Planning a Journey in an Uncertain Environment: The Stochastic Shortest Path Problem Revisited

Mickael Randour
ULB, Computer Science Department
UMONS, Theoretical Computer Science Unit, Infortech – Complexys

April 21, 2016

Infortech Scientific Day

UNIVERSITÉ LIBRE DE BRUXELLES

UMONS Université de Mons

Controller synthesis

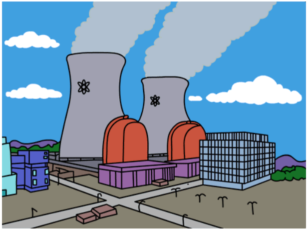

- Setting:
  - a reactive **system** to *control*,
  - an *interacting* **environment**,
  - a **specification** to *enforce*.

Controller synthesis

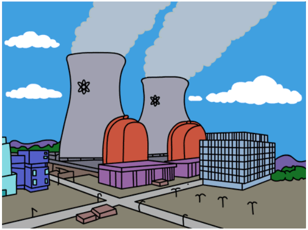

- Setting:
  - ▷ a reactive **system** to *control*,
  - ▷ an *interacting* **environment**,
  - ▷ a **specification** to *enforce*.
- For **critical** systems (e.g., airplane controller, power plants, ABS), testing is not enough!
  - ⇒ Need **formal methods**.

Controller synthesis

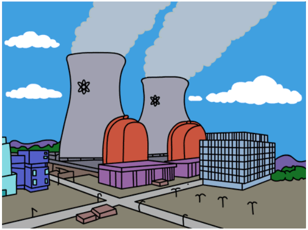

- Setting:
  - a reactive **system** to *control*,
  - an *interacting* **environment**,
  - a **specification** to *enforce*.
- For **critical** systems (e.g., airplane controller, power plants, ABS), testing is not enough!
  - ⇒ Need **formal methods**.
- **Automated synthesis** of provably-correct and efficient controllers:
  - mathematical frameworks,
    - ↪ e.g., game theory [GTW02, Ran13, Ran14]
  - software tools.

Strategy synthesis in stochastic environments

Strategy = formal model of how to control the system

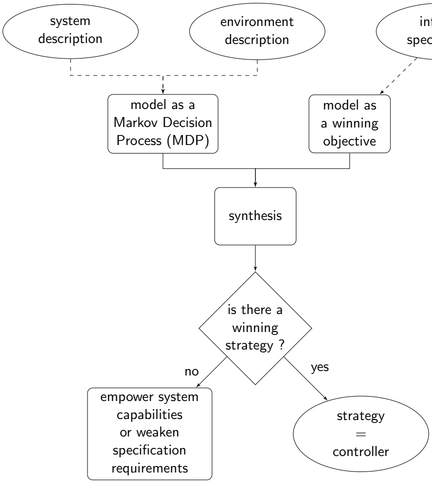

Strategy synthesis in stochastic environments

Strategy = formal model of how to control the system

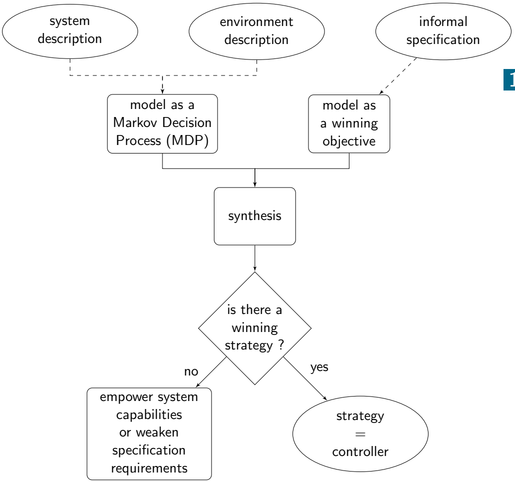

1. How complex is it to decide if a winning strategy exists?

Strategy synthesis in stochastic environments

Strategy = formal model of how to control the system

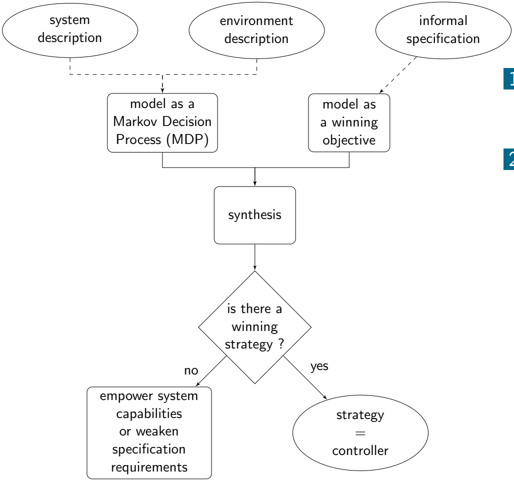

1. How complex is it to decide if a winning strategy exists?
2. How complex such a strategy needs to be? **Simpler is better**.

Strategy synthesis in stochastic environments

Strategy = formal model of how to control the system

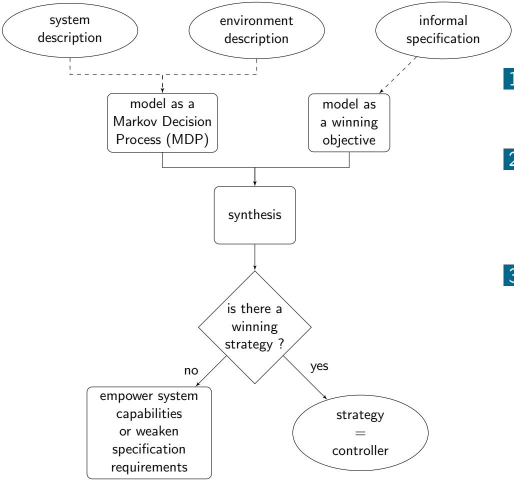

1. How complex is it to decide if a winning strategy exists?
2. How complex such a strategy needs to be? **Simpler is better**.
3. Can we synthesize one efficiently?

Strategy synthesis in stochastic environments

Strategy = formal model of how to control the system

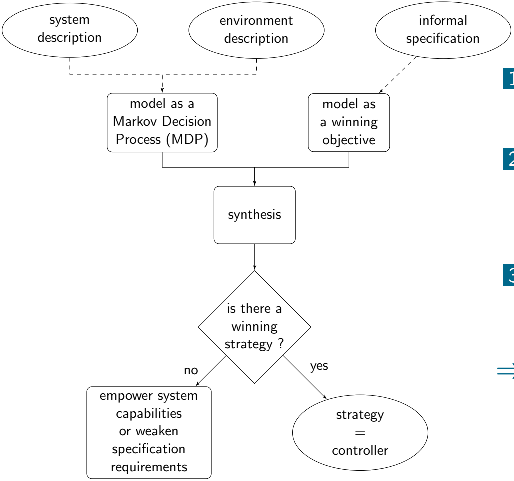

1. How complex is it to decide if a winning strategy exists?
2. How complex such a strategy needs to be? **Simpler is better**.
3. Can we synthesize one efficiently?

⇒ Depends on the winning objective, the exact type of interaction, etc.

Aim of this talk

Flavor of $\neq$ types of **useful strategies** in stochastic environments.

- ▷ Joint paper[1] with J.-F. Raskin (ULB) and O. Sankur (IRISA, Rennes) [RRS15b]
- ▷ Full paper available on arXiv: abs/1411.0835

[1] Invited talk in VMCAI 2015 - 16th International Conference on Verification, Model Checking, and Abstract Interpretation.

Aim of this talk

Flavor of $\neq$ types of **useful strategies** in stochastic environments.

- ▷ Joint paper[1] with J.-F. Raskin (ULB) and O. Sankur (IRISA, Rennes) [RRS15b]
- ▷ Full paper available on arXiv: abs/1411.0835

Applications to the **shortest path problem**.

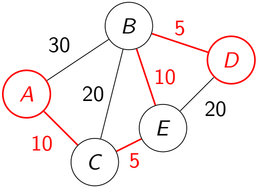

↪ Find a **path of minimal length** in a weighted graph (Dijkstra, Bellman-Ford, etc) [CGR96].

[1] Invited talk in VMCAI 2015 - 16th International Conference on Verification, Model Checking, and Abstract Interpretation.

Aim of this talk

Flavor of $\neq$ types of **useful strategies** in stochastic environments.

- ▷ Joint paper[1] with J.-F. Raskin (ULB) and O. Sankur (IRISA, Rennes) [RRS15b]
- ▷ Full paper available on arXiv: abs/1411.0835

Applications to the **shortest path problem**.

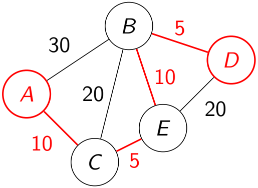

What if the environment is **uncertain**? E.g., in case of heavy traffic, some roads may be crowded.

[1] Invited talk in VMCAI 2015 - 16th International Conference on Verification, Model Checking, and Abstract Interpretation.

Planning a journey in an uncertain environment

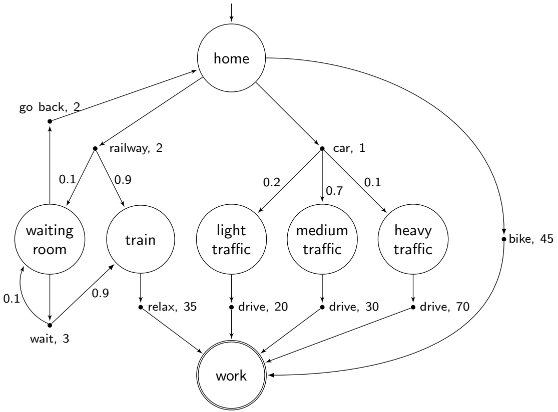

Each action takes time, target = work.

- ▷ What kind of strategies are we looking for when the environment is **stochastic** (MDP)?

Solution 1: minimize the *expected* time to work

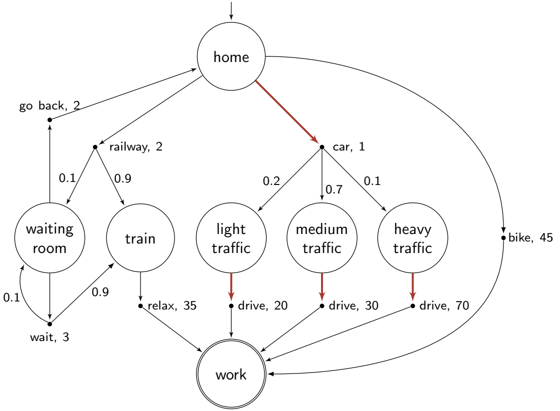

- ▷ "Average" performance: meaningful when you journey often.
- ▷ **Simple strategies** suffice: no memory, no randomness.
- ▷ Taking the **car** is optimal: $\mathbb{E}^{\sigma}_{D}(\mathsf{TS}^{\text{work}}) = 33$.

Solution 2: traveling without taking too many risks

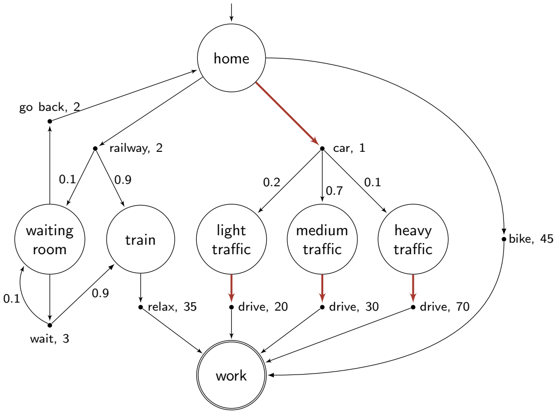

Minimizing the *expected time* to destination makes sense **if** we travel often and **it is not a problem to be late**.

With car, in 10% of the cases, the journey takes 71 minutes.

Solution 2: traveling without taking too many risks

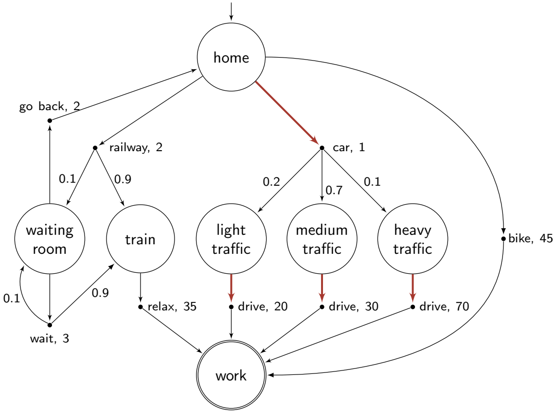

Most bosses will not be happy if we are late too often...

$\rightsquigarrow$ what if we are risk-averse and want to avoid that?

Solution 2: maximize the *probability* to be on time

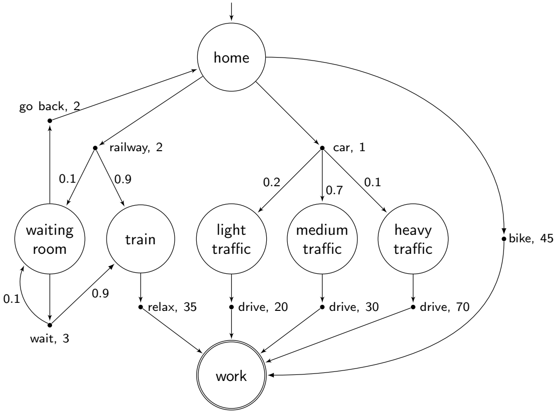

Specification: reach work within 40 minutes with 0.95 probability

Solution 2: maximize the *probability* to be on time

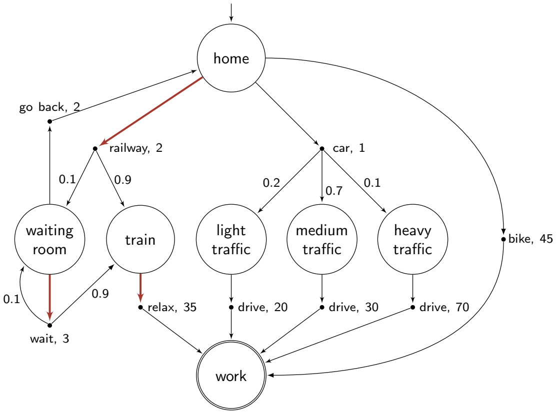

Specification: reach work within 40 minutes with 0.95 probability

Sample strategy: take the **train** $\rightsquigarrow \mathbb{P}^{\sigma}_{D}\left[\mathsf{TS}^{\mathsf{work}} \leq 40\right] = 0.99$

Bad choices: car (0.9) and bike (0.0)

Solution 3: strict worst-case guarantees

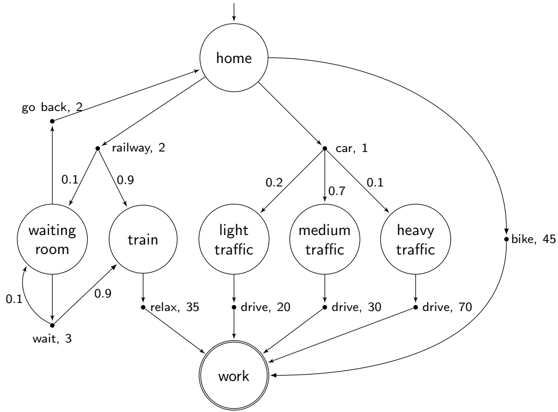

Specification: *guarantee* that work is reached within 60 minutes (to avoid missing an important meeting)

Solution 3: strict worst-case guarantees

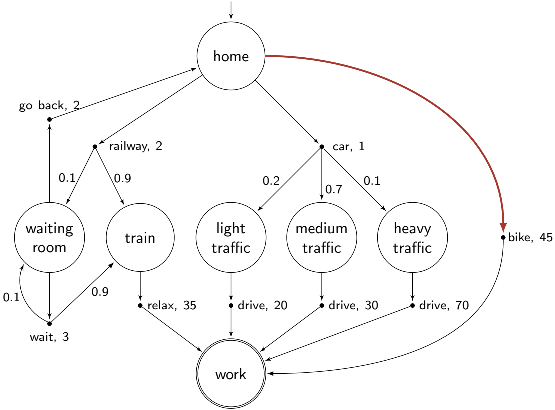

Specification: *guarantee* that work is reached within 60 minutes (to avoid missing an important meeting)

Sample strategy: **bike** $\rightsquigarrow$ worst-case reaching time = 45 minutes.
Bad choices: train ($wc = \infty$) and car ($wc = 71$)

Solution 3: strict worst-case guarantees

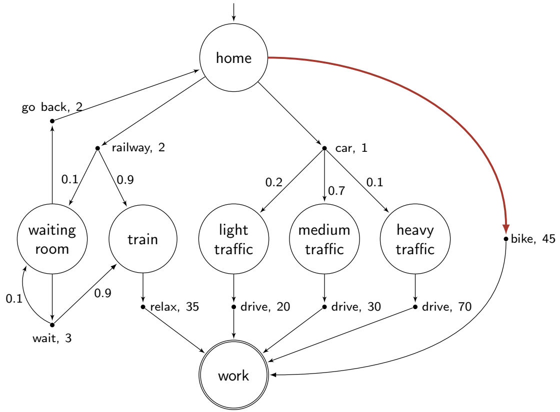

Worst-case analysis $\rightsquigarrow$ **two-player game** against an antagonistic adversary (*bad guy*)

- ▷ forget about probabilities and give the choice of transitions to the adversary

Solution 4: minimize the *expected* time under strict worst-case guarantees

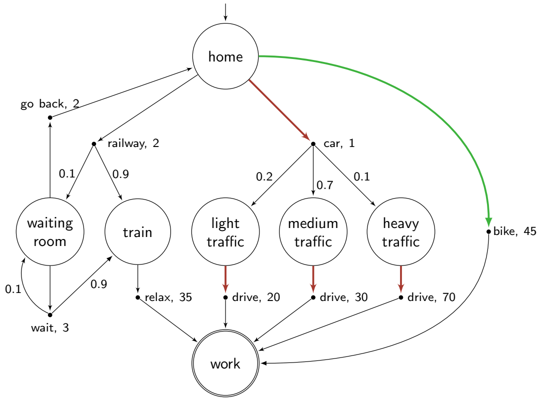

- Expected time: car $\rightsquigarrow \mathbb{E} = 33$ but $wc = 71 > 60$
- Worst-case: bike $\rightsquigarrow wc = 45 < 60$ but $\mathbb{E} = 45 >>> 33$

Solution 4: minimize the *expected* time under strict worst-case guarantees

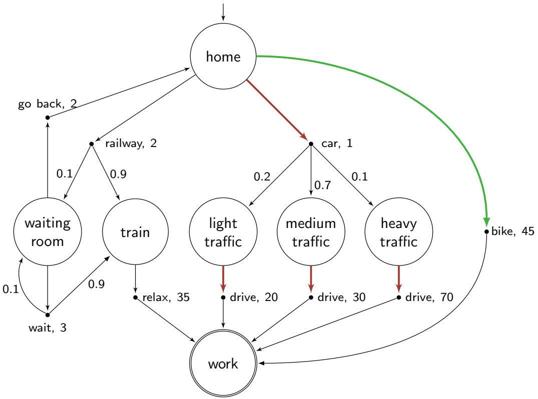

In practice, we want both! Can we do better?

▷ **Beyond worst-case synthesis** [BFRR14b, BFRR14a]: minimize the expected time under the worst-case constraint.

Solution 4: minimize the *expected* time under strict worst-case guarantees

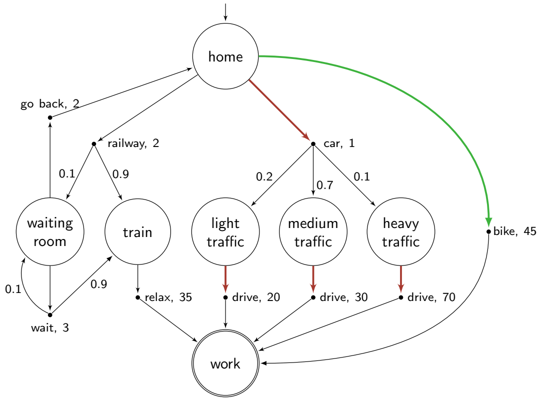

Sample strategy: try train up to 3 delays then switch to bike.

$\rightsquigarrow$ $wc = 58 < 60$ and $\mathbb{E} \approx 37.34 << 45$

$\rightsquigarrow$ Strategies need **memory** $\rightsquigarrow$ more complex!

Solution 5: multiple objectives ⇒ trade-offs

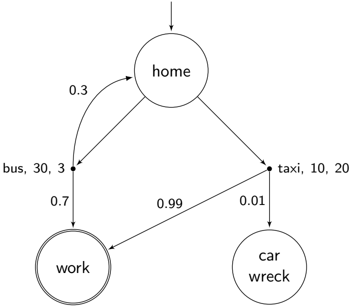

Two-dimensional weights on actions: *time* and *cost*.

Often necessary to consider trade-offs: e.g., between the probability to reach work in due time and the risks of an expensive journey.

Solution 5: multiple objectives $\Rightarrow$ trade-offs

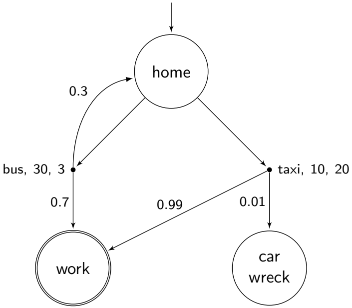

Solution 2 (probability) can only ensure a **single constraint**.

- **C1**: 80% of runs reach work in at most 40 minutes.
  - Taxi $\rightsquigarrow \leq 10$ minutes with probability $0.99 > 0.8$.

Solution 5: multiple objectives $\Rightarrow$ trade-offs

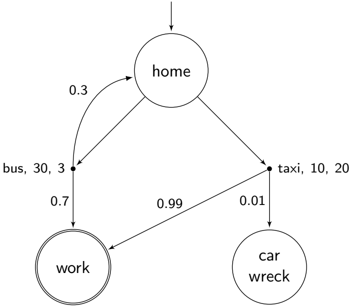

Solution 2 (probability) can only ensure a **single constraint**.

- **C1**: 80% of runs reach work in at most 40 minutes.
  - ▷ Taxi $\leadsto \leq 10$ minutes with probability $0.99 > 0.8$.
- **C2**: 50% of them cost at most 10$ to reach work.
  - ▷ Bus $\leadsto \geq 70\%$ of the runs reach work for 3$.

Solution 5: multiple objectives ⇒ trade-offs

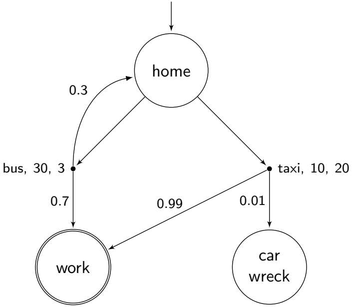

Solution 2 (probability) can only ensure a **single constraint**.

- **C1**: 80% of runs reach work in at most 40 minutes.
  - ▷ Taxi $\rightsquigarrow \leq 10$ minutes with probability $0.99 > 0.8$.
- **C2**: 50% of them cost at most 10$ to reach work.
  - ▷ Bus $\rightsquigarrow \geq 70\%$ of the runs reach work for 3$.

Taxi $\not\models$ C2, bus $\not\models$ C1. What if we want C1 $\wedge$ C2?

Solution 5: multiple objectives ⇒ trade-offs

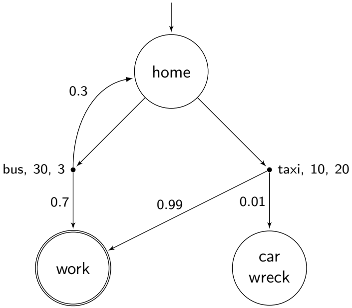

- **C1**: 80% of runs reach work in at most 40 minutes.
- **C2**: 50% of them cost at most 10$ to reach work.

Study of **multi-constraint percentile queries** [RRS15a].

- ▷ Sample strategy: bus once, then taxi. Requires *memory*.
- ▷ Another strategy: bus with probability 3/5, taxi with probability 2/5. Requires *randomness*.

Solution 5: multiple objectives ⇒ trade-offs

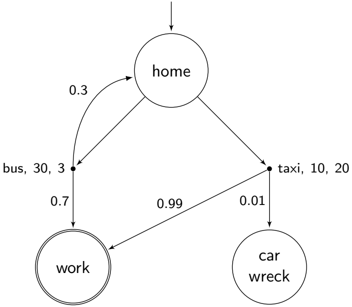

- **C1**: 80% of runs reach work in at most 40 minutes.
- **C2**: 50% of them cost at most 10$ to reach work.

Study of **multi-constraint percentile queries** [RRS15a].

In general, *both* memory *and* randomness are required.

$\neq$ previous problems $\rightsquigarrow$ more complex!

Conclusion (1/3)

This talk was about **shortest path objectives**, but there are many more! Some examples based on energy applications.

- **Energy**: operate with a (bounded) fuel tank and never run out of fuel [BFL+08].
- **Mean-payoff**: average cost/reward (or energy consumption) per action in the long run [EM79].
- **Average-energy**: energy objective + optimize the long-run average amount of fuel in the tank [BMR+15].

Conclusion (2/3)

Our research aims at:

- defining meaningful *strategy concepts*,
- providing *algorithms* and *tools* to compute those strategies,
- classifying the *complexity* of the different problems from a theoretical standpoint.
  - ↪ Is it mathematically possible to obtain efficient algorithms?

Conclusion (3/3)

Algorithmic complexity: hierarchy of problems (for shortest path)

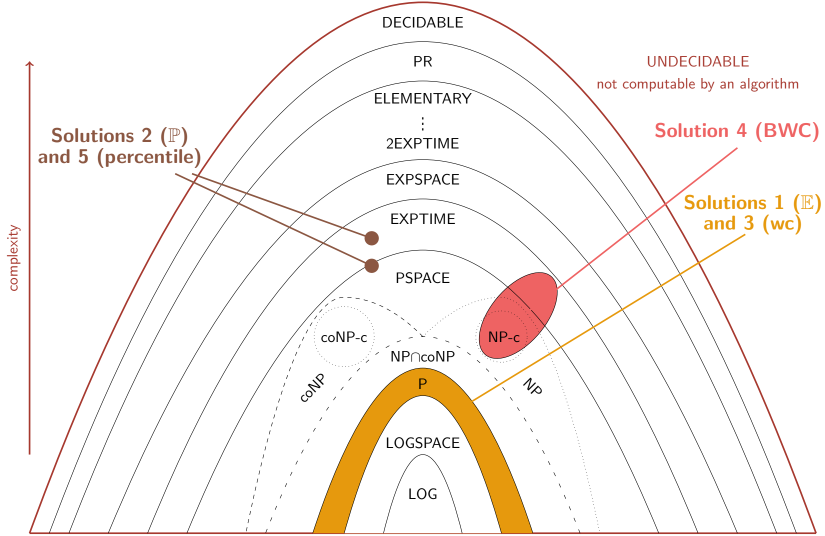

Thank you! Any question?

References I

P. Bouyer, U. Fahrenberg, K.G. Larsen, N. Markey, and J. Srba.
Infinite runs in weighted timed automata with energy constraints.
In Proc. of FORMATS, LNCS 5215, pages 33–47. Springer, 2008.

V. Bruyère, E. Filiot, M. Randour, and J.-F. Raskin.
Expectations or guarantees? I want it all! A crossroad between games and MDPs.
In Proc. of SR, EPTCS 146, pages 1–8, 2014.

V. Bruyère, E. Filiot, M. Randour, and J.-F. Raskin.
Meet your expectations with guarantees: Beyond worst-case synthesis in quantitative games.
In Proc. of STACS, LIPIcs 25, pages 199–213. Schloss Dagstuhl - LZI, 2014.

P. Bouyer, N. Markey, M. Randour, K.G. Larsen, and S. Laursen.
Average-energy games.
In Proc. of GandALF, EPTCS 193, pages 1–15, 2015.

B.V. Cherkassky, A.V. Goldberg, and T. Radzik.
Shortest paths algorithms: Theory and experimental evaluation.
Math. programming, 73(2):129–174, 1996.

A. Ehrenfeucht and J. Mycielski.
Positional strategies for mean payoff games.
International Journal of Game Theory, 8:109–113, 1979.

E. Grädel, W. Thomas, and T. Wilke, editors.
Automata, Logics, and Infinite Games: A Guide to Current Research, LNCS 2500. Springer, 2002.

References II

M. Randour.
Automated synthesis of reliable and efficient systems through game theory: A case study.
In Proceedings of the European Conference on Complex Systems 2012, Springer Proceedings in Complexity XVII, pages 731–738. Springer, 2013.

M. Randour.
Synthesis in Multi-Criteria Quantitative Games.
PhD thesis, Université de Mons, Belgium, 2014.

M. Randour, J.-F. Raskin, and O. Sankur.
Percentile queries in multi-dimensional Markov decision processes.
In Proc. of CAV, LNCS 9206, pages 123–139. Springer, 2015.

M. Randour, J.-F. Raskin, and O. Sankur.
Variations on the stochastic shortest path problem.
In Proc. of VMCAI, LNCS 8931, pages 1–18. Springer, 2015.