Playing Games with Counter Automata

Antonín Kučera

IFIP WG 2.2 Meeting, Amsterdam, September 2012

Aims

- Model discrete systems with multiple resources that can be consumed, produced, or “reloaded”.
- User control over the system is only partial, some states are adversarial or stochastic.
- Stochastic multi-counter games.
- Questions of interest:
  - Can a user control the system so that the resources are never exhausted?
  - What is the minimal “storage capacity” for achieving the above objective?
  - Many others. . .
- Can we answer such questions effectively/efficiently ?

Multi-counter stochastic games

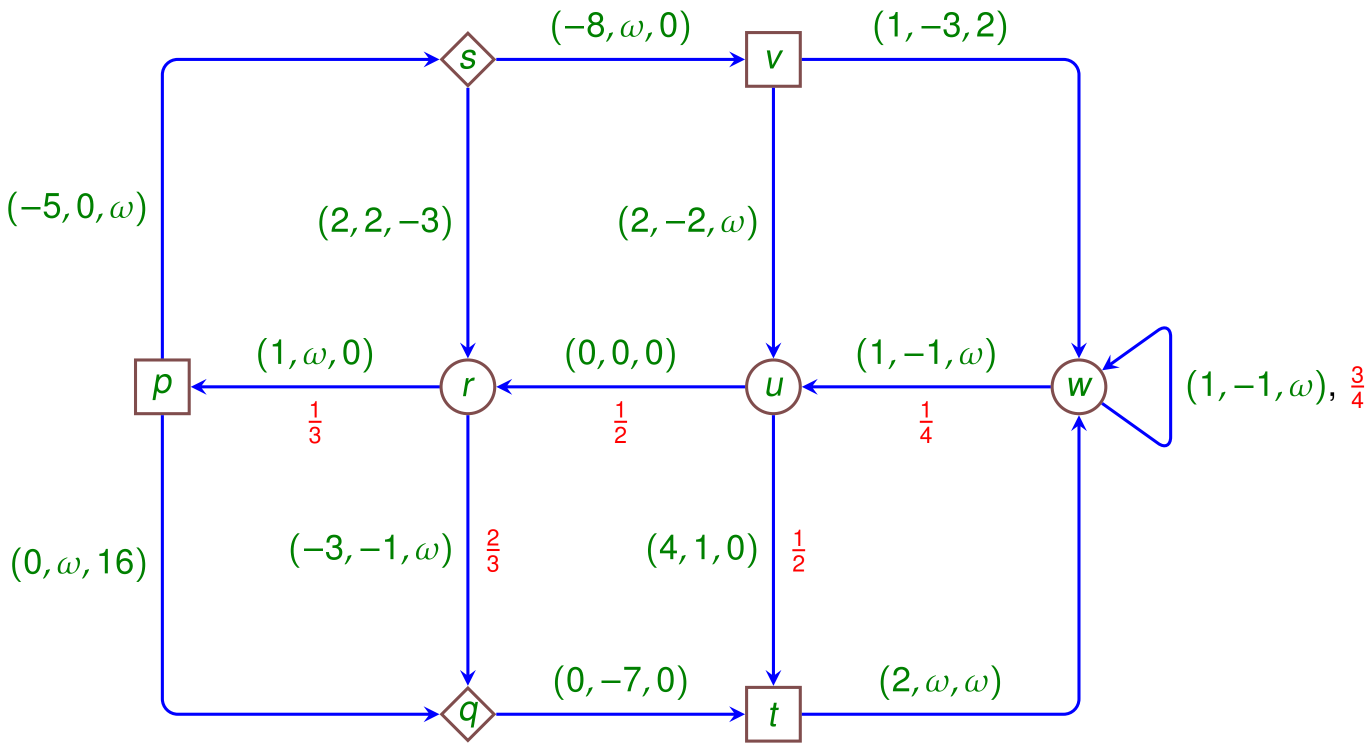

The semantics of multi-counter stochastic games

- Each MCSG determines an infinite game graph which can be infinitely-branching.
- Vertices are configurations of the form $p\vec{v}$.
- $\mathbb{Z}$-semantics: $\vec{v} \in \mathbb{Z}^k$, non-blocking transitions.
- $\mathbb{N}$-semantics: $\vec{v} \in \mathbb{N}^k$, blocking transitions (Petri net style).

An example

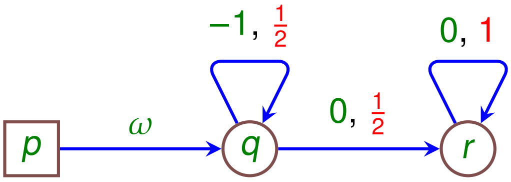

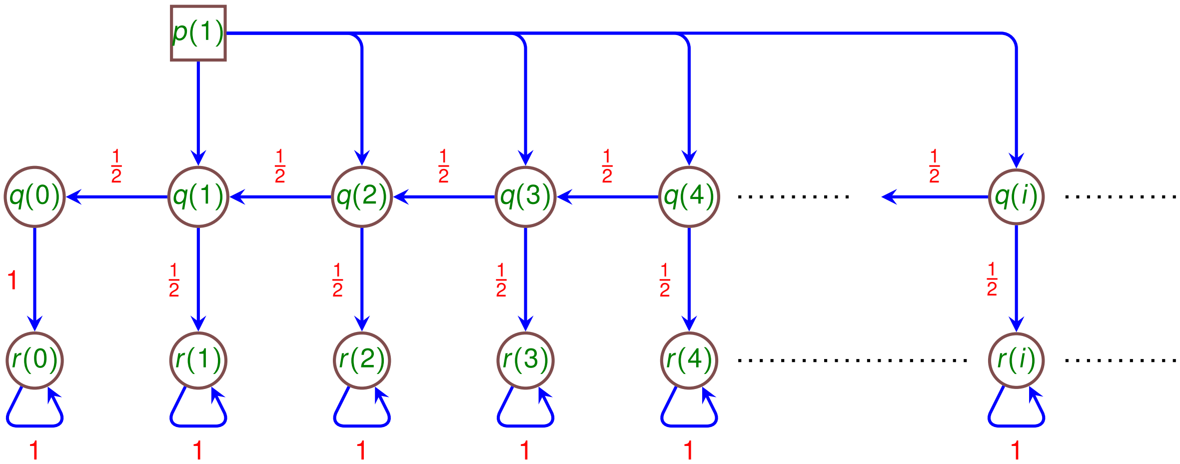

Winning objectives in stochastic games (1)

- We deal with infinite-state and possibly infinitely-branching stochastic games.
- Some of very basic properties of finite-state games do not carry over to infinite-state games.
- In particular, this applies to the existence and type of optimal/winning strategies, even for simple objectives such as reachability or safety.

Winning objectives in stochastic games (2)

- Let f be a function assigning a real-valued payoff to every run if a given game graph G.
- The aim of Player $\Box/\Diamond$ is to maximize/minimize $\mathbb{E}[f]$.
- If f is Borel and bounded, then every vertex v has a value $val(v)$ given by

$$\sup_{\sigma\in\mathrm{HR}_\Box}\ \inf_{\pi\in\mathrm{HR}_\Diamond}\ \mathbb{E}[f] \quad = \quad \inf_{\pi\in\mathrm{HR}_\Diamond}\ \sup_{\sigma\in\mathrm{HR}_\Box}\ \mathbb{E}[f]$$

  This holds also for infinitely-branching games.
- The existence of $val(v)$ induces the notions of optimal and ε-optimal strategy.

Reachability objective in stochastic games

- The aim of Player $\Box/\Diamond$ is to maximize/minimize the probability of visiting a target vertex.
- Let $Reach$ be a function which to every run assigns either 1 or 0 depending on whether or not the run visits a target vertex.
- In finite-state stochastic games, both players have optimal MD strategies. Hence, finite-state reachability games are determined for every strategy class C that subsumes MD strategies:

$$\sup_{\sigma\in C_\Box}\ \inf_{\pi\in C_\Diamond}\ \mathbb{E}[Reach] \quad = \quad \inf_{\pi\in C_\Diamond}\ \sup_{\sigma\in C_\Box}\ \mathbb{E}[Reach]$$

Reachability in infinite-state stochastic games

In infinitely-branching games:

- Optimal strategies do not necessarily exist.
- Even if they do exist, they may require infinite memory.
- The games are not determined for finite-memory strategies.

In finitely-branching games:

- An optimal strategy for Player □ (Max) does not necessarily exist.
- Player ◇ (Min) has an optimal MD strategy.
- The games are determined for any strategy type that subsumes MD strategies, and the value is the same.

Similar results hold also for the (unbounded!) total accumulated reward payoff function (Brázdil, K, Novotný; MEMICS 2012).

Counterexamples (1)

Player □ (Max) does not necessarily have an optimal strategy, even in finitely-branching MDPs.

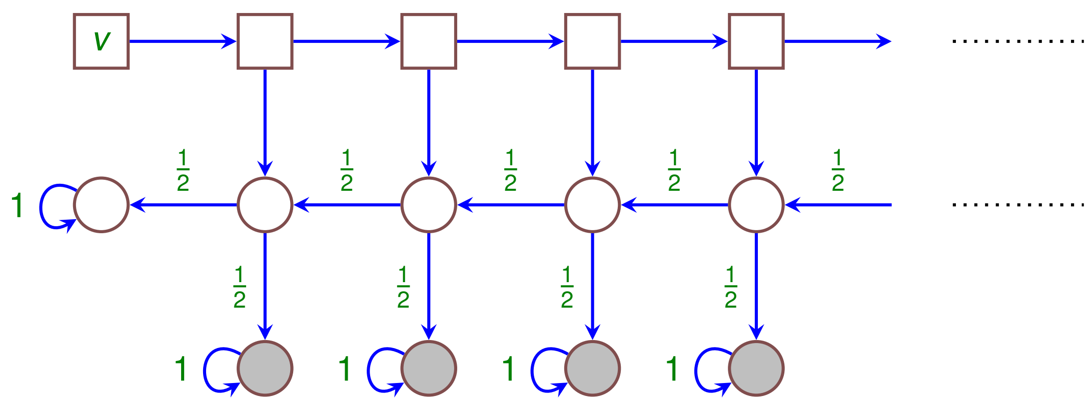

Counterexamples (2)

An optimal strategy for Player □ (Max) may require infinite memory, even in finitely-branching games.

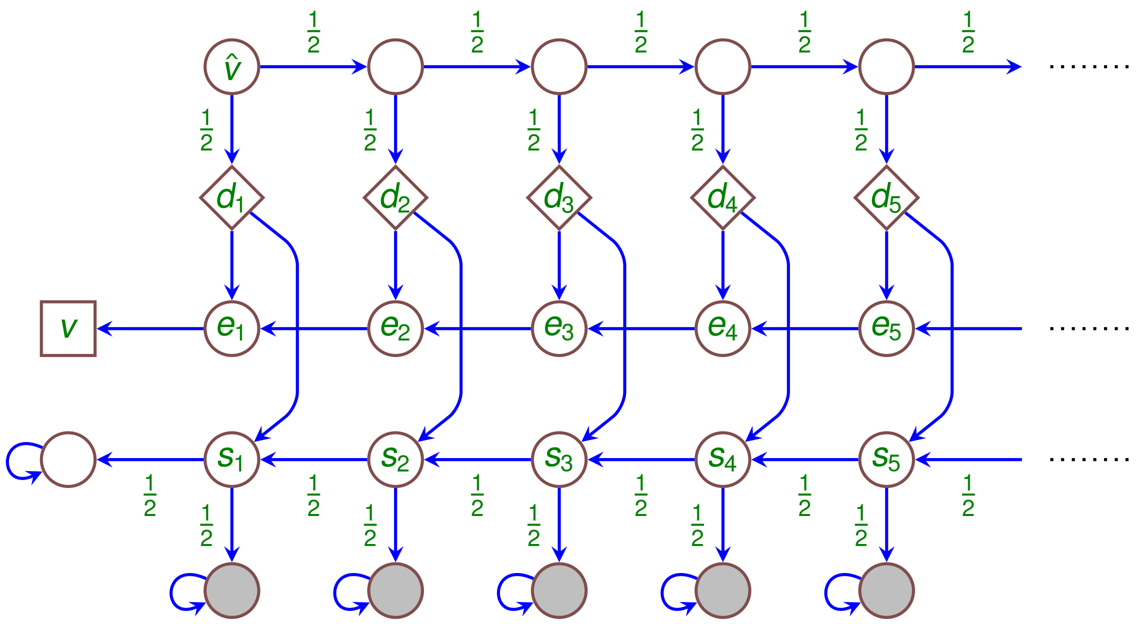

Counterexamples (3)

Optimal minimizing strategies do not necessarily exist, and (ε-) optimal minimizing strategies may require infinite memory.

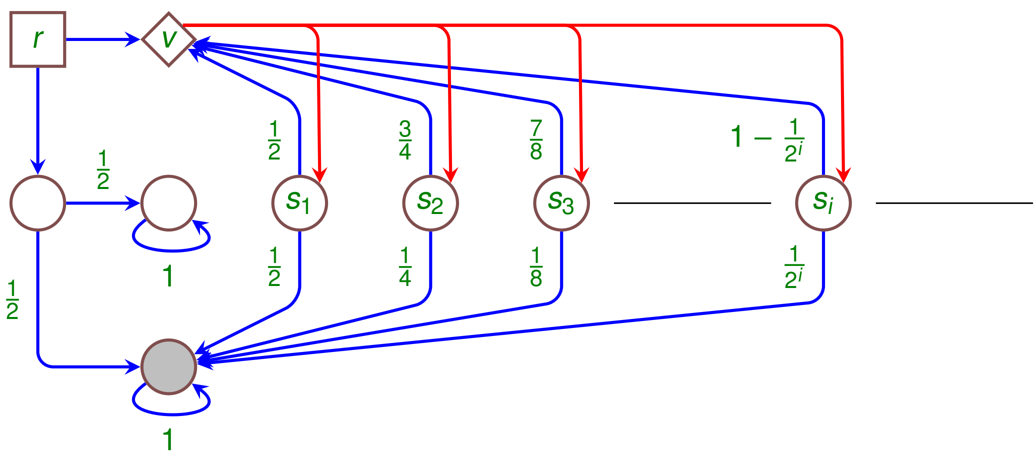

Counterexamples (4)

Infinitely-branching games are not determined for finite-memory strategies.

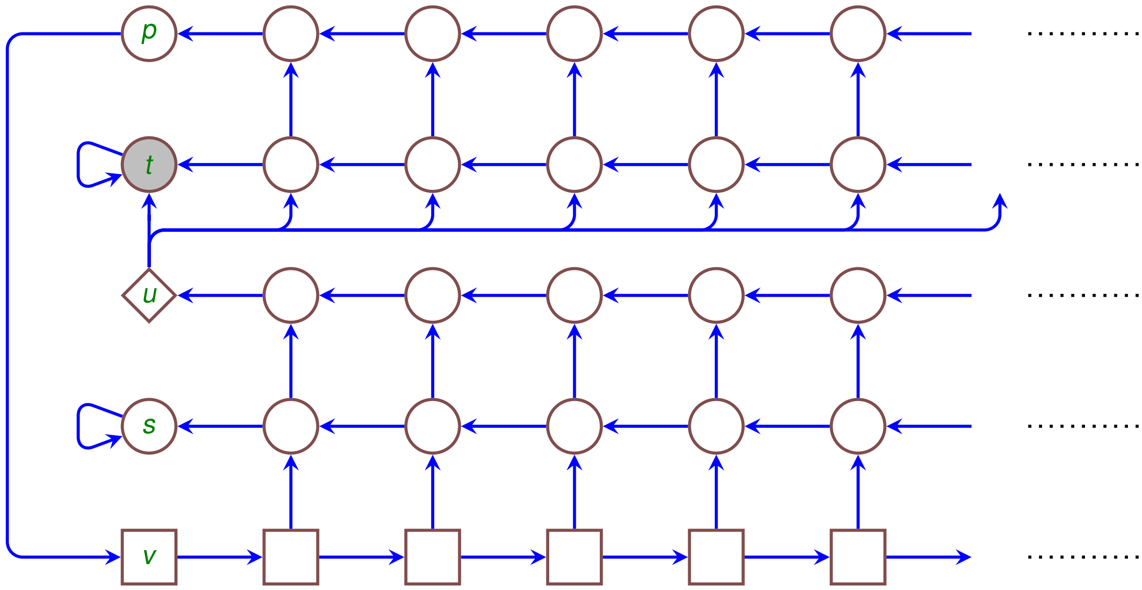

- $\sup_{\sigma \in \mathrm{MD}_\Box} \inf_{\pi \in \mathrm{MD}_\Diamond} \mathbb{E}[Reach] = 0$
- $\inf_{\pi \in \mathrm{MD}_\Diamond} \sup_{\sigma \in \mathrm{MD}_\Box} \mathbb{E}[Reach] = 1$

Existing results about multi-counter games

- Only for restricted models:
  - non-stochastic games with multiple counters;
  - stochastic games with only one counter.
- Objectives:
  - zero-safety
  - (selective) zero-reachability
  - expected zero-termination time

Non-stochastic eVASS games

- Let $\mathcal{G}$ be an eVASS game, i.e., a non-stochastic multi-counter game with k counters where update vectors belong to $\{0, 1, -1, \omega\}^k$. We consider $\mathbb{N}$-semantics.
- The aim of Player $\Box$ is to avoid visiting configurations with zero in some counter. Player $\Diamond$ aims at the opposite.
- Let *safe* be the set of all configurations where Player $\Box$ has a winning strategy.

Non-stochastic eVASS games (2)

Observation 1

The set safe is upwards-closed w.r.t. component-wise ordering. Hence, it is fully characterized by its finitely many minimal elements. Consequently, the membership to safe is semi-decidable.

Observation 2

Suppose that the update vectors of $\mathcal{G}$ do not contain any ω's. Then the game graph of $\mathcal{G}$ is finitely-branching, and a winning strategy for Player $\Diamond$ can be encoded by a finite tree. Consequently, the non-membership to safe is semi-decidable.

Non-stochastic eVASS games (2)

Theorem 3 (Brázdil, Jančar, K.; Icalp 2010)

*The set of minimal elements of safe is computable in $k-1$-**EXPTIME**. The membership problem for safe is **EXPSPACE**-hard.*

- A symbolic configuration is a pair $p\vec{v}$ where $\vec{v}_i \in \mathbb{N} \cup \{*\}$. The precision of $p\vec{v}$ is the number of non-$*$ elements in $\vec{v}$.
- An instance of a symbolic configuration $p\vec{v}$ is obtained by substituting all $*$-elements in $\vec{v}$ with concrete values. The component-wise ordering is extended to symbolic configurations by stipulating $n < *$ for all $n \in \mathbb{N}$.
- For $j = 0, 1, 2, \ldots, k$, we inductively compute
  - the set C_j of all minimal symbolic configurations $p\vec{v}$ of precision j such that some instance of $p\vec{v}$ is safe.
  - a bound B_j such that for every $p\vec{v} \in C_j$ we have that the instance obtained by substituting every $*$ in $\vec{v}$ with B_j is safe.
- Clearly, C_k is the set of all minimal safe configurations.

Non-stochastic eVASS games (3)

Further remarks.

- For $k = 2$ and no ω components, the complexity has been improved from **EXPTIME** to **P** by Chaloupka (RP 2010).
- The exact complexity classification for a fixed number of counters is open.

Consumption games

- Let $\mathcal{G}$ be an consumption game, i.e., a non-stochastic multi-counter game with k counters where the components of update vectors are non-positive integers or ω's.
- Let S be the set of states of $\mathcal{G}$, and ℓ the maximal absolute value of a counter update.

Theorem 4 (Brázdil, Chatterjee, K., Novotný; CAV 2012)

*The problem whether $p\vec{v} \in safe$ is **PSPACE**-hard and solvable in time $|\vec{v}| \cdot (k \cdot \ell \cdot |S|)^{O(k)}$. Further, the set of minimal elements of $safe$ is computable in time $(k \cdot \ell \cdot |S|)^{O(k)}$.*

Consumption games (2)

- Let $\mathcal{G}$ be an consumption game.
- Let $cover$ be the set of all $p\vec{v}$ where Player $\Box$ can play so that
  - all counters stay positive;
  - every visited configuration $q\vec{u}$ satisfies $\vec{u} \leq \vec{v}$.
- Clearly, $cover \subseteq safe$, and the set $cover$ is upwards closed.

Theorem 5 (Brázdil, Chatterjee, K., Novotný; CAV 2012)

*The problem whether $p\vec{v} \in cover$ is **PSPACE**-hard and solvable in $\mathcal{O}(\Lambda^2 \cdot |S|^2)$ time, where $\Lambda = \Pi_{i=1}^{k} \vec{v}_i$. Further, the set of all minimal elements of $cover$ is computable in $(k \cdot \ell \cdot |S|)^{O(k \cdot k!)}$ time.*

Multiweighted energy games

- Let $\mathcal{G}$ be a multiweighted (or generalized) energy game, i.e., a non-stochastic multicounter game with $\mathbb{Z}$-semantics with no ω-components in update vectors.
- Let $b \in \mathbb{N}$. For every control state p, let $bound_b(p)$ be the set of all $\vec{v} \in \mathbb{N}^k$ where Player $\Box$ can play so that all counters stay non-negative and bounded by b.

Theorem 6 (Fahrenberg, Juhl, Larsen, Srba; ICTAC 2011)

The problem whether $\vec{v} \in bound_b(p)$ is ***EXPTIME****-complete. If all control states belong to Player $\Box$, then the problem is* ***PSPACE****-complete.*

Stochastic games with one counter

Solvency games (Berger, Kapur, Schulman, Vazirani; FST& TCS 2008)

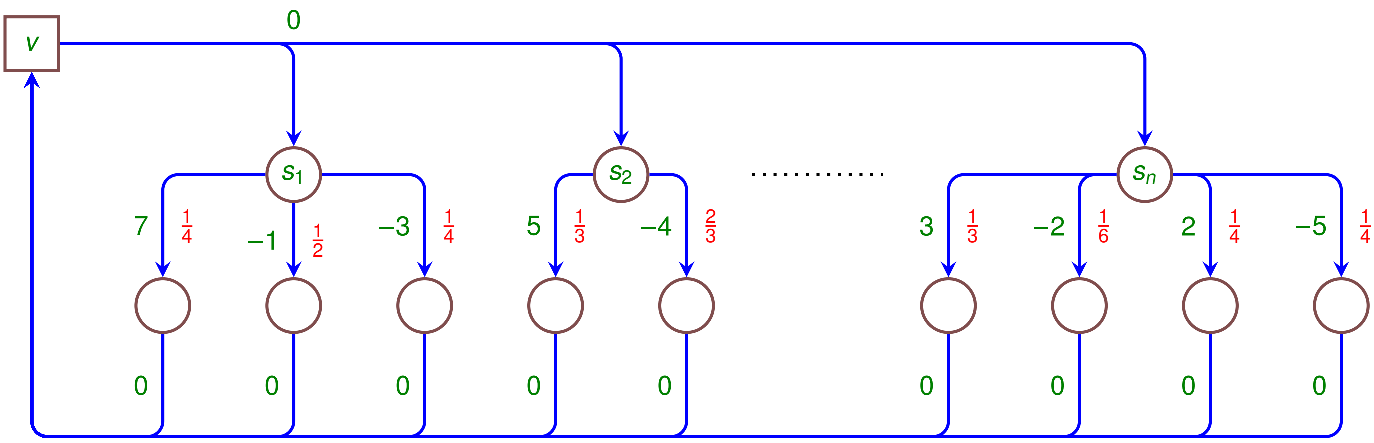

One-counter games and MDPs

- Let $\mathcal{G}$ be a one-counter game, i.e., a multi-counter game with one counter and $\mathbb{Z}$-semantics where the counter updates range over $\{0, 1, -1\}$.
- So far, the following payoff functions have been studied for one-counter games:
  - Cover negatives (CN), which to every run assigns 1 or 0 depending on whether or not $\liminf$ of all counter values visited along the run is equal to $-\infty$.
  - Zero reachability (Z), which to every run assigns 1 or 0 depending on whether or not the run visits a configuration with zero counter.
  - Selective zero reachability (SZ), which to every run assigns 1 or 0 depending on whether or not the run visits a configuration with zero counter in one of the selected control states, and the counter stays positive in all preceding configurations.
  - Termination time (T), which to every run assigns the number of transitions performed before visiting a configuration with zero counter.

One-counter games and MDPs (2)

The results concern the following problems:

- Qualitative questions:
  - Is the value of a given configuration equal to one?
  - Is there a strategy for Player $\Box$ which achieves the outcome one against every strategy of Player $\Diamond$?
- Quantitative questions:
  - Can we compute/approximate the value of a given configuration?
  - Can we compute (ε-optimal) strategies?
- Many answers exist and many questions remain open. The underlying analytical tools are nontrivial.

Open problems

- The tractability frontier for non-stochastic multi-counter games is still not well understood, and only very basic payoff functions have been studied.
- The structure of optimal strategies in one-counter games and MDPs is unclear. Are optimal strategies in solvency games “ultimately periodic”?
- What can be done for stochastic games with multiple counters?