Active Diagnosis for Probabilistic Systems

Nathalie Bertrand[1] Eric Fabre[1] Stefan Haar[2] Serge Haddad[2] Loïc Hélouët[1]

[1] Inria, France

[2] LSV, ENS Cachan & CNRS & Inria, France

LIP6 seminar: Hiding and Disclosing Information

December the 2nd 2013

Outline

Outline

Observing a LTS

Events (Σ) are either *observable* (Σ_o) or *unobservable* (Σ_u).

Faults (f) are unobservable.

An execution sequence yields an *observed sequence* by $\mathcal{P} : \mathcal{L}^\omega(\mathcal{A}) \to \Sigma_o^\omega$.

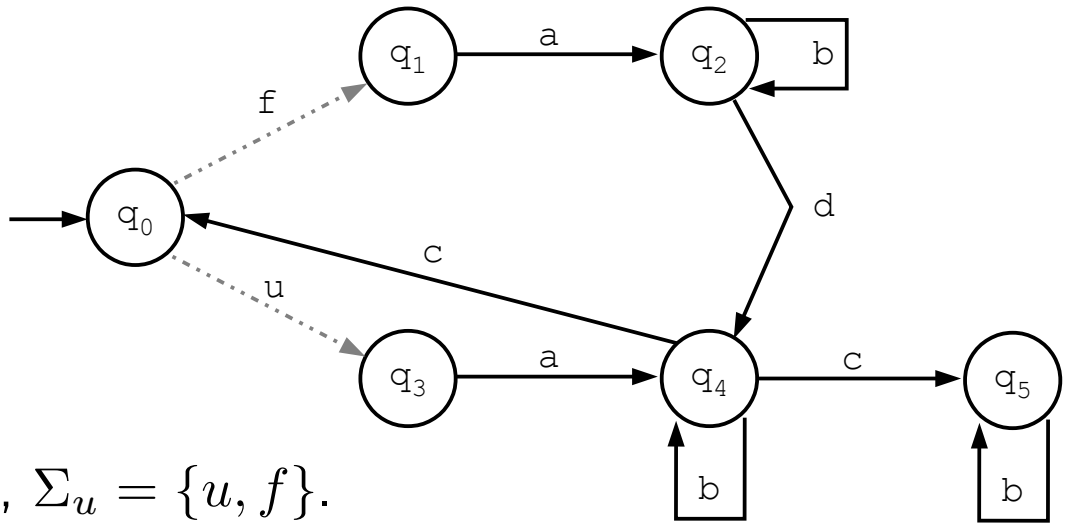

$\Sigma_o = \{a, b, c, d\}$, $\Sigma_u = \{u, f\}$.

Let $\sigma = uacfab^\omega$. Then $\mathcal{P}(\sigma) = acab^\omega$.

We only consider *convergent* systems: $\mathcal{L}^\omega(\mathcal{A}) \cap \Sigma^* \Sigma_u^\omega = \emptyset$

There is no infinite sequence of unobservable events from any reachable state.

Classification of observed sequences

An execution sequence is *faulty* if it contains a fault otherwise it is *correct*.

An observed sequence σ is *surely faulty* if for all $\sigma' \in \mathcal{P}^{-1}(\sigma)$, σ' is faulty.

An observed sequence σ is *surely correct* if for all $\sigma' \in \mathcal{P}^{-1}(\sigma)$, σ' is correct.

An observed sequence σ is *ambiguous* if it is neither surely faulty nor surely correct.

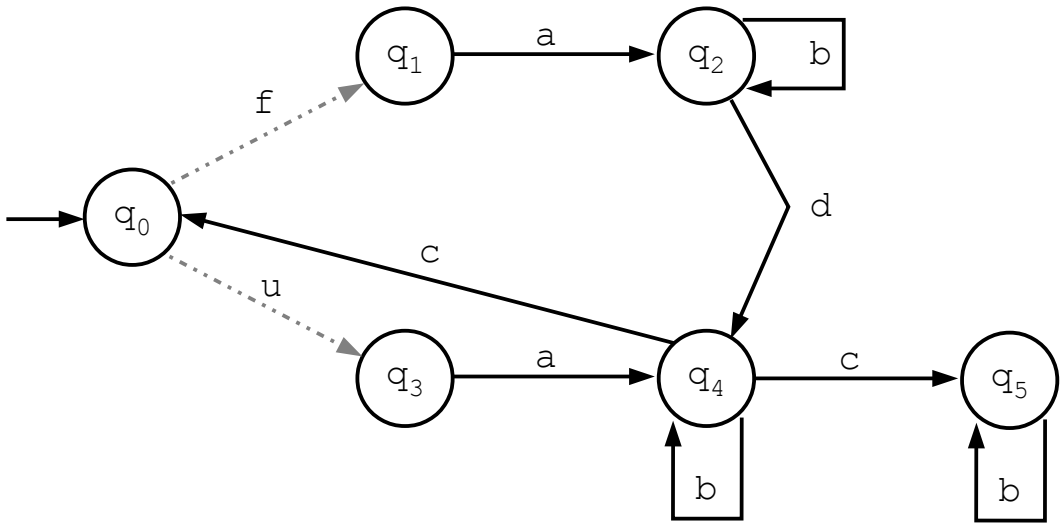

Sequence $adcb^\omega$ is surely faulty.

Sequence acb^ω is surely correct.

Sequence ab^ω is ambiguous.

Characterization of ambiguity

One can build an "optimal" deterministic Büchi automaton that accepts the unambiguous observed sequences (Haar, H, Melliti, Schwoon, FSTTCS 2013).

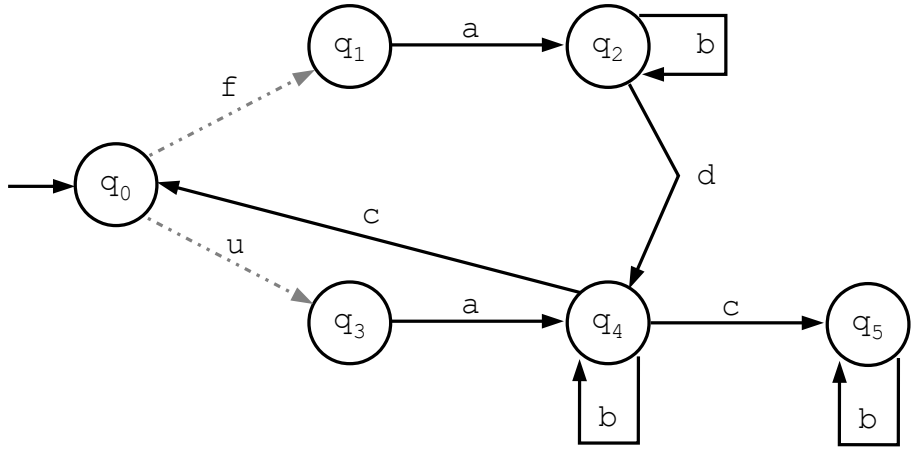

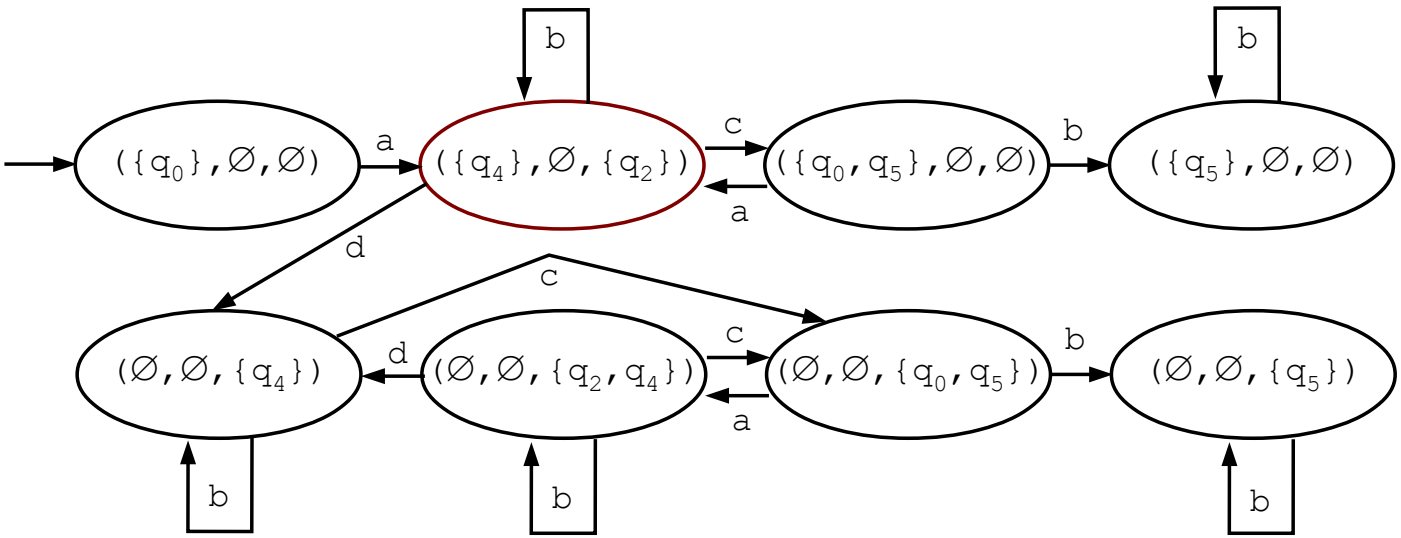

.

pLTS

A probabilistic labelled transition system (pLTS) is a *live* LTS with a transition probability matrix **P**.

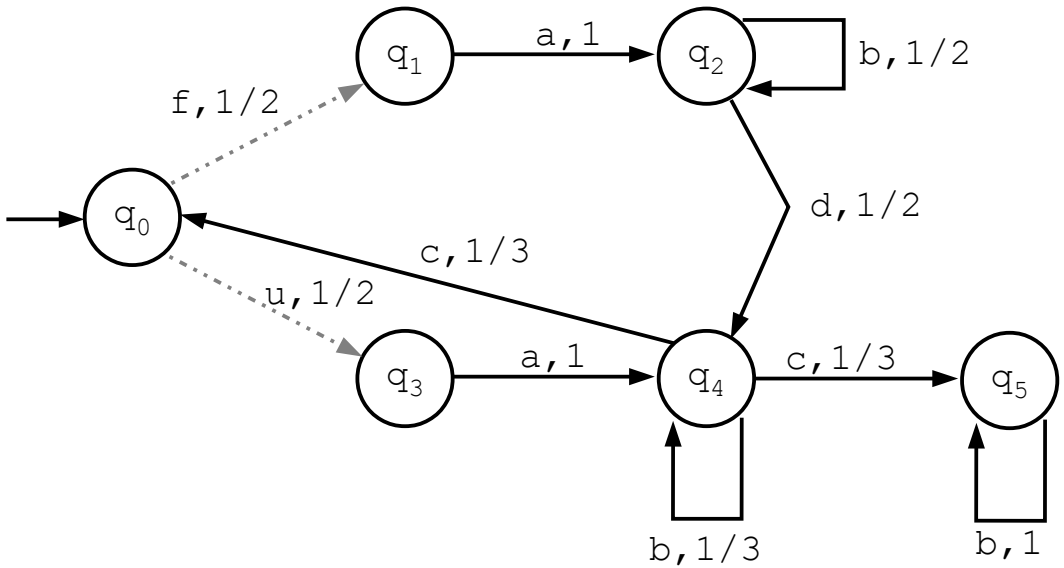

Without labels, a pLTS is a discrete time Markov chain.

Without transition probabilities, a pLTS is a LTS.

(Safe) Diagnosability

A pLTS is *diagnosable* if the set of sequences yielding ambiguous observed sequences has null measure.

A pLTS is *safely diagnosable* if it is diagnosable and the set of correct sequences has positive measure.

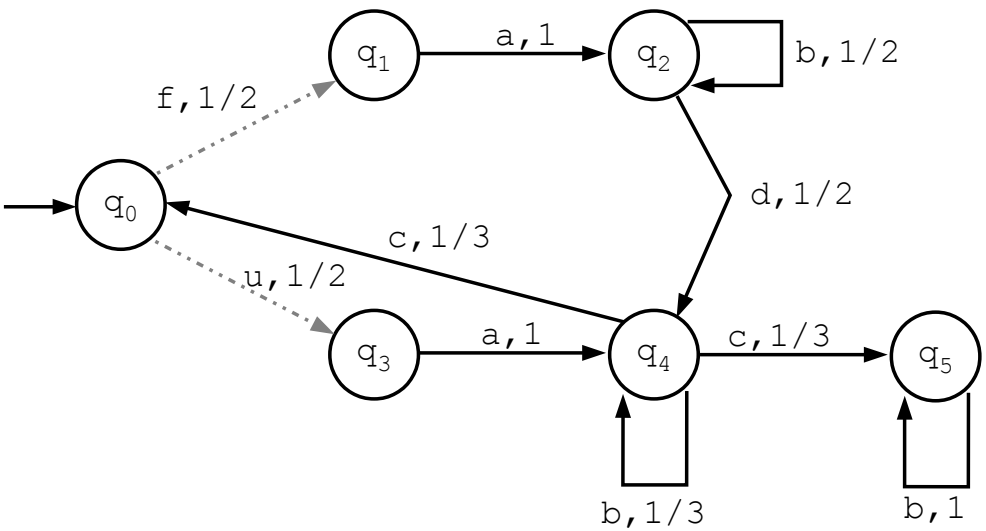

safely diagnosable

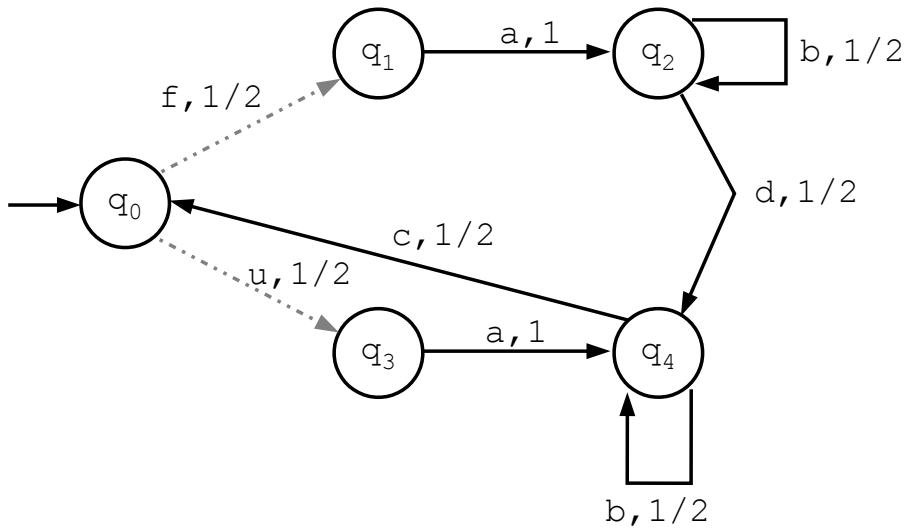

diagnosable but not safely diagnosable

cLTS

A *controllable labelled transition system* (cLTS) is a live LTS
with integer weights on transitions T
and another partition between *controllable* (Σ_c) and *uncontrollable* (Σ_e) events.

A *strategy* is a mapping $\pi : \Sigma_o^* \to \text{Dist}(2^\Sigma)$ such that:
for every $\sigma \in \Sigma_o^*$, for every $\Sigma^\bullet \in \text{Supp}(\pi(\sigma))$, $\Sigma^\bullet \supseteq \Sigma_e$.
π is *live* if it does not introduce deadlocks.

Let $\mathcal{C}$ be a cLTS and π be a live strategy. Then $\mathcal{C}_\pi$ is a pLTS where:

- the states are tuples $(\sigma, q, \Sigma^\bullet)$ (with an additional initial state) where:
  σ is the observed sequence, q is the state, $\Sigma^\bullet$ are the allowed events.
- When $a \in \Sigma^\bullet \cap \Sigma_u$ is an allowed unobservable event,
  $\mathbf{P}_\pi\left((\sigma, q, \Sigma^\bullet), a, (\sigma, q', \Sigma^\bullet)\right) = T^{\Sigma^\bullet}(q, a, q')$
- When $a \in \Sigma^\bullet \cap \Sigma_o$ is an allowed observable event,
  $\mathbf{P}_\pi\left((\sigma, q, \Sigma^\bullet), a, (\sigma a, q', \Sigma^{\bullet\prime})\right) = T^{\Sigma^\bullet}(q, a, q').\pi(\sigma.a)(\Sigma^{\bullet\prime})$.

Here $T^{\Sigma^\bullet}$ denotes the normalization of the weights w.r.t. $\Sigma^\bullet$.

Illustration

A deterministic strategy π: Forbid two consecutive b after an a.

With a finite memory abstraction, π leads to the following pLTS.

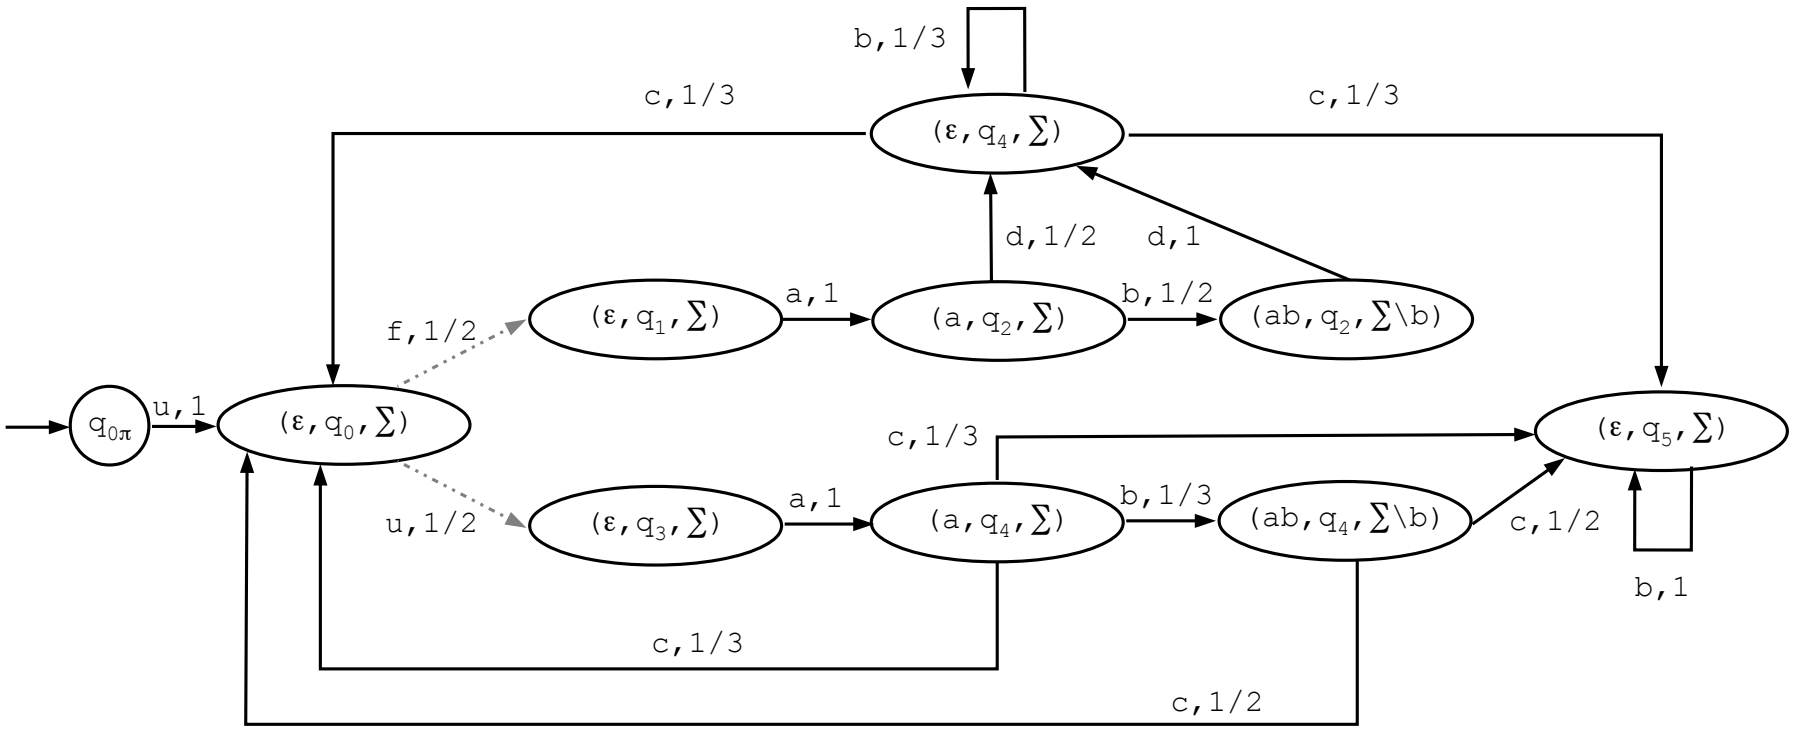

Active probabilistic diagnosis

The *active probabilistic diagnosis problem* asks, whether
there exists a live strategy π in $\mathcal{C}$ such that the pLTS $\mathcal{C}_\pi$ is diagnosable.

The *safe active probabilistic diagnosis problem* asks whether
there exists a live strategy π in $\mathcal{C}$ such that the pLTS $\mathcal{C}_\pi$ is safely diagnosable.

The *synthesis problems* consist in building a live strategy π in $\mathcal{C}$
such that the pLTS $\mathcal{C}_\pi$ is (safely) diagnosable.

Outline

Partially observed Markov decision process

A *partially observable* Markov decision process (POMDP) is a tuple $\mathsf{M}_{\mathcal{C}} = \langle Q, q_0, \mathsf{Obs}, \mathsf{Act}, T \rangle$ where

- Q is a finite set of states with q_0 the initial state;
- $\mathsf{Obs} : Q \to \mathcal{O}$ assigns an observation $\mathsf{O} \in \mathcal{O}$ to each state.
- Act is a finite set of actions;
- $T : Q \times \mathsf{Act} \to \mathsf{Dist(Q)}$ is a partial transition function. Letting $\mathsf{Ena}(q) = \{a \in \mathsf{Act} \mid T(q,a) \text{ is defined}\}$, we assume that:
  - for all $q \in Q$, $\mathsf{Ena}(q) \neq \emptyset$, and
  - whenever $\mathsf{Obs}(q) = \mathsf{Obs}(q') = \mathsf{O}$, then $\mathsf{Ena}(q) = \mathsf{Ena}(q') = \mathsf{Ena}(\mathsf{O})$.

A *decision rule* is an item of $\mathsf{Dist}(\mathsf{Act})$.

A *strategy* π maps histories of observations to decision rules: $\pi : \mathcal{O}^+ \to \mathsf{Dist}(\mathsf{Act})$ such that for all $\mathsf{O}_1 \cdots \mathsf{O}_i$, $\mathsf{Supp}(\pi(\mathsf{O}_1 \cdots \mathsf{O}_i)) \subseteq \mathsf{Ena}(O_i)$.

Given a strategy π and an initial distribution δ over states, a POMDP M becomes a (possibly infinite) pLTS denoted $\mathsf{M}(\pi)$.

From cLTS diagnosis to POMDP problems

Let $\mathcal{C}$ be a cLTS and its Büchi automaton $\mathcal{B}$, $\mathsf{M}_{\mathcal{C}}$ is built as follows.

States are pairs (l, q) with l a state of $\mathcal{B}$ and q a state of $\mathcal{C}$ with $\mathsf{Obs}(l, q) = l$.

Actions of $\mathsf{M}_{\mathcal{C}}$ are subset of events that includes the uncontrollable events.

Given some action $\Sigma^{\bullet}$, the transition probability of $\mathsf{M}_{\mathcal{C}}$ from (l, q) to (l', q') is:

- the sum of probabilities of paths in $\mathcal{C}$ from q to q';
- labelled by unobservable events of $\Sigma^{\bullet}$;
- ending with an observable event $b \in \Sigma^{\bullet}$ such that $l \xrightarrow{b}_{\mathcal{B}} l'$.

The probability of any such path is the product of the individual step probabilities.

The latter are then defined by the normalization of weights w.r.t. $\Sigma^{\bullet}$.

When in $\mathcal{C}$, some path reaches a state where no event of $\Sigma^{\bullet}$ is possible,
one reaches in $\mathsf{M}_{\mathcal{C}}$ an additional state `lost`.

Illustration

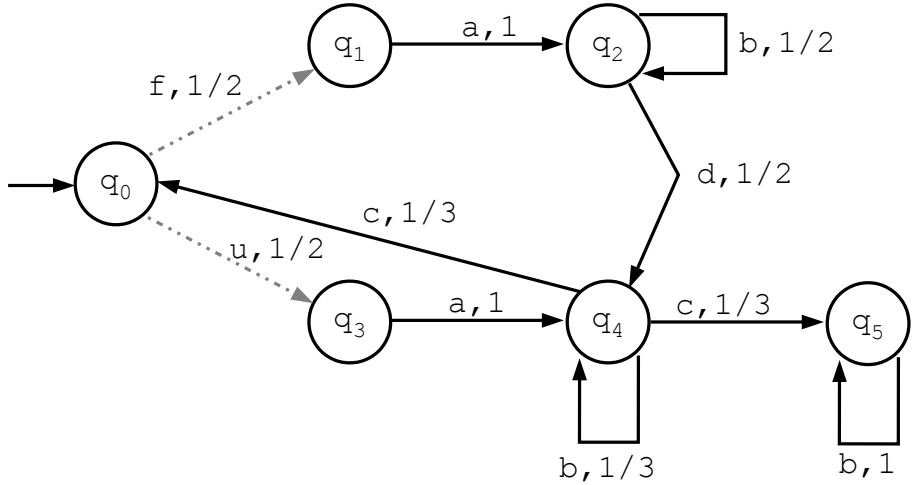

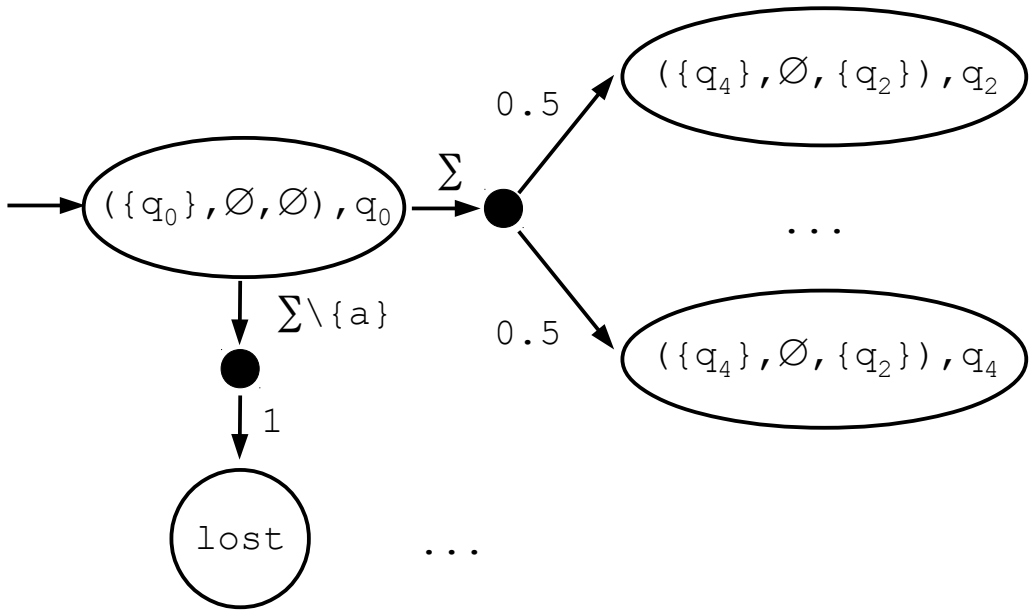

Decidability of the active diagnosis problem

$\mathcal{C}$ is actively diagnosable iff there exists a strategy π in $\mathsf{M}_{\mathcal{C}}$ such that:

$$\mathbb{P}_{\pi}^{l_0,q_0}(\mathsf{M}_{\mathcal{C}} \models \Box\Diamond(W = \emptyset \vee U = \emptyset)) = 1$$

$\mathcal{C}$ is safely actively diagnosable iff there exists a strategy π in $\mathsf{M}_{\mathcal{C}}$ such that:

$$\mathbb{P}_{\pi}^{l_0,q_0}(\mathsf{M}_{\mathcal{C}} \models \Box\Diamond(W = \emptyset \vee U = \emptyset)) = 1 \text{ and } \mathbb{P}_{\pi}^{l_0,q_0}(\mathsf{M}_{\mathcal{C}} \models \Box(U \neq \emptyset)) > 0$$

The existence of a strategy in a POMDP for almost surely satisfying a Büchi objective is decidable (Baier, Bertrand, Größer, FoSSaCS 2008).

The proof in (Bertrand, Genest, Gimbert, LICS 2009) is more general and elegant.

Complexity of the active diagnosis problem

The decision procedure for POMDP in EXPTIME is based on the construction of a *belief graph* where a *belief* is a subset of states leading to the same observation.

The initial belief is $\{s_0\}$ where s_0 is the initial state of the POMDP.

There is an edge $B \xrightarrow{\delta,\mathsf{O}} B'$ with δ a decision rule and O an observation iff

$$B' = \bigcup_{s \in B, a \in \mathsf{Supp}(\delta)} \mathsf{Supp}(T(s,a)) \cap \mathsf{Obs}^{-1}(\mathsf{O}) \neq \emptyset$$

For the graph, one restricts δ to be a deterministic decision rule.

The (possible) winning strategy is a randomized belief-based strategy.

A state $((U,V,W),q)$ of the POMDP already contains the belief, i.e. $\{((U,V,W),q') \mid q' \in U \cup V \cup W\}$. So there is a single exponential blowup.

The problem is EXPTIME-hard.

by adaptation of a lower bound from (Haar, H, Melliti, Schwoon, FSTTCS 2013).

Outline

Belief-based strategies are not enough (1)

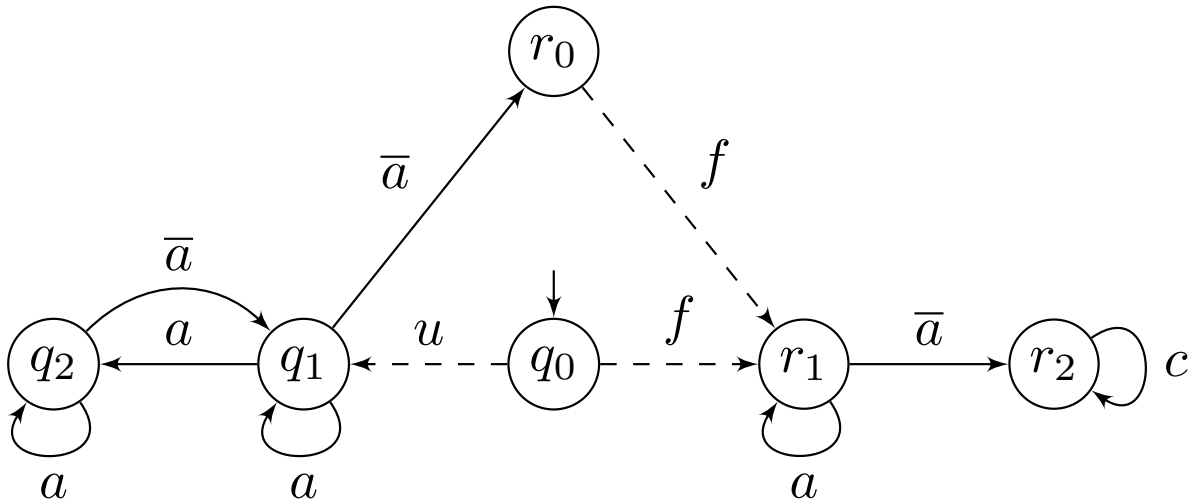

$$\Sigma_u = \{u, f\}, \Sigma_e = \{u, f, c\}$$

An observed sequence σ is surely faulty iff $\sigma \in \Sigma_o^* c^\omega$.

An observed sequence σ is surely correct iff $\sigma \in (a^+\overline{a})^\omega$.

Belief-based strategies are not enough (2)

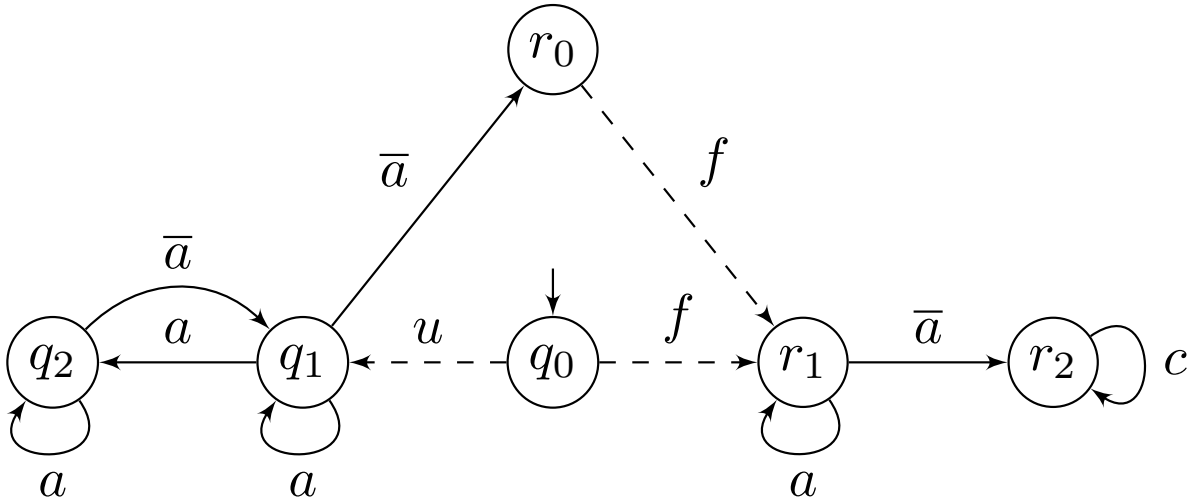

A winning strategy

Pick any sequence of positive integers $\{\alpha_i\}_{i\geq 1}$ such that $\prod_{i\geq 1} 1 - 2^{-\alpha_i} > 0$.

Let $A = \{a\} \cup \Sigma_e$ and $\overline{A} = \{\overline{a}\} \cup \Sigma_e$.

Let π be the strategy that consists in selecting, at instant n, the n^{th} subset in the following sequence $A^{\alpha_1}\overline{A}A^{\alpha_2}\overline{A}\ldots$.

Then π is winning:

- All observed sequences are either surely faulty or surely correct.
- The probability that a sequence is correct is $\frac{1}{2}\prod_{i\geq 1} 1 - 2^{-\alpha_i} > 0$.

Belief-based strategies are not enough (3)

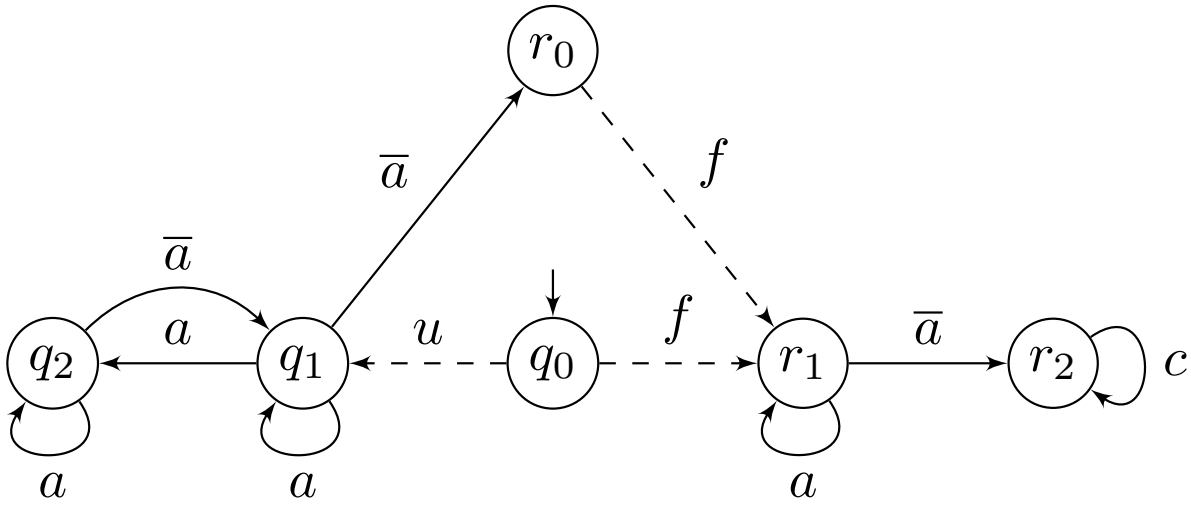

A belief-based strategy π

Three possible controls: A, $\overline{A}$ and Σ.

Let $\pi(\{q_1, q_2, r_1\}) = \mathbf{p} \stackrel{\text{def}}{=} p_A \cdot A + p_{\overline{A}} \cdot \overline{A} + p_\Sigma \cdot \Sigma$.

If $p_A = 1$ then the possible first fault remains undetected.

Let $\pi(\{q_1, r_0, r_2\}) = \mathbf{p}' \stackrel{\text{def}}{=} p'_A \cdot A + p'_{\overline{A}} \cdot \overline{A} + p'_\Sigma \cdot \Sigma$.

If $p'_{\overline{A}} = 1$ then at the next instant there is no possible correct sequence.

Let $\alpha q_1 + \beta r_0 + (1 - \alpha - \beta) r_2$ be a distribution then after the next occurrence of $\overline{a}$ the distribution is: $\alpha_{\mathbf{p},\mathbf{p}'} \alpha q_1 + (1 - \alpha_{\mathbf{p},\mathbf{p}'}) \alpha r_0 + (1 - \alpha) r_2$, where $\alpha_{\mathbf{p},\mathbf{p}'} < 1$.

After n $\overline{a}$ the probability of a correct sequence is bounded by $\alpha^n_{\mathbf{p},\mathbf{p}'}$. So π is losing.

From blind POMDP to safe active diagnosis

The existence of an infinite word accepted by a Büchi probabilistic automaton with positive probability is undecidable (Baier, Bertrand, Größer, Fossacs 2008).

The existence of a winning strategy with positive probability for a Büchi objective in a *blind* POMDP (i.e. without observation) is undecidable (Chatterjee, Doyen, Gimbert, Henzinger, MFCS 2010).

We reduce the latter problem to a safe active diagnosability problem.

Corollary.

The problem whether, given a POMDP M with subsets of states F and I, there exists a strategy π with $\mathbb{P}_\pi(\mathsf{M} \models \Box\Diamond F) = 1$ and $\mathbb{P}_\pi(\mathsf{M} \models \Box I) > 0$, is undecidable.

Observation: The existence of a strategy for each objective is decidable.

Scheme of the reduction

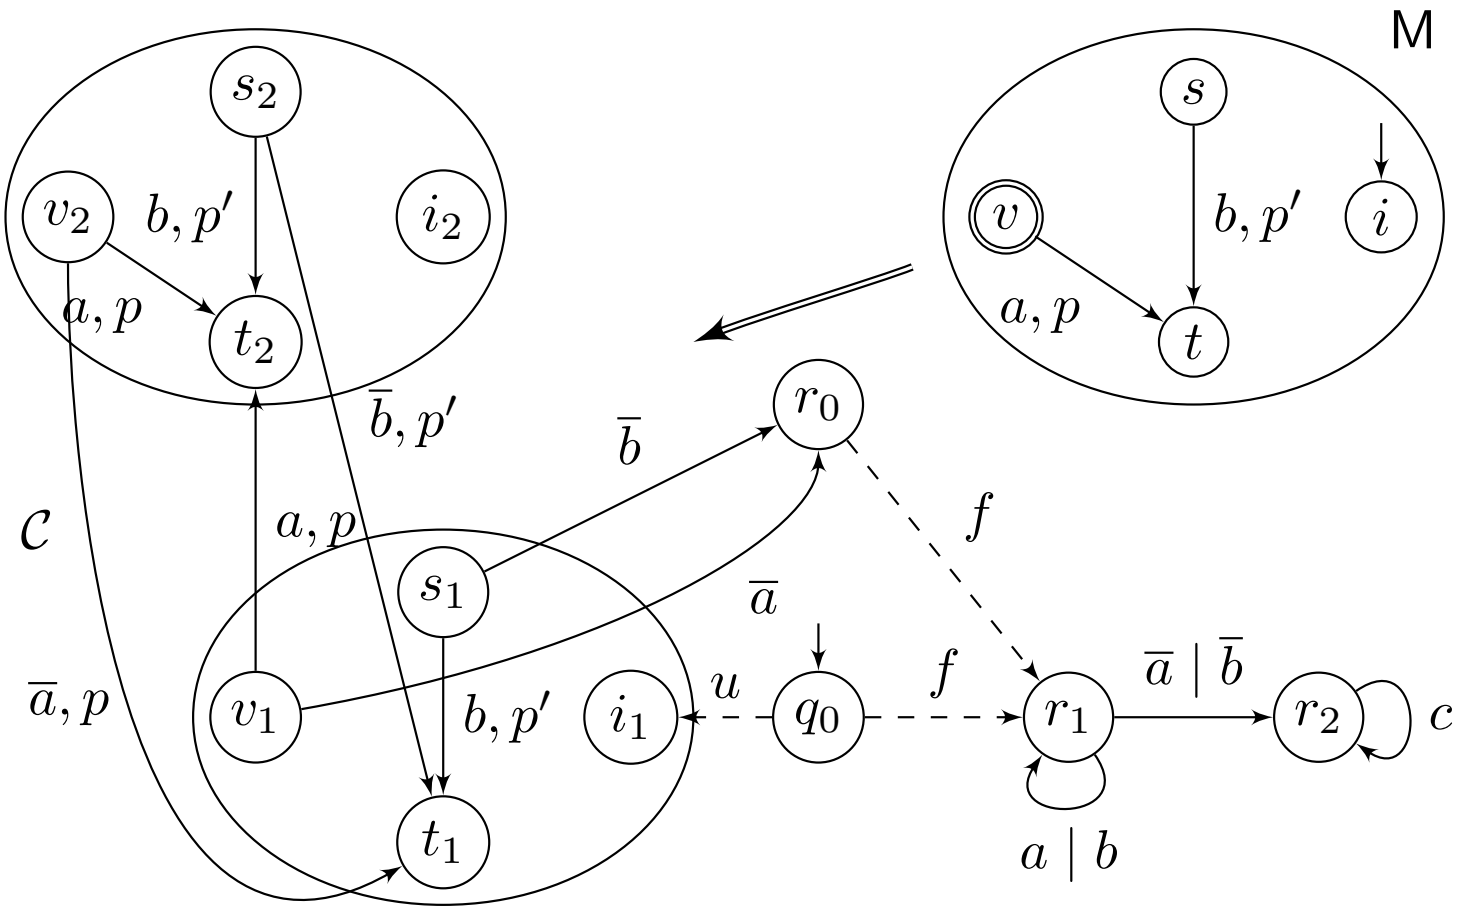

An observed sequence σ is surely faulty iff $\sigma \in \Sigma_o^* c^\omega$.

An observed sequence σ is surely correct iff $\sigma \in ((a+b)^+(\overline{a}+\overline{b}))^\omega$.

Correctness of the reduction (1)

From POMDP to cLTS

Let $\pi = \sigma_1\sigma_2\ldots$ be a deterministic winning strategy of M for the Büchi goal with positive probability.

Let $p \stackrel{\text{def}}{=} \mathbb{P}^i_\sigma(\mathsf{M} \models \Box\Diamond F) > 0$. Pick $(\beta_j)_{j\in\mathbb{N}}$ with $0 < \beta_j < 1$ and $\prod_{j\geq 0}\beta_j > 0$.

Iteratively build an infinite increasing sequence $(n_j)_{j\in\mathbb{N}}$ such that:

$$\mathbb{P}^i_\sigma\left(\mathsf{M} \models \Diamond^{[n_j+1,n_{j+1}]}F \mid \mathsf{M} \models \bigwedge_{k=0}^{j-1}\Diamond^{[n_k+1,n_{k+1}]}F \wedge \Box\Diamond F\right) \geq \beta_j$$

By construction: $\mathbb{P}^i_\sigma\left(\mathsf{M} \models \bigwedge_{j\geq 0}\Diamond^{[n_j+1,n_{j+1}]}F\right) \geq p\prod_{j\geq 0}\beta_j > 0$

Let π' the strategy of $\mathcal{C}$ that:

- at time instant k different from any n_j, π selects X with $x = \sigma_k$,
- at time instant n_j, π selects $\overline{X}$ with $x = \sigma_{n_j}$.

with $X = \{x\} \cup \Sigma_e$ and $\overline{X} = \{\overline{x}\} \cup \Sigma_e$ for $x \in \{a, b\}$.

Then π' is winning:

- All observed sequences are either surely faulty or surely correct.
- The probability that a sequence is correct is $\geq \frac{p}{2}\prod_{j\geq 0}\beta_j > 0$.

Correctness of the reduction (2)

From cLTS to POMDP

Let π be a winning strategy in $\mathcal{C}$.

Let Ex be the set of executions associated with unambiguous sequences.

Let $\mathsf{Ex}' \subseteq \mathsf{Ex}$ the subset of correct executions. $\mathbb{P}^{q_0}_{\pi}(\mathsf{Ex}') > 0$.

An execution of Ex' only visits states of $Q_1^{\mathsf{M}} \cup Q_2^{\mathsf{M}}$ with infinitely often actions in $\{\overline{a}, \overline{b}\}$ and so visits infinitely often F_1.

Define $\mathsf{proj}(x) = \mathsf{proj}(\overline{x}) = x$ for $x \in \{a, b\}$ and $\mathsf{proj}(q_x) = q$ for $x \in \{1, 2\}$.

Any sequence of $\mathsf{proj}(\mathsf{Ex}')$ visits infinitely often F.

Let $\sigma \in \{a, b\}^*$ be a sequence such that $\mathbb{P}^{q_0}_{\pi}(\mathcal{C} \models \rho \in \mathsf{proj}^{-1}(\sigma)) > 0$
where ρ is a random sequence of length $|\sigma|$.

Let $\sigma' \in \mathsf{proj}^{-1}(\sigma)$. Define $p_{\sigma'} \stackrel{\mathsf{def}}{=} \mathbb{P}^{q_0}_{\pi}(\mathcal{C} \models \rho = \sigma' \mid \mathsf{proj}(\rho) = \sigma)$.

Let π' be a winning strategy for M defined by $\pi'(\sigma) = \sum_{\sigma' \in \mathsf{proj}^{-1}(\sigma)} p_{\sigma'} \pi(\sigma')$.

By construction and induction: $\mathbb{P}^{i}_{\pi'}(\mathsf{M} \models \rho = \sigma) = \sum_{\sigma' \in \mathsf{proj}^{-1}(\sigma)} \mathbb{P}^{q_0}_{\pi}(\mathcal{C} \models \rho = \sigma')$.

So $\mathbb{P}^{i}_{\pi'}(\mathsf{M} \models \mathsf{proj}(\mathsf{Ex}')) = \mathbb{P}^{q_0}_{\pi}(\mathcal{C} \models \mathsf{Ex}') > 0$. Thus π' is a winning strategy.

Restriction to belief-based strategies

Observation

The status of a randomized strategy only depends on the support of the distributions.

Decision procedure in NEXPTIME

- Guessing the supports of a potential winning strategy.
- Checking the status of this strategy in the underlying graph of an appropriate DTMC.

The problem is EXPTIME-hard.

by adaptation of a lower bound from (Haar, H, Melliti, Schwoon, FSTTCS 2013).

Conclusion and perspectives

Contributions

- Introduction of (safe) active diagnosis problems for probabilistic systems.
- Analysis of the problems using a POMDP framework.

Perspectives

- Closing the gap between lower and upper bounds related to the existence of winning belief-based strategies for safe active diagnosability.
- Introducing the active predictability problem.
- Investigating further POMDP problems with multiple objectives.