Controlling a population of identical MDP

Nathalie Bertrand

Inria Rennes

ongoing work with Miheer Dewaskar (CMI), Blaise Genest (IRISA) and Hugo Gimbert (LaBRI)

Trends and Challenges in Quantitative Verification

Motivation

Control of gene expression for a population of cells

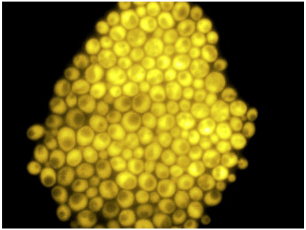

credits: G. Batt

Motivation

Control of gene expression for a population of cells

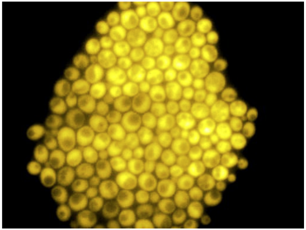

credits: G. Batt

- cell population
- gene expression monitored through fluorescence level
- drug injections affect all cells
- response varies from cell to cell
- obtain a large proportion of cells with desired gene expression level

Motivation

Control of gene expression for a population of cells

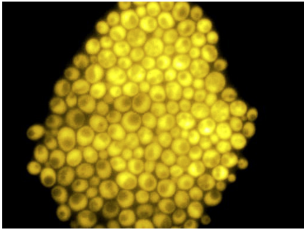

credits: G. Batt

- cell population
- gene expression monitored through fluorescence level
- drug injections affect all cells
- response varies from cell to cell
- obtain a large proportion of cells with desired gene expression level

- arbitrary nb of components
- full observation
- uniform control
- MDP model for single cell
- global quantitative reachability objective

Markov decision processes

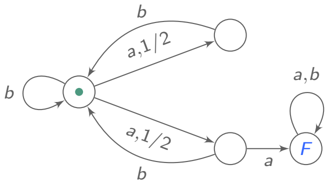

- non-deterministic actions: $\{a, b\}$
- prob. distribution over successors

Markov decision processes

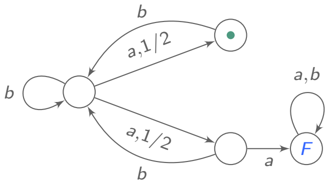

- non-deterministic actions: $\{a, b\}$
- prob. distribution over successors

Markov decision processes

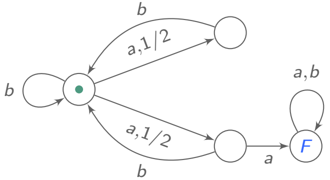

- non-deterministic actions: $\{a, b\}$
- prob. distribution over successors

Markov decision processes

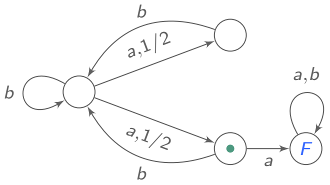

- non-deterministic actions: $\{a, b\}$
- prob. distribution over successors

Markov decision processes

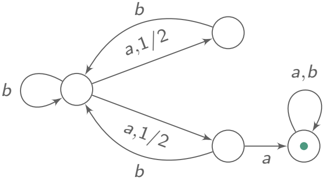

- non-deterministic actions: $\{a, b\}$
- prob. distribution over successors

Markov decision processes

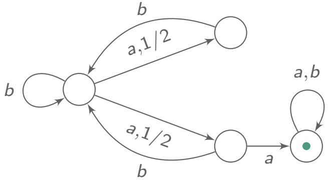

- non-deterministic actions: $\{a, b\}$
- prob. distribution over successors

Scheduler $\sigma : S^+ \to \Sigma$ resolves non-determinism
induces Markov chain with probability measure $\mathbb{P}_\sigma$

Markov decision processes

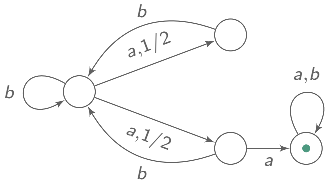

- non-deterministic actions: $\{a, b\}$
- prob. distribution over successors

Scheduler $\sigma : S^+ \to \Sigma$ resolves non-determinism
induces Markov chain with probability measure $\mathbb{P}_\sigma$

Theorem: reachability checking for MDP
The following problems are in PTIME
$\exists\sigma,\ \mathbb{P}_\sigma(\Diamond F) = 1?$ $\exists\sigma,\ \mathbb{P}_\sigma(\Diamond F) > .7?$ compute $\max_\sigma \mathbb{P}_\sigma(\Diamond F)$.

Back to our motivating application

Control of gene expression for a population of cells

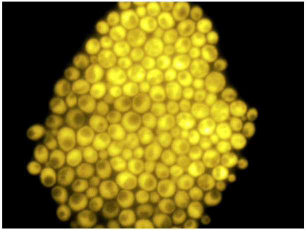

credits: G. Batt

- arbitrary nb of components
- full observation
- uniform control
- MDP model for single cell
- global quantitative reachability objective

Modelling

- population of N identical MDP $\mathcal{M}$
- uniform control policy under full observation

Modelling

- population of N identical MDP $\mathcal{M}$
- uniform control policy under full observation

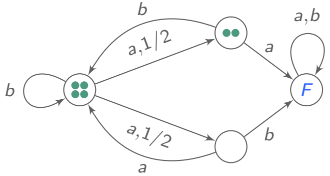

Modelling

- population of N identical MDP $\mathcal{M}$
- uniform control policy under full observation

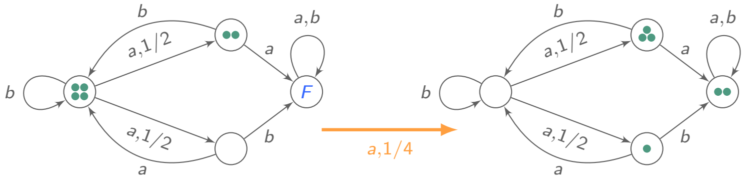

Modelling

- population of N identical MDP $\mathcal{M}$
- uniform control policy under full observation

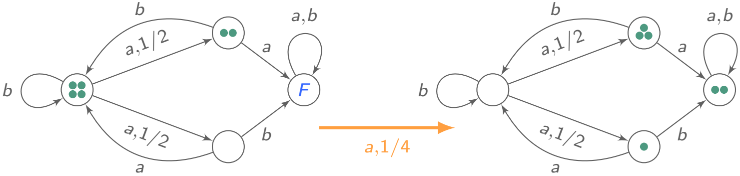

Verification question does the maximum probability that a given proportion of MDPs reach a target set of states meet a threshold?

Modelling

- population of N identical MDP $\mathcal{M}$
- uniform control policy under full observation

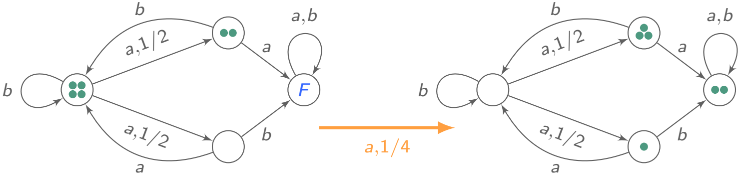

Verification question does the maximum probability that a given proportion of MDPs reach a target set of states meet a threshold?

Fixed N: build the product MDP $\mathcal{M}^N$, identify global target states, compute optimal scheduler

Parameterized verification

Verification question does the maximum probability that a given proportion of MDPs reach a target set of states meet a threshold?

Parameterized verification

Verification question does the maximum probability that a given proportion of MDPs reach a target set of states meet a threshold?

Parameter N: check the global objective for all population sizes N

$$\forall N \max_{\sigma} \mathbb{P}_{\sigma}(\mathcal{M}^N \models \Diamond \text{ at least 80\% of MDPs in F }) \geq .7?$$

Parameterized verification

Verification question does the maximum probability that a given proportion of MDPs reach a target set of states meet a threshold?

Parameter N: check the global objective for all population sizes N

$$\forall N \max_{\sigma} \mathbb{P}_{\sigma}(\mathcal{M}^N \models \Diamond \text{ at least 80\% of MDPs in F }) \geq .7?$$

Restricted cases

- qualitative: almost-sure convergence

$$\forall N \max_{\sigma} \mathbb{P}_{\sigma}(\mathcal{M}^N \models \Diamond F^N) = 1?$$

- Boolean: sure convergence

$$\forall N \; \exists \sigma, \; \mathcal{M}^N \models \forall_{\sigma} \Diamond F^N?$$

This talk

Problem setting

- **Boolean** parameterized verification questions
- uniform control for population of **NFA** $\equiv$ 2-player turn-based game
  - controller chooses the action (*e.g.* a)
  - opponent chooses how to move each individual copy (a-transition)
- **convergence** objective: all copies in a target set $F \subseteq Q$

$$\forall N\ \exists\sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N?$$

$$\forall N\ \exists\sigma,\ \forall\tau,\ (\mathcal{M}^N, \sigma, \tau) \models \Diamond F^N?$$

This talk

Problem setting

- **Boolean** parameterized verification questions
- uniform control for population of **NFA** $\equiv$ 2-player turn-based game
  - controller chooses the action (*e.g.* a)
  - opponent chooses how to move each individual copy (a-transition)
- **convergence** objective: all copies in a target set $F \subseteq Q$

$$\forall N\ \exists \sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N?$$

$$\forall N\ \exists \sigma,\ \forall \tau,\ (\mathcal{M}^N, \sigma, \tau) \models \Diamond F^N?$$

Questions addressed

- decidability
- memory requirements for controller σ
- admissible values for N

Monotonicity property

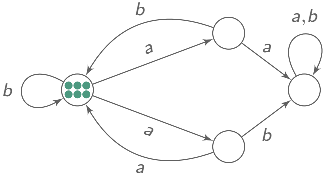

$$\forall N\ \exists\sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N?$$

Monotonicity property

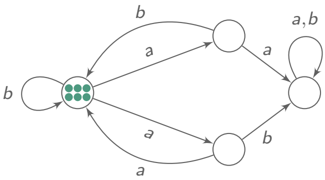

$$\forall N\ \exists \sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N?$$

Monotonicity: harder when N grows

$$\exists \sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N \quad \Longrightarrow \quad \forall M \leq N,\ \exists \sigma,\ \mathcal{M}^M \models \forall_\sigma \Diamond F^M$$

Monotonicity property

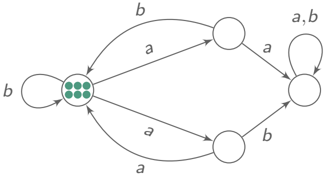

$$\forall N\ \exists\sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N?$$

Monotonicity: harder when N grows

$$\exists\sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N \quad\Longrightarrow\quad \forall M \leq N,\ \exists\sigma,\ \mathcal{M}^M \models \forall_\sigma \Diamond F^M$$

Cutoff: smallest N for which there is no admissible controller σ

A first example and a first question

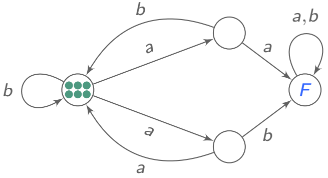

$$\forall N,\ \exists\sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N$$

$\sigma(k, 0, 0, \star) = a \qquad \sigma(0, k_u, k_d, \star) = a \qquad \sigma(0, 0, k_d, \star) = b$

memoryless support-based controllers suffice on this example

A first example and a first question

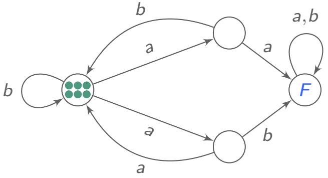

$$\forall N,\ \exists\sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N$$

$\sigma(k, 0, 0, \star) = a \qquad \sigma(0, k_u, k_d, \star) = a \qquad \sigma(0, 0, k_d, \star) = b$

memoryless support-based controllers suffice on this example

Question 1 Are memoryless support-based controllers enough in general?

A second example and a second question

$A = \{a_1, \cdots, a_M\}$ unspecified edges lead to a sink state

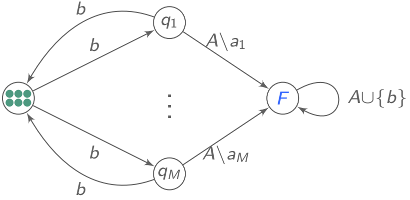

$$\forall N < M,\ \exists\sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N$$

A second example and a second question

$A = \{a_1, \cdots, a_M\}$ unspecified edges lead to a sink state

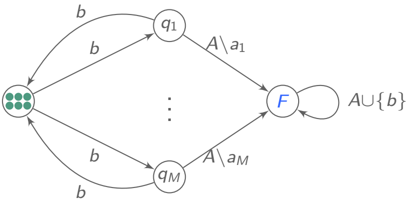

$$\forall N < M,\ \exists \sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N$$

Cutoff $\min\{N \mid \forall \sigma,\ \mathcal{M}^N \not\models \forall_\sigma \Diamond F^N\}$ here $\mathcal{O}(|\mathcal{M}|)$

A second example and a second question

$A = \{a_1, \cdots, a_M\}$ unspecified edges lead to a sink state

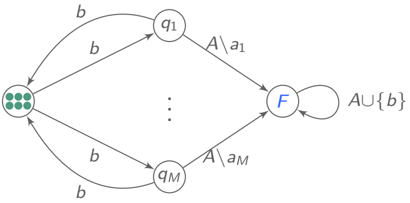

$$\forall N < M,\ \exists \sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N$$

Cutoff $\min\{N \mid \forall \sigma,\ \mathcal{M}^N \not\models \forall_\sigma \Diamond F^N\}$ here $\mathcal{O}(|\mathcal{M}|)$

Question 2 Are cutoffs always polynomial in $|\mathcal{M}|$?

A first answer

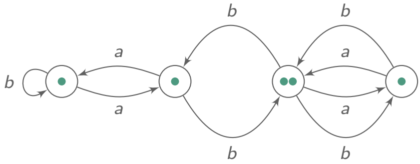

Assumption: if at least one state is empty, the controller ensures convergence

gadget similar to previous example with actions $\{a_1, \cdots, a_4\}$

A first answer

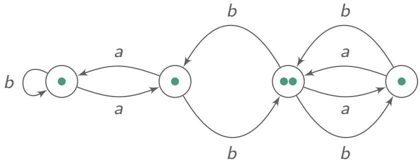

Assumption: if at least one state is empty, the controller ensures convergence

gadget similar to previous example with actions $\{a_1, \cdots, a_4\}$

Possible controllers

- always a: deterministic behaviour, full support is maintained
- always b: splitting the copies in third state allows opponent to win

A first answer

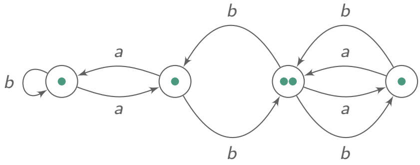

Assumption: if at least one state is empty, the controller ensures convergence

gadget similar to previous example with actions $\{a_1, \cdots, a_4\}$

Possible controllers

- always a: deterministic behaviour, full support is maintained
- always b: splitting the copies in third state allows opponent to win
- a and b in alternation: leak from first/second states to third

A first answer

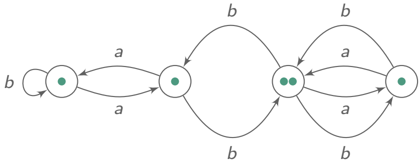

Assumption: if at least one state is empty, the controller ensures convergence

gadget similar to previous example with actions $\{a_1, \cdots, a_4\}$

Possible controllers

- always *a*: deterministic behaviour, full support is maintained
- always *b*: splitting the copies in third state allows opponent to win
- *a* and *b* in alternation: leak from first/second states to third

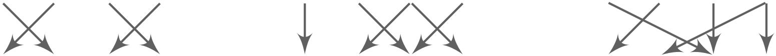

A first answer

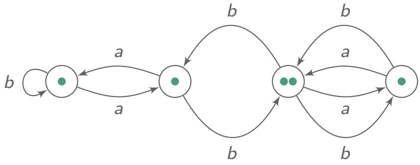

Assumption: if at least one state is empty, the controller ensures convergence

gadget similar to previous example with actions $\{a_1, \cdots, a_4\}$

Possible controllers

- always *a*: deterministic behaviour, full support is maintained
- always *b*: splitting the copies in third state allows opponent to win
- *a* and *b* in alternation: leak from first/second states to third

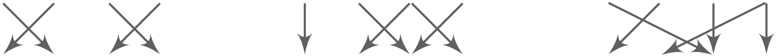

Memoryless support-based controllers are not enough!
Exponential memory on top of support may even be needed.

A second answer

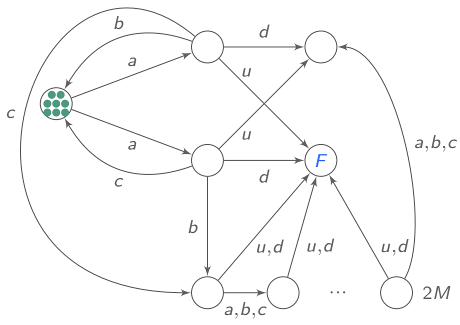

- $\forall N \leq 2^M,\ \exists \sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N$
  accumulate copies in bottom states, then u/d to converge
- for $N > 2^M$ controller cannot avoid reaching the sink state

Cutoff $\mathcal{O}(2^{|\mathcal{M}|})$

A second answer

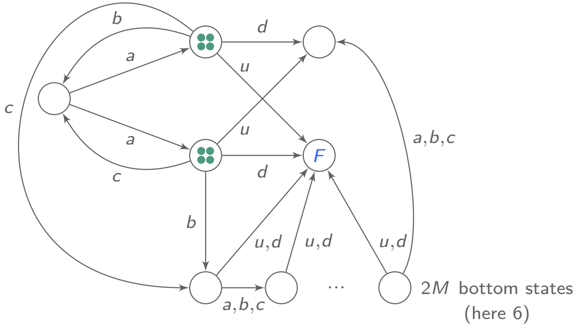

- $\forall N \leq 2^M,\ \exists \sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N$
  accumulate copies in bottom states, then u/d to converge
- for $N > 2^M$ controller cannot avoid reaching the sink state

Cutoff $\mathcal{O}(2^{|\mathcal{M}|})$

A second answer

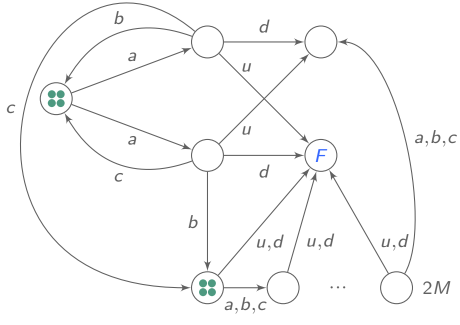

- $\forall N \leq 2^M,\ \exists \sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N$
  accumulate copies in bottom states, then u/d to converge
- for $N > 2^M$ controller cannot avoid reaching the sink state

Cutoff $\mathcal{O}(2^{|\mathcal{M}|})$

A second answer

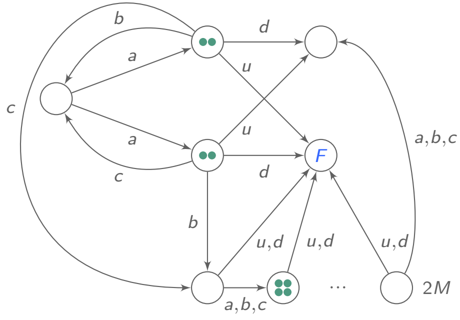

- $\forall N \leq 2^M,\ \exists \sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N$
  accumulate copies in bottom states, then u/d to converge
- for $N > 2^M$ controller cannot avoid reaching the sink state

Cutoff $\mathcal{O}(2^{|\mathcal{M}|})$

A second answer

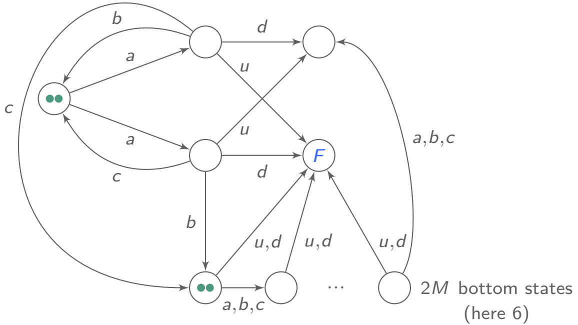

- $\forall N \leq 2^M,\ \exists\sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N$
  accumulate copies in bottom states, then u/d to converge
- for $N > 2^M$ controller cannot avoid reaching the sink state

Cutoff $\mathcal{O}(2^{|\mathcal{M}|})$

A second answer

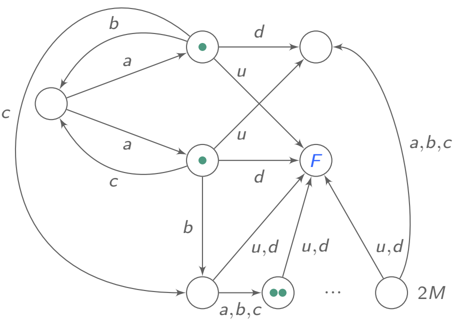

- $\forall N \leq 2^M,\ \exists \sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N$
  accumulate copies in bottom states, then u/d to converge
- for $N > 2^M$ controller cannot avoid reaching the sink state

Cutoff $\mathcal{O}(2^{|\mathcal{M}|})$

A second answer

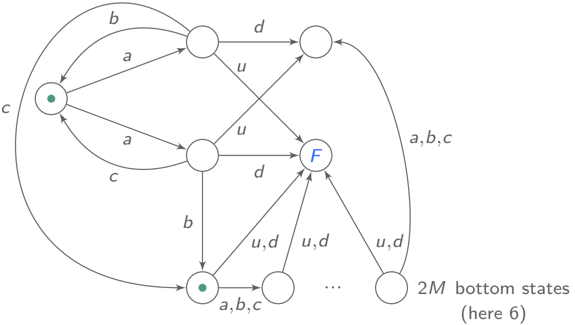

- $\forall N \leq 2^M,\ \exists \sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N$
  accumulate copies in bottom states, then u/d to converge
- for $N > 2^M$ controller cannot avoid reaching the sink state

Cutoff $\mathcal{O}(2^{|\mathcal{M}|})$

A second answer

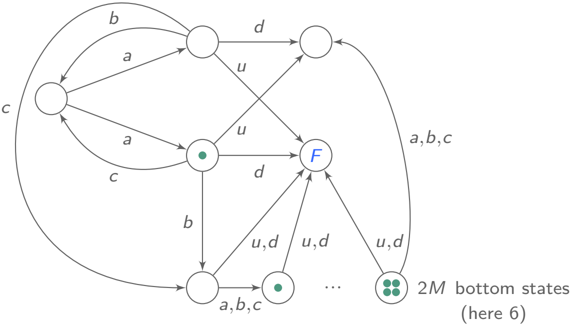

- $\forall N \leq 2^M,\ \exists \sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N$
  accumulate copies in bottom states, then u/d to converge
- for $N > 2^M$ controller cannot avoid reaching the sink state

Cutoff $\mathcal{O}(2^{|\mathcal{M}|})$

A second answer

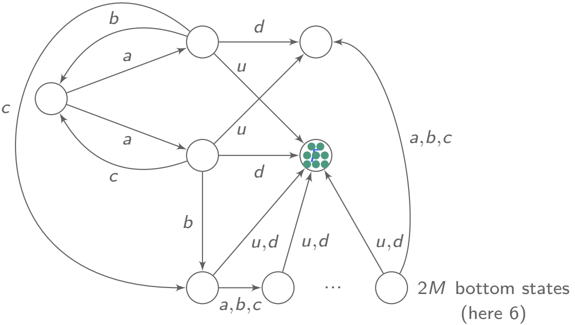

- $\forall N \leq 2^M,\ \exists \sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N$
  accumulate copies in bottom states, then u/d to converge
- for $N > 2^M$ controller cannot avoid reaching the sink state

Cutoff $\mathcal{O}(2^{|\mathcal{M}|})$

A second answer

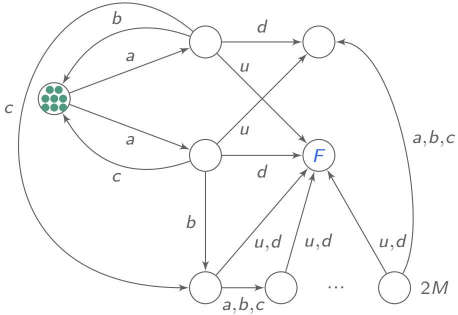

2M bottom states
(here 6)

- $\forall N \leq 2^M,\ \exists\sigma,\ \mathcal{M}^N \models \forall_\sigma \Diamond F^N$
  accumulate copies in bottom states, then u/d to converge
- for $N > 2^M$ controller cannot avoid reaching the sink state

Cutoff $\mathcal{O}(2^{|\mathcal{M}|})$

Cutoff can even be doubly exponential!

Lessons learnt so far

Boolean problem is harder than expected

Lessons learnt so far

Boolean problem is harder than expected

- supports are not enough
- doubly exponential lower bound on cutoffs

  somehow prevents from building the product MDP

Lessons learnt so far

Boolean problem is harder than expected

- supports are not enough
- doubly exponential lower bound on cutoffs
  somehow prevents from building the product MDP
- the more copies the harder, the larger support the harder
- looking at whether supports can be maintained seems promising

Support game

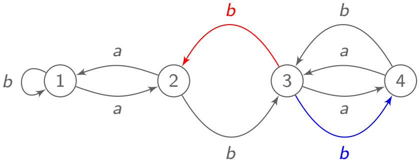

Support game

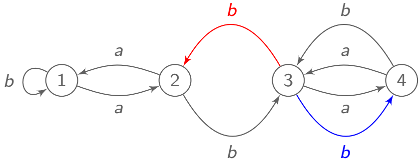

2-player game on possible supports

- □ Eve chooses action
- ◇ Adam chooses transfer relation

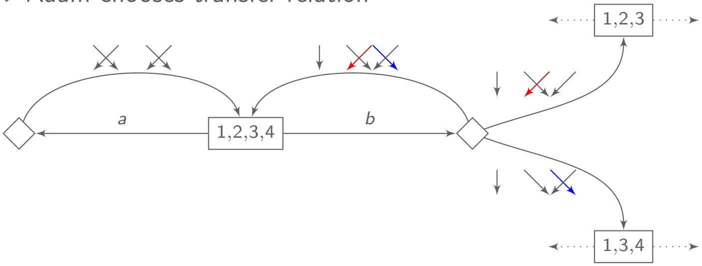

Support game

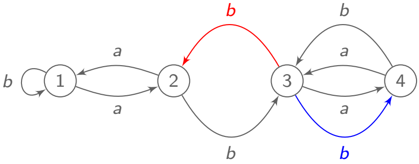

2-player game on possible supports

- □ Eve chooses action
- ◇ Adam chooses transfer relation

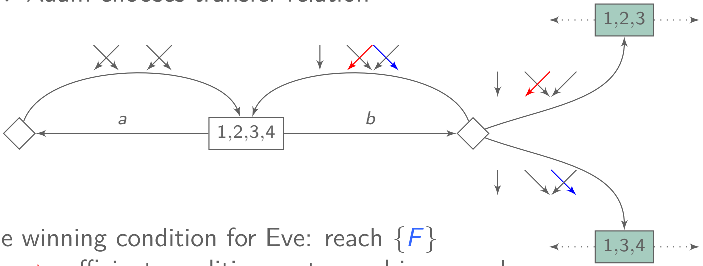

simple winning condition for Eve: reach $\{F\}$

$\rightarrow$ sufficient condition, not sound in general

Refined winning condition

Intuition: allow Eve to monitor some copies and pinpoint leaks
$\rightarrow$ along a play only finitely many leaks are possible

Refined winning condition

Intuition: allow Eve to monitor some copies and pinpoint leaks
$\rightarrow$ along a play only finitely many leaks are possible

Play $\rho = S_0 \xrightarrow{a_1}\xrightarrow{R_1} S_1 \cdots$ **winning for Eve** if there exists $(T_i)_{i\in\mathbb{N}}$ s.t.

(1) $\forall i,\ \emptyset \neq T_i \subseteq S_i$

(2) $\forall i,\ \mathrm{Pre}[R_{i+1}](T_{i+1}) \subseteq T_i$

(3) $\exists^\infty j,\ T_{j+1} \subsetneq \mathrm{Post}[R_{j+1}](T_j)$

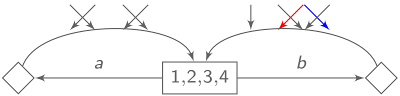

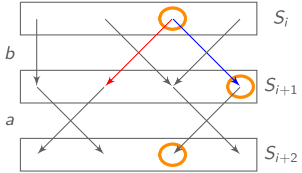

Refined winning condition

Intuition: allow Eve to monitor some copies and pinpoint leaks
$\rightarrow$ along a play only finitely many leaks are possible

Play $\rho = S_0 \xrightarrow{a_1}\xrightarrow{R_1} S_1 \cdots$ **winning for Eve** if there exists $(T_i)_{i \in \mathbb{N}}$ s.t.

(1) $\forall i,\ \emptyset \neq T_i \subseteq S_i$

(2) $\forall i,\ \mathrm{Pre}[R_{i+1}](T_{i+1}) \subseteq T_i$

(3) $\exists^\infty j,\ T_{j+1} \subsetneq \mathrm{Post}[R_{j+1}](T_j)$

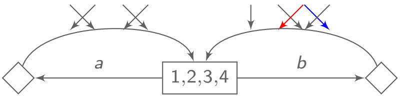

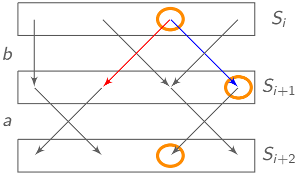

Eve wins support game with refined winning condition iff
$\forall N$ controller has a strategy to reach winning supports

Solving support game w. refined winning condition

Transformation into 2-player partial observation game
with Büchi winning condition

- exponential blowup of game arena

  states (S, T) for all possible $T \subseteq S$
- Adam shall not observe the subsets monitored by Eve

  he only observes S-component of state (S, T)

Solving support game w. refined winning condition

Transformation into 2-player partial observation game
with Büchi winning condition

- exponential blowup of game arena

  states (S, T) for all possible $T \subseteq S$
- Adam shall not observe the subsets monitored by Eve

  he only observes S-component of state (S, T)

Theorem: Decidability and complexity (still to be checked)
Boolean parameterized convergence is decidable in 3EXPTIME.
Cutoff is at most triply exponential in $|\mathcal{M}|$.

Solving support game w. refined winning condition

Transformation into 2-player partial observation game
with Büchi winning condition

- exponential blowup of game arena

  states (S, T) for all possible $T \subseteq S$
- Adam shall not observe the subsets monitored by Eve

  he only observes S-component of state (S, T)

Theorem: Decidability and complexity (still to be checked)
Boolean parameterized convergence is decidable in 3EXPTIME.
Cutoff is at most triply exponential in $|\mathcal{M}|$.

Theorem: (far from matching) **Lower-bounds**
PSPACE-hardness for Boolean parameterized convergence.
Doubly exponential lower-bound on the cutoff.

Contributions

Uniform control of a population of identical MDP

- parameterized verification problem
- Boolean convergence: bring all MDP at the same time in F
  - surprisingly quite involved!
  - beyond support-based optimal controllers
  - 3EXPTIME-decision procedure
  - cutoff between doubly exponential and triply exponential

Back to motivations

Motivation 1: practical motivation

Control of gene expression for a population of cells

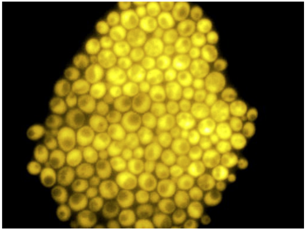

credits: G. Batt

Back to motivations

Motivation 1: practical motivation

Control of gene expression for a population of cells

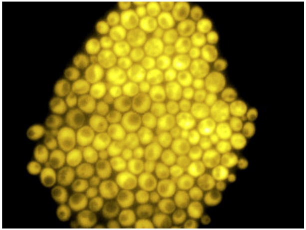

credits: G. Batt

- need for truely probabilistic model
  $\rightarrow$ MDP instead of NFA
- need for truely quantitative questions
  $\rightarrow$ proportions and probabilities instead of convergence and (almost)-sure

$$\forall N \max_{\sigma} \mathbb{P}_\sigma(\mathcal{M}^N \models \Diamond \text{ at least 80\% of MDPs in } F\ ) \geq .7?$$

Back to motivations

Motivation 2: theoretical motivation

Discrete approximation of probabilistic automata

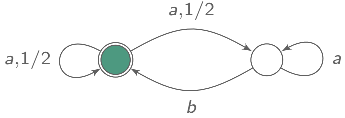

Back to motivations

Motivation 2: theoretical motivation

Discrete approximation of probabilistic automata

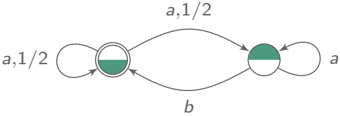

Back to motivations

Motivation 2: theoretical motivation

Discrete approximation of probabilistic automata

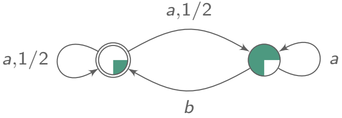

Back to motivations

Motivation 2: theoretical motivation

Discrete approximation of probabilistic automata

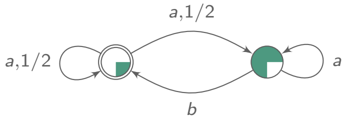

Back to motivations

Motivation 2: theoretical motivation

Discrete approximation of probabilistic automata

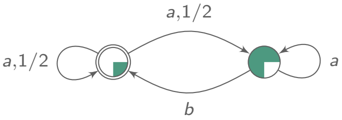

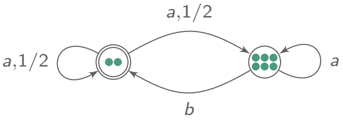

Arguable: optimal reachability probability not continuous when $N \to \infty$

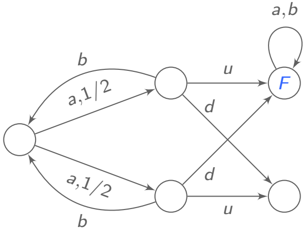

Back to motivations

Motivation 2: theoretical motivation

Discrete approximation of probabilistic automata

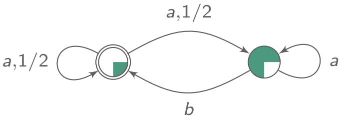

Arguable: optimal reachability probability not continuous when $N \to \infty$

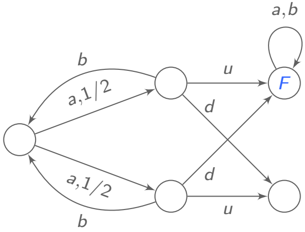

- $\forall N, \exists \sigma,\ \mathbb{P}_\sigma(\Diamond F^N) = 1.$
- In the PA, the maximum probability to reach F is .5.

Back to motivations

Motivation 2: theoretical motivation

Discrete approximation of probabilistic automata

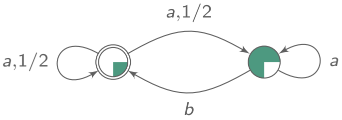

Arguable: optimal reachability probability not continuous when $N \to \infty$

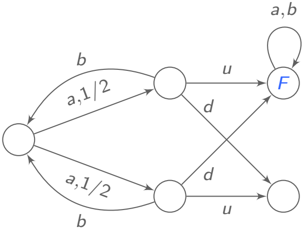

- $\forall N, \exists \sigma,\ \mathbb{P}_\sigma(\Diamond F^N) = 1$.
- In the PA, the maximum probability to reach F is .5.

Good news? hope for alternative **more decidable** semantics for PA

Thanks for your attention!