

Fault-Tolerant Supervisory Control in Terms of Formal Languages

Thomas Moor

Friedrich-Alexander Universität-Erlangen-Nürnberg, Germany

IFAC WC 2020, Berlin, Germany

Discrete-Event Systems

Supervisory Control

Naive Fault-Tolerant Control

Active Fault-Tolerant Control

Post-Fault Recovery

Fault-Hiding Approach

Formal Languages

- Σ^* set of all finite strings over Σ
- empty string $\epsilon \in \Sigma^*$, $\epsilon \notin \Sigma$
- *-language $M \subseteq \Sigma^*$
- prefix operator $\text{pre } s := \{ t \mid \exists r : tr = s \}$
- note that $\text{pre } M = \{ t \mid \exists r : tr \in M \}$
- M is closed $:\Leftrightarrow M = \text{pre } M$

A natural domain for the interpretation of liveness properties are ω -languages, i.e., sets of infinite-length strings $w \in \Sigma^\omega$.

If there are no deadlocks, we may use

$$\mathcal{M} := \{ w \in \Sigma^\omega \mid \text{pre } w \subseteq \text{pre } M \}$$

to model the process w.r.t. infinite time.

If, in addition, there are no livelocks, we may consider

$$\mathcal{L} := \{ w \in \Sigma^\omega \mid \|(\text{pre } w) \cap L\| = \infty \}.$$

to model the process w.r.t. infinite time.

Control patterns

$$\Gamma := \{ \gamma \subseteq \Sigma \mid \Sigma_{uc} \subseteq \gamma \}$$

Projection

- natural projection $p_o: \Sigma^* \rightarrow \Sigma_o^*$
read “removes all symbols not from Σ_o ”
- for languages take point-wise images
- set-valued inverse $p_o^{-1}: \Sigma_o^* \rightsquigarrow \Sigma^*$
read “arbitrarily inserts symbols from Σ_{uo} ”

Language Quotient

$$K/E := \{ s \mid \exists t \in E : st \in K \}.$$

Language Convergence

K *finitely converges* to E if there exists a uniform bound k such that every $s \in K$ can be decomposed

$$s = vw, \quad w \in E, \quad \text{and } |v| \leq k.$$

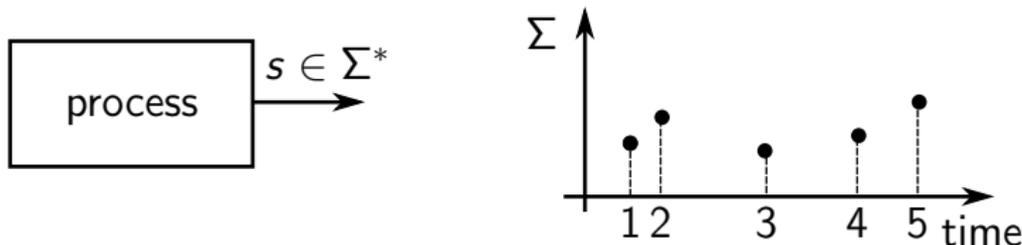
This is written $E \Leftarrow K$.

For not-uniformly bounded convergence, one refers to the respective ω -languages and requires

$$\lim K \subseteq \lim(\Sigma^* E).$$

**A discrete-event system is a model of a process
... with a particular focus on the occurrence of events**

- finite set Σ of symbols $\sigma \in \Sigma$
- only event ordering is regarded relevant (logic time)
- within finite physical time a finite sequence $s \in \Sigma^*$ is generated
- set $M \subseteq \Sigma^*$ of sequences that can be generated
- write $\text{pre } M$ to emphasise that $M = \text{pre } M$ (local behaviour)



A closed language $\text{pre } M \subseteq \Sigma^*$ is a discrete-event system.

Properties

- safety – bad things never happen

with $\text{pre } E \subseteq \Sigma^*$, require

$$\text{pre } M \subseteq \text{pre } E$$

- liveness – good things do happen

free of deadlocks

$$(\forall s \in \text{pre } M)(\exists \sigma \in \Sigma)[s\sigma \in \text{pre } M]$$

free of livelocks w.r.t. $L \subseteq \text{pre } M$

$$(\forall s \in \text{pre } M)(\exists t \in \Sigma^*)[st \in L \cap \text{pre } M]$$

For systems with liveness properties:

A language $L \subseteq \Sigma^*$ is a discrete-event system.

Discrete-Event Systems

Supervisory Control

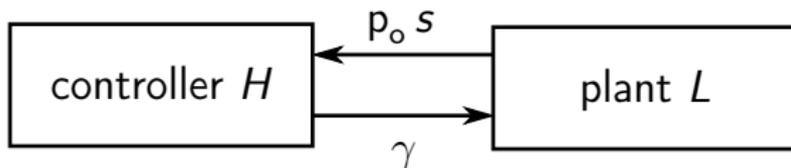
Naive Fault-Tolerant Control

Active Fault-Tolerant Control

Post-Fault Recovery

Fault-Hiding Approach

W.r.t. the partitioning $\Sigma = \Sigma_c \dot{\cup} \Sigma_{uc} = \Sigma_o \dot{\cup} \Sigma_{uo}$ consider a plant $L \subseteq \Sigma^*$ and a controller H in closed-loop configuration:



- at any time, the controller is provided $p_o s \in \Sigma_o^*$ where $s \in \Sigma^*$ is the sequence generated so far;
- in turn, the controller applies a control pattern $\gamma \in \Gamma$ of enabled events, where $\Sigma_{uc} \subseteq \gamma$;
- represent the controller as a discrete-event system

$$H \subseteq \Sigma^* .$$

$$\text{i.e. } \gamma = \{ \sigma \mid s\sigma \in H \} .$$

Def. A controller $H \subseteq \Sigma^*$ is *admissible w.r.t. the plant* $L \subseteq \Sigma^*$, if

$$[H0] H = \text{pre } H,$$

$$[H1] (\text{pre } H)\Sigma_{uc} \subseteq \text{pre } H,$$

$$[H2] \text{pre } H = p_o^{-1} p_o \text{pre } H \quad (\dots \text{ assuming } \Sigma_c \subseteq \Sigma_o),$$

[H3] $(\text{pre } L) \cap (\text{pre } H)$ does not deadlock, and

$$[H4] (\text{pre } L) \cap (\text{pre } H) = \text{pre } (L \cap H).$$

Then $K := L \cap H$ represents the closed-loop behaviour. □

Structural requirement [H4&5]:
liveness properties of the plant shall be retained.

Thm [SCT]: For a plant $L \subseteq \Sigma^*$ and an admissible controller $H \subseteq \Sigma^*$ let $K = L \cap H$. Then

[K0] K is relatively closed w.r.t. L ,

[K1] K is controllable w.r.t. L ,

[K2] K prefix-normal w.r.t. L , and

[K3] K does not deadlock.

- K is rel. closed w.r.t. L iff

$$K = (\text{pre } K) \cap L$$
- K is controllable w.r.t. L iff

$$((\text{pre } K)\Sigma_{uc}) \cap (\text{pre } L) \subseteq \text{pre } K$$
- K is prefix normal w.r.t. L iff

$$\text{pre } K = (p_o^{-1} p_o \text{pre } K) \cap (\text{pre } L)$$
- K does not deadlock iff

$$\forall s \in \text{pre } K \exists \sigma \in \Sigma: s\sigma \in \text{pre } K$$

Vice versa, if K satisfies [K0]-[K3], then there exists an admissible controller H such that $K = L \cap H$. □

Control Problem: given (L, E) with plant $L \subseteq \Sigma^*$ and a specification $E \subseteq \Sigma^*$ construct an admissible controller $H \subseteq \Sigma^*$ such that

$$K := L \cap H \subseteq E.$$

Solution: all closed-loop properties are retained under arbitrary union; thus

$$K^\dagger = \sup\{K \subseteq L \cap E \mid K \text{ satisfies [K0]–[K3]}\}$$

itself satisfies [K0]–[K3] and is used to extract a maximally permissive controller.

Note: E can be substituted by a closed language without affecting solutions – it is effectively a pure safety specification. This becomes a different story when considering ω -languages.

Discrete-Event Systems

Supervisory Control

Naive Fault-Tolerant Control

Active Fault-Tolerant Control

Post-Fault Recovery

Fault-Hiding Approach

Fault-Tolerant Control

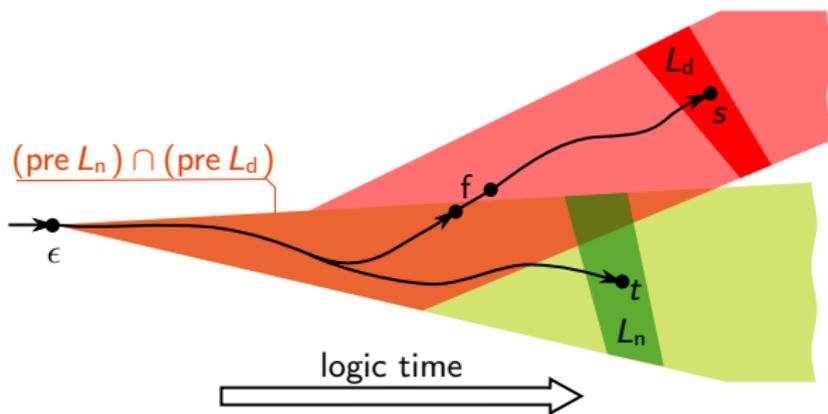
- a fault is a sudden change of behaviour
- passive approach: have a single controller that can handle pre-fault and post-fault behaviour (robust control)
- active approach: detect the fault and switch to another controller (adaptive control)

Core challenge for continuous control systems: switching of plant and controller dynamics and transient behaviour. However, for discrete-event systems:

Sudden change of behaviour and switching in the control scheme are the very nature of discrete-event systems. Hence, fault-tolerant control can be synthesised by the same methods as nominal control [??]

Naive approach to fault-tolerant control:

- nominal plant $L_n \subseteq \Sigma_n^*$
- fault event $f \notin \Sigma_n$, unctrl. and unobs., let $\Sigma_f := \Sigma_n \dot{\cup} \{f\}$
- degraded post-fault behaviour $L_d \subseteq (\text{pre } L_n)f\Sigma_f^*$
- *fault-accommodating model* $L_f := L_n \cup L_d$



Algebraic consequence:

Prop. The prerequisite $L_d \subseteq (\text{pre } L_n) \text{f} \Sigma_n^*$ implies that:

$$\begin{aligned} L_d \cap \Sigma_n^* &= \emptyset, & (\text{pre } L_d) \cap \Sigma_n^* &\subset \text{pre } L_n, \\ L_f \cap \Sigma_n^* &= L_n, & (\text{pre } L_f) \cap \Sigma_n^* &= \text{pre } L_n. \end{aligned}$$

From the last line we obtain $(\text{pre } L_f) \cap \Sigma_n^* = \text{pre } L_n = \text{pre } (L_f \cap \Sigma_n^*)$.
 I.e., the fault-accommodating model and the hypothesis the fault not to occur are non-conflicting. More general we require the fault to never become an inevitable consequence of the past event sequence.

Def. The fault-accommodating model L_f is *well-posed*, if

$$\begin{aligned} \forall s \in \text{pre } L_f \exists \sigma \in \Sigma_n : s\sigma \in \text{pre } L_f, \\ \forall s \in \text{pre } L_f \exists t \in \Sigma_n^* : st \in L_f. \end{aligned}$$

Naive approach to fault-tolerant control (cnt.)

- fault-accommodating specs. $E_f := E_n \cup E_d$, same spirit as L_f
- invoke std. synthesis procedure for (L_f, E_f)
- obtain a controller H_f with sup. closed loop $K_f^\uparrow = L_f \cap H_f$

However: we may encounter

$$\begin{aligned} \exists s \in \text{pre } K_f \quad \forall \sigma \in \Sigma_f : \quad s\sigma \in \text{pre } K_f &\Rightarrow \sigma = f, \\ \exists s \in \text{pre } K_f \quad \forall t \in \Sigma_f^* : \quad st \in K_f &\Rightarrow t \notin \Sigma_n^*, \end{aligned}$$

This is not desirable – impose additional requirements:

$$\begin{aligned} \text{[K4]} \quad &\forall s \in \text{pre } L_f \quad \exists \sigma \in \Sigma_n \quad : \quad s\sigma \in \text{pre } L_f, \\ \text{[K5]} \quad &\forall s \in \text{pre } L_f \quad \exists t \in \Sigma_n^* \quad : \quad st \in L_f. \end{aligned}$$

Naive approach to fault-tolerant control (cnt.)

- fault-accommodating specs. $E_f := E_n \cup E_d$, same spirit as L_f
- invoke synthesis procedure for (L_f, E_f) incl. [K4] and [K5]
- obtain a controller H_f with sup. closed loop $K_f^\uparrow = L_f \cap H_f$

Thm. [N-FTC]: Consider a persistent fault, $L_f \subseteq \Sigma_n^* \{ \epsilon, \mathbf{f} \} \Sigma_n^*$. Then there exists a controller H_f with $K_f = L_f \cap H_f$ that is admissible to both L_f and L_n if and only if K_f satisfies [K0]-[K5].

- diagnosability not required, passive fault-tolerant control
- in general, we have $L_n \cap H_f \subseteq K_n^\uparrow$ — may compute K_n^\uparrow and test for equality.

Discrete-Event Systems

Supervisory Control

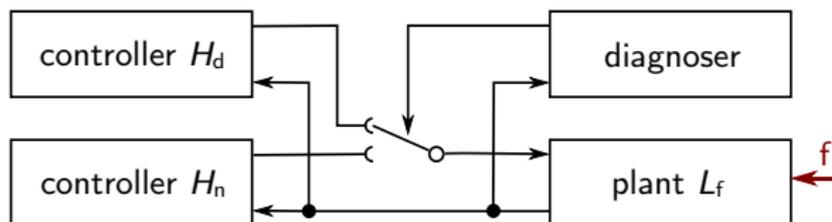
Naive Fault-Tolerant Control

Active Fault-Tolerant Control

Post-Fault Recovery

Fault-Hiding Approach

Active fault-tolerant control

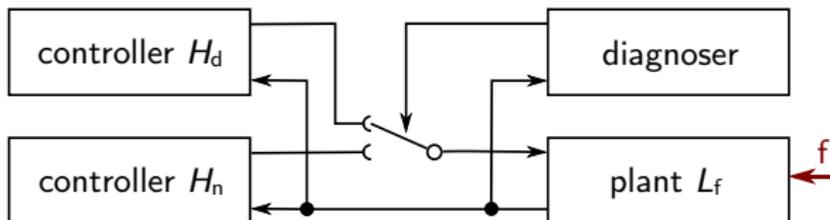


- require the fault to be diagnosable, denote $D \subseteq L_d$ the strings corresponding to f -certain diagnoser states

Diagnosis of DES (Sampath et al 1995)

- *diagnoser*: observer automaton with dedicated state labels
- *f -certain state*: state in which the fault must have occurred some time ago
- *diagnosability*: require the plant to after the fault attain an f -certain state after a bounded number of transitions.

Active fault-tolerant control



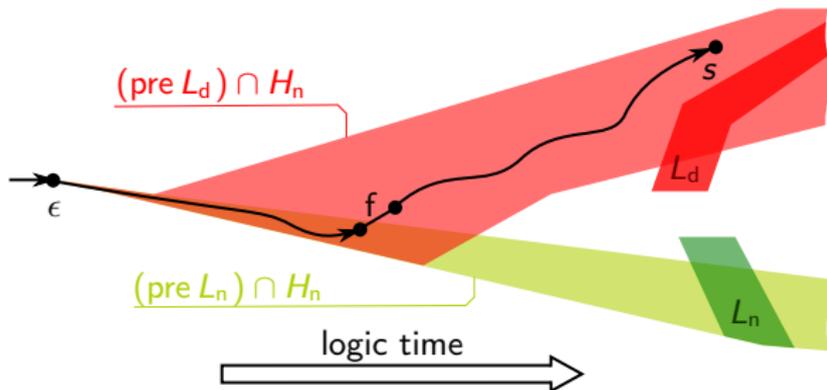
- require the fault to be diagnosable, denote $D \subseteq L_d$ the strings corresponding to f-certain diagnoser states
- require/test that the post-fault-pre-detection behaviour satisfies a safety specification (safe diagnosibility)
- design H_d to take over H_n when the plant first enters D
- note: nominal pre-fault behaviour is guaranteed
- option: synthesise H_d online once the fault has been diagnosed

- associate $H_n \leftrightarrow p_n^{-1} H_n$ and consider the local-closed loop under nominal control $K_{loc} := (\text{pre } L_f) \cap H_n$
- safe diagnosability condition:

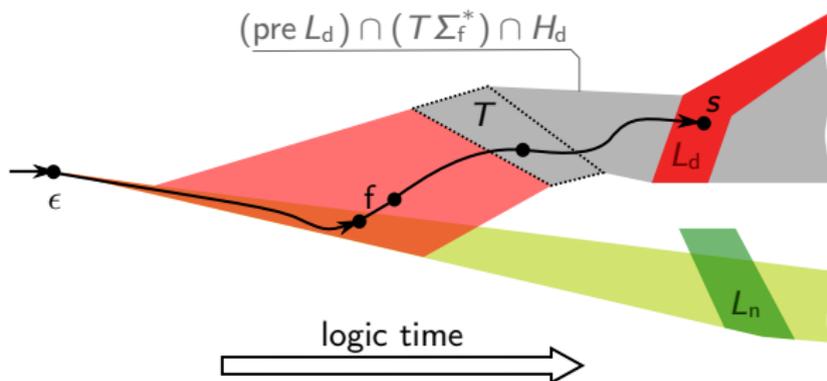
$$D := \{ s \in \Sigma_f^* \mid K_{loc} \cap (p_o^{-1} p_o s) \subseteq \Sigma_n^* f \Sigma_f^* \},$$

$$K_{loc} \cap \Sigma_n^* f \Sigma_f^k \subseteq D \text{ for some } k \in \mathbb{N},$$

$$T := \{ s \in K_{loc} \mid (\text{pre } s) \cap D = s \} \subseteq E_{\text{phi}};$$



- local post-fault-detection behaviour $(\text{pre } L_d) \cap (T\Sigma_n^*)$
- post-fault-detection controller H_d requirements:
 - [A1] admissible w.r.t. $L_d \cap (T\Sigma_n^*)$
 - [A2] enforces post fault specs. $L_d \cap (T\Sigma_n^*) \cap H_d \subseteq E_d$
 - [A3] passive before fault-detection $T \subseteq (\text{pre } L_d) \cap (T\Sigma_n^*) \cap H_d$



- local post-fault-detection behaviour $(\text{pre } L_d) \cap (T\Sigma_n^*)$
- post-fault-detection controller H_d requirements:
 - [A1] admissible w.r.t. $L_d \cap (T\Sigma_n^*)$
 - [A2] enforces post fault specs. $L_d \cap (T\Sigma_n^*) \cap H_d \subseteq E_d$
 - [A3] passive before fault-detection $T \subseteq (\text{pre } L_d) \cap (T\Sigma_n^*) \cap H_d$
- require that E_d and E_{phi} relate by

$$(\text{pre } T) \cap (\Sigma_n^* f \Sigma_n^*) \subseteq E_d \subseteq E_{\text{phi}}$$
- post-fault-detection controller H_d , synthesis:
 - for [A1] and [A2] use std. procedure on $(L_d \cap (T\Sigma_n^*), E_d)$
 - then test for [A3]
 - if test fails, no solution exists

Re-interpret active FTC as naive FTC

- formally construct overall controller H_f :

$$H_f := \{\epsilon\} \cup \{s\sigma \in H_n \mid s \notin C\} \cup \{s\sigma \in H_d \mid s \in C\},$$

$$C := (D \cap H_n \cap H_d) \cap \Sigma_f^*.$$

Thm. [A-FTC]: Given a fault-accomodating model L_f and a nominal controller H_n admissible to $L_n = L_f \cap \Sigma_n^*$, assume that local closed loop $K_{loc} := (\text{pre } L_f) \cap H_n$ is safe diagnosable. If a post-fault-detection controller H_d satisfies conditions [A1]–[A3], then the overall controller H_f defined above is admissible to both L_f and L_n with $L_n \cap H_f = L_n \cap H_n$.

- by Thm. [N-FTC] the conclusion is equivalent to K_f being a closed-loop behaviour achievable by naive FTC.

Discrete-Event Systems

Supervisory Control

Naive Fault-Tolerant Control

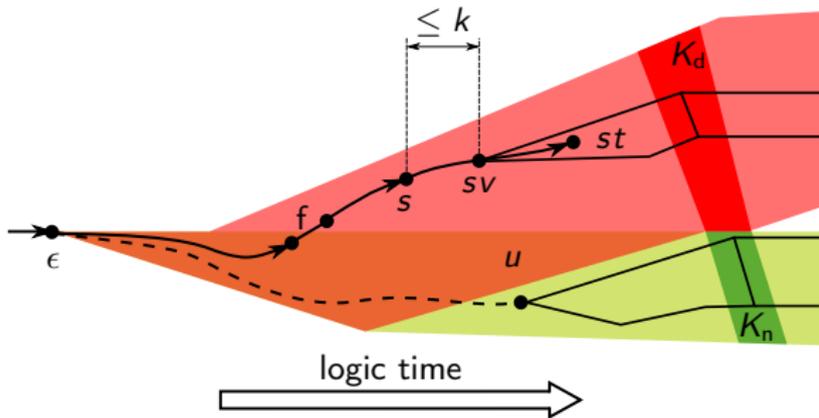
Active Fault-Tolerant Control

Post-Fault Recovery

Fault-Hiding Approach

Post-Fault Recovery: safety

Def. [K6]: A closed loop $K_f = K_n \dot{\cup} K_d$ is *weakly recovering* if there exists a uniform bound k such that for all s, t , $|t| \geq k$ with $s \in (\text{pre } K_f) \cap (\Sigma_n^* f \Sigma_n^*)$ and $st \in \text{pre } K_f$ there exists $u \in \text{pre } K_n$, $v \in \text{pre } t$, $|v| \leq k$ with $K_f / sv \subseteq K_n / u$.



Post-Fault Recovery: safety

Def. [K6]: A closed loop $K_f = K_n \dot{\cup} K_d$ is *weakly recovering* if there exists a uniform bound k such that for all s, t , $|t| \geq k$ with $s \in (\text{pre } K_f) \cap (\Sigma_n^* f \Sigma_n^*)$ and $st \in \text{pre } K_f$ there exists $u \in \text{pre } K_n$, $v \in \text{pre } t$, $|v| \leq k$ with $K_f/sv \subseteq K_n/u$.

- synthesis problem: given $L_f = L_n \cup L_d$ and E_f , compute an admissible controller H_f such that the closed loop K_f satisfies [K6].
- the property is not retained under union; synthesis procedure exists for $\Sigma_o = \Sigma$

Post-Fault Recovery: safety

- Weakly recovering [K6] implies conditional finite convergence:

$$[K6'] \quad K_n/\Sigma^* \Leftarrow K_f/(\Sigma_n^*f)$$

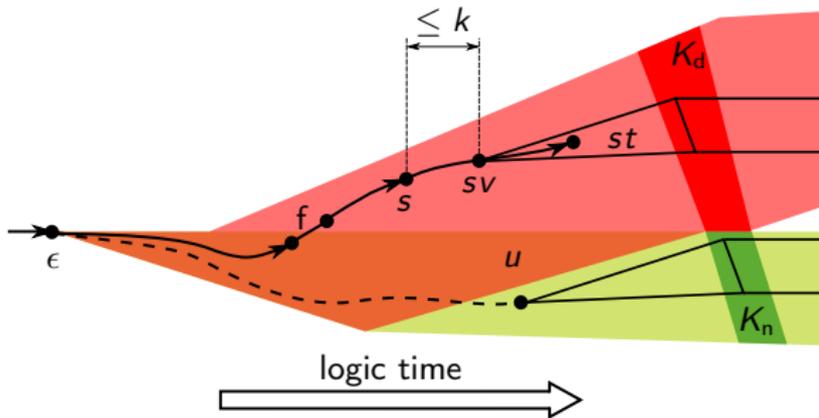
- formally generalise to:

$$[K6''] \quad E_f \Leftarrow K_f/(\Sigma_n^*f)$$

- synthesis problem: given $L_f = L_n \cup L_d$ and E_f , compute an admissible controller H_f such that the closed loop K_f satisfies $[K6']/[K6'']$.
- neither $[K6'']$ nor $[K6']$ are retained under union; synthesis procedure exists.

Post-Fault Recovery: liveness

Def. [K7]: A closed loop $K_f = K_n \dot{\cup} K_d$ is **weakly** recovering if there exists a uniform bound k such that for all $s, t, |t| \geq k$ with $s \in (\text{pre } K_f) \cap (\Sigma_n^* f \Sigma_n^*)$ and $st \in \text{pre } K_f$ there exists $u \in \text{pre } K_n, v \in \text{pre } t, |v| \leq k$ with $K_f / sv = K_n / u$.



Post-Fault Recovery: liveness

Def. [K7]: A closed loop $K_f = K_n \dot{\cup} K_d$ is **weakly recovering** if there exists a uniform bound k such that for all s, t , $|t| \geq k$ with $s \in (\text{pre } K_f) \cap (\Sigma_n^* f \Sigma_n^*)$ and $st \in \text{pre } K_f$ there exists $u \in \text{pre } K_n$, $v \in \text{pre } t$, $|v| \leq k$ with $K_f / sv = K_n / u$.

- synthesis problem: given $L_f = L_n \cup L_d$ and E_f , compute an admissible controller H_f such that the closed loop K_f satisfies [K7].
- the property is not retained under union; synthesis procedure exists for $\Sigma_o = \Sigma$

Discrete-Event Systems

Supervisory Control

Naive Fault-Tolerant Control

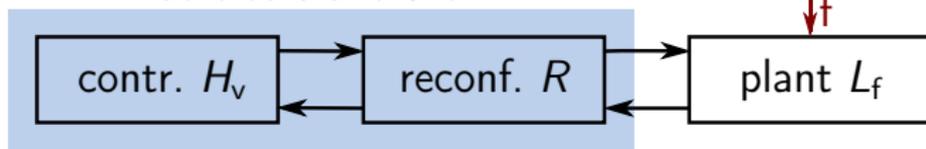
Active Fault-Tolerant Control

Post-Fault Recovery

Fault-Hiding Approach

Fault Hiding

Given $L_f = L_n \cup L_d$, $E_f = E_n \cup E_d$, and a solution H_n to (L_n, E_n)
 fault-tolerant cntrl.



- disconnect nominal controller, i.e., $H_v = h(H_n) \subseteq \Sigma_v^*$ with $\Sigma_v \cap \Sigma_f = \emptyset$, h bijective and applied per event.
- synthesise reconfiguration dynamics $R \subseteq (\Sigma_v \cup \Sigma_o)^*$ to re-connect
- do so by interpreting $H_v \parallel L_f$ as plant and use std. procedures on adapted language inclusion specification, extract R from K
- obtain an overall fault-tolerant controller from H_v and R

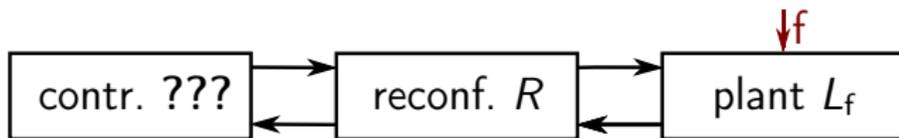
- when using a minimal restrictive solution H_n^\uparrow resp. $H_v^\uparrow \parallel L_f$ for the design, and if the closed loop K^\uparrow satisfies [K0]-[K3] and additionally

$$\begin{aligned}
 \text{[K8]} \quad & (\forall s \in \text{pre } K) [((p_v s)h(\Sigma_{uc})) \cap (\text{pre } h(L_n)) \neq \emptyset \\
 & \Rightarrow s(\Sigma - h(\Sigma_c))^* h(\Sigma_{uc}) \cap (\text{pre } K) \neq \emptyset]
 \end{aligned}$$

then the corresponding R is admissible to $H_v \parallel L_f$ for any nominal controller H_n that solves (L_n, E_n) .

- [K8] is retained under union, synthesis procedures are available.

Note: Nominal controller does not need to be known.



Discrete-Event Systems

Supervisory Control

Naive Fault-Tolerant Control

Active Fault-Tolerant Control

Post-Fault Recovery

Fault-Hiding Approach

Summary

Fault-tolerant supervisory control is addressed by the recent literature in various ways, including passive and active approaches, post-fault recovery and fault-hiding.

Conclusions

- switching is addressed by the common modelling framework — any method for fault-tolerant supervisory control should be interpretable within this framework
- additional features of individual approaches amount to additional closed-loop properties — and novel to synthesis problems
- insisting in uniform bounds for diagnosability and language convergence may be too strict for particular applications — discussion in terms of ω -languages may turn out beneficial

References

- Blanke, M., Kinnaert, M., Lunze, J., Staroswiecki, M., Schröder, J. (2006). *Diagnosis and Fault-Tolerant Control*. Springer.
- Paoli, A. and Lafortune, S. (2005). Safe diagnosability for fault-tolerant supervision of discrete-event systems. *Automatica*, 41(8), 1335–1347.
- Paoli, A., Sartini, M., and Lafortune, S. (2008). A fault tolerant architecture for supervisory control of discrete event systems. *Proceedings of the 17th IFAC world congress*, 6542–6547.
- Paoli, A., Sartini, M., and Lafortune, S. (2011). Active fault tolerant control of discrete event systems using online diagnostics. *Automatica*, 47(4), 639–649.
- Sülek, A.N. and Schmidt, K.W. (2014). Computation of supervisors for fault-recovery and repair for discrete event systems. *WODES*, 428–438.
- Watanabe, T.Y, Leal, A.B, Cury, J.E.R, Queiroz, M.H.de. (2017) Safe controllability using online prognosis, *IFAC WC 2017*.
- Wen, Q., Kumar, R., and Huang, J. (2014). Framework for optimal fault-tolerant control synthesis: maximize pre-fault while minimize post-fault behaviors for discrete event systems. *IEEE Trans. Syst. Man Cybern. Syst.*, 44, 1056–1066.
- Wen, Q., Kumar, R., Huang, J., and Liu, H. (2008). A framework for fault-tolerant control for discrete event systems. *IEEE TAC*, 53, 1839–1849.
- Wittmann, T., Richter, J., and Moor, T. (2013). Fault-hiding control reconfiguration for a class of discrete-event systems. *IFAC DCDS*.
- Wittmann, T., Richter, J., and Moor, T. (2012). Fault-tolerant control of discrete event systems based on fault-accommodating models. *SAFEPROCESS*, 854–859.

More references given in

- Moor, T. (2016). A discussion of fault-tolerant supervisory control in terms of formal languages. *Annual Reviews in Control*, 159-169.
- Schmuck, A.-K., Moor, T., Majumdar, (2020). On the relation between reactive synthesis and supervisory control of non-terminating processes, *Discrete Event Dynamic Systems*, 30, 81–124.