The Controllability Prefix for Supervisory Control under Partial Observation with an Application to Fault-Tolerant Control

Thomas Moor [*] **Klaus Werner Schmidt** [**]

[*] *Lehrstuhl für Regelungstechnik*
Friedrich-Alexander Universität Erlangen-Nürnberg, Germany
(e-mail: lrt@fau.de)
[**] *Mechatronics Engineering Department*
Çankaya University, Ankara, Turkey
(e-mail: schmidt@cankaya.edu.tr)

Abstract: The *controllability prefix* is known as a useful concept for the discussion and solution of synthesis problems in supervisory control of ω-languages, i.e., formal languages of infinite-length words. There, the controllability prefix is defined as the set of all finite-length prefixes that can be controlled to satisfy prescribed liveness and safety properties. In this paper, we discuss a variation of the controllability prefix to address supervisory control under partial observation for regular $*$-languages, i.e., formal languages of finite-length words. We derive algebraic properties that are useful for a quantitative analysis on how an upper-bound language-inclusion specification affects achievable lower-bound specifications. Our study is motivated by the synthesis of fault-tolerant supervisory controllers, where the possible occurrence of a fault may restrict the achievable pre-fault behaviour so severe, that a relaxation of the upper-bound specification becomes a practical option. As our study shows, such a relaxation can be systematically constructed in terms of the controllability prefix.

Keywords: Discrete-event systems, supervisory control, fault-tolerant control, partial observation.

INTRODUCTION

Given the behaviour of a plant and an upper-bound specification, both in terms of formal languages, Thistle and Wonham (1994b) define the *controllability prefix* as the set of strings, from which on a supervisor can take over the plant in order to enforce the specification. Interpreting control as a game between the supervisor and the plant, the controllability prefix characterises the *winning configurations* for the supervisor. Thus, if the empty string is within the controllability prefix, the supervisor wins and the control problem has a solution. In their study, Thistle and Wonham (1994b) address languages of infinite-length words also known as *ω-languages* or *sequential behaviours*. In this setting, a supremal achievable closed-loop behaviour in general fails to exist, however, the authors establish a tight upper bound in terms of the controllability prefix. Together with the computational procedure provided in (Thistle and Wonham, 1994a), this effectively solves the synthesis problem for supervisory control of ω-languages.

In this paper, we address supervisory control under partial observation where the plant and the upper-bound specification are represented as $*$-languages and under the assumption that all controllable events are observable. For this situation, Lin and Wonham (1988) establish the unique existence of the supremal achievable closed-loop behaviour, characterised by controllability, prefix-normality and relative closedness. For practical purposes, it is common to compute the supremal closed-loop behaviour and to test whether it also satisfies an additional lower-bound specification. If the test passes, the supremal closed-loop behaviour is used to extract the supervisor. If the test fails, no acceptable solution exists. It is the latter case, for which we propose to consider the controllability prefix as a guidance on how the upper-bound specification can be relaxed in order to meet the lower bound. More specifically, we characterise those configurations, in which the supervisor may enable additional events in favour of the lower bound while risking to fail on the original upper bound, but to do so only if it is known by observation that there is still the chance to win. Technically, the present paper is a further development of (Moor and Schmidt, 2015) to account for partial observation.

While we believe that the concept of the controllability prefix is of general interest, our study is motivated by a specific problem of fault-tolerant supervisory control; see e.g. (Paoli and Lafortune, 2005; Wittmann et al., 2012; Wen et al., 2014; Sülek and Schmidt, 2014; Acar and Schmidt, 2015), with an overview given in (Moor, 2016). In this setting, a nominal plant and a nominal specification are extended to model the effect of a fault. The task is then to design a control scheme that initially enforces the nominal specification and that, after the occurrence of a fault, continues to faithfully operate the plant. A first step for a practical solution is to compare the supremal closed-loop behaviours obtained for the nominal plant and the extended plant, both for the nominal specification. Here, it typically shows that considering the possible occurrence of the fault, the supervisor turns out more restrictive even in the pre-fault behaviour. Thus, there is a trade-off between maintaining pre-fault performance and allowance for post-fault degradation. This trade-off can be analysed in terms of the controllability prefix, making explicit when a supervisor may risk to fail on the nominal specification while maintaining the chance to win.

The paper is organised as follows. Section 1 provides common notation and a concise review of supervisory control under partial observation. A motivating example from fault-tolerant control is presented in Section 2. As our main contribution, we formally define the controllability prefix in Section 3 and elaborate algebraic properties. The latter are used in Section 4 for a systematic relaxation of a given upper-bound specification. Turning back to the example, Section 5 demonstrates how our results can be utilised in the context of fault-tolerant control.

1. PRELIMINARIES

We give a summary of common notation and elementary facts, and recall essential results regarding supervisory control under partial observation, as relevant for the present paper.

1.1 Notation

Let Σ be a *finite alphabet*, i.e., a finite set of symbols $\sigma \in \Sigma$. The *Kleene-closure* Σ^* is the set of finite strings $s = \sigma_1\sigma_2\cdots\sigma_n$, $n \in \mathbb{N}$, $\sigma_i \in \Sigma$, and the *empty string* $\epsilon \in \Sigma^*$, $\epsilon \notin \Sigma$. If, for two strings $s, r \in \Sigma^*$, there exists $t \in \Sigma^*$ such that $s = rt$, we say r is a *prefix* of s, and write $r \leq s$; if in addition $r \neq s$, we say r is a *strict prefix* of s and write $r < s$.

A $*$-*language* (or short a *language*) over Σ is a subset $L \subseteq \Sigma^*$. The *prefix* of a language $L \subseteq \Sigma^*$ is defined by $\operatorname{pre} L := \{r \in \Sigma^* \mid \exists\, s \in L : r \leq s\}$. The prefix operator is also referred to as the *prefix-closure* (or short *closure*), and, a language L is *closed* if $L = \operatorname{pre} L$. A language K is *relatively closed w.r.t. L* if $K = (\operatorname{pre} K) \cap L$. The prefix operator distributes over arbitrary unions of languages. However, for the intersection of two languages L and K, we have $\operatorname{pre}(L \cap K) \subseteq (\operatorname{pre} L) \cap (\operatorname{pre} K)$. If equality holds, L and K are said to be *non-conflicting*.

For the *observable events* $\Sigma_o \subseteq \Sigma$, the *natural projection* $\mathrm{p}_o\colon \Sigma^* \to \Sigma_o^*$ is defined iteratively: (1) let $\mathrm{p}_o\epsilon := \epsilon$; (2) for $s \in \Sigma^*$, $\sigma \in \Sigma$, let $\mathrm{p}_o(s\sigma) := (\mathrm{p}_o s)\sigma$ if $\sigma \in \Sigma_o$, or, if $\sigma \notin \Sigma_o$, let $\mathrm{p}_o(s\sigma) := \mathrm{p}_o s$. The set-valued inverse p_o^{-1} of p_o is defined by $\mathrm{p}_o^{-1}(r) := \{s \in \Sigma^* \mid \mathrm{p}_o(s) = r\}$ for $r \in \Sigma_o^*$. When applied to languages, the projection distributes over unions, and the inverse projection distributes over unions and intersections. The prefix operator commutes with projection and inverse projection.

Given two languages L, $K \subseteq \Sigma^*$, and a set of *uncontrollable events* $\Sigma_{uc} \subseteq \Sigma$, we say K is *controllable w.r.t. L*, if $(\operatorname{pre} K)\Sigma_{uc} \cap (\operatorname{pre} L) \subseteq \operatorname{pre} K$. With the set of observable events $\Sigma_o \subseteq \Sigma$, we say that K is *normal w.r.t. L*, if $K = (\mathrm{p}_o^{-1}\mathrm{p}_o K) \cap L$. Furthermore, K is *prefix-normal w.r.t. L*, if $\operatorname{pre} K = (\mathrm{p}_o^{-1}\mathrm{p}_o \operatorname{pre} K) \cap (\operatorname{pre} L)$. If K is relatively closed and prefix-normal w.r.t. L, then K is also normal w.r.t. L. Controllability, normality, prefix-normality, closedness and relative closedness are each retained under arbitrary union; see (Ramadge and Wonham, 1987) regarding controllability, and (Lin and Wonham, 1988) regarding normality. Note that closedness and relative closedness are also retained under arbitrary intersection.

1.2 Supervisory control

We refer to supervisory control as originally introduced by Ramadge and Wonham (1987), and extended to address partial observation by Lin and Wonham (1988).

For a given alphabet Σ, consider a formal language $L \subseteq \Sigma^*$ to represent the *plant behaviour*. The prefix $\operatorname{pre} L$ is also referred to as the *local behaviour*. It is the set of all event sequences that can be generated by the physical plant while time passes. In contrast, the *accepted behaviour* L is commonly used to indicate task completion.

For the purpose of control, the *common partition* $\Sigma = \Sigma_c \dot\cup \Sigma_{uc} = \Sigma_o \dot\cup \Sigma_{uo}$ is used to distinguish *controllable events*, *uncontrollable events*, *observable events* and *unobservable events*, respectively, where we assume $\Sigma_c \subseteq \Sigma_o$ throughout this paper. The supervisor is then represented as a *causal feedback map* $f\colon \operatorname{pre} L \to \Gamma$ with the set of *control-patterns* $\Gamma := \{\gamma \mid \Sigma_{uc} \subseteq \gamma \subseteq \Sigma\}$ and with $f(s)$ the events enabled after the occurrence of $s \in \Sigma^*$. For partial observation, the feedback map must satisfy the *observability condition* $f(s') = f(s'')$ for all s', $s'' \in \operatorname{pre} L$ with $\mathrm{p}_o s' = \mathrm{p}_o s''$. The *local closed-loop behaviour* $K_{\mathrm{loc}} \subseteq \Sigma^*$ is obtained by restricting $\operatorname{pre} L$ according to the feedback map f, and, for *non-blocking supervision*, one requires that K_{loc} and L are non-conflicting. Finally, the *accepted closed-loop behaviour* is defined $K := K_{\mathrm{loc}} \cap L$.

The commonly studied control problem is parameterised by a plant $L \subseteq \Sigma^*$ together with lower- and upper *language inclusion specifications* $A \subseteq \Sigma^*$ and $E \subseteq \Sigma^*$, respectively. A solution to the control problem is a non-blocking supervisor f that satisfies the observability condition and that operates the plant with an accepted closed-loop behaviour K satisfying the prescribed bounds, i.e.,

$$A \subseteq K \subseteq E\,. \tag{1}$$

Without loss of generality, we may assume that $\emptyset \neq A \subseteq E \subseteq L \subseteq \Sigma^*$.

Except for cosmetic differences in notation, a constructive solution to the above problem has been presented by Lin and Wonham (1988). It is based on two technical results. First, it is observed that a language $K \neq \emptyset$ can be obtained as a closed-loop behaviour if and only if K is controllable, prefix-normal and relatively closed w.r.t. L. Second, controllability, prefix-normality and relative closedness are retained under arbitrary union. Thus, denoting the set of all achievable closed-loop behaviours that satisfy the upper bound as

$$\begin{aligned}\mathsf{CNF}(L,E) := \{K \subseteq E \mid\ & \\ & K \text{ is controllable w.r.t. } L\,,\\ & K \text{ is prefix-normal w.r.t. } L\,,\\ & K \text{ is relatively closed w.r.t. } L\,\}, \end{aligned} \tag{2}$$

it holds that the supremum

$$K^{\uparrow} := \sup \mathsf{CNF}(L,E) := \cup\{K \subseteq \Sigma^* \mid K \in \mathsf{CNF}(L,E)\,\} \tag{3}$$

is itself an achievable closed-loop behaviour with $K^{\uparrow} \subseteq E$. For regular parameters E and L, the supremum $K^{\uparrow}$ is also regular and an automaton representation can be obtained by well known algorithms; see e.g. Cho and Marcus (1989) or the more recent variations given by Moor et al. (2012) and by Cai et al. (2015).

If $K^{\uparrow}$ happens to also satisfy the lower-bound specification $A \subseteq K^{\uparrow}$, then a feedback map f to solve the control problem can be extracted from $K^{\uparrow}$. If, on the other hand, $K^{\uparrow}$ does not satisfy the inclusion $A \subseteq K^{\uparrow}$, then neither does any other achievable closed-loop behaviour. Then, the control problem has no solution.

2. FAULT-TOLERANT SUPERVISORY CONTROL

We provide a simple example that illustrates the results presented so far in the context of fault-tolerant control. More specifically, we follow the *naive approach* proposed by Moor (2016).

Consider a processing machine with the physical behaviour realised by the automaton in Fig. 1. Referring to the additional external events A, B and X, the task is to design a controller that accepts the commands A and B to select a particular processing scheme and to provide feedback X upon completion. Here, all events are considered observable, and, except a, b and t, also controllable. The formal specification is given by the two automata in Fig. 2, where the overall specification E is obtained by parallel composition.

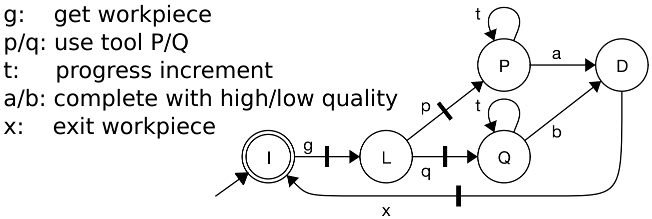

Fig. 1. plant behaviour L

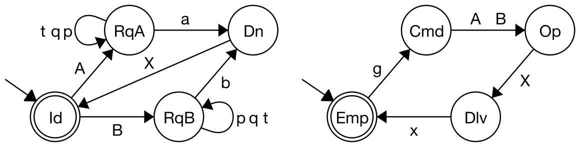

Fig. 2. upper-bound specifications E_1 (left) and E_2 (right)

The synthesis procedure then needs to figure when to accept commands A or B, which tool to choose in order to achieve the requested quality, and when to provide acknowledgment X. A realisation of $K^{\uparrow} = \sup \mathsf{CNF}(L, E)$, as shown in Fig. 3, is readily obtained by available software tools.

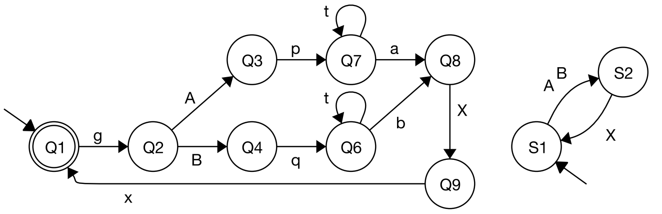

Fig. 3. closed loop $K^{\uparrow}$ (left) with external behaviour L_{hi} (right)

We now assume that the processing tool P is subject to wear-out in that it degrades to only produce low quality output, indicated by b. The transition from normal operation to wear-out can be modelled by the distinguished fault event f, to obtain the overall *fault-accommodating model* $L_\mathrm{f} \subseteq \Sigma_\mathrm{f}^*$, $\Sigma_\mathrm{f} := \Sigma\dot{\cup}(\mathrm{f})$; see Fig. 4.

The fault event f is considered uncontrollable, and, in contrast to the predecessor paper (Moor and Schmidt, 2015), also as unobservable. For a first step in a fault-tolerant design, we lift the nominal specification E to Σ_f, i.e., we consider the specification $E_\mathrm{f} := \mathrm{p}_\mathrm{f}^{-1}E$, with $\mathrm{p}_\mathrm{f}^{-1}$ the set-valued inverse of the natural projection $\mathrm{p}_\mathrm{f} : \Sigma_\mathrm{f}^* \to \Sigma^*$. However, the closed-loop behaviour $K_\mathrm{f}^{\uparrow} := \sup \mathsf{CNF}(L_\mathrm{f}, E_\mathrm{f})$ is considered inappropriate: in order to unconditionally satisfy E_f, the supervisor is required never to use tool P. For practical reasons, we would prefer the supervisor to optimistically risk the use of tool P until the fault actually occurs. Since $K_\mathrm{f}^{\uparrow}$ is supremal, this effectively requires us to relax E_f. In the following two sections we will develop a method to further analyse this situation.

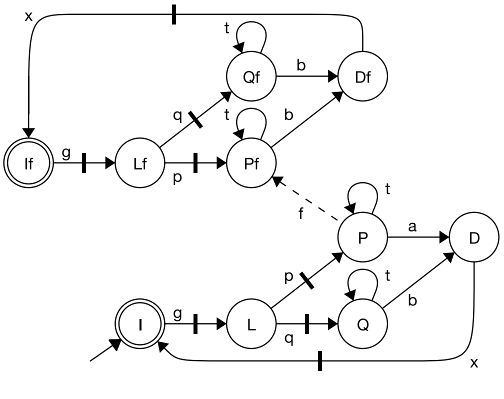

Fig. 4. fault-accommodating model L_f

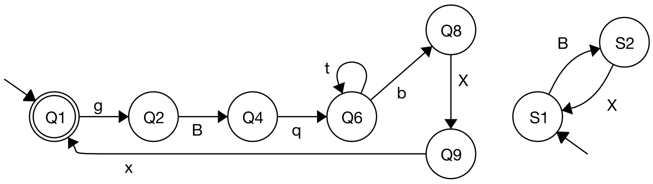

Fig. 5. Fault tolerant $K_\mathrm{f}^{\uparrow}$ (left) with ext. behaviour L_{hif} (right)

3. CONTROLLABILITY PREFIX

We provide a quantitative analysis of the supremal closed-loop behaviour for an upper-bound inclusion specification. Our discussion uses the notion of the *controllability prefix* as proposed by Thistle and Wonham (1994b) in the context of ω-languages, which we adapt to the problem at hand, i.e., the supervision of $*$-languages under partial observation.

Given a plant $L \subseteq \Sigma^*$ and an upper-bound specification $E \subseteq L$, consider a sequence $s \in \operatorname{pre} L$ from the local plant behaviour and define

$$L_s := L \cap (\mathrm{p}_\mathrm{o}^{-1}\mathrm{p}_\mathrm{o} s\Sigma^*)\,, \qquad (4)$$

$$E_s := E \cap (\mathrm{p}_\mathrm{o}^{-1}\mathrm{p}_\mathrm{o} s\Sigma^*)\,, \qquad (5)$$

$$K_s^{\uparrow} := \sup \mathsf{CNF}(L_s, E_s)\,. \qquad (6)$$

The plant L_s and the upper-bound specification E_s are obtained form the original parameters L and E, respectively, under the additional assumption that the plant will start by tracking a string that is observed as $\mathrm{p}_\mathrm{o}s$. If pre $K_s^{\uparrow}$ accounts for every such string, i.e., if

$$(\operatorname{pre} L) \cap (\mathrm{p}_\mathrm{o}^{-1}\mathrm{p}_\mathrm{o} s) \subseteq \operatorname{pre} K_s^{\uparrow}\,, \qquad (7)$$

then there exists a supervisor that can take over the plant after the observation $\mathrm{p}_\mathrm{o}s$ in order to enforce the upper-bound specification E. By the below proposition, Eq. (7) is equivalent to $K_s^{\uparrow} \neq \emptyset$ and we formally define controllability prefix for the control problem at hand as follows.

Definition 1. Given two languages $E \subseteq L \subseteq \Sigma^*$, the *controllability prefix of E w.r.t. L* is defined by

$$T := \{s \in \operatorname{pre} L \,|\, K_s^{\uparrow} \neq \emptyset\}\,, \qquad (8)$$

referring to Eqs. (4), (5) and (6), and to the common partition $\Sigma = \Sigma_\mathrm{c}\dot{\cup}\Sigma_\mathrm{uc} = \Sigma_\mathrm{o}\dot{\cup}\Sigma_\mathrm{uo}$ with $\Sigma_\mathrm{c} \subseteq \Sigma_\mathrm{o}$. □

From a game-theoretic perspective, T is seen as the set of *winning configurations* from which, once attained, a supervisor can "win" in the sense of being able to enforce E when "playing against the plant".

Proposition 2. Given $E \subseteq L \subseteq \Sigma^*$, denote the controllability prefix T. For any $s \in \operatorname{pre} L$ consider L_s, E_s and $K_s^{\uparrow}$ defined by Eqs. (4), (5) and (6), respectively. Then

(i) $s \in \operatorname{pre} K_s^{\uparrow} \quad \Leftrightarrow \quad K_s^{\uparrow} \neq \emptyset$,

(ii) $(\operatorname{pre} L) \cap (\mathrm{p_o^{-1}p_o} s) \subseteq \operatorname{pre} K_s^{\uparrow} \quad \Leftrightarrow \quad K_s^{\uparrow} \neq \emptyset$,

(iii) T is normal w.r.t. $\operatorname{pre} L$.

Proof (outline) [1] **.** All three claims follow by elementary considerations and the respective definitions. At (i). "$\Rightarrow$" is trivial. For "$\Leftarrow$", pick $t \in K_s^{\uparrow} \subseteq L_s \subseteq (\mathrm{p_o^{-1}p_o} s\Sigma^*)$, decompose with $r \leq t$, $\mathrm{p_o} r = \mathrm{p_o} s$, and refer to prefix-normality of $K_s^{\uparrow}$ to obtain $s \in \operatorname{pre} K_s^{\uparrow}$. At (ii). "$\Rightarrow$" is trivial. For "$\Leftarrow$" pick any $s' \in \operatorname{pre} L \cap (\mathrm{p_o^{-1}p_o} s)$, to observe that $s' \in \operatorname{pre} L_s$ and, by (i), $s' \in \mathrm{p_o^{-1}p_o} \operatorname{pre} K_s^{\uparrow}$. Then prefix-normality of $K_s^{\uparrow}$ implies $s' \in \operatorname{pre} K_s^{\uparrow}$. Ad (iii). Pick $s \in T$ and $s' \in \operatorname{pre} L$ such that $\mathrm{p_o} s = \mathrm{p_o} s'$. Observe $K_{s'}^{\uparrow} = K_s^{\uparrow} \neq \emptyset$, and, hence, $s' \in T$. □

Intuitively, once the plant attains a winning configuration $s \in T$, it can be controlled to maintain this status, i.e., to continue to evolve within T.

Proposition 3. Given $E \subseteq L \subseteq \Sigma^*$, denote the controllability prefix T. For any $s \in \operatorname{pre} L$ consider L_s, E_s and $K_s^{\uparrow}$ defined by Eqs. (4), (5) and (6), respectively. Then for any $t \in \Sigma^*$

$$s \leq t \in \operatorname{pre} K_s^{\uparrow} \quad \Rightarrow \quad t \in T\,. \tag{9}$$

Proof (outline). Consider the candidate

$$K := K_s^{\uparrow} \cap (\mathrm{p_o^{-1}p_o} t\Sigma^*)\,, \tag{10}$$

i.e., the restriction of $K_s^{\uparrow}$ to strings begin with the observation $\mathrm{p_o} t$. By the prerequisite we have $t \in \operatorname{pre} K$ and $K \subseteq E_t := E \cap (\mathrm{p_o^{-1}p_o} t\Sigma^*)$. As it turns out, K can be verified to be relatively closed, controllable and prefix-normal w.r.t. $L_t := L \cap (\mathrm{p_o^{-1}p_o} t\Sigma^*)$. Hence, $\emptyset \neq K \in \mathsf{CNF}(L_t, E_t)$, with $K_t^{\uparrow} \neq \emptyset$ as immediate consequence. This implies $t \in T$. □

As a technical consequence of the above proposition, we note that L and T are non-conflicting.

Lemma 4. Let $E \subseteq L \subseteq \Sigma^*$. Then L and the controllability prefix T of E w.r.t. L are non-conflicting. In particular, we have $\operatorname{pre} T = \operatorname{pre}(L \cap T)$.

Proof (outline). Pick any $s \in (\operatorname{pre} L) \cap (\operatorname{pre} T)$ and extend it by t such that $st \in T$. This implies, by Proposition 2, part (i), $st \in \operatorname{pre} K_{st}^{\uparrow}$. Thus, we can further extent st by w such that $stw \in K_{st}^{\uparrow} \subseteq L$ and conclude by Proposition 3 that $stw \in T$. Hence, $s \in \operatorname{pre}(L \cap T)$. □

Any admissible controller that enforces the upper-bound specification E must control the plant to evolve within T. However, since T may not be prefix-closed, there may exist winning configurations that do not contribute to the supremal closed-loop behaviour. The following proposition manifests this intuition.

Proposition 5. Given $E \subseteq L \subseteq \Sigma^*$, denote the controllability prefix T and let $K^{\uparrow} := \sup \mathsf{CNF}(L, E)$. Then

$$\operatorname{pre} K^{\uparrow} \subseteq T\,, \tag{11}$$

where equality holds if and only if T is closed.

Proof (outline). The inclusion (11) is an immediate consequence of Proposition 3, where we consider the special case of $s = \epsilon$ with $K_{\epsilon}^{\uparrow} = K^{\uparrow}$. Moreover, if equality holds in (11), closedness of T is trivial. For the converse implication, T is assumed to be closed and $K := L \cap T$ is considered a closed loop candidate. Indeed, each of the properties relative closedness, controllability and prefix normality of K w.r.t. L can be verified. Now pick any $s \in K$ to observe $s \in L_s \cap \operatorname{pre} K_s^{\uparrow} = K_s^{\uparrow} \subseteq E_s \subseteq E$. Hence, $K \subseteq E$ and we note that $K \in \mathsf{CNF}(L, E)$. To obtain equality in (11), we refer to Lemma 4 and observe that $T = \operatorname{pre} T \cap \operatorname{pre} L = \operatorname{pre}(L \cap T) = \operatorname{pre} K \subseteq \operatorname{pre} K^{\uparrow}$. □

Thus, if T is closed, it solves the control problem with accepted closed-loop behaviour $K^{\uparrow}$ and provides no further insight. More interestingly, if T fails to be closed, there must exist sequences $s \notin \operatorname{pre} K^{\uparrow}$ but $s \in T$, i.e., even under maximally permissive control, the closed-loop is constrained not to attain certain winning configurations. However, by the following proposition, $s \notin \operatorname{pre} K^{\uparrow}$ can only be the case if s exits $\operatorname{pre} K^{\uparrow}$ via a non-winning configuration.

Proposition 6. Given $E \subseteq L \subseteq \Sigma^*$, denote the controllability prefix T and let $K^{\uparrow} := \sup \mathsf{CNF}(L, E)$. Then, and for any $s \in \Sigma^*$, $\sigma \in \Sigma$, we have

$$s \in \operatorname{pre} K^{\uparrow},\ s\sigma \notin \operatorname{pre} K^{\uparrow} \quad \Rightarrow \quad s\sigma \notin T\,. \tag{12}$$

Proof (outline). It can be verified that for any $K \in \mathsf{CNF}(L, E)$, $s \in \operatorname{pre} K$, $\sigma \in \Sigma$, $K' \in \mathsf{CNF}(L_{s\sigma}, E_{s\sigma})$, controllers can be merged in the sense of $K \cup K' \in \mathsf{CNF}(L, E)$. The claim is then proven by contradiction, i.e., we assume that we can pick $s \in \Sigma^*$ and $\sigma \in \Sigma$ such that $s \in \operatorname{pre} K^{\uparrow}$, $s\sigma \notin \operatorname{pre} K^{\uparrow}$, and $s\sigma \in T$. By Proposition 2, this implies $s\sigma \in \operatorname{pre} K_{s\sigma}^{\uparrow}$ with $K_{s\sigma}^{\uparrow} := \sup \mathsf{CNF}(L_{s\sigma}, E_{s\sigma})$. By our preliminary consideration, we obtain $K^{\uparrow} \cup K_{s\sigma}^{\uparrow} \in \mathsf{CNF}(L, E)$. Then supremality of $K^{\uparrow}$ implies $K_{s\sigma}^{\uparrow} \subseteq K^{\uparrow}$ and, hence, $s\sigma \in \operatorname{pre} K^{\uparrow}$. This contradicts the choice of s and σ. □

4. RELAXING THE UPPER BOUND

Insisting on the upper-bound specification E implicitly imposes a restriction on the achievable lower bound, namely $A \subseteq K^{\uparrow} := \mathsf{CNF}(L, E)$, which by the analysis of Section 3 implies

$$\epsilon \in \operatorname{pre} A \subseteq T\,, \tag{13}$$

with T the controllability prefix of E w.r.t. L. However, for some applications the imposed restriction may not be acceptable; see the example in Section 2.

Note that the dual approach, namely to consider the infimum $K^{\downarrow}$ over all achievable closed-loop behaviours

$$K^{\downarrow} := \cap\{K \subseteq \Sigma^* \,|\, A \subseteq K,\ K \in \mathsf{CNF}(L, L)\} \tag{14}$$

and thereby giving preference to the lower bound A, in general fails to solve the problem. This is because only in the special case when L is closed, $K^{\downarrow}$ turns out controllable and prefix-normal; see also (Lafortune and Chen, 1990). In the following discussion, we therefore propose an alternative approach that addresses not-necessarily closed plants L and that is based on relaxing the interpretation of the upper-bound parameter E. While in the original setting, the local closed-loop behaviour must consist of winning configurations only, we consider the weaker requirement of the chance to attain a winning configuration.

To this end, we define the class of languages

[1] Observing page limitations, we only give proof outlines. A technical report including full proofs is available from the authors upon request.

$$\begin{aligned}\mathsf{MN}(L,T) := \{K \subseteq \operatorname{pre} T \mid &\\ &\text{(a)}\quad (\forall\, s\in K\ \exists\, t\in\Sigma^*)[\, st\in T,\ s \operatorname{pre} t\subseteq K\,]\,,\ \text{and}\\ &\text{(b)}\quad K \text{ is normal w.r.t. } \operatorname{pre} L\ \}\,. \end{aligned} \tag{15}$$

Condition (a) ensures the chance to attain a winning configuration while remaining within K. The latter clause in (a), to remain within K, together with normality condition (b) ensures that there persistently is the chance not only to attain a winning configuration but also that this status is known by observation. We provide some technical properties of the class $\mathsf{MN}(L,T)$.

Proposition 7. Given $E \subseteq L \subseteq \Sigma^*$ and the controllability prefix T of E w.r.t. L, consider $\mathsf{MN}(L,T)$ as defined by (15) and the supremum $M := \sup \mathsf{MN}(L,T) := \cup\{K \mid K \in \mathsf{MN}(L,T)\,\}$. Then

(i) $T \in \mathsf{MN}(L,T)$,

(ii) $M \in \mathsf{MN}(L,T)$, and

(iii) $(\forall\, s\in M\ \exists\, t\in\Sigma^*)[\, s \operatorname{pre} t\subseteq M,\ st\in L\,]$, and

(iv) M and L are non-conflicting.

Proof (outline). Ad (i). For (a), pick $s \in T$ and refer to Proposition 2, part (i), to obtain $s \in \operatorname{pre} K_s^{\uparrow}$. Then, extend s by t such that $st \in K_s^{\uparrow}$ and refer to Proposition 3 to obtain $s \operatorname{pre} t \subseteq T$. Regarding (b), we refer to Proposition 2, part (iii). Ad (ii). It is well known that normality is retained under arbitrary union. This is also readily verified for the condition in part (a) of (15). Ad (iii). Pick any $s \in M$. By (ii) we have $M \in \mathsf{MN}(L,T)$ and can therefore extend s by t such that $st \in T$ and $s \operatorname{pre} t \subseteq M$. By Proposition 2, part (i), we can further extend by u such that $stu \in K_{st}^{\uparrow} \subseteq L$. Referring to Proposition 3, this implies $st \operatorname{pre} u \subseteq T$ and, by (i), $st \operatorname{pre} u \subseteq T \subseteq M$. Ad (iv). Pick $s \in (\operatorname{pre} L) \cap (\operatorname{pre} M)$, and extend s by v such that $sv \in M$. Refer to (iii) to obtain t such that $sv \operatorname{pre} t \subseteq M$ and $svt \in L$. □

Referring to $T \subseteq M$ implied by (i) and (ii) above, we propose

$$M := \sup \mathsf{MN}(L,T) \tag{16}$$

as an optimistic variant of the controllability prefix T and ask for the controller to maintain containment in $\operatorname{pre} M$ whenever possible. This is expressed by the following construction of the relaxed upper bound E'.

$$\begin{aligned} N &:= \{s \mid (\forall\, r,\ t\in\Sigma^*,\ \sigma\in\Sigma)\\ &\qquad [\, s = r\sigma t,\ r\in M,\ r\sigma\notin M \ \Rightarrow\ \sigma\in\Sigma_{\mathrm{uc}}\,]\,\}\,, && (17)\\ E' &:= E\ \cup\ (N\cap L)\,. && (18)\end{aligned}$$

In other words: the closed-loop behaviour may only sacrifice the persistent chance to attain a winning configuration if this can not be prevented by control. As a preliminary observation, the following proposition establishes relevant closed-loop properties for the overall relaxation N.

Proposition 8. If $\epsilon \in M$ then N is closed, L and N are non-conflicting, and

$$N\cap L \in \mathsf{CNF}(L,\ N\cap L)\,. \tag{19}$$

Proof (outline). The claim can be proved by verifying the individual properties in a specific order. (1) By inspection of the respective definition, N is seen to be closed and, thus, $N \cap L$ is relatively closed w.r.t. L. (2) Non-conflictingness can be established by picking an arbitrary $s \in (\operatorname{pre} L) \cap (\operatorname{pre} N)$ and considering distinct cases. All cases can be dealt with elementary, except when $s \notin M$ and any extension v with svL passes through M. Here, Proposition 7, part (iii), can be applied. For (3) and (4), where one needs to prove controllability and prefix-normality, respectively, non-conflictingness (2) can be conveniently used. □

The main result from this section is an immediate consequence of the above proposition.

Theorem 9. Given a control problem parameterised by $A \subseteq E \subseteq L \subseteq \Sigma^*$, consider the relaxed upper bound E' in Eqs. (17) and (18), referring to the controllability prefix T. Write $K' := \sup \mathsf{CNF}(L,\ E')$ for the supremal closed-loop behaviour implied by E'. Then,

$$\epsilon \in \operatorname{pre} A \subseteq M \tag{20}$$

implies that $\operatorname{pre} A \subseteq \operatorname{pre} K'$. If, in addition, A is relatively closed w.r.t. L, then

$$A \subseteq K' \subseteq E'\,. \tag{21}$$

Proof (outline). From Proposition 8 one can derive that $N \cap L \subseteq \sup \mathsf{CNF}(L,\ E') = K'$ and, $(\operatorname{pre} N) \cap (\operatorname{pre} L) \subseteq \operatorname{pre} K'$. With $\epsilon \in M$ we have $M \subseteq N$ and thus obtain $\operatorname{pre} A \subseteq \operatorname{pre} M \subseteq \operatorname{pre} K'$. Intersecting both sides with L establishes the claim. □

Note that the requirement of A being relatively closed w.r.t. L can always be satisfied by substituting A by its infimal relatively closed superset. This is not restrictive, since any closed-loop behaviour is relatively closed.

5. EXAMPLE (CNT.)

Recall from Section 2 that the supremal closed-loop behaviour $K_{\mathrm{f}}^{\uparrow}$ obtained by applying the nominal upper-bound specification E_{f} to the fault-accommodating plant L_{f} was an unacceptably small subset of the supremal closed-loop behaviour $K^{\uparrow}$ obtained for the nominal plant L. In particular, no high-level command A is ever accepted by $K_{\mathrm{f}}^{\uparrow}$. We are now in the position to formally identify problematic configurations by inspecting the controllability prefix T_{f} of E_{f} w.r.t. L_{f}; see Fig. 6.

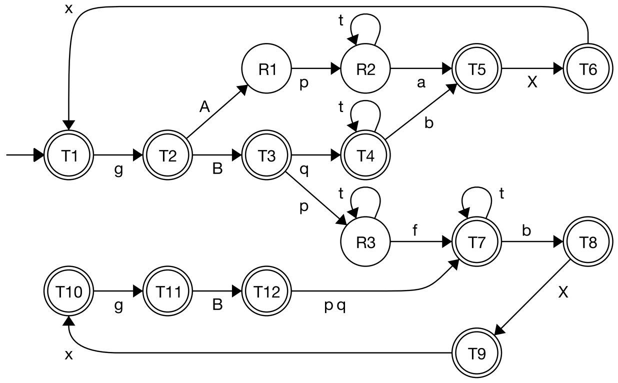

Fig. 6. T_{f} for plant L_{f} and original upper bound E_{f}

The state set $\{\mathsf{T1},\ \mathsf{T2},\ \ldots,\ \mathsf{T6}\}$ corresponds to the closed-loop behaviour $K_{\mathrm{f}}^{\uparrow}$. Note that each of the latter states is marked and, hence, corresponds to strings within T_{f}; i.e., the supervisor does not *need* to risk to become unable to enforce E_{f}. However, when in state $\mathsf{T2}$ or $\mathsf{T3}$, there is the option to optimistically enable additional events at the price to leave the guaranteed safe region. Considering the particular semantics of the fault event f, states $\mathsf{T2}$ and $\mathsf{T3}$ differ. If A becomes enabled while the plant is in state $\mathsf{T2}$, the closed loop can still comply with E_{f}, provided that the fault does not occur before p, i.e., provided that the wear-out does not show while the present workpiece is being processed. In contrast, when B becomes enabled while the plant is in state $\mathsf{T3}$, the closed-loop can only comply with E_{f} if the fault actually occurs. However, for practical reasons, a supervisor should not operate the plant such that the occurrence of the fault is required to satisfy liveness properties. Moor

(2016) therefore proposes to intersect $\mathsf{CNF}(L_\mathrm{f}, E_\mathrm{f})$ with the class

$$\mathsf{FF}(L_\mathrm{f}) := \{K_\mathrm{f} \subseteq L_\mathrm{f} \,|\, (\forall\, s \in \mathrm{pre}\, K_\mathrm{f}\ \exists\, t \in \Sigma^*)[\, st \in K_\mathrm{f}\ ]\}, \quad (22)$$

for controller synthesis. In particular, the cited literature verifies that the class $\mathsf{FF}(\cdot)$ is closed under arbitrary union. For the example at hand, we pragmatically propose to restrict T_f by $T'_\mathrm{f} := \sup \mathsf{FF}(T_\mathrm{f})$; see Fig. 7.

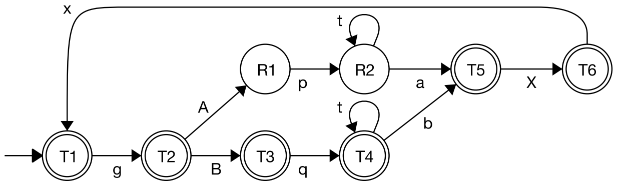

Fig. 7. $T'_\mathrm{f} := \mathsf{FF}(T_\mathrm{f})$ for plant L_f and original upper bound E_f

To this end, we denote $M'_\mathrm{f} := \sup \mathsf{MN}(L_\mathrm{f}, T'_\mathrm{f})$ and relax the nominal specification by $E' := E_\mathrm{f} \cup (N'_\mathrm{f} \cap L_\mathrm{f})$, referring to Eq. (17) and substituting M by M'_f. Note that, by monotonicity, we have $M'_\mathrm{f} \subseteq \sup \mathsf{MN}(L_\mathrm{f}, T_\mathrm{f})$. The resulting supremal closed-loop behaviour for the present input data is given in Fig. 8.

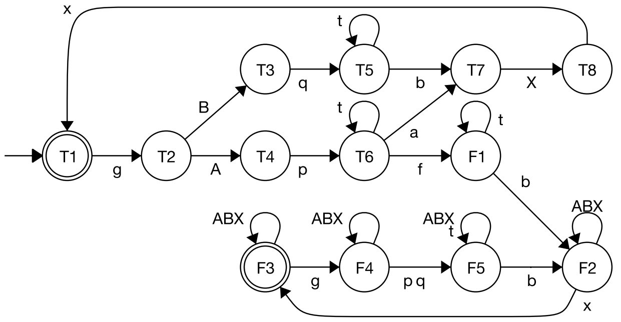

Fig. 8. closed-loop behaviour for relaxed upper bound E'

As a further refinement, we propose to adapt N'_f in order to pass on when it is known by observation that compliance with the nominal specification can no longer be achieved. For a distinguished external event $\mathsf{F} \notin \Sigma_\mathrm{f}$, $\Sigma_\mathrm{F} := \Sigma_\mathrm{f} \dot\cup \{\mathsf{F}\}$, and with the low-level alphabet $\Sigma_\mathrm{lo} := \{\mathsf{g}, \mathsf{p}, \mathsf{q}, \mathsf{t}, \mathsf{f}, \mathsf{a}, \mathsf{b}, \mathsf{x}\}$ let

$$\begin{aligned} N''_\mathrm{f} := \{s \,|\, (\forall\, r,t \in \Sigma_\mathrm{F}^*,\ \sigma \in \Sigma_\mathrm{F})[\, s = r\sigma t,\ r \in M'_\mathrm{f},\ r\sigma \notin M'_\mathrm{f} \\ \Rightarrow \quad \sigma \in \Sigma_\mathrm{uc} \text{ and } t \in \Sigma_\mathrm{lo}^* \mathsf{F} \Sigma_\mathrm{f}^*\ ]\}. \quad (23)\end{aligned}$$

Proceeding as above, we end up with the same closed loop as is in Fig. 8, expect that when leaving state $\mathsf{F1}$ by event b, the indicator event F is issued before entering state $\mathsf{F2}$, i.e., there is one extra state "between" $\mathsf{F1}$ and $\mathsf{F2}$). Projecting to the external alphabet $\Sigma_\mathrm{hi} := \{\mathsf{A}, \mathsf{B}, \mathsf{X}, \mathsf{F}\}$, the external closed-loop behaviour turns out as given in Fig. 9.

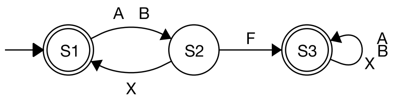

Fig. 9. external behaviour L''_hif obtained with relaxation N''_f

CONCLUSION

Referring to a plant and an upper-bound specification, the controllability prefix is defined as the set of words from which a supervisor can take over the plant to enforce the specification. This concept is well established for the supervision of ω-languages and we have elaborated a variation to address the supervision of $*$-languages under partial observation. Our study establishes algebraic properties of the controllability prefix that can be used to systematically relax a given upper-bound specification by allowing the supervisor to take the risk in failing on the upper bound while there is still the chance to win, but to do so only if this is known by observation. We demonstrate by example how this concept can be used in the context of fault-tolerant supervisory control, where a core challenge is how to relax a nominal specification to accommodate for a fault. Preliminary studies towards a software implementation include an early prototype to handle the example for the present paper.

REFERENCES

Acar, A.N. and Schmidt, K.W. (2015). Discrete event supervisor design and application for manufacturing systems with arbitrary faults and repairs. In *IEEE International Conference on Automation Science and Engineering*, 825–830.

Cai, K., Zhang, R., and Wonham, W. (2015). Relative observability of discrete-event systems and its supremal sublanguages. *IEEE Trans. Autom. Control*, 60, 659–670.

Cho, H. and Marcus, S.I. (1989). On supremal languages of classes of sublanguages that arise in supervisor synthesis problems with partial observation. *Maths. of Control, Signals & Systems*, 2, 47–69.

Lafortune, S. and Chen, E. (1990). The infimal closed controllable superlanguage and its application in supervisory control. *IEEE Trans. Autom. Control*, 35, 398–405.

Lin, F. and Wonham, W.M. (1988). On observability of discrete-event systems. *Information Sciences*, 44, 173–198.

Moor, T. (2016). A discussion of fault-tolerant supervisory control in terms of formal languages. *Annual Reviews in Control*, 41, 159 – 169.

Moor, T., Baier, C., Yoo, T.S., Lin, F., and Lafortune, S. (2012). On the computation of supremal sublanguages relevant to supervisory control. *Workshop on Discrete Event Systems (WODES)*, 175–180.

Moor, T. and Schmidt, K.W. (2015). Fault-tolerant control of discrete-event systems with lower-bound specifications. *Workshop on Dependable Control of Discrete Systems*.

Paoli, A. and Lafortune, S. (2005). Safe diagnosability for fault-tolerant supervision of discrete-event systems. *Automatica*, 41(8), 1335–1347.

Ramadge, P.J. and Wonham, W.M. (1987). Supervisory control of a class of discrete event processes. *SIAM J. Control and Optimization*, 25, 206–230.

Sülek, A.N. and Schmidt, K.W. (2014). Computation of supervisors for fault-recovery and repair for discrete event systems. In *Workshop on Discrete Event Systems*, 428–438.

Thistle, J.G. and Wonham, W.M. (1994a). Control of infinite behavior of finite automata. *SIAM J. Control and Optimization*, 32, 1075–1097.

Thistle, J.G. and Wonham, W.M. (1994b). Supervision of infinite behavior of discrete event systems. *SIAM J. Control and Optimization*, 32, 1098–1113.

Wen, Q., Kumar, R., and Huang, J. (2014). Framework for optimal fault-tolerant control synthesis: maximize prefault while minimize post-fault behaviors for discrete event systems. *IEEE Transactions on Systems, Man and Cybernetics: Systems*, 44, 1056–1066.

Wittmann, T., Richter, J., and Moor, T. (2012). Fault-tolerant control of discrete event systems based on fault-accommodating models. *Symposium on Fault Detection, Supervision and Safety of Technical Processes*, 854–859.