Markov Decision Processes with Energy-Parity Objectives

Richard Mayr, Patrick Totzke,
Sven Schewe, Dominik Wojtczak

Edinburgh/Liverpool, UK

LICS'17
June 22, 2017

MDPs

Finite graphs, partitioned into *controlled* □ and *random* ○ states;
A prob. dist. over successors for every random state.

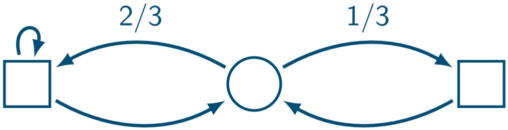

Controller *Strategies*

resolve choice of successor for controlled states to induce a Markov Chain with associated probability space over infinite runs.

The Almost-sure Problem
Does there exist a strategy with $\mathbb{P}^{\sigma}(Obj) = 1$?

Some Objective Functions

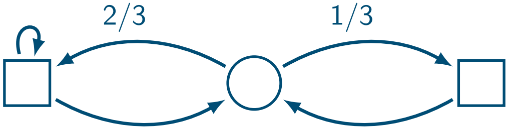

Some Objective Functions

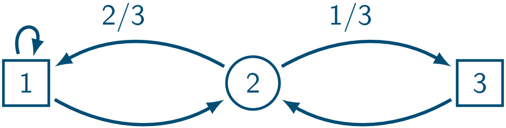

PARITY
maximal colour visited infinitely often is even.

Some Objective Functions

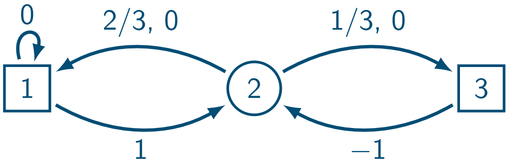

PARITY

maximal colour visited infinitely often is even.

ENERGY

$\forall n \in \mathbb{N}.\ \sum_{i=0}^{n} \mathrm{cost}(e_i) \geq 0$

Some Objective Functions

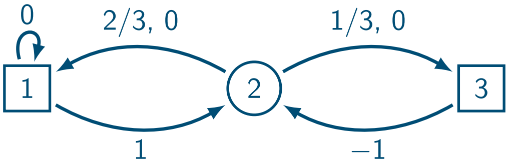

PARITY

maximal colour visited infinitely often is even.

ENERGY

$\forall n \in \mathbb{N}.\ \sum_{i=0}^{n} \mathrm{cost}(e_i) \geq 0$

Positive Mean Payoff

$\lim_{n\to\infty} \sum_{i=0}^{n} \mathrm{cost}(e_i)/n > 0$

Almost-sure Problems for finite MDPs

ENERGY

Almost-sure Problems for finite MDPs

ENERGY

- FD determined
- NP $\cap$ coNP
- pseudo P

Almost-sure Problems for finite MDPs

ENERGY

- FD determined
- NP ∩ coNP
- pseudo P

Positive Mean-Payoff

Almost-sure Problems for finite MDPs

ENERGY

- FD determined
- NP $\cap$ coNP
- pseudo P

Positive Mean-Payoff

- MD determined
- P-time

Almost-sure Problems for finite MDPs

ENERGY

- FD determined
- NP $\cap$ coNP
- pseudo P

Positive Mean-Payoff

- MD determined
- P-time

PARITY

Almost-sure Problems for finite MDPs

ENERGY

- FD determined
- NP $\cap$ coNP
- pseudo P

Positive Mean-Payoff

- MD determined
- P-time

PARITY

- MD determined
- P-time

Almost-sure Problems for finite MDPs

ENERGY

- FD determined
- NP $\cap$ coNP
- pseudo P

Positive Mean-Payoff

- MD determined
- P-time

PARITY

- MD determined
- P-time

Almost-sure Problems for finite MDPs

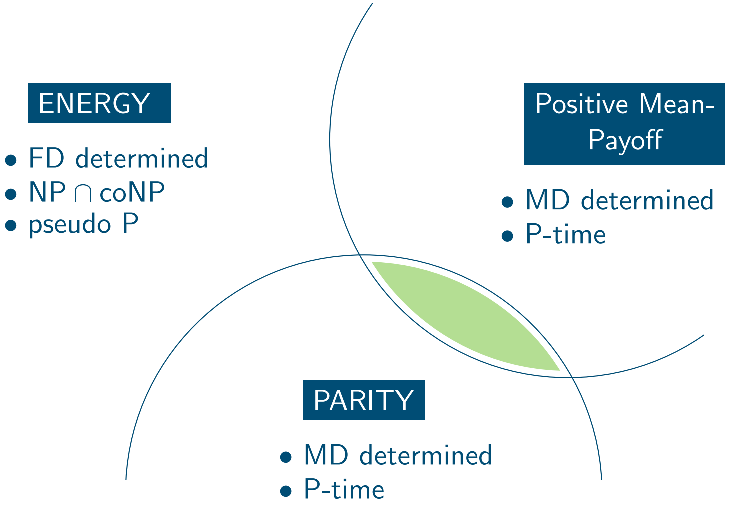

Almost-sure Problems for finite MDPs

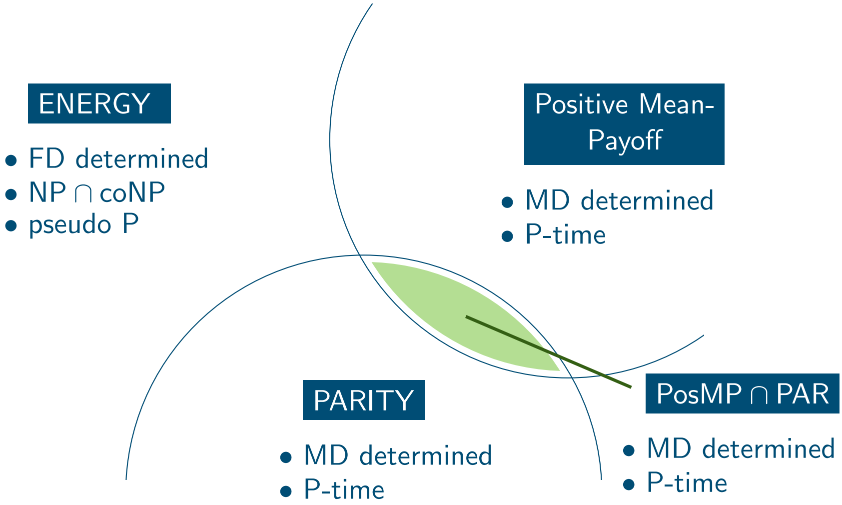

Chatterjee/Doyen, MFCS'11

Almost-sure Problems for finite MDPs

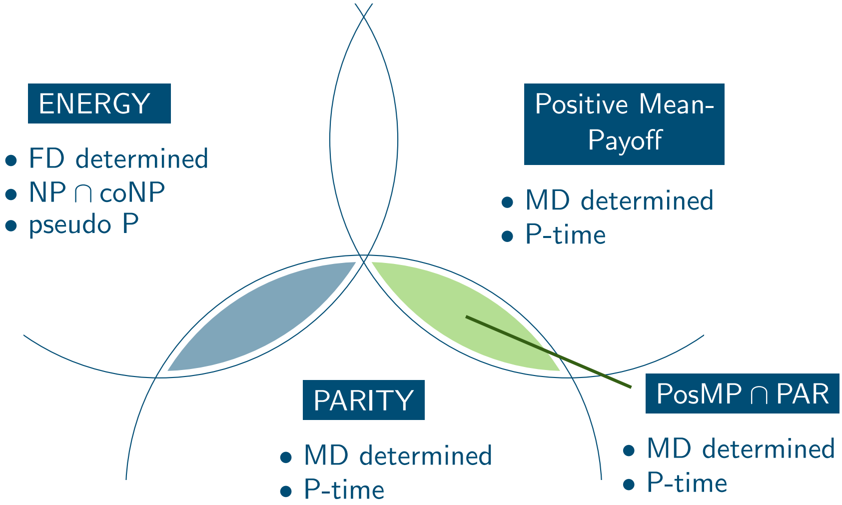

Chatterjee/Doyen, MFCS'11

Almost-sure Problems for finite MDPs

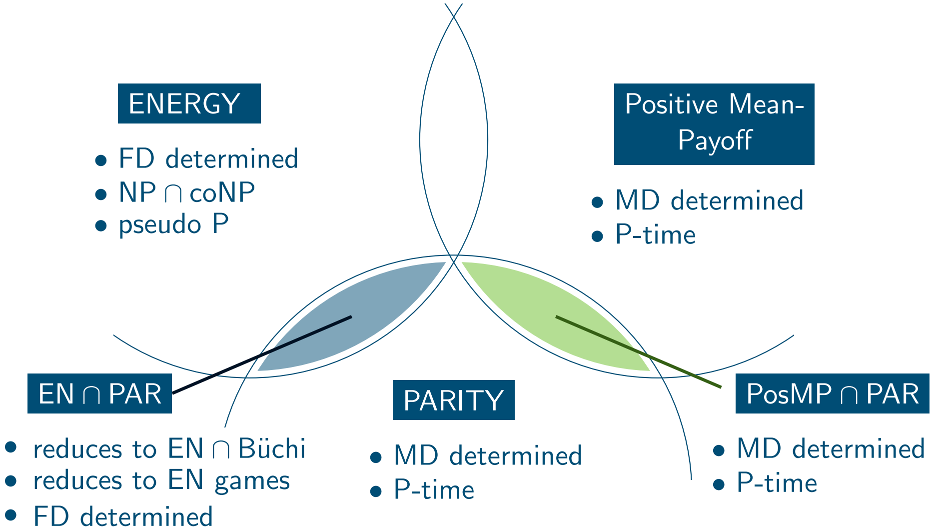

Chatterjee/Doyen, MFCS'11

Almost-sure Problems for finite MDPs

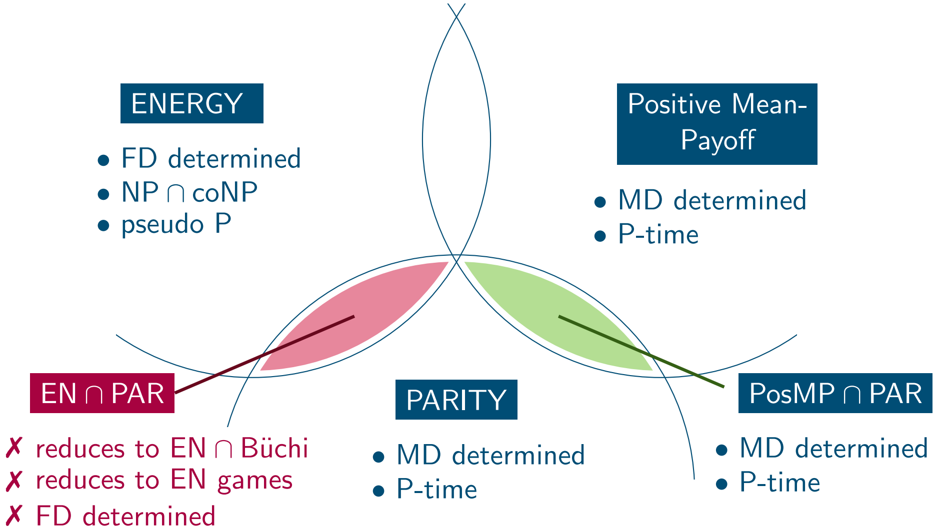

Chatterjee/Doyen, MFCS'11

ENERGY $\cap$ PARITY objectives for finite MDPs:

1. Almost-sure optimal strategies need infinite memory.

ENERGY $\cap$ PARITY objectives for finite MDPs:

1. Almost-sure optimal strategies need infinite memory.
2. A.s. winning sets are computable in NP $\cap$ coNP and (pseudo) P-time, by (a new!) reduction to Mean-Payoff games.

ENERGY $\cap$ PARITY objectives for finite MDPs:

1. Almost-sure optimal strategies need infinite memory.
2. A.s. winning sets are computable in NP $\cap$ coNP and (pseudo) P-time, by (a new!) reduction to Mean-Payoff games.
3. Same bounds hold for the limit-sure problem;

What's the Problem?

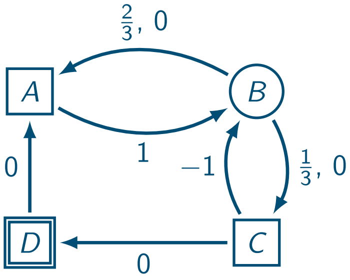

What's the Problem?

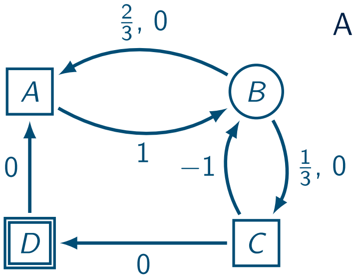

Aim: Satisfy ENERGY and avoid D

What's the Problem?

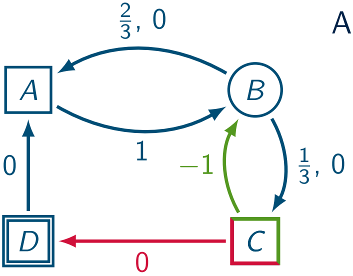

Aim: Satisfy ENERGY and avoid D

What's the Problem?

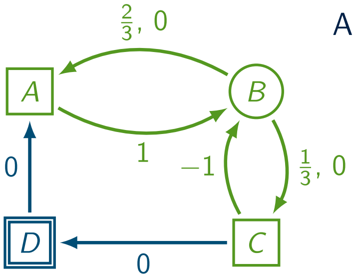

Aim: Satisfy ENERGY and avoid D

- stay in $\{A, B, C\}$ and lose ENERGY

What's the Problem?

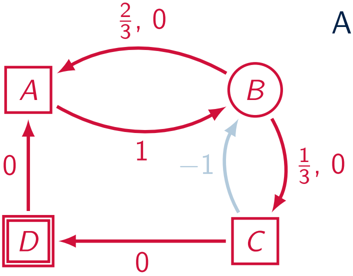

Aim: Satisfy ENERGY and avoid D

- stay in $\{A, B, C\}$ and lose ENERGY
- prefer D over B and lose PARITY

What's the Problem?

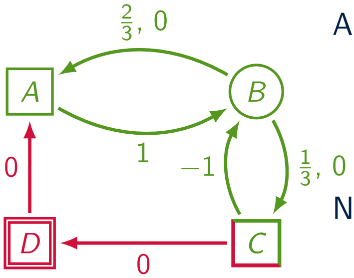

Aim: Satisfy ENERGY and avoid D

- stay in $\{A, B, C\}$ and lose ENERGY
- prefer D over B and lose PARITY

No FM-strategy wins (a.s.)

- eventually commits to red or green
- maintain bounded distance to D

What's the Problem?

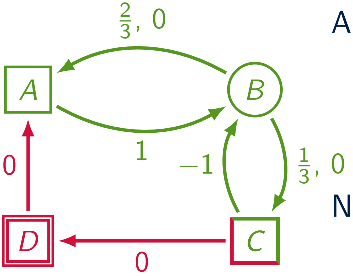

Aim: Satisfy ENERGY and avoid D

- stay in $\{A, B, C\}$ and lose ENERGY
- prefer D over B and lose PARITY

No FM-strategy wins (a.s.)

- eventually commits to red or green
- maintain bounded distance to D

An (a.s.) winning strategy

- "move to D only if energy level is 0"
- works because $\mathbb{P}^{green}(\text{always} > 0) > 1/2$

What's the Problem?

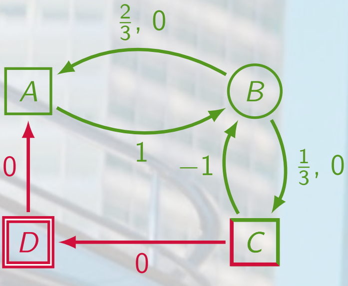

Almost-sure Problems for finite MDPs

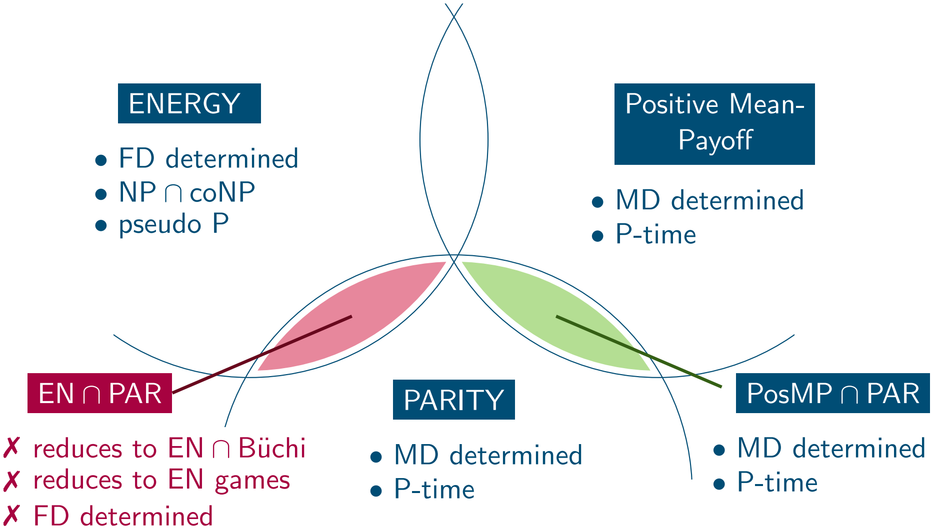

Chatterjee/Doyen, MFCS'11

Almost-sure Problems for finite MDPs

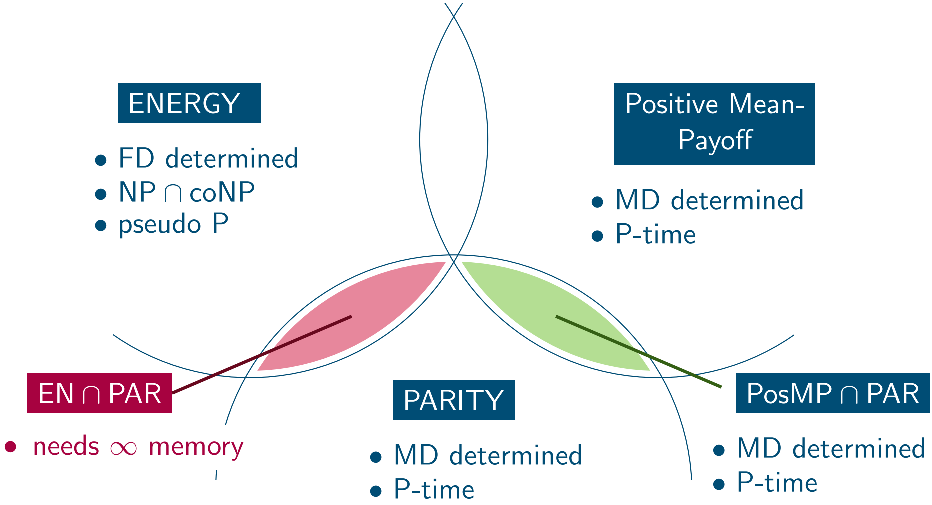

The Storage Objective

A path satiesfies the k-Storage condition ST(k) if

$$k + \sum_{i=n}^{m} \mathrm{cost}(e_i) \geq 0$$

for all indices $n < m$.

The Storage Objective

A path satiesfies the k-Storage condition $\mathrm{ST}(k)$ if

$$k + \sum_{i=n}^{m} \mathrm{cost}(e_i) \geq 0$$

for all indices $n < m$.

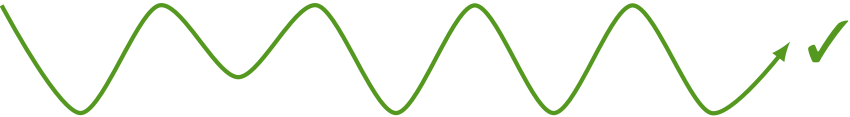

The Storage Objective

A path satiesfies the k-Storage condition $\mathrm{ST}(k)$ if

$$k + \sum_{i=n}^{m} \mathrm{cost}(e_i) \geq 0$$

for all indices $n < m$.

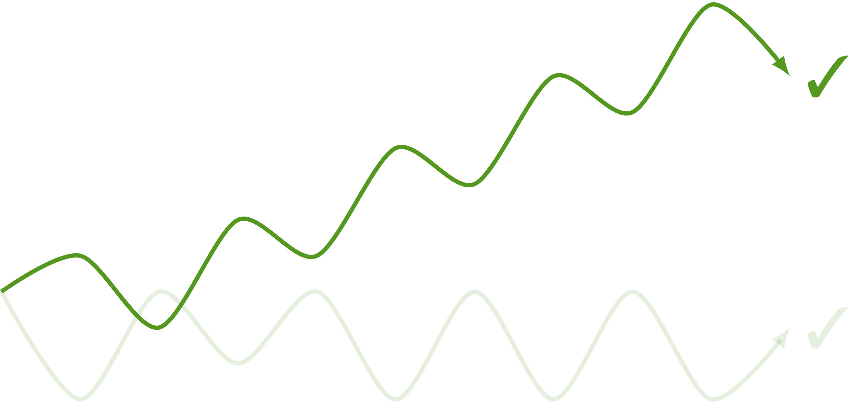

The Storage Objective

A path satiesfies the k-Storage condition ST(k) if

$$k + \sum_{i=n}^{m} \text{cost}(e_i) \geq 0$$

for all indices $n < m$.

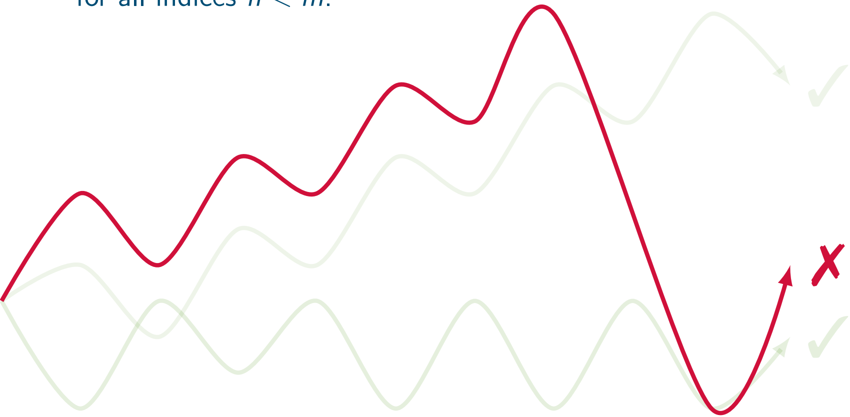

The Storage Objective

A path satiesfies the k-Storage condition $\mathsf{ST}(k)$ if

$$k + \sum_{i=n}^{m} \mathsf{cost}(e_i) \geq 0$$

for all indices $n < m$.

The STORAGE objective asks for ENERGY and $\mathsf{ST}(k)$ for some k.

The Storage Objective

A path satiesfies the k-Storage condition $\mathsf{ST}(k)$ if

$$k + \sum_{i=n}^{m} \mathsf{cost}(e_i) \geq 0$$

for all indices $n < m$.

The STORAGE objective asks for ENERGY and $\mathsf{ST}(k)$ for some k.
(it holds a.s. if $\exists \sigma\ \exists k.\ \mathbb{P}^{\sigma}(\mathsf{ST}(k)) = 1 \wedge \mathbb{P}^{\sigma}(\mathsf{ENERGY}) = 1$)

The Storage Objective

A path satiesfies the k-Storage condition ST(k) if

$$k + \sum_{i=n}^{m} \mathsf{cost}(e_i) \geq 0$$

for all indices $n < m$.

The STORAGE objective asks for ENERGY and ST(k) for some k.
(it holds a.s. if $\exists\sigma\ \exists k.\ \mathbb{P}^{\sigma}(\mathrm{ST}(k)) = 1 \wedge \mathbb{P}^{\sigma}(\mathrm{ENERGY}) = 1$)

CAUTION

(a.s.) ENERGY $\iff$ (a.s.) STORAGE
but
(a.s.) ENERGY $\cap$ PARITY $\not\iff$ (a.s.) STORAGE $\cap$ PARITY

Almost-sure Problems for finite MDPs

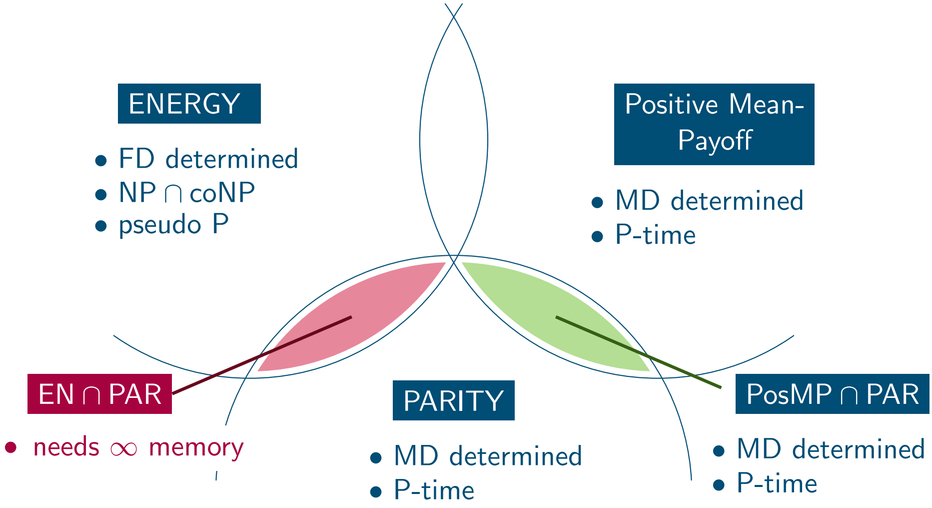

Almost-sure Problems for finite MDPs

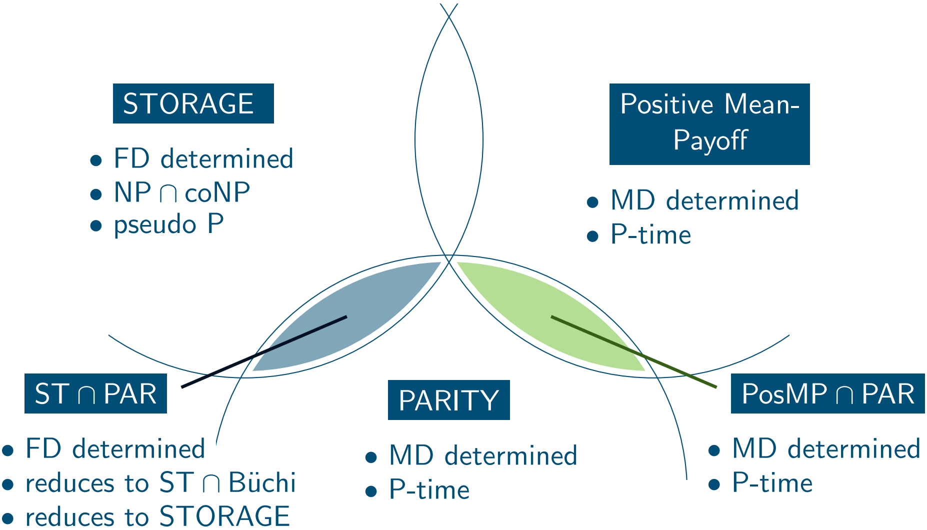

Almost-sure Problems for finite MDPs

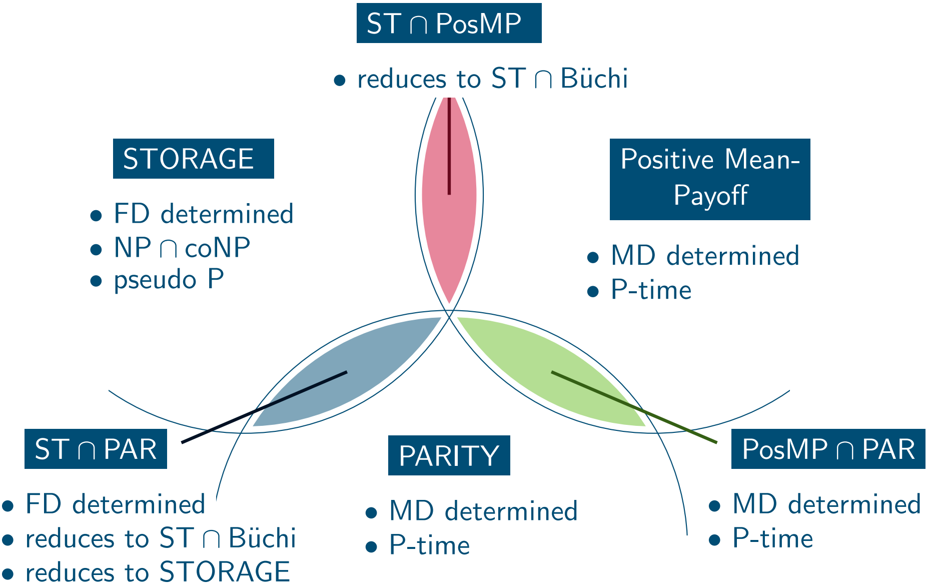

Computing Winning Sets

ENERGY $\cap$ PARITY holds almost-surely iff

1. ST $\cap$ PARITY holds a.s., or

Computing Winning Sets

ENERGY $\cap$ PARITY holds almost-surely iff

1. ST $\cap$ PARITY holds a.s., or
2. There are strategies σ, σ' such that
   - 2.1 σ witnesses a.s. ST $\cap$ PosMP
   - 2.2 σ stays in the winning set of σ'
   - 2.3 σ' witnesses a.s. PAR $\cap$ PosMP
   - 2.4 σ' stays in the winning set of σ

Computing Winning Sets

ENERGY $\cap$ PARITY holds almost-surely iff

Bailout: arbitrary increase

1. ST $\cap$ PARITY holds a.s., or
2. There are strategies σ, σ' such that
   - 2.1 σ witnesses a.s. ST $\cap$ PosMP
   - 2.2 σ stays in the winning set of σ'
   - 2.3 σ' witnesses a.s. PAR $\cap$ PosMP
   - 2.4 σ' stays in the winning set of σ

Computing Winning Sets

ENERGY $\cap$ PARITY holds almost-surely iff

1. ST $\cap$ PARITY holds a.s., or
2. There are strategies σ, σ' such that
   - 2.1 σ witnesses a.s. ST $\cap$ PosMP
   - 2.2 σ stays in the winning set of σ'
   - 2.3 σ' witnesses a.s. PAR $\cap$ PosMP
   - 2.4 σ' stays in the winning set of σ

$\delta > 0$ chance of success

Computing Winning Sets

ENERGY $\cap$ PARITY holds almost-surely iff

1. ST $\cap$ PARITY holds a.s., or
2. There are strategies σ, σ' such that
   - 2.1 σ witnesses a.s. ST $\cap$ PosMP
   - 2.2 σ stays in the winning set of σ'
   - 2.3 σ' witnesses a.s. PAR $\cap$ PosMP
   - 2.4 σ' stays in the winning set of σ

switch and try again

Summary & Outlook

ENERGY $\cap$ PARITY objectives for finite MDPs

- Almost-sure optimal strategies need infinite memory.

Summary & Outlook

ENERGY $\cap$ PARITY objectives for finite MDPs

- Almost-sure optimal strategies need infinite memory.
- A.s. winning sets are in NP $\cap$ coNP and (pseudo) P-time, by (a new!) reduction to Mean-Payoff games.

Summary & Outlook

ENERGY $\cap$ PARITY objectives for finite MDPs

- Almost-sure optimal strategies need infinite memory.
- A.s. winning sets are in NP $\cap$ coNP and (pseudo) P-time, by (a new!) reduction to Mean-Payoff games.
- Same bounds for the limit-sure problem; (ε-optimal strategies can be chosen FD)

Summary & Outlook

ENERGY $\cap$ PARITY objectives for finite MDPs

- Almost-sure optimal strategies need infinite memory.
- A.s. winning sets are in NP $\cap$ coNP and (pseudo) P-time, by (a new!) reduction to Mean-Payoff games.
- Same bounds for the limit-sure problem; (ε-optimal strategies can be chosen FD)

What next?

- generalize to $2\frac{1}{2}$ player games?
- multiple dimensions?

thank you.