Tropical polyhedra are equivalent to mean payoff games

Marianne Akian

(INRIA Saclay - Île-de-France and CMAP, École Polytechnique)

joint work with Stéphane Gaubert (INRIA Saclay and CMAP) and Alexander Guterman (Moscow State Univ.), see arXiv:0912.2462

EPSRC Symposium Workshop on Game theory for finance, social and biological sciences (GAM)
14-17 April 2010

Max-plus or tropical algebra (semiring)

$$
\begin{array}{cccccccccccc}
\mathbb{R}_{\max} := (& \mathbb{R} \cup \{-\infty\} & , & \max & , & + & , & -\infty & , & 0 &) \\
 & & & \oplus & & \otimes & & \mathbb{0} & & \mathbb{1} & \\
 & & & \vee & & + & & & & & \\
 & & & \text{“}+\text{”} & & \text{concatenation} & & & & &
\end{array}
$$

- $2 \oplus 3 = 3$, $2 \otimes 3 = 5$.
- $a \oplus b = a \vee b =$“$a + b$”;
- $a \otimes b = a + b =$“ab”.
- $\mathbb{R}_{\max}$ is idempotent: $a \oplus a = a$.
- Hence there are no opposites,
- The natural order ($a \leq b$ if $a \oplus b = b$) is the usual order and all numbers are ≥ 0.

- A max-plus linear operator $A : \mathbb{R}^n_{\max} \to \mathbb{R}^m_{\max}$ can be represented by a matrix $A \in \mathbb{R}^{m\times n}_{\max}$:

$$(Ax)_i = \max_{j\in[n]}(A_{ij} + x_j), \quad i \in [n] := \{1, \dots n\} \ .$$

- Several ways to define a hypersurface:
  - with two-sided equations:

$$S = \{x \in \mathbb{R}^n_{\max} \mid \max_{j\in[n]}(a_j + x_j) = \max_{j\in[n]}(b_j + x_j)\}$$

  - with "one-sided" equations, as in tropical geometry:

$$S = \{x \in \mathbb{R}^n_{\max} \mid \text{the max in } \max_{j\in[n]}(a_j + x_j) \text{ is attained at least twice}\} \ ,$$

denoted "$\sum_j a_j x_j = \mathbf{0}$" or "$\max_j(a_j + x_j) = \mathbf{0}$".

Example

- The tropical line "$x + y + \mathbf{1} = \mathbf{0}$" is the set of points where $\max(x, y, 0)$ is attained at least twice:

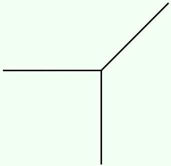

- this is the limit of the amoeba:
$\lim_{t \to 0^+} \{ -\frac{1}{\log t} (\log(|x|, \log|y|);\ ax + by + c = 0 \}$ where $a, b, c \in \mathbb{C}$.

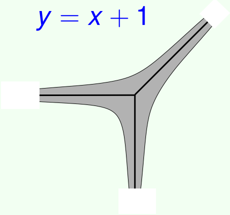

See Gelfand, Kapranov & Zelevinsky, Passare ...

Tropical segments:

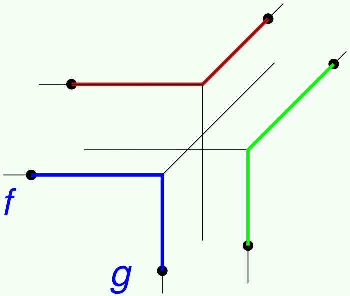

$[f, g] := \{(\lambda + f) \vee (\mu + g) \mid \lambda, \mu \in \mathbb{R}_{\max},\ \lambda \vee \mu = 0\}$.

$C \subset \mathbb{R}^n_{\max}$ is a tropical convex set if $f, g \in C \implies [f, g] \in C$

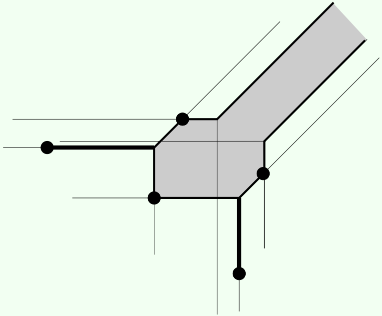

Tropical convex cones $\Leftrightarrow$ subsemimodules over $\mathbb{R}^n_{\max}$.

A *tropical half-space* is a set of the form

$$H = \{x \in \mathbb{R}^n_{\max} \mid \max_j (a_j + x_j) \leq \max_j (b_j + x_j)\}$$

It is also the union of sectors (usual convex sets) delimited by the tropical hyperplane: "$\max_j(c_j + x_j) = \mathbf{0}$" with $c_j = a_j \vee b_j$.

From the *separation theorem*, we have:

Theorem

Every closed tropical convex cone of $\mathbb{R}^n_{\max}$ is the intersection of tropical half-spaces:

$$C = \{x \in \mathbb{R}^n_{\max} \mid Ax \leq Bx\}$$

with $A, B \in \mathbb{R}^{I \times [n]}_{\max}$, and I possibly infinite.

See for instance [Zimmermann 77], [Cohen, Gaubert, Quadrat 01 and LAA04].

Tropical polyhedral cones are defined as the intersection of finitely many tropical half-spaces ($I = [m]$), or equivalently, the convex hull of finitely many rays.
See the works of [Gaubert, Katz, Butkovič, Sergeev, Schneider,...].
See also the tropical geometry point of view [Sturmfels, Develin, Joswig, Yu,...].

Tropical convex cones and games

- $Ax \leq Bx \Leftrightarrow x \leq f(x)$ with $f(x) = A^\sharp Bx$:

$$(f(x))_j = \inf_{i \in I}(-A_{ij} + \max_{k \in [n]}(B_{ik} + x_k)) \enspace .$$

- f is the *dynamic programming operator* of a zero-sum two player deterministic game: the states and actions are in I and $[n]$, Min plays in states $j \in [n]$, choose a state $i \in I$ and receive A_{ij} from Max, Max plays in states $i \in I$, choose a state $j \in [n]$ and receive B_{ij} from Min.
The vector of values v_j^N of the game after N turns (Min + Max) starting in state j satisfies:

$$v^N = f(v^{N-1}),\ v^0 = 0 \enspace .$$

- f is a min-max function [Olsder 91] when I is finite, and $f : \mathbb{R}^n \to \mathbb{R}^n$ when the columns of A and the rows of B are not $\equiv -\infty$.
- f is order preserving ($x \leq y \Rightarrow f(x) \leq f(y)$) and additively homogeneous ($f(\lambda + x) = \lambda + f(x)$).

Tropical convex cones and games

- Every order preserving and additively homogeneous map $g : \mathbb{R}^n \to \mathbb{R}^n$ can be written as the dynamic programming operator of a zero-sum two player deterministic game (with infinite action space I):
$$[g(x)]_j = \inf_{i \in I} \max_{k \in [n]} (r_{jik} + x_k)$$
(take $I = \mathbb{R}^n$ and $r_{jyk} = g(y)_j - y_k$) [Kolokoltsov; Gunawardena, Sparrow; Rubinov, Singer].
- Every dynamic programming operator g as above can be written as $g(x) = A^\sharp Bx$ for some (infinite) matrices $A, B \in \mathbb{R}_{\max}^{I' \times [n]}$ (take $I' = I \times [n]$, $A_{(i,\ell),j} = \delta_{\ell,j}$, $B_{(i,\ell),j} = r_{\ell,i,j}$).
- Thus $C := \{x \in (\mathbb{R} \cup \{-\infty\})^n \mid x \leq g(x)\}$ is a tropical convex cone.

Corollary

Every dynamic programming operator of a deterministic game (resp. every order preserving additively homogeneous map) yields an external representation of a closed tropical convex cone, and vice versa. In this correspondence, games with finite action spaces, or equivalently min-max functions, are mapped to polyhedral cones.

Perron-Frobenius tools for order preserving homogeneous maps

$\exp : x \mapsto (\exp(x_j))_{j\in[n]}$ maps $\mathbb{R}^n_{\max}$ to the positive cone $\mathbb{R}^n_+$ of $\mathbb{R}^n$, and send order preserving additively homogeneous maps of $(\mathbb{R} \cup \{-\infty\})^n$ into order preserving positively homogenous maps of $\mathbb{R}^n_+$.

Spectral radius, Collatz-Wielandt number, and dual CW number:

$$\begin{aligned}
\rho(f) &:= \max\{\lambda \in \mathbb{R}_{\max} \mid \exists u \in \mathbb{R}^n_{\max} \setminus \{-\infty\},\ f(u) = \lambda + u\} \ , \\
\mathrm{cw}(f) &:= \inf\{\mu \in \mathbb{R} \mid \exists w \in \mathbb{R}^n, f(w) \leq \mu + w\} \ , \\
\mathrm{cw}'(f) &:= \sup\{\lambda \in \mathbb{R}_{\max} \mid \exists u \in \mathbb{R}^n_{\max} \setminus \{-\infty\},\ f(u) \geq \lambda + u\} \ .
\end{aligned}$$

Theorem (see [Nussbaum, LAA 86] for general cones of $\mathbb{R}^n$)

Let f be a continuous, order preserving and additively homogeneous self-map of $(\mathbb{R} \cup \{-\infty\})^n$, then

$$\rho(f) = \mathrm{cw}(f) \ .$$

Proposition

The following limit exists and is independent of the choice of x:

$$\bar{\chi}(f) := \lim_{N\to\infty} \max_{j\in[n]} f_j^N(x)/N \enspace ,$$

and we have:

$$\mathrm{cw}'(f) = \rho(f) = \mathrm{cw}(f) = \bar{\chi}(f) \enspace .$$

Moreover, there is at least one coordinate $j \in [n]$ such that $\chi_j(f) := \lim_{N\to\infty} f_j^N(x)/N$ exists and is equal to $\bar{\chi}(f)$.

See [Vincent 97, Gunawardena, Keane 95, Gaubert, Gunawardena 04] for the existence of $\bar{\chi}$ when f preserves $\mathbb{R}^n$.

$\chi_j(f)$ is the *mean payoff* of the game starting in state j.

When f is a min-max function which preserves $\mathbb{R}^n$, this can be shown also using Kohlberg's theorem (80) on the existence of invariant half-lines $f(u + t\eta) = u + (t+1)\eta$ for t large. Then $\chi_j(f)$ exists for all j and $\bar{\chi}(f) = \max_{j\in[n]} \chi_j(f)$.

$$C := \{x \mid \max_{j \in [n]}(A_{ij} + x_j) \leq \max_{j \in [n]}(B_{ij} + x_j), \quad i \in I\}$$

Theorem

$\exists x \in C \setminus \{0\}$ *iff Max has at least one winning position in the mean payoff game with dynamic programming operator*

$$f_j(x) = (A^\sharp Bx)_j = \inf_{i \in I}(-A_{ij} + \max_{k \in [n]}(B_{ik} + x_k)) \ ,$$

i.e., $\exists j \in [n]$, $\chi_j(f) \geq 0$.

$$A = \begin{pmatrix} 2 & -\infty \\ 8 & -\infty \\ -\infty & 0 \end{pmatrix} \qquad B = \begin{pmatrix} 1 & -\infty \\ -3 & -12 \\ -9 & 5 \end{pmatrix}$$

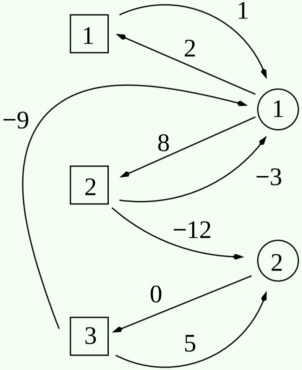

players receive the weight of the arc

$$\begin{aligned} 2 + x_1 &\leq 1 + x_1 \\ 8 + x_1 &\leq \max(-3 + x_1, -12 + x_2) \\ x_2 &\leq \max(-9 + x_1, 5 + x_2) \end{aligned}$$

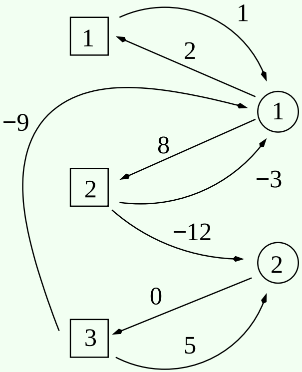

$$
\begin{aligned}
2 + x_1 &\leq 1 + x_1 \\
8 + x_1 &\leq \max(-3 + x_1, -12 + x_2) \\
x_2 &\leq \max(-9 + x_1, 5 + x_2)
\end{aligned}
$$

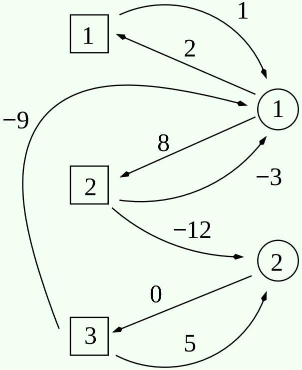

$\chi(g) = (-1, 5), x = (-\infty, 0)$ solution

Theorem

When C is a polyhedron, the set of winning initial positions $\{j \in [n] \mid \chi_j(f) \geq 0\}$ is exactly the union of supports (indices of finite entries) of the elements of C.

The proof relies on Kohlberg's theorem of existence of invariant half-lines.

Corollary

Whether an (affine) tropical polyhedron

$$\{x \mid \max(\max_{j \in [n]}(A_{ij} + x_j), c_i) \leq \max(\max_{j \in [n]}(B_{ij} + x_j), d_i),\ i \in [m]\}$$

is non-empty reduces to whether a specific state of a mean payoff game is winning.

Corollary

Each of the following problems:

1. *Is an (affine) tropical polyhedron empty?*
2. *Is a prescribed initial state in a mean payoff game winning?*

can be transformed in linear time to the other one.

One can then compute $\chi(f)$ by dichotomy solving the emptyness problem for convex polyhedra.

It follows that all these problems

- belong to NP $\cap$ co-NP ([Condon 92], [Zwick and Paterson 96])
- can be solved in pseudo-polynomial time (value iteration).
- other algorithms with experimentally fast average execution time:
  - pumping algorithm [Gurvich, Karzanov, and Khachiyan 88],
  - policy iteration ([Cochet, Gaubert, Gunawardena 98],....), parity game algorithm of [Jurdziński and Vöge 00], but the number of iterations may be exponential, see [Fridman, 2009].
- the existence of a polynomial algorithm is an open problem.

Mean payoff games associated to linear independence

Let $A \in M_{m,n}(\mathbb{R}_{\max})$. The columns of A are *tropically linearly dependent* if we can find scalars $x_1, \dots, x_n \in \mathbb{R}_{\max}$, not all equal to $-\infty$, such that "$Ax = \mathbf{0}$", that is for all $i \in [m]$, when evaluating the expression

$$\max_{j \in [n]} (A_{ij} + x_j)$$

the maximum is attained (at least) twice.

Equivalently, the rows of A belongs to the tropical hyperplane

$$\{z \mid \max_{j \in [n]} z_j + x_j \text{ attained twice}\} \ .$$

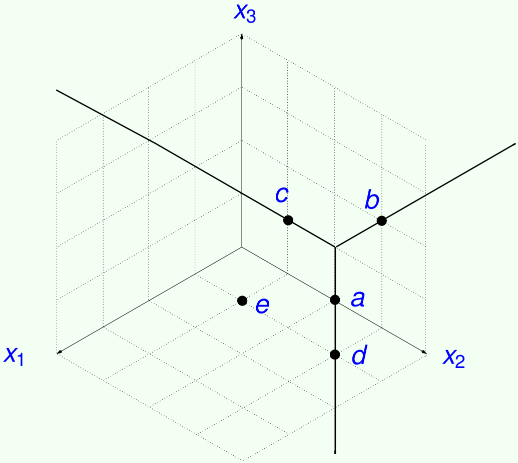

x_3
c
b
a
e
d
x_1
x_2

We define the game with dynamic programming operator

$$g_j(x) = \min_{i \in [m],\ (i,j) \in E} (-A_{ij} + \max_{k \in [n],\ k \neq j} (A_{ik} + x_k)) \ ,$$

where $E = \{(i,j) \mid A_{ij} \neq -\infty\}$.
$k \neq j$: the backspace move is forbidden for Max. So $\chi(g) \leq 0$.

Theorem

The columns of A are linearly dependent if and only if Max has at least one winning position in the game with operator g.

We define the game with dynamic programming operator

$$g_j(x) = \min_{i\in[m],\ (i,j)\in E} \left(-A_{ij} + \max_{k\in[n],\ k\neq j} (A_{ik} + x_k)\right) ,$$

where $E = \{(i,j) \mid A_{ij} \neq -\infty\}$.
$k \neq j$: the backspace move is forbidden for Max. So $\chi(g) \leq 0$.

Theorem

The columns of A are linearly dependent if and only if Max has at least one winning position in the game with operator g.

Idea of the proof. If in $(Au)_i$ the max is attained only once, then, there is an index j such that $A_{ij} + u_j > \max_{k\neq j} A_{ik} + u_k$. We deduce that $u_j > g_j(u)$. □

$$a = \begin{pmatrix} 0 & 2 & 0 \end{pmatrix} \; b = \begin{pmatrix} 0 & 3 & 2 \end{pmatrix} \; c = \begin{pmatrix} 0 & 1 & 1 \end{pmatrix} \; d = \begin{pmatrix} 1 & 3 & 0 \end{pmatrix}$$

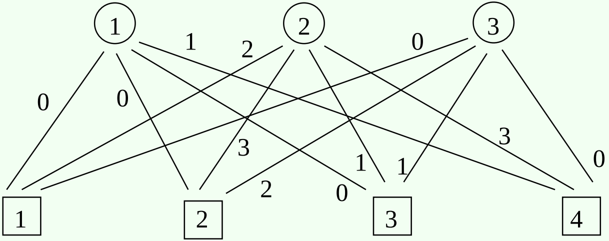

$a = \begin{pmatrix} 0 & 2 & 0 \end{pmatrix}$ $b = \begin{pmatrix} 0 & 3 & 2 \end{pmatrix}$ $c = \begin{pmatrix} 0 & 1 & 1 \end{pmatrix}$ $e = \begin{pmatrix} 1 & 1 & 0 \end{pmatrix}$

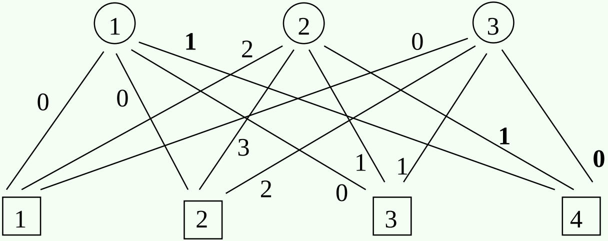

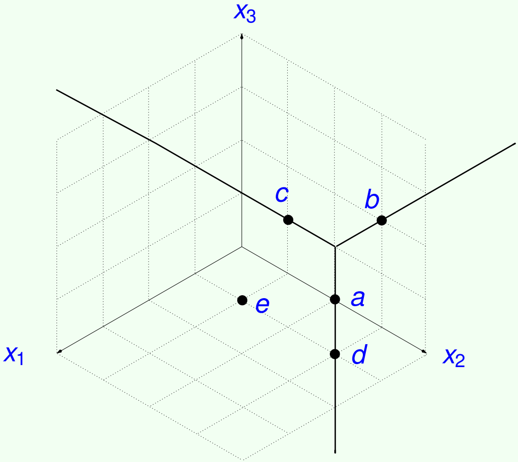

If one replaces d by e, we leave it to the reader to check that Max looses at all states.

A $n \times n$ matrix B is tropically nonsingular iff the optimal assignment problem

$$\max_{\sigma} \sum_{i \in [n]} B_{i\sigma(i)}$$

has a unique optimal solution. We get a game proof of what follows:

Corollary

If $m \geq n$, the columns of A are linearly independent if and only if there is a $n \times n$ tropically nonsingular submatrix (unique optimal assignment).

[Develin, Santos, Sturmels 05]: mixed subdivision proof (special case finite entries), see also [Izhakian, Rowen 09]. Can we find this matrix in polynomial time ?

Concluding remarks

- Tropical convexity yields a geometrical point of view on mean payoff games.
- Several tropical problems reduce to mean payoff games. See also current works of Gaubert and co-authors.
- Mean payoff deterministic games with finite action spaces $\Longleftrightarrow$ tropical linear programming...
- Can we find new algorithms for mean payoff games using the correspondance with tropical polyhedra?
- Can we find polynomial algorithms for all these problems?