Stake-governed random tug-of-war

Gábor Pete
Rényi Institute and TU Budapest
http://www.math.bme.hu/∼gabor

joint work with Alan Hammond
UC Berkeley

Why random-turn games?

Instead of alternating moves, flip a fair coin before each turn. Introduced by Peres, Schramm, Sheffield, Wilson in 2006-09.

0. Possibly interesting **new games**.

1. Random-turn selection games: new algorithms to compute Boolean functions, often with low revealment, à la OSSS '05 and Schramm-Steif '10, important in sharp thresholds and noise sensitivity.

2. Random tug-of-war: a discrete game to study continuum ∞-harmonic functions, the solutions to the degenerate elliptic PDE

$$0 = \Delta_\infty u := \|\nabla u\|^{-2} \sum_{i,j} u_{x_i} u_{x_i x_j} u_{x_j} \, ,$$

vanishing second derivative in the direction of the gradient. Degenerate, because viscosity solutions are often not twice differentiable.

Also: p-harmonic functions are the minimizers of the L^p-norm of the gradient. Then let $p \to \infty$: Absolutely Minimizing Lipschitz Extensions.

The game of Hex

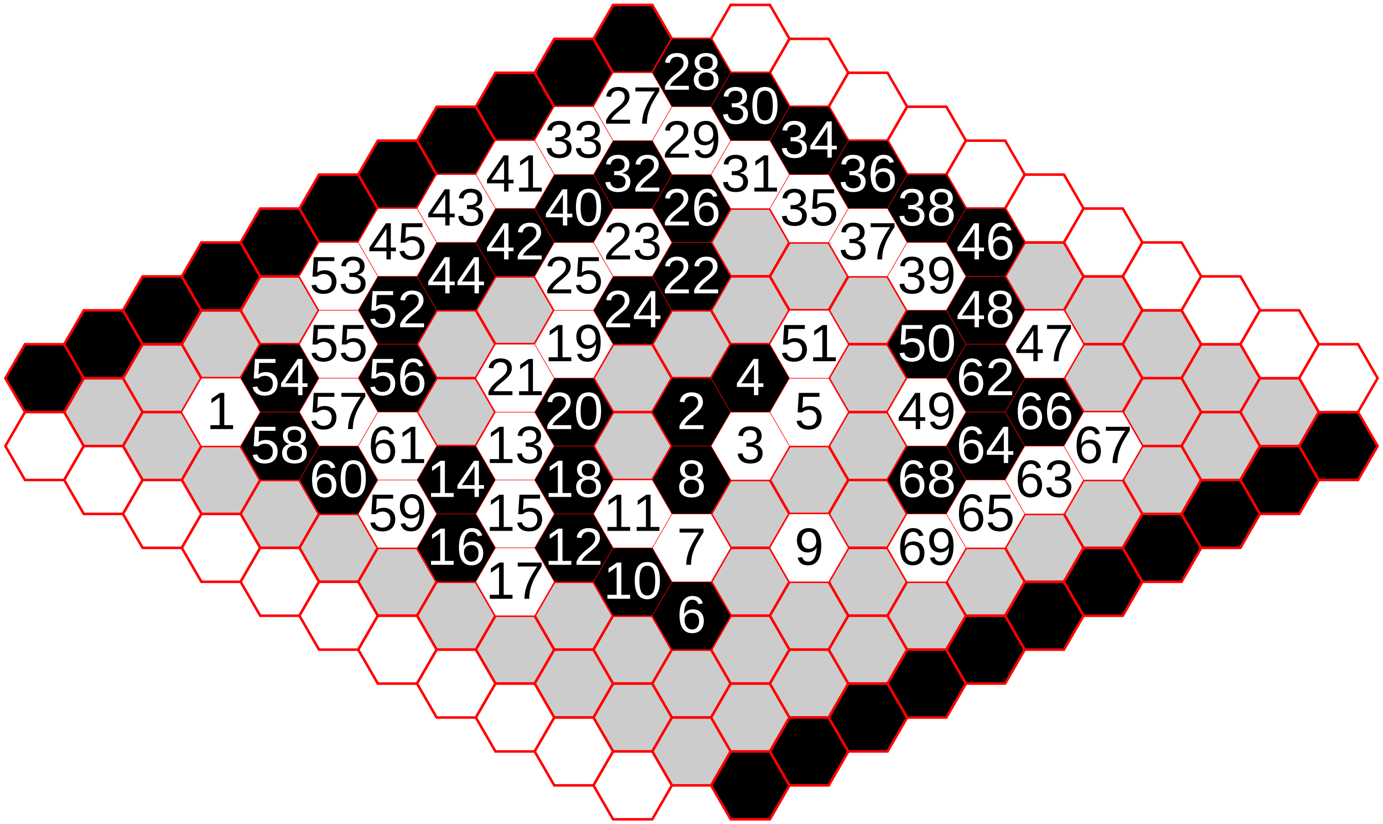

A game from the 5th Computer Olympiad, London, 2000.

Random-turn Hex

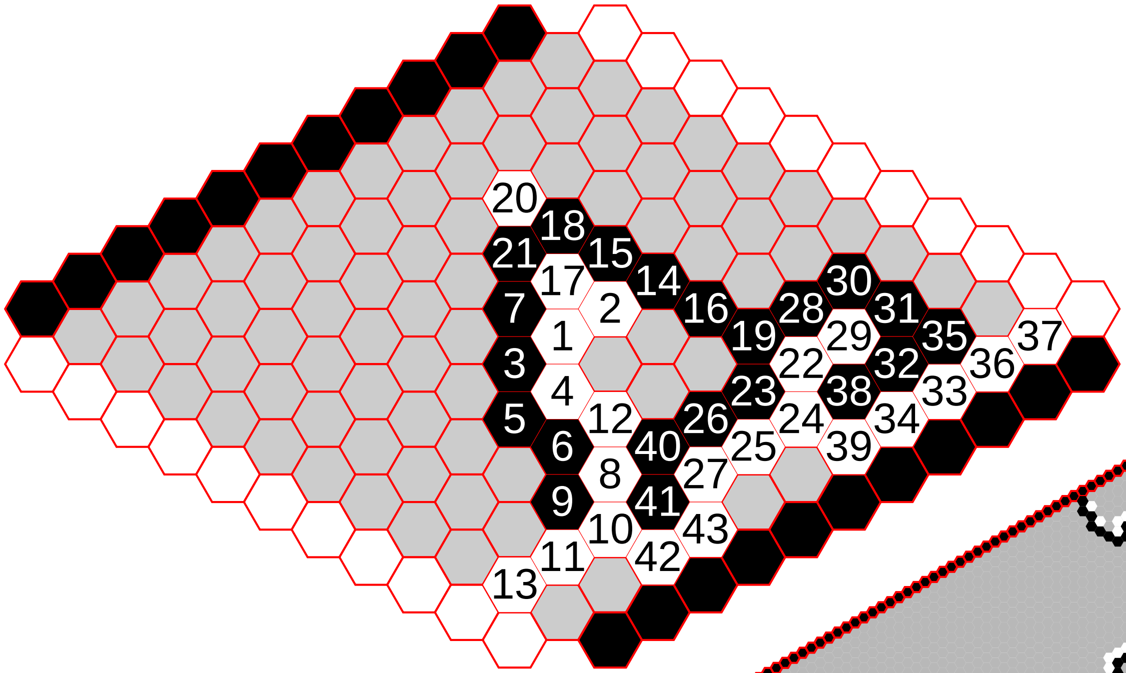

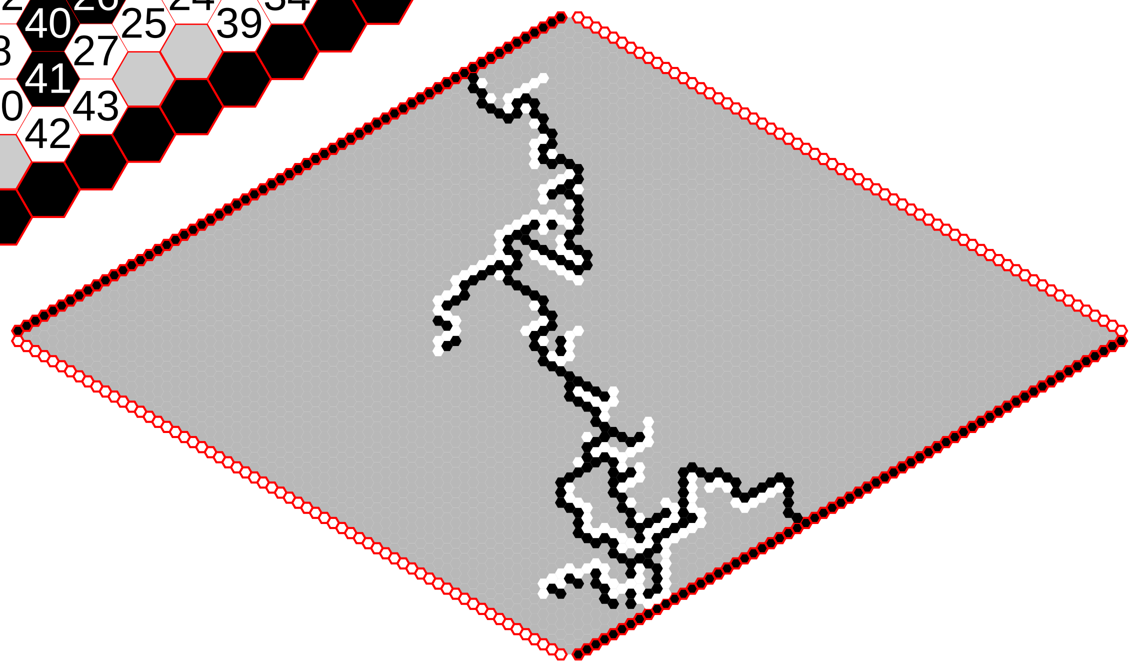

Random-turn selection games in general

For any monotone function $f : \{-1, 1\}^n \longrightarrow \mathbb{R}$, can play random-turn game with payoff f: Maxine and Mina flip a coin, winner colors one of the bits to $+$ or $-$, repeat until function is determined, and Mina pays $f(\omega)$ to Maxine. They want to minimize/maximize expectation.

Full-info finite zero-sum game $\Longrightarrow$ there is an optimal strategy pair

Theorem (Peres, Schramm, Sheffield, Wilson '07).

- For any f, the value of the game is $\mathbf{E}f(\omega)$, iid input!
- Because both players can achieve this by strategy stealing.
- For any monotone $f : \{-1, 1\}^n \longrightarrow \{-1, 1\}$, the optimal strategy for both players: choose the bit that is most likely to be pivotal if the remaining bits are colored randomly.

Zero-sum games

$M(S_-, S_+) := \mathbf{E}[\mathrm{Pay}(S_-, S_+)]$.

The value for Mina is the best she can do if she has to announce first:

$$V_- := \inf_{S_- \in \mathcal{S}_-} \sup_{S_+ \in \mathcal{S}_+} M(S_-, S_+)\,.$$

Similarly for Maxine:

$$V_+ := \sup_{S_+ \in \mathcal{S}_+} \inf_{S_- \in \mathcal{S}_-} M(S_-, S_+)\,.$$

Note that $V_+ \leqslant V_-$ always. The game has a value if $V_+ = V_-$.

A pair (S_-, S_+) is a Nash equilibrium if, for any $S'_- \in \mathcal{S}_-$ and $S'_+ \in \mathcal{S}_+$,

$$M(S_-, S'_+) \leqslant M(S_-, S_+) \leqslant M(S'_-, S_+)\,.$$

At any Nash equilibrium, the payoff is the value of the game.

Number of steps in random-turn games

Optimal play asks ω_J, $j \subseteq [n]$.

$M_i := \mathbf{E}\big[\, f(\omega) \,\big|\, \omega_1, \dots, \omega_i \,\big]$ is a martingale, $i = 0, 1, \dots, |J|$.

Pythagorean theorem for martingales:

$$\mathrm{Var} f(\omega) = \mathrm{Var}\big(M_{|J|}\big) = \mathrm{Var} M_1 + \mathrm{Var}\big(M_2 - M_1\big) + \cdots + \mathrm{Var}\big(M_{|J|} - M_{|J|-1}\big)$$

Optimal play minimizes remaining variance in every step: some sort of *greedy algorithm* to finish the game as soon possible.

When is it optimal for the number of steps? Nobody knows.

Hard to analyze: in random-turn hex, seems $n^{\approx 1.6}$, but no $n^{2-\epsilon}$ is known.

In iterated 3-majority, it is *not* always optimal (Mátyás Susits, BSc '19).

Lower bounds on length come from discrete Fourier analysis.

Random tug-of-war

Given a graph $G(V,E)$, with boundary function $f : \partial V \longrightarrow \mathbb{R}$. A token starts at $X_0 = v \in V \setminus \partial V$. Maxine and Mina flip a coin, the winner can move from X_i to a neighbour X_{i+1}. Game ends when $X_F \in \partial V$, and Mina pays $f(X_F)$ to Maxine.

If the token never reaches ∂V, then Mina pays some fixed amount f_∞ to Maxine. They may use extra randomness in their choices.

Given any pair of strategies, (S_-, S_+), Mina wants to minimize $M(S_-, S_+) := \mathbf{E}[\mathrm{Pay}(S_-, S_+)]$, Maxine wants to maximize it.

Theorem (PSSW '09). On any finite graph, the value of the game exists, and it is the discrete ∞-harmonic extension $h(v) = \frac{1}{2}\big(\max_{w\sim v} h(w) + \min_{w\sim v} h(w)\big)$. In the optimal strategy pair, Maxine wants to move to v_+, Mina to v_-.

On reasonable domains in $\mathbb{R}^d$, with neighbouring relation given by dist $\leqslant \delta$, the value function exists again, and for $\delta \to 0$, it converges to the unique continuum ∞-harmonic extension.

Random tug-of-war

On finite graphs, there is a fast algorithm to compute the discrete ∞-harmonic function: find $x, y \in \partial V$ with *maximal slope* $\frac{f(y)-f(x)}{d(x,y)}$, take linear extension on a graph geodesic, add this segment to the boundary.

Extension by Peres-Šunić '19 for λ-biased discrete ∞-harmonic function

$$h(v) = \frac{\lambda}{\lambda+1} \max_{w \sim v} h(w) + \frac{1}{\lambda+1} \min_{w \sim v} h(w) :$$

define the λ-slope between x and y by

$$\frac{f(y) - \lambda^{-d(x,y)} f(x)}{1 + \lambda^{-1} + \cdots + \lambda^{1-d(x,y)}},$$

then do the same iterative procedure of extensions along line segments.

Extension of Euclidean $\delta \to 0$ result for a biased coin $\frac{1}{2} \pm \lambda\delta$ by Peres-P-Somersille '10.

Stake-governed random tug-of-war

Token starts at $X_0 = v \in V \setminus \partial V$.

Mina has fortune 1, Maxine has λ.

Before each turn, Maxine stakes $a \in [0, \lambda]$, Mina stakes $b \in [0, 1]$.

Maxine wins coin flip with probability $\frac{a}{a+b}$.

Then she moves the token to a neighbour $X_{i+1} \sim X_i$ of her choice.

The new fortunes are 1 for Mina, and $\frac{\lambda - a}{1-b}$ for Maxine.

Game ends when $X_F \in \partial V$.

Mina pays $f(X_F)$ to Maxine (or some fixed f_∞ if $F = \infty$). The remaining fortunes are irrelevant.

Game value in a dream world

Proposition. On any finite graph, if $\Delta(\lambda, v) := \max_{w\sim v} h(\lambda, w) - \min_{w\sim v} h(\lambda, w) > 0$ for every inner v, and there is a pure Nash equilibrium, then the game value exists, and it is the λ-biased ∞-harmonic function $h(v) = h(\lambda, v)$.

Proof (stake strategy stealing). Let (S_-^0, S_+^0) be an equilibrium. Now let S_- be the strategy for Mina where she stakes the same proportion as Maxine in S_+^0, and wants to move to an h-minimizing neighbour. The resulting $h(X_i)$ is a bounded super-martingale, whose limiting value is the payoff, because the game ends a.s. in finite time. Thus, $M(S_-, S_+^0) \leqslant h(X_0) = h(v)$. On the other hand, $M(S_-^0, S_+^0) \leqslant M(S_-, S_+^0)$ because of being an equilibrium.

Similar strategy for Maxine gives $M(S_-^0, S_+^0) \geqslant h(v)$. $\square$

Still the questions: Is there a pure Nash equilibrium?
If so, what are the optimal deterministic stake amounts?

First guess for the optimal stake

Instead of optimal $(\lambda S, S)$, assume that Maxine stakes $\lambda(S+\eta)$. Increased probability of winning this turn:

$$\frac{\lambda(S+\eta)}{\lambda(S+\eta)+S} - \frac{\lambda}{\lambda+1} = \eta S^{-1}\tfrac{\lambda}{(1+\lambda)^2} + O(\eta^2)\,.$$

However, smaller fortune for the future: $\lambda_{\text{alt}} = \frac{\lambda-\lambda(S+\eta)}{1-S} = \lambda - \lambda(1-S)^{-1}\eta\,,$ which reduces the chance of winning each future step by

$$\lambda/(1+\lambda) - \lambda_{\text{alt}}/(1+\lambda_{\text{alt}}) = \eta(1-S)^{-1}\tfrac{\lambda}{(1+\lambda)^2} + O(\eta^2)\,.$$

Expected gain and loss, at order η, should balance out:

$$S^{-1}\tfrac{\lambda}{(1+\lambda)^2}\cdot\Delta(\lambda,v) = (1-S)^{-1}\tfrac{\lambda}{(1+\lambda)^2}\cdot\mathbf{E}\sum_{i=1}^{F-1}\Delta(\lambda,X_i)\,,$$

hence the optimal stake could be (based on this very local analysis!):

$$S = \frac{\Delta(\lambda,v)}{\mathbf{E}\sum_{i=0}^{F-1}\Delta(\lambda,X_i)}\,.$$

Ouch: no pure Nash equilibrium!

Path $\{0, 1, 2, 3\}$, payoff 1 at vertex 3 and 0 at vertex 0.
Maxine $\lambda = 1/2$. Non-lazy $\epsilon = 1$. Start at vertex 2.

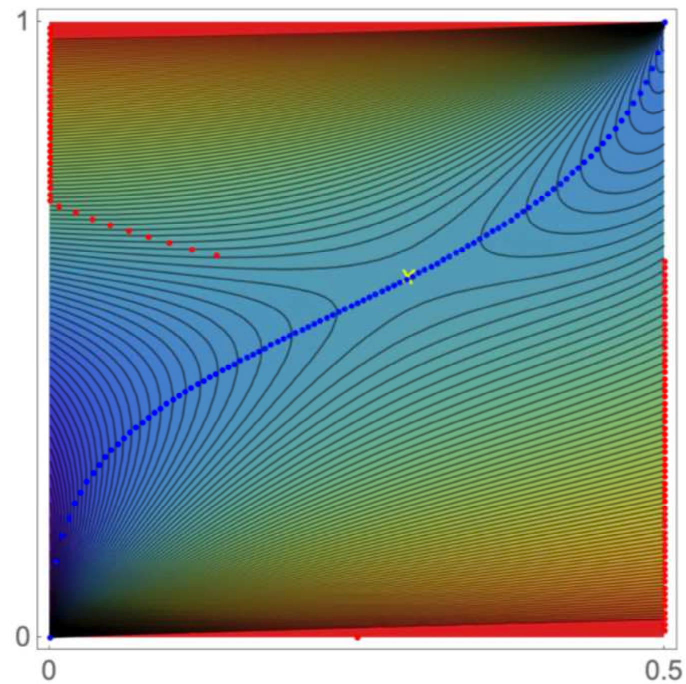

Assuming that the value is always the current $h(\lambda, v)$, this would be the value given Maxine's first-turn stake $a \in [0, 1/2]$ and Mina's $b \in [0, 1]$.

Blue dots are Mina's best responses to stakes a of Maxine, red dots are Maxine's best responses to stakes b of Mina. No global saddle point.

Introducing laziness, and the Poisson game

Maxine's previous go-for-broke strategy becomes nonsense if we introduce laziness: after the stakes are made, a move takes place only with a small probability $\epsilon > 0$. But the stakes are always deducted.

Extreme version, the Poisson game: in continuous time, the stakes are $a(t)$ and $b(t)$, measurable w.r.t. everything before time t, new fortune is $\lambda(t + \mathrm{d}t) = \frac{\lambda(t) - a(t)\mathrm{d}t}{1 - b(t)\mathrm{d}t}$, moves happen at Poisson times.

Laziness helps, and randomizing the stakes seems to make less sense, so it seems plausible that pure Nash equilibria exist, and hence the value of the game is $h(\lambda, v)$.

However, the game is hard to define properly. . . :-)

Second guess for the optimal stake

Nevertheless, the "Poisson game" suggests a second formula for the stake value. Starting at vertex v at time 0, if $a(t) = a$ and $b(t) = b$ for $t \in [0, \mathrm{d}t]$, then the value of the game at dt, written as $h(\lambda, v) + \Phi(a, b)\mathrm{d}t$, is

$$(1 - \mathrm{d}t)h\big(\lambda(\mathrm{d}t), v\big) + \mathrm{d}t\tfrac{a}{a+b}h\big(\lambda(\mathrm{d}t), v_+\big) + \mathrm{d}t\tfrac{b}{a+b}h\big(\lambda(\mathrm{d}t), v_-\big)\,,$$

where $v_\pm$ are maximizer/minimizer neighbours of v for $h(\lambda, \cdot)$. Rearranging,

$$\Phi(a, b) = -h(\lambda, v) - (a - b\lambda)h'(\lambda, v) + \tfrac{a}{a+b}h(\lambda, v_+) + \tfrac{b}{a+b}h(\lambda, v_-).$$

If (a_0, b_0) are the stakes in a Nash equilibrium, then should have $\Phi(a_0, b_0) = 0$, and $\frac{\partial}{\partial a}\Phi(a_0, b_0) = \frac{\partial}{\partial b}\Phi(a_0, b_0) = 0$, and $\frac{\partial^2}{\partial a^2}\Phi(a_0, b_0) < 0 < \frac{\partial^2}{\partial b^2}\Phi(a_0, b_0)$.

First of these gives $a_0 = b_0\lambda$, the first derivatives give

$$b_0 = \frac{h\big(\lambda, v_+\big) - h\big(\lambda, v_-\big)}{(\lambda+1)^2 h'(\lambda, v)} = \frac{\Delta(\lambda, v)}{(\lambda+1)^2 h'(\lambda, v)}\,,$$

and the second derivatives have the right signs for all a, b: a global saddle!

Even bigger ouch: not just the stakes but also the moves could be random

The previous calculations make little sense if $h(\lambda, v)$ changes drastically at some λ: the neighbours $v_{\pm}(\lambda)$ could depend on λ, and $h(\lambda, v)$ may not be differentiable in λ.

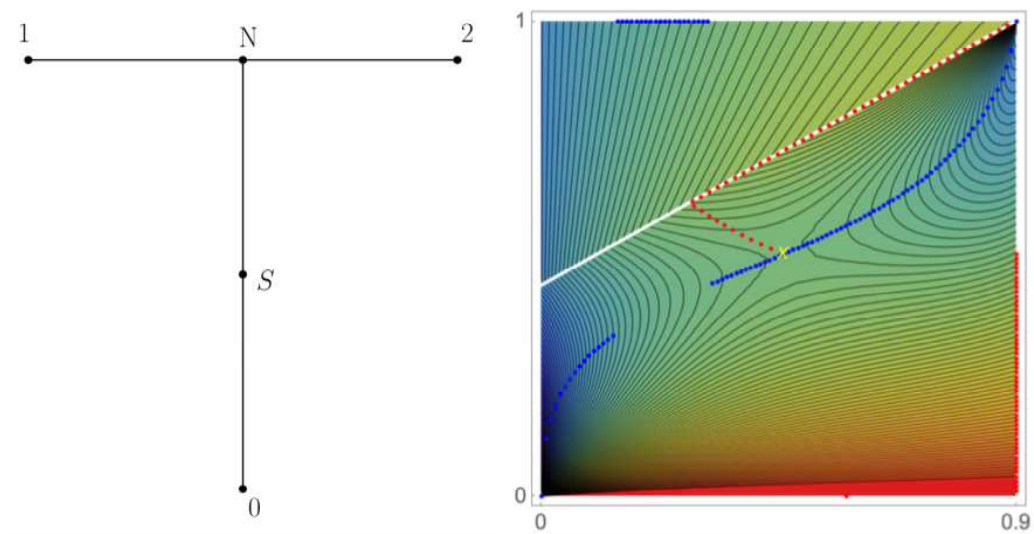

On this T-graph, the Peres-Šunić decomposition changes at $\lambda_c = \frac{\sqrt{5}+1}{2}$, giving exotic value plots for most λ values.

Root-reward trees

The leafs are the boundary vertices. One leaf, the root r, has payoff 1, all other leafs have payoff 0.

Not hard to see that the Peres-Šunić decomposition into basic trees is independent of λ: the first basic tree is the one with the smallest possible diameter, then iterate.

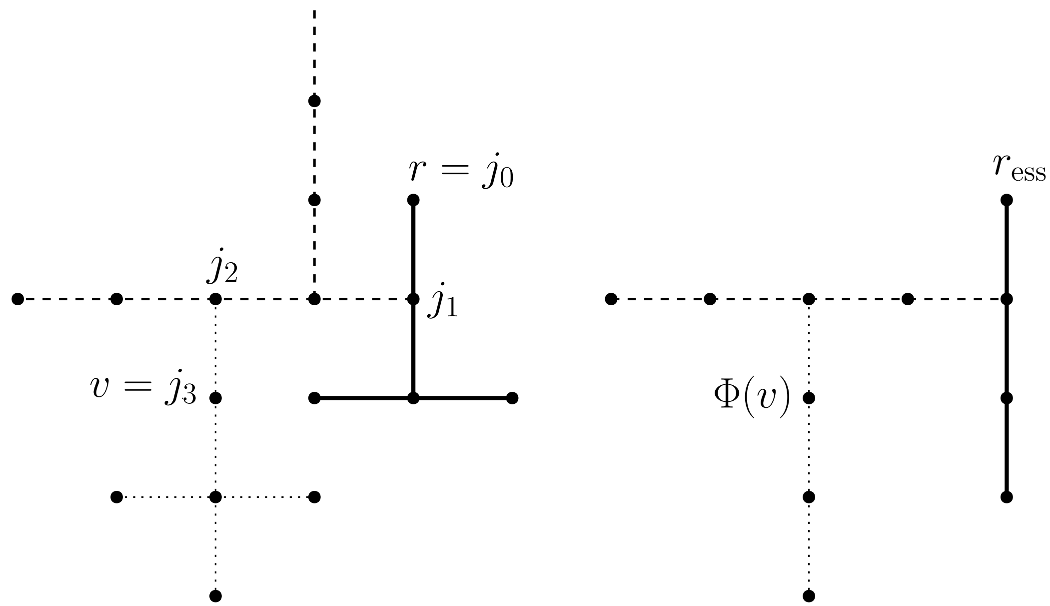

From any tree, can produce an essence tree.

$h(\lambda, v)$ is a product of λ-biased ∞-harmonic functions on segments.

Game value and Nash equilibria on root-reward trees

Theorem (Hammond & P). On any root-reward tree, any compact $K \subset (0, \infty)$, there is $\epsilon_K > 0$ s.t. for any fortune $\lambda \in K$ and $0 < \epsilon < \epsilon_K$, in the ϵ-lazy game started at any vertex v, the value is $h(\lambda, v)$, every Nash equilibrium consists of $h(\lambda, \cdot)$-maximizing/minimizing moves, and the stake values are $(\lambda S, S)$ with

$$S = \frac{\epsilon\,\Delta(\lambda, v)}{(\lambda+1)^2 \frac{\partial}{\partial \lambda} h(\lambda, v)} = \frac{\epsilon\,\Delta(\lambda, v)}{\mathbf{E}\,\mathrm{TotVar}(1, \lambda, v)} = \frac{\Delta(\lambda, v)}{\mathbf{E}\,\mathrm{TotVar}(\epsilon, \lambda, v)},$$

where $\mathrm{TotVar}(\epsilon, \lambda, v) = \sum_{i=0}^{F-1} \Delta\big(\lambda, X_i\big)$ with $X_0 = v$.

Game value and Nash equilibria on root-reward trees

Theorem (Hammond & P). On any root-reward tree, any compact $K \subset (0, \infty)$, there is $\epsilon_K > 0$ s.t. for any fortune $\lambda \in K$ and $0 < \epsilon < \epsilon_K$, in the ϵ-lazy game started at any vertex v, the value is $h(\lambda, v)$, every Nash equilibrium consists of $h(\lambda, \cdot)$-maximizing/minimizing moves, and the stake values are $(\lambda S, S)$ with

$$S = \frac{\epsilon\,\Delta(\lambda, v)}{(\lambda+1)^2 \frac{\partial}{\partial\lambda} h(\lambda, v)} = \frac{\epsilon\,\Delta(\lambda, v)}{\mathbf{E}\,\mathrm{TotVar}(1, \lambda, v)} = \frac{\Delta(\lambda, v)}{\mathbf{E}\,\mathrm{TotVar}(\epsilon, \lambda, v)},$$

where $\mathrm{TotVar}(\epsilon, \lambda, v) = \sum_{i=0}^{F-1} \Delta\big(\lambda, X_i\big)$ with $X_0 = v$.

Caveat. Here the payoff for infinite play is $f_\infty = 1$. This is important:

On the long T-graph $\{0, 1, \ldots, n-1, n, n^*\}$, large λ, started at $n-1$, if Maxine plays as dictated, a stake of order λ^{3-n}, but Mina always stakes $1/2$ at $n-1$, then λ goes up exponentially, and the game lasts forever with positive probability.

Open problems

1. For what monotone functions is playing the random-turn game optimally also an optimal low revealment algorithm?

2. How do stake-governed random-turn selection games look like?

3. Define properly and analyse the stake-governed Poisson tug-of-war on finite graphs.

4. Study stake-governed Euclidean tug-of-war with fixed small $\delta > 0$ and in the limit $\delta \to 0$.