Semantical and Algorithmic Features of Diagnosis and Prediction

Serge Haddad

Université Paris-Saclay, LMF, ENS Paris-Saclay, CNRS

Séminaire Mefosyloma

March the 13th 2026

Outline

Formalisms

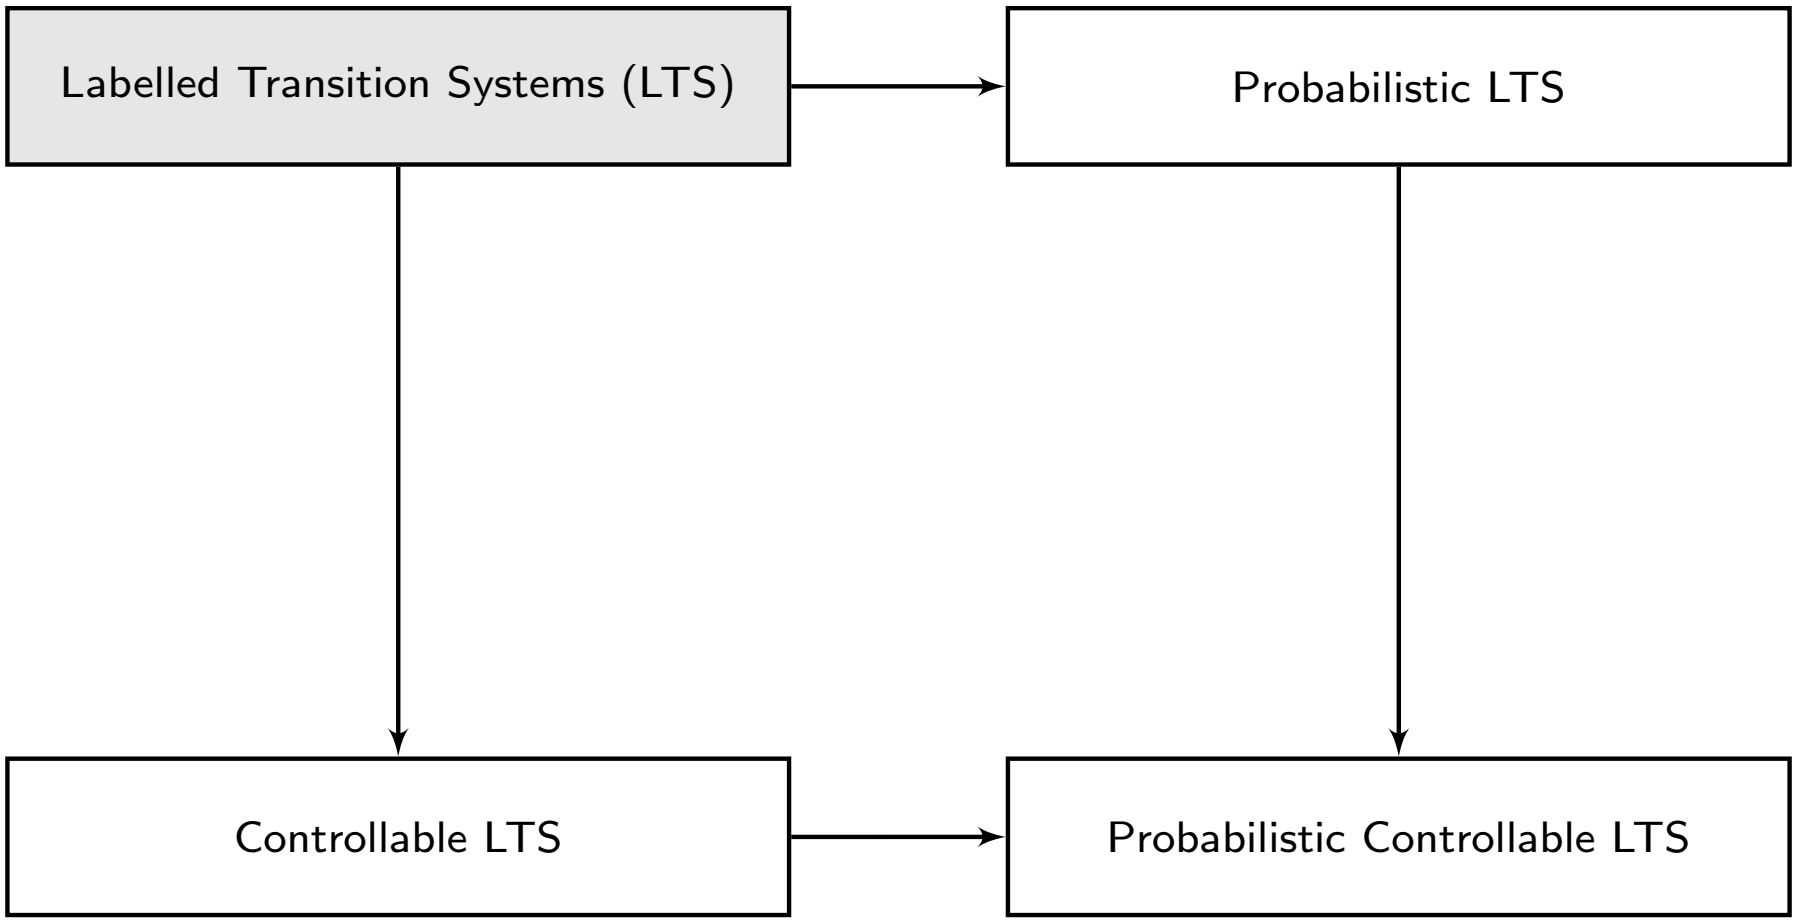
Labelled Transition Systems (LTS)
Probabilistic LTS
Controllable LTS
Probabilistic Controllable LTS

Observing a labelled transition system

Formalization.

- In a labelled transition system (LTS), states are either *correct* or *faulty*;
- There is at least one outgoing transition from any state *(liveness)*;
- There is no transition from a faulty state to a correct one;
- *Events* label transitions and they are *observable*;
- States are *unobservable*.

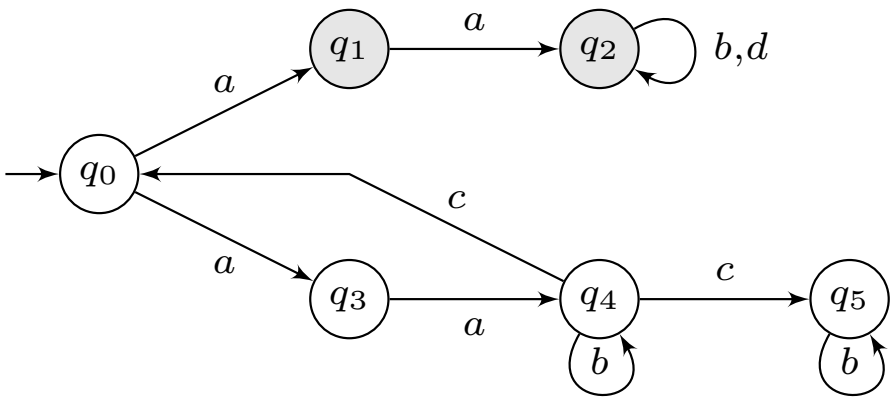

A run yields an (observed) sequence.

- Let $\rho = q_0 a q_1 a q_2 d$. Then $\mathcal{P}(\rho) = aad$;
- Let $\rho = q_0 a q_3 a q_4 c q_0 a q_1 a (q_2 b)^\omega$. Then $\mathcal{P}(\rho) = aacaab^\omega$.

Classification of runs and sequences

Runs.

A run is *faulty* if it reaches a faulty state otherwise it is *correct*.

The set of faulty runs is denoted F and the set of correct runs is denoted C.

A faulty run $\rho = q_0 a_1 \dots q_{n-1} a_n q_n$ is *minimal* if q_n is faulty and q_{n-1} is correct.

Sequences.

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

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

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

Runs w.r.t. their sequences.

A run ρ is *ambiguous* (resp. *surely faulty*, *surely correct*)

if $\mathcal{P}(\rho)$ is ambiguous (resp. surely faulty, surely correct).

Illustration

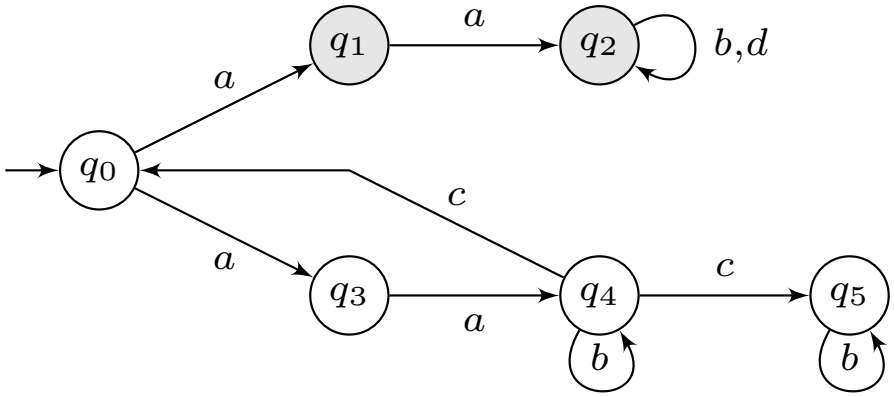

aad is surely faulty: the occurrence of d implies the occurrence of q_2.

$aacb$ is surely correct: $\mathcal{P}^{-1}(aacb) = \{q_0aq_3aq_4cq_5bq_5\}$.

aab^ω is ambiguous: $\mathcal{P}^{-1}(aab^\omega) = \{q_0aq_3a(q_4b)^\omega, q_0aq_1a(q_2b)^\omega\}$.

Infinite runs and diagnosability

Let $\mathcal{A}$ be a LTS.

- $\mathsf{F}^{\infty}_{\mathcal{A}}$ (resp. $\mathsf{SF}^{\infty}_{\mathcal{A}}$) is the set of infinite (resp. surely) faulty runs of $\mathcal{A}$;
- $\mathsf{C}^{\infty}_{\mathcal{A}}$ (resp. $\mathsf{SC}^{\infty}_{\mathcal{A}}$) is the set of infinite (resp. surely) correct runs of $\mathcal{A}$;
- $\mathsf{SF}^{n}_{\mathcal{A}}$ is the set of infinite faulty runs of $\mathcal{A}$ whose subrun of length n is surely faulty;
- $\mathsf{SC}^{n}_{\mathcal{A}}$ is the set of infinite correct runs of $\mathcal{A}$ whose subrun of length n is surely correct.

Several definitions of diagnosability.

- $\mathcal{A}$ is IF-diagnosable if $\mathsf{F}^{\infty}_{\mathcal{A}} = \mathsf{SF}^{\infty}_{\mathcal{A}}$;
- $\mathcal{A}$ is IA-diagnosable if $\mathsf{F}^{\infty}_{\mathcal{A}} = \mathsf{SF}^{\infty}_{\mathcal{A}} \wedge \mathsf{C}^{\infty}_{\mathcal{A}} = \mathsf{SC}^{\infty}_{\mathcal{A}}$;
- $\mathcal{A}$ is FF-diagnosable if $\mathsf{F}^{\infty}_{\mathcal{A}} = \bigcup_{n\in\mathbb{N}} \mathsf{SF}^{n}_{\mathcal{A}}$;
- $\mathcal{A}$ is FA-diagnosable if $\mathsf{F}^{\infty}_{\mathcal{A}} = \bigcup_{n\in\mathbb{N}} \mathsf{SF}^{n}_{\mathcal{A}} \wedge \mathsf{C}^{\infty}_{\mathcal{A}} = \bigcap_{n\in\mathbb{N}} \bigcup_{m\geq n} \mathsf{SC}^{m}_{\mathcal{A}}$.

Relations between definitions

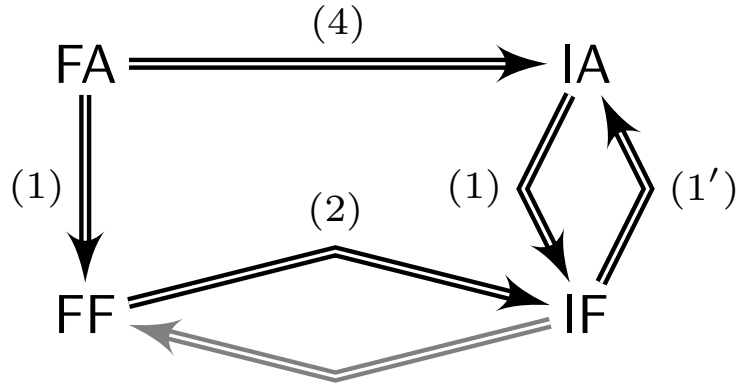

(3) for finitely branching LTS

Sketch of proof.

(1) By definition, and (1') since $\mathsf{F}^{\infty}_{\mathcal{A}} \neq \mathsf{SF}^{\infty}_{\mathcal{A}}$ implies $\mathsf{C}^{\infty}_{\mathcal{A}} \neq \mathsf{SC}^{\infty}_{\mathcal{A}}$.

(2) Since for all n, $\mathsf{SF}^{n}_{\mathcal{A}} \subseteq \mathsf{SF}^{\infty}_{\mathcal{A}}$, $\bigcup_{n\in\mathbb{N}} \mathsf{SF}^{n}_{\mathcal{A}} \subseteq \mathsf{SF}^{\infty}_{\mathcal{A}}$.

(3) Let $\rho \in \mathsf{SF}^{\infty}_{\mathcal{A}} \setminus \bigcup_{n\in\mathbb{N}} \mathsf{SF}^{n}_{\mathcal{A}}$. Define a tree using ρ and $\{\rho_n\}_{n\in\mathbb{N}}$ where ρ_n of length n is correct and such that $\mathcal{P}(\rho_n)$ is a prefix of $\mathcal{P}(\rho)$. Using Koenig's lemma deduce ρ' a correct run with $\mathcal{P}(\rho') = \mathcal{P}(\rho)$.

(4) Let $\rho \in \mathsf{C}^{\infty}_{\mathcal{A}} \setminus \mathsf{SC}^{\infty}_{\mathcal{A}}$. There exists a faulty run ρ' with $\mathcal{P}(\rho') = \mathcal{P}(\rho)$. Let ρ'_n some faulty prefix of ρ' with length n. Then $\rho \notin \bigcup_{m\geq n} \mathsf{SC}^{m}_{\mathcal{A}}$.

FF $\not\Rightarrow$ FA **and** IA $\not\Rightarrow$ FA

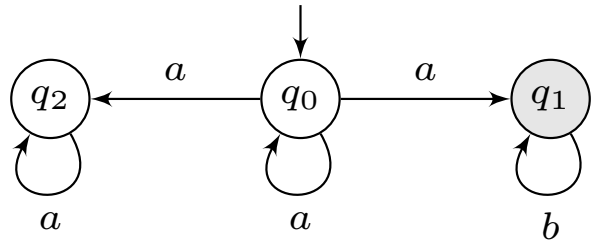

Sketch of proof.

- $\mathsf{F}^{\infty}_{\mathcal{A}} = (q_0 a)^* q_0 a (q_1 b)^{\omega}$.

  For all $n \geq 2$, $\mathsf{SF}^{n}_{\mathcal{A}} = \left\{(q_0 a)^{m-2} q_0 a (q_1 b)^{\omega}\right\}_{2 \leq m \leq n}$.

  Whence $\mathsf{F}^{\infty}_{\mathcal{A}} = \bigcup_{n \in \mathbb{N}} \mathsf{SF}^{n}_{\mathcal{A}}$.

- $\mathsf{C}^{\infty}_{\mathcal{A}} = (q_0 a)^{\omega} + (q_0 a)^* q_0 a (q_2 a)^{\omega}$.

  For all $n \geq 1$, $\mathsf{SC}^{n}_{\mathcal{A}} = \emptyset$.

  Whence $\mathsf{C}^{\infty}_{\mathcal{A}} \neq \bigcap_{n \in \mathbb{N}} \bigcup_{m \geq n} \mathsf{SC}^{m}_{\mathcal{A}}$

- Since FF implies IA, IA does not imply FA.

IF $\not\Rightarrow$ FF

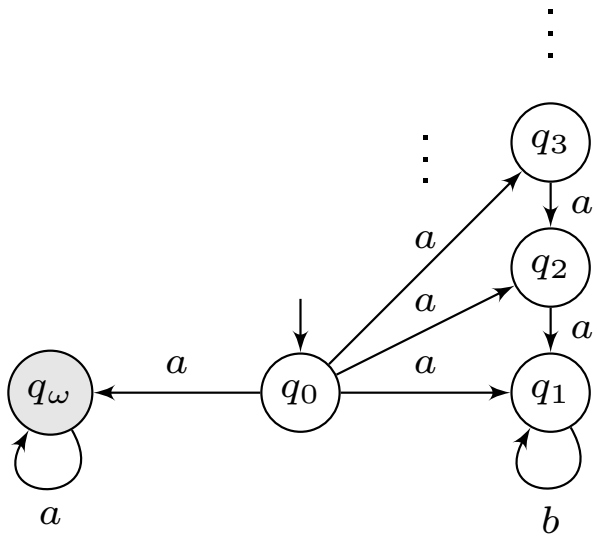

Sketch of proof.

$\mathsf{F}^{\infty}_{\mathcal{A}} = \mathsf{SF}^{\infty}_{\mathcal{A}} = q_0 a (q_\omega a)^\omega$.

For all n, $\mathsf{SF}^{n}_{\mathcal{A}} = \emptyset$.

Thus $\mathsf{F}^{\infty}_{\mathcal{A}} \neq \bigcup_{n \in \mathbb{N}} \mathsf{SF}^{n}_{\mathcal{A}}$.

Formalisms

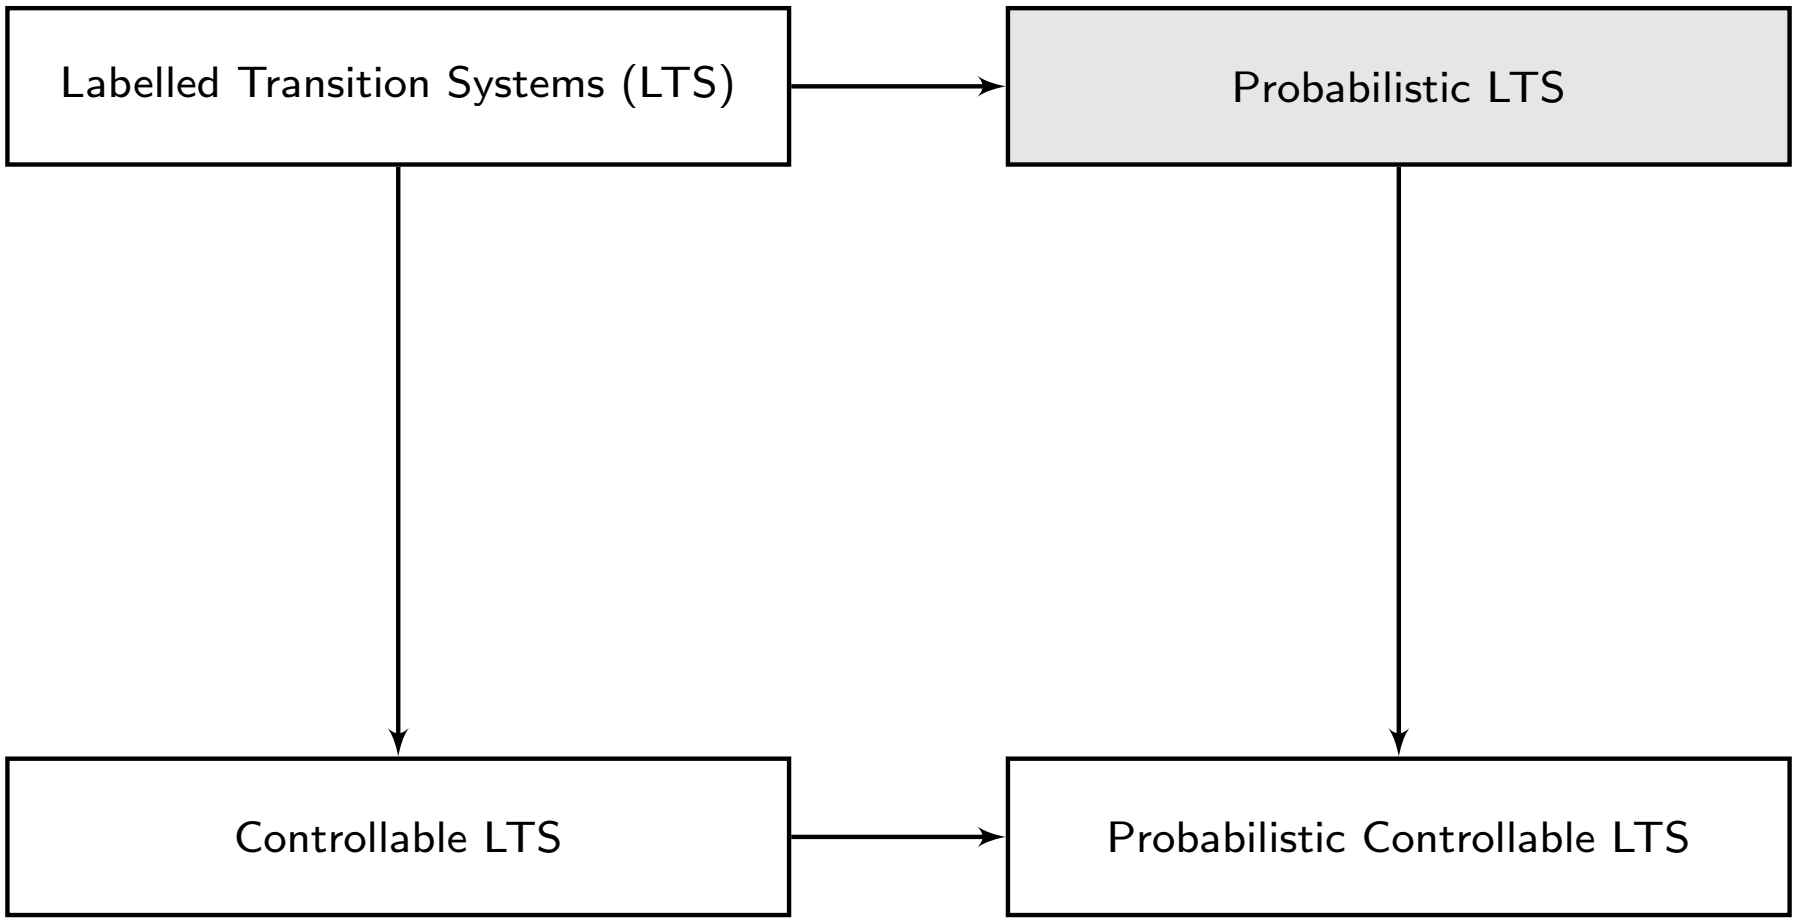
Labelled Transition Systems (LTS)
Probabilistic LTS
Controllable LTS
Probabilistic Controllable LTS

Randomizing a LTS

Formalization.

A randomized LTS is a labelled Markov chain, denoted pLTS.

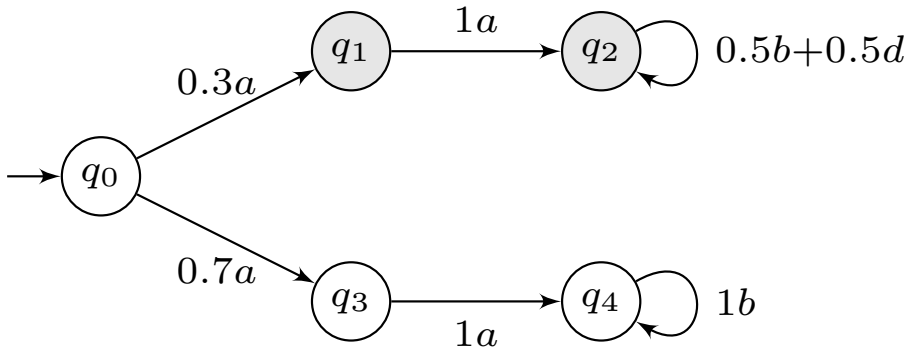

Observations.

- As a LTS, $\mathcal{A}$ is not IF-diagnosable due to the sequence a^2b^ω.
- As a PLTS, it seems to become IF-diagnosable but not IA-diagnosable.

How to formalize diagnosability?

Diagnosability for pLTS

Considering null subsets of runs as irrelevant.

- $\mathcal{A}$ is IF-diagnosable if $\mathbf{Pr}(\mathsf{F}^{\infty}_{\mathcal{A}} \setminus \mathsf{SF}^{\infty}_{\mathcal{A}}) = 0$;
- $\mathcal{A}$ is IA-diagnosable if $\mathbf{Pr}(\mathsf{F}^{\infty}_{\mathcal{A}} \setminus \mathsf{SF}^{\infty}_{\mathcal{A}}) + \mathbf{Pr}(\mathsf{C}^{\infty}_{\mathcal{A}} \setminus \mathsf{SC}^{\infty}_{\mathcal{A}}) = 0$;
- $\mathcal{A}$ is FF-diagnosable if $\mathbf{Pr}(\mathsf{F}^{\infty}_{\mathcal{A}} \setminus \bigcup_{n\in\mathbb{N}} \mathsf{SF}^{n}_{\mathcal{A}}) = 0$
  equivalent to $\lim_{n\to\infty} \mathbf{Pr}(\mathsf{SF}^{n}_{\mathcal{A}}) = \mathbf{Pr}(\mathsf{F}^{\infty}_{\mathcal{A}})$;
- $\mathcal{A}$ is FA-diagnosable if $\mathbf{Pr}(\mathsf{F}^{\infty}_{\mathcal{A}} \setminus \bigcup_{n\in\mathbb{N}} \mathsf{SF}^{n}_{\mathcal{A}}) + \mathbf{Pr}(\mathsf{C}^{\infty}_{\mathcal{A}} \setminus \bigcap_{n\in\mathbb{N}} \bigcup_{m\geq n} \mathsf{SC}^{m}_{\mathcal{A}}) = 0$
  equivalent to $\lim_{n\to\infty} \mathbf{Pr}(\mathsf{SF}^{n}_{\mathcal{A}}) = \mathbf{Pr}(\mathsf{F}^{\infty}_{\mathcal{A}}) \wedge \lim_{n\to\infty} \mathbf{Pr}(\bigcup_{m\geq n} \mathsf{SC}^{m}_{\mathcal{A}}) = \mathbf{Pr}(\mathsf{C}^{\infty}_{\mathcal{A}})$.

New relations between the definitions.

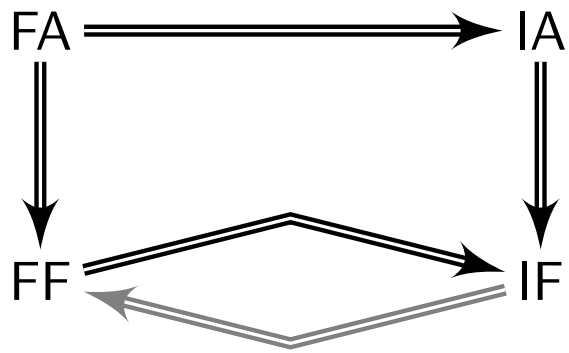

for finitely branching LTS

FF $\not\Rightarrow$ IA

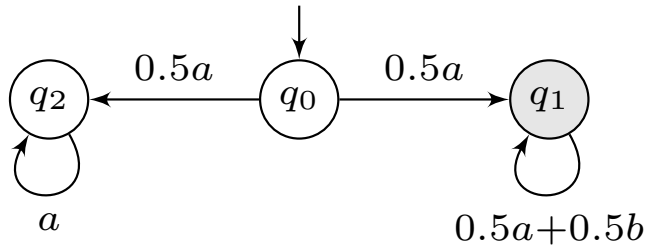

Sketch of proof.

$\mathsf{F}^{\infty}_{\mathcal{A}} = q_0 a (q_1 \Sigma)^{\omega}$.

For all $n \geq 2$, $\mathsf{SF}^{n}_{\mathcal{A}} = \left\{ q_0 a (q_1 \Sigma)^{m-2} q_1 b (q_1 \Sigma)^{\omega} \right\}_{2 \leq m \leq n}$.

Thus $\mathsf{F}^{\infty}_{\mathcal{A}} \setminus \bigcup_{n \in \mathbb{N}} \mathsf{SF}^{n}_{\mathcal{A}} = \{ q_0 a (q_1 a)^{\omega} \}$ with $\mathbf{Pr}(\{ q_0 a (q_1 a)^{\omega} \}) = 0$.

$\mathsf{C}^{\infty}_{\mathcal{A}} = q_0 a (q_2 a)^{\omega}$ with $\mathbf{Pr}(\{ q_0 a (q_2 a)^{\omega} \}) = \frac{1}{2}$ and $\mathsf{SC}^{\infty}_{\mathcal{A}} = \emptyset$.

Consequence: FF $\not\Rightarrow$ FA and IF $\not\Rightarrow$ IA.

Still relaxing the requirements

Observation. Using pLTS instead of LTS relaxes the diagnosability requirement.

But there are other kinds of relaxation in the randomized framework.

Let ρ be a finite run and $Cyl(\rho)$ the set of infinite runs with ρ as prefix.
By definition $\mathbf{Pr}(\rho) = \mathbf{Pr}(Cyl(\rho))$.

Let σ be a finite sequence. Then the *correctness proportion* of σ is defined by:

$$\mathsf{CorP}(\sigma) = \frac{\mathbf{Pr}(\mathcal{P}^{-1}(\sigma) \cap \mathsf{C})}{\mathbf{Pr}(\mathcal{P}^{-1}(\sigma))}$$

$\mathsf{CorP}(\sigma)$ is the probability of an error when observing σ, one declares a fault.

Other definitions of diagnosability.

Let $\varepsilon \geq 0$. $\mathcal{A}$ is εFF-*diagnosable* if for all minimal faulty run ρ and $\alpha > 0$, there exists $n_{\rho,\alpha}$ such that for all $n \geq n_{\rho,\alpha}$:

$$\mathbf{Pr}(\{\rho' \in Cyl(\rho) \mid |\rho'| = n + |\rho| \wedge \mathsf{CorP}(\mathcal{P}(\rho')) > \varepsilon\}) \leq \alpha \mathbf{Pr}(\rho)$$

$\mathcal{A}$ is AFF-*diagnosable* if for all $\varepsilon > 0$, $\mathcal{A}$ is εFF-diagnosable.

Relations between definitions

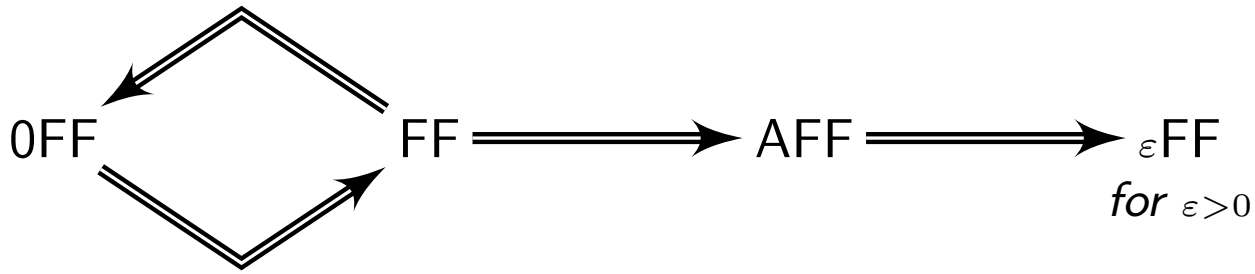

- A non implication: AFF $\not\Rightarrow$ FF.

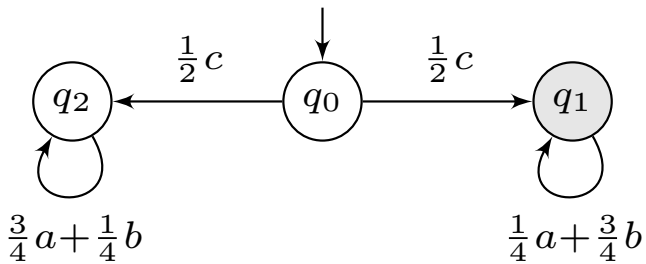

Sketch of proof

- $\mathsf{F}^{\infty}_{\mathcal{A}} = q_0 c (q_1 \Sigma)^{\omega}$ with $\mathbf{Pr}(\mathsf{F}^{\infty}_{\mathcal{A}}) = \frac{1}{2}$. $\mathsf{SF}^{\infty}_{\mathcal{A}} = \emptyset$.

  Thus $\mathcal{A}$ is not IF-diagnosable and so not FF-diagnosable.

- There is a single minimal faulty run $\rho = q_0 c q_1$ with $\mathbf{Pr}(\rho) = \frac{1}{2}$.

  Let $\sigma = c(a+b)^*$ with $na = |\sigma|_a$ and $nb = |\sigma|_b$.

  Then $\mathsf{CorP}(\sigma) = \frac{1}{1+3^{nb-na}}$.

  Let $\varepsilon > 0$. Select some n_ε such that $\frac{1}{1+3^{n_\varepsilon}} \leq \varepsilon$.

  Let $\alpha > 0$.

  By the law of large numbers there exists n_α such that for all $n \geq n_\alpha$,

$$\mathbf{Pr}(\rho' \in Cyl(\rho) \,||\rho'| = n \wedge |\rho'|_b \geq 2|\rho'|_a) \geq \frac{1}{2}(1-\alpha)$$

  Fix $n_{\varepsilon,\alpha} = \max(3n_\varepsilon, n_\alpha)$. Then for all $n \geq n_{\varepsilon,\alpha}$,

$$\mathbf{Pr}(\{\rho' \in Cyl(\rho) \mid |\rho'| = n \wedge \mathsf{CorP}(\mathcal{P}(\rho')) > \varepsilon\}) \leq \alpha \mathbf{Pr}(\rho)$$

Formalisms

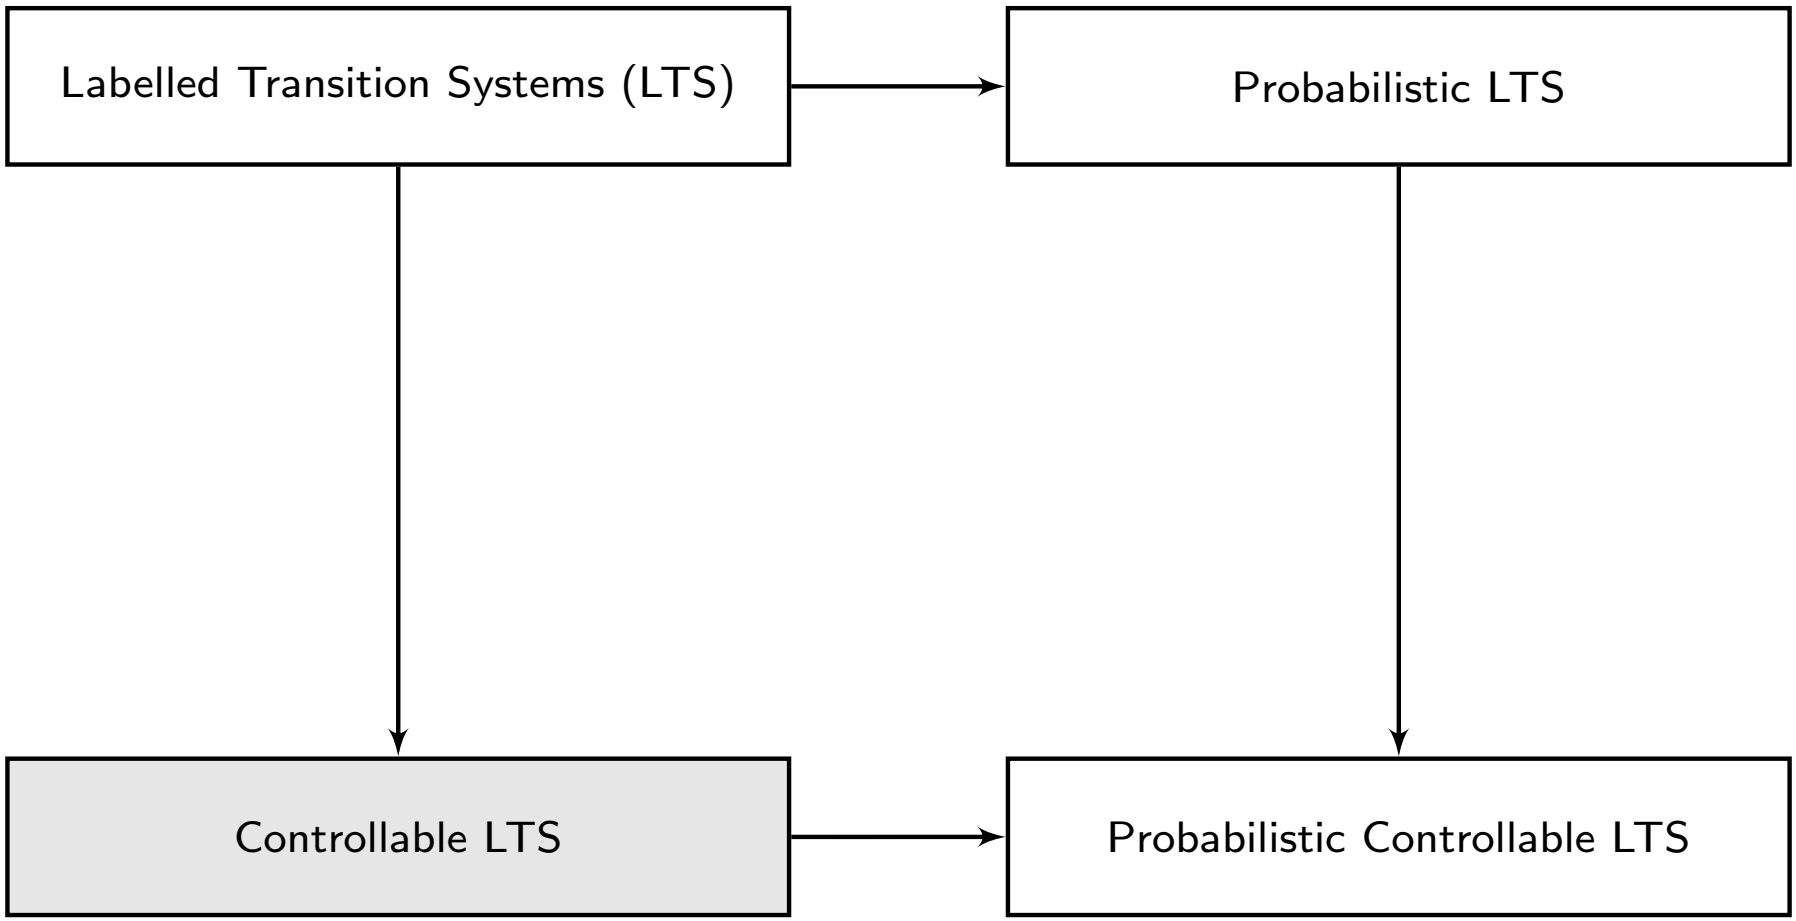
Labelled Transition Systems (LTS)
Probabilistic LTS
Controllable LTS
Probabilistic Controllable LTS

Controllable LTS and active diagnoser

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

A *controller* forbids some controllable events

depending on the current observed sequence.

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

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

The *delay* of an active diagnoser is the maximal number of event occurrences between a minimal faulty run

and the possible extending minimal surely faulty sequences.

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).

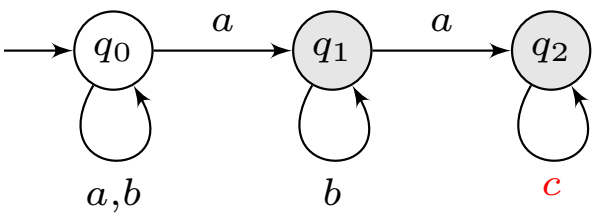

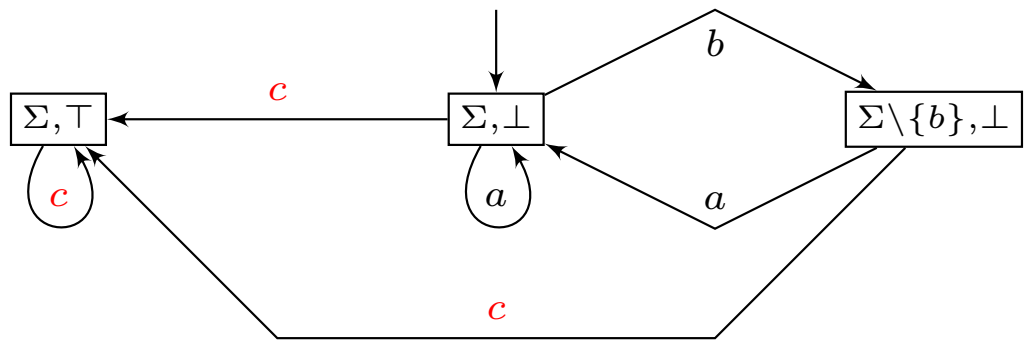

Active diagnosis problems

- The *active diagnosis decision problem*, i.e. decide
  whether a cLTS is actively (FF- or FA- or IA-) 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.

Revisiting liveness in cLTS (1)

Let $\mathbf{Idle} \subseteq Q$ be a subset of states (possibly faulty or correct)

where the system may stay indefinitely without executing a transition.

An active diagnoser is allowed to block all the (controllable) events

outgoing from an idle state.

The special event δ reports to the active diagnoser

that no transition has been triggered.

Illustration.

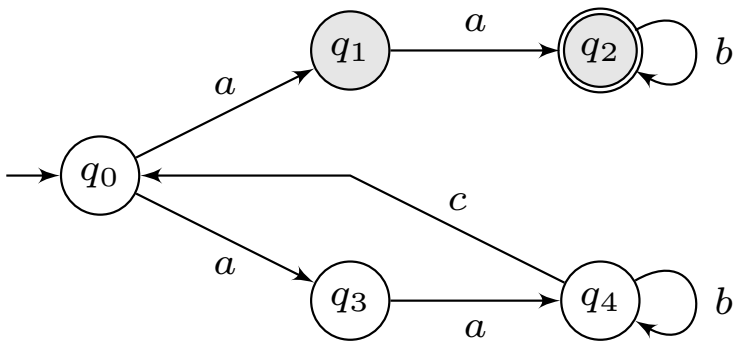

If $q_2 \in \mathbf{Idle}$, the active diagnoser blocks b after two a's

and either observes c (triggered by q_4) or δ (triggered by q_2).

Revisiting liveness in cLTS (2)

There are two kinds of transitions: *eager* or *lazy* ones.

All transitions outgoing from an idle state are lazy.

In a state,

- if there is at least one outgoing eager transition labelled by an allowed event then the system must trigger some transition (eager or lazy);
- otherwise the system may not execute a transition reported by δ;
- however if the state is not idle, it must eventually trigger some transition.

Illustration.

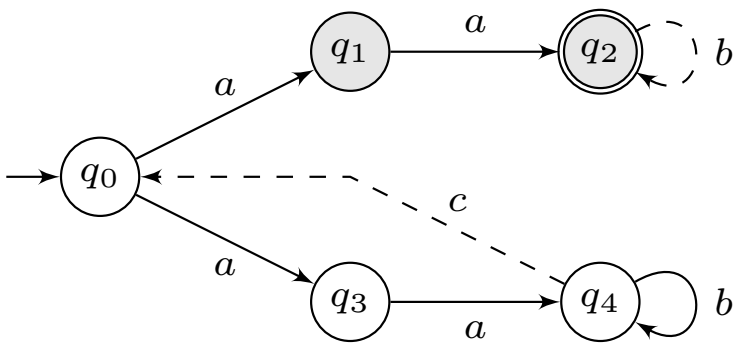

Consider a controller that blocks b after two a's.

Then the controlled system is IA-diagnosable but not FA-diagnosable.

Formalisms

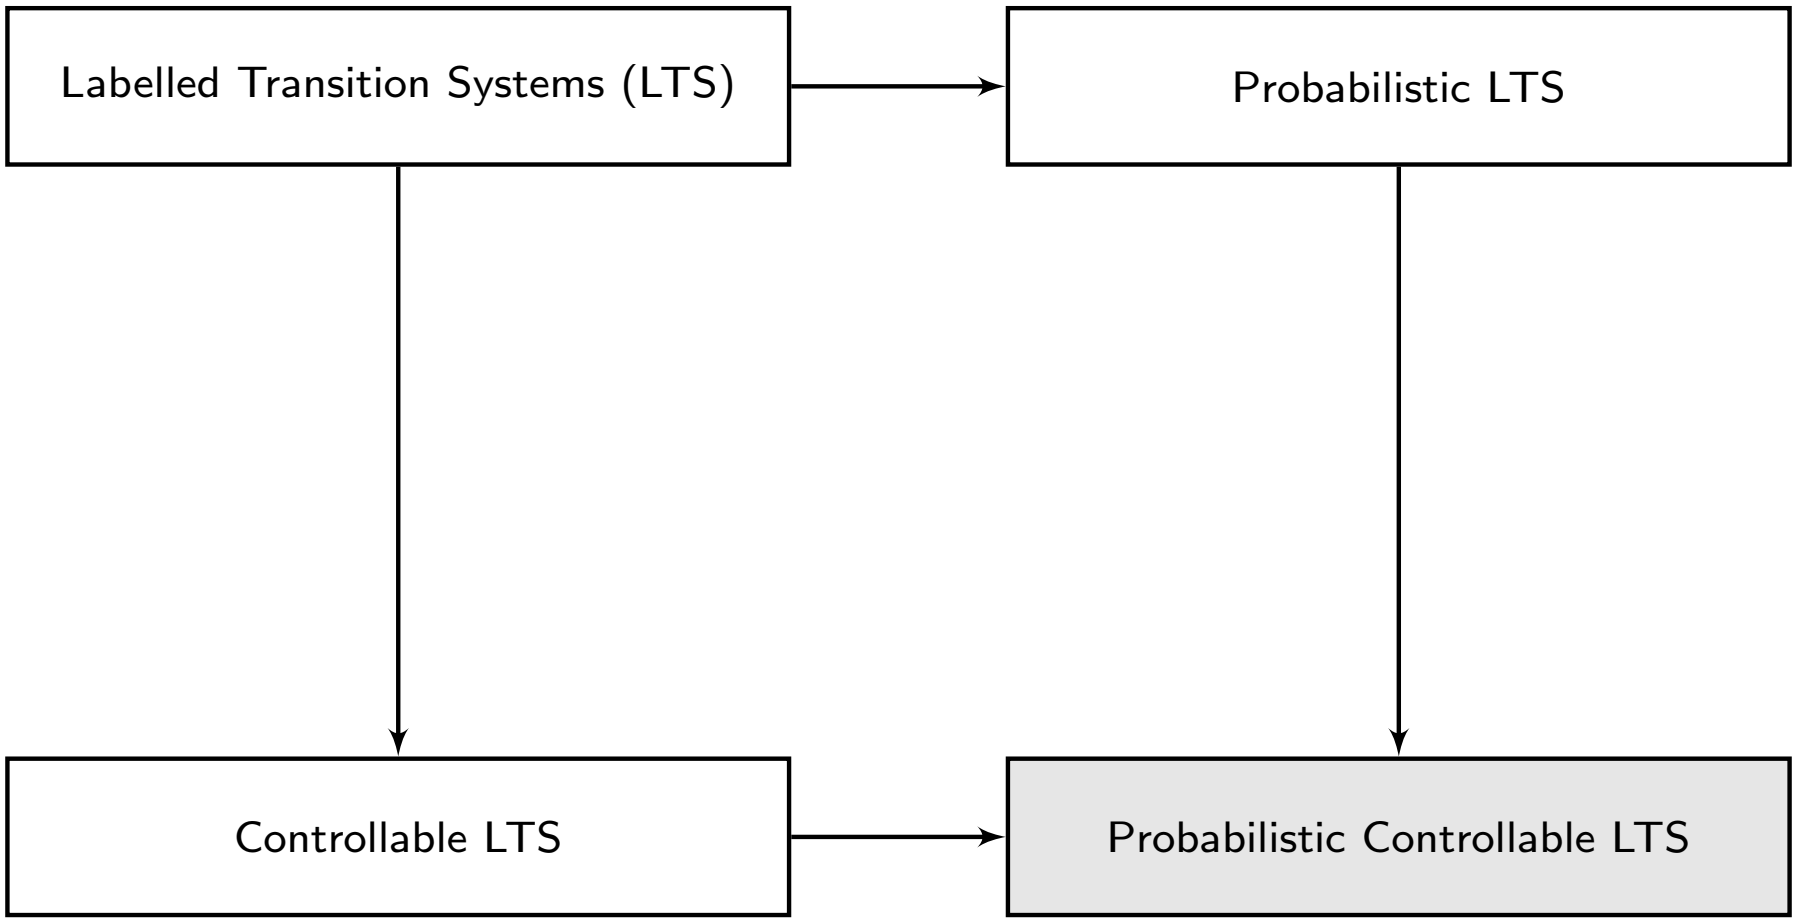
Labelled Transition Systems (LTS)
Probabilistic LTS
Controllable LTS
Probabilistic Controllable LTS

pcLTS

A *probabilistic controllable labelled transition system* (pcLTS) is a cLTS with positive integer weights on transitions.

A controller forbids controllable events depending on the observed sequence.
It can *randomly* select the forbidden events.
As before, a controller must not introduce deadlocks.

Let $\mathcal{A}$ be a pcLTS and τ be a controller.
Then $\mathcal{A}_\tau$ is a pLTS where the probability are obtained by normalization among the allowed events.

Controller τ is an active diagnoser if $\mathcal{A}_\tau$ is diagnosable.

Illustration

A *deterministic* active diagnoser τ:

Forbid two consecutive b after two a.

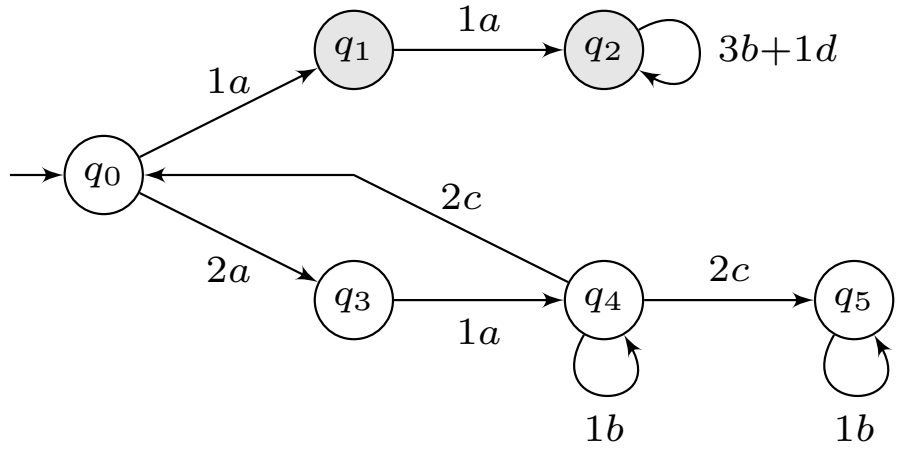

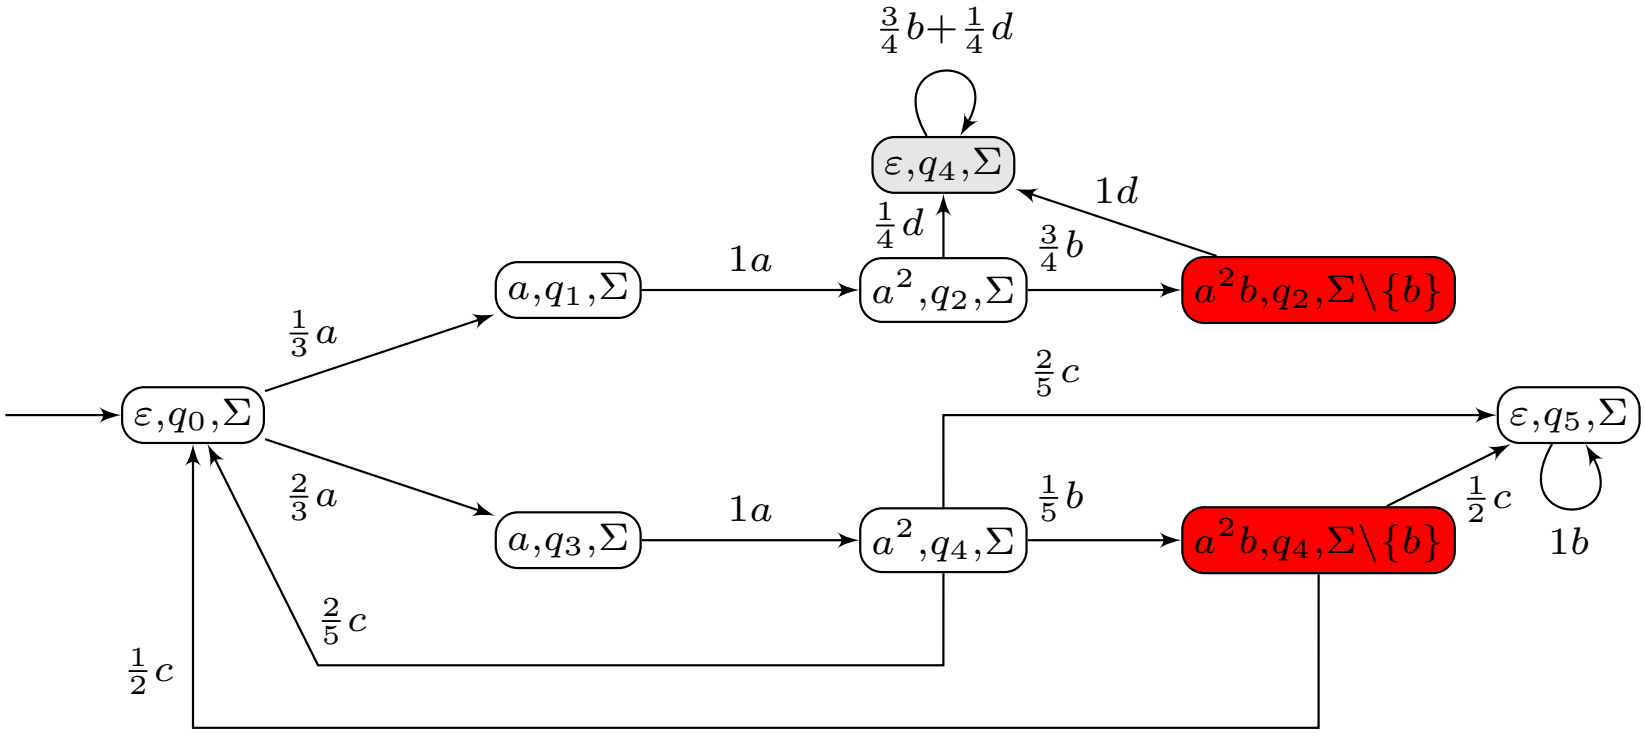

Revisiting correctness in pLTS

One wants to study the evolution of correct runs along the execution.

Let us propose several formalizations for a pLTS $\mathcal{A}$.

$\mathcal{A}$ is *safe* if $\mathbf{Pr}(\mathsf{C}^{\infty}_{\mathcal{A}}) > 0$.

Let $\mathsf{C}^{n}_{\mathcal{A}}$ the set of correct runs of length n.

Let $0 < \alpha < 1$. $\mathcal{A}$ is α-*resilient* if $\limsup_{n\to\infty} \frac{\alpha^n}{\mathbf{Pr}(\mathsf{C}^{n}_{\mathcal{A}})} = 0$.

$\mathcal{A}$ is *strongly resilient* if for all $0 < \alpha < 1$, $\mathcal{A}$ is α-*resilient*.

$\mathcal{A}$ is *weakly resilient* if there exists $0 < \alpha < 1$ such that $\mathcal{A}$ is α-*resilient*.

Active probabilistic diagnosis problems

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

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

The *resilient active probabilistic diagnosis problems* ask whether
there exists a (α-, strongly, weakly) resilient active diagnoser τ for $\mathcal{A}$.

The *synthesis problems* consist in building
a (safe, α-resilient, strongly resilient, weakly resilient) active diagnoser τ for $\mathcal{A}$
in the positive case.

Outline

Diagnosability in finite LTS

$\mathcal{A}$ is not IA-diagnosable if $\exists \rho = (q_n a_n)_{n\in\mathbb{N}}, \rho' = (q'_n a_n)_{n\in\mathbb{N}}$ fulfilling:

- for all n, q_n is correct;
- there exists n_0 such that q'_{n_0} is faulty.

In finite LTS this is equivalent to $\exists \rho = (q_i a_i)_{i\leq n}, \rho' = (q'_i a_i)_{i\leq n}$ fulfilling:

- for all $i \leq n$, q_i is correct;
- q'_n is faulty;
- there exists $m < n$ such that $q_m = q_n$ and $q'_m = q'_n$.

The twin plant construction $\widehat{\mathcal{A}}$.

- The states of $\widehat{\mathcal{A}}$ are $\{(q_1, q_2) \in Q^2 \mid q_1 \text{ is correct}\}$;
- $(q_1, q_2) \xrightarrow{a} (q'_1, q'_2)$ if for all i, $q_i \xrightarrow{a} q'_i$.

$\mathcal{A}$ is not IA-diagnosable
if there exists (q_1, q_2) belonging to a non trivial SCC of $\widehat{\mathcal{A}}$ with q_2 faulty.

The IA-diagnosability problem for LTS belongs to NLOGSPACE.

Illustration

$\mathcal{A}$

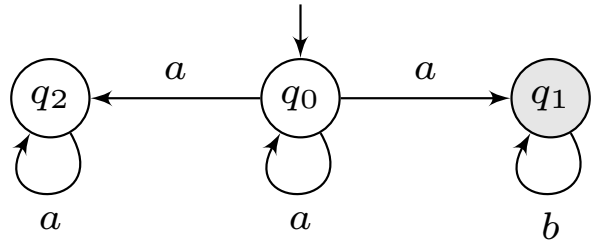

$\widehat{\mathcal{A}}$

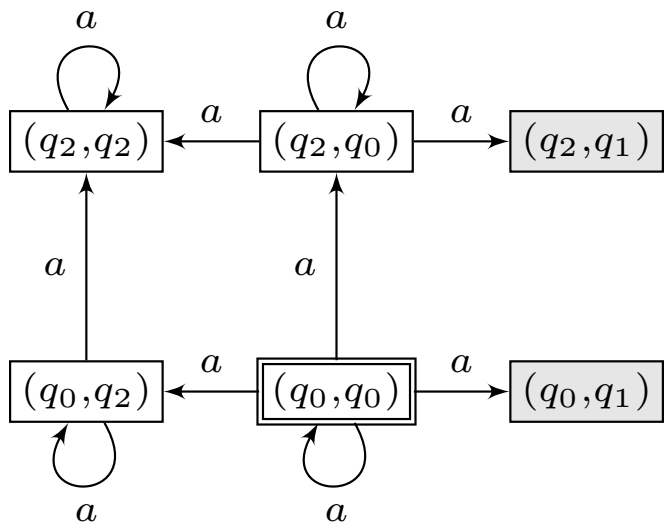

Diagnosability in finite pLTS

Let $\mathcal{A}$ be a pLTS and Cr be its set of correct states.

Then the pLTS $\mathcal{B}_{\mathcal{A}}$ is defined as follows.

- The set of states of $\mathcal{B}_{\mathcal{A}}$ is $Q \times 2^{Cr}$;
- The initial state is $(q_0, \{q_0\})$;
- $(q, U) \xrightarrow{\mathbf{Pr}(q \xrightarrow{a} q') \ a} (q', \delta(U, a) \cap Cr)$.

Observations.

- $\mathcal{B}_{\mathcal{A}}$ has the same behaviour as $\mathcal{A}$;
- The second component of a state of $\mathcal{B}_{\mathcal{A}}$
  is the set of possible correct states according to the observed sequence.

Since almost surely $\mathcal{B}_{\mathcal{A}}$ reaches a BSCC, $\mathcal{A}$ is not IF-diagnosable
iff there exists a state (q, U) of a BSCC with $U \neq \emptyset$.

By repeated guesses, explorations of $\mathcal{B}_{\mathcal{A}}$ without building it
and applications of Savitch theorem, one gets:

The IF-diagnosability problem for pLTS belongs to PSPACE.

Hardness of diagnosability

Let $\mathcal{A}$ be a non deterministic finite live automaton where all states are final.

Then the eventual universal problem:

"Does there exists $v \in \Sigma^*$ such that $v^{-1}\mathcal{L}(\mathcal{A}) = \Sigma^*$?"

is PSPACE-complete.

Let $\mathcal{A}'$ be the pLTS be defined as follows.

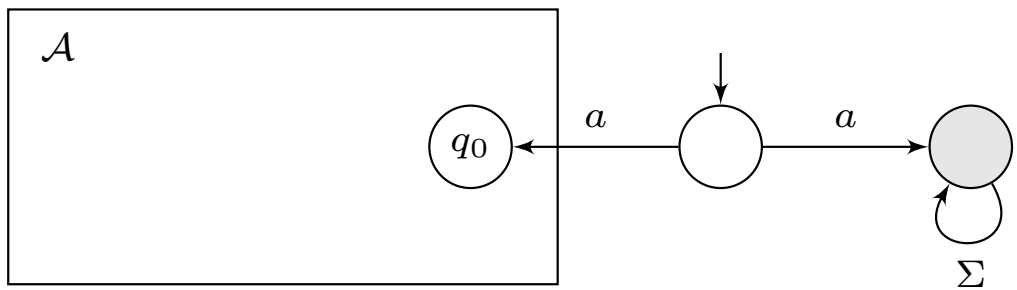

(the exact values of probabilities of $\mathcal{A}'$ are irrelevant.)

Then $\mathcal{A}'$ is not IF-diagnosable iff $\mathcal{A}$ is eventually universal. So:

The IF-diagnosability problem for pLTS is PSPACE-complete.

Turn-based Büchi games

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

- A live graph (V, E) with $v_0 \in V = V_{\mathbf{I}} \uplus V_{\mathbf{II}}$ and $F \subseteq V$;
- In a vertex $v \in V_{\mathbf{P}}$, **P** selects an edge (v, w) and the game goes on with w as current vertex;
- Player **I** wins if 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 of cLTS

Vertices of Player I: $\langle U, V, W \rangle$.

- U is the set of possible correct states;
- V is the set of possible faulty states due to some *recent* fault;
- W is the set of possible faulty states due to some *older* fault;

Player **I** selects an *admissible* control $\Sigma^\bullet$ for $U \cup V \cup W$
leading to vertex $(\langle U, V, W \rangle, \Sigma^\bullet)$ of Player **II**.

Then player **II** selects an action $a \in \Sigma^\bullet$ such that $\delta(U \cup V \cup W, a) \neq \emptyset$
leading to vertex $\langle U', V', W' \rangle$ where:

- $U' = \delta(U, a) \cap Cr$;
- If $W \neq \emptyset$ then $V' = \delta(U \cup V, a) \setminus Cr$ else $V' = \emptyset$;
- If $W \neq \emptyset$ then $W' = \delta(W, a)$ else $W' = \delta(U \cup V, a) \setminus Cr$.

The set of accepting states is $F = \{\langle U, V, \emptyset \rangle\}_{U,V} \cup \{\langle \emptyset, V, W \rangle\}_{V,W}$.

Illustration

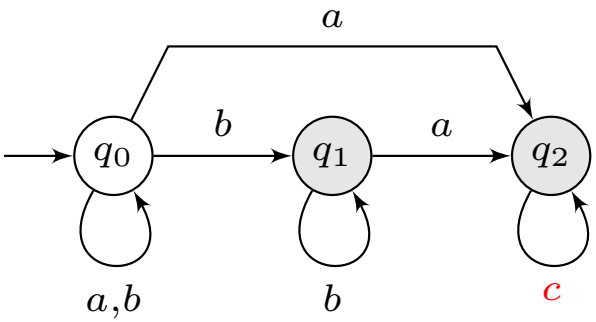

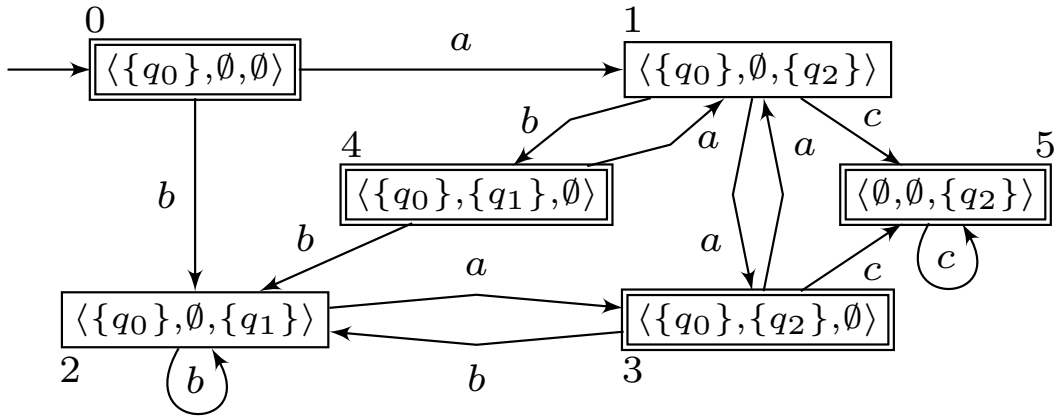

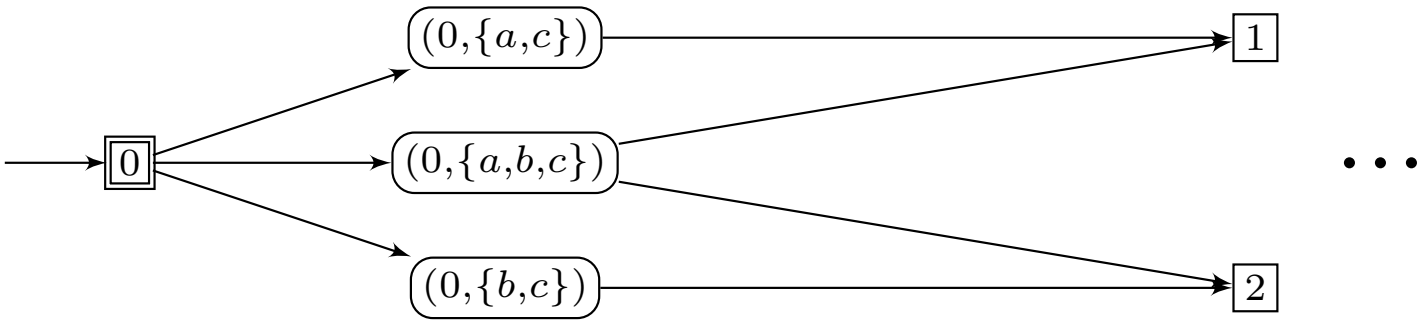

Results of this construction

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 Player **I**.

Consequences

- The decision problem belongs to EXPTIME;
- The synthesis algorithm yields an active diagnoser with $2^{\mathcal{O}(|Q|)}$ states.

Lower bounds.

- The decision problem belongs to EXPTIME-hard
  (*obtained by a reduction from safety games with partial observation*);
- For all $n \in \mathbb{N}$, there is a cLTS with n states
  such that any active diagnoser requires $2^{\Omega(n)}$ states.

Partially observed Markov decision process

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

- Q is a finite set of states with q_0 the initial state;
- Obs : $Q \to \mathcal{O}$ assigns an observation O $\in \mathcal{O}$ to each state;
- Act is a finite set of actions;
- $T : Q \times \text{Act} \to Dist(Q)$ is a partial transition function.

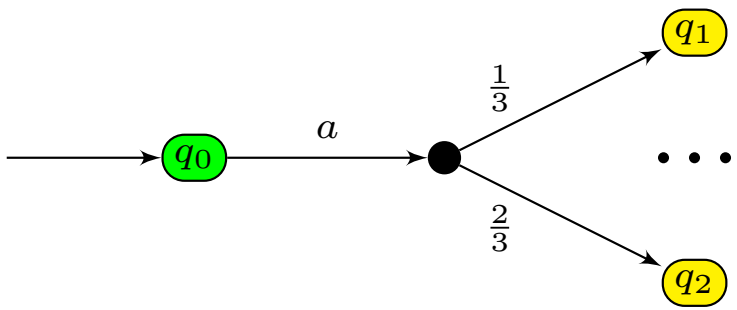

Given a sequence of observations,

a *strategy* randomly selects an action to be performed.

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

From pcLTS diagnosis to POMDP problems

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

- States are pairs (ℓ, q) with ℓ a state of $\mathcal{B}$ and q a state of $\mathcal{A}$ with $\mathsf{Obs}(\ell, q) = \ell$.
- Actions of $\mathcal{M}_{\mathcal{A}}$ are subset of events that includes the uncontrollable events.
- Given some event $b \in \Sigma^{\bullet}$ and a transition $\ell \xrightarrow{b}_{\mathcal{B}} \ell'$
  the transition probability of $\mathcal{M}_{\mathcal{A}}$ from (ℓ, q) to (ℓ', q')
  is the weight of $q \xrightarrow{b} q'$ appropriately normalized.
- In a state (ℓ, q) where no event of $\Sigma^{\bullet}$ is possible
  one reaches in $\mathcal{M}_{\mathcal{A}}$ an additional state `lost`.

Illustration

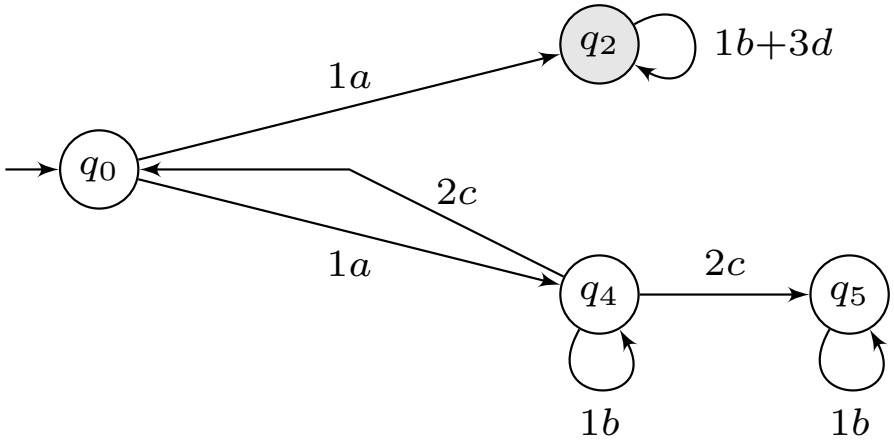

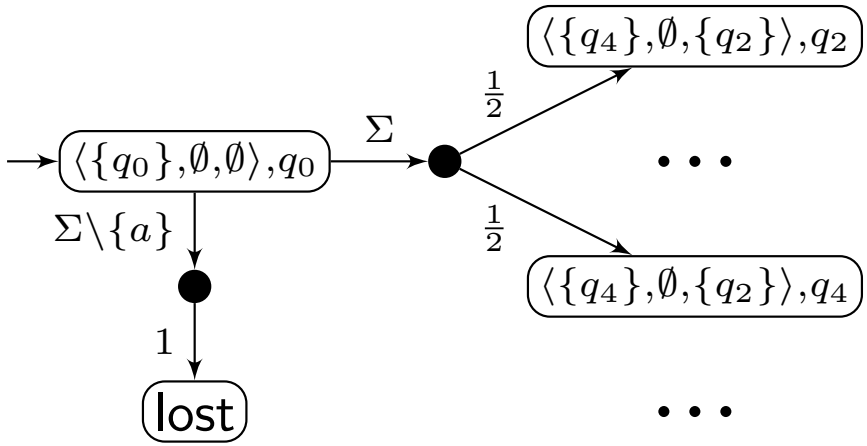

Complexity of the active diagnosis problem

- $\mathcal{A}$ is actively diagnosable iff there exists a strategy in $\mathcal{M}_{\mathcal{A}}$ such that:
$$\text{almost surely } \Box\Diamond(W = \emptyset \vee U = \emptyset)$$

- The existence of a strategy almost surely satisfying a Büchi objective in a POMDP is decidable.

- Analyzing the reduction to the POMDP problem, one gets that the active diagnosis problem is EXPTIME-complete.

The safe active diagnosis problem for pcLTS is undecidable.

Outline

Predictability in LTS

Let $\mathcal{A}$ be a LTS and ρ be a run.

- for $k \leq |\rho|$, $\rho_{\downarrow k}$ is the prefix of length k of ρ;
- for all k, $\rho_{\overleftarrow{k}} = \rho_{\downarrow \max(0,|\rho|-k)}$.

Let $\sigma \in \Sigma^*$. Then σ is *ultimately possibly correct* (UPC) if

$$\{\rho \in \mathsf{C}^\infty_{\mathcal{A}} \mid \sigma \text{ is a prefix of } \mathcal{P}(\rho)\} \neq \emptyset$$

Let $k \geq 1$. Then $\mathcal{A}$ is *k-predictable* if for all ρ minimal faulty run,

$$\mathcal{P}(\rho_{\overleftarrow{k}}) \text{ is not UPC.}$$

Illustration. $\mathcal{A}$ is 1-predictable.

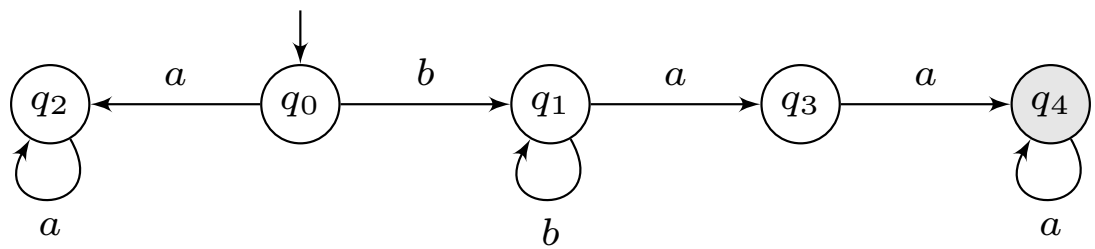

Informally in q_3, one can predict that the system will be faulty

while in q_1, there is still a possible infinite correct run.

Solving predictability in LTS

A characterization of predictability in finite LTS $\mathcal{A}$.

$\mathcal{A}$ is not k-predictable if there exist correct runs $q_0 \xrightarrow{\rho_0} q_1$ and $q_0 \xrightarrow{\rho'_0} q'_1$ such that:

- $\mathcal{P}(\rho_0) = \mathcal{P}(\rho'_0)$;
- there exists $q_1 \xrightarrow{\rho_1} q_2$ with $|\rho_1| \leq k$ and q_2 faulty;
- there exists $q'_1 \xrightarrow{\rho'_2} q'_2$ with q_2 correct belonging to a non trivial SCC.

By non deterministic explorations and applications of Immerman theorem one gets:

The predictability problem in LTS belongs to NLOGSPACE.

Predictability in pLTS

- One first strengthen the definition of UPC.

  Let $\sigma \in \Sigma^*$. Then σ is *ultimately possibly correct* (UPC) if

$$\mathbf{Pr}(\{\rho \in \mathsf{C}_{\mathcal{A}}^{\infty} \mid \sigma \text{ is a prefix of } \mathcal{P}(\rho)\}) > 0$$

- With this new definition, the formalization of predictability is unchanged.

  Let $k \geq 1$. Then $\mathcal{A}$ is *k-predictable* if for all ρ minimal faulty run,

$$\mathcal{P}(\rho_{\overleftarrow{k}}) \text{ is not UPC.}$$

Illustration. $\mathcal{A}$ is 2-predictable.

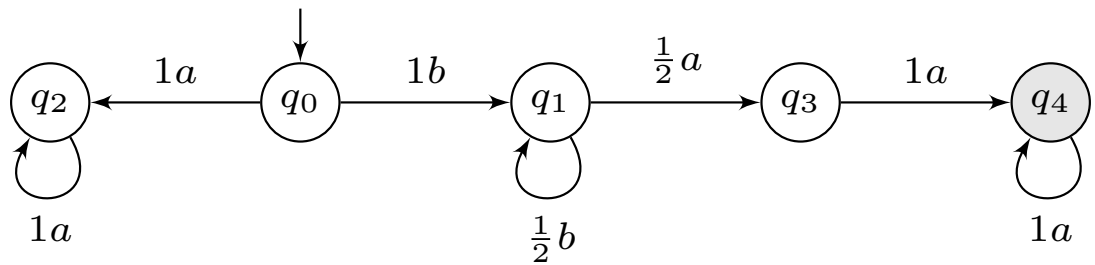

In q_1, the single infinite correct run has a null probability.

Solving predictability in pLTS

A characterization of predictability in finite pLTS $\mathcal{A}$.

$\mathcal{A}$ is not k-predictable if there exist correct runs $q_0 \xrightarrow{\rho_0} q_1$ and $q_0 \xrightarrow{\rho'_0} q'_1$ such that:

- $\mathcal{P}(\rho_0) = \mathcal{P}(\rho'_0)$;
- there exists $q_1 \xrightarrow{\rho_1} q_2$ with $|\rho_1| \leq k$ and q_2 faulty;
- there exists $q'_1 \xrightarrow{\rho'_2} q'_2$ with q_2 correct belonging to a BSCC.

By non deterministic explorations and applications of Immerman theorem one gets:

The predictability problem in pLTS belongs to NLOGSPACE.

Hardness of predictability

The reachability problem in directed acyclic graph (DAG)

is NLOGSPACE-complete.

Let G be a DAG and s, t be two vertices

and $\mathcal{A}$ the following LTS (or pLTS with uniform distribution) with $|\Sigma| = 1$.

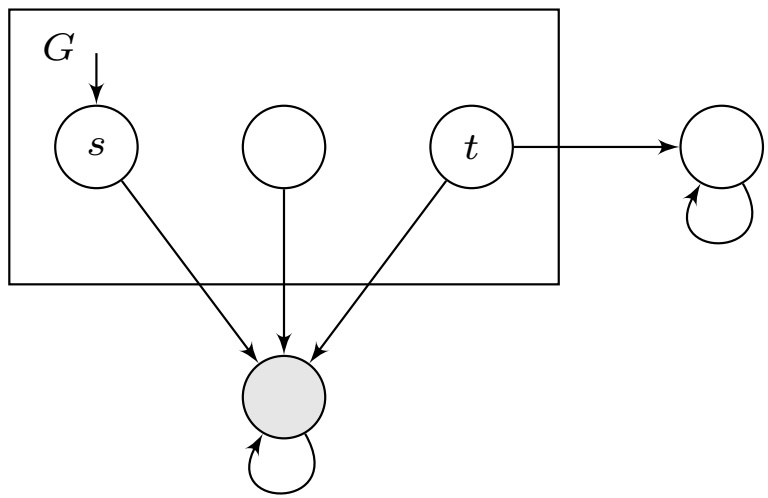

- If t is reachable from s in G then $\mathcal{A}$ is not k-predictable for any k.
- Otherwise $\mathcal{A}$ is k-predictable for all k.

The predictability problem is NLOGSPACE-hard in LTS and pLTS.

Refining Predictability in LTS

Refining predictability in a LTS $\mathcal{A}$.

Let $\sigma \in \Sigma^*$ and $\ell \in \mathbb{N}_\omega$. Then σ is *ℓ-possibly correct* (ℓ-PC) if

$$\{\rho \in \mathsf{C}_\mathcal{A} \mid \sigma \text{ is a prefix of } \mathcal{P}(\rho) \wedge |\rho| = |\sigma| + \ell\} \neq \emptyset$$

Let $1 \leq k \leq \ell$. Then $\mathcal{A}$ is *k-ℓ-predictable* if for all ρ minimal faulty run,

$$\mathcal{P}(\rho_{\overleftarrow{k}}) \text{ is not } \ell\text{-PC.}$$

Illustration.

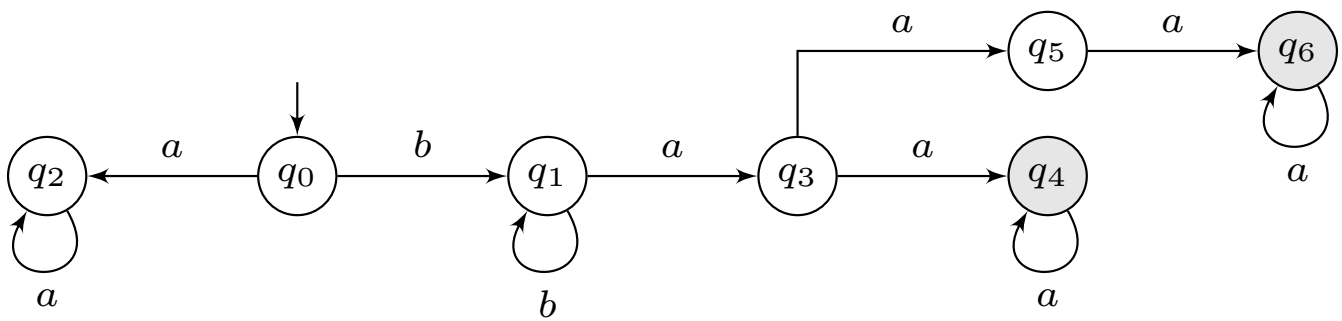

In q_3, the system will be faulty in at least 1 step and at most 2 steps.

Observation. In a finite k-predictable $\mathcal{A}$,

there exists $\ell < \omega$ such that $\mathcal{A}$ is k-ℓ-predictable.

Solving Active Predictability in cLTS

The active predictability problem.

Let $\mathcal{A}$ be a cLTS and $k \leq \ell$.

Does there exist a controller such that the controlled system is k-ℓ-predictable?

The active predictability problem is EXPTIME-conplete.

Observation. As for active diagnosability, one builds and solve a Büchi game.

However the kind of Büchi game that are built can decided in linear time

contrary the general case decidable in quadratic time.

Conclusion

- The specification of diagnosability is quite intricate
  with numerous possible formalizations.
- Furthermore every formalism triggers new semantical and algorithmic issues.
- In particular, the cLTS framework has led to optimization problems
  related to memory sizes and delays *(not discussed here)*.
- Predictability problems seem similar to diagnosability problems
  but their computational complexities may be different.

Bibliography

2013

- S. Haar, S. Haddad, T. Melliti, S. Schwoon. *Optimal Constructions for Active Diagnosis.* In FSTTCS'13, volume 24 of LIPIcs, pages 527-539

2014

- N. Bertrand, E. Fabre, S. Haar, S. Haddad, L. Hélouët. *Active Diagnosis for Probabilistic Systems.* In FoSSaCS'14, volume 8412 of LNCS, pages 29-42
- N. Bertrand, S. Haddad, E. Lefaucheux. *Foundation of Diagnosis and Predictability in Probabilistic Systems.* In FSTTCS'14, volume 29 of LIPIcs, pages 417-429

2015

- S. Boehm, S. Haar, S. Haddad, P. Hofman, S.Schwoon. *Active Diagnosis with Observable Quiescence.* In CDC'15, pages 1663-1668

2016

- N. Bertrand, S. Haddad, E. Lefaucheux. *Accurate Approximate Diagnosability of Stochastic Systems.* In LATA'16, volume 9618 of LNCS, pages 549-561

2017

- S. Haar, S. Haddad, T. Melliti, S. Schwoon. *Optimal Constructions for Active Diagnosis.* Journal of Computer and System Sciences, 83(1):101-120

2019

- N. Bertrand, S. Haddad, E. Lefaucheux. *A Tale of Two Diagnoses in Probabilistic Systems.* Information and Computation, 269.

2020

- N. Bertrand, S. Haddad, E. Lefaucheux. *Diagnosis and Degradation Control for Probabilistic Systems.* Discrete Event Dynamic Systems: Theory and Applications, 30:695-723
- S. Haar, S. Haddad, S. Schwoon, L. Ye. *Active Prediction for Discrete Event Systems.* In FSTTCS'20, volume 182 of LIPIcs, pages 48:1-48:16