From max-plus algebra to non-linear Perron-Frobenius theory

Stéphane Gaubert
INRIA

Stephane.Gaubert@inria.fr

CODE 2007 / Paris, April 18-20

A survey, gathering joint works with M. Akian, J. Gunawardena, B. Lemmens, R. Nussbaum, C. Walsh,. . . illustrated by applications done with E. Goubault, S. Zennou, A. Taly, L. Ninove.

Historical background: classical Perron-Frobenius theory

Perron (1907) proved the following.

Let $A \in \mathbb{R}^{n\times n}$, with $A_{ij} > 0 \; \forall i, j$. Then,

1. $\exists u \in \mathbb{R}^n$, $u_i > 0 \; \forall i$, $Au = \rho(A)u$, with $\rho(A) := \max\{|\lambda| \mid \lambda$ eigenval. of $A\}$.

2. The eigenvalue $\rho(A)$ is algebraically simple, a fortiori, u is unique up to a multiplicative constant.

Frobenius (1912) showed that the same is true when $A_{ij} \geq 0$, with $G := \{(i, j) \mid A_{ij} > 0\}$ strongly connected (A irreducible).

Let c be the cyclicity of G (=gcd lengths of circuits), and $\omega = \exp(i2\pi/c)$.

Then,

3. the whole spectrum of A is invariant by multiplication by ω, and $\omega^j \rho(A)$, $j = 0, \ldots, c-1$ are the only eigenvalues of maximal modulus (all algebraically simple)

4. so $\rho_{\max}(A)^{-kc}A^{kc}$ converges as k tends to ∞.

Kreĭn and Rutman (1948) considered more generally a linear operator A leaving a (closed, convex, pointed) cone C invariant in a Banach space E.

So A preserves the order: $x \leq y \iff y - x \in C$. When $E = \mathbb{R}^n$ and $C = \mathbb{R}_+^n$, we recover Perron-Frobenius theory.

(Garrett) Birkhoff (1957) approached Perron-Frobenius theory by means of Hilbert's geometry. *Hilbert's projective metric* is defined by

$$d_H(x, y) = \log \inf\{\frac{\beta}{\alpha} \mid \ \alpha y \leq x \leq \beta y,\ \alpha, \beta > 0\} \ .$$

It defines a metric on the set of rays included in the interior of C, i.e. on $\{\mathbb{R}_+ u \mid \ u \in \operatorname{int} C\}$, because $d(x, y) = 0$ iff $x = \beta y$. If C is normal, meaning that $0 \leq x \leq y \implies \|x\| \leq \gamma \|y\|$ for some constant γ, the latter metric space is complete.

Here is the intersection of a ball in Hilbert metric with the simplex, when $C = \mathbb{R}^3_+$:

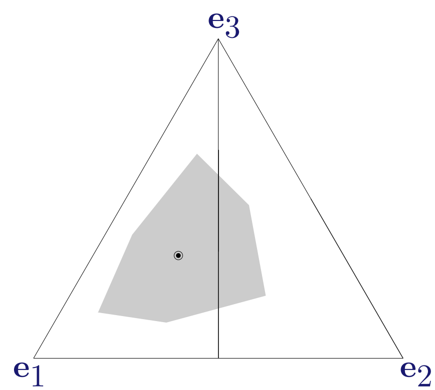

When $C = \mathbb{R}^n_+$, the Hilbert's metric can be understood by taking *logarithmic glasses*, setting $\log(x) := (\log(x_i))_{1 \le i \le n}$,

$$d_H(x, y) = \| \log x - \log y\|_H \text{ where } \|z\|_H := \max_i z_i - \min_i z_i$$

One may also consider *Thompson's metric*:

$$d_T(x, y) := \| \log x - \log y\|_\infty$$

Birkhoff showed that if A is a linear self-map of the interior of C, and if the diameter Δ of $A(C)$ in Hilbert's metric is finite, then

$$d_H(Ax, Ay) \leq \gamma d_H(x, y) \qquad \forall x, y \in \operatorname{int} C, \qquad \gamma := \tanh(\frac{\Delta}{4}) \ .$$

Perron's theorem is a corollary.

Max-plus spectral theory

There are remarkable similarities between the Perron-Frobenius problem and the *max-plus spectral problem*

Given $A = (A_{ij}) \in (\mathbb{R} \cup \{-\infty\})^{n \times n}$, find $u \in \mathbb{R} \cup \{-\infty\}^n,\ u \not\equiv -\infty$, $\lambda \in \mathbb{R}$, such that

$$\max_j A_{ij} + u_j = \lambda + u_i$$

The inhabitants of the max-plus world consider the semiring structure $\mathbb{R}_{\max}$, consisting of $\mathbb{R} \cup \{-\infty\}$, equipped with

$$a \oplus b := \max(a, b),\ \ a \otimes b := a + b\ \ , \qquad \mathbb{0} = -\infty, \mathbb{1} = 0$$

so $2 \oplus 3 = 3,\ 2 \otimes 3 = 5$ and the spectral problem reads $Au = \lambda u$.

$Au = \lambda u$ is the dynamic programming equation of the *simplest ergodic control problem*.

Example (N. Bacaer). Consider a field with n possible cultures, like wheat, oat,. . . or even fallow (no crop). Let A_{ij} denote the income of the land if crop j follows crop i. If the field is initialised with crop i, and if a bonus u_j is given for the final crop j, the optimal income is:

$$(A^k u)_i = \max_{i_1,\dots,i_k} A_{ii_1} + \cdots + A_{i_{k-1}i_k} + u_{i_k}$$

If u is an eigenvector, $(A^k u)_i = \lambda^k u_i$ in $\mathbb{R}_{\max}$, i.e. $k\lambda + u_i$, so the eigenvalue λ is the optimal reward per year, and u is a fair bonus, avoiding the "après nous le déluge" effect (Yakovenko, Kontorer).

Let $G := \{(i,j) \mid \ A_{ij} > \mathbb{0}\}$.

Max-plus spectral theorem. Assume that G is strongly connected.

1. $\rho_{\max}(A) := \max_{(i_1,\ldots,i_k)\text{circuit}} \dfrac{A_{i_1 i_2} + \cdots + A_{i_k i_1}}{k}$ is the only eigenvalue of A (so crop rotation is optimal).

The circuits which realise the maximum are called *critical*. Normalise rewards: $\tilde{A}_{ij} := -\rho_{\max}(A) + A_{ij}$, and define

$$\tilde{A}^+ := \tilde{A} \oplus \tilde{A}^2 \oplus \tilde{A}^3 \oplus \cdots$$

2. The columns $\tilde{A}_{\cdot i}$, with i in a critical circuit, generate the eigenspace.

Proved independently by several researchers, with various degree of generality and precision, including: Cuninghame-Green (61), Romanovskii (67), Vorobyev (67), Gondran-Minoux (77). For infinite dimensional versions, see Maslov's school (in particular the book by Maslov and Kolokoltsov, 97).

Cyclicity theorem (Cohen, Dubois, Quadrat, Viot 83, Nussbaum 88). There exists $c \geq 1$ such that the sequence $(\rho_{\max}(A)^{-kc}A^{kc})_{k=1,2,\ldots}$ is ultimately stationnary.

The smallest c is obtained as follows: define the critical graph to be the union of critical circuits, and take the lcm of the cyclicities of its strongly connected components.

Simpler than Perron-Frobenius, but more degenerate (several nonproportional eigenvectors), so contraction in Hilbert type metric would not work.

The analogy between both results can be explained by "Maslov's dequantisation". Define:

$$a \oplus_h b := h \log(e^{a/h} + e^{b/h}) \ .$$

Then, $(\mathbb{R} \cup \{-\infty\}, \oplus_h, +) \simeq (\mathbb{R}_+, +, \times)$, but

$$\lim_{h \to 0^+} a \oplus_h b = \max(a, b)$$

Same idea has been used recently in "tropical algebraic geometry", in relation with the theory of amoebas (Viro, Passare, Mikhalkin, Sturmfels,. . .)

This deformation transforms problems of algebra or analysis to "polyhedral problems" (combinatorial or discrete optimisation).

Can we embed Perron-Frobenius theory and max-plus spectral theory in a common perspective?

Answer:

nonlinear Perron-Frobenius theory

Keep A order preserving:

$$x \leq y \implies A(x) \leq A(y) \ ,$$

but drop linearity, and replace it by milder assumptions, like homogeneity, nonexpansiveness, convexity or concavity.

Motivations: 4 problems where nonlinear Perron-Frobenius theory arise

1. Lagrange problem / calculus of variations

$$v(t,x) = \sup_{X(0)=x,\ X(\cdot)} \int_0^t L(X(s),\dot{X}(s))ds + \phi(X(t)) \ .$$

$$\frac{\partial v}{\partial t} = H(x, \frac{\partial v}{\partial x}) \qquad v(0,\cdot) = \phi \qquad H(x,p) = \sup_u (L(x,u) + p\cdot u) \ .$$

Let S^t denote the Lax-Oleinik semigroup, so that $S^t\phi := v(t,\cdot)$.

$S^t(\sup(\phi_1,\phi_2)) = \sup(S^t\phi_1, S^t\phi_2)$, and $S^t(\lambda + \phi_1) = \lambda + S^t\phi_1$, so S^t is maxplus linear (infinite dimension now).

Max-plus eigenproblem. Find a function ϕ and $\lambda \in \mathbb{R}$ such that $S^t\phi = \lambda t + \phi$, for all $t \geq 0$.

2. Self-referential phenomena in web ranking

Google's page rank is based on the following idea:

the rank r_i of page i is the frequency of visit of this page by a random walker on the web graph.

Simplest model: let W denote the adjacency matrix of the web, so that $W_{ij} = 1$ if there is a link from page i to page j, and $W_{ij} = 0$ otherwise.

$$r = rP, \quad r \geq 0, \quad \sum_i r_i = 1, \qquad P_{ij} = \frac{W_{ij}}{\sum_k W_{ik}}$$

The actual google page rank relies on

$$r = r(1-\gamma)P + \gamma f$$

$\gamma = 0.15$ zapping probability, f vector of “preferences”: with probability γ, the user resets his exploration, and moves to the next page according to the probabilities in f.

Some users believe (foolishly?) that the pagerank measures quality. So the pagerank influences the behaviour of web users . . . which ultimately determines the pagerank.

A simple model of this circular effect (Akian, Ninove, SG).

Let T denote a social temperature, measuring the insensitivity of the user to the web rank.

If the current pagerank is r, the user moves from page i to page j with probability:

$$P_T(r)_{ij} = \frac{W_{ij} e^{r_j/T}}{\sum_k W_{ik} e^{r_k/T}}$$

If $T = \infty$, we recover the basic pagerank model.

The pagerank of tomorrow r_T^{k+1} is obtained as follows from the pagerank of today, r_T^k:

$$r_T^{k+1} = r_T^{k+1} P_T(r^k)$$

Does r_T^k converge ?

Let u_T denote the map $r_T^k \to r_T^{k+1}$.

By Tutte's matrix tree theorem, assuming that W is irreducible:

$$u_T(x) = h_T(x)/(\sum_k h_T(x)_k)$$

where

$$h_T(x)_l = (\sum_k W_{ik} e^{x_k/T})(\sum_{R \to l} \prod_{(i,j) \in R} W_{ij} e^{x_j/T}) \ ,$$

the latter sum being taken over river networks R with sea l.

So h is order preserving.

3. Zero-sum repeated games

Dynamic operators of zero-sum repeated games with state space $\{1, \ldots, n\}$ are of the form:

$$f : \mathbb{R}^n \to \mathbb{R}^n \qquad f_i(x) = \inf_{a \in A(i)} \sup_{b \in B(i,a)} (r_i^{ab} + P_i^{ab} x)$$

$P_i^{ab} := (P_{ij}^{ab})$, proba. of moving $i \to j$, $\sum_j P_{ij}^{ab} \leq 1$.

r_i^{ab}: payment of Player I to player II.

Questions: $f^k(0)$ value of the game in horizon k, as $k \to \infty$? $x = f(x)$?

f is order preserving, and nonexpansive in the sup-norm: $\|f(x) - f(y)\|_\infty \leq \|x - y\|_\infty$

Neyman, Sorin, Rosenberg,. . . have developed an "operator approach" of games using these properties.

Rubinov and Singer have shown that any order preserving sup-norm nonexpansive map can be represented by such a game (with deterministic transition probabilities P_i^{ab}).

4. Static analysis of programs by abstract interpretation

Goal: determine automatically the invariants of critical programs (airplanes, . . .), well established problem since Ariane 501.

Unsolvable in full generality and precision, because the halting of Turing machines is undecidable.

P. Cousot's abstract interpretation method gives an approached solution. Let $\mathcal{T}$ denote a complete lattice of subsets of $\mathbb{R}^d$, e.g.: (products of) intervals, polyhedra, convex sets,. . .

To each breakpoint i of the program, is associated a set $x_i \in \mathcal{T}$ which is an overapproximation of the set of reachable values of the variables, at this breakpoint. So $d =$ number of variables of the program.

The best x is the smallest solution of a fixed point problem $x = f(x)$ with f order preserving.

```
void main() {
  int x=0;          // 1
  while (x<100) {   // 2
    x=x+1;          // 3
  }                 // 4
}
```

$$
\begin{aligned}
x_1 &= [0,0] \\
x_2 &=]-\infty, 99] \cap (x_1 \cup x_3) \\
x_3 &= x_2 + [1,1] \\
x_4 &= [100, +\infty[\cap (x_1 \cup x_3)
\end{aligned}
$$

Let $x_2^+ := \max x_2$. After some elimination, we arrive at

$$x_2^+ = \min(99, \max(0, x_2^+ + 1)) \ .$$

The smallest x_2^+ is 99, it is the value of a zero-sum game.

Problem: standard fixed point iteration takes 99 iterations to converge, we need to compute x more rapidly and accurately.

In a joint work with E. Goubault, S. Zennou, . . . at CEA, we are applying game algorithms (policy iteration) to address this, more latter in this talk.

```
void main() {
 i = 1; j = 10;
 while (i <= j){
  i = i + 2;
  j = j - 1; }
}
```

Policy iteration:

$$5 \leq i \leq 10,\ 4 \leq j \leq 8,\ -3 \leq j - i \leq -1$$

Kleene on octagons:

$$6 \leq i \leq 12,\ \frac{9}{2} \leq j \leq 10,\ -3 \leq j - i \leq -1$$

When the arithmetics of the program is linear (no product or division of variables), and when the lattice of Manna's "templates" ($\sim$ discretised support functions of polyhedra) is chosen, f can essentially be written as:

$$f_i(x) = \inf_{a \in A(i)} \sup_{b \in B(i,a)} (r_i^{ab} + P_i^{ab} x)$$

with $P_i^{ab} := (P_{ij}^{ab})$, $P_{ij}^{ab} \geq 0$, but possibly $\sum_j P_{ij}^{ab} > 1 \to$ game in infinite horizon with "negative discount rate".

Some results of non-linear Perron-Frobenius theory

1. Convergence to periodic orbits.

Theorem (Akian, SG, Lemmens, Nussbaum, Math. Proc. Camb. Phil. Soc.,06)**.** *Let C be a polyhedral cone with N facets in a finite dimensional vector space X. If $f : C \to C$ is a continuous order preserving subhomogeneous map and the orbit of $x \in C$ is bounded, then $\lim_{k\to\infty} f^{kp}(x)$ exists with*

$$p \leq \max_{q+r+s=N} \frac{N!}{q!r!s!} = \frac{N!}{\lfloor \frac{N}{3} \rfloor ! \lfloor \frac{N+1}{3} \rfloor ! \lfloor \frac{N+2}{3} \rfloor !} .$$

Subhomogeneous means that $f(\lambda x) \leq \lambda f(x)$ for $\lambda > 1$.

This comes after a long series of works on periodic orbits of nonexpansive maps when the norm is polyhedral: Ackoglu and Krengel, Weller, Martus, Nussbaum, Sine, Scheutzow, Verdyun-Lunel,. . .

Some ingredients of the proof: Reduce $C = \mathbb{R}^n_+$. If the orbit of x stays in the interior of C, we can look at it with logarithmic glasses, i.e., consider $F := \log \circ f \circ \exp$, which is order preserving and nonexpansive in the sup-norm. Then a result of Lemmens and Scheutzow (Erg. Th. Dyn. S. 05) shows that the orbit length is at most

$$\frac{N!}{\lfloor \frac{N}{2} \rfloor ! \lfloor \frac{N+1}{2} \rfloor !}$$

Conjecture (Nussbaum). If F is non-expansive in the sup-norm (but not order preserving) the bound becomes 2^N.

Application. Assume that $f : \mathbb{R}^n \to \mathbb{R}^n$ is order preserving and additively homogeneous: $f(\alpha + x) = \alpha + f(x)$ (think to undiscounted games). Here, $\alpha + x := (\alpha + x_i)_{1\leq i\leq n}$.

Assume that f has an additive eigenvector, so $f(u) = \lambda + u$, $\lambda \in \mathbb{R}$, $u \in \mathbb{R}^n$, then, for all $x \in \mathbb{R}^n$,

$$f^k(x) = k\lambda + \text{asymp. periodic. term in } k$$

$\neq$ Difficult case in which $f_i^k(x) - f_j^k(x) \to \infty$: Neyman, Sorin, Rosenberg. . .

2. Existence of eigenvectors.

Theorem (SG, Gunawardena, TAMS 04). *Assume that* f *is order preserving and additively homogeneous, and that the recession function*

$$\hat{f}(x) := \lim_{t\to\infty} t^{-1} f(tx)$$

exists. If

$$\hat{f}(x) = x \implies x_1 = \cdots = x_n$$

then

$$\exists u \in \mathbb{R}^n,\ \exists \lambda \in \mathbb{R}, \qquad f(u) = \lambda + u$$

Application. A game inspired by Richman games / discretisation of infinity Laplacian.

Let

$$f_i(x) = \frac{1}{2}\Big(\min_{(i,j)\in G} A_{ij} + x_j + \max_{(i,j)\in G} A_{ij} + x_j\Big) \ .$$

Two players. One flips a coin to decide who plays. Player MIN plays A_{ij} to Player MAX if the move is $i \to j$.

$$\hat{f}(x) = \frac{1}{2}\Big(\min_{(i,j)\in G} x_j + \max_{(i,j)\in G} x_j\Big)$$

If $x_i = m := \max_k x_k$, and $x = \hat{f}(x)$, $(i,j) \in G \implies x_j = m$. So $x = \hat{f}(x) \implies x_1 = \cdots = x_n$ if G is strongly connected.

3. Uniqueness of eigenvector / fixed point.

The uniqueness of the T-pagerank can be derived from the following finite dimensional version of (Nussbaum, 1988):

Theorem (Nussbaum). *Let* $\Sigma := \{x \in \mathbb{R}_n \mid \ x \in \mathbb{R}^n_+ \ \sum_i x_i = 1\}$. *Assume that* $h : \mathbb{R}^n_+ \to \mathbb{R}^n_+$ *is continuous, order preserving and subhomogeneous on* $\operatorname{int}\Sigma$, *meaning that* $h(\alpha x) \leq \alpha h(x)$, *for* $\alpha > 1$ *and* $x \in \operatorname{int}(\Sigma)$. *Assume that* $h(x) = \lambda x$ *with* $x \in \operatorname{int}\Sigma$, *that* h *is* C^1, *and that* $h'(x)$ *is irreducible. Then,* x *is the only eigenvector of* h *in* $\in \Sigma$.

Corollary (Akian, SG, Ninove, Posta'06). *For* $T \geq n$, *the* T*-pagerank is unique.*

For small values of T, one can show that there are always several T-pageranks.

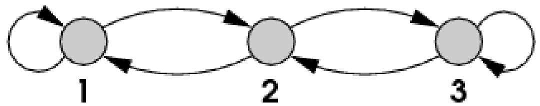

If $T \approx 0$ and $r_T^0 = (\frac{1+\varepsilon}{3}\ \frac{1}{3}\ \frac{1-\varepsilon}{3})$ then $\lim_{t\to\infty} r_T^k \approx (1\ 0\ 0)$.

Provocative interpretation: if one believes that the pagerank measures quality, then the pagerank might become meaningless.

In a recent work with Nussbaum, we extend the latter general uniqueness result. It is enough to assume that h is *semidifferentiable* at point x, meaning that there exists a continuous positively homogeneous map h'_x such that

$$h(x+y) = h(x) + h'_x(y) + o(\|y\|)$$

This works in infinite dimension (normal cone), under mild compactness assumptions (Fredholm type).

4. Representation of the fixed point set

Precise results are available when $f : \mathbb{R}^n \to \mathbb{R}^n$ is order preserving, nonexpansive in the sup-norm, and *convex*.

Then, by Legendre-Fenchel duality

$$f(x) = \sup_{P \in S_n^+} (Px - f^*(P))$$

where S_n^+ denote the set of substochastic matrices, and $f^*(P) \in (\mathbb{R} \cup \{+\infty\})^n$.

Compare with the dynamic programming operator of a stochastic control problem with state space $\{1, \ldots, n\}$:

$$f_i(x) = \sup_{a \in A(i)} (r_i^a + P_i^a x)$$

The previous expression is a canonical form of f: when in state i, the player chooses the substochastic vector $P_i \in \operatorname{dom} f_i^*$, receives the payment $-f_i^*(P)$ when in state i, and moves to j with probability P_{ij}.

The ergodic control problem consists in finding $P \in \operatorname{dom} f^*$ such that $-mf^*(P)$ is maximal, where m is an invariant measure of P. If $f(u) = \lambda + u$ with $\lambda \in \mathbb{R}$ and $u \in \mathbb{R}^n$, this maximum is equal to λ.

Normalise f and assume that $f(u) = u$. We say that i is *critical* if it belongs to a recurrence class of some matrix P such that $f(u) = Pu - f^*(P)$ (in other words, if it is recurrent for a stationnary strategy which is optimal for the ergodic problem). Let C denote the set of critical nodes.

Theorem (Akian, SG, NLA TMA 03). *The restriction $x \mapsto x_C$ is a sup-norm isometry from $\{x \mid f(x) = x\}$ to a convex set.*

“An harmonic function is defined uniquely by its value on the boundary”. Here, C plays the role of the boundary, it is a discrete version of the Aubry set arising in Fathi weak KAM's theory.

Continuous time/second order PDE special case: current work with Akian and David.

Policy iteration for static analysis

Policy iteration was introduced by Hoffman and Karp (66) for stochastic games with mean payoff with irreducible transition matrices.

In static analyis, transition matrices may be degenerate. We are interested in the smallest fixed point of

$$f = \inf \mathcal{G}$$

where $\mathcal{G}$ is a set of "simpler" self-maps of a lattice $\mathcal{L}$ (for instance, $\mathcal{L} = \overline{\mathbb{R}}^n$).

We say that a set $\mathcal{G}$ of maps from a set X to a lattice $\mathcal{L}$ *admits a lower selection* if for all $x \in X$, there exists a map $g \in \mathcal{G}$ such that $g(x) \leq h(x)$, for all $h \in \mathcal{G}$.

Example. Take $\mathcal{L} = \overline{\mathbb{R}}$, and consider the self-map of $\mathcal{L}$, $f(x) = \inf_{1\leq i\leq m} \max(a_i + x, b_i)$, where $a_i, b_i \in \mathbb{R}$. The set $\mathcal{G}$ consisting of the m maps $x \mapsto \max(a_i + x, b_i)$ admits a lower selection.

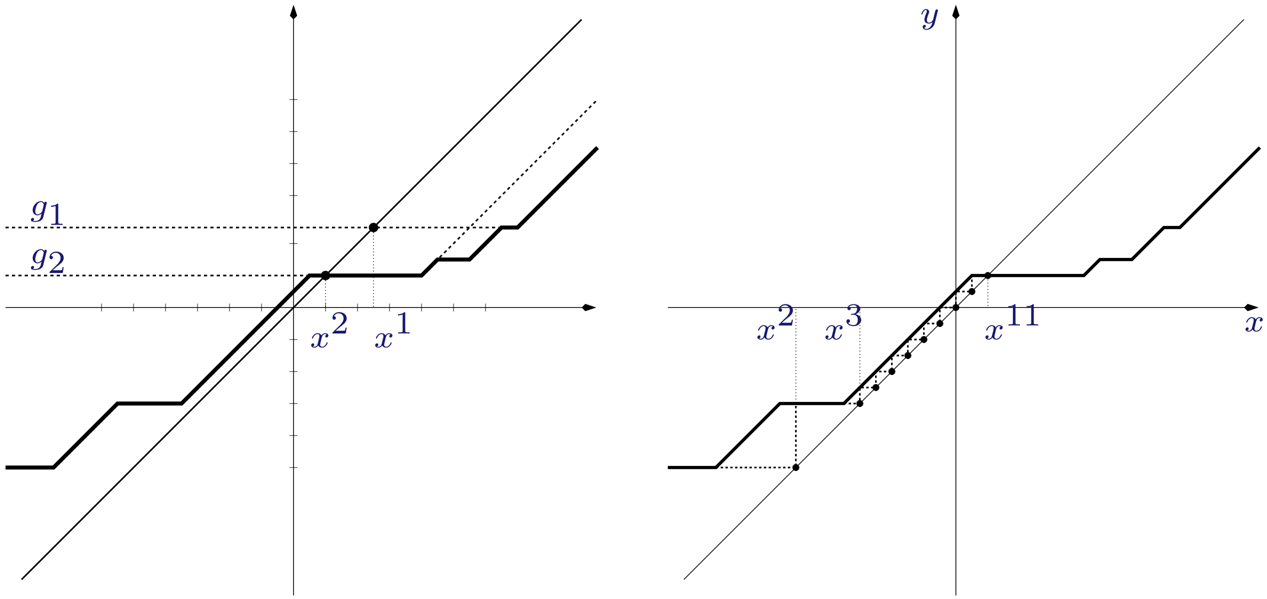

We denote by f^- the smallest fixed point of a monotone self-map f of a complete lattice $\mathcal{L}$, whose existence is guaranteed by Tarski's fixed point theorem.

Theorem (Costan, SG, Goubault, Martel, Putot CAV'05). *Let $\mathcal{G}$ denote a family of monotone self-maps of a complete lattice $\mathcal{L}$ with a lower selection, and let $f = \inf \mathcal{G}$. Then $f^- = \inf_{g \in \mathcal{G}} g^-$.*

The input of the following algorithm consists of a finite set $\mathcal{G}$ of monotone self-maps of a lattice $\mathcal{L}$ with a lower selection. When the algorithm terminates, its output is a fixed point of $f = \inf \mathcal{G}$.

1. *Initialization*. Set $k = 1$ and select any map $g_1 \in \mathcal{G}$.

2. *Value determination*. Compute a fixed point x^k of g_k.

3. Compute $f(x^k)$.

4. If $f(x^k) = x^k$, return x^k.

5. *Policy improvement*. Take g_{k+1} such that $f(x^k) = g_{k+1}(x^k)$. Increment k and goto Step 2.

The algorithm does terminate when at each step, the smallest fixed-point of g_k, $x^k = g_k^-$ is selected.

Open problem 1: the algorithm may return a nonminimal fixed point of f, what to do next if f is not sup-norm nonexpansive.

Open problem 2: the worst case complexity is not known (but Condon showed: mean payoff games is in NP $\cap$ co-NP).

In general, every g is convex, and the smallest finite fixed point of g is typically computed by solving a convex program: $\min \sum_i x_i; \quad g(x) \leq x$. Difficulties arise when there is no smallest finite fixed point (unlike in the case of classical games, x_i may be in $\overline{\mathbb{R}}$).

```
i = 150;
j = 175;
while (j >= 100){
  i++;
  if (j<= i){
      i = i - 1;
      j = j - 2;
  }
}
```

$$
\begin{aligned}
M_0 &= \mathit{context\_initialization} \\
M_2 &= (\mathit{Assignment}\ (i \leftarrow 150,\ \ j \leftarrow 175)(M_0))^* \\
M_3 &= ((M_2 \sqcup M_8) \sqcap (j \geq 100))^* \\
M_4 &= (\mathit{Assignment}\ (i \leftarrow i + 1)(M_3))^* \\
M_5 &= (M_4 \sqcap (j \leq i))^* \\
\\
M_7 &= (\mathit{Assignment}\ (i \leftarrow i - 1,\ \ j \leftarrow j - 2)(M_5))^* \\
M_8 &= ((M_4 \sqcap (j > i))^* \sqcup M_7 \\
M_9 &= ((M_2 \sqcup M_8) \sqcap (j < 100))^*
\end{aligned}
$$

$$
\text{IP} \begin{cases} 150 \leq i \leq 174 \\ 98 \leq j \leq 99 \\ -76 \leq j - i \leq -51 \end{cases}
\qquad
\text{Mine's Octogon} \begin{cases} 150 \leq i \\ 98 \leq j \leq 99 \\ j - i \leq -51 \\ 248 \leq j + i \end{cases}
$$

Related work: policy iteration for repeated stochastic games with mean payoff and degenerate transition probabilities (with Cochet CRAS'06, stochastic case, and Dhingra, Valuetools'06, fast combinatorial implementation / deterministic).

$G = (V, E)$ directed bipartite graph, r_{ij} weight of arc $(i, j) \in E$.

Two players, "Max", and "Min", move a pawn.

The pawn is initially at a given node $i_0 \in V$. The player who plays first, chooses an arc (i_0, i_1) in E, moves the pawn from i_0 to i_1, and Min pays $r_{i_0 i_1}$ to him. Then, the other player chooses an arc (i_1, i_2) in E, moves the pawn from i_1 to i_2, and pays $r_{i_1 i_2}$ to Max, etc.

The reward of Max (or the loss of Min) after k turns is

$$r_{i_0 i_1} + \cdots + r_{i_{2k-1} i_{2k}} \ .$$

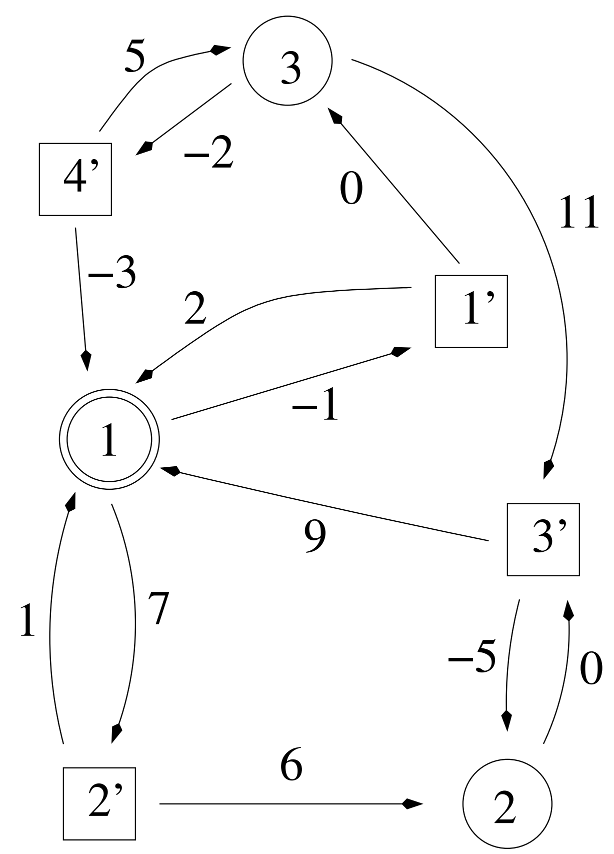

The circles (resp. squares) represent the nodes at which Max (resp. Min) can play. The initial node, "1", is indicated by a double circle.

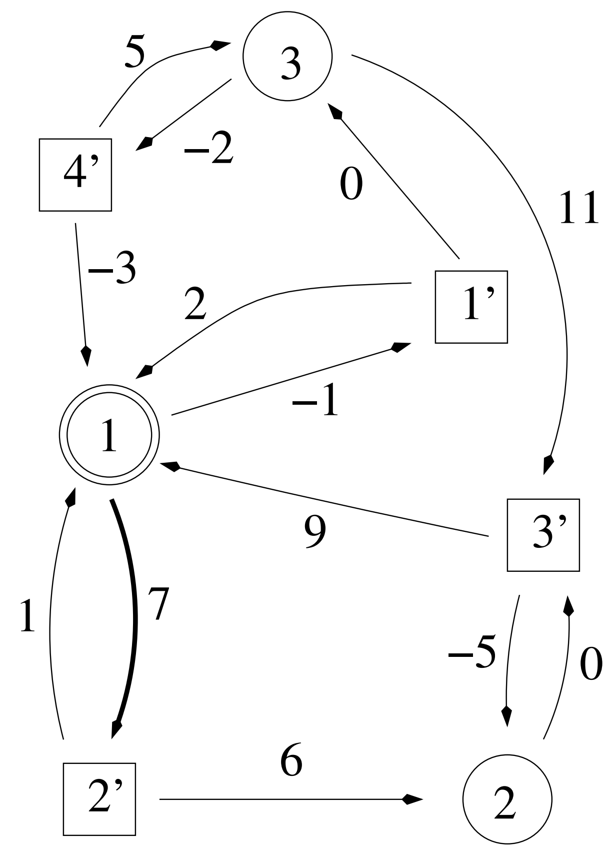

If Max initialy moves to $2'$. . .

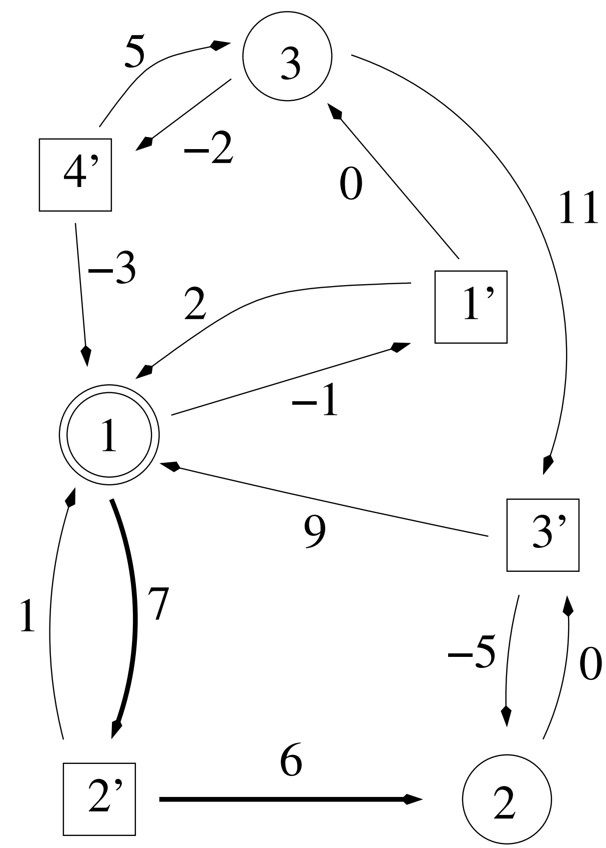
5
3
−2
4’
0
11
−3
2
1’
1
−1
9
3’
1
7
−5
0
6
2’
2

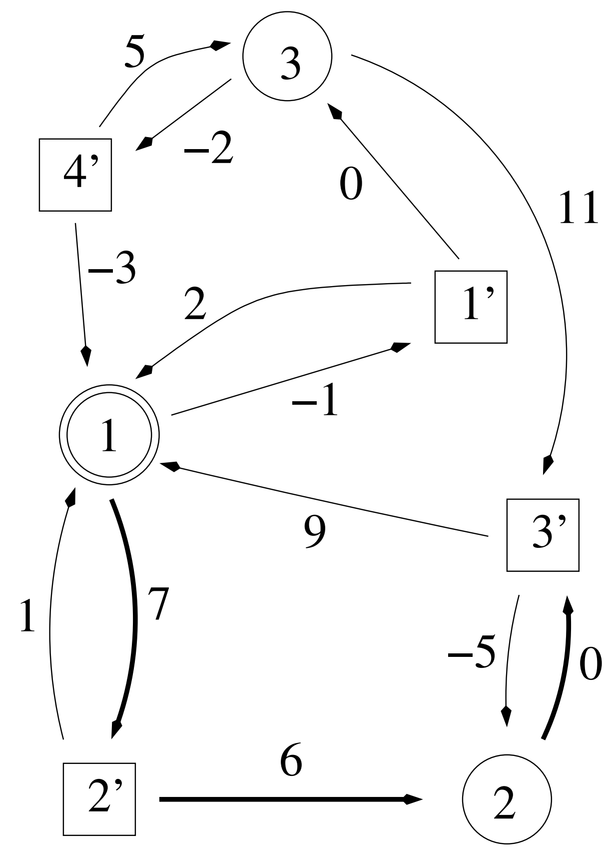
5
3
−2
4’
0
11
−3
2
1’
−1
1
9
3’
1
7
−5
0
6
2’
2

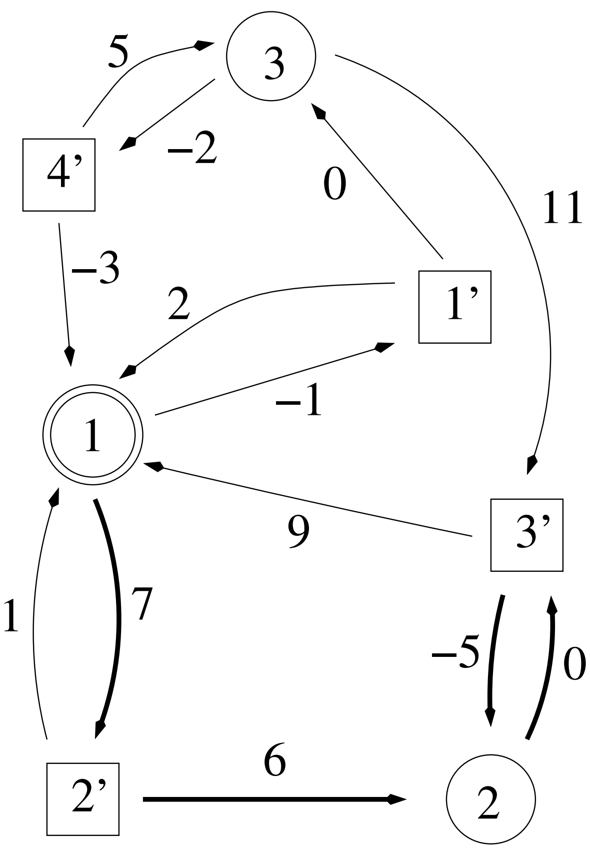

he eventually looses 5 per turn.

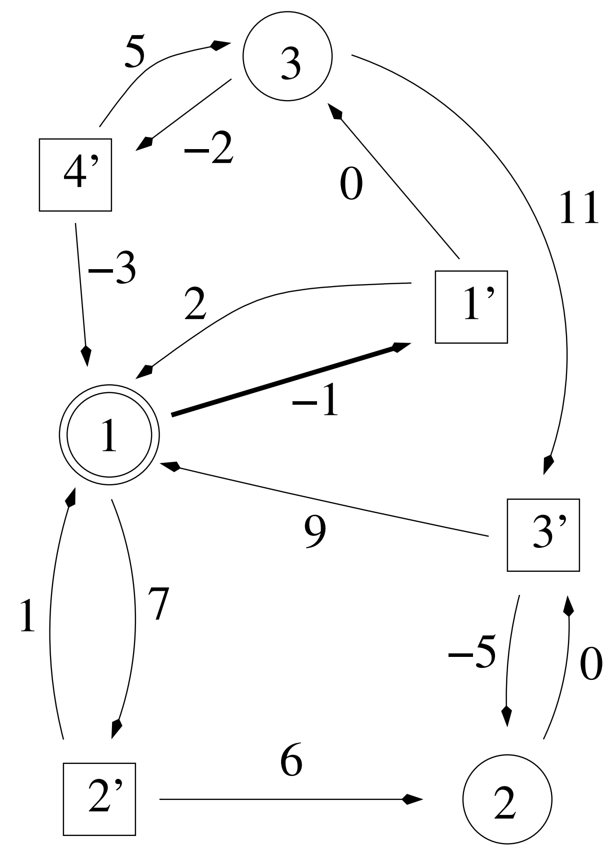

But if Max initialy moves to $1'$. . .

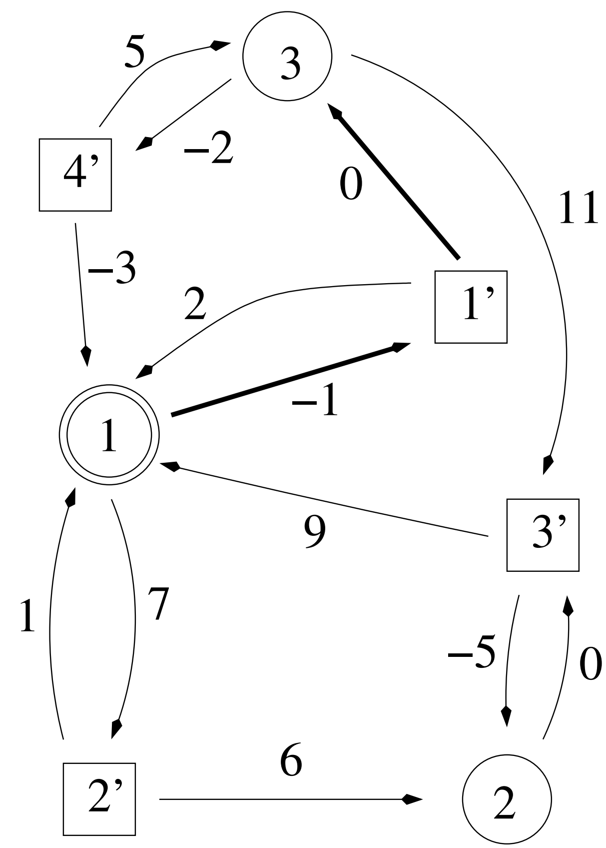
5
3
−2
4’
0
11
−3
2
1’
1
−1
9
3’
1
7
−5
0
6
2’
2

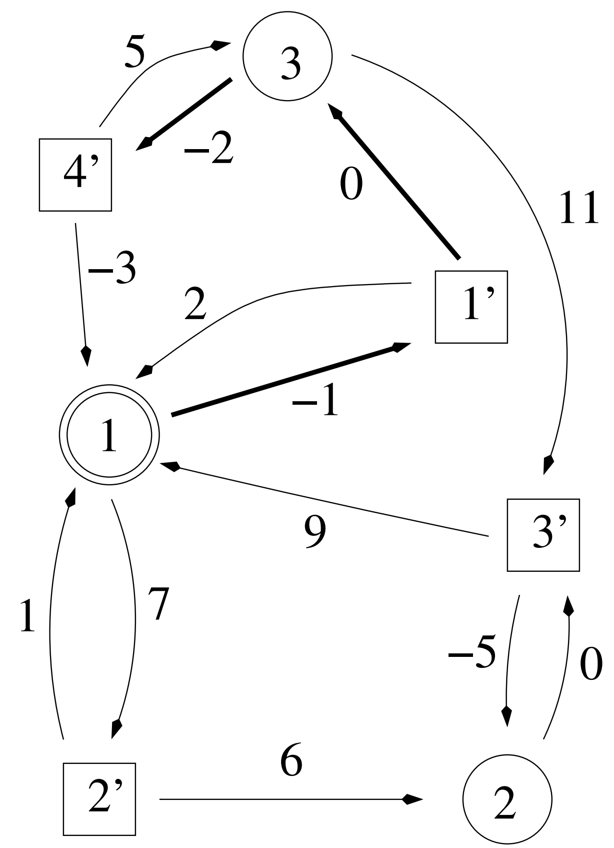
5
3
−2
4’
0
11
−3
2
1’
1
−1
9
3’
1
7
−5
0
6
2’
2

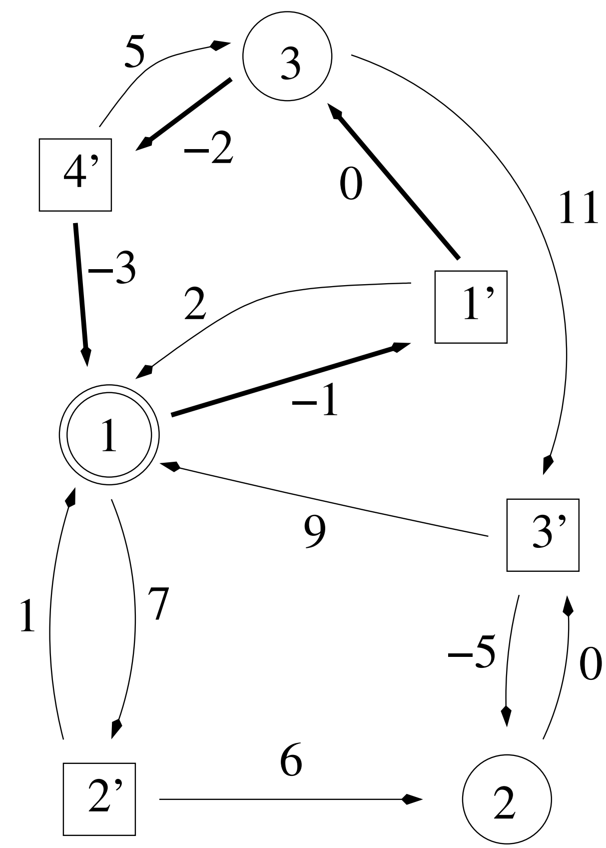

he only looses eventually $(1 + 0 + 2 + 3)/2 = 3$ per turn.

1) We write the dynamic programming operator as an infimum of g, where g is a "one player map".

2) For each g, we compute an invariant half-line, $w(t) = v + t\eta$ such that $f(w(t)) = w(t+1)$ for $t \in \mathbb{R}$ large enough. (Done by linear programming, or even policy iteration).

η gives the vector of mean payoff per time unit.

Naive policy iteration does not work, because w is not unique. The key idea of the algorithm is to use the critical graph of g, to control the nonuniqueness of w, similarity with the reduction of super-harmonic functions, CRAS'06. Each step requires solving a one player stochastic optimal stopping problem.

Experiments

Complete bipartite graphs, in which $n = p$. Random weights (uniform).

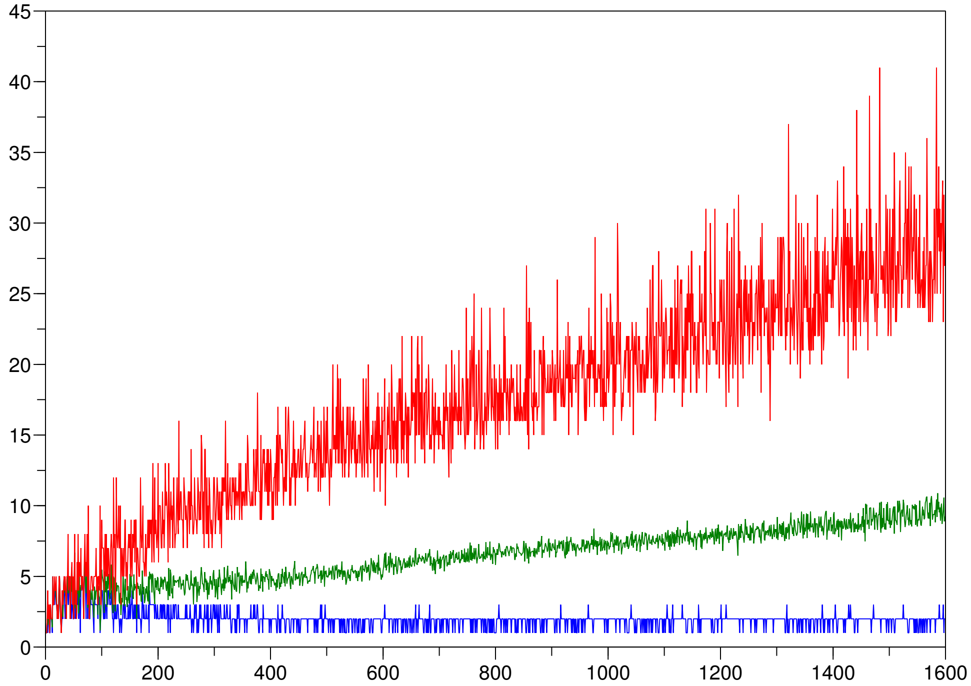

N_{min}= number of strategies of Min before the algorithm terminates. 100 graphs for each n, max, average, and min of N_{min} shown.

Sparse bipartite graphs. n nodes of each kind, every node has exactly 2 successors drawn at random; random weights.

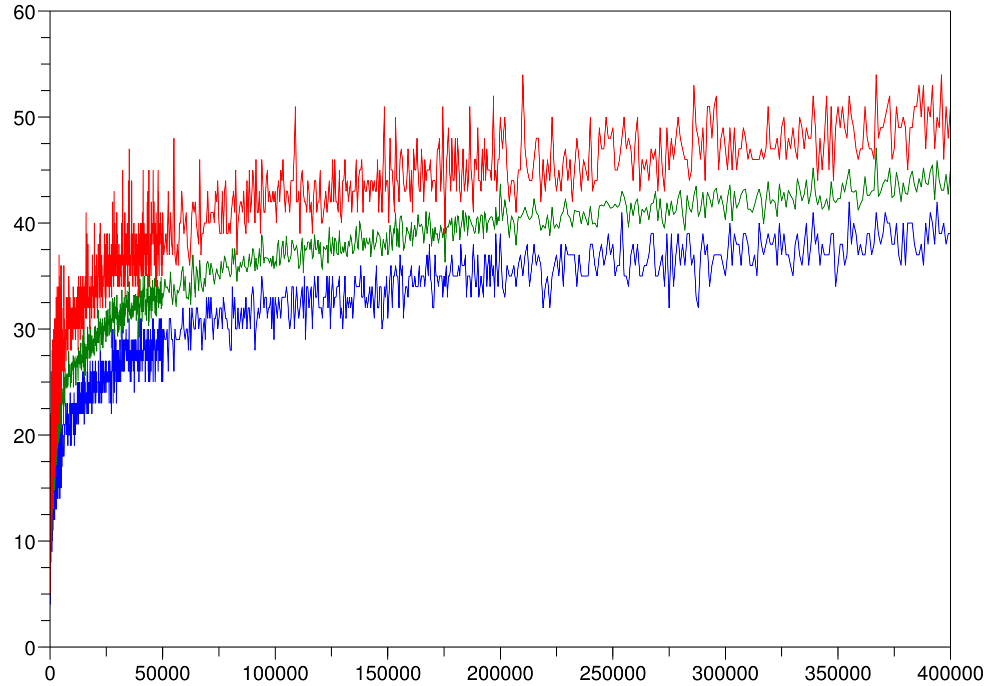

Number of iterations of minimizer $N_{\min}$

The max-plus Martin boundary

The analogy with classical potential theory is particularly visible for the spectral problem over a non compact state space (Work with Akian and Walsh, 04).

Recall the classical Martin boundary theory, discrete case for simplicity (Dynkin).

Let P_{xy} denote a Markov kernel, over a discrete infinite set E. We wish to find all nonnegative harmonic functions: $u = Pu$.

1) Define the Green kernel: $G = P^0 + P + P^2 + \cdots$

2) The Martin kernel is:

$$K_{xy} = \frac{G_{xy}}{G_{by}}$$

where $b \in E$ is a basepoint.

3) Let $\mathcal{K} := \{K_{\cdot y} \mid \; y \in E\}$

4) The Martin space $\mathcal{M}$ is the closure of $\mathcal{K}$ in the product topology.

5) The Martin boundary is $\mathcal{B} := \mathcal{M} \setminus \mathcal{K}$.

Theorem (Martin representation). *Every harmonic function u can be written as a positive linear combination of functions from the boundary:*

$$u = \int_{\mathcal{B}} w\mu(dw) \; .$$

μ can be choosen to be supported by a subset of $\mathcal{B}$, the minimal Martin boundary. (We recognise Choquet's theorem!).

Theorem. Akian, SG, Walsh 04. A similar representation theorem holds for discrete max-plus harmonic functions:

$$u_i = \sup_j A_{ij} + u_j$$

and harmonic functions of Lax-Oleinik semigroups

$$u = S^t u, \ \ \forall t \geq 0$$

The Martin kernel reads: $K_{xy} = A^*_{xy} - A^*_{by}$.

The max-plus Martin boundary has been defined by Gromov (78) when $A^*_{xy} = -d(x,y)$, its elements are horofunctions or Busemann functions.

The Martin representation should be understood in the max-plus sense:

$$u = \sup_{w \in \mathcal{M}_m} w + \mu(w), \qquad \mu : \mathcal{M}_m \to \mathbb{R} \cup \{-\infty\}$$

$\mathcal{M}_m$ minimal Martin space (may include contain non boundary -recurrent-points).

This extends the representation of weak KAM solutions by Fathi (compact case). A similar representation has been found in the non compact case by Ishii and Mitake (06), working with viscosity techniques - we work directly with the Lax Oleinik semigroup.

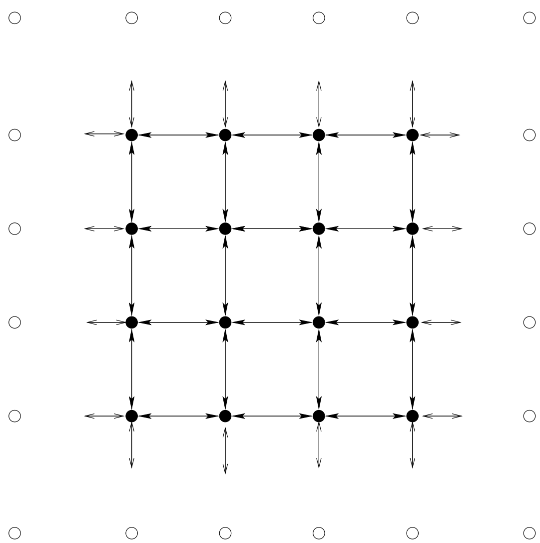

$$S := \mathbb{Z}^2$$

$$A_{ij} := \begin{cases} -1, & \text{if } i \text{ and } j \text{ are nearest neighbours,} \\ -\infty, & \text{otherwise.} \end{cases}$$

The boundary is the square at infinity, in the probabilistic case, it is the circle (c.f. Ney & Spitzer '65).

Example: Linear quadratic control

Hamilton–Jacobi equation

$$\lambda = -|\mathbf{x}|^2 + \frac{1}{4}|\nabla w|^2$$

Maximise reward:

$$-\int_0^T \left(|\gamma(t)|^2 + |\dot{\gamma}(t)|^2 + \lambda\right) dt,$$

If $\lambda > 0$, solutions are

$$w(\mathbf{x}) = \sup_{\mathbf{n}} (\nu(\mathbf{n}) + h_{\mathbf{n}}(\mathbf{x})),$$

where ν is an upper semi–continuous map from the unit vectors to $\mathbb{R} \cup \{-\infty\}$.

When $\lambda = 0$, there is a horofunction for each direction $\mathbf{n}$:

$$h_{\mathbf{n}}(\mathbf{x}) = \begin{cases} -|\mathbf{x}|^2 + 2(\mathbf{x}\cdot\mathbf{n})^2, & \text{if } \mathbf{x}\cdot\mathbf{n} > 0, \\ -|\mathbf{x}|^2, & \text{otherwise.} \end{cases}$$

The function $-|\mathbf{x}|^2$ is also a horofunction.

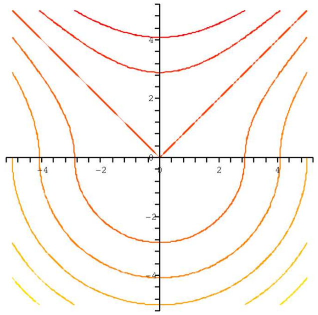

Horospheres of $h_{\mathbf{n}}$ with $\mathbf{n} = (0, 1)$.

When $\lambda > 0$: for each direction $\mathbf{n}$,

$$h_{\mathbf{n}}(\mathbf{x}) = -\lambda\frac{|\mathbf{x}|^2}{R^2} + \mathbf{x}\cdot\mathbf{n}\frac{\lambda + 2|\mathbf{x}|^2}{R} - \lambda\log\frac{R}{\sqrt{\lambda}},$$

where $R := \sqrt{(\mathbf{x}\cdot\mathbf{n})^2 + \lambda} - \mathbf{x}\cdot\mathbf{n}$.

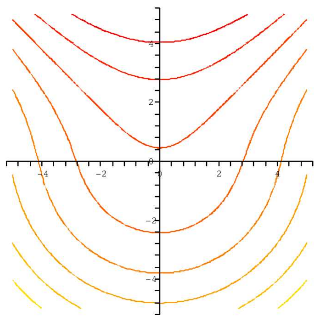

Horospheres of $h_{\mathbf{n}}$ with $\mathbf{n} = (0, 1)$.

That's all...