Active Diagnosis

Serge Haddad

LSV, ENS Cachan & CNRS & Inria, France

GT-Vérif 2014, LIP6

June the 16th 2014

joint work with Nathalie Bertrand[2], Eric Fabre[2], Stefan Haar[1,2], Loïc Hélouët[2], Tarek Melliti[1], Stefan Schwoon[1]

(1) FSTTCS 2013 and (2) FOSSACS 2014

Diagnosis: from failures to faults

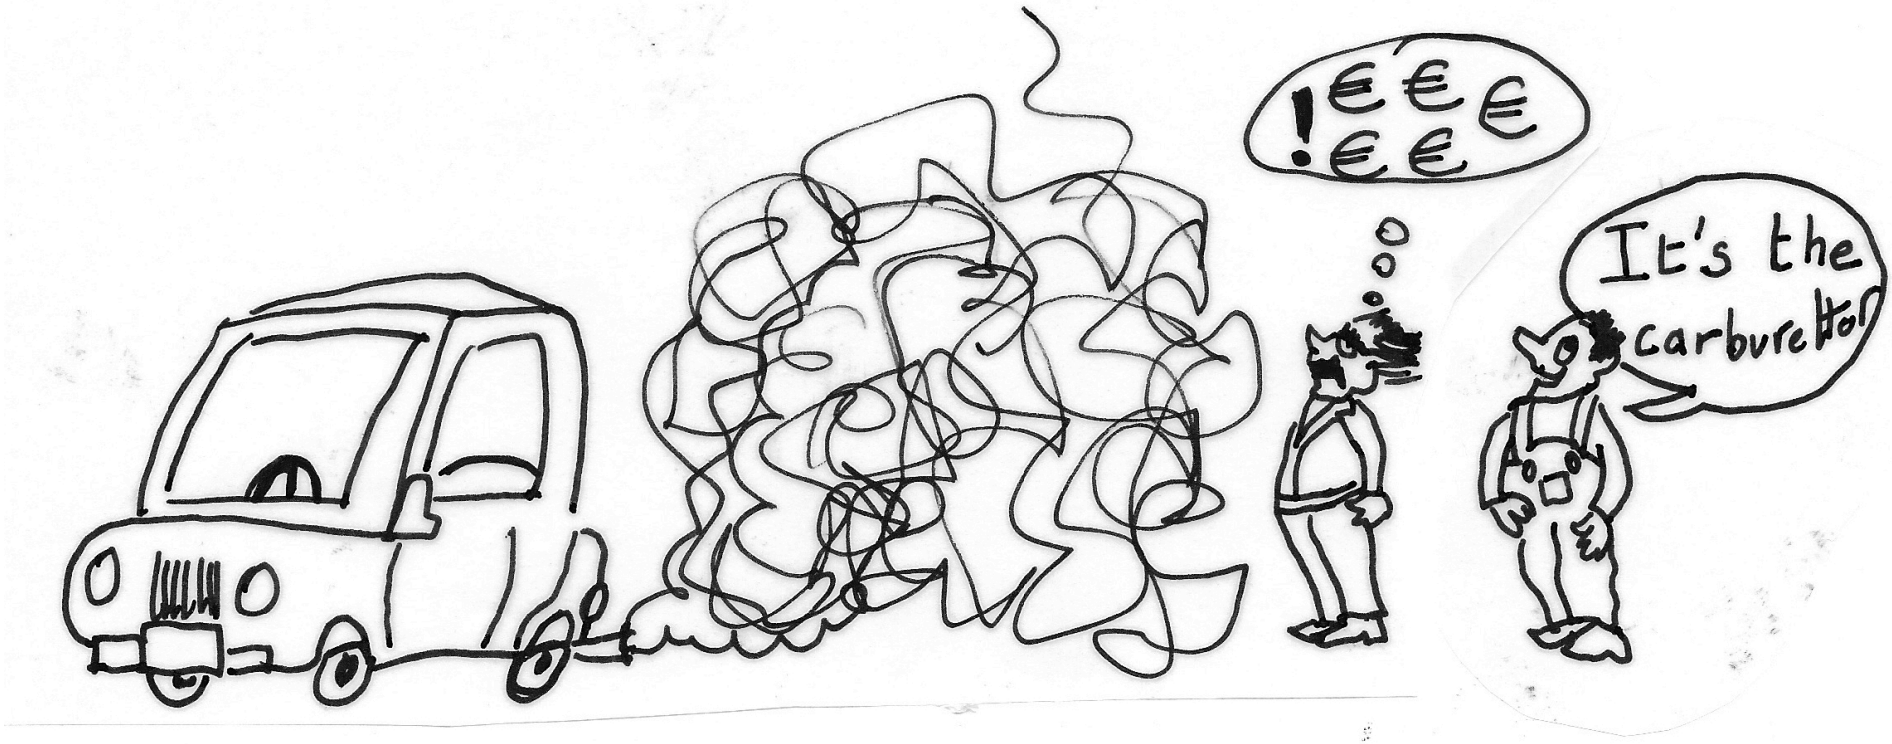

Example: MYCIN, an expert system, that used *artificial intelligence* to identify bacteria causing severe infections (1975).

Diagnosis: detecting faults

Fault detection: a subfield of *control engineering* which concerns itself with monitoring a system, identifying when a fault has occurred, and pinpointing the type of fault and its location.

Diagnosis: predicting faults

Enhancing reactivity: *see the talk of Engel Lefaucheux.*

Active diagnosis: forcing detection

Combining control and diagnosis

Outline

Observing a Labelled Transition System

States are *unobservable*.

Events are either observable or unobservable.

Faults $(\mathbf{f})$ are unobservable.

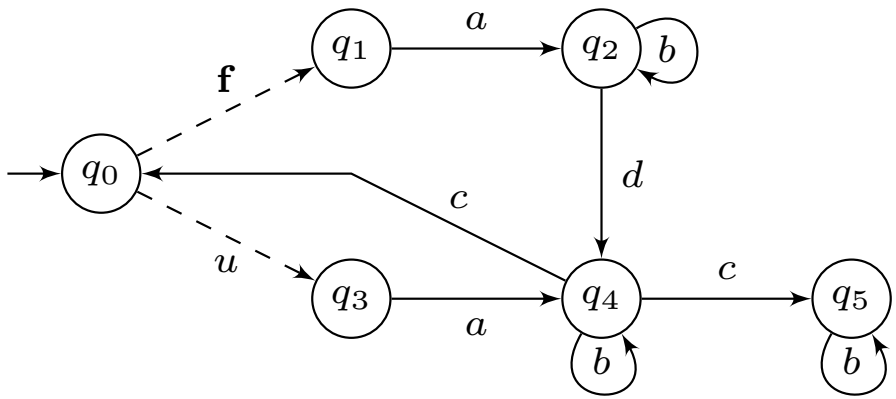

An *execution sequence* yields an *observed sequence*.

Let $\sigma = q_0 u q_3 a q_4 c q_0 \mathbf{f} q_1 a (q_2 b)^\omega$. Then $\mathcal{P}(\sigma) = acab^\omega$.

We only consider *live* and *convergent* systems:

- There is at least an event from any state.
- 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.

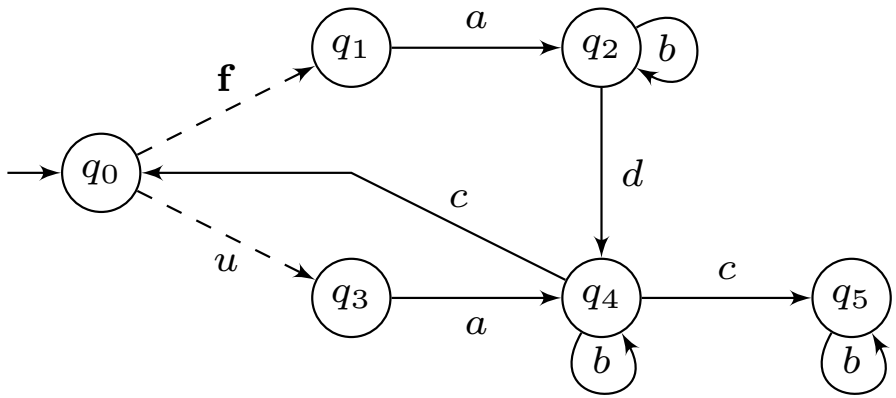

$adcb^\omega$ is surely faulty: the occurrence of d implies the occurrence of $\mathbf{f}$.

acb^ω is surely correct: $\mathcal{P}^{-1}(acb) = \{q_0 u q_3 a q_4 c q_5 b q_5\}$.

ab^ω is ambiguous: $\mathcal{P}^{-1}(ab^\omega) = \{q_0 u q_3 a (q_4 b)^\omega, q_0 \mathbf{f} q_1 a (q_2 b)^\omega\}$.

How to determine unambiguous sequences?

• Build a Büchi automaton as a synchronized product of the LTS with fault memory and the LTS without faults.

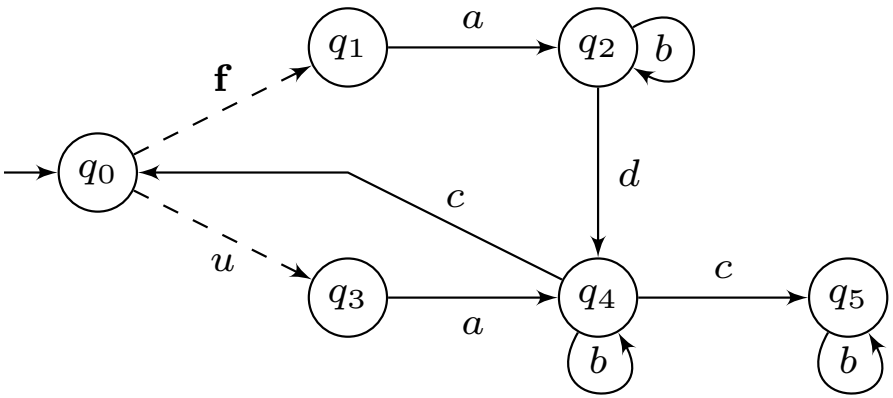

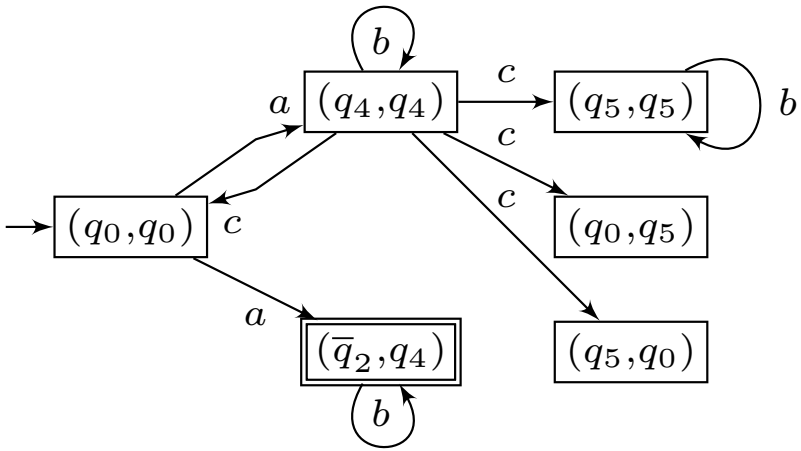

• Determinize and complement it as:

- a Street automaton with $2^{\mathcal{O}(n^2 \log(n))}$ states where n is the number of states of the LTS.
- a Büchi automaton with 3^{2n^2} states using the breakpoint construction of Miyano and Hayashi appropriate for the initial Büchi automaton.

An optimal characterization

Build a deterministic Büchi automaton whose states are triples (U, V, W) with:

- U the set of possible states reached by a correct sequence;
- W the set of possible states reached by an earliest faulty sequence;
- V the set of other possible states reached by faulty sequences.

The accepting states are (U, V, W) with:

- $U = \emptyset$, i.e. the observed sequence is (and will remain) surely faulty;
- $W = \emptyset$, i.e. the earliest faulty sequences are discarded.

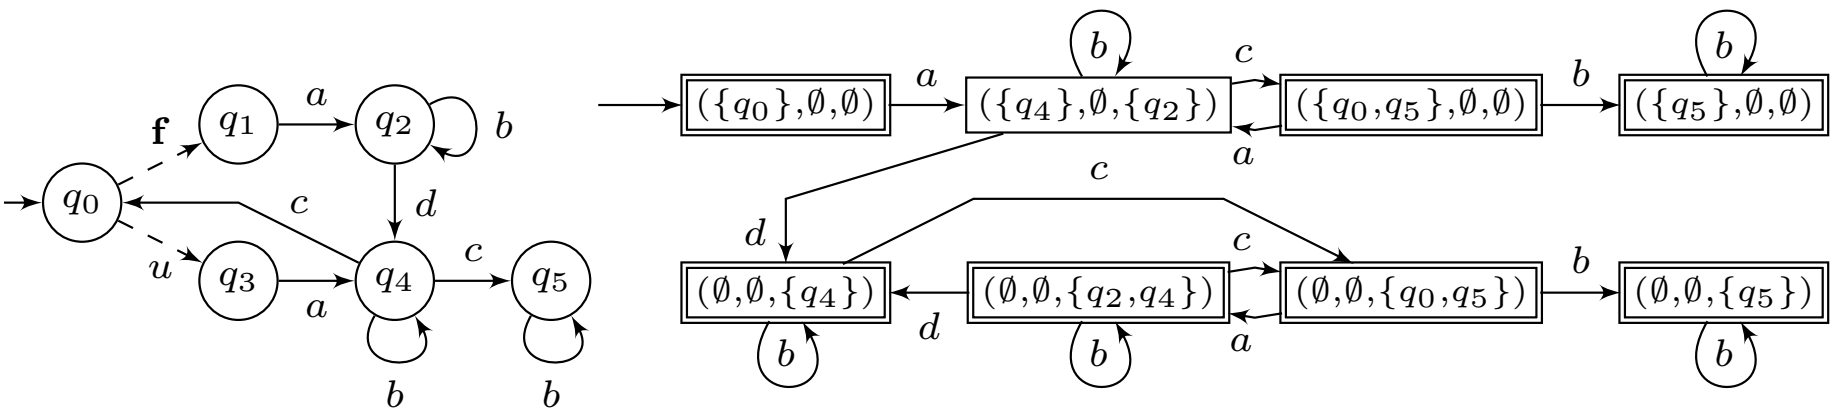

The number of states is at most 7^n.

A lower bound for ambiguity

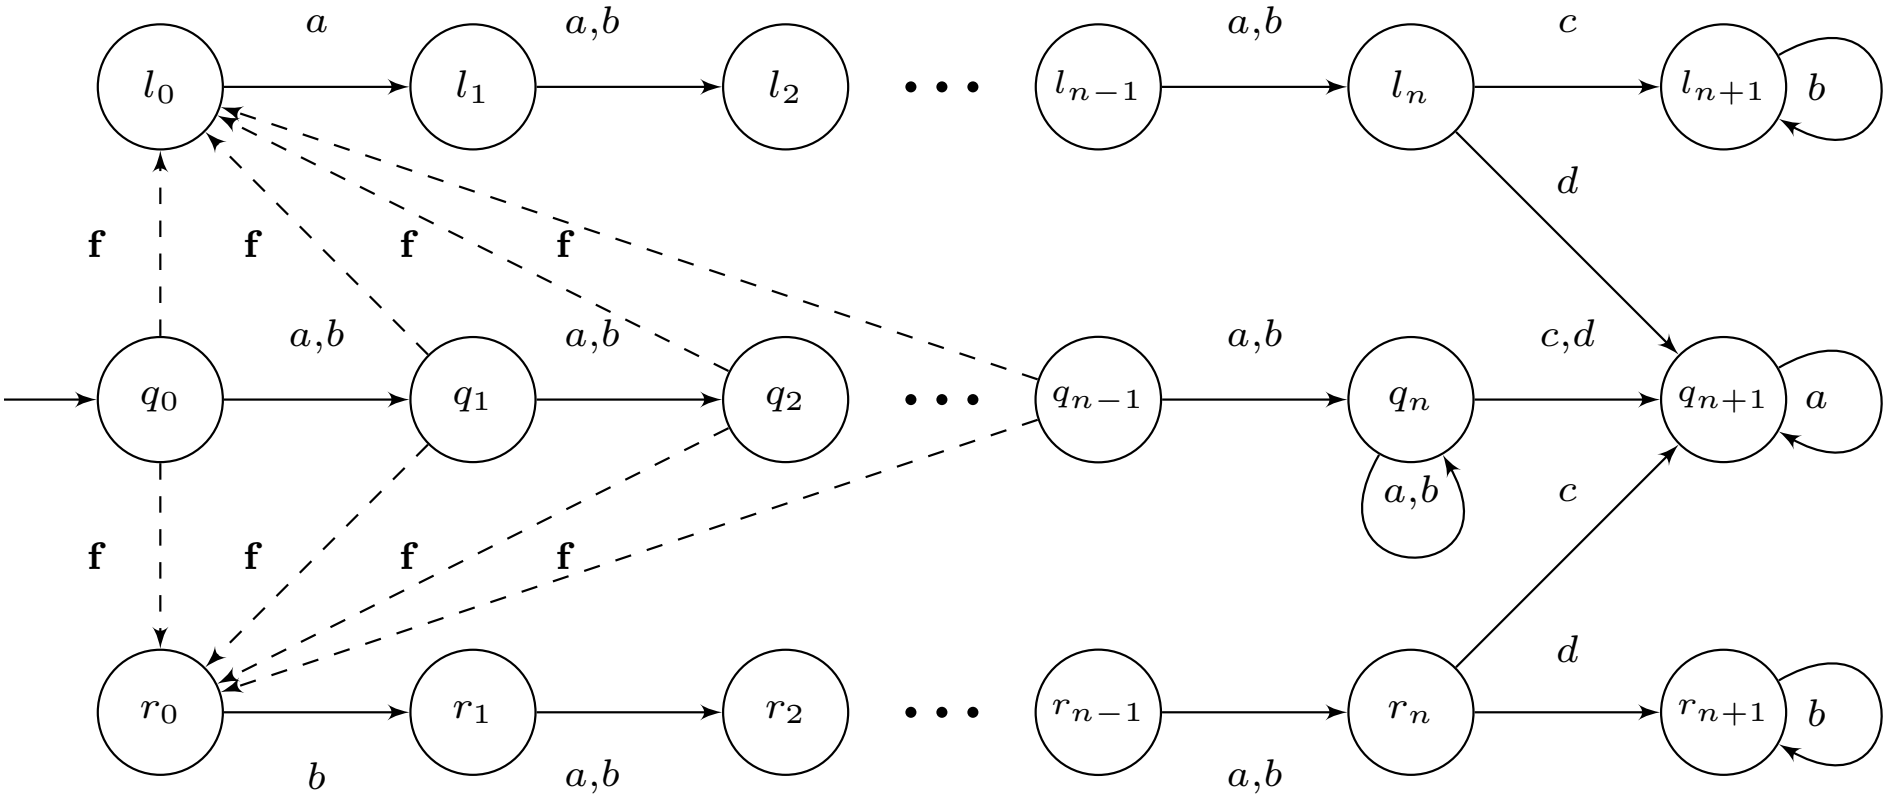

Ambiguous sequences are either $\{a,b\}^k a \{a,b\}^{n-1} d a^\omega$ or $\{a,b\}^k b \{a,b\}^{n-1} c a^\omega$ (with $0 \leq k \leq n-1$).

So an automaton for ambiguity must have (at least) 2^n states reachable after n events.

Outline

Controllable LTS and active diagnoser

Events are also partitioned in *controllable* and *uncontrollable* events.

A *controller* forbids controllable events depending on the current observed sequence.

An *active diagnoser* is a controller such that the controlled LTS:

- is still live;
- does not contain ambiguous sequences.

The *delay* of an active diagnoser is the maximal number of event occurrences between a execution sequence is faulty and an observed sequence is surely faulty.

An example of active diagnoser

The ambiguous sequences are $\{a, b\}^* b^\omega$.

The (finite-state) active diagnoser forbids two consecutive 'b'.

Its delay is 3 (at most an occurrence of bac).

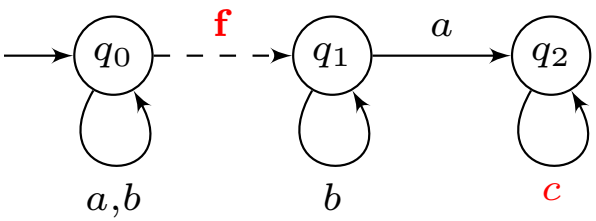

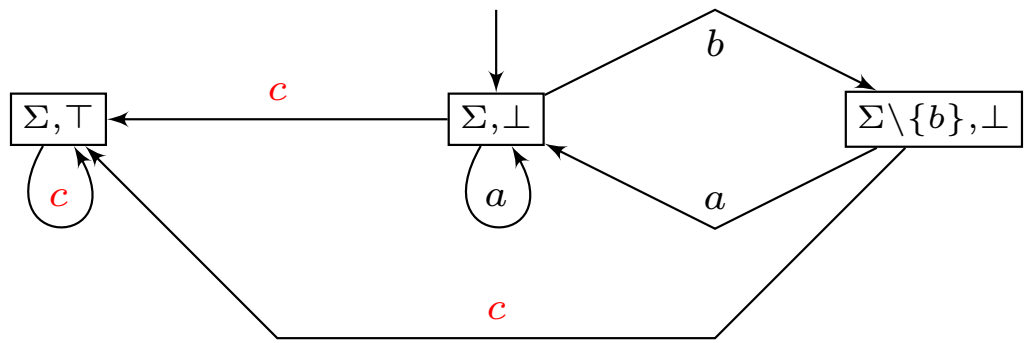

Active diagnosis problems

• The *active diagnosis decision problem*, i.e. decide whether a LTS is actively diagnosable.

• The *synthesis problem*, i.e. decide whether a LTS is actively diagnosable and in the positive case build an active diagnoser.

• The *minimal-delay synthesis problem*, i.e. decide whether a LTS is actively diagnosable and in the positive case build an active diagnoser with minimal delay.

Büchi games

A two-player (**I** and **II**) *Büchi game* is defined by:

- A graph (V, E) whose vertices are owned by players with accepting vertices F;
- In a vertex v owned by a player, he selects an edge (v, w) and the game goes on with w as current vertex.
- Player **I** wins if Player **II** is stuck in a dead vertex or the infinite path infinitely often visits F.

Game problems:

- Does there exists a *winning strategy* for Player **I**?
- In the positive case how to build such a strategy?

Classical results:

- The decision problem is PTIME-complete.
- In the positive case, there is a *positional* winning strategy.

A Büchi game for active diagnosis

Vertices of the game

- The vertices of Player **I** are the states of the Büchi automaton.
- The vertices of Player **II** are pairs of states of the Büchi automaton and (subsets of) events of the LTS.
- The accepting vertices are the accepting states of the Büchi automaton.

Edges of the game

- There is an edge $((U, V, W), ((U, V, W), \Sigma^\bullet))$ if $\Sigma^\bullet$ is a subset of events (including the uncontrollable ones) such that from all state of $U \cup V \cup W$, there is an observed sequence labelled by some $e \in \Sigma^\bullet$.
- There is an edge $(((U, V, W), \Sigma^\bullet), ((U, V, W), e)$ if $e \in \Sigma^\bullet$.
- There is an edge $(((U, V, W), e), (U', V', W')$ if there is a transition $(U, V, W) \xrightarrow{e} (U', V', W')$ in the Büchi automaton.

Example of a Büchi game

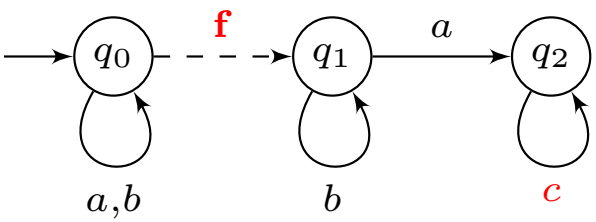

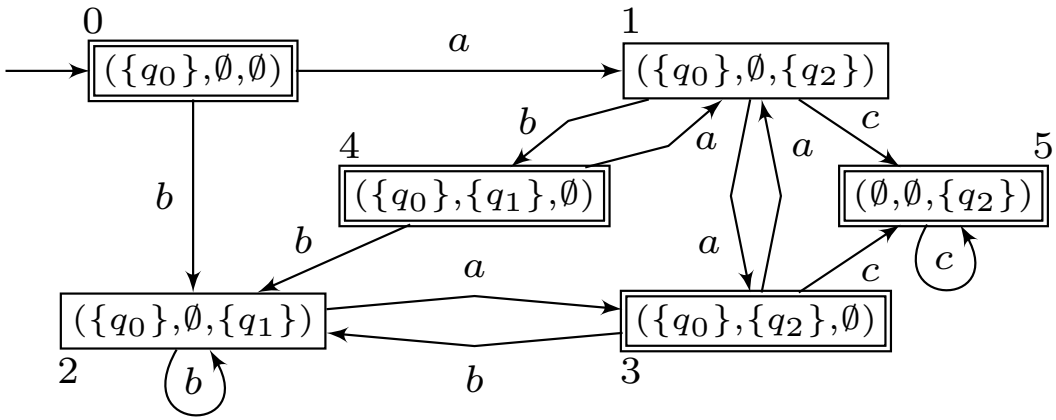

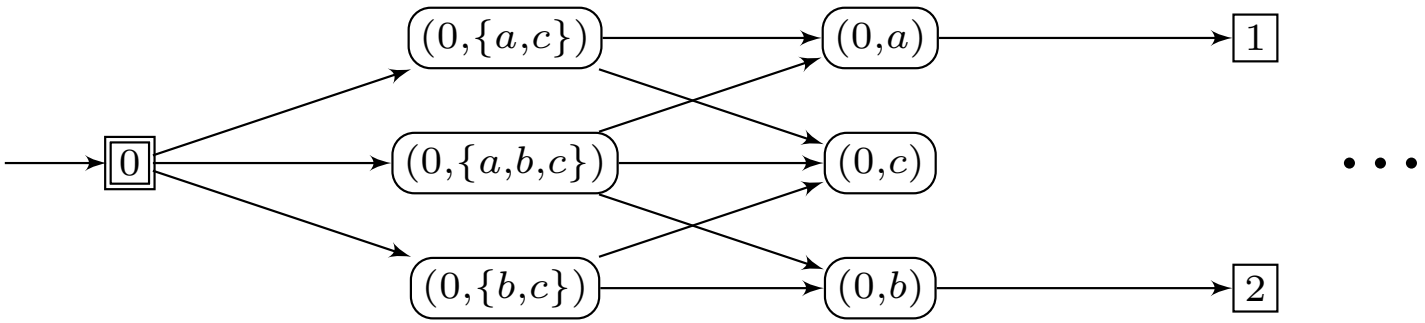

Results of this construction

Correspondence between problems

- There is a winning strategy for Player **I**
  if and only if there is an active diagnoser.
- The states of this active diagnoser are the states of the Büchi automaton.

Consequences

- The decision problem is EXPTIME-complete (*the lower bound holds by reduction from safety games with partial observation* D. Berwanger and L. Doyen FSTTCS 2008).
- The synthesis algorithm yields an active diagnoser with $2^{\mathcal{O}(n)}$ states.
  The previous synthesis algorithm yields a doubly exponential number of states (M. Sampath, S. Lafortune, and D. Teneketzis, IEEE TAC 1998).
- For all $n \in \mathbb{N}$, there is a LTS with n states such that any active diagnoser requires $2^{\Omega(n)}$ states.

A lower bound for the synthesis problem

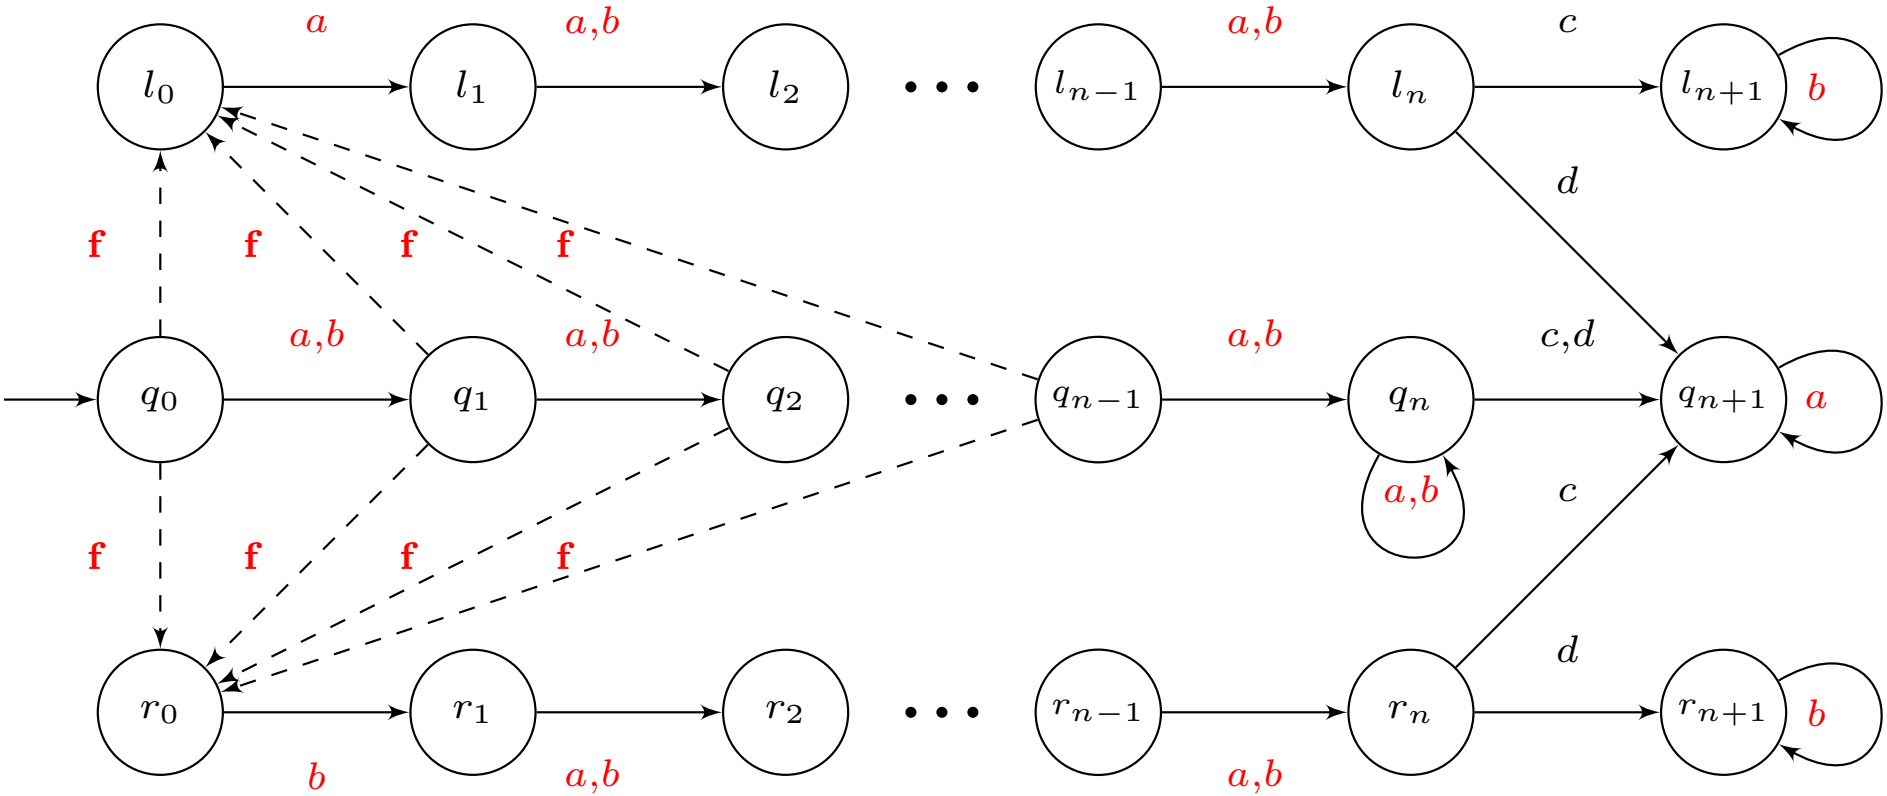

An active diagnoser must forbid a d (resp. c) if it has observed an a (resp. b) n times before.

So an active diagnoser must have (at least) 2^n states reachable after n observable events.

What about minimal delay synthesis?

Our synthesis algorithm provides a delay at most twice the minimal delay.

For all $n \in \mathbb{N}$, there is a LTS with n states such that
any active diagnoser with minimal delay requires $2^{\Omega(n \log(n))}$ states.

We have designed a synthesis algorithm of an active diagnoser with minimal delay that requires $2^{\mathcal{O}(n^2)}$ states.

Outline

pLTS

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

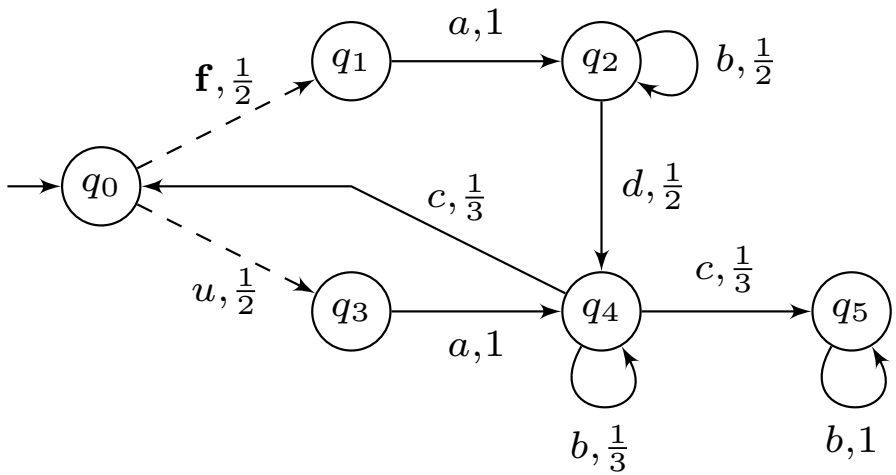

Without labels, a pLTS is a discrete time Markov chain.

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.

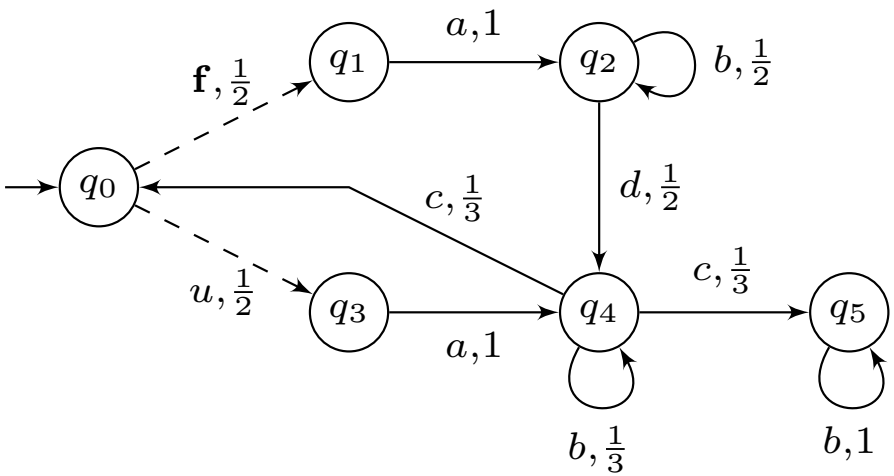

safely diagnosable

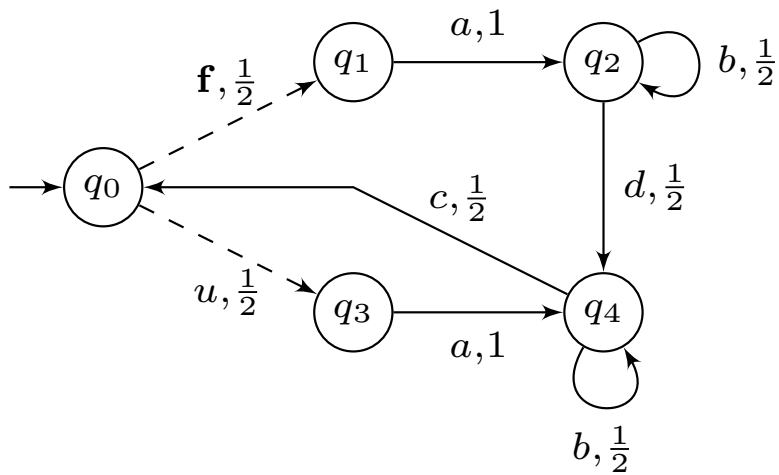

diagnosable but not safely diagnosable

cLTS

A *controllable probabilistic labelled transition system* (cLTS) is a live pLTS with integer weights on transitions.
and a partition between controllable and uncontrollable events.

An controller forbids controllable events depending on the current observed sequence. It can *randomly* select the forbidden events.
A controller must not introduce deadlocks.

Let $\mathcal{C}$ be a cLTS and π be a controller. Then $\mathcal{C}_\pi$ is a pLTS where the probability are obtained by normalization among the allowed events.

Controller π is a (safe) active diagnoser if $\mathcal{C}_\pi$ is (safely) diagnosable.

Illustration

A *deterministic* active diagnoser π:
Forbid two consecutive b after an a.

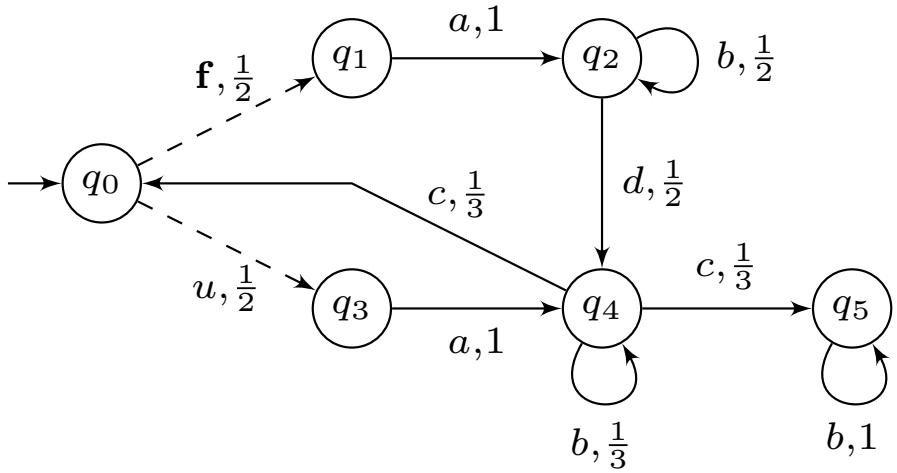

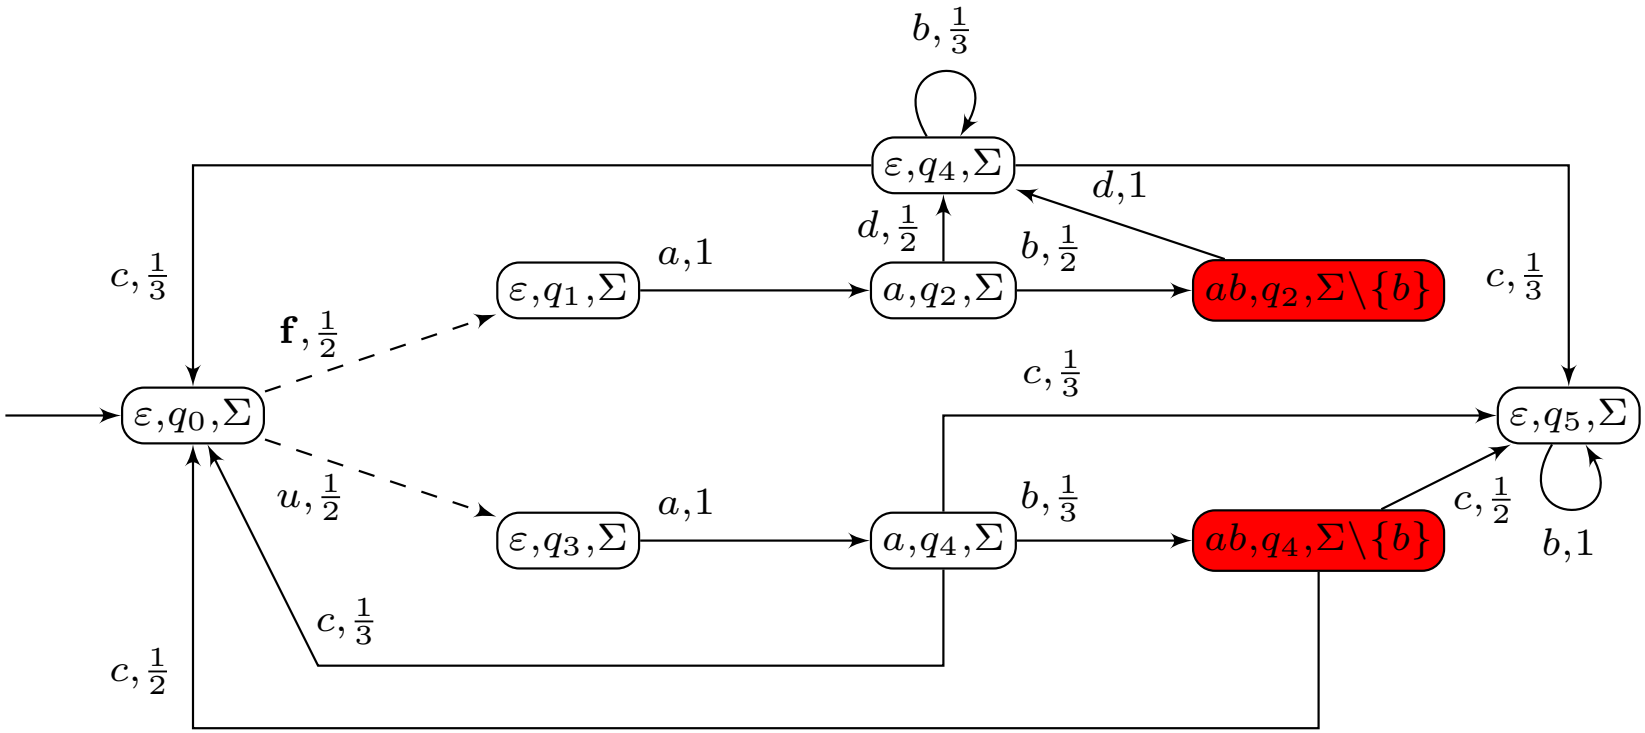

Active probabilistic diagnosis

The *active probabilistic diagnosis problem* asks whether there exists an active diagnoser π for $\mathcal{C}$.

The *safe active probabilistic diagnosis problem* asks whether there exists a safe active diagnoser π for $\mathcal{C}$.

The *synthesis problems* consist in building a (safe) active diagnoser π for $\mathcal{C}$ in the positive case.

Outline

Partially observed Markov decision process

A *partially observable* Markov decision process (POMDP) is a tuple $\mathsf{M} = \langle Q, q_0, \mathsf{Obs}, \mathsf{Act}, T \rangle$ 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;
- $T : Q \times \mathsf{Act} \to \mathsf{Dist}(\mathsf{Q})$ is a partial transition function.

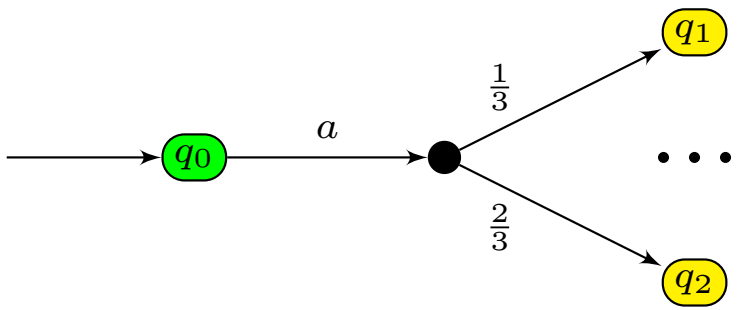

Given a sequence of observations, a *strategy* randomly selects an action to be performed.

Given a strategy, a POMDP becomes a (possibly infinite) pLTS.

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

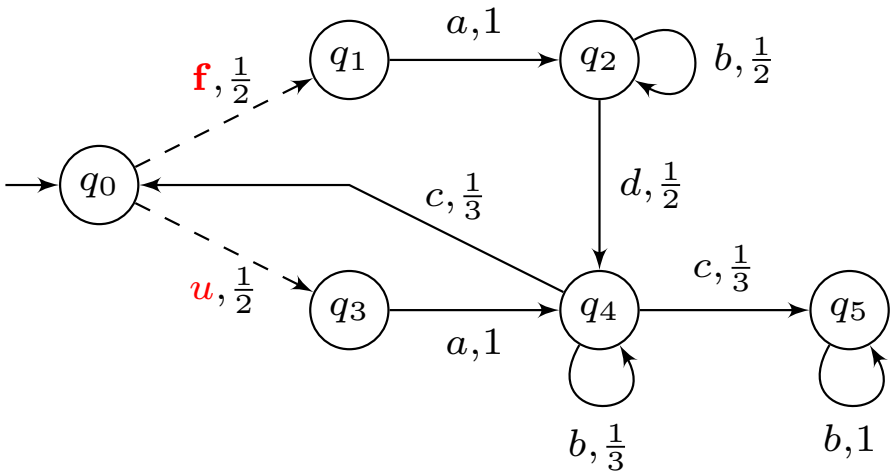

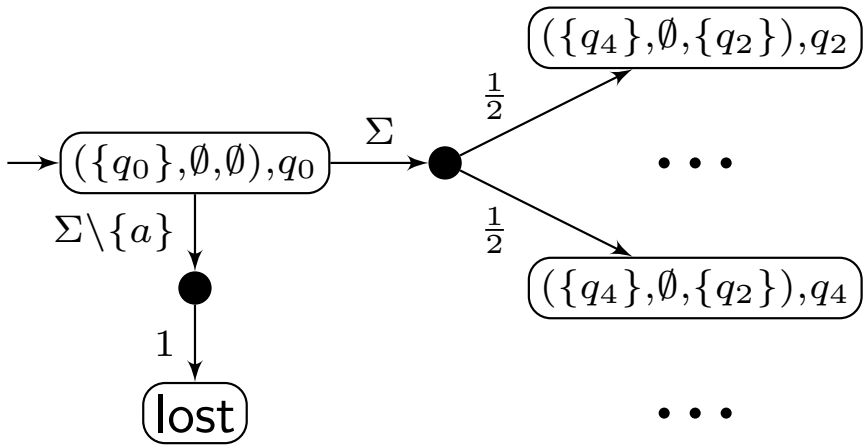

Decidability of the active diagnosis problem

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

$$\text{almost surely } \Box\Diamond(W = \emptyset \vee U = \emptyset)$$

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.

Analyzing the reduction to the POMDP problem, we get that the active diagnosis problem is EXPTIME-complete.

- $\mathcal{C}$ is safely actively diagnosable iff there exists a strategy in $\mathsf{M}_{\mathcal{C}}$ such that:
  - almost surely $\Box\Diamond(W = \emptyset \vee U = \emptyset)$;
  - with positive probability $\Box U \neq \emptyset$.

Belief-based diagnosers are not enough

In our context, the *belief* is the current state of the Büchi automaton.

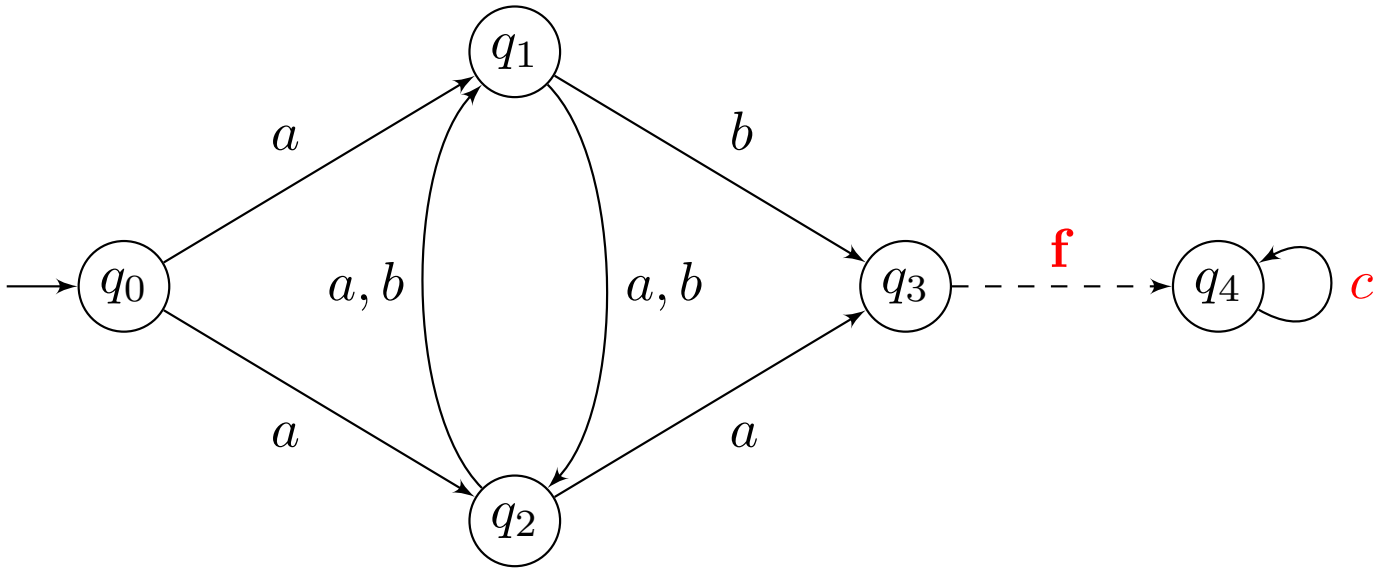

The cLTS is straightforwardly diagnosable but it is not safe.

A safe active diagnoser must perform a guess and keep in memory one bit:

- forbidding a after an odd number of observations;
- and forbidding b after an even number of observations.

Finite-memory diagnosers are not enough

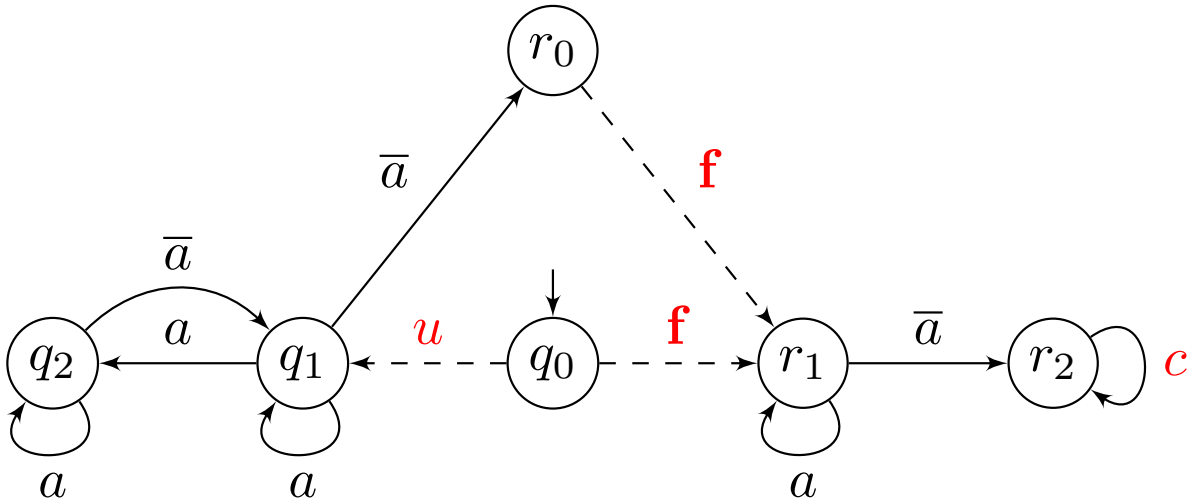

An observed sequence σ is surely faulty iff $\sigma \in \Sigma^* c^\omega$.

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

Finite-memory diagnosers are not enough

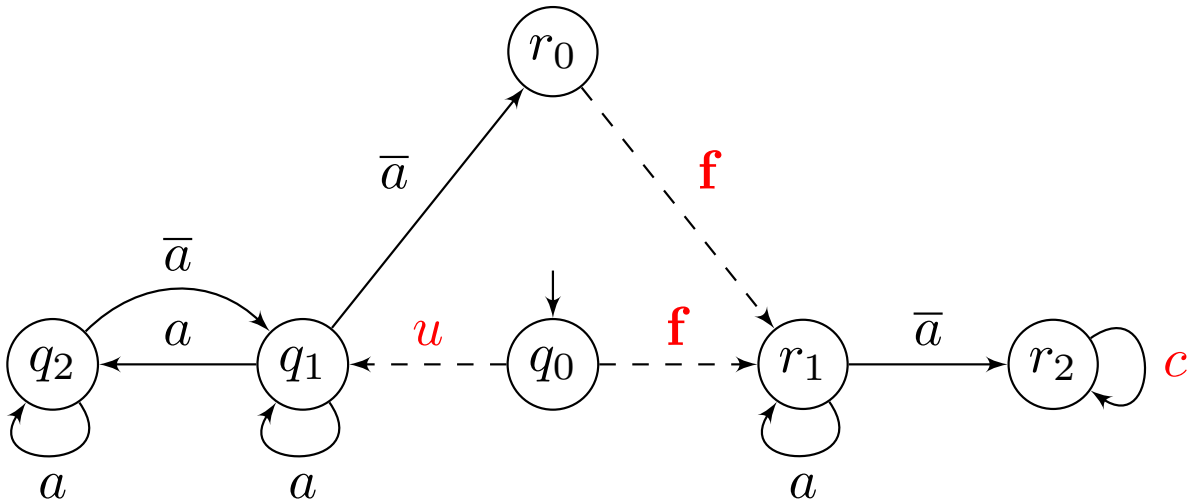

A safe active diagnoser

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

Let $A = \{a, u, \mathbf{f}, c\}$ and $\overline{A} = \{\overline{a}, u, \mathbf{f}, c\}$.

Let π be the controller 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 a safe active diagnoser:

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

There is no finite-memory safe active diagnoser.

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

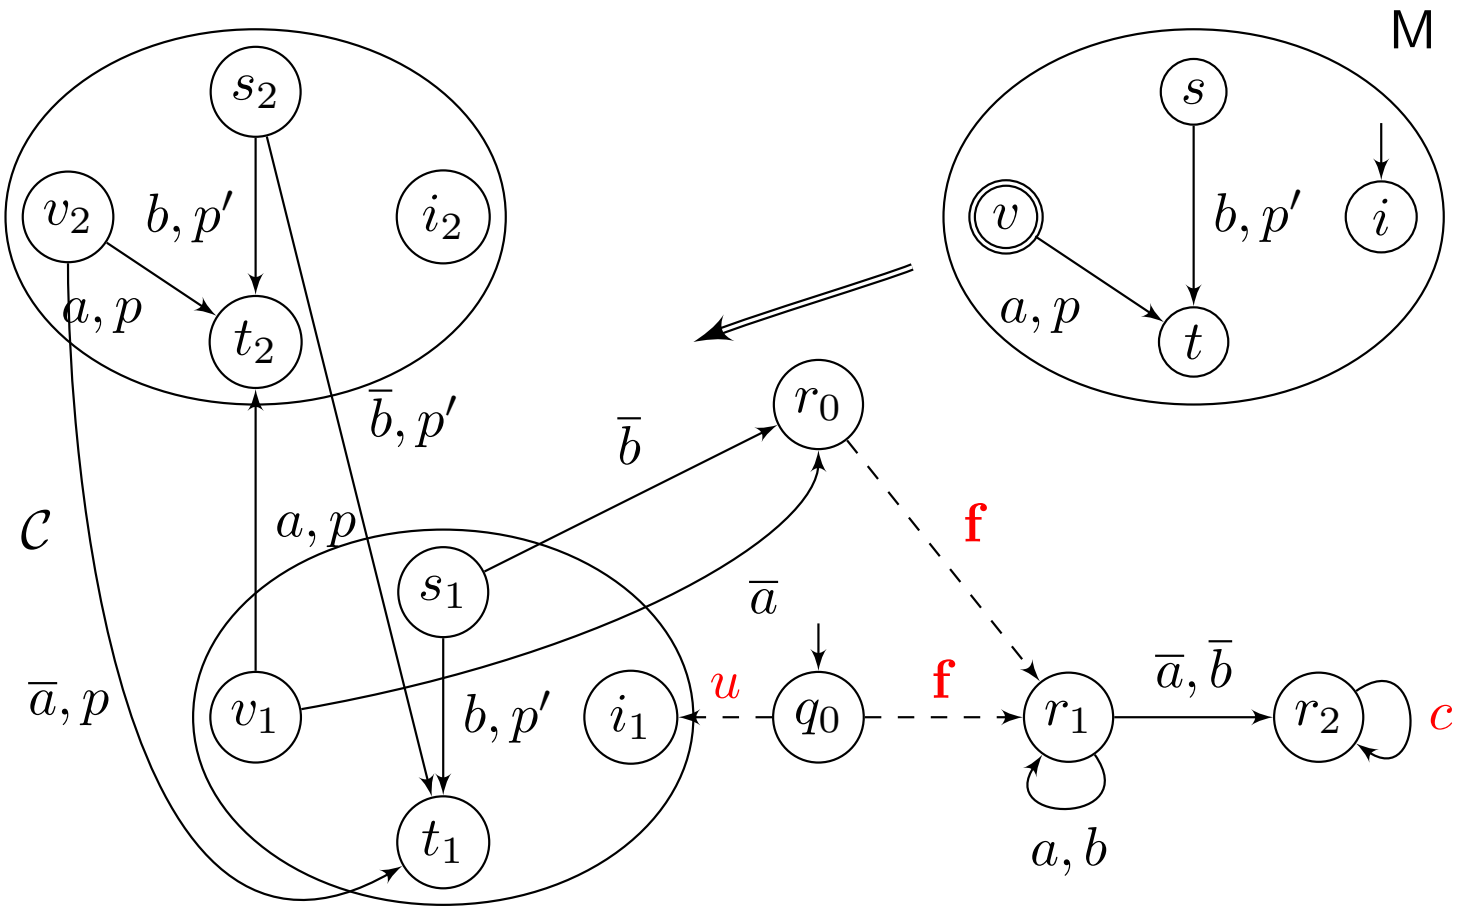

An observed sequence σ is surely faulty iff $\sigma \in \Sigma^* c^\omega$.

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

Restriction to finite-memory diagnosers

Observation

A priori the finite-memory requirement does not ensure decidability.

A decision procedure in EXPTIME:

- Computing the *safe beliefs* that ensure the existence of an active diagnoser surely yielding correct sequences.
- Checking the existence of a diagnoser that ensure active diagnosability almost surely and reaching a belief including a safe belief with positive probability.

The active diagnoser only requires an additional boolean (for switching its mode).

The problem is EXPTIME-hard (using the same reduction as before).

Conclusion and perspectives

Contributions

- Strong improvement of the active diagnosis procedures for transition systems.
- Almost matching lower bounds of the active diagnosis problems for transition systems.
- Introduction of (safe) active diagnosis problems for probabilistic systems.
- Analysis of the problems for probabilistic systems using a POMDP framework.

Perspectives

- Closing the gap between lower and upper bounds related to the minimal delay synthesis problem.
- Introducing the active predictability problem (and other related issues).
- Investigating further POMDP problems with multiple objectives.
- Modelling and analyzing diagnosis with stochastic games.