq -\infty \text{ is solvable}\} \quad (29)$$

have equal solutions if the following conditions are satisfied: (i) for any solution y of $Uy \leq V(\lambda)y$, $y_1 \neq -\infty$ implies $y_2 \neq -\infty$, (ii) the variable y_2 appears (with a real coefficient) on the r.h.s. of the last inequality of $Uy \leq V(\lambda)y$ (i.e., in $uy \leq \lambda + vy$), (iii) there exists a solution \tilde{y} of the first m inequalities of $Uy \leq V(\lambda)y$ such that $\tilde{y}_1 \neq -\infty$.

If we set $V^\sigma(\lambda)$ instead of $V(\lambda)$ in $Uy \leq V(\lambda)y$, the least possible value of λ is exactly the minimal zero of $\phi^\sigma(\lambda)$. The main idea is to compute this minimal zero by considering the system $Uy \leq V^\sigma(\lambda)y$ directly. Denote $l := \sigma(m+1)$. Then the last equation of $Uy \leq V^\sigma(\lambda)y$ reads

$$\bigvee_{i=1}^{n+1} u_i + y_i \leq \lambda + v_l + y_l, \quad (30)$$

and we require $y_{n+1} \neq -\infty$ (instead of $y_l \neq -\infty$). Since we have $\phi^\sigma(\lambda) \geq 0$, we know by Theorem 5 that the rest of the inequalities are satisfied by certain $z \in \mathbb{R}_{\max}^{n+1}$ with $z_{n+1} \neq -\infty$.

Assume now that for each λ and each solution y of $Uy \leq V^\sigma(\lambda)y$, $y_{n+1} \neq -\infty$ implies $y_l \neq -\infty$. This case will be referred to as non-degenerate in the sequel (we shall see that the degenerate case is not more difficult). Then, the conditions of Lemma 11 are satisfied with $n+1$ for 1, l for 2 and z for \tilde{y} , and we can switch to the formulation with $y_l \neq -\infty$ and further, set $y_l = 0$. Hence, our problem can be rewritten as

$$\text{minimize } px \vee r \quad \text{subject to} \quad Ax \vee c \leq B^\sigma x \vee d^\sigma, \quad x \in \mathbb{R}_{\max}^n, \quad (31)$$

where $p \in \mathbb{R}_{\max}^n$, $r \in \mathbb{R}_{\max}$, $c, d^\sigma \in \mathbb{R}_{\max}^m$ and $A, B^\sigma \in \mathbb{R}_{\max}^{m \times n}$ are such that

$$U = \begin{pmatrix} A & c \\ p + v_l & r + v_l \end{pmatrix}, \quad V^\sigma(\lambda) = \begin{pmatrix} B^\sigma & d^\sigma \\ -\infty & \lambda + v_l \end{pmatrix}. \quad (32)$$

Here, column l is in the place of column $n+1$ just for the simplicity of the presentation. First note that there are two subsystems:

$$\begin{aligned} (Ax)_i \vee c_i &\leq b_{i\sigma(i)} + x_{\sigma(i)}, & i \in I, \\ (Ax)_i \vee c_i &\leq d_i, & i \notin I, \end{aligned} \quad (33)$$

where $I \subseteq [m]$. The second subsystem can be dispensed with. Indeed, let x be a solution of both subsystems. If the first subsystem has the least solution \underline{x} , then $\underline{x} \leq x$ and it is also a solution to the second subsystem. Computing \underline{x} is easily seen to reduce to finding the least solution \underline{t} of a system $t \geq Ft \vee h$, where the vector t collects the variables x_j arising at the right-hand side of the first system in (33). By Proposition 4, $\underline{t} = F^*h$, from which we determine \underline{x} . The following proposition summarizes the previous analysis of the nondegenerate case.

Proposition 12. *Checking the existence of a feasible solution $x \in \mathbb{R}_{\max}^n$ of the auxiliary problem (31), and computing the least feasible solution \underline{x} , reduces to solving a shortest path problem (including the detection of circuits of negative weights), which can be done in $O(mn) + O(n^3)$ time. The value of the problem is $p\underline{x} \vee r$ (or $+\infty$ if there is no feasible solution). \square*

Now we consider the degenerate case where for $Uy \leq V^\sigma(\lambda)y$, there may be solutions with $y_{n+1} \neq -\infty$ but $y_l = -\infty$. This may happen only if: (i) $n+1 \notin \text{supp}(u)$, and (ii) there exists a solution of the first m inequalities with $y_i = -\infty$ for all $i \in \text{supp}(u)$, but $y_{n+1} \neq -\infty$.

Conditions (i) and (ii) imply that the ‘‘ratio’’ $uy - vy$ with $y_{n+1} \neq -\infty$ is unbounded from below.

To check Condition (ii), we set $y_i = -\infty$ for $i \in \text{supp}(u)$, and $y_{n+1} = 0$. We arrive at two subsystems like in (33), where the free terms correspond to the last coordinate $n+1$. Next,

we find the least solution of the first subsystem as described above, and we check whether it belongs to \mathbb{R}_{\max}^n and satisfies the second subsystem. The overall complexity of this step does not exceed $O(mn) + O(n^3)$.

At each step of Algorithm 1 we need to verify conditions (i) and (ii) above, and if these conditions do not hold, we solve the auxiliary problem (31) by Proposition 12.

We now summarize the worst-case complexity of solving tropical linear programming problems by this method. Denote by $\text{MPG}(m, n, M)$ the worst-case complexity of an oracle which yields a left-optimal strategy at a point λ . Here the constant M measures the greatest absolute size of entries of U and $V(\lambda)$, over all λ computed by Newton iterations. To eliminate this dependence on λ , we notice that all values λ_k , except for the initial value λ_1 , are found by means of computations with Kleene star and are sums of no more than n entries of $V(0)$ and n entries of U (with minus sign). For the choice of λ_1 see Subsection 3.5, it depends on the choice of a vector x that satisfies $Ax \vee b \leq Cx \vee d$. We can assume that this vector can be chosen concisely, in terms of A, b, C and d . Summarizing all worst-time complexities, we obtain the following result.

Proposition 13. *Let $A, C \in \mathbb{R}_{\max}^{m \times n}$, $b, d \in \mathbb{R}_{\max}^m$, $p, q \in \mathbb{R}_{\max}^n$, $r, s \in \mathbb{R}_{\max}$, and $|S|$ be the number of strategies of Player Max. Algorithm 1 solves a tropical linear programming problem (3) in no more than $(\text{MPG}(m, n, M) + O(mn) + O(n^3)) \times |S|$ operations.*

3.5 Remaining ingredients of the positive Newton iteration

Computing a left-optimal strategy This is not more difficult than solving a mean payoff game problem. Actually, assume that the entries $a_{ij}, b_{ij}, c_j, d_j, p_j, q_j, r, s$ are either infinite or belong to \mathbb{Z} , and let M denote the maximal absolute value of these finite entries. Then, every linear piece of the spectral function $\phi^\sigma(\lambda)$ must be of the form $(\alpha + \beta\lambda)/k$, and it can be checked that $1 \leq k \leq n + 1$, $|\alpha| \leq 2M(n + 1)$, and $\beta \in \{0, 1\}$. It follows that every nondifferentiability point of ϕ^σ , which is the intersection of two linear pieces, is a rational with numerator $O(Mn)$ and denominator $O(n)$. Given $\lambda^* \in \mathbb{R}$, it suffices to evaluate the greatest rational $\bar{\lambda}$ of this form smaller than λ^* , and to evaluate $\phi(\mu)$, for any $\mu \in]\bar{\lambda}, \lambda^*[$, to get a left-optimal strategy. Indeed, computing $\phi(\mu)$ is the same as solving a non-parametric mean payoff game, and any strategy of Max that is optimal in this non-parametric game is left-optimal at point λ^* .

Phase I: Computing an initial λ_0 First, we shall always assume that the system $Cy \leq Dy$, $y_{n+1} > -\infty$ is feasible, which can be checked by computing $\chi_{n+1}(C^\sharp D)$ (i.e. solving a mean payoff game). Otherwise, the tropical linear programming problem (22) is infeasible and its value is $+\infty$.

We shall need the following elementary result, which follows readily from Theorem 5.

Proposition 14. *There is a feasible solution x of $Cy \leq Dy$ such that vx and x_{n+1} are finite if, and only if, $\chi_i(C^\sharp D) \geq 0$ for $i = n + 1$ and for some i such that v_i is finite.*

Thus, such a vector x can be obtained (or shown not to exist) by computing the vector $\chi(C^\sharp D)$, which is the same as computing the value, for all initial states i , of the mean payoff game whose rewards are given by the matrices C and D .

If such an x exists, since vx is finite, we can find a finite λ such that $ux \leq \lambda + vx$, and since $Cx \leq Dx$, we deduce that $\phi(\lambda) \geq 0$, which gives an initial λ_0 for Algorithm 1.

If there is no such x , we know that vx is $-\infty$ for all x satisfying $Cx \leq Dx$ and x_{n+1} finite. Now we need to determine whether there exists x such that $Cx \leq Dx$, x_{n+1} finite and $x_i = -\infty$ for all $i \in \text{supp } u$, because if this holds the tropical linear programming problem (22) is clearly




Figure 1: Minimization example: the tropical polyhedron and the spectral function

unbounded. With this aim, we eliminate the variables x_i , for $i \in \text{supp } u$, from the system $Cx \leq Dx$ and consider the resulting reduced system (in fewer variables). Here we assume that $n+1 \notin \text{supp } u$ because otherwise our problem is infeasible. As above, by solving a mean payoff game we can determine whether this reduced system, together with the condition that x_{n+1} is finite, is feasible. If it is, as we already said our original problem is unbounded. However, if it is not, then any solution of the system $Cx \leq Dx$ with x_{n+1} finite satisfies that vx is $-\infty$ and that x_i is finite for some $i \in \text{supp } u$. Therefore, in the latter case we conclude that the tropical linear programming problem (22) is infeasible and its value is $+\infty$.

4 Examples

Minimization Let us minimize the max-linear combination $\max(2 + x_1, -4 + x_2)$ over the tropical polyhedron in \mathbb{R}_{\max}^2 defined by the system of inequalities $Ax \vee c \leq Bx \vee d$, where

$$(A \mid c) = \left(\begin{array}{cc|c} -\infty & -\infty & 0 \\ -\infty & -\infty & 0 \\ -\infty & -\infty & 0 \\ -\infty & -3 & 0 \\ -\infty & -4 & -\infty \\ -\infty & -5 & -\infty \\ -\infty & -6 & -\infty \end{array} \right), \quad (B \mid d) = \left(\begin{array}{cc|c} -2 & 0 & -\infty \\ 0 & -1 & -\infty \\ 1 & -2 & -\infty \\ 2 & -\infty & -\infty \\ 0 & -\infty & 0 \\ -2 & -\infty & 0 \\ -4 & -\infty & 0 \end{array} \right). \quad (34)$$

This polyhedron is displayed on the left hand side of Figure 1. The direction of minimization of $\max(2 + x_1, -4 + x_2)$ is shown there by a dotted line above the polyhedron, together with the optimal hyperplane $\max(2 + x_1, -4 + x_2) = 0$.

We start the algorithm with $\lambda_1 = 15$, where $\phi(\lambda_1) = 5.5$. The map $\sigma_1(1) = 1, \sigma_1(2) = 1, \dots, \sigma_1(7) = 1$ and $\sigma_1(8) = 3$ is an optimal strategy of Max. To perform the first step of the Newton iteration, we find the minimal solution of the system

$$0 \leq x_1 - 2, \quad 0 \leq x_1, \quad 0 \leq 1 + x_1, \quad \max(x_2 - 3, 0) \leq x_1 + 2, \quad x_2 - 4 \leq x_1, \quad x_2 - 3 \leq x_1, \quad x_2 - 2 \leq x_1,$$

which is $(x_1, x_2) = (2, -\infty)$. The next value is $\lambda_2 = x_1 + 2 = 4$. Then, $\phi(\lambda_2) = 1.5$ and $\sigma_2(1) = 2, \sigma_2(2) = 2, \sigma_2(3) = 1, \sigma_2(4) = 1, \sigma_2(5) = 1, \sigma_2(6) = 3, \sigma_2(7) = 3$ and $\sigma_2(8) = 3$ is a




Figure 2: Maximization example: the tropical polyhedron and the spectral function

new optimal strategy of Max. For the next Newton iteration, we find the minimal solution of the system

$$0 \leq x_2, 0 \leq x_2 - 1, 0 \leq 1 + x_1, \max(x_2 - 3, 0) \leq x_1 + 2, x_2 - 4 \leq x_1, x_2 - 5 \leq 0, x_2 - 6 \leq 0,$$

which is $(x_1, x_2) = (-1, 1)$. The next value is $\lambda_3 = x_1 + 2 = 1$. Now $\phi(\lambda_3) = 0.5$ and $\sigma_3(1) = 2, \sigma_3(2) = 2, \sigma_3(3) = 2, \sigma_3(4) = 1, \sigma_3(5) = 3, \sigma_3(6) = 3, \sigma_3(7) = 3$ and $\sigma_3(8) = 3$ is the optimal strategy of Max. For the next Newton iteration, we find the minimal solution of the system

$$0 \leq x_2, 0 \leq x_2 - 1, 0 \leq x_2 - 2, \max(x_2 - 3, 0) \leq x_1 + 2, x_2 - 4 \leq 0, x_2 - 5 \leq 0, x_2 - 6 \leq 0,$$

which is $(x_1, x_2) = (-2, 2)$. This gives the optimal value $\lambda_4 = 0$. Vectors $(2, -\infty), (-1, 1)$ and $(-2, 2)$ found by the Newton iterations are indicated on the left hand side of Figure 1 as “1”, “2” and “3”.

The right hand side of Figure 1 displays the graph of $\phi(\lambda)$, together with the Newton iterations. The graphs of partial spectral functions $\phi^\sigma(\lambda)$ are given by red dashed lines.

Maximization We maximize $\max(1 + x_1, 3 + x_2)$ over the subset of \mathbb{R}_{\max}^3 defined by the system $Ax \vee c \leq Bx \vee d$, where

$$(A \mid c) = \left(\begin{array}{cc|c} -\infty & -1 & -\infty \\ -2 & -2 & -\infty \\ -1 & -\infty & -\infty \\ 0 & -\infty & -\infty \end{array} \right), \quad (B \mid d) = \left(\begin{array}{cc|c} 0 & -\infty & 0 \\ -\infty & -\infty & 0 \\ -\infty & 0 & 0 \\ -\infty & 2 & 0 \end{array} \right). \quad (35)$$

This maximization problem is equivalent to finding a minimal λ such that $0 \leq \lambda + \max(1 + x_1, 3 + x_2)$ is consistent with $Ax \vee c \leq Bx \vee d$. Indeed, the maximum of $\max(1 + x_1, 3 + x_2)$ is given by the opposite (tropical inverse) of that minimal λ .

Let us take $\lambda_1 = 3$ for the first time, then we obtain that $\sigma_1(1) = 3, \sigma_1(2) = 3, \sigma_1(3) = 3, \sigma_1(4) = 3$ and $\sigma_1(5) = 2$ is an optimal strategy of Max. To perform the Newton iteration we first notice that $\sigma_1(5) = 2$, which means that the second columns of $(A \mid c)$ and $(B^\sigma \mid d^\sigma)$ will correspond to free terms in (31). We have to find the minimal solution of the following system:

$$y_3 \geq -1, y_3 \geq \max(-2 + y_1, -2), y_3 \geq -1 + y_1, y_3 \geq y_1,$$

which is $(y_1, y_3) = (-\infty, -1)$. The full vector $y = [-\infty; 0; -1]$ is a translate of $y + 1 = [x_1; x_2; 0]$, where $(x_1, x_2) = (-\infty, 1)$ is marked as “1” at the left of Figure 2. Meanwhile we obtain $\lambda_2 = y_3 - 3 = -4$, and $\sigma_2(1) = 1$, $\sigma_2(2) = 3$, $\sigma_2(3) = 2$, $\sigma_2(4) = 2$ and $\sigma_2(5) = 2$ is now a new optimal strategy of Max. Here again $\sigma(5) = 2$ and the second columns of $(A \mid c)$ and $(B^\sigma \mid d^\sigma)$ are again free coefficients. We have to find the minimal solution of the following system:

$$y_1 \geq -1, y_3 \geq \max(-2 + y_1, -2), 0 \geq -1 + y_1, 2 \geq y_1,$$

which is $(y_1, y_3) = (-1, -2)$. We obtain $\lambda_3 = y_3 - 3 = -5$, which is the optimal value. The full vector $y = [-1; 0; -2]$ is a translate of $y + 2 = [x_1; x_2; 0]$, where $(x_1, x_2) = (1, 2)$ is marked as “2” at the left of Figure 2.

As in the case of Figure 1, the right hand side of Figure 2 displays the graph of $\phi(\lambda)$, together with the Newton iterations. The graphs of partial spectral functions $\phi^\sigma(\lambda)$ are given by red dashed lines.

References

- [AGG08] X. Allamigeon, S. Gaubert, and É. Goubault. Inferring min and max invariants using max-plus polyhedra. In *Proceedings of the 15th International Static Analysis Symposium (SAS’08)*, volume 5079 of *Lecture Notes in Comput. Sci.*, pages 189–204. Springer, Valencia, Spain, 2008.
- [AGG09] M. Akian, S. Gaubert, and A. Guterman. Tropical polyhedra are equivalent to mean payoff games. Eprint arXiv:0912.2462, 2009.
- [AGG10a] A. Adj, S. Gaubert, and E. Goubault. Coupling policy iteration with semi-definite relaxation to compute accurate numerical invariants in static analysis. In A. D. Gordon, editor, *Programming Languages and Systems, 19th European Symposium on Programming, ESOP 2010*, number 6012 in LNCS, pages 23–42. Springer, 2010.
- [AGG10b] X. Allamigeon, S. Gaubert, and É. Goubault. The tropical double description method. In *Proceedings of the 27th International Symposium on Theoretical Aspects of Computer Science (STACS’2010, March 4-6, Nancy, France)*. Leibniz Center in Informatics, 2010.
- [AGK10a] X. Allamigeon, S. Gaubert, and R. D. Katz. The number of extreme points of tropical polyhedra. *J. Comb. Theory Series A*, 2010. Published on line, doi:10.1016/j.jcta.2010.04.003.
- [AGK10b] X. Allamigeon, S. Gaubert, and R. D. Katz. Tropical polars, hypergraph transversals, and mean payoff games. Eprint arXiv:1004.2778, 2010.
- [All09] X. Allamigeon. *Static analysis of memory manipulations by abstract interpretation — Algorithmics of tropical polyhedra, and application to abstract interpretation*. PhD thesis, École Polytechnique, Palaiseau, France, November 2009. <http://www.lix.polytechnique.fr/Labo/Xavier.Allamigeon/papers/thesis.pdf>.
- [BA08] P. Butkovič and A. Aminu. Introduction to max-linear programming. *IMA Journal of Management Mathematics*, 2008.
- [BCOQ92] F. L. Baccelli, G. Cohen, G.-J. Olsder, and J.-P. Quadrat. *Synchronization and Linearity: an Algebra for Discrete Event Systems*. Wiley, 1992.
- [BH04] W. Briec and C. Horvath. \mathbb{B} -convexity. *Optimization*, 53:103–127, 2004.
- [BNgC08] M. Bezem, R. Nieuwenhuis, and E. Rodríguez Carbonell. The max-atom problem and its relevance. In *Proceedings of the 15th International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR’08)*, volume 5330 of LNCS, Doha (Qatar), November 2008. Springer.
- [BNgC10] M. Bezem, R. Nieuwenhuis, and E. Rodríguez Carbonell. Hard problems in max-algebra, control theory, hypergraphs and other areas. *Information processing letters*, 110:113–138, 2010.

- [BV07] H. Bjorklund and S. Vorobyov. A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. *Discrete Appl. Math.*, 155:210–229, 2007.
- [CG79] R. A. Cuninghame-Green. *Minimax Algebra*, volume 166 of *Lecture Notes in Economics and Mathematical Systems*. Springer, Berlin, 1979.
- [CGQ04] G. Cohen, S. Gaubert, and J. P. Quadrat. Duality and separation theorems in idempotent semimodules. *Linear Algebra Appl.*, 379:395–422, 2004.
- [CTCG⁺98] J. Cochet-Terrasson, G. Cohen, S. Gaubert, M. M. Gettrick, and J. P. Quadrat. Numerical computation of spectral elements in max-plus algebra. In *Proceedings of the IFAC conference on systems structure and control*, pages 699–706, IRCT, Nantes, France, 1998.
- [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.
- [DG06] V. Dhingra and S. Gaubert. How to solve large scale deterministic games with mean payoff by policy iteration. In *Proceedings of the 1st international conference on Performance evaluation methodologies and tools (VALUETOOLS)*, volume 180, Pisa, Italy, 2006. article No. 12.
- [DS04] M. Develin and B. Sturmfels. Tropical convexity. *Documenta Math.*, 9:1–27, 2004.
- [Fri09] O. Friedmann. An exponential lower bound for the parity game strategy improvement algorithm as we know it. In *Proceedings of the Twenty-Fourth Annual IEEE Symposium on Logic in Computer Science (LICS 2009)*, pages 145–156. IEEE Computer Society Press, August 2009.
- [GG98] S. Gaubert and J. Gunawardena. The duality theorem for min-max functions. *C. R. Acad. Sci. Paris.*, 326, Série I:43–48, 1998.
- [GK09] S. Gaubert and R. D. Katz. The tropical analogue of polar cones. *Linear Algebra Appl.*, 431(5-7):608–625, 2009. E-print arXiv:0805.3688.
- [GP97] S. Gaubert and M. Plus. Methods and applications of $(\max,+)$ linear algebra. In R. Reischuk and M. Morvan, editors, *Proceedings of the 14th Annual Symposium on Theoretical Aspects of Computer Science (STACS'97)*, number 1200 in *Lecture Notes in Comput. Sci.*, Lübeck, March 1997. Springer.
- [HOvdW05] B. Heidergott, G.-J. Olsder, and J. van der Woude. *Max-plus at Work*. Princeton Univ. Press, 2005.
- [Jos05] M. Joswig. Tropical halfspaces. In *Combinatorial and computational geometry*, volume 52 of *Math. Sci. Res. Inst. Publ.*, pages 409–431. Cambridge Univ. Press, Cambridge, 2005.
- [Kat07] R. D. Katz. Max-plus (A, B) -invariant spaces and control of timed discrete event systems. *IEEE Trans. Aut. Control*, 52(2):229–241, 2007.
- [Koh80] E. Kohlberg. Invariant half-lines of nonexpansive piecewise-linear transformations. *Math. Oper. Res.*, 5(3):366–372, 1980.
- [LL69] T. M. Liggett and S. A. Lippman. Stochastic games with perfect information and time average payoff. *SIAM Rev.*, 11:604–607, 1969.
- [LMS01] G. L. Litvinov, V. P. Maslov, and G. B. Shpiz. Idempotent functional analysis: An algebraic approach. *Math. Notes (Moscow)*, 69(5):758–797, 2001.
- [Min04] A. Miné. *Weakly Relational Numerical Abstract Domains*. PhD thesis, École Polytechnique, Palaiseau, France, December 2004. <http://www.di.ens.fr/~mine/these/these-color.pdf>.
- [MSS04] R. H. Möhring, M. Skutella, and F. Stork. Scheduling with AND/OR precedence constraints. *SIAM J. Comput.*, 33(2):393–415 (electronic), 2004.
- [Pur95] A. Puri. *Theory of Hybrid Systems and Discrete Event Systems*. PhD thesis, University of Berkeley, 1995.
- [SCSM06] S. Sankaranarayanan, M. Colon, H. Sipma, and Z. Manna. Efficient strongly relational polyhedral analysis. In *Verification, Model Checking, and Abstract Interpretation: 7th*

- International Conference, (VMCAI)*, volume 3855 of *LNCS*, pages 111–125, Charleston, SC, January 2006. Springer.
- [Zim77] K. Zimmermann. A general separation theorem in extremal algebras. *Ekonom.-Mat. Obzor (Prague)*, 13:179–201, 1977.
- [ZP96] U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158(1-2):343–359, 1996.