Accurate Approximate Diagnosability of Stochastic Systems

Nathalie Bertrand[1], Serge Haddad[2], Engel Lefaucheux[1,2]

1 Inria, France
2 LSV, ENS Cachan & CNRS & Inria, France

LATA, March 17th 2016

Diagnosis Framework

LTS: Labelled transition system.

Diagnoser: must tell whether a fault **f** occurred, based on observations.

Convergence hypothesis: no infinite sequence of unobservable events.

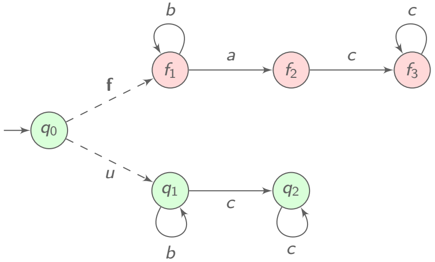

Diagnosis Framework

LTS: Labelled transition system.

Diagnoser: must tell whether a fault **f** occurred, based on observations.

Convergence hypothesis: no infinite sequence of unobservable events.

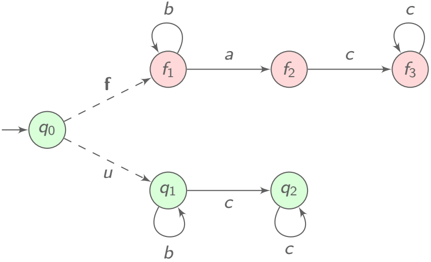

A *run* $\rho = q_0 \xrightarrow{u} q_1 \xrightarrow{c} q_2$ has an *observation sequence* $\mathcal{P}(\rho) = c$.

Diagnosis Framework

LTS: Labelled transition system.

Diagnoser: must tell whether a fault **f** occurred, based on observations.

Convergence hypothesis: no infinite sequence of unobservable events.

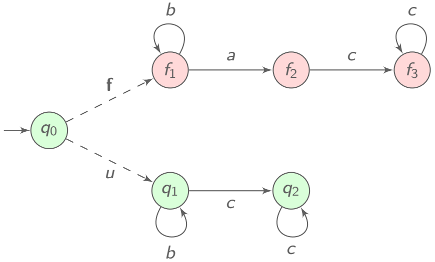

A *run* $\rho = q_0 \xrightarrow{u} q_1 \xrightarrow{c} q_2$ has an *observation sequence* $\mathcal{P}(\rho) = c$.

✓ c is *surely correct* as $\mathcal{P}^{-1}(c) = \{q_0 \xrightarrow{u} q_1 \xrightarrow{c} q_2\}$.

Diagnosis Framework

LTS: Labelled transition system.

Diagnoser: must tell whether a fault **f** occurred, based on observations.

Convergence hypothesis: no infinite sequence of unobservable events.

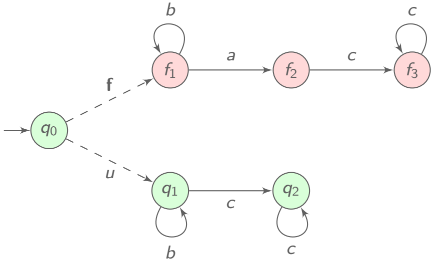

A *run* $\rho = q_0 \xrightarrow{u} q_1 \xrightarrow{c} q_2$ has an *observation sequence* $\mathcal{P}(\rho) = c$.

✓ c is *surely correct* as $\mathcal{P}^{-1}(c) = \{q_0 \xrightarrow{u} q_1 \xrightarrow{c} q_2\}$.

✗ ac is *surely faulty* as $\mathcal{P}^{-1}(ac) = \{q_0 \xrightarrow{\mathbf{f}} f_1 \xrightarrow{a} f_2 \xrightarrow{c} f_3\}$.

Diagnosis Framework

LTS: Labelled transition system.

Diagnoser: must tell whether a fault **f** occurred, based on observations.

Convergence hypothesis: no infinite sequence of unobservable events.

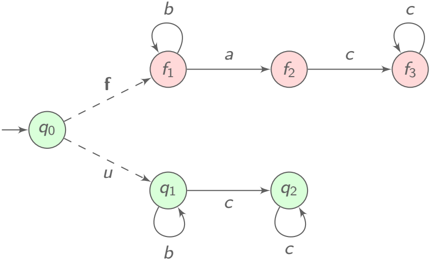

A *run* $\rho = q_0 \xrightarrow{u} q_1 \xrightarrow{c} q_2$ has an *observation sequence* $\mathcal{P}(\rho) = c$.

✓ c is *surely correct* as $\mathcal{P}^{-1}(c) = \{q_0 \xrightarrow{u} q_1 \xrightarrow{c} q_2\}$.

✗ ac is *surely faulty* as $\mathcal{P}^{-1}(ac) = \{q_0 \xrightarrow{\mathbf{f}} f_1 \xrightarrow{a} f_2 \xrightarrow{c} f_3\}$.

? b is *ambiguous* as $\mathcal{P}^{-1}(b) = \{q_0 \xrightarrow{\mathbf{f}} f_1 \xrightarrow{b} f_1, q_0 \xrightarrow{u} q_1 \xrightarrow{b} q_1\}$.

Diagnosis Problems

Diagnoser requirements:

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

Diagnosis Problems

Diagnoser requirements:

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

A decision problem (*diagnosability*): does there exist a diagnoser?

A synthesis problem: how to build a diagnoser?

Diagnosis Problems

Diagnoser requirements:

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

A decision problem (*diagnosability*): does there exist a diagnoser?

A synthesis problem: how to build a diagnoser?

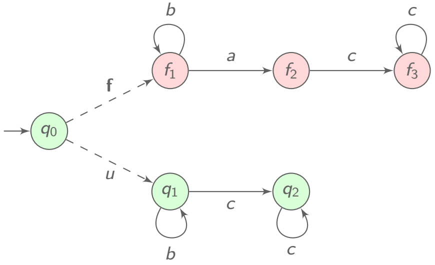

A sound but not reactive diagnoser : claiming a fault when *a* occurs.

Diagnosis of Probabilistic Systems [TT05]

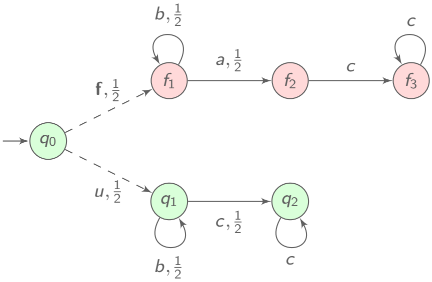

[TT05] Thorsley and Teneketzis
Diagnosability of stochastic discrete-event systems, IEEE TAC, 2005.

Diagnosis of Probabilistic Systems [TT05]

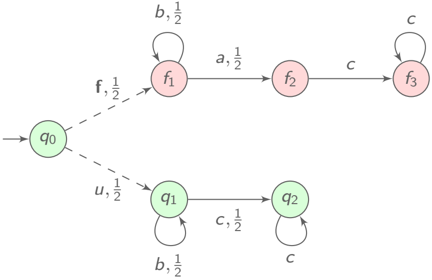

b^n ambiguous but...

[TT05] Thorsley and Teneketzis
Diagnosability of stochastic discrete-event systems, IEEE TAC, 2005.

Diagnosis of Probabilistic Systems [TT05]

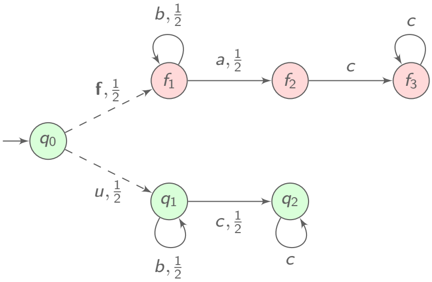

b^n ambiguous but...

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

[TT05] Thorsley and Teneketzis
Diagnosability of stochastic discrete-event systems, IEEE TAC, 2005.

Diagnosis of Probabilistic Systems [TT05]

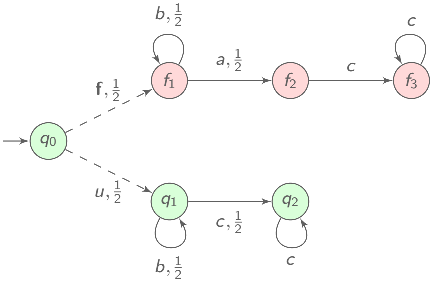

b^n ambiguous but...

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

How to adapt soundness and reactivity?

[TT05] Thorsley and Teneketzis
Diagnosability of stochastic discrete-event systems, IEEE TAC, 2005.

Exact Diagnosis [BHL14]

An *exact diagnoser* fulfills

- Soundness: if a fault is claimed, a fault happened.

[BHL14] Bertrand, Haddad, Lefaucheux

Foundation of Diagnosis and Predictability in Probabilistic Systems, FSTTCS'14.

Exact Diagnosis [BHL14]

An *exact diagnoser* fulfills

- Soundness: if a fault is claimed, a fault happened.
- Reactivity: the diagnoser will provide information almost surely.

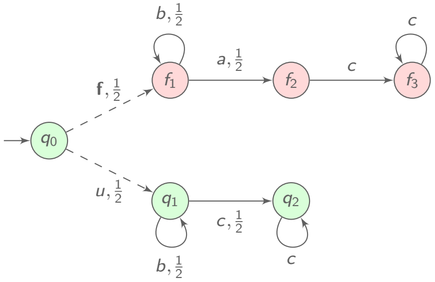

Exactly diagnosable.

[BHL14] Bertrand, Haddad, Lefaucheux

Foundation of Diagnosis and Predictability in Probabilistic Systems, FSTTCS'14.

Exact Diagnosis [BHL14]

An *exact diagnoser* fulfills

- Soundness: if a fault is claimed, a fault happened.
- Reactivity: the diagnoser will provide information almost surely.

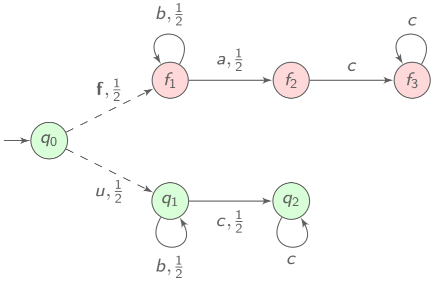

Exactly diagnosable.

Exact diagnosability is PSPACE-complete.

Also studied : exact prediction and prediagnosis.
[BHL14] Bertrand, Haddad, Lefaucheux

Foundation of Diagnosis and Predictability in Probabilistic Systems, FSTTCS'14.

Exact Diagnosis versus Approximate Diagnosis

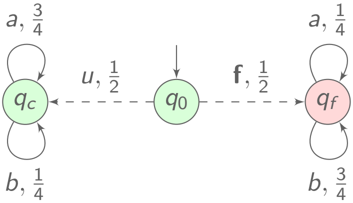

Exact Diagnosis versus Approximate Diagnosis

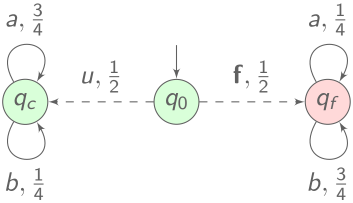

Not exactly diagnosable

Exact Diagnosis versus Approximate Diagnosis

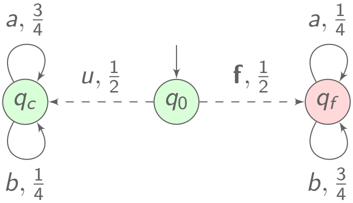

Not exactly diagnosable

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

Exact Diagnosis versus Approximate Diagnosis

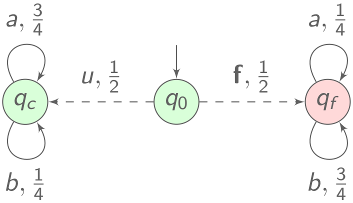

Not exactly diagnosable

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.

Outline

Specification of Approximate Diagnosis

AA-diagnosis is Easy

Other Approximate Diagnoses are Hard

Outline

Specification of Approximate Diagnosis

AA-diagnosis is Easy

Other Approximate Diagnoses are Hard

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 $\sigma \in \Sigma_o^*$,

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

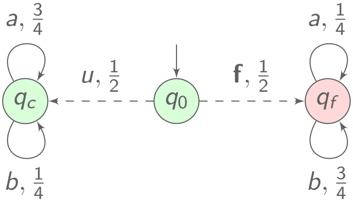

$\mathrm{CorP}(a) = 3/4$,

Proportion of Correct Runs

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

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

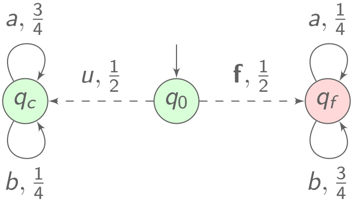

$\mathrm{CorP}(a) = 3/4$, $\mathrm{CorP}(ab) = 1/2$,

Proportion of Correct Runs

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

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

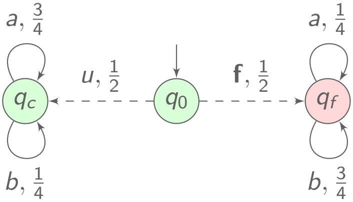

$\mathrm{CorP}(a) = 3/4$, $\mathrm{CorP}(ab) = 1/2$, $\mathrm{CorP}(abb) = 1/4$, $\mathrm{CorP}(abbb) = 1/10$.

Relaxing Soundness

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

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

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 $\mathsf{CorP}(\sigma) \leq \varepsilon$.
- 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.

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 diagnosers.

Approximate Diagnosis Problems

	Reactivity	
Accuracy	*ε-diagnosability* Given $\varepsilon > 0$, does there exist an ε-diagnoser?	*uniform ε-diagnosability* Given $\varepsilon > 0$, does there exist a uniform ε-diagnoser?
	AA-diagnosability For all $\varepsilon > 0$, does there exist an ε-diagnoser?	*uniform AA-diagnosability* For all $\varepsilon > 0$, does there exist a uniform ε-diagnoser?

AA-diagnosability allows to select ε depending on external requirements.

Illustration

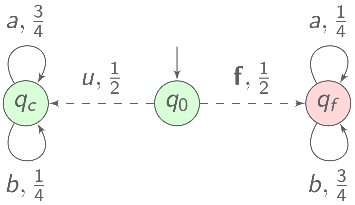

AA-diagnosable but
not uniformly AA-diagnosable

Illustration

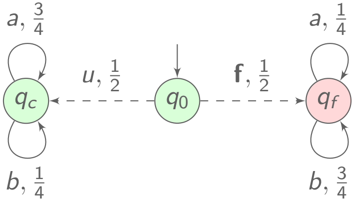

AA-diagnosable but
not uniformly AA-diagnosable

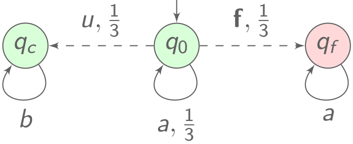

Uniformly AA-diagnosable but
not exactly diagnosable

Establishing relations between the Specifications

$$
\begin{tikzcd}
\begin{array}{c}\text{uniformly}\\ \text{exact diagnosis}\end{array} \arrow[r, Rightarrow] \arrow[d, Leftrightarrow] & \begin{array}{c}\text{uniformly}\\ \text{AA-diagnosable}\end{array} \arrow[r, Rightarrow, "\text{for all } \varepsilon"] \arrow[d, Rightarrow] & \begin{array}{c}\text{uniformly}\\ \varepsilon\text{-diagnosable}\end{array} \arrow[d, Rightarrow] \\
\text{exact diagnosis} & \text{AA-diagnosable} \arrow[r, Rightarrow, "\text{for all } \varepsilon"] & \varepsilon\text{-diagnosable}
\end{tikzcd}
$$

Complexity of the Problems

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

Outline

Specification of Approximate Diagnosis

AA-diagnosis is Easy

Other Approximate Diagnoses are Hard

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.

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.

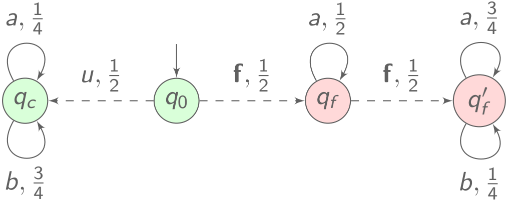

- an initial state, q_0;
- an arbitrary pLTS with states $\{q_f, q'_f\}$;
- a correct pLTS with state q_c.

Solving AA-diagnosability for Initial-Fault pLTS

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

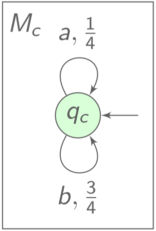

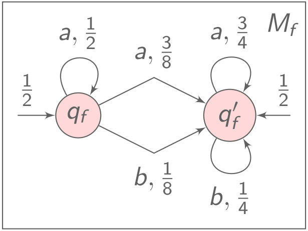

Solving AA-diagnosability for Initial-Fault pLTS

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

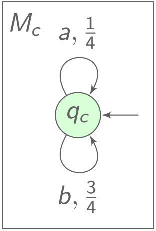

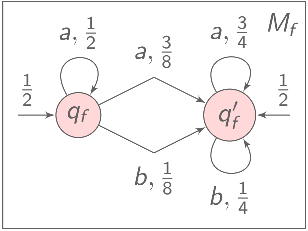

- $\mathbb{P}^M(E)$ = measure of infinite runs of M with observation in E.
Distance 1 problem: $\exists E$ (measurable) $\subseteq \Sigma_o^\omega, \mathbb{P}^{M_c}(E) - \mathbb{P}^{M_f}(E) = 1$?
- Illustration: $E = \{\sigma \mid \limsup_{n\to\infty} \frac{|\sigma_{\downarrow n}|_b}{|\sigma_{\downarrow n}|_a} > 1\}$

[CK14] Chen and Kiefer
On the Total Variation Distance of Labelled Markov Chains, CSL-LICS'14.

Solving AA-diagnosability for Initial-Fault pLTS

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

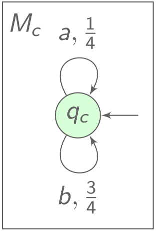

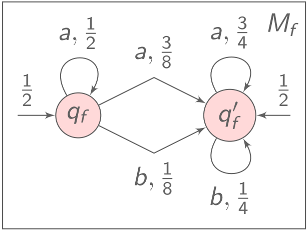

• $\mathbb{P}^M(E)$ = measure of infinite runs of M with observation in E.
Distance 1 problem: $\exists E$ (measurable) $\subseteq \Sigma_o^\omega, \mathbb{P}^{M_c}(E) - \mathbb{P}^{M_f}(E) = 1$?

• 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.

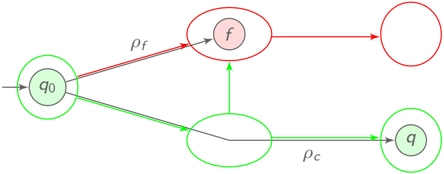

$$\mathcal{P}(\rho_c) = \mathcal{P}(\rho_f)$$

Solving AA-diagnosability

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

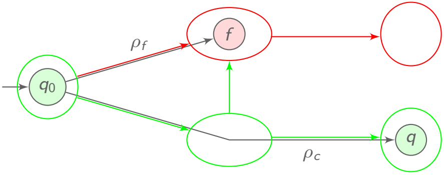

$$\mathcal{P}(\rho_c) = \mathcal{P}(\rho_f)$$

- Checking distance 1 for all relevant pairs.

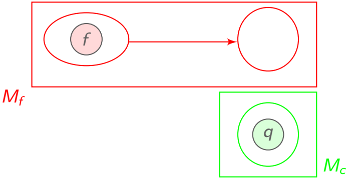

Solving AA-diagnosability

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

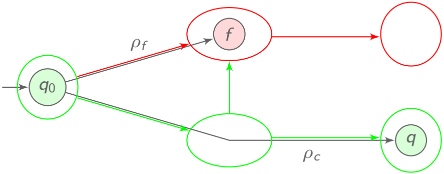

$$\mathcal{P}(\rho_c) = \mathcal{P}(\rho_f)$$

- Checking distance 1 for all relevant pairs.

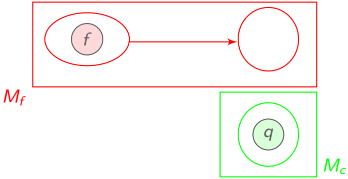

AA-diagnosability is decidable in PTIME.

Diagnoser Synthesis

An ε-diagnoser may need infinite memory.

Diagnoser Synthesis

An ε-diagnoser may need infinite memory.

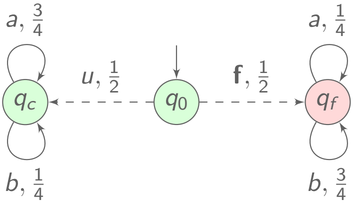

Diagnoser Synthesis

An ε-diagnoser may need infinite memory.

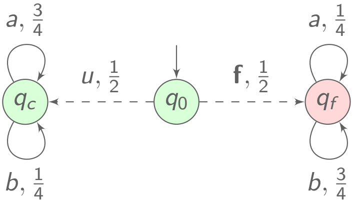

For all $k < n$, a^k and a^n lead to different states of the diagnoser.

Diagnoser Synthesis

An ε-diagnoser may need infinite memory.

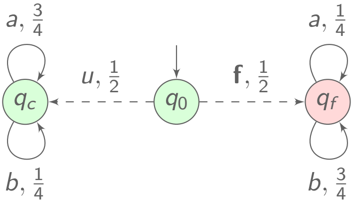

For all $k < n$, a^k and a^n lead to different states of the diagnoser.

Otherwise for all $i \in \mathbb{N}$, $a^{k+i(n-k)}$ lead to the same state.

Diagnoser Synthesis

An ε-diagnoser may need infinite memory.

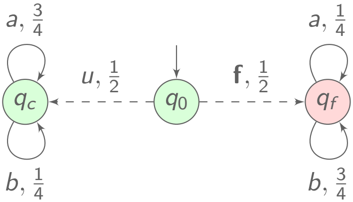

For all $k < n$, a^k and a^n lead to different states of the diagnoser.

Otherwise for all $i \in \mathbb{N}$, $a^{k+i(n-k)}$ lead to the same state.

By reactivity for some σ, the diagnoser must claim a fault after $a^k\sigma$ and thus after all $a^{k+i(n-k)}\sigma$.

Diagnoser Synthesis

An ε-diagnoser may need infinite memory.

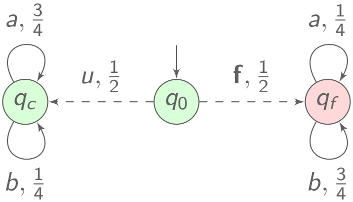

For all $k < n$, a^k and a^n lead to different states of the diagnoser.

Otherwise for all $i \in \mathbb{N}$, $a^{k+i(n-k)}$ lead to the same state.

By reactivity for some σ, the diagnoser must claim a fault after $a^k\sigma$ and thus after all $a^{k+i(n-k)}\sigma$.

But $\lim_{i\to\infty} \mathrm{CorP}(a^{k+i(n-k)}\sigma) = 1$.

Diagnoser Synthesis

An ε-diagnoser may need infinite memory.

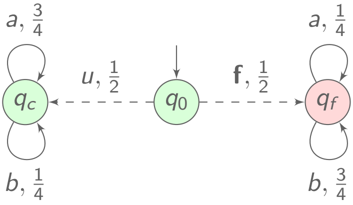

For all $k < n$, a^k and a^n lead to different states of the diagnoser.

Otherwise for all $i \in \mathbb{N}$, $a^{k+i(n-k)}$ lead to the same state.

By reactivity for some σ, the diagnoser must claim a fault after $a^k\sigma$ and thus after all $a^{k+i(n-k)}\sigma$.

But $\lim_{i\to\infty} \mathrm{CorP}(a^{k+i(n-k)}\sigma) = 1$.

For exact diagnosis, one can build a diagnoser exponential in the size of the pLTS [BHL14].

Outline

Specification of Approximate Diagnosis

AA-diagnosis is Easy

Other Approximate Diagnoses are Hard

The Emptiness Problem for Probabilistic Automata (PA)

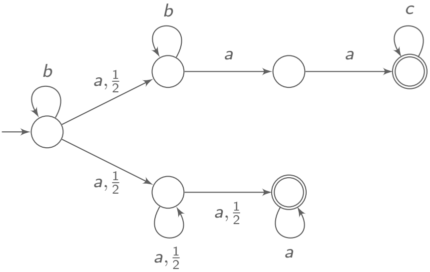

$\mathbb{P}(b) = 0, \mathbb{P}(baa) = \frac{1}{4}, \mathbb{P}(baaa) = \frac{7}{8}$

The Emptiness Problem for Probabilistic Automata (PA)

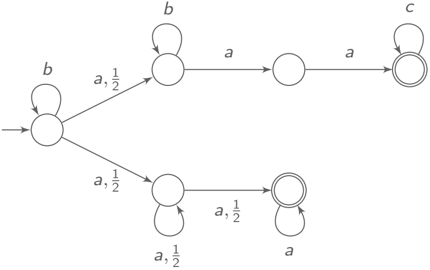

$\mathbb{P}(b) = 0, \mathbb{P}(baa) = \frac{1}{4}, \mathbb{P}(baaa) = \frac{7}{8}$

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)

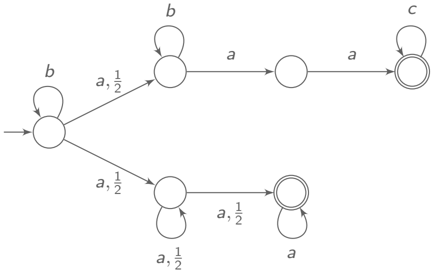

$\mathbb{P}(b) = 0, \mathbb{P}(baa) = \frac{1}{4}, \mathbb{P}(baaa) = \frac{7}{8}$

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

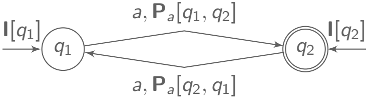

From PA to Uniform AA-diagnosability

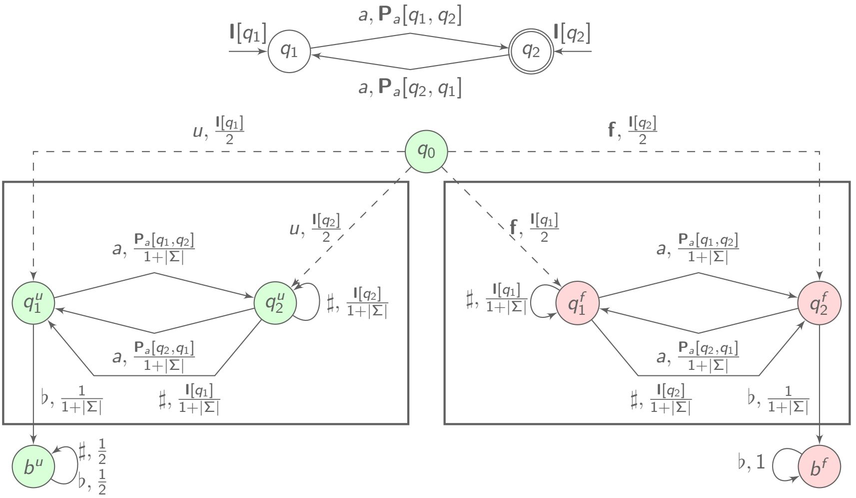

From PA to Uniform AA-diagnosability

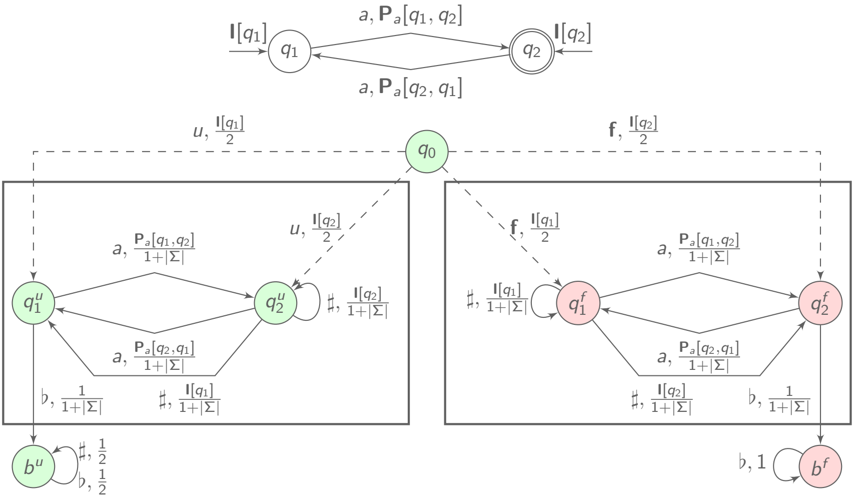

If $\exists w \in \Sigma_o^*, \mathbb{P}_{\mathcal{A}}(w) > 1/2$ then $\lim_{n \to \infty} \mathrm{CorP}((w\sharp)^n \flat) = 1$.

From PA to Uniform AA-diagnosability

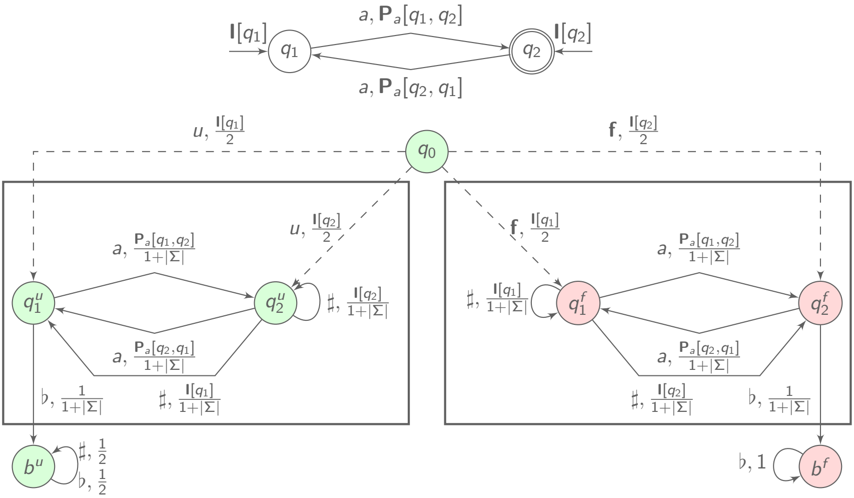

If $\exists w \in \Sigma_o^*, \mathbb{P}_{\mathcal{A}}(w) > 1/2$ then $\lim_{n\to\infty} \mathrm{CorP}((w\sharp)^n\flat) = 1$.

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

Conclusion

Contributions

- Investigation of semantical issues
- Complexity of the notions of approximate diagnosis
  - A PTIME algorithm for AA-diagnosability
  - Undecidability of other approximate diagnosability

Conclusion

Contributions

- Investigation of semantical issues
- Complexity of the notions of approximate diagnosis
  - A PTIME algorithm for AA-diagnosability
  - Undecidability of other approximate diagnosability

Future work

- Approximate prediction and prediagnosis
- Diagnosis of infinite state stochastic systems