Fault Diagnosis in Probabilistic Systems

Serge Haddad

LSV, ENS Cachan & CNRS & Inria, France

Intelligent Distributed Computing 2016

October the 11th 2016

joint work with Nathalie Bertrand[2,3,5,6], Stanislav Böhm[4], Eric Fabre[2], Stefan Haar[1,2,4], Loïc Hélouët[2], Piotr Hofman[4], Engel Lefaucheux[3,5,6], Tarek Melliti[1], Stefan Schwoon[1,4]

(1) FSTTCS 2013 (2) FOSSACS 2014 (3) FSTTCS 2014
(4) CDC 2015 (5) LATA 2016 (6) CONCUR 2016

Why diagnosis?

Faults and/or failures are unavoidable for some systems:

- Components have a finite lifetime;
- Reactive systems suffer pathological behaviour of the environment.

Why diagnosis?

Faults and/or failures are unavoidable for some systems:

- Components have a finite lifetime;
- Reactive systems suffer pathological behaviour of the environment.

Consequences of unhandled faults may be critical:

- Human casualties (e.g. pacemaker);
- Financial losses (e.g. mission to Mars).

Why diagnosis?

Faults and/or failures are unavoidable for some systems:

- Components have a finite lifetime;
- Reactive systems suffer pathological behaviour of the environment.

Consequences of unhandled faults may be critical:

- Human casualties (e.g. pacemaker);
- Financial losses (e.g. mission to Mars).

Necessity of a reactive and sound diagnoser.

Diagnosis: from failures to faults

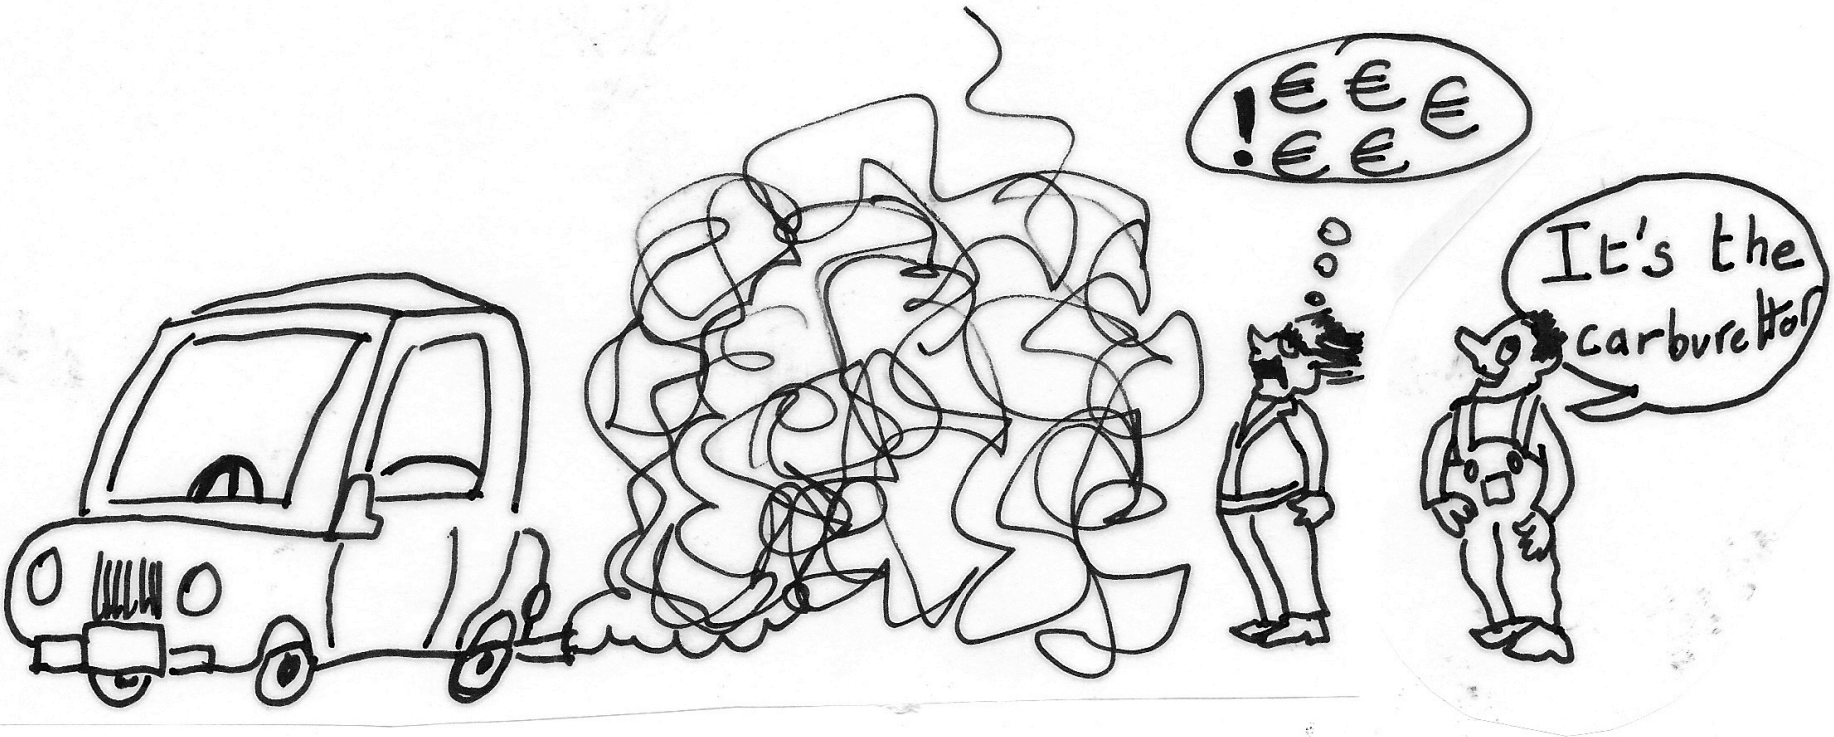

Example: MYCIN, an expert system, that used *artificial intelligence* to identify bacteria causing severe infections (1975).

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

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.

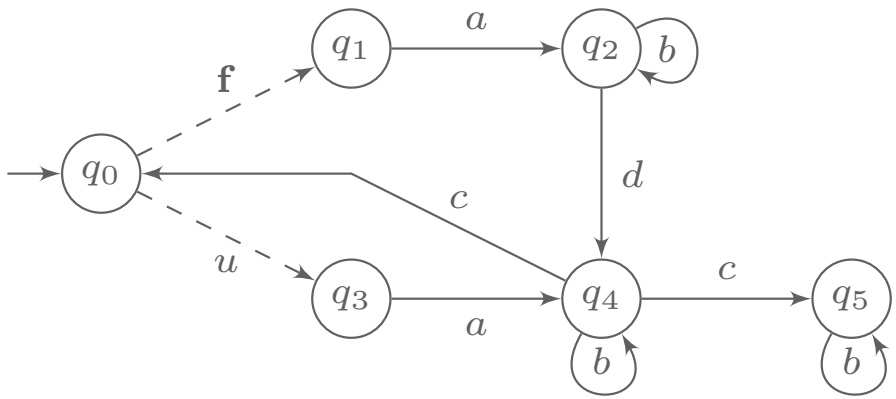

An *execution sequence* yields an *observed sequence*.

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.

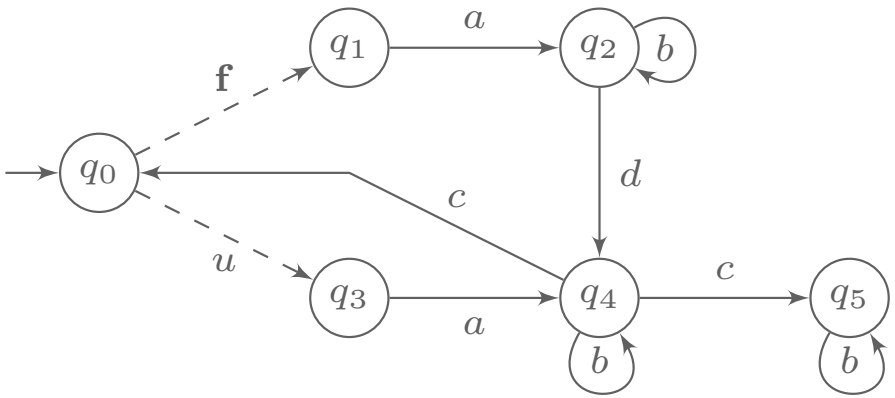

$adcb^\omega$ is surely faulty: the occurrence of d implies the occurrence of $\mathbf{f}$.

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

Diagnosis of discrete event systems

Objective: tell whether a fault $\mathbf{f}$ occurred, based on observations.

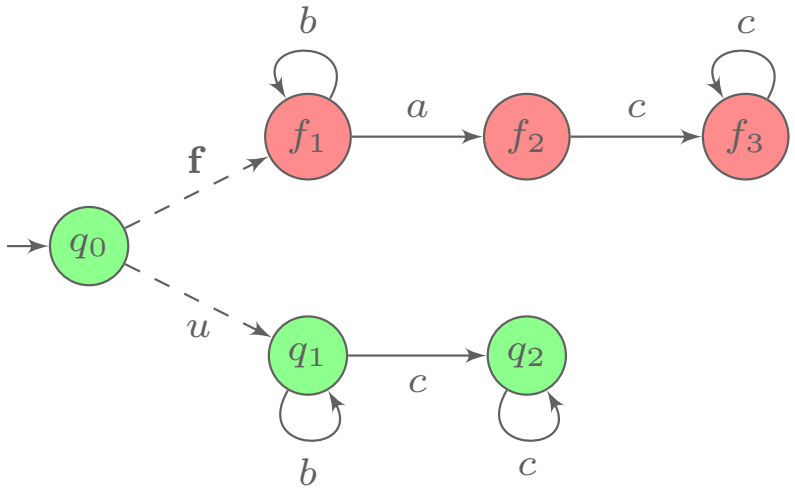

Diagnosis of discrete event systems

Objective: tell whether a fault $\mathbf{f}$ occurred, based on observations.

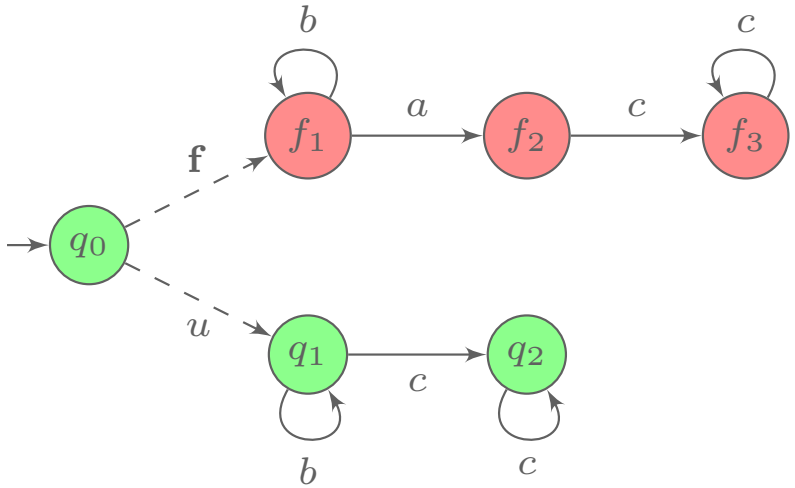

c^+	✓	correct
ac^+	✗	faulty
b^+	?	ambiguous

Diagnosis of discrete event systems

Objective: tell whether a fault $\mathbf{f}$ occurred, based on observations.

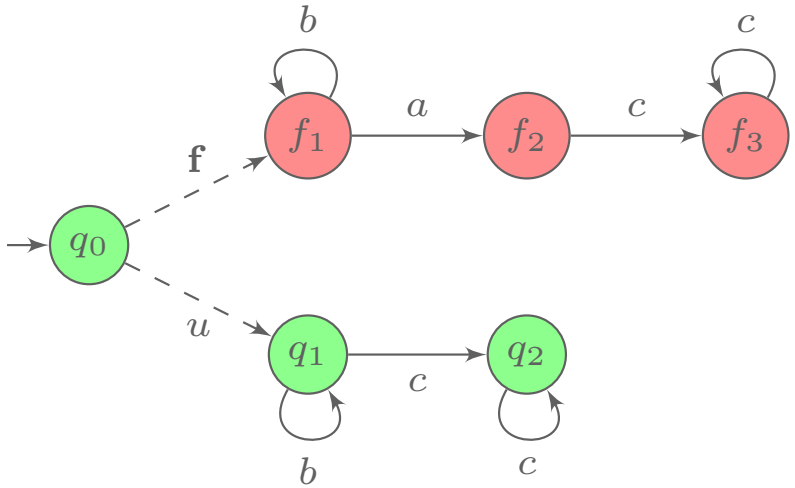

c^+	✓	correct
ac^+	✗	faulty
b^+	?	ambiguous

Diagnosability (in this context): all observed sequences are unambiguous.

Diagnosis of discrete event systems

Objective: tell whether a fault $\mathbf{f}$ occurred, based on observations.

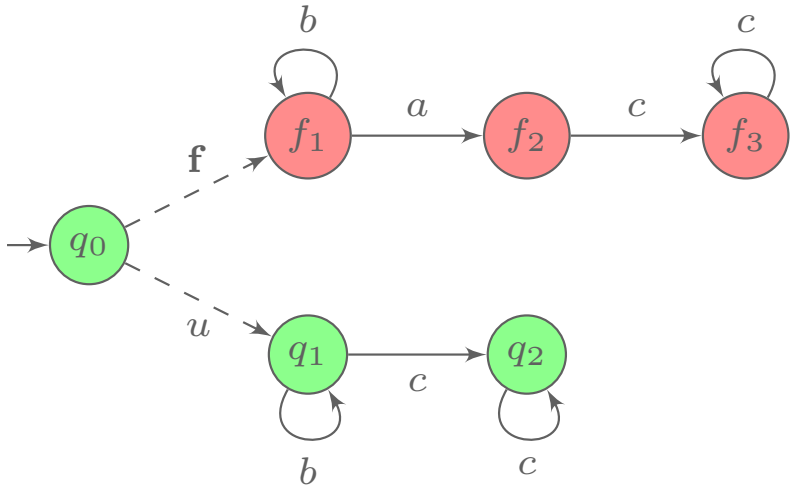

c^+	✓	correct
ac^+	✗	faulty
b^+	?	ambiguous

Diagnosability (in this context): all observed sequences are unambiguous.

Diagnoser: assigns verdicts to observed sequences $D : \Sigma_o^* \to \{✓, ✗, ?\}$

- **Soundness**: if a fault is claimed, a fault occurred.
- **Reactivity**: every fault will be detected.

pLTS

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

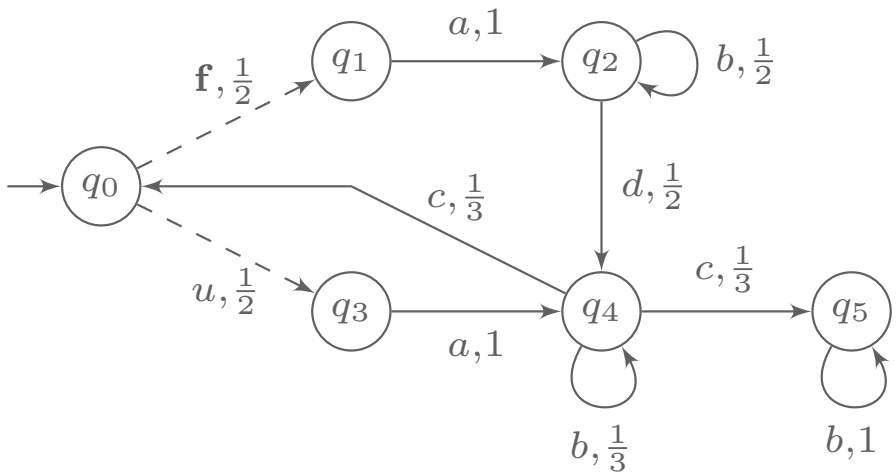

Without labels, a pLTS is a discrete time Markov chain.

Without transition probabilities, a pLTS is a LTS.

Diagnosis of probabilistic systems

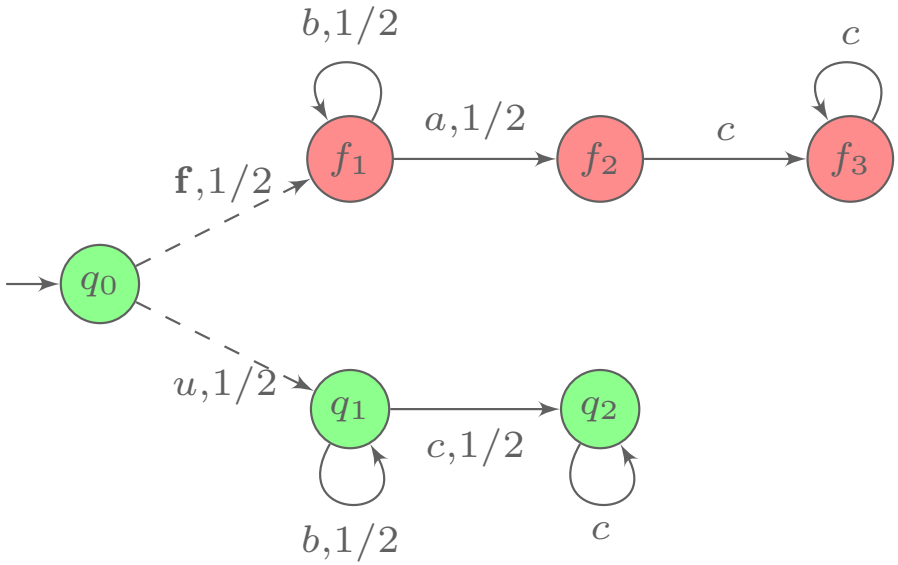

Diagnosis of probabilistic systems

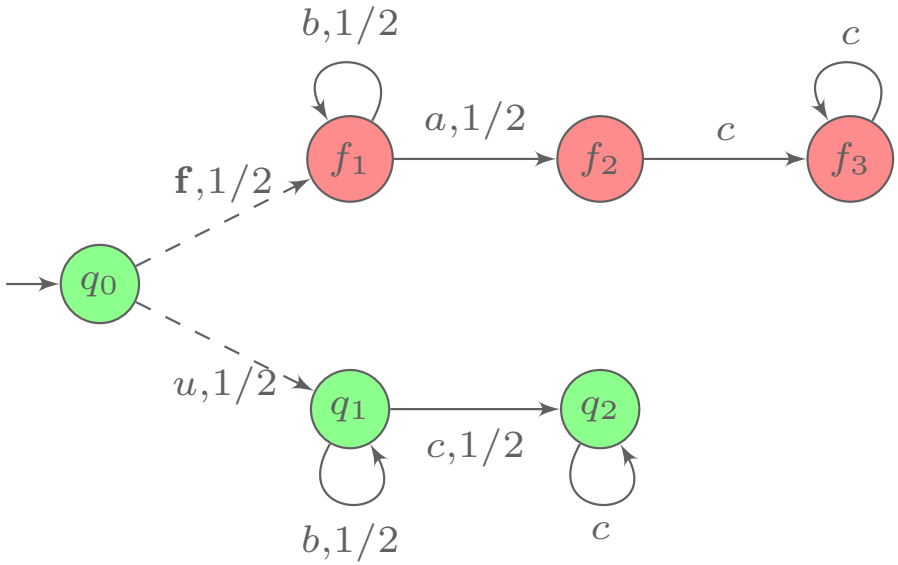

b^+ ambiguous but...

Diagnosis of probabilistic systems

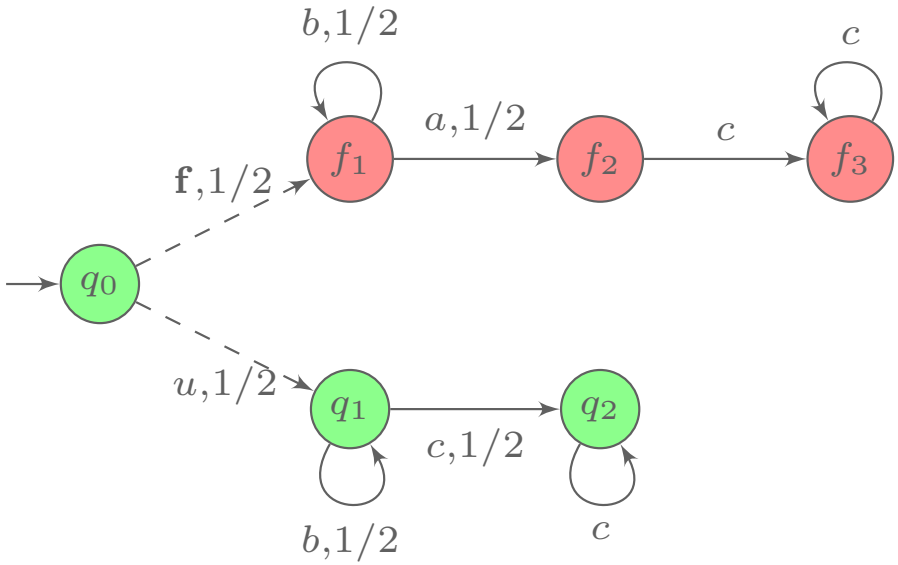

b^+ ambiguous but...

$$\lim_{n\to\infty} \mathbb{P}(\mathbf{f}b^n + ub^n) = 0$$

Diagnosis of probabilistic systems

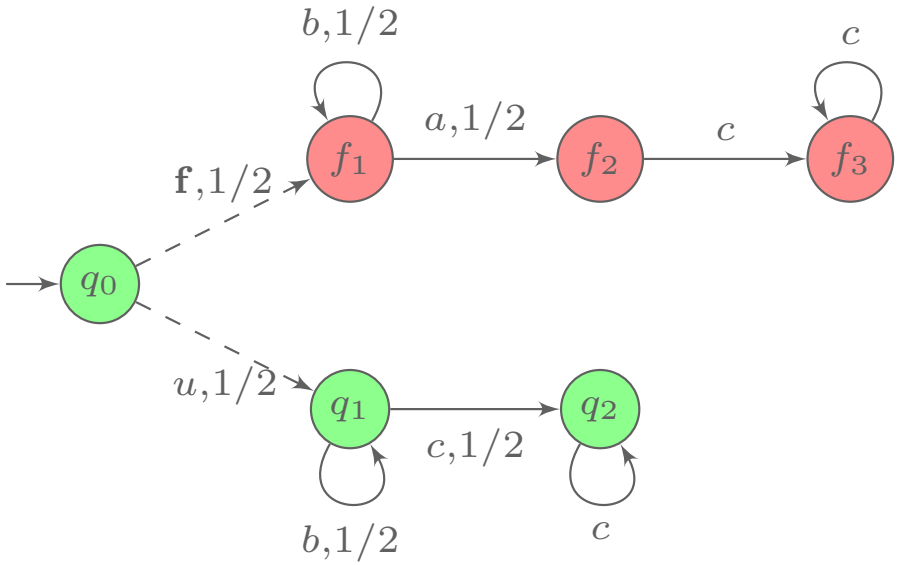

b^+ ambiguous but...

$$\lim_{n\to\infty} \mathbb{P}(\mathbf{f}b^n + ub^n) = 0$$

How to handle probabilities?

Diagnosis of probabilistic systems

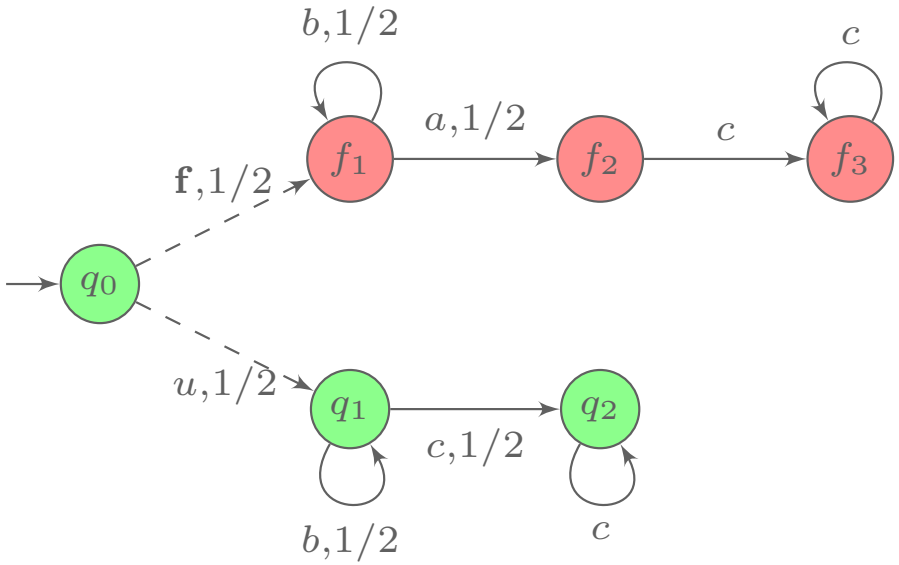

b^+ ambiguous but...

$$\lim_{n\to\infty} \mathbb{P}(\mathbf{f}b^n + ub^n) = 0$$

How to handle probabilities?

A first answer: discard pathologic behaviours (i.e. with null probability).

Outline

All runs or faulty runs?

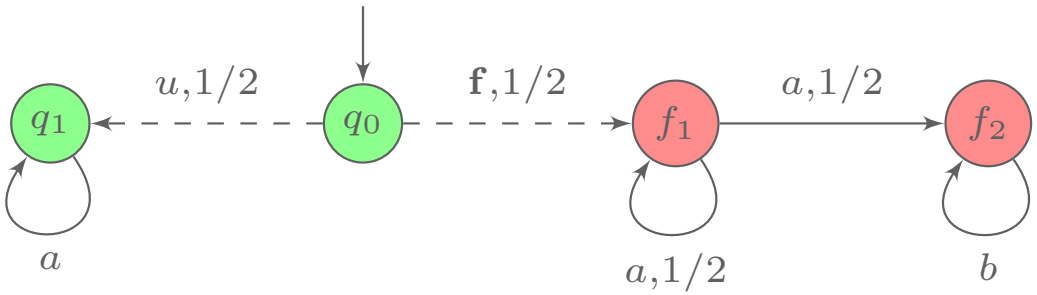

All runs or faulty runs?

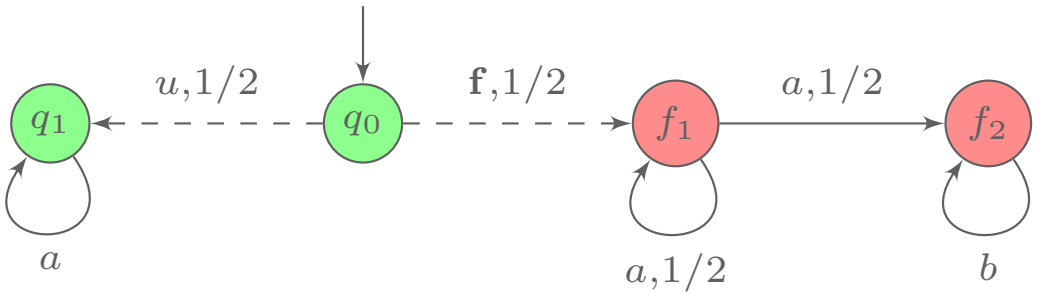

a^+ is ambiguous

All runs or faulty runs?

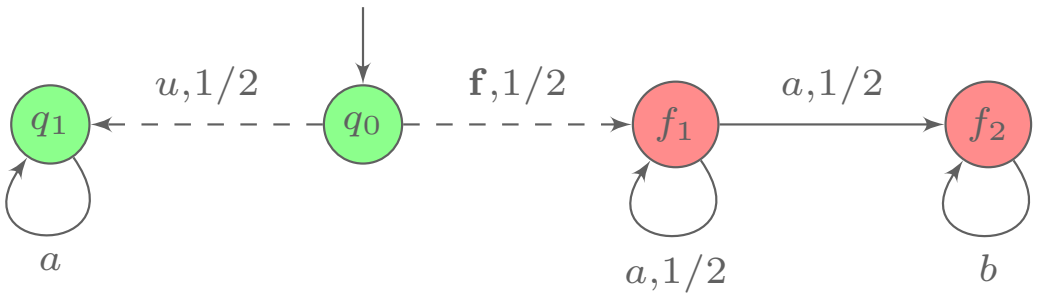

a^+ is ambiguous

$$\lim_{n\to\infty} \mathbb{P}(\mathbf{f}a^n) = 0$$

All runs or faulty runs?

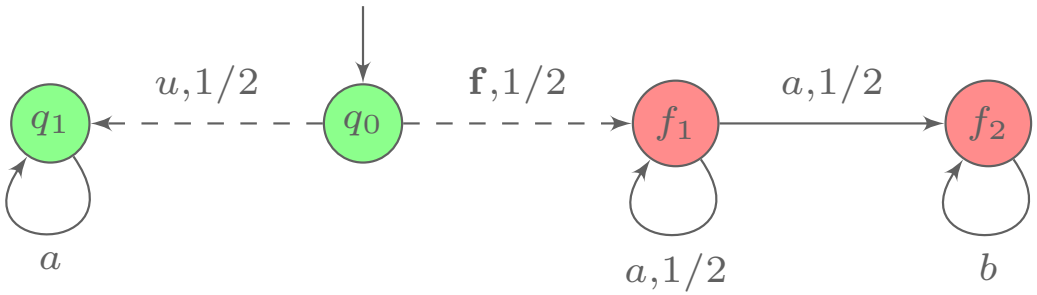

a^+ is ambiguous

$$\lim_{n\to\infty} \mathbb{P}(\mathbf{f}a^n) = 0$$

$$\lim_{n\to\infty} \mathbb{P}(ua^n) = \frac{1}{2}$$

All runs or faulty runs?

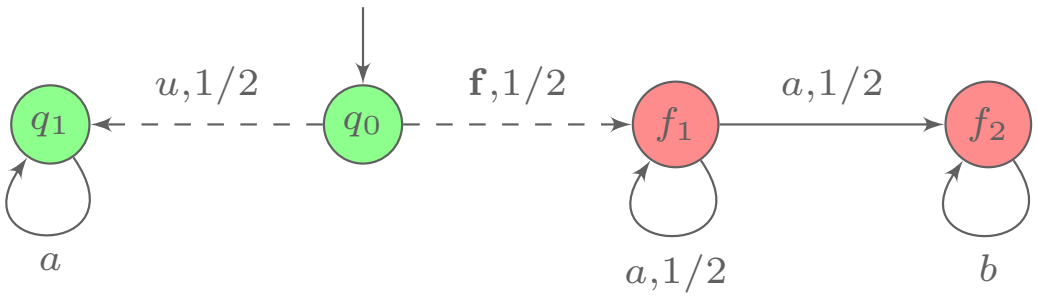

a^+ is ambiguous

$$\lim_{n\to\infty} \mathbb{P}(\mathbf{f}a^n) = 0$$

$$\lim_{n\to\infty} \mathbb{P}(ua^n) = \frac{1}{2}$$

Reactivity specifications:

- Detect a fault, almost surely.

All runs or faulty runs?

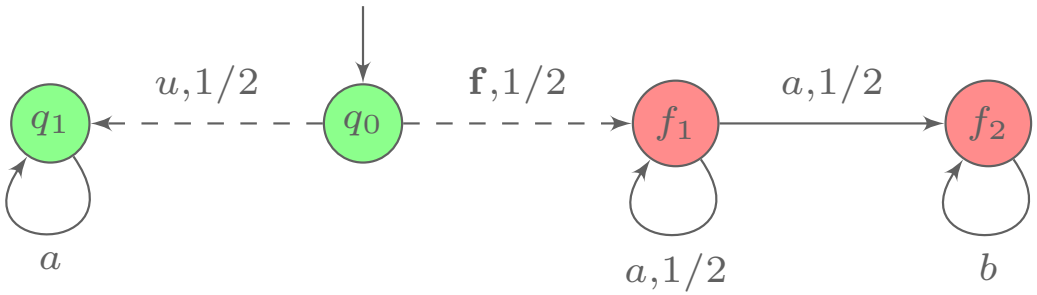

a^+ is ambiguous

$$\lim_{n\to\infty} \mathbb{P}(\mathbf{f}a^n) = 0$$

$$\lim_{n\to\infty} \mathbb{P}(ua^n) = \frac{1}{2}$$

Reactivity specifications:

- Detect a fault, almost surely.
- Detect if a run is faulty or correct, almost surely.

Infinite sequences or their finite prefixes?

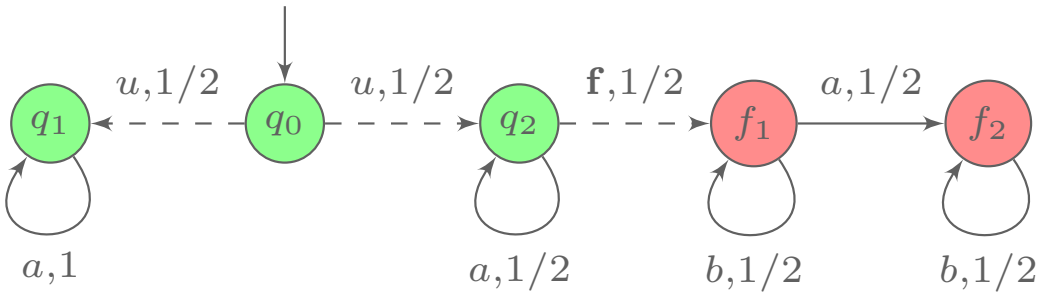

Infinite sequences or their finite prefixes?

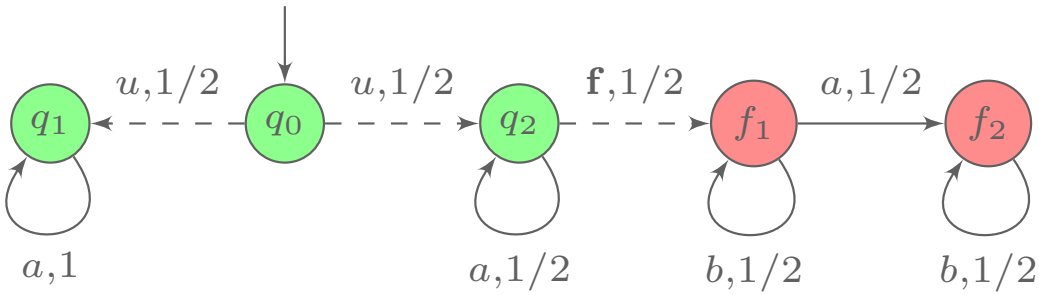

a^ω is surely correct.

Infinite sequences or their finite prefixes?

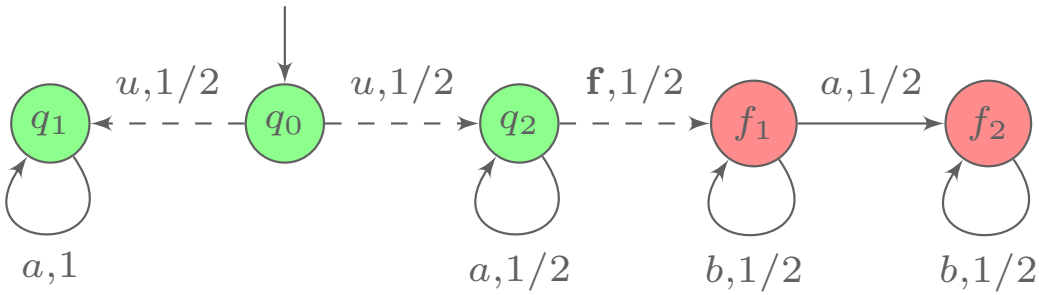

a^ω is surely correct.

a^n is ambiguous and $\mathbb{P}(q_0 u (q_1 a)^n) = \frac{1}{2}$.

Infinite sequences or their finite prefixes?

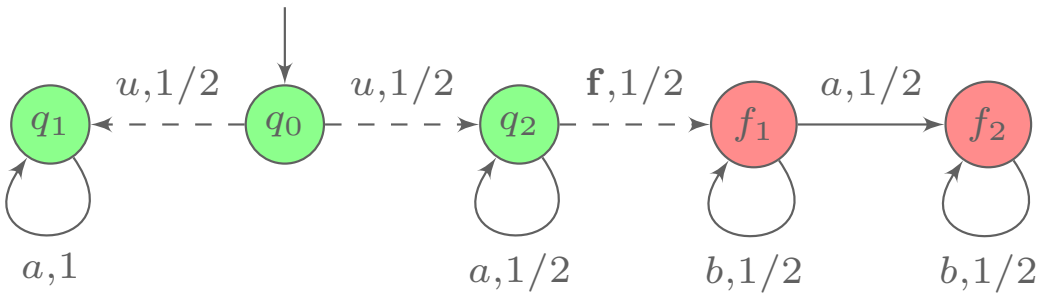

a^ω is surely correct.

a^n is ambiguous and $\mathbb{P}(q_0u(q_1a)^n) = \frac{1}{2}$.

- Infinite sequences are almost surely non ambiguous.

Infinite sequences or their finite prefixes?

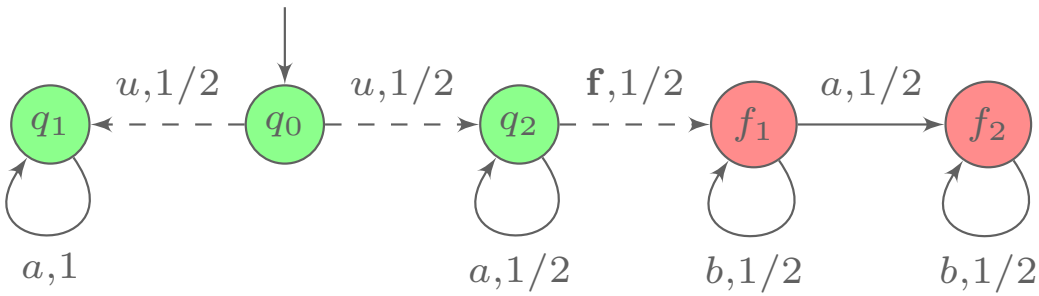

a^ω is surely correct.

a^n is ambiguous and $\mathbb{P}(q_0u(q_1a)^n) = \frac{1}{2}$.

- Infinite sequences are almost surely non ambiguous.
- The probability of ambiguous prefixes tends to 0.

Four diagnosability notions

Diagnosability	All runs		Faulty runs
Finite prefixes	FA	$\Rightarrow$ $\not\Leftarrow$	FF
	$\Downarrow$ $\not\Uparrow$		$\Downarrow\Uparrow^*$
Infinite sequences	IA	$\Rightarrow$ $\not\Leftarrow$	IF

* assuming finitely-branching models

Quest for a characterisation

Objective: a characterisation that can be split between the qualitative and the quantitative parts of the system.

$\mathcal{N}$ is diagnosable iff $\mathbb{P}_{\mathcal{N}}(B) \bowtie p$, where:

- $p \in \{0, 1\}$, $\bowtie \in \{<, =, >\}$;
- ($\star$) B belongs to a low level of Borel hierarchy and
  ($\star$) B only depends on the underlying LTS.

Quest for a characterisation

Objective: a characterisation that can be split between the qualitative and the quantitative parts of the system.

$\mathcal{N}$ is diagnosable iff $\mathbb{P}_{\mathcal{N}}(B) \bowtie p$, where:

- $p \in \{0, 1\}$, $\bowtie \in \{<, =, >\}$;
- ($\star$) B belongs to a low level of Borel hierarchy and
  ($\star$) B only depends on the underlying LTS.

Definitions are not directly applicable:

•	IA	$\mathbb{P}(\mathsf{Amb}_\infty) = 0$	Amb_∞ analytic set, a priori not Borel
•	IF	$\mathbb{P}(\mathsf{FAmb}_\infty) = 0$	FAmb_∞ analytic set, a priori not Borel

Quest for a characterisation

Objective: a characterisation that can be split between the qualitative and the quantitative parts of the system.

$\mathcal{N}$ is diagnosable iff $\mathbb{P}_{\mathcal{N}}(B) \bowtie p$, where:

- $p \in \{0, 1\}$, $\bowtie \in \{<, =, >\}$;
- $(\star)$ B belongs to a low level of Borel hierarchy and
  $(\star)$ B only depends on the underlying LTS.

Definitions are not directly applicable:

• IA	$\mathbb{P}(\mathsf{Amb}_\infty) = 0$	Amb_∞ analytic set, a priori not Borel
• IF	$\mathbb{P}(\mathsf{FAmb}_\infty) = 0$	FAmb_∞ analytic set, a priori not Borel
• FA	$\lim_{n\to\infty} \mathbb{P}(\mathsf{Amb}_n) = 0$	$(\mathsf{Amb}_n)_{n\in\mathbb{N}}$ family of Borel sets
• FF	$\lim_{n\to\infty} \mathbb{P}(\mathsf{FAmb}_n) = 0$	$(\mathsf{FAmb}_n)_{n\in\mathbb{N}}$ family of Borel sets

Characterisation in pathL: an expressive linear temporal logic

$\phi ::= \alpha \mid \neg\phi \mid \phi \wedge \phi \mid \Diamond\phi$ where α is a finite path formula

Characterisation in pathL: an expressive linear temporal logic

$\phi ::= \alpha \mid \neg\phi \mid \phi \wedge \phi \mid \Diamond\phi$ where α is a finite path formula

pathL subsumes all ω-regular linear specification languages

Characterisation in pathL: an expressive linear temporal logic

$\phi ::= \alpha \mid \neg\phi \mid \phi \wedge \phi \mid \Diamond\phi$ where α is a finite path formula

pathL subsumes all ω-regular linear specification languages

- $\mathfrak{f}(\rho) \equiv \rho$ faulty
- $\mathfrak{U}(\rho) \equiv \exists \rho'$ correct s.t. $\mathcal{P}(\rho) = \mathcal{P}(\rho')$

$\mathcal{N}$ is FF-diagnosable iff $\mathcal{N} \models \mathbb{P}^{=0}(\Diamond\Box(\mathfrak{U} \wedge \mathfrak{f}))$.

also valid for IF*-diagnosability if* $\mathcal{N}$ *is finitely-branching*

Characterisation in pathL: an expressive linear temporal logic

$\phi ::= \alpha \mid \neg\phi \mid \phi \wedge \phi \mid \Diamond\phi$ where α is a finite path formula

pathL subsumes all ω-regular linear specification languages

- $\mathfrak{f}(\rho) \equiv \rho$ faulty
- $\mathfrak{U}(\rho) \equiv \exists \rho'$ correct s.t. $\mathcal{P}(\rho) = \mathcal{P}(\rho')$

$\mathcal{N}$ is FF-diagnosable iff $\mathcal{N} \models \mathbb{P}^{=0}(\Diamond\Box(\mathfrak{U} \wedge \mathfrak{f}))$.

also valid for IF*-diagnosability if* $\mathcal{N}$ *is finitely-branching*

- $\mathfrak{W}(\rho) \equiv$ last observation does not change the time of the earliest possible fault

$\mathcal{N}$, finitely branching, is IA-diagnosable iff $\mathcal{N} \models \mathbb{P}^{=0}(\Diamond\Box(\mathfrak{U} \wedge \mathfrak{W}))$.

Expressing FA-diagnosability is hard!

There does not exist a F_σ set B only depending on the underlying LTS such that $\mathbb{P}(B) = 0$ characterises FA-diagnosability.

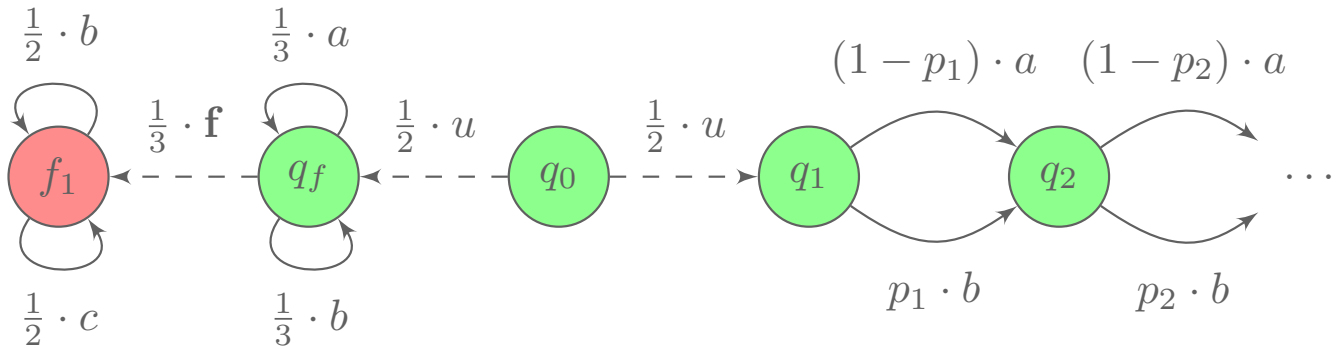

Expressing FA-diagnosability is hard!

There does not exist a F_σ set B only depending on the underlying LTS such that $\mathbb{P}(B) = 0$ characterises FA-diagnosability.

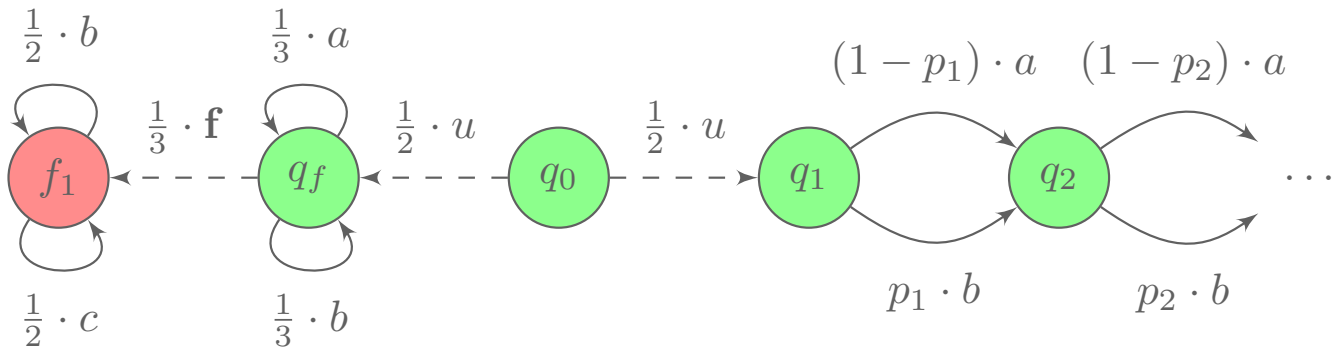

There does not exist a Borel set B only depending on the underlying LTS such that $\mathbb{P}(B) > 0$ characterises FA-diagnosability.

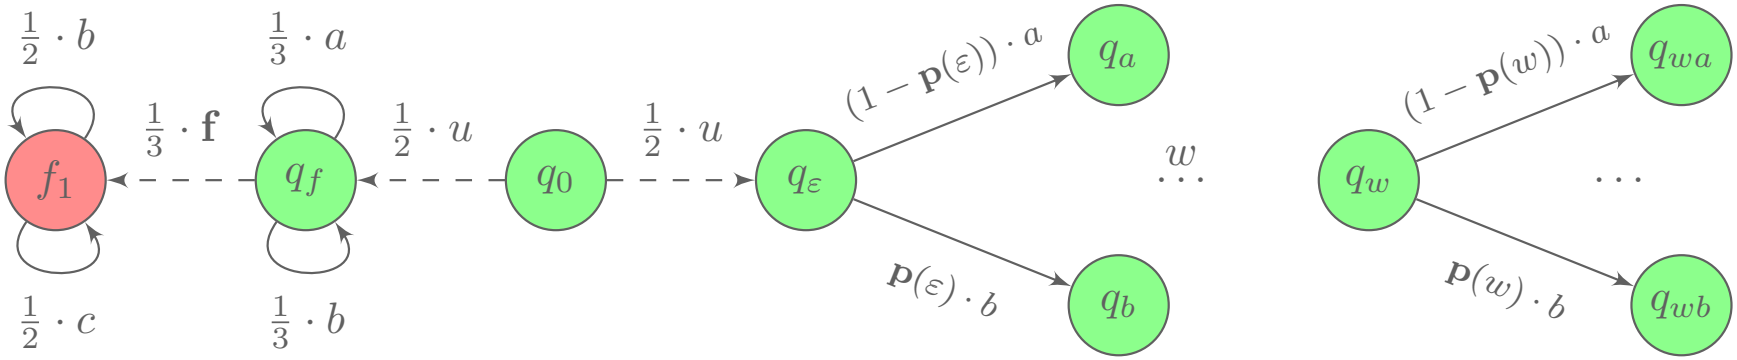

Predictability

Objective: tell whether a fault *will* occur, based on observations.

Predictability

Objective: tell whether a fault *will* occur, based on observations.

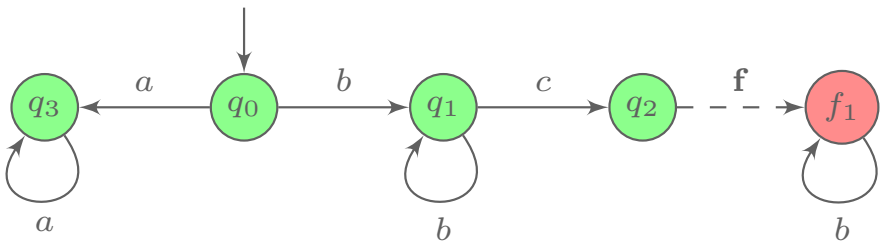

a^+	✓	correct
b^+c	✗	surely eventually faulty
b^+	✗	a.s. eventually faulty

Predictability

Objective: tell whether a fault *will* occur, based on observations.

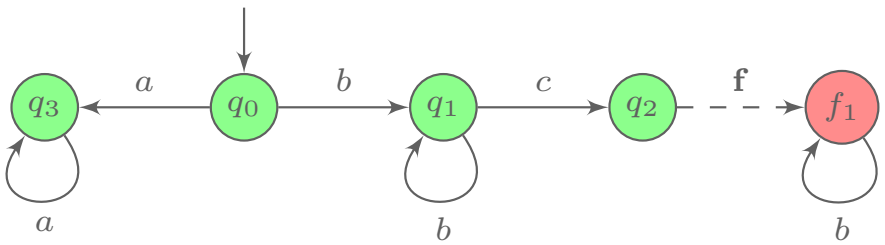

a^+	✓	correct
b^+c	✗	surely eventually faulty
b^+	✗	a.s. eventually faulty

Two notions of **soundness**:

- sure: if a fault is claimed, a fault will occur
- almost-sure: if a fault is claimed, a fault will almost-surely occur

Reactivity: a fault is detected at least k steps before occurrence.

Predictability

Objective: tell whether a fault *will* occur, based on observations.

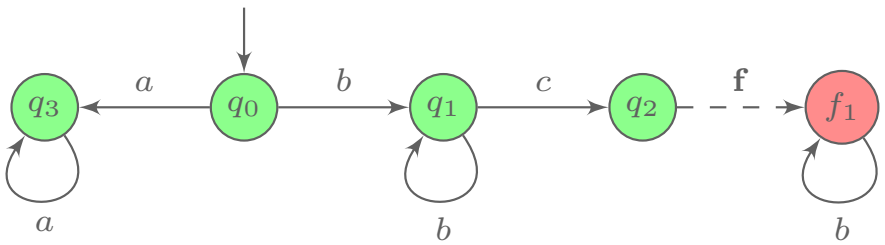

a^+	✓	correct
b^+c	✗	surely eventually faulty
b^+	✗	a.s. eventually faulty

surely 0-predictable almost surely 1-predictable not 2-predictable

Two notions of **soundness**:

- sure: if a fault is claimed, a fault will occur
- almost-sure: if a fault is claimed, a fault will almost-surely occur

Reactivity: a fault is detected at least k steps before occurrence.

Prediagnosability

Objective: detect and foresee faults analysing past and future

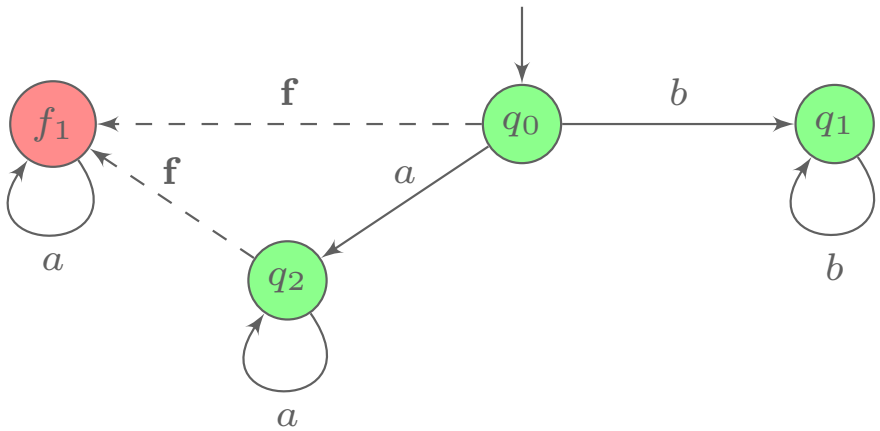

b^+ ✓ correct

a^+ ✗ a.s. eventually faulty

Prediagnosability

Objective: detect and foresee faults analysing past and future

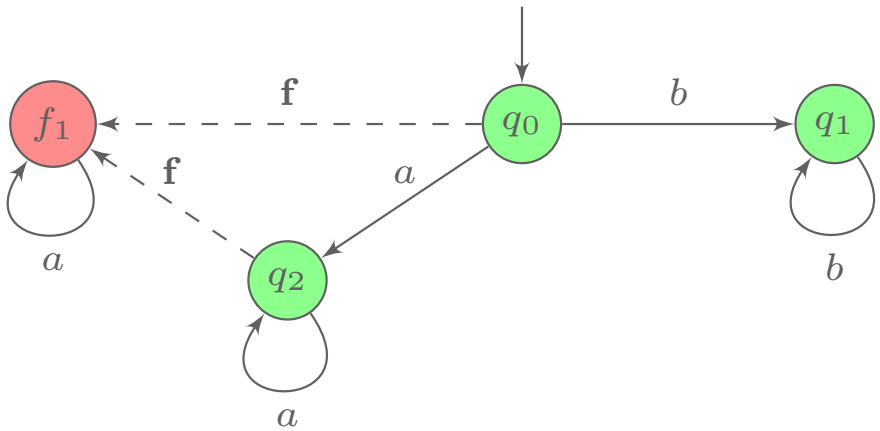

b^+	✓	correct
a^+	✗	a.s. eventually faulty

Soundness: If a fault is claimed, a fault happened or (almost) surely will.

Reactivity: Faults are almost surely claimed.

Prediagnosability

Objective: detect and foresee faults analysing past and future

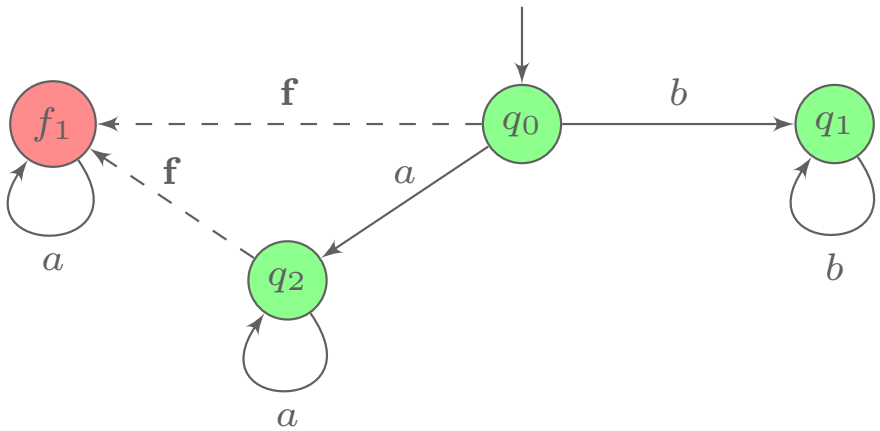

b^+	✓	correct
a^+	✗	a.s. eventually faulty

Soundness: If a fault is claimed, a fault happened or (almost) surely will.

Reactivity: Faults are almost surely claimed.

Relations between the specifications

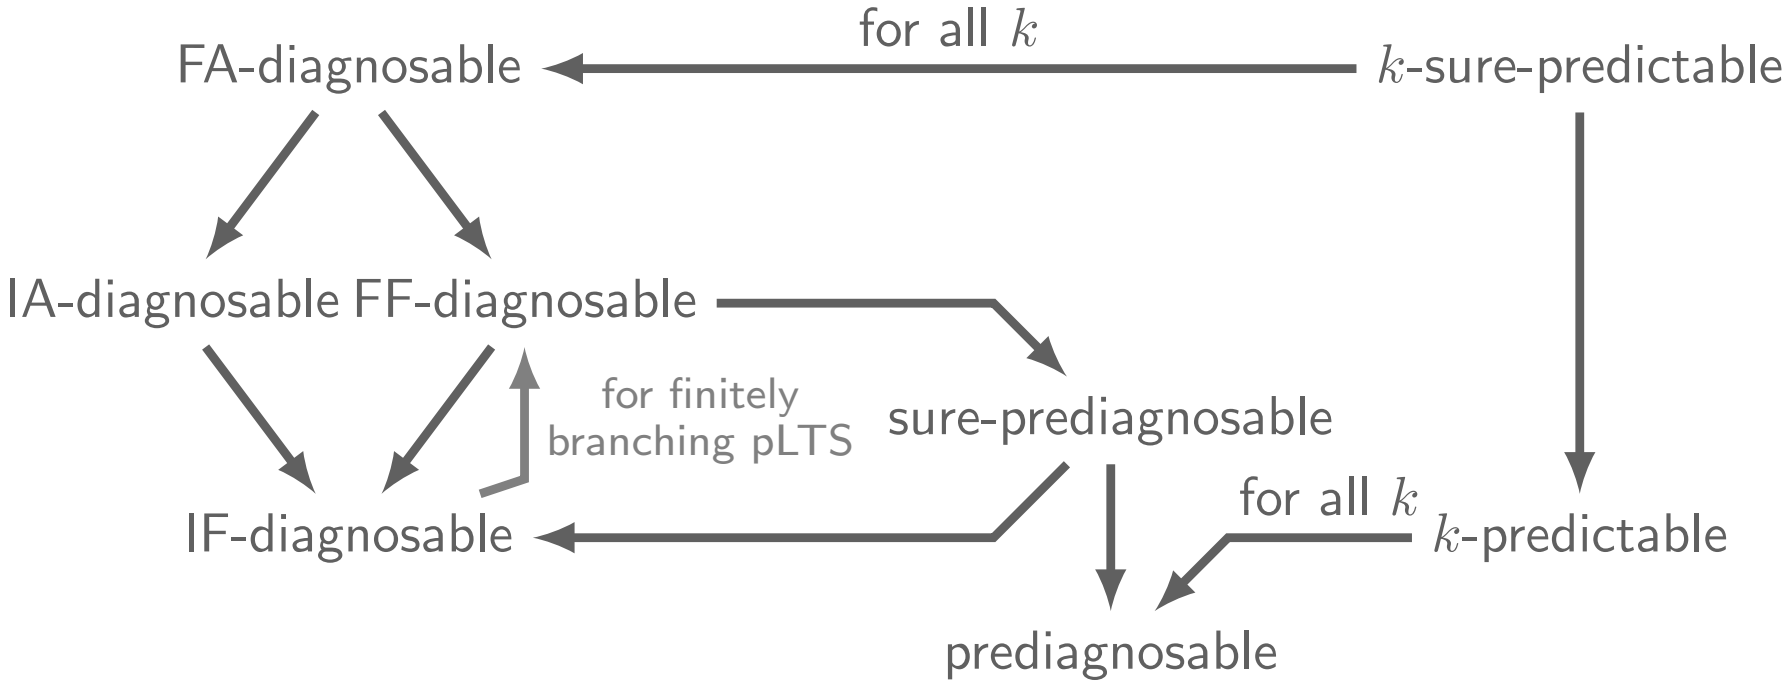

Outline

Exact diagnosis versus approximate diagnosis

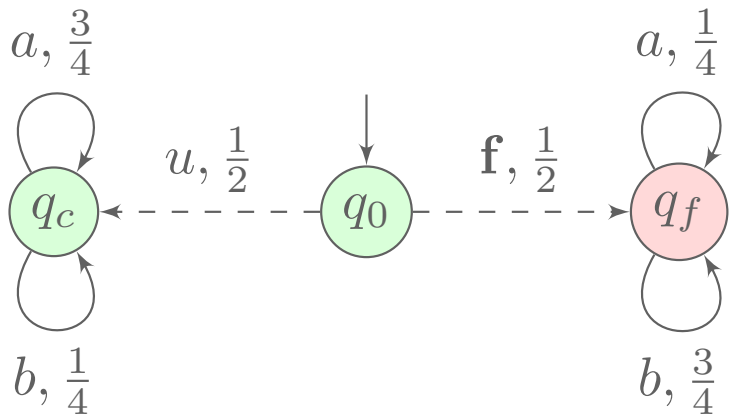

Exact diagnosis versus approximate diagnosis

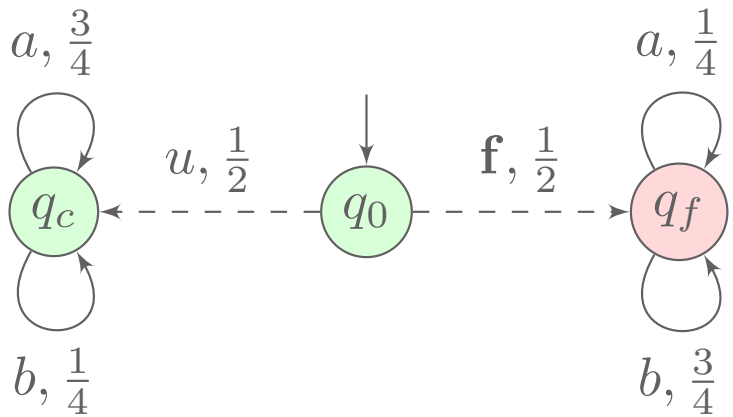

Not exactly diagnosable

Exact diagnosis versus approximate diagnosis

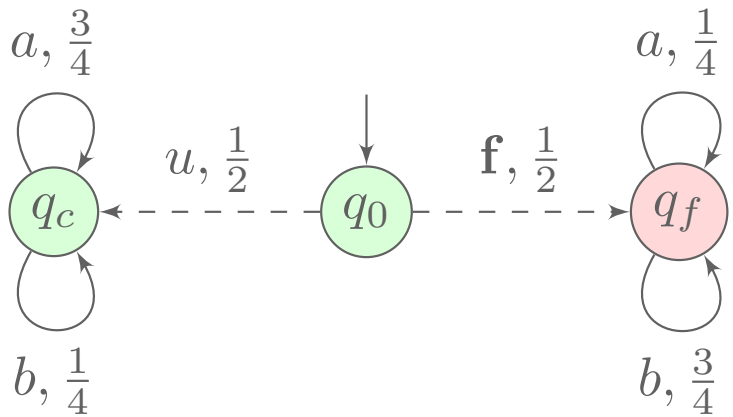

Not exactly diagnosable

However a high proportion of b implies a highly probable faulty run.

Exact diagnosis versus approximate diagnosis

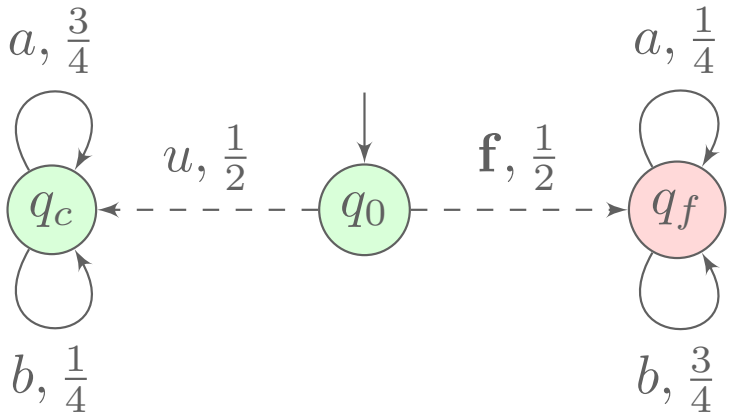

Not exactly diagnosable

However a high proportion of b implies a highly probable faulty run.

Relaxed Soundness: if a fault is claimed the probability of error is small.

Proportion of correct runs

Given an observation sequence $\sigma \in \Sigma_o^*$,

$$\mathsf{CorP}(\sigma) = \frac{\mathbb{P}(\{\pi^{-1}(\sigma) \cap \mathsf{correct}\})}{\mathbb{P}(\{\pi^{-1}(\sigma)\})}$$

Proportion of correct runs

Given an observation sequence $\sigma \in \Sigma_o^*$,

$$\mathsf{CorP}(\sigma) = \frac{\mathbb{P}(\{\pi^{-1}(\sigma) \cap \mathsf{correct}\})}{\mathbb{P}(\{\pi^{-1}(\sigma)\})}$$

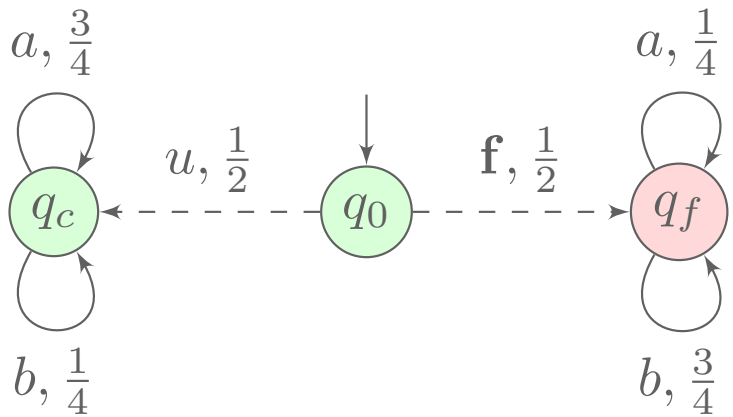

$\mathsf{CorP}(a) = 3/4$,

Proportion of correct runs

Given an observation sequence $\sigma \in \Sigma_o^*$,

$$\mathsf{CorP}(\sigma) = \frac{\mathbb{P}(\{\pi^{-1}(\sigma) \cap \mathsf{correct}\})}{\mathbb{P}(\{\pi^{-1}(\sigma)\})}$$

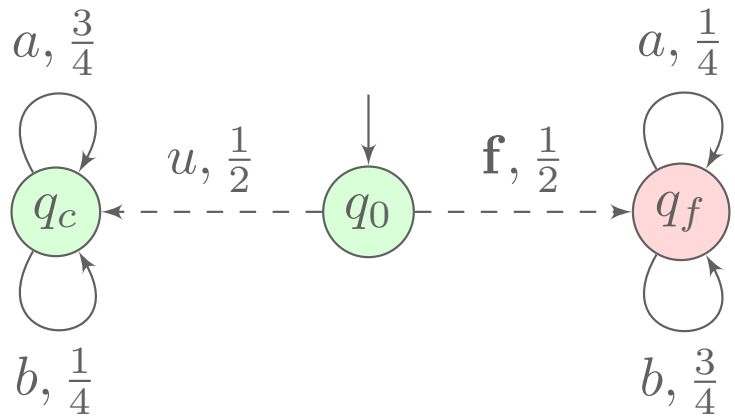

$\mathsf{CorP}(a) = 3/4$, $\mathsf{CorP}(ab) = 1/2$,

Proportion of correct runs

Given an observation sequence $\sigma \in \Sigma_o^*$,

$$\mathsf{CorP}(\sigma) = \frac{\mathbb{P}(\{\pi^{-1}(\sigma) \cap \mathsf{correct}\})}{\mathbb{P}(\{\pi^{-1}(\sigma)\})}$$

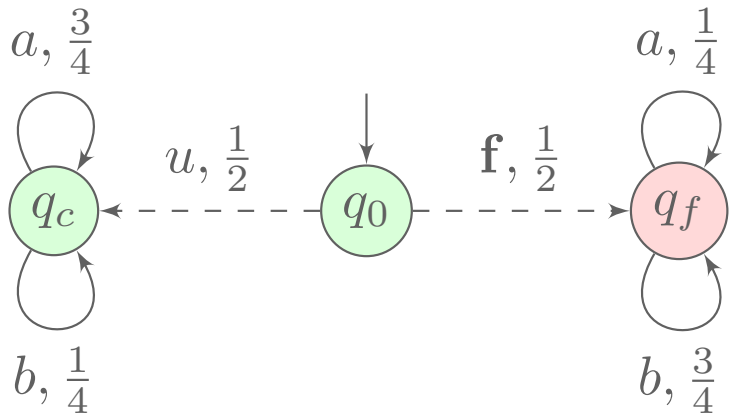

$\mathsf{CorP}(a) = 3/4$, $\mathsf{CorP}(ab) = 1/2$, $\mathsf{CorP}(abb) = 1/4$, $\mathsf{CorP}(abbb) = 1/10$.

Relaxing Soundness

Given $\varepsilon \geq 0$, an ε-*diagnoser* fulfills

- Soundness: If a fault is claimed after an observation sequence σ, then $\mathrm{CorP}(\sigma) \leq \varepsilon$.

Relaxing Soundness

Given $\varepsilon \geq 0$, an *ε-diagnoser* fulfills

- Soundness: If a fault is claimed after an observation sequence σ, then $\mathsf{CorP}(\sigma) \leq \varepsilon$.
- Reactivity: Given a faulty run ρ, the measure of undetected runs extending ρ converges to 0.

Relaxing Soundness

Given $\varepsilon \geq 0$, an *ε-diagnoser* fulfills

- Soundness: If a fault is claimed after an observation sequence σ, then $\mathrm{CorP}(\sigma) \leq \varepsilon$.
- Reactivity: Given a faulty run ρ, the measure of undetected runs extending ρ converges to 0.

A *uniform ε-diagnoser* ensures for reactivity a uniform convergence over the faulty runs.

Relaxing Soundness

Given $\varepsilon \geq 0$, an *ε-diagnoser* fulfills

- Soundness: If a fault is claimed after an observation sequence σ, then $\mathsf{CorP}(\sigma) \leq \varepsilon$.
- Reactivity: Given a faulty run ρ, the measure of undetected runs extending ρ converges to 0.

A *uniform ε-diagnoser* ensures for reactivity a uniform convergence over the faulty runs.

0-diagnosers correspond to (exact) FF-diagnosers.

Relaxing Soundness

Given $\varepsilon \geq 0$, an *ε-diagnoser* fulfills

- Soundness: If a fault is claimed after an observation sequence σ, then $\mathrm{CorP}(\sigma) \leq \varepsilon$.
- Reactivity: Given a faulty run ρ, the measure of undetected runs extending ρ converges to 0.

A *uniform ε-diagnoser* ensures for reactivity a uniform convergence over the faulty runs.

0-diagnosers correspond to (exact) FF-diagnosers.

A model is (uniformly) AA-diagnosable, for accurately approximately diagnosable, if it is (uniformly) ε-diagnosable for all $\varepsilon > 0$.

Illustration

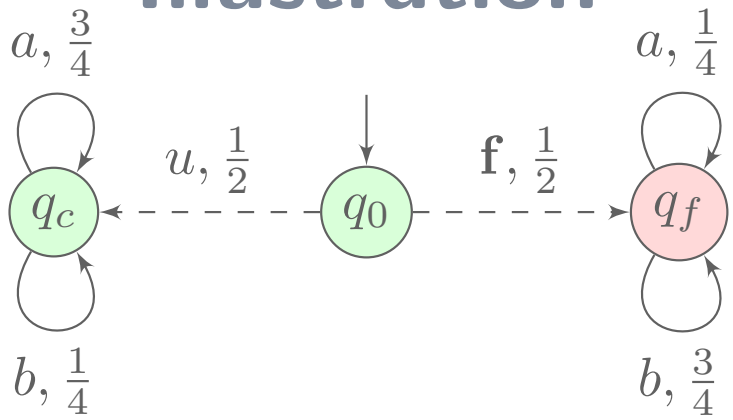

Illustration

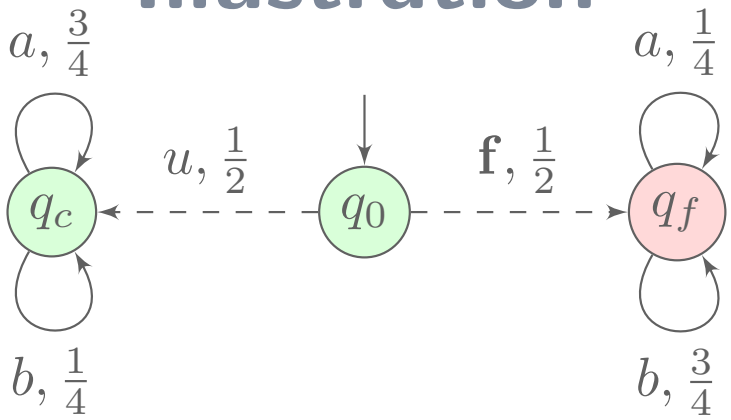

- **AA-diagnosable**

Let $\rho = q_0 \xrightarrow{\mathbf{f}} q_f \cdots q_f$. Let ρ_f extending ρ.

Almost surely, when $|\rho_f| \to \infty$, ρ_f has more b's than a's.

Illustration

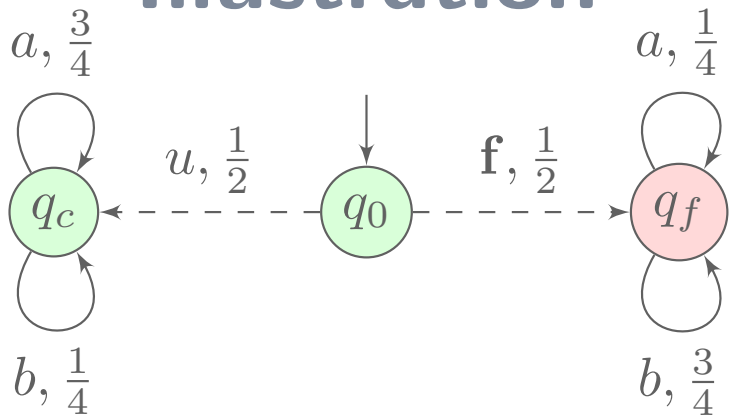

- **AA-diagnosable**

Let $\rho = q_0 \xrightarrow{\mathbf{f}} q_f \cdots q_f$. Let ρ_f extending ρ.

Almost surely, when $|\rho_f| \to \infty$, ρ_f has more b's than a's.

Let ρ_c the correct run with $\mathcal{P}(\rho_f) = \mathcal{P}(\rho_c)$.

Almost surely, when $|\rho_c| \to \infty$, ρ_c has less b's than a's.

Illustration

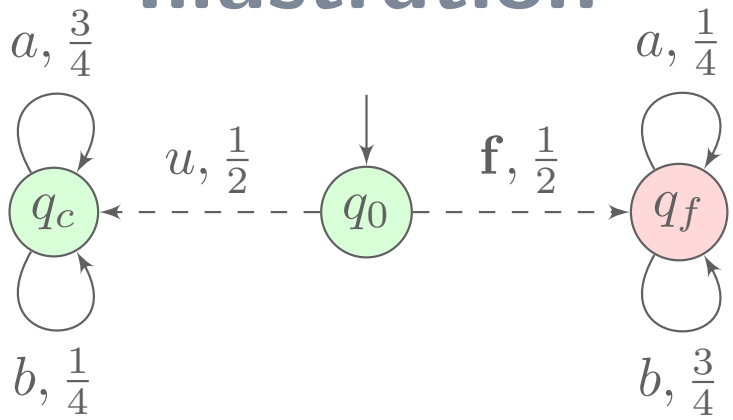

- **AA-diagnosable**

Let $\rho = q_0 \xrightarrow{\mathbf{f}} q_f \cdots q_f$. Let ρ_f extending ρ.

Almost surely, when $|\rho_f| \to \infty$, ρ_f has more b's than a's.

Let ρ_c the correct run with $\mathcal{P}(\rho_f) = \mathcal{P}(\rho_c)$.

Almost surely, when $|\rho_c| \to \infty$, ρ_c has less b's than a's.

- **Not uniformly AA-diagnosable**

Let $\varepsilon = \frac{1}{2}$, $\alpha > 0$ and ρ be the faulty run with $\mathcal{P}(\rho) = a^n$.

Illustration

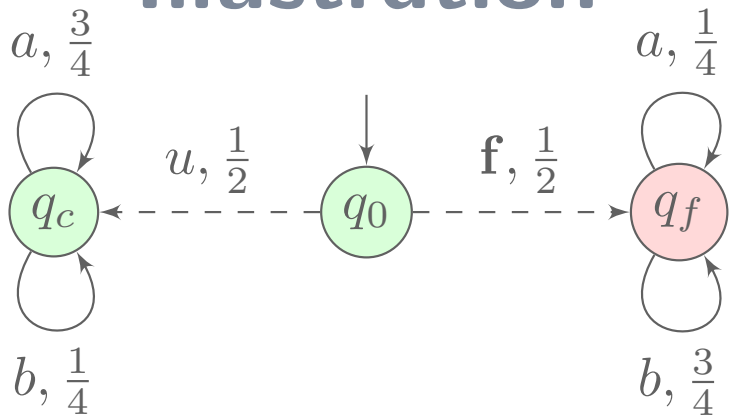

- **AA-diagnosable**

Let $\rho = q_0 \xrightarrow{\mathbf{f}} q_f \cdots q_f$. Let ρ_f extending ρ.

Almost surely, when $|\rho_f| \to \infty$, ρ_f has more b's than a's.

Let ρ_c the correct run with $\mathcal{P}(\rho_f) = \mathcal{P}(\rho_c)$.

Almost surely, when $|\rho_c| \to \infty$, ρ_c has less b's than a's.

- **Not uniformly AA-diagnosable**

Let $\varepsilon = \frac{1}{2}$, $\alpha > 0$ and ρ be the faulty run with $\mathcal{P}(\rho) = a^n$.

Then for all ρ_f extending ρ with $|\rho_f| \leq n + |\rho|$, $\mathsf{CorP}(\mathcal{P}(\rho)) \geq \frac{1}{2}$.

Illustration

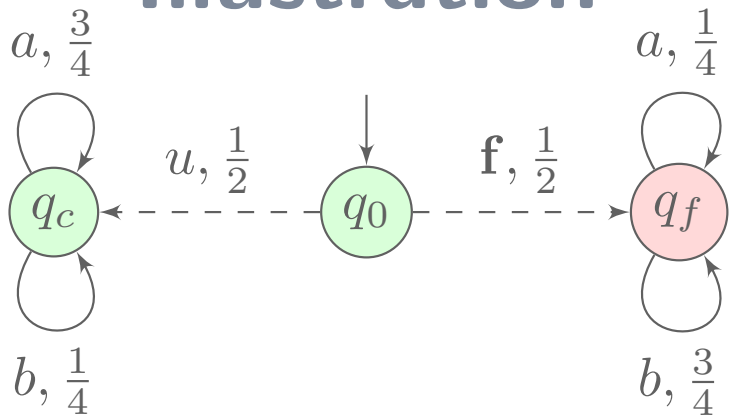

- **AA-diagnosable**

Let $\rho = q_0 \xrightarrow{\mathbf{f}} q_f \cdots q_f$. Let ρ_f extending ρ.

Almost surely, when $|\rho_f| \to \infty$, ρ_f has more b's than a's.

Let ρ_c the correct run with $\mathcal{P}(\rho_f) = \mathcal{P}(\rho_c)$.

Almost surely, when $|\rho_c| \to \infty$, ρ_c has less b's than a's.

- **Not uniformly AA-diagnosable**

Let $\varepsilon = \frac{1}{2}$, $\alpha > 0$ and ρ be the faulty run with $\mathcal{P}(\rho) = a^n$.

Then for all ρ_f extending ρ with $|\rho_f| \leq n + |\rho|$, $\mathsf{CorP}(\mathcal{P}(\rho)) \geq \frac{1}{2}$.

$\mathbb{P}(\rho_f | \rho \preceq \rho_f \wedge |\rho_f| = k + |\rho| \wedge \mathsf{CorP}(\mathcal{P}(\rho_f)) \leq \varepsilon) \leq \alpha \mathbb{P}(\rho)$ implies $k \geq n$.

Relations between the specifications

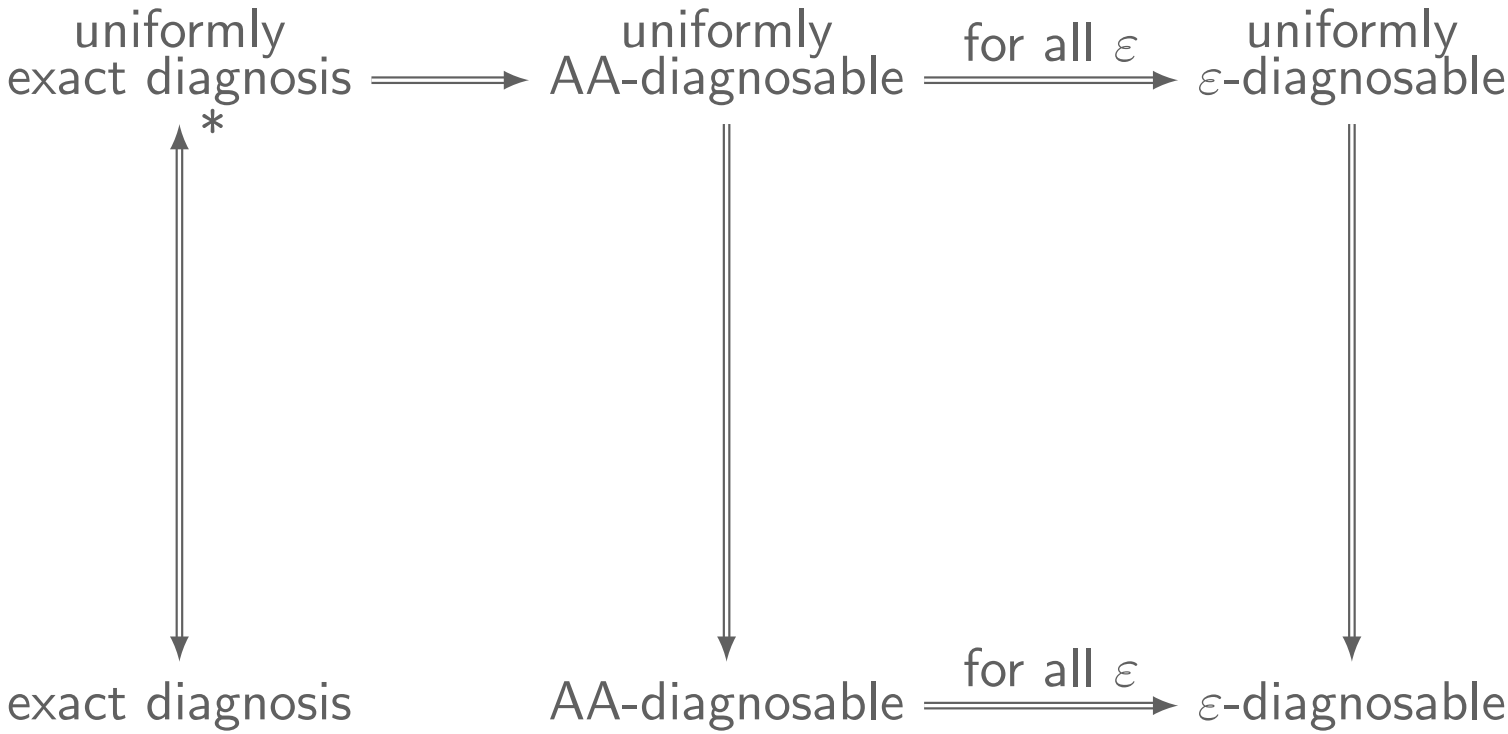

* assuming finite models

Outline

Outline

Characterisation of diagnosability

Specification of IF-diagnosability: Infinite sequences, Fault diagnosis

$\mathcal{A}$

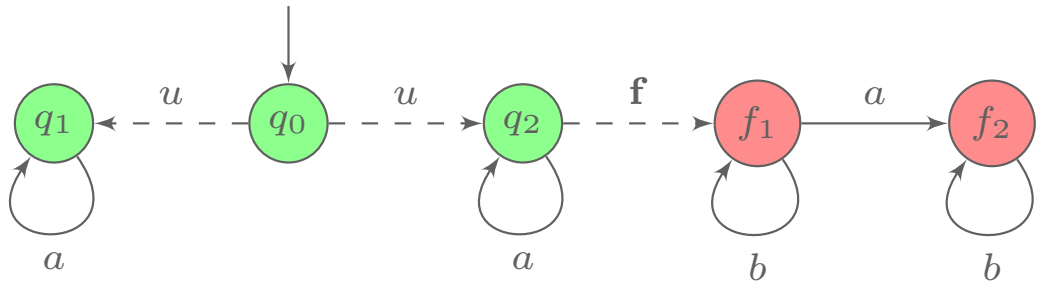

Characterisation of diagnosability

Specification of IF-diagnosability: Infinite sequences, Fault diagnosis

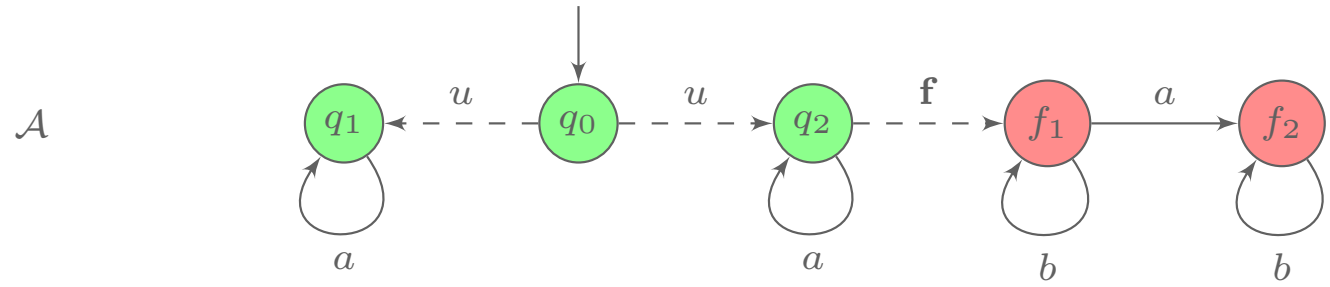

Observer: tracks possible correct states after given observed sequence.

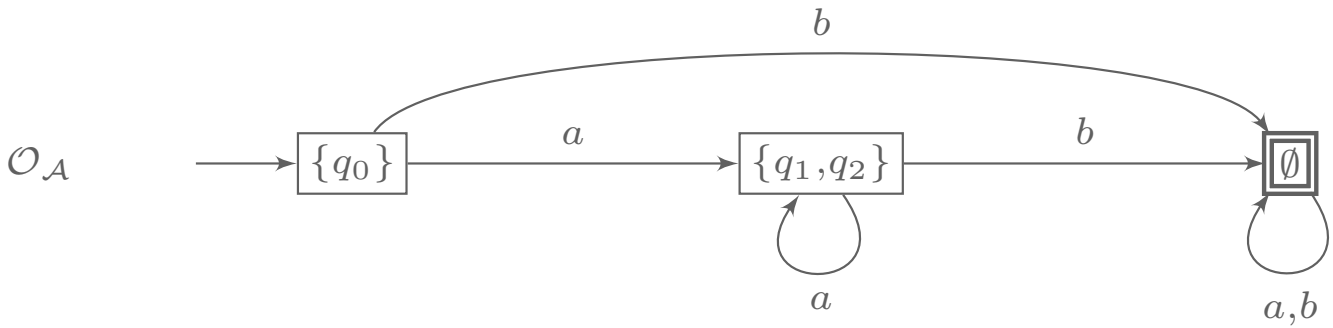

Characterisation of diagnosability

Specification of IF-diagnosability: Infinite sequences, Fault diagnosis

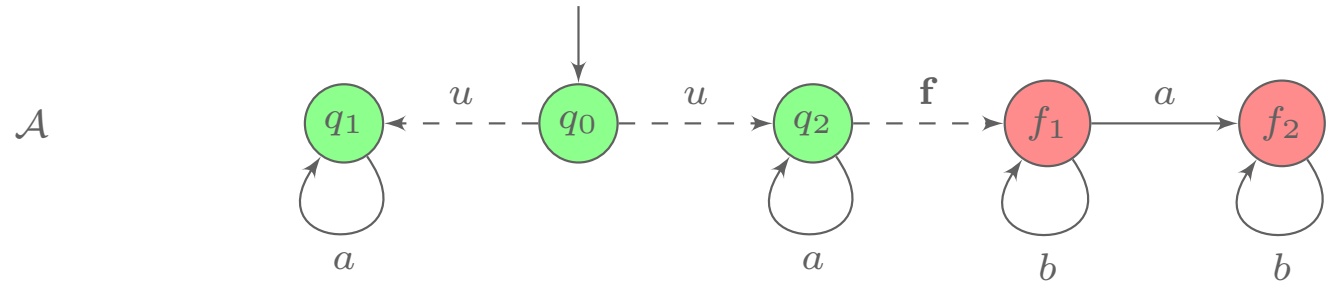

Observer: tracks possible correct states after given observed sequence.

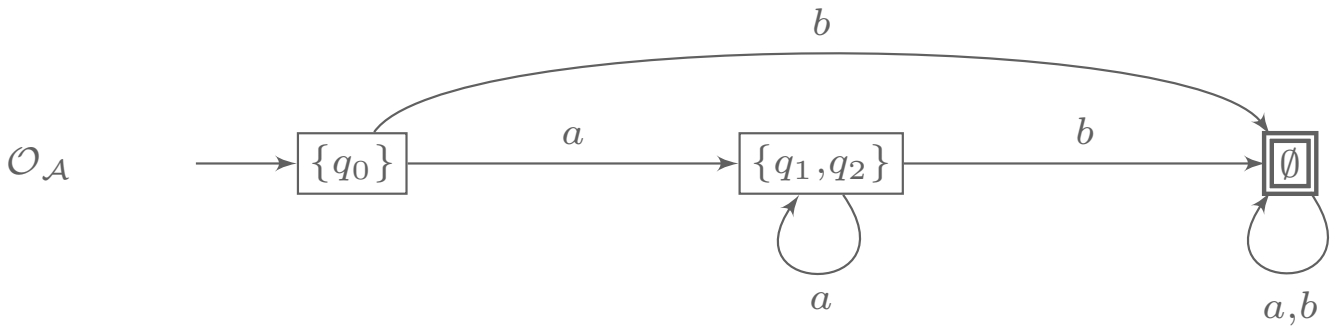

$\mathcal{A}$ is not IF-diagnosable
iff
there exists a state (q, U) in a BSCC of $\mathcal{A} \times \mathcal{O}_\mathcal{A}$ with q faulty and $U \neq \emptyset$.

Diagnoser synthesis

For every IF-diagnosable system with n correct states
one can build an IF-diagnoser with at most 2^n states.

Diagnoser derived from observer $\mathcal{O}_{\mathcal{A}}$: emits ✗ in state $\emptyset$.

Diagnoser synthesis

For every IF-diagnosable system with n correct states
one can build an IF-diagnoser with at most 2^n states.

Diagnoser derived from observer $\mathcal{O}_{\mathcal{A}}$: emits ✗ in state $\emptyset$.

There is a family $(\mathcal{A}_n)$ of IF-diagnosable systems such that
$\mathcal{A}_n$ has $n+1$ correct states and any IF-diagnoser needs 2^n states.

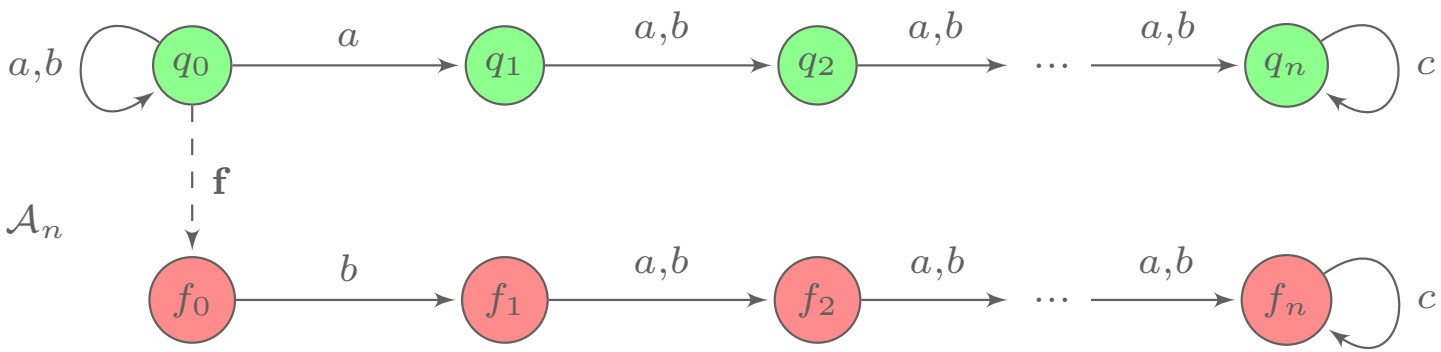

Diagnosability is in PSPACE

Diagnosability is decidable in PSPACE for probabilistic systems.

Diagnosability is in PSPACE

Diagnosability is decidable in PSPACE for probabilistic systems.

Sketch of proof

- relies on the characterisation on $\mathcal{A} \times \mathcal{O}_{\mathcal{A}}$
- avoids building the product
- uses Savitch's theorem for appropriate guesses

Diagnosability is PSPACE-hard

$\mathcal{L} \subseteq \Sigma^*$ is *eventually universal* if $\exists v \in \Sigma^*, v^{-1}\mathcal{L} = \Sigma^*$.

The eventual universality problem for NFA is PSPACE-hard.

Diagnosability is PSPACE-hard

$\mathcal{L} \subseteq \Sigma^*$ is *eventually universal* if $\exists v \in \Sigma^*, v^{-1}\mathcal{L} = \Sigma^*$.

The eventual universality problem for NFA is PSPACE-hard.

Diagnosability is PSPACE-hard.

Reduction from eventual universality to diagnosability.

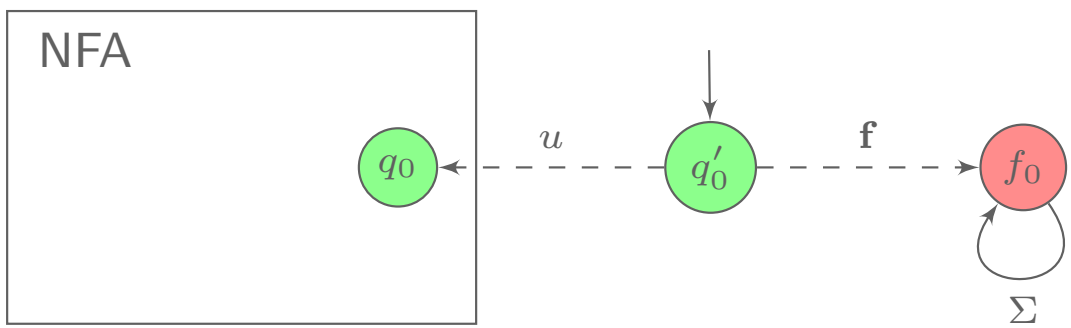

$\mathcal{A}$ not diagnosable iff
$\mathcal{A} \times \mathcal{O}_{\mathcal{A}}$ contains a BSCC where each state has the form (f_0, U) with $U \neq \emptyset$

Comparison with non-probabilistic discrete event systems

Diagnosability is PSPACE-complete for probabilistic systems.

Comparison with non-probabilistic discrete event systems

Diagnosability is PSPACE-complete for probabilistic systems.

Diagnosability is decidable in PTIME for non-probabilistic systems.
[Jiang, Huang, Chandra, Kumar TAC 2001]

Sketch of proof

- build the twin-product with a copy restricted to correct states
- check for SCC with faulty states in the first component

Comparison with non-probabilistic discrete event systems

Diagnosability is PSPACE-complete for probabilistic systems.

Diagnosability is decidable in PTIME for non-probabilistic systems.
[Jiang, Huang, Chandra, Kumar TAC 2001]

Sketch of proof

- build the twin-product with a copy restricted to correct states
- check for SCC with faulty states in the first component

Erroneous adaptation to probabilistic case in [Chen, Kumar TASE 2013].

Revisiting the relations

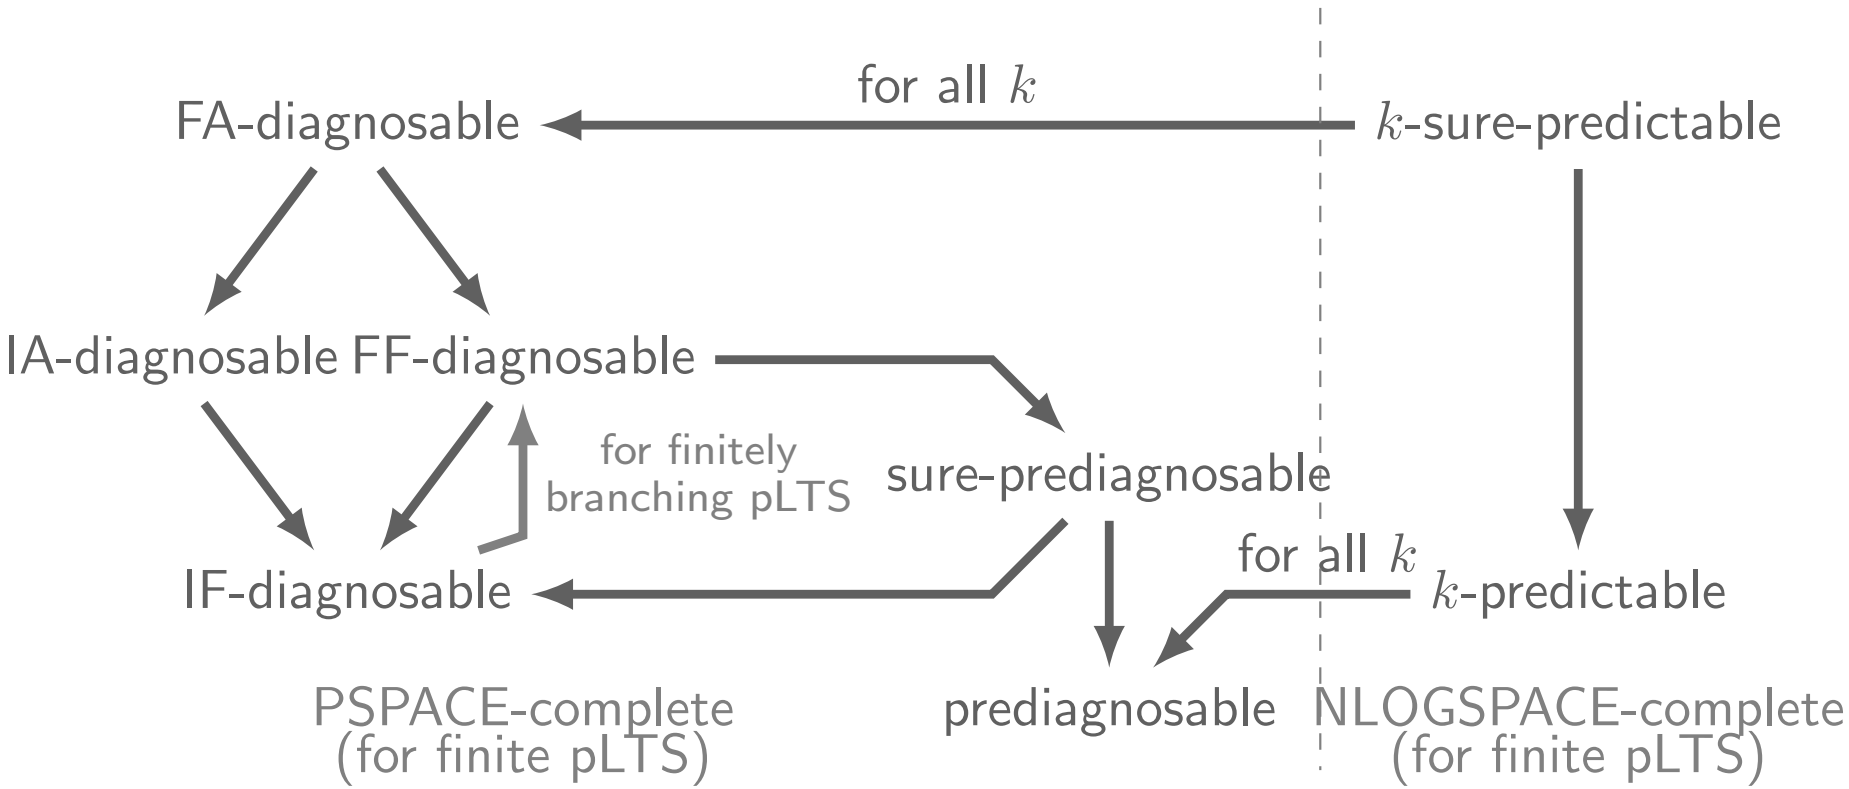

Outline

Complexity of the problems

	Simple	Uniform
ε-diagnosability	undecidable	undecidable
AA-diagnosability	PTIME	undecidable

AA-diagnosability: a simple case

Initial fault pLTS. Initially, an unobservable split towards two subpLTS:

- a *correct* event u leads to a *correct* subpLTS;
- a *faulty* event $\mathbf{f}$ leads to an *arbitrary* subpLTS.

AA-diagnosability: a simple case

Initial fault pLTS. Initially, an unobservable split towards two subpLTS:

- a *correct* event u leads to a *correct* subpLTS;
- a *faulty* event $\mathbf{f}$ leads to an *arbitrary* subpLTS.

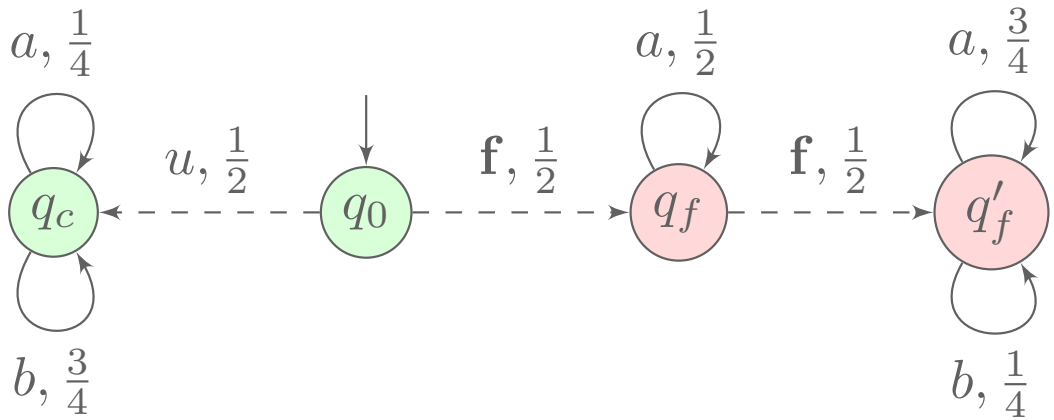

- an initial state, q_0;
- an arbitrary pLTS with states $\{q_f, q'_f\}$;
- a correct pLTS with state q_c.

AA-diagnosability for initial-fault pLTS

• Transform the correct and arbitrary subpLTS in *labelled Markov chains* by merging the unobservable transitions.

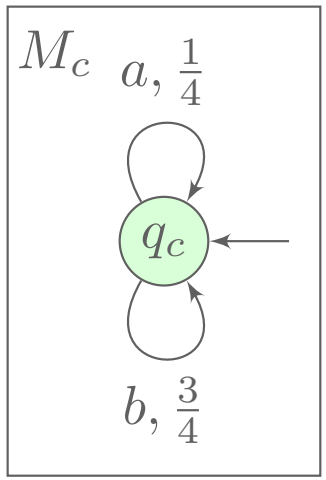

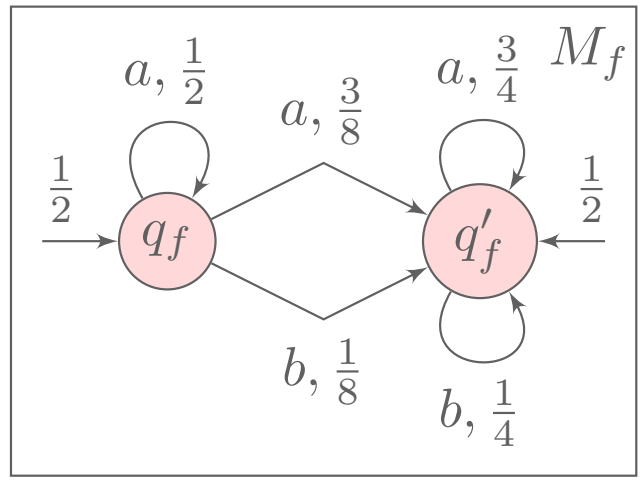

AA-diagnosability for initial-fault pLTS

• Transform the correct and arbitrary subpLTS in *labelled Markov chains* by merging the unobservable transitions.

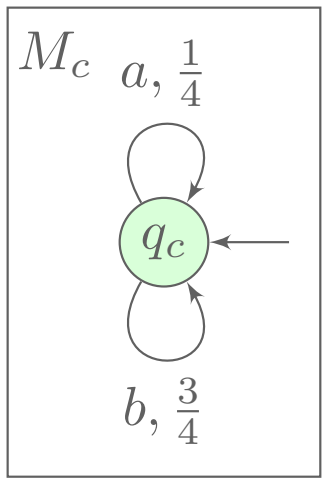

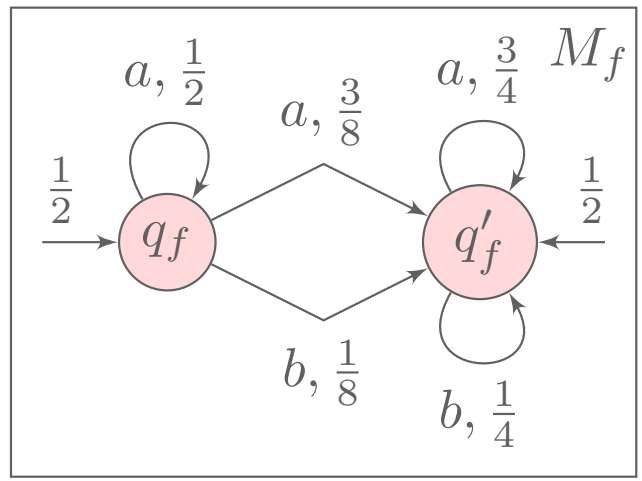

• $\mathbb{P}^M(E) =$ measure of infinite runs of M with observation in E.

Distance 1 problem: $\exists$ (measurable) $E \subseteq \Sigma_o^\omega, \mathbb{P}^{M_c}(E) - \mathbb{P}^{M_f}(E) = 1$?

• Illustration: $E = \{\sigma \mid \limsup_{n\to\infty} \frac{|\sigma_{\downarrow n}|_b}{|\sigma_{\downarrow n}|_a} > 1\}$

AA-diagnosability for initial-fault pLTS

• Transform the correct and arbitrary subpLTS in *labelled Markov chains* by merging the unobservable transitions.

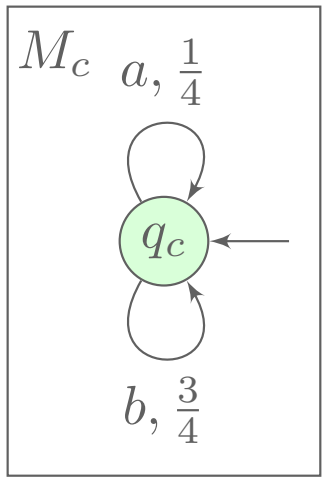

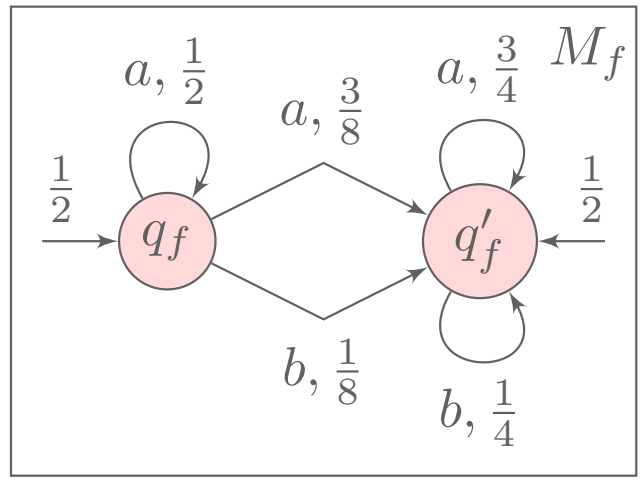

• $\mathbb{P}^M(E)$ = measure of infinite runs of M with observation in E.

Distance 1 problem: $\exists$ (measurable) $E \subseteq \Sigma_o^\omega, \mathbb{P}^{M_c}(E) - \mathbb{P}^{M_f}(E) = 1$?

• Illustration: $E = \{\sigma \mid \limsup_{n\to\infty} \frac{|\sigma_{\downarrow n}|_b}{|\sigma_{\downarrow n}|_a} > 1\}$

The distance 1 problem is decidable in PTIME.

[CK14] Chen and Kiefer

On the Total Variation Distance of Labelled Markov Chains, CSL-LICS'14.

Solving AA-diagnosability

• Identifying relevant pairs of states by reachability analysis in the synchronised self-product.

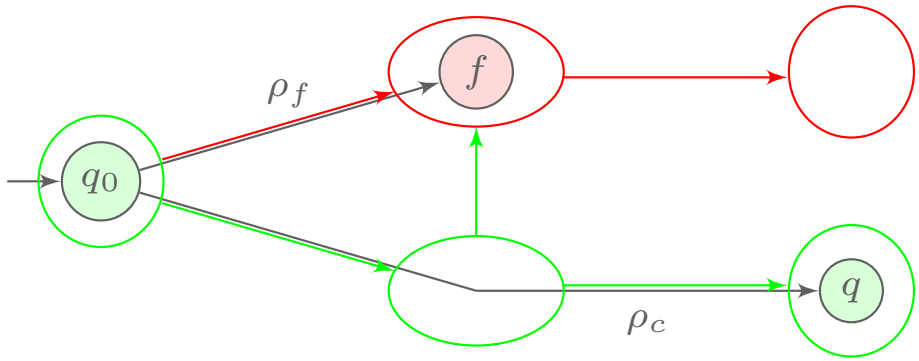

$\mathcal{P}(\rho_c)=\mathcal{P}(\rho_f)$

Solving AA-diagnosability

• Identifying relevant pairs of states by reachability analysis in the synchronised self-product.

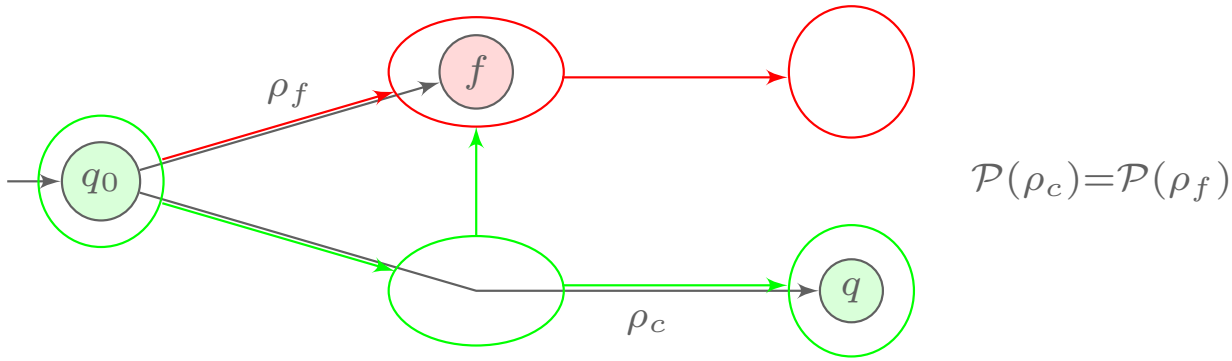

• Checking distance 1 for all relevant pairs.

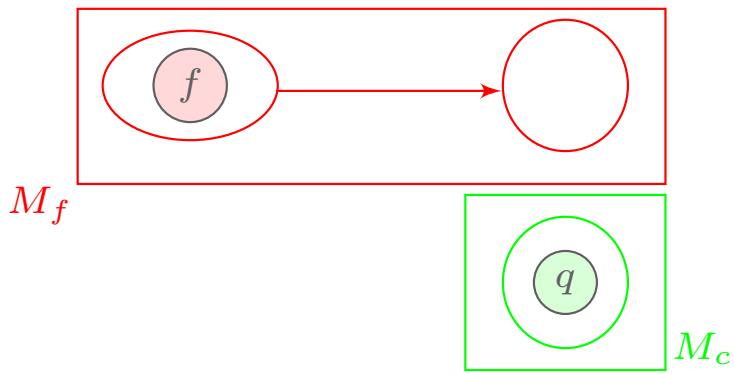

Solving AA-diagnosability

• Identifying relevant pairs of states by reachability analysis in the synchronised self-product.

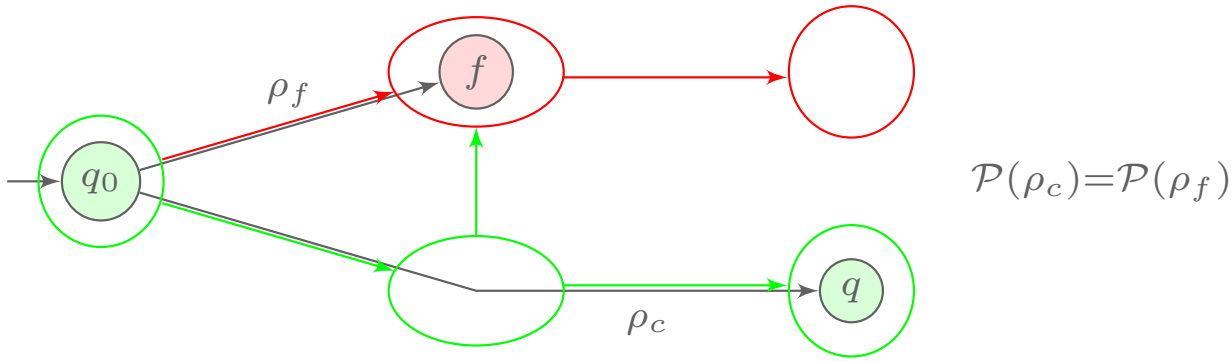

• Checking distance 1 for all relevant pairs.

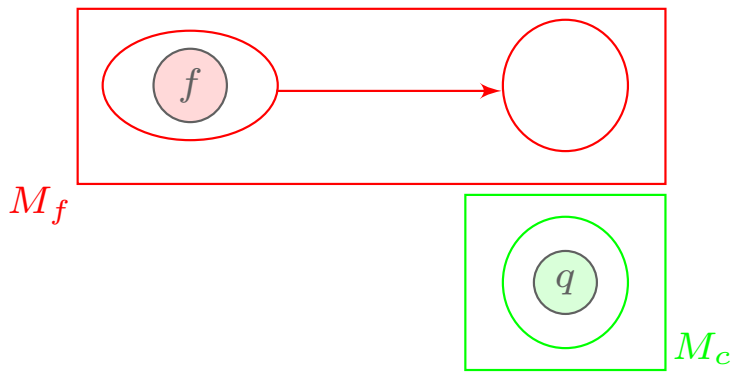

AA-diagnosability is decidable in PTIME.

Solving AA-diagnosability

• Identifying relevant pairs of states by reachability analysis in the synchronised self-product.

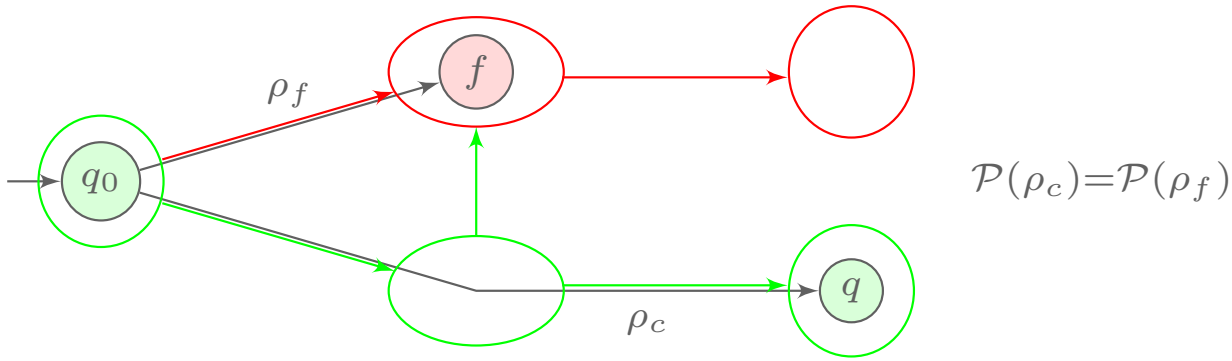

• Checking distance 1 for all relevant pairs.

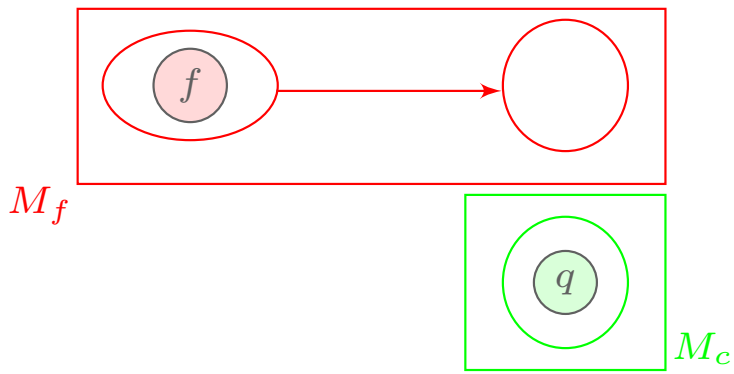

AA-diagnosability is decidable in PTIME.

However an ε-diagnoser may need infinite memory.

The emptiness problem for probabilistic automata (PA)

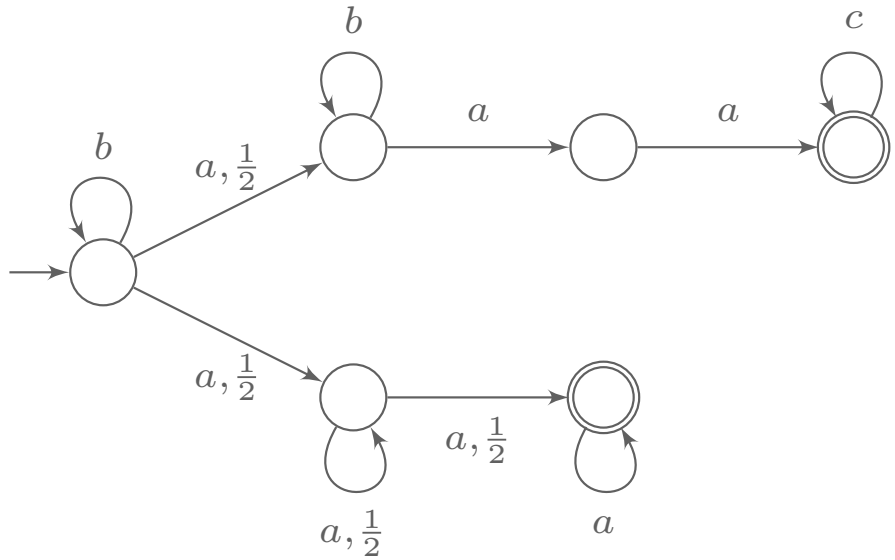

$\mathbb{P}(b) = 0, \mathbb{P}(baa) = \frac{1}{4}, \mathbb{P}(baaa) = \frac{7}{8}$

The emptiness problem for probabilistic automata (PA)

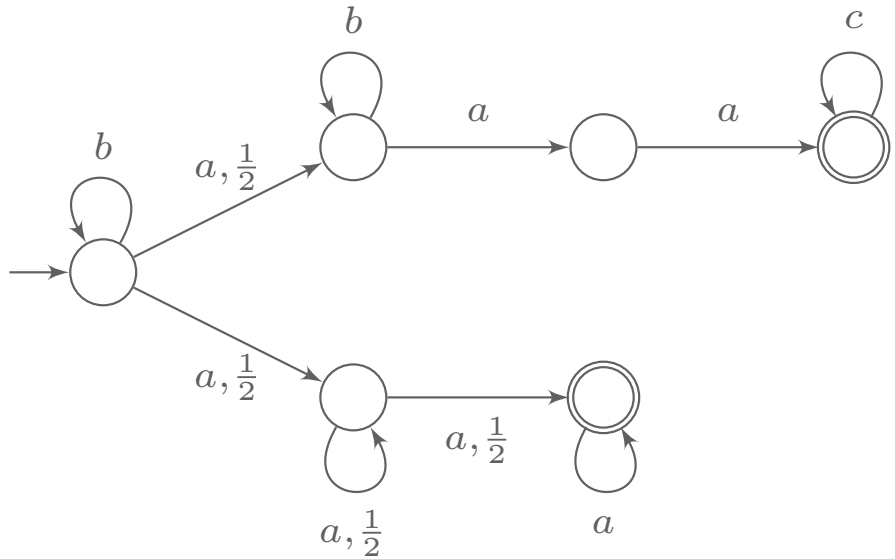

$\mathbb{P}(b) = 0, \mathbb{P}(baa) = \frac{1}{4}, \mathbb{P}(baaa) = \frac{7}{8}$

Emptiness problem: Given a PA $\mathcal{A}$,
$\exists w \in \Sigma^*, \mathbb{P}_{\mathcal{A}}(w) > \frac{1}{2}$?

The emptiness problem for probabilistic automata (PA)

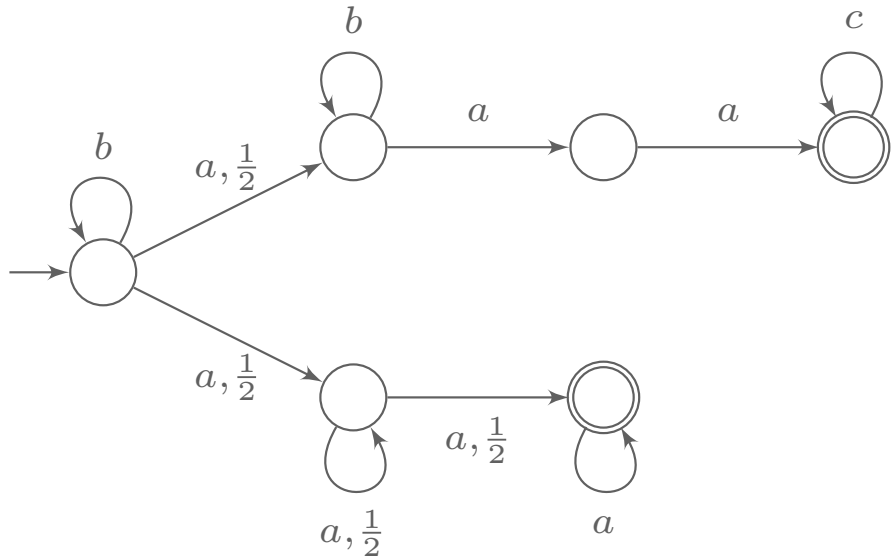

$\mathbb{P}(b) = 0, \mathbb{P}(baa) = \frac{1}{4}, \mathbb{P}(baaa) = \frac{7}{8}$

Emptiness problem: Given a PA $\mathcal{A}$,
$\exists w \in \Sigma^*, \mathbb{P}_{\mathcal{A}}(w) > \frac{1}{2}$?

The emptiness problem for PA is undecidable even when for all w, $\frac{1}{4} \leq \mathbb{P}_{\mathcal{A}}(w) \leq \frac{3}{4}$.

[P71] Paz, *Introduction to Probabilistic Automata*, Academic Press 1971.

From PA to uniform AA-diagnosability

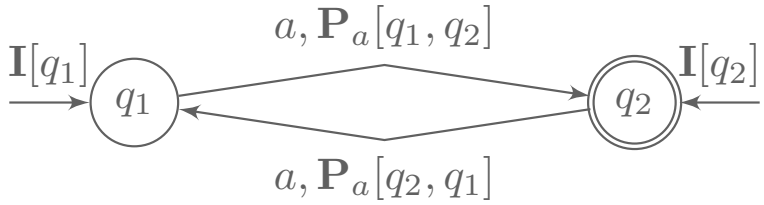

From PA to uniform AA-diagnosability

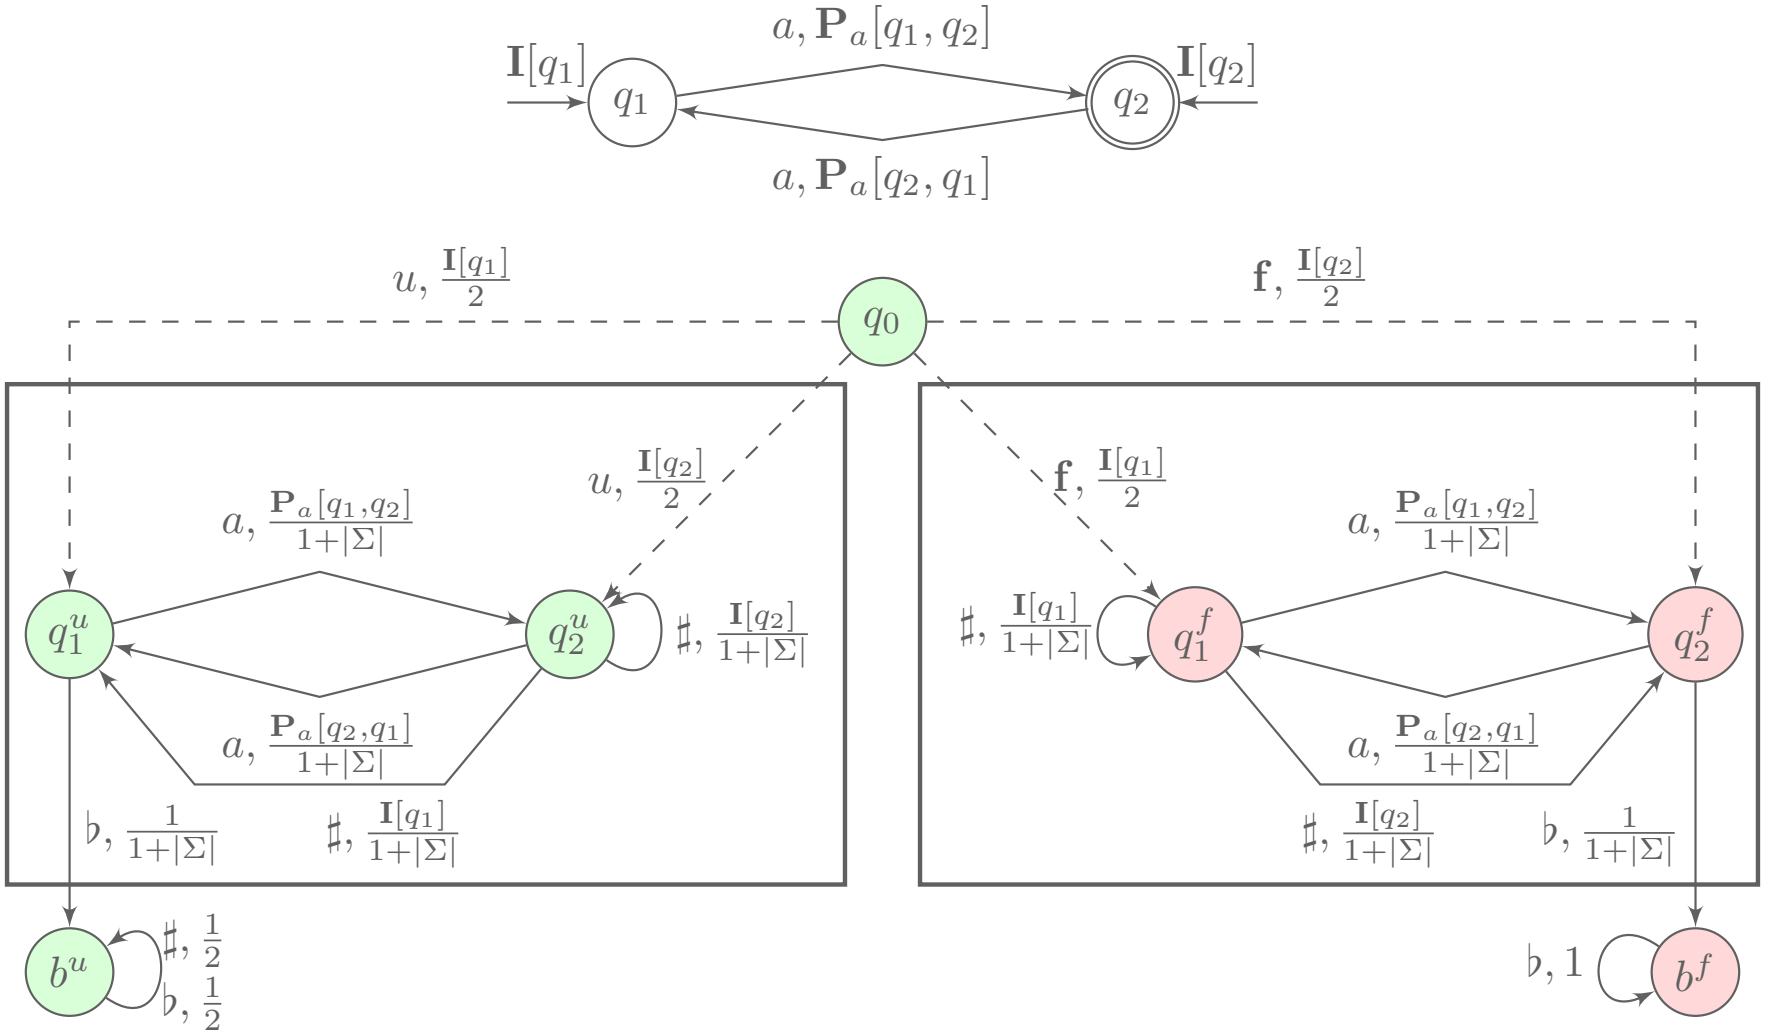

From PA to uniform AA-diagnosability

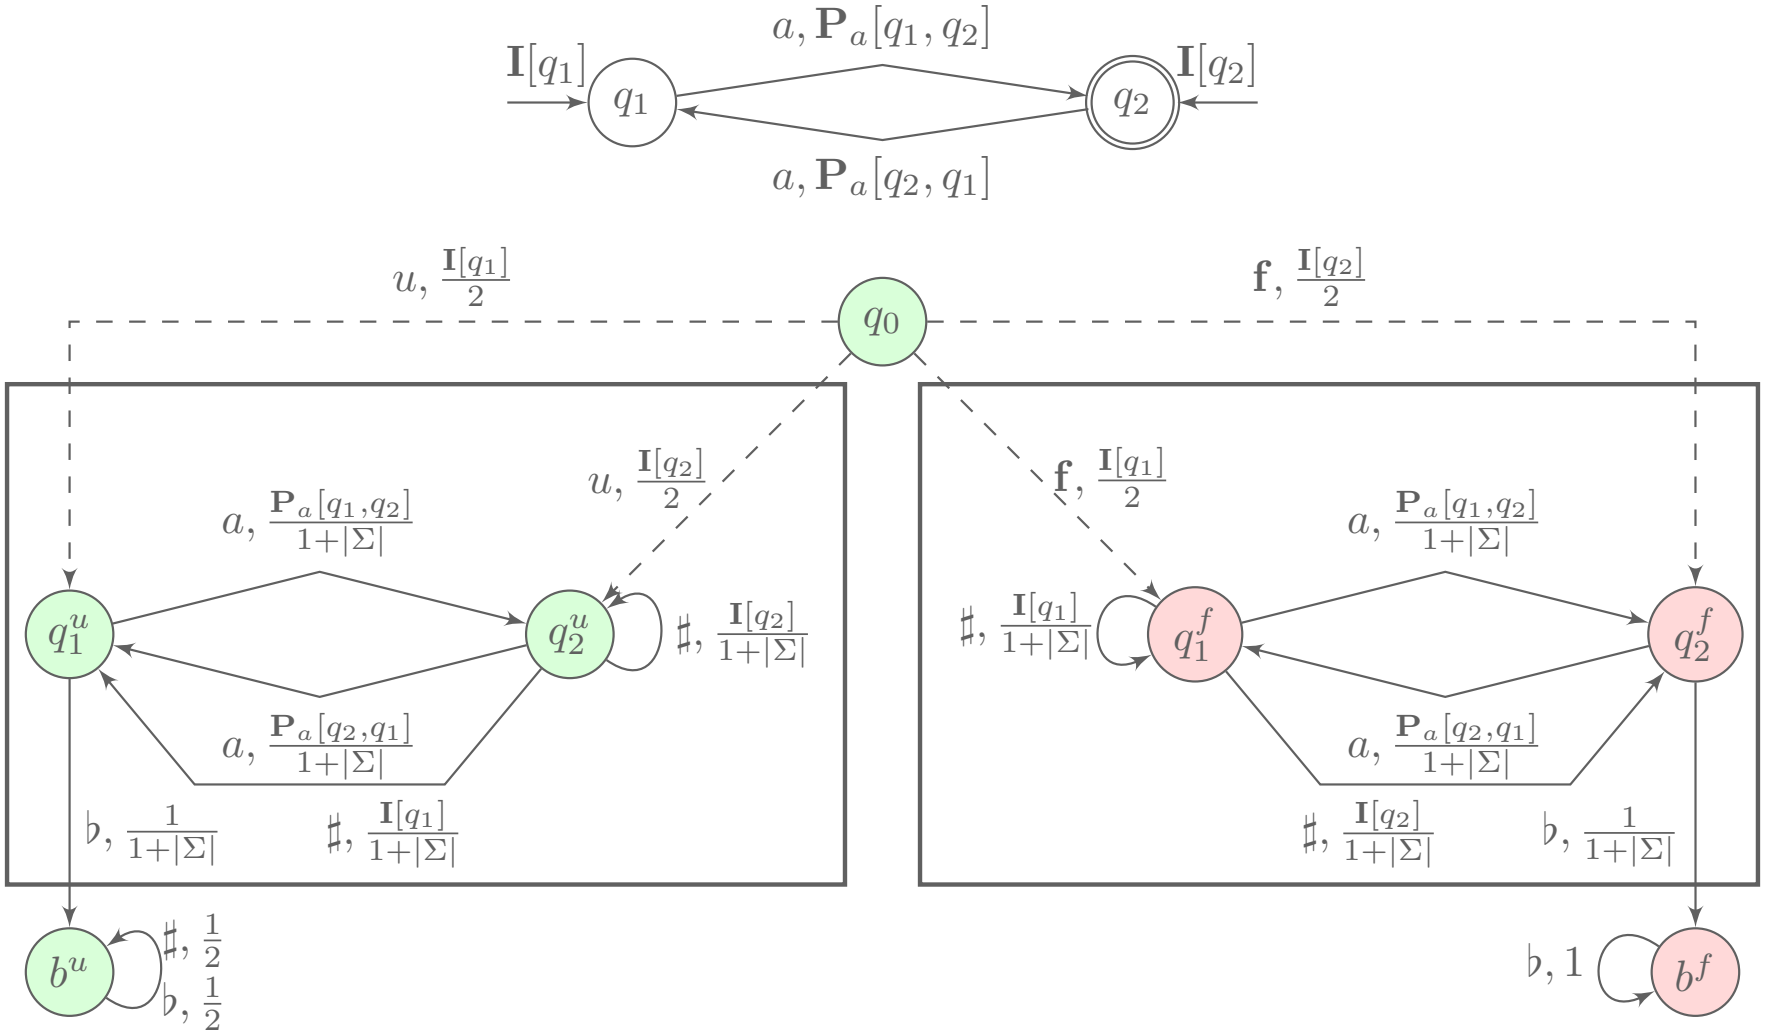

If $\exists w \in \Sigma_o^*, \mathbb{P}_{\mathcal{A}}(w) > 1/2$ then $\lim_{n\to\infty} \mathsf{CorP}((w\sharp)^n\flat) = 1$.

From PA to uniform AA-diagnosability

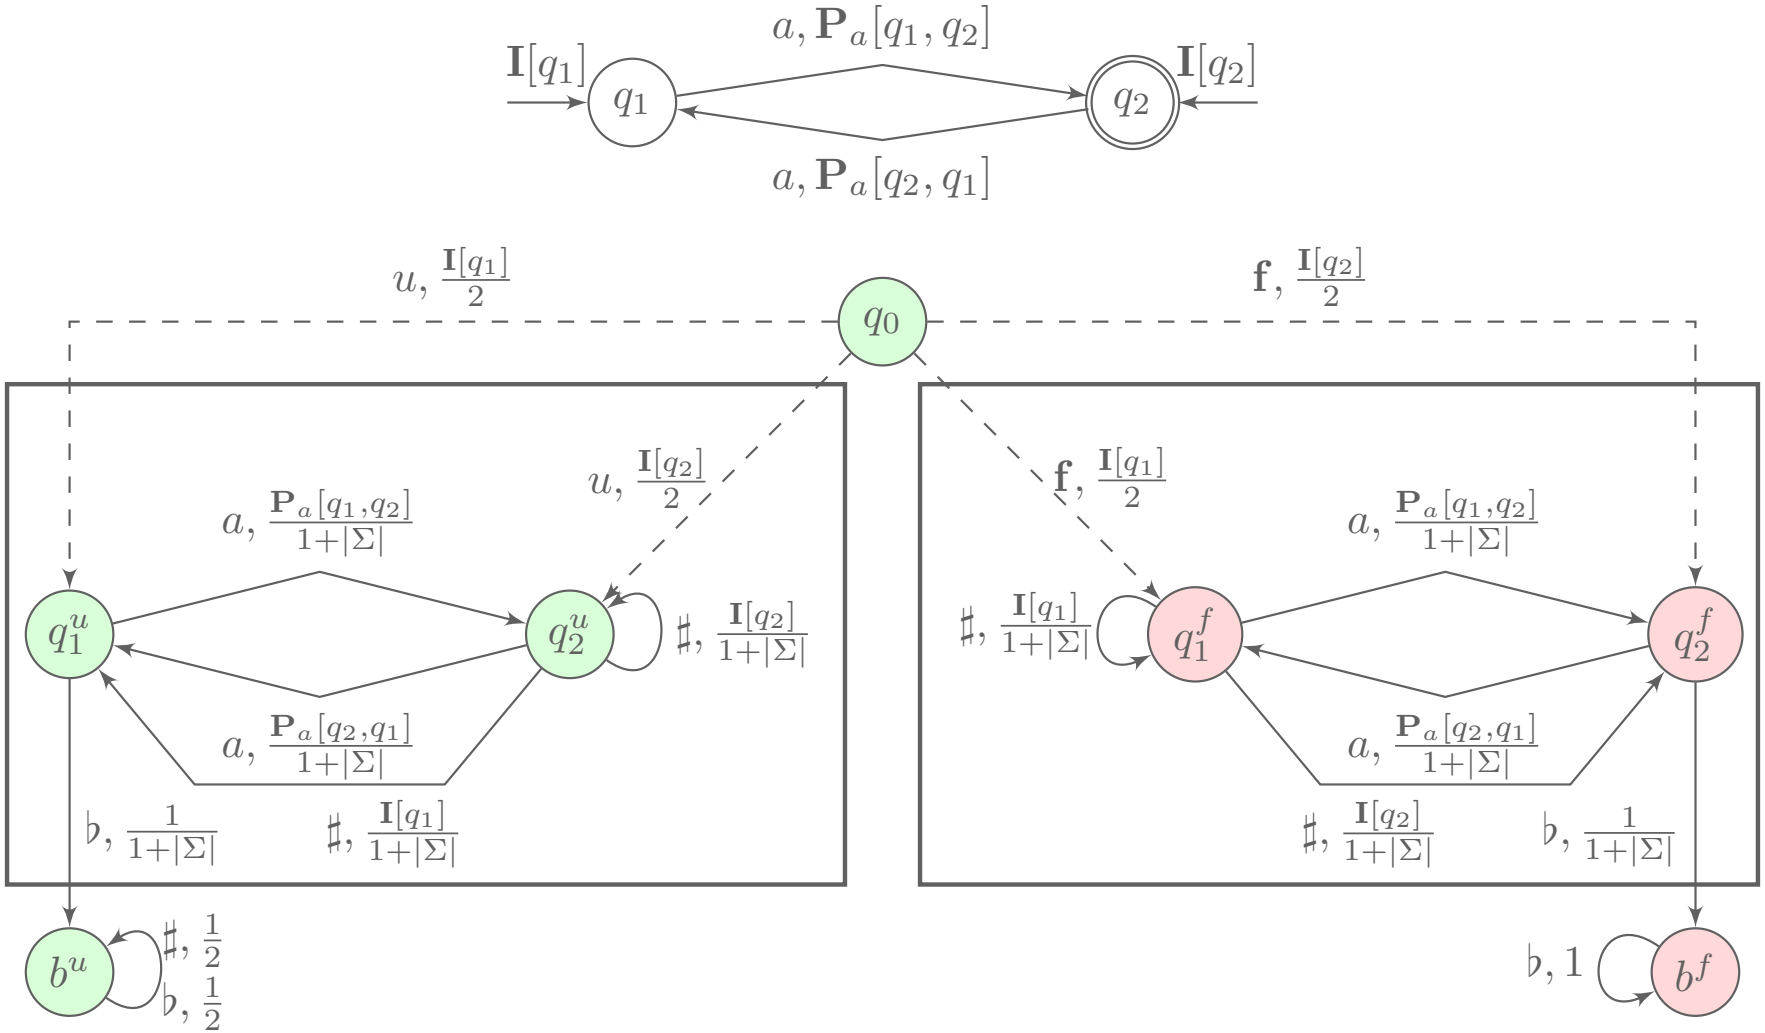

If $\exists w \in \Sigma_o^*, \mathbb{P}_{\mathcal{A}}(w) > 1/2$ then $\lim_{n\to\infty} \mathsf{CorP}((w\sharp)^n\flat) = 1$.

If $\forall w \in \Sigma_o^*, \mathbb{P}_{\mathcal{A}}(w) \leq 1/2$ then $\forall n\ \mathsf{CorP}((w\sharp)^n\flat) \leq \frac{3}{4}$.

Outline

Outline

Controllable LTS and active diagnoser

Events are also partitioned in *controllable* and *uncontrollable* events.
Controllable events must be observable.

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).

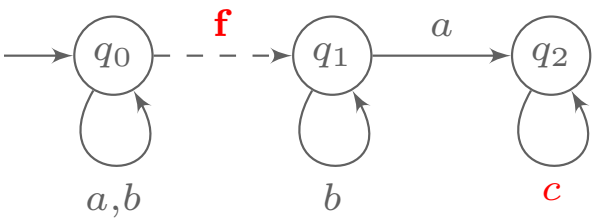

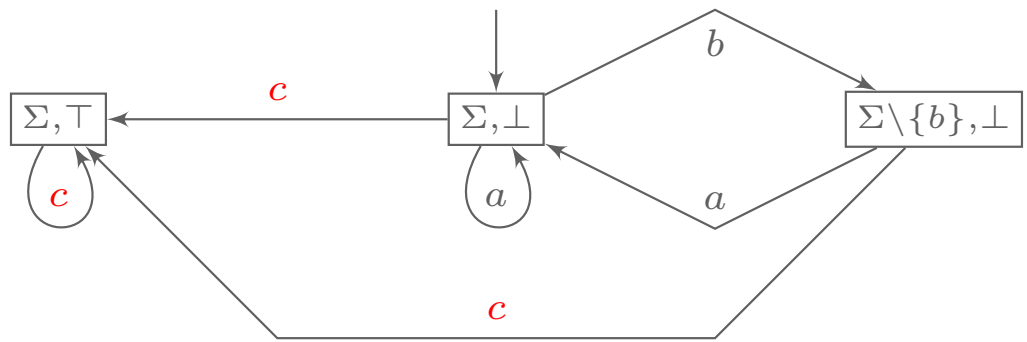

Active diagnosis problems

• 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.

Unambiguous sequences of the LTS

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

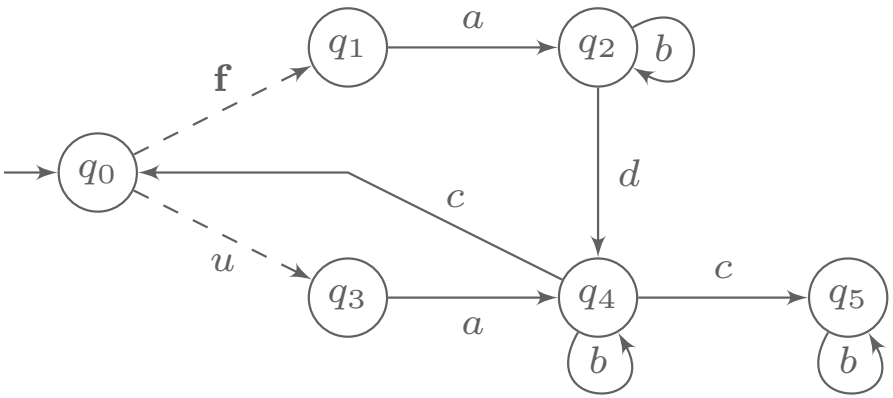

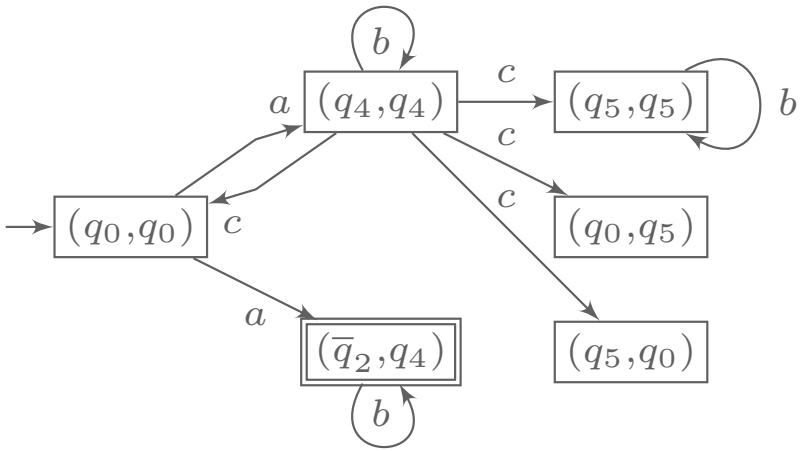

• Determinize and complement it as:

- 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.

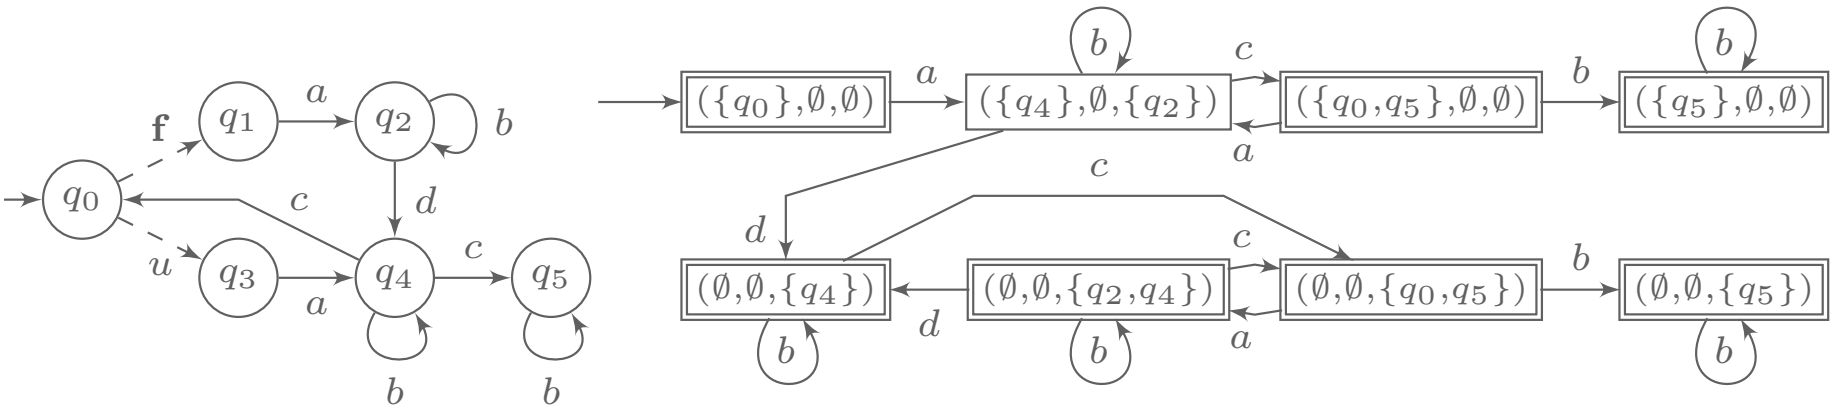

The number of states is at most 7^n.

A lower bound for ambiguity

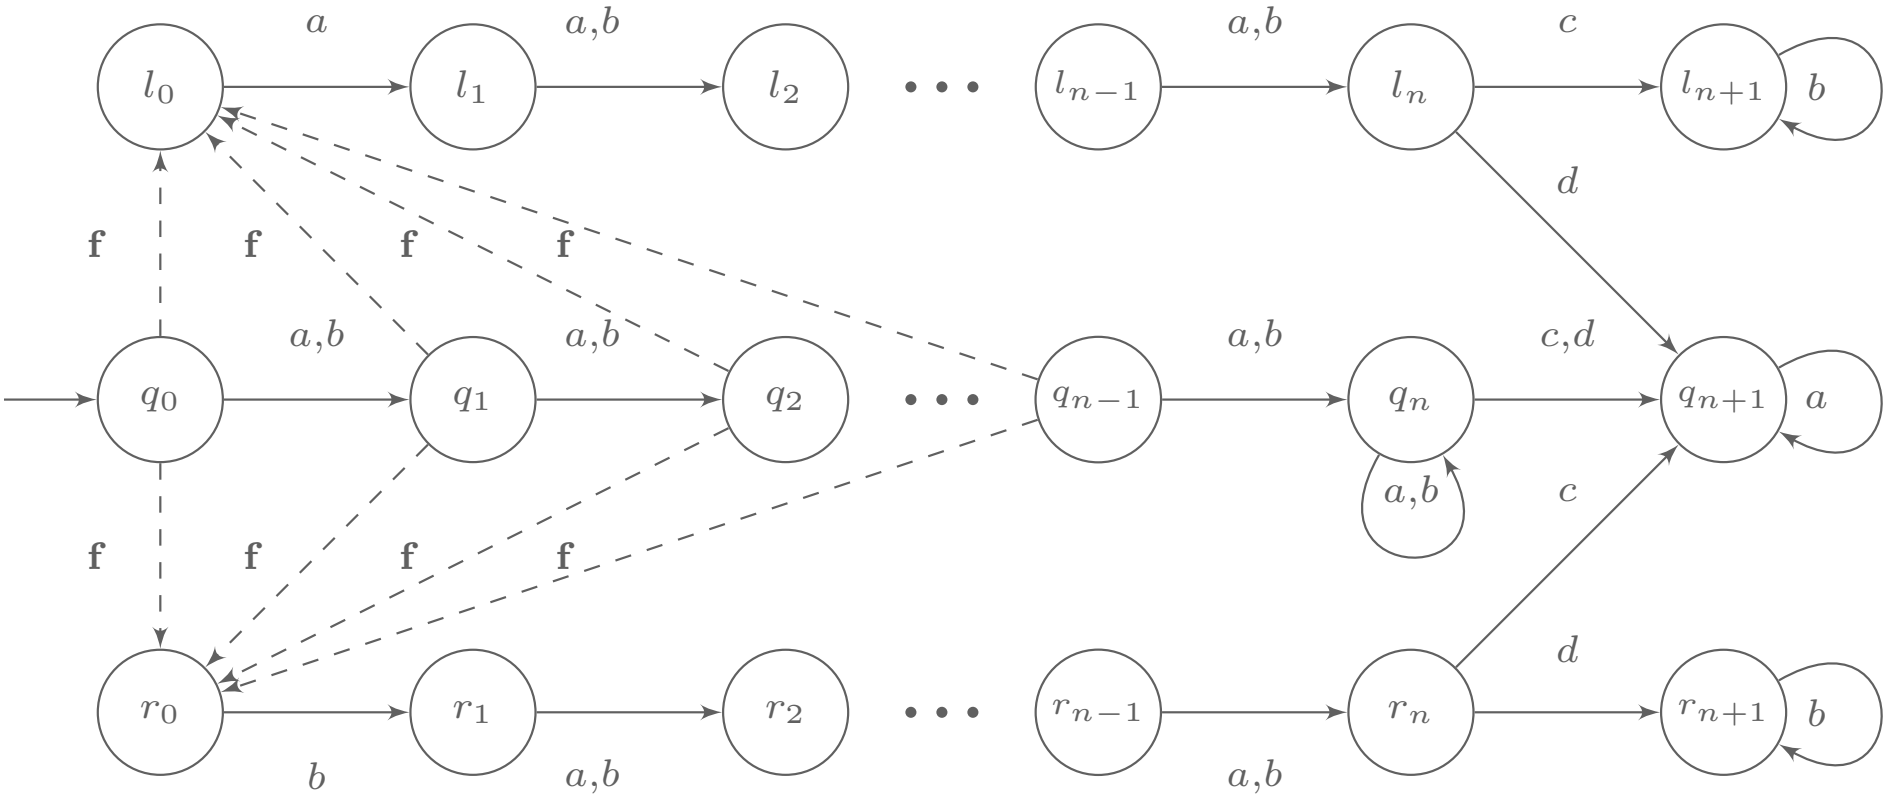

Ambiguous sequences are either $\{a,b\}^k a \{a,b\}^{n-1} d a^\omega$ or $\{a,b\}^k b \{a,b\}^{n-1} c a^\omega$ (with $0 \leq k \leq n-1$).

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

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

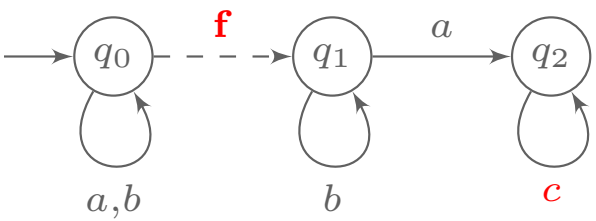

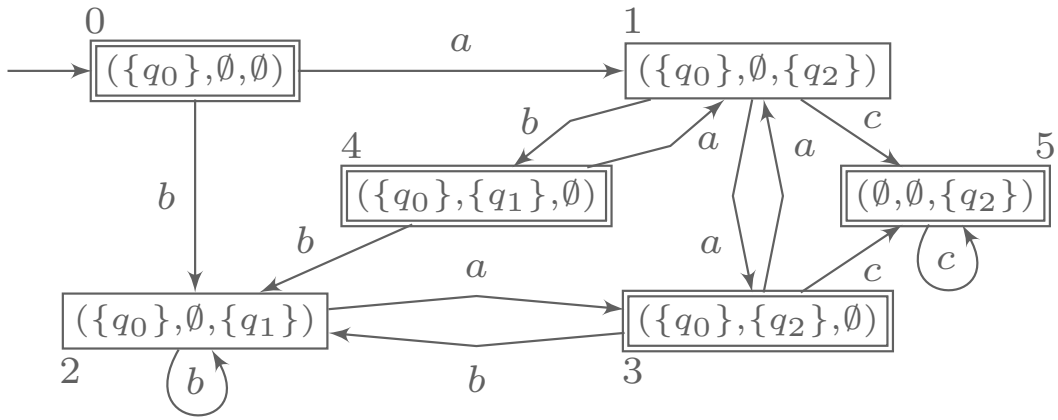

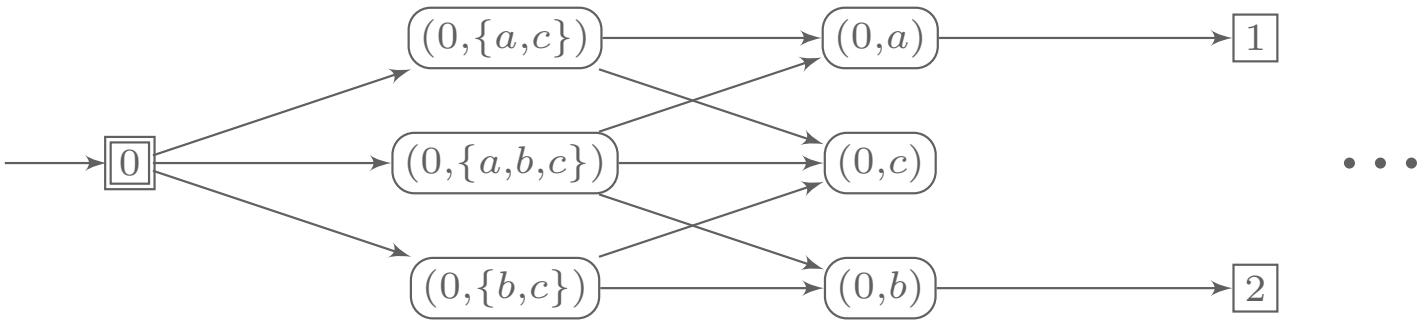

Results of this construction

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

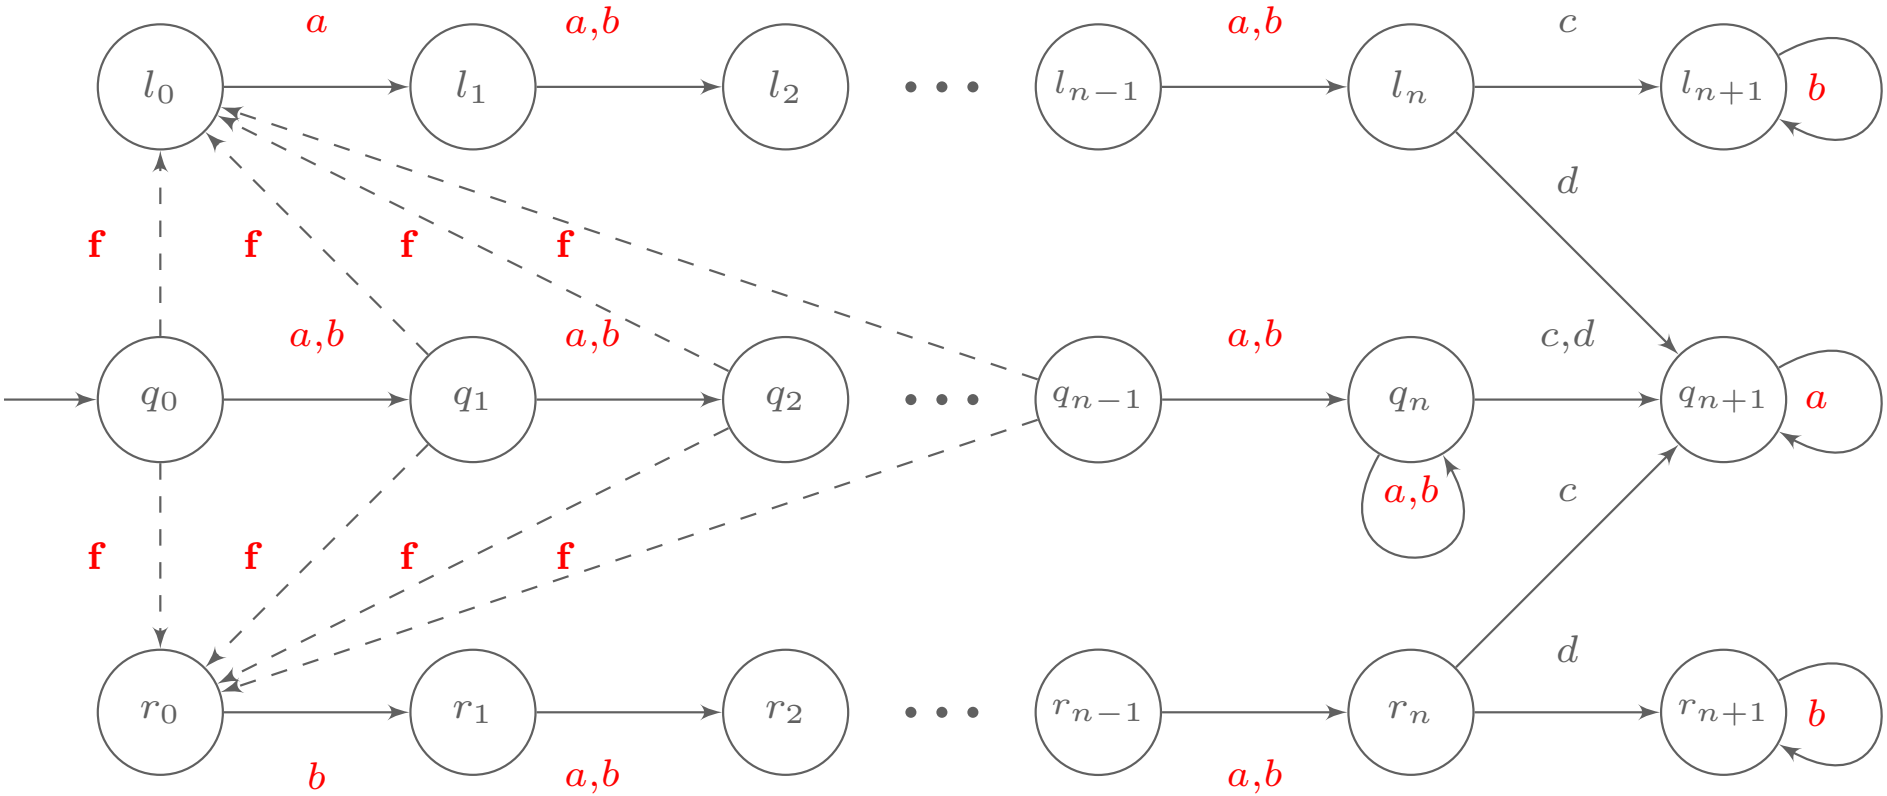

An active diagnoser must forbid a d (resp. c) if it has observed an a (resp. b) n times before.

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

Safe diagnosability

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

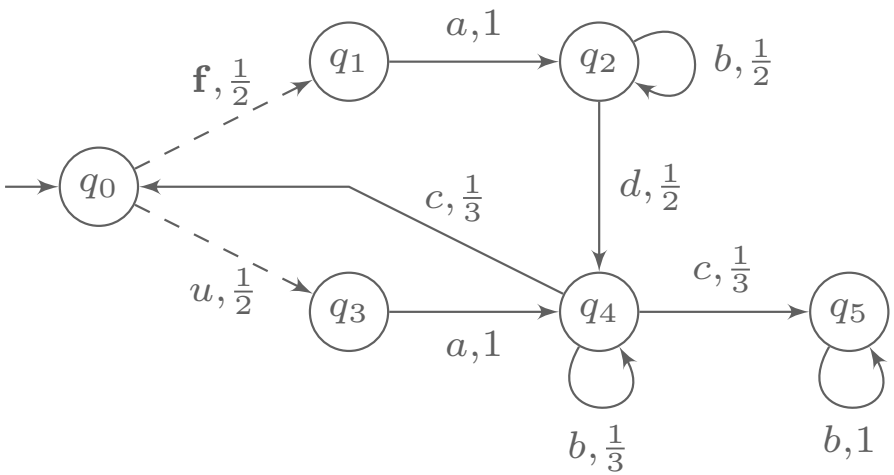

safely diagnosable

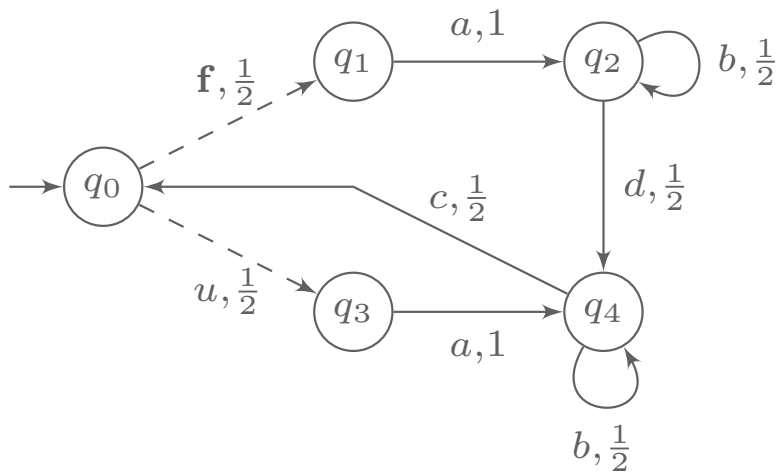

diagnosable but not safely diagnosable

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.

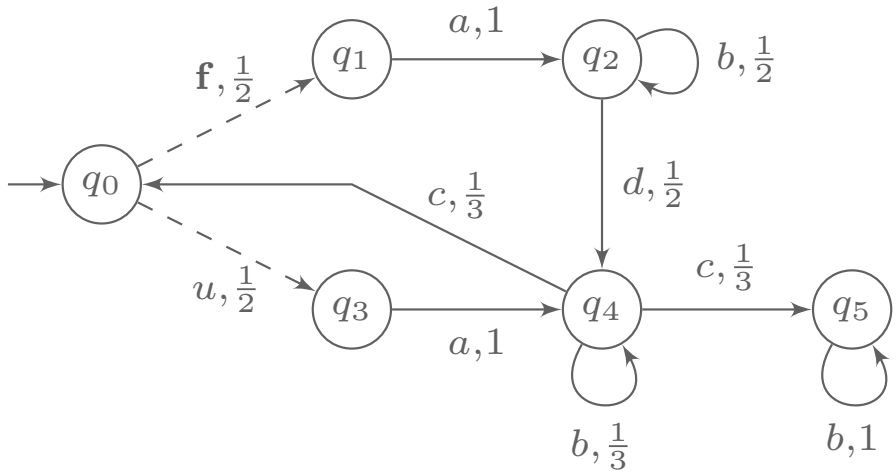

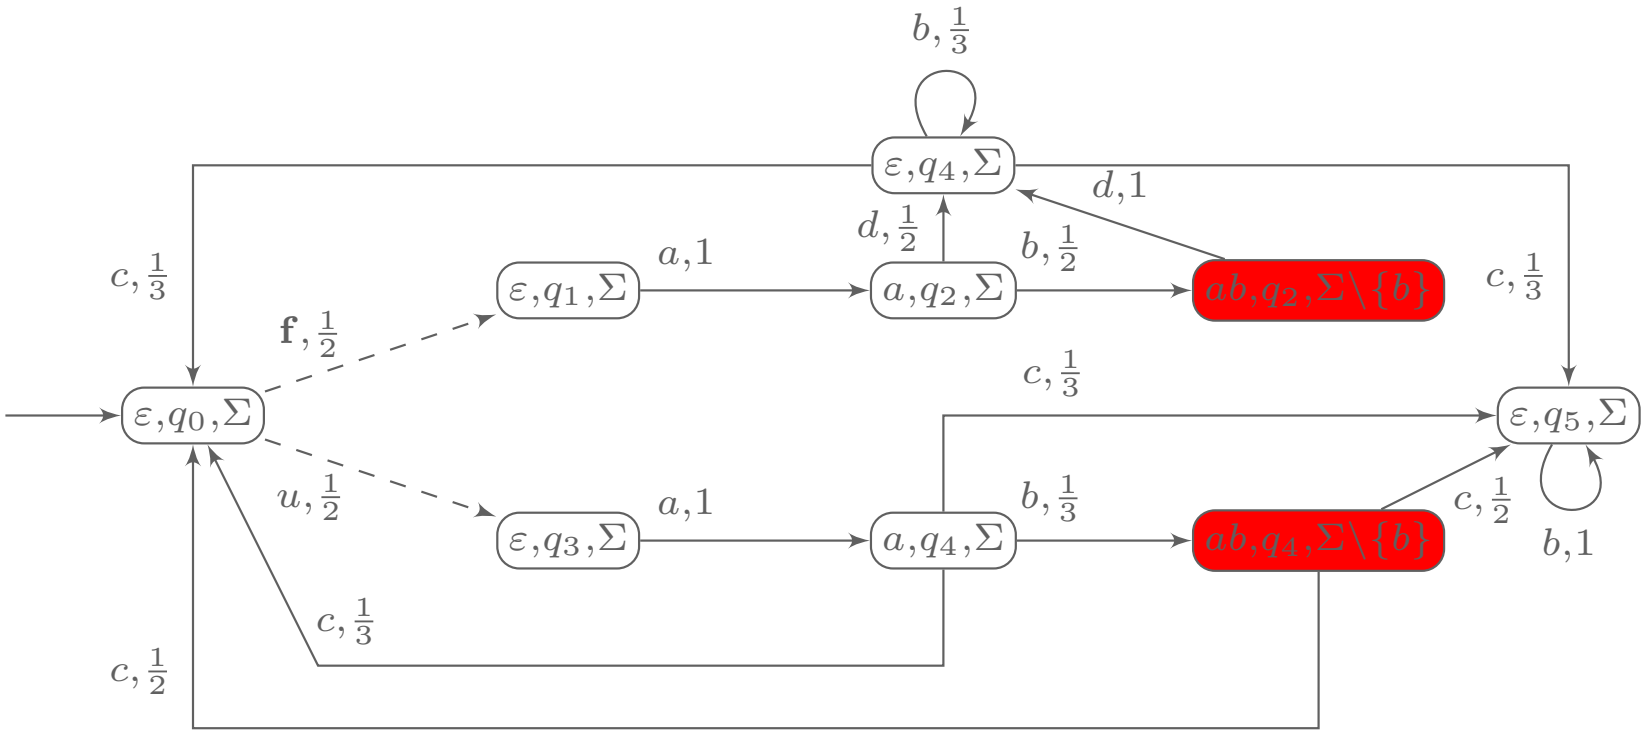

Active probabilistic diagnosis

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.

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.

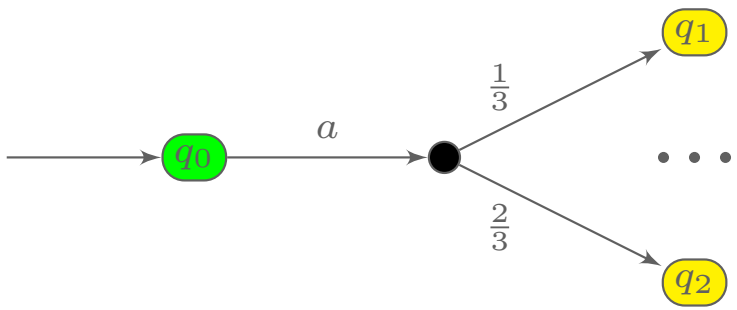

Given a sequence of observations, a *strategy* randomly selects an action to be performed.

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 $\mathcal{B}$, $\mathrm{M}_{\mathcal{C}}$ is built as follows.

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

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

Given some action $\Sigma^{\bullet}$, the transition probability of $\mathrm{M}_{\mathcal{C}}$ 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 $\mathrm{M}_{\mathcal{C}}$ an additional state `lost`.

Illustration

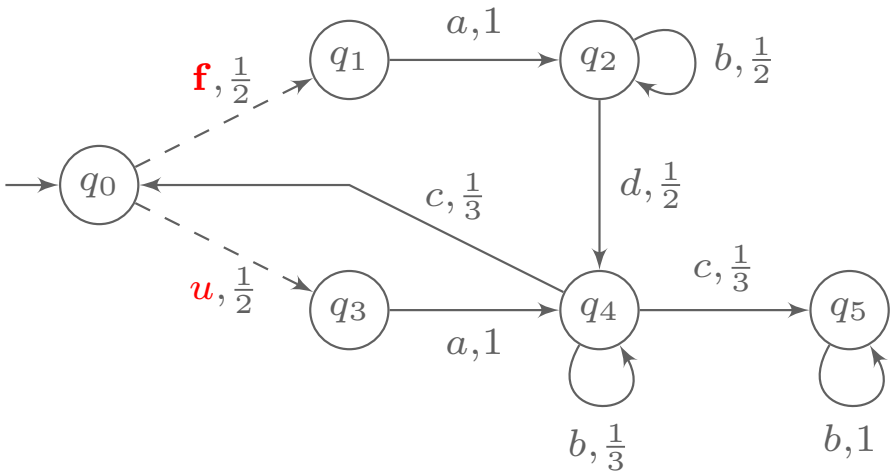

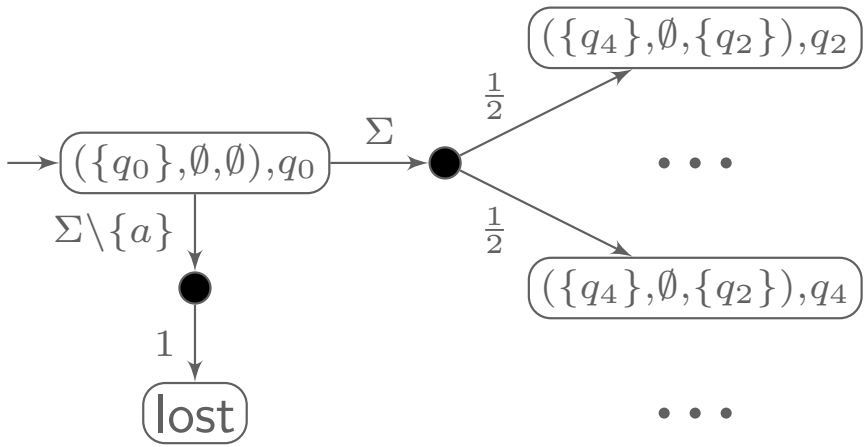

Decidability of the active diagnosis problem

- $\mathcal{C}$ 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.

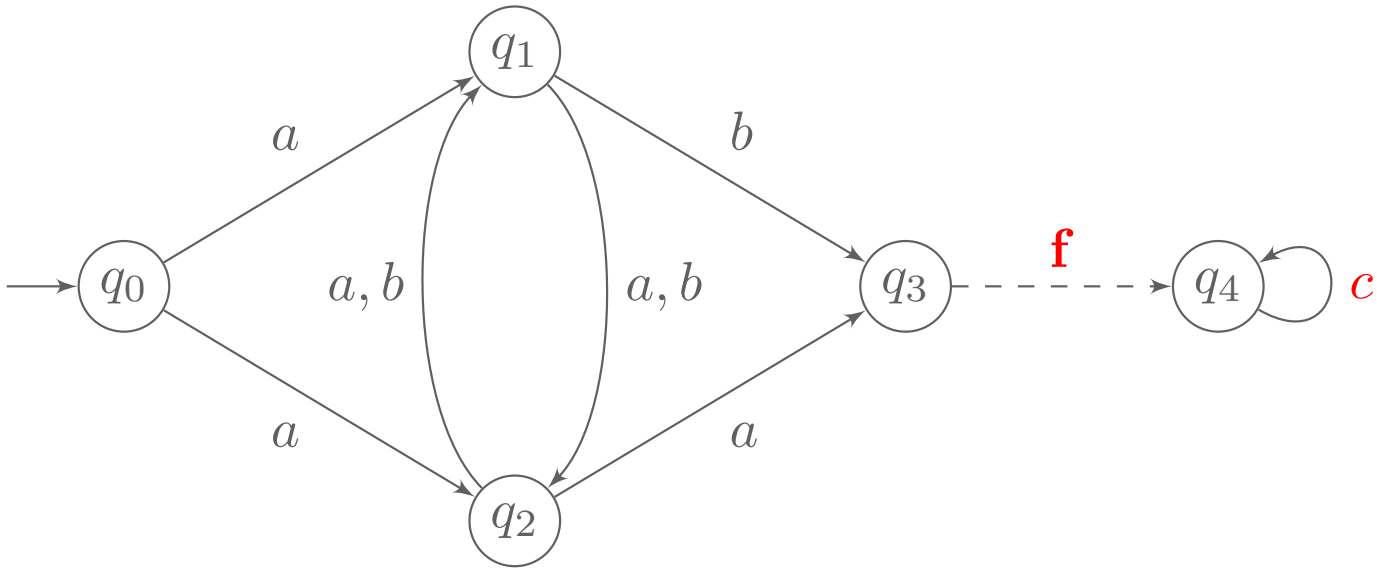

The cLTS is straightforwardly diagnosable but it is not safe.

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

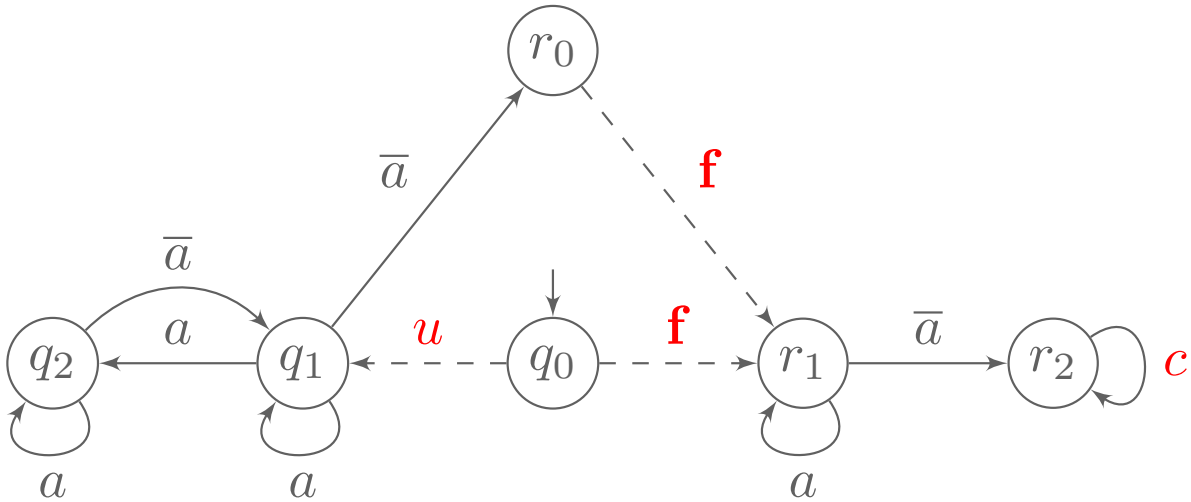

An observed sequence σ is surely faulty iff $\sigma \in \Sigma^* c^\omega$.

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

Finite-memory diagnosers are not enough

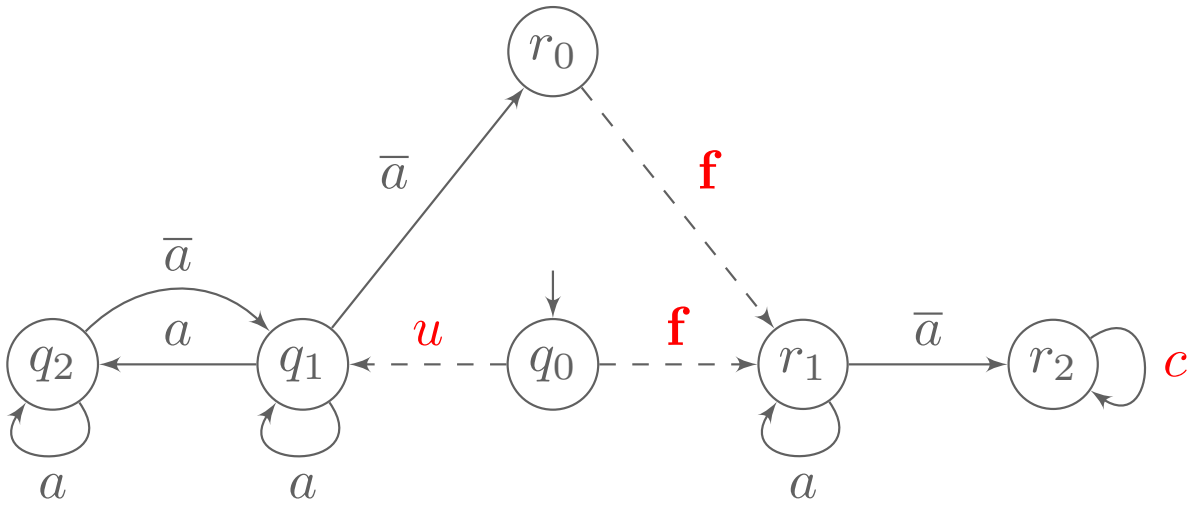

A safe active diagnoser

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

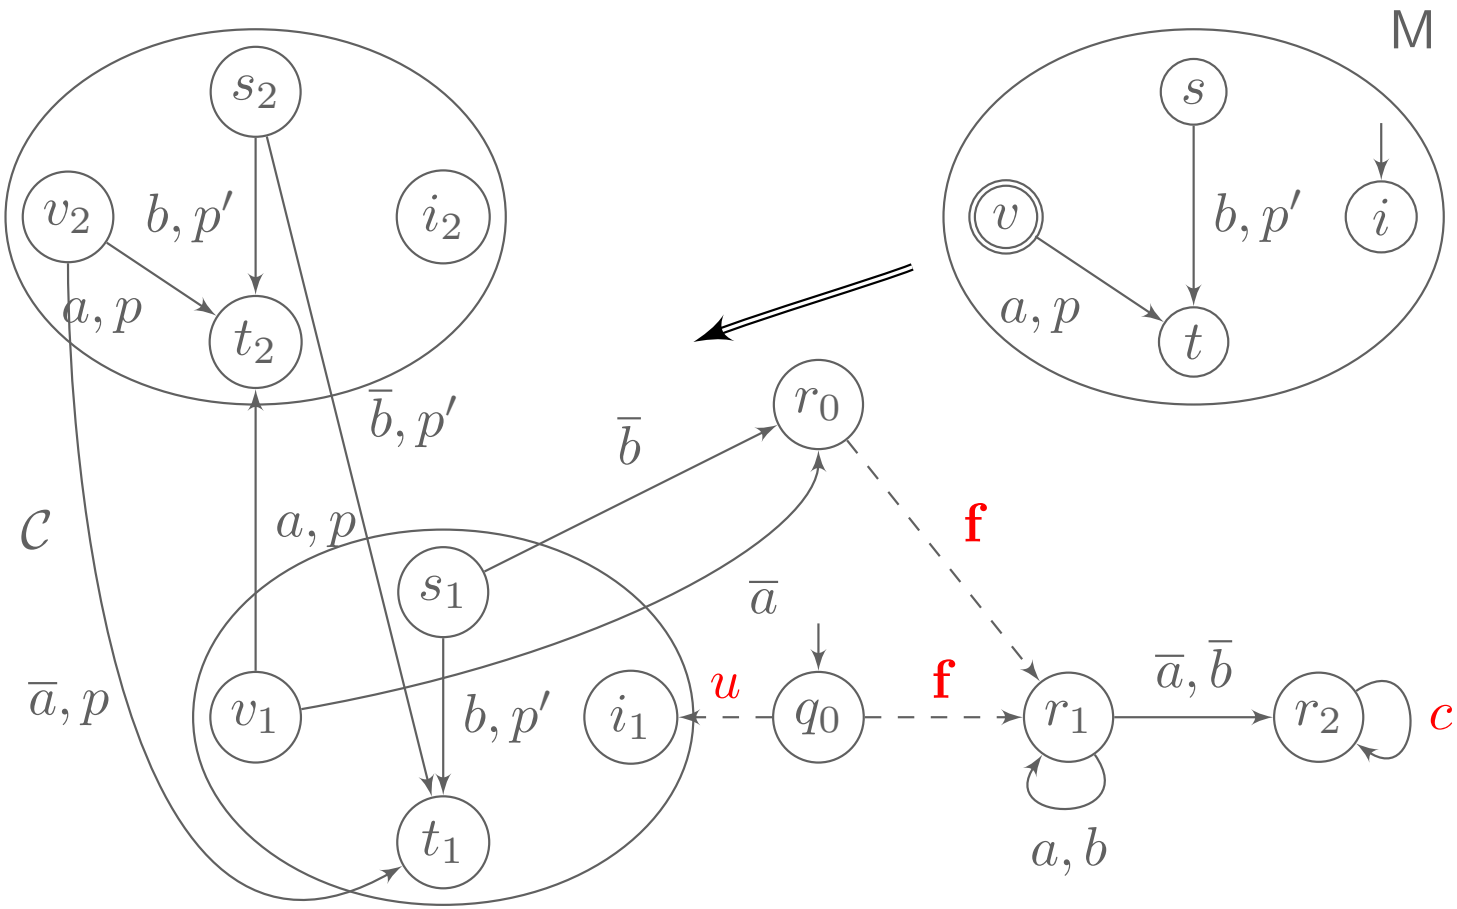

An observed sequence σ is surely faulty iff $\sigma \in \Sigma^* c^\omega$.

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

Revisiting concepts introduced by the "control theory" community

both from a semantical and algorithmic points of view.

- Complete classification of the specifications;
- Useful characterisations using logic, automata theory, games, etc.
- Design of new or more efficient algorithms (even for infinite-state systems);
- (Almost) matching complexity lower bounds and undecidability results.

Perspectives

Short-term

- Closing the gap between lower and upper bounds related to the minimal delay synthesis problem;
- Refining the safety requirement;

Long-term

- Studying the related problems (e.g. opacity, privacy, etc.);
- Investigating further POMDP problems with multiple objectives.
- Modelling and analysing diagnosis with stochastic games.