Specification and Verification of Diagnosis and Predictability in Probabilistic Systems

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

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

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.

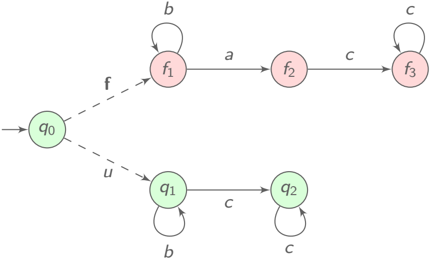

Diagnosis Framework

LTS: Labelled transition system.

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

Convergence hypothesis: no infinite sequence of unobservable events.

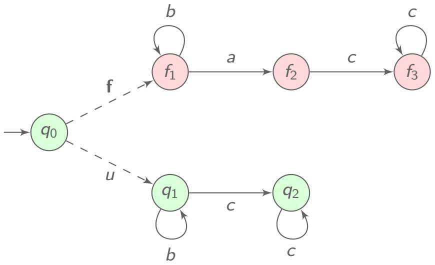

A *run* $\rho = q_0 \xrightarrow{u} q_1 \xrightarrow{c} q_2$ has an *observation sequence* $\mathcal{P}(\rho) = c$.

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.

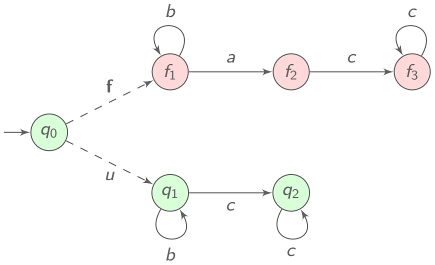

A *run* $\rho = q_0 \xrightarrow{u} q_1 \xrightarrow{c} q_2$ has an *observation sequence* $\mathcal{P}(\rho) = c$.

✗ 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.

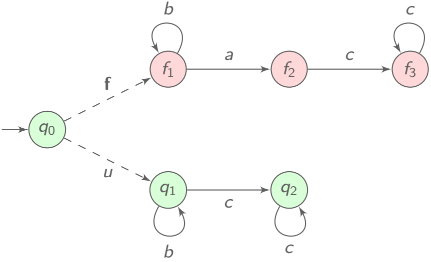

A *run* $\rho = q_0 \xrightarrow{u} q_1 \xrightarrow{c} q_2$ has an *observation sequence* $\mathcal{P}(\rho) = c$.

✗ 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?

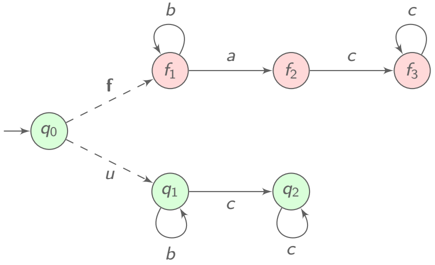

A sound but not reactive diagnoser : claiming a fault when *a* occurs.

Diagnosis of Probabilistic Systems

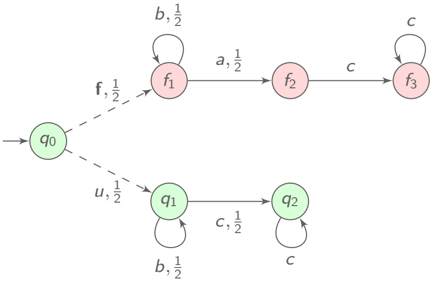

[TT05] Thorsley and Teneketzis

Diagnosability of stochastic discrete-event systems, IEEE TAC, 2005.

Diagnosis of Probabilistic Systems

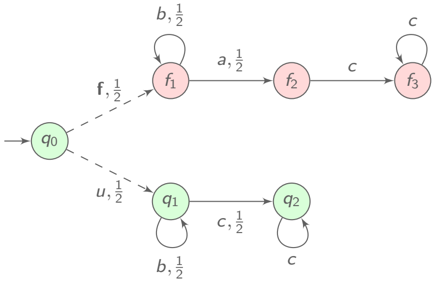

b^n ambiguous but...

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

Diagnosis of Probabilistic Systems

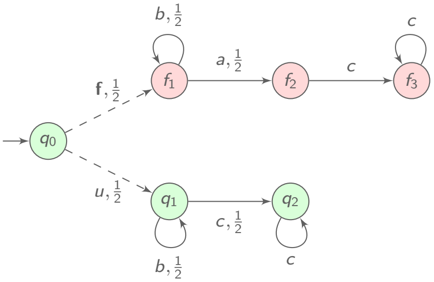

b^n ambiguous but...

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

[TT05] Thorsley and Teneketzis

Diagnosability of stochastic discrete-event systems, IEEE TAC, 2005.

Diagnosis of Probabilistic Systems

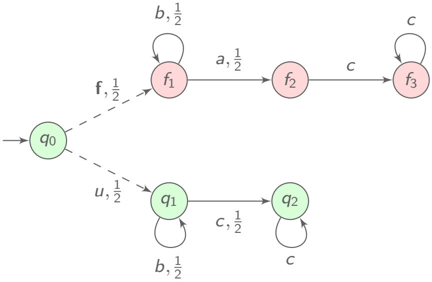

b^n ambiguous but...

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

How to adapt soundness and reactivity?

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

Outline

Outline

All runs or faulty runs?

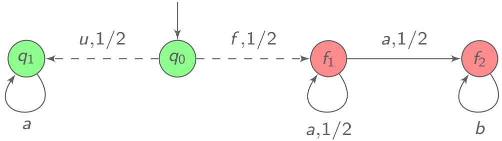

All runs or faulty runs?

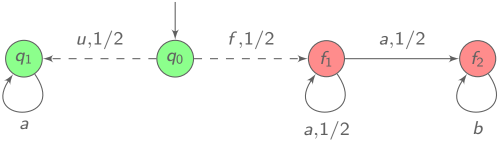

a^+ is ambiguous

All runs or faulty runs?

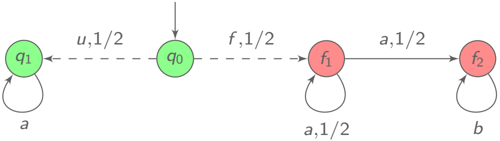

a^+ is ambiguous

$$\lim_{n\to\infty} \mathbb{P}(fa^n) = 0$$

All runs or faulty runs?

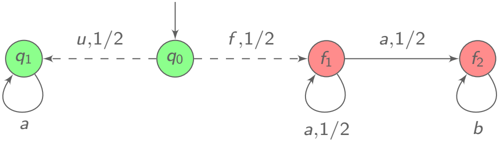

a^+ is ambiguous

$$\lim_{n\to\infty} \mathbb{P}(fa^n) = 0$$

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

All runs or faulty runs?

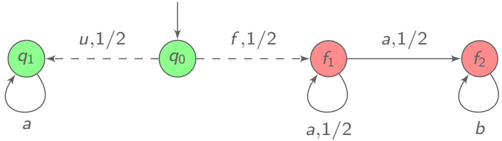

a^+ is ambiguous

$$\lim_{n\to\infty} \mathbb{P}(fa^n) = 0$$

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

Two reactivity specifications:

- Detect a fault almost surely.

All runs or faulty runs?

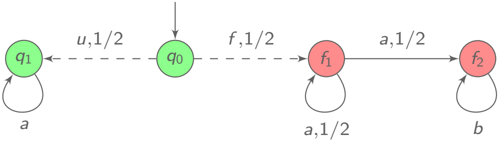

a^+ is ambiguous

$$\lim_{n\to\infty} \mathbb{P}(fa^n) = 0$$

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

Two reactivity specifications:

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

Infinite sequences or their finite prefixes?

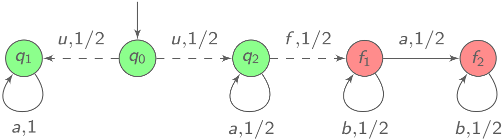

Infinite sequences or their finite prefixes?

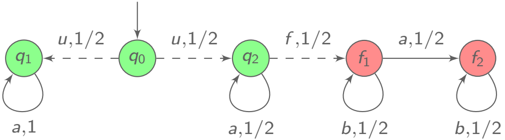

a^n is ambiguous and $\mathbb{P}(q_0 u (q_1 a)^n) = \frac{1}{2}$.

Infinite sequences or their finite prefixes?

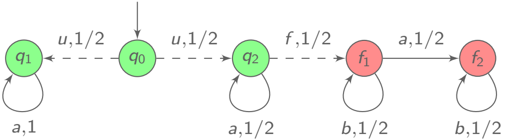

a^n is ambiguous and $\mathbb{P}(q_0 u (q_1 a)^n) = \frac{1}{2}$.

a^ω is correct.

Infinite sequences or their finite prefixes?

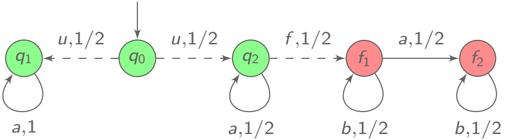

a^n is ambiguous and $\mathbb{P}(q_0 u (q_1 a)^n) = \frac{1}{2}$.

a^ω is correct.

- The probability of ambiguous sequences goes to 0 when the length goes to infinity.

Infinite sequences or their finite prefixes?

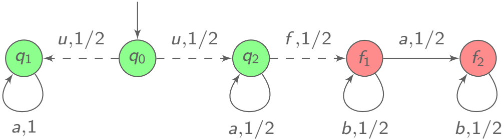

a^n is ambiguous and $\mathbb{P}(q_0 u (q_1 a)^n) = \frac{1}{2}$.

a^ω is correct.

- The probability of ambiguous sequences goes to 0 when the length goes to infinity.
- An infinite sequence is almost surely non ambiguous.

Four specifications of diagnosis

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

Outline

IF-diagnosability

Specification of IF: 1) Infinite sequences
2) Fault diagnosis

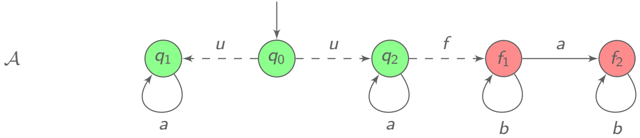

IF-diagnosability

Specification of IF: 1) Infinite sequences
2) Fault diagnosis

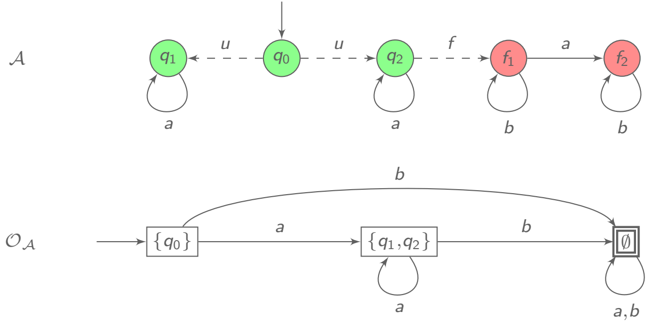

IF-diagnosability

Specification of IF: 1) Infinite sequences
2) Fault diagnosis

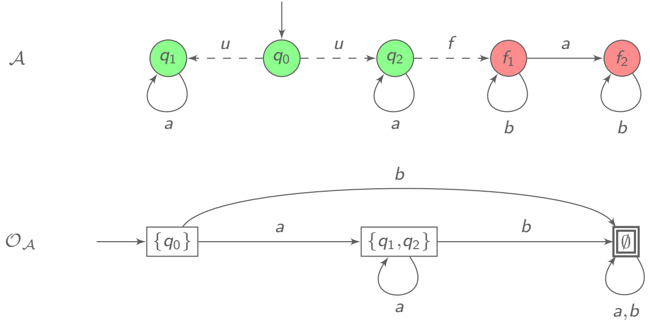

$\mathcal{A}$ is not IF-diagnosable
iff
there exists a state (q, U) in a BSCC of $\mathcal{A} \times \mathcal{O}_\mathcal{A}$ with q faulty and $U \neq \emptyset$.

IF-diagnosability

Specification of IF: 1) Infinite sequences
2) Fault diagnosis

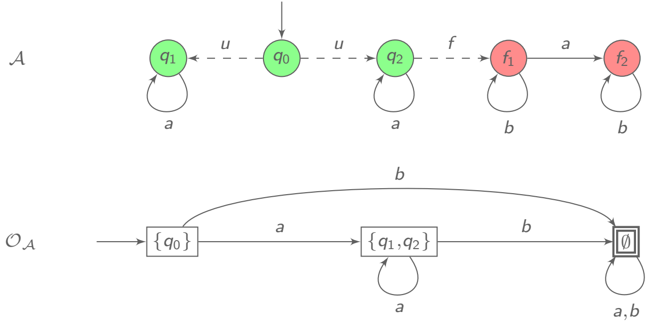

$\mathcal{A}$ is not IF-diagnosable
iff
there exists a state (q, U) in a BSCC of $\mathcal{A} \times \mathcal{O}_\mathcal{A}$ with q faulty and $U \neq \emptyset$.

If $\mathcal{A}$ is a IF-diagnosable pLTS with n correct states
then $\mathcal{O}_\mathcal{A}$ is a IF-diagnoser of $\mathcal{A}$ with at most 2^n states.

IA-diagnosability

Specification of IA: 1) Infinite sequences
2) All sequences disambiguation

$\mathcal{A}$

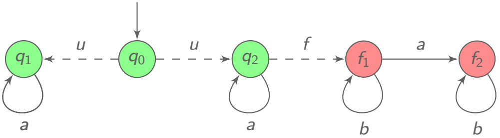

IA-diagnosability

Specification of IA: 1) Infinite sequences
2) All sequences disambiguation

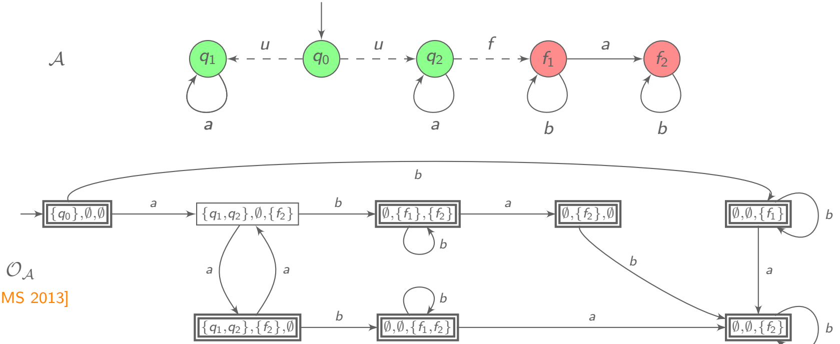

[HHMS 2013]

IA-diagnosability

Specification of IA: 1) Infinite sequences
2) All sequences disambiguation

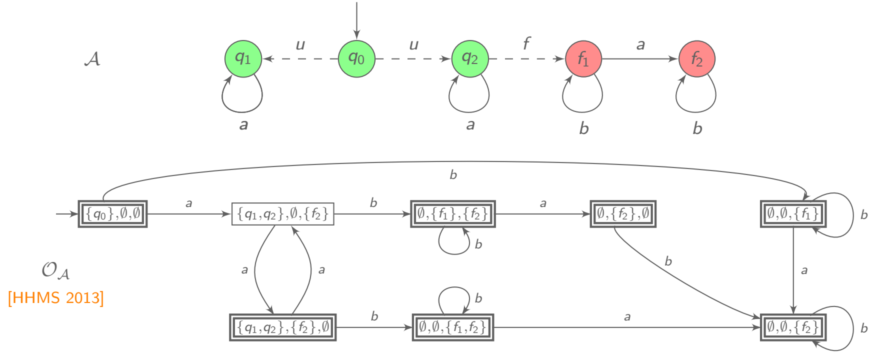

$\mathcal{A}$ is not IA-diagnosable iff
there exists a BSCC of $\mathcal{A} \times \mathcal{O}_\mathcal{A}$ where every state (q, U, V, W) satisfies
q faulty and $U \neq \emptyset$ or q correct and $W \neq \emptyset$.

Outline

Diagnosability is in PSPACE

The IA/FA/IF-diagnosability problems are in PSPACE.

Diagnosability is in PSPACE

The IA/FA/IF-diagnosability problems are in PSPACE.

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

Outline

Diagnosability is PSPACE-hard

First step: We define and analyse a new universality notion.

$\mathcal{L} \subseteq \Sigma^*$ 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

First step: We define and analyse a new universality notion.

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

The eventual universality problem for NFA is PSPACE-hard.

Second step: We reduce eventual universality to diagnosability.

The IA/FA/IF-diagnosability problems are PSPACE-hard.

Diagnosability is PSPACE-hard

First step: We define and analyse a new universality notion.

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

The eventual universality problem for NFA is PSPACE-hard.

Second step: We reduce eventual universality to diagnosability.

The IA/FA/IF-diagnosability problems are PSPACE-hard.

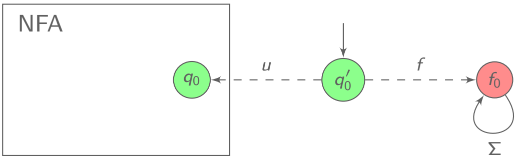

Diagnosability is PSPACE-hard

First step: We define and analyse a new universality notion.

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

The eventual universality problem for NFA is PSPACE-hard.

Second step: We reduce eventual universality to diagnosability.

The IA/FA/IF-diagnosability problems are PSPACE-hard.

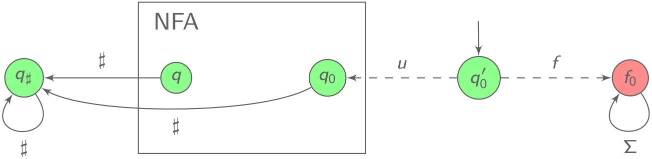

Comparison with diagnosability of non-stochastic systems

The IA/FA/IF-diagnosability problems are PSPACE-complete.

Comparison with diagnosability of non-stochastic systems

The IA/FA/IF-diagnosability problems are PSPACE-complete.

The diagnosability problem is in PTIME [Jiang/Huang/Chandra/Kumar-tac01].

Sketch of proof

- builds the product of the LTS with a copy restricted to correct states
- checks for SCC with faulty states in the first component

Comparison with diagnosability of non-stochastic systems

The IA/FA/IF-diagnosability problems are PSPACE-complete.

The diagnosability problem is in PTIME [Jiang/Huang/Chandra/Kumar-tac01].

Sketch of proof

- builds the product of the LTS with a copy restricted to correct states
- checks for SCC with faulty states in the first component

Erroneous adaptation to stochastic systems in [Chen/Kumar-tase13].

Outline

Sure and almost-sure k-predictability

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

Sure and almost-sure k-predictability

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

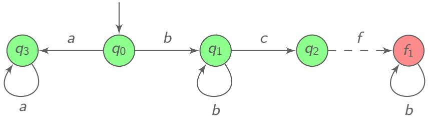

a^+	never faulty
b^+c	surely eventually faulty
b^+	almost surely eventually faulty

Sure and almost-sure k-predictability

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

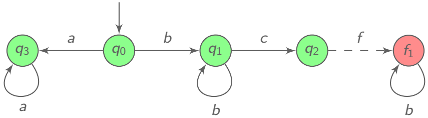

a^+	never faulty
b^+c	surely eventually faulty
b^+	almost surely eventually faulty

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.

Sure and almost-sure k-predictability

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

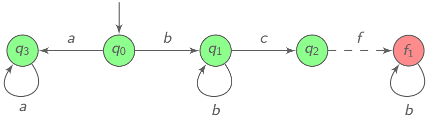

a^+	never faulty
b^+c	surely eventually faulty
b^+	almost surely eventually faulty

- 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.

Outline

Complexity of predictability

The (almost) sure predictability problem is NLOGSPACE-complete.

Complexity of predictability

The (almost) sure predictability problem is NLOGSPACE-complete.

Sketch of proof

- **Upper bound:** graph characterisation

Complexity of predictability

The (almost) sure predictability problem is NLOGSPACE-complete.

Sketch of proof

- **Upper bound:** graph characterisation
- **Lower bound:** reduction of reachability in acyclic graphs

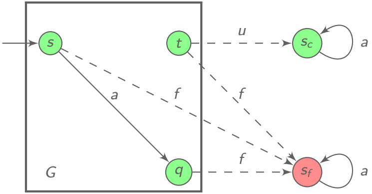

Predictor synthesis

An (almost) surely k-predictable pLTS with n correct states admits an (almost) sure k-predictor of size 2^n.

Construction similar to the previous observers.

Predictor synthesis

An (almost) surely k-predictable pLTS with n correct states admits an (almost) sure k-predictor of size 2^n.

Construction similar to the previous observers.

$\mathcal{A}_n$ has $2n+2$ correct states and every (almost) sure 0-predictor of $\mathcal{A}_n$ needs at least 2^n states.

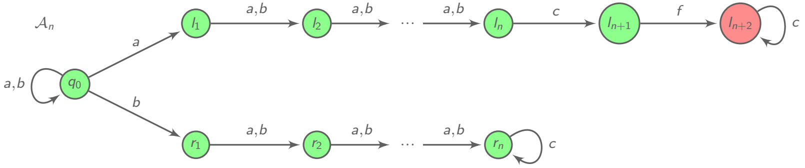

Outline

Prediagnosability

Prediagnoser: detects and foresees faults analysing past and future

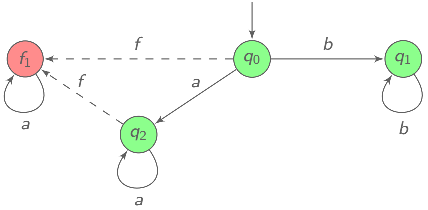

b^+	never faulty
a^+	almost surely eventually faulty

Prediagnosability

Prediagnoser: detects and foresees faults analysing past and future

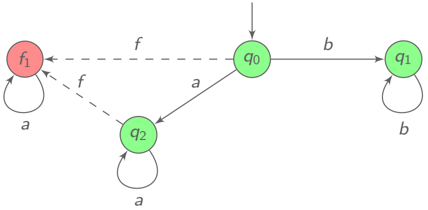

b^+ never faulty
a^+ almost surely eventually faulty

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

Prediagnosability

Prediagnoser: detects and foresees faults analysing past and future

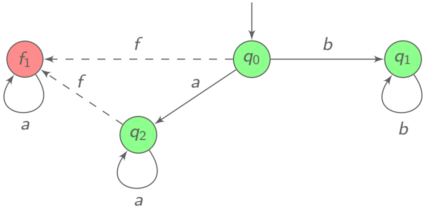

b^+	never faulty
a^+	almost surely eventually faulty

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

Prediagnosability is PSPACE-complete.

Conclusion: Foundation of stochastic diagnosis

Summary of contributions

- Investigation of semantical issues
- Introduction of prediagnosability
- Tight complexity bounds for diagnosability and diagnoser synthesis problems

Conclusion: Foundation of stochastic diagnosis

Summary of contributions

- Investigation of semantical issues
- Introduction of prediagnosability
- Tight complexity bounds for diagnosability and diagnoser synthesis problems

Future work

- Approximate diagnosis (relaxing soundness)
- Other paradigms related to partial observation (detectability, opacity, etc.)
- Space and time optimisation of observations