

# Limited Lookahead Policies for the Control of Discrete-Event Systems: A Tutorial

Richard Hugh Moulton,<sup>1</sup> Anthony J. Marasco,<sup>2</sup>  
and Karen Rudie<sup>3</sup>

June 11, 2020

---

Some problems in discrete-event systems (DES) model large, time-varying state spaces with complex legal languages. To address these problems, Chung et al. introduced limited lookahead policies (LLP) to provide online supervisory control for discrete-event systems. This seminal paper, along with an addendum of technical results, provided the field with a series of very important and powerful results, but in a notationally- and conceptually-dense manner. In this tutorial, we present Chung et al.'s problem formulation for online control and unravel the formal definitions and proofs from their original work with the aim of making the ideas behind limited lookahead accessible to all DES researchers. Finally, we introduce the Air Traffic Control problem as an example of an online control problem and demonstrate the synthesis of LLP supervisors.

---

<sup>1</sup> richard.moulton@queensu.ca  
Department of Electrical and Computer  
Engineering, Queen's University,  
Kingston ON, Canada

<sup>2</sup> anthony.marasco@rmc-cmr.ca  
School of Computing, Queen's Univer-  
sity, Kingston ON, Canada  
Department of Electrical and Computer  
Engineering, Royal Military College of  
Canada, Kingston ON, Canada

<sup>3</sup> karen.rudie@queensu.ca  
Department of Electrical and Computer  
Engineering and Ingenuity Labs Re-  
search Institute, Queen's University,  
Kingston ON, Canada

## 1 INTRODUCTION

THE MAIN PROBLEM in control theory is the problem of providing a control signal to a process to ensure that only desirable behaviour is produced. In discrete event systems (DES) this is done through supervisory control: enabling or disabling events that cause the process to transition from one state to another. Although DES theory provides supervisors that are correct by construction, there are real-world applications whose characteristics make constructing such supervisors infeasible. These applications include processes with a very large state space, processes that are time-varying, and processes whose desirable behaviours are difficult to specify.

Chung, Lin and Lafortune's 1992 paper "Limited Lookahead Policies in Supervisory Control of Discrete Event Systems" addressed these challenges by providing a theoretical base for the development of DES supervisors that performed control in an online manner rather than computing a monolithic control policy offline.<sup>1</sup> This work inspired a number of follow-on works and was recognized with the George S. Axelby Outstanding Paper Award for the *IEEE Transactions on Automatic Control*. Nonetheless, we believe that wider use of these limited lookahead policy (LLP) supervisors has been prevented by the conceptually dense material that must be mastered.

<sup>1</sup> Chung et al., 1992b

In this tutorial we present Chung et al.'s problem formulation, supervisor characterization and formal proofs of correctness in a way that emphasizes the underlying ideas. Our aim is to make the theoretical results in "Limited Lookahead Policies in Supervisory Control of Discrete Event Systems" usable for researchers and practitioners alike. We begin by introducing the *Air Traffic Control (ATC) problem* as an illustrative example of the online control problem we use throughout the tutorial (Section 2). Next we present Chung et al.'s seminal results (Sections 3 through 6) and demonstrate how these results are used to synthesize LLP supervisors (Section 7). Finally, we provide an overview of works that directly extend or apply ideas from the original LLP supervisor (Sections 8 and 9).

## 2 THE AIR TRAFFIC CONTROL PROBLEM

WE WILL USE THE ATC problem as a concrete example throughout this paper to illustrate the theoretical concepts we are discussing. This problem requires the formulation of a supervisor to control air traffic in the vicinity of an airfield. It is useful because it is conceptually easy to understand and because some of its characteristics make it challenging for traditional DES approaches.




Figure 1: The Air Traffic Control problem. The airspace to be controlled is shown: the airfield is in the centre square; aircraft can arrive in any of the four corners, but can only leave from the top left and bottom right corners; and once an aircraft is in the airspace it must travel in a clockwise manner around the airfield. The supervisor that we synthesize must ensure that aircraft can take off and land from this airfield with minimal restrictions while guaranteeing flight safety.

## 2.1 The Alphabet

As a discrete event system model, the ATC problem uses the alphabet  $\Sigma = \Sigma_c \cup \Sigma_{uc}$ . The controllable events, denoted by  $\Sigma_c$  (1), represent aircraft movements into a new section of airspace, as well as take-off ( $\delta_i$ ), landing ( $\eta_i$ ), entering the airfield's vicinity ( $\rho_i^j$ ), and leaving the airfield's vicinity ( $\chi_i^j$ ). The uncontrollable events, denoted by  $\Sigma_{uc}$  (2), represent notifications of a pending arrival or take off and are distinguished by a tilde,  $\sim$ . All events are subscripted with the index of the plane performing the action.

$$\Sigma_c = \{\epsilon, \alpha_i, \beta_i, \gamma_i, \theta_i, \kappa_i, \lambda_i, \mu_i, \nu_i, \eta_i, \delta_i, \rho_i^1, \rho_i^2, \rho_i^3, \rho_i^4, \chi_i^1, \chi_i^3\} \quad i \in \mathbb{N}^+ \quad (1)$$

$$\Sigma_{uc} = \{\tilde{\rho}_i^1, \tilde{\rho}_i^2, \tilde{\rho}_i^3, \tilde{\rho}_i^4, \tilde{\delta}_i\} \quad i \in \mathbb{N}^+ \quad (2)$$

In order to simplify the problem, we assume that all events are observable, although subsequent work allowed for unobserved events in the plant (see Section 8).

## 2.2 The Plant

The plant representing the uncontrolled airspace can be modelled as the synchronous product of multiple subplants. This will allow us to reason about each component of the plant separately, without having to consider the effects of the whole plant.

First, there is a permanent subplant that represents the supervisors for adjacent spaces, which is itself a synchronous product between subplants for departures and arrivals (Figure 2).




Figure 2: Supervisors for adjacent spaces; the first index in each state represents the departures controller and the second index represents the arrivals controller. The initial state is the only marked state in  $G_0$  and it represents the state where there are no aircraft waiting for permission to enter the airspace or take off.

Second, there is a subplant for each aircraft that is airborne or ready to enter the airspace. These subplants will need to be added to or removed from the overall plant as aircraft arrive and depart from

the airspace. An example of an aircraft subplant is shown in Figure 3; these subplants are denoted  $G_i$  where  $i$  is the aircraft's index.




Figure 3: A subplant representing an aircraft departing towards  $\chi^3$ . The initial state is before the aircraft has requested to depart and the marked state is when the aircraft has successfully departed the airspace.

As in Chung et al.'s original paper, we assume that the language generated by the plant is equal to the prefix-closure of the plant's marked language:  $L(G) = \overline{L_m(G)}$ .<sup>2</sup>

<sup>2</sup> Chung et al., 1992b

### 2.3 The Specification

Both the large state space of the problem and the plant's dynamic nature will make it difficult to write down the language generated by the plant, but we will also need to consider how we will describe the legal language,  $K$ . As in the original paper, we consider the legal language to be  $L_m(G)$ -closed, that is  $K = \overline{K} \cap L_m(G)$ .<sup>3</sup>

<sup>3</sup> Chung et al., 1992b

At a high level, the supervisor must ensure that aircraft can take off from and land at this airfield with minimal restrictions and while guaranteeing flight safety. In plain English, the specifications captured by  $K$  are:

- Two aircraft can't occupy the same section of airspace;
- An arriving aircraft can't be delayed for more than one controllable action;
- A departing aircraft can't be delayed for more than ten controllable actions; and
- Aircraft departing through the same gate must be separated by at least five actions.

### 2.4 Summary

It is clear from its formulation that the ATC problem meets all three of Chung et al.'s criteria for problems requiring LLP control.

First, there is an explosion in the size of the state space. If we consider the case of five aircraft airborne at once, the synchronous product of the subplants representing the adjacent supervisors and each aircraft has on the order of  $6 \times 10^5$  states, i.e. even with a grossly simplified problem we have an impractical number of states.

Second, the plant is dynamic. Its composition changes every time an aircraft enters (or exits) the airspace. In the DES model this results in a subplant being added (or dropped) from the overall plant's synchronous product. This dynamic nature would require that an offline supervisor be recomputed for every new aircraft.

Third, the legal behaviour is difficult to capture as a regular language. This results from the plant's dynamic and modular natures: although event legality can be easily determined for any given point in time, determining this for all possible futures requires the supervisor to keep track of all possible legal languages.

### 3 DES OPERATIONS FOR LLP SUPERVISORS

ALL OF THE CHARACTERISTICS of DES problems that make LLP supervisors necessary also make the analysis of the problems themselves somewhat difficult. We begin, therefore, by defining and illustrating some DES operations that will prove useful in our reasoning about the problem as well as the supervisor's own operations (all definitions from Chung et al.<sup>4</sup>).

**Definition 1** (Active Set). The active set in a language  $L(G)$  after the trace  $s$  is denoted  $\Sigma_{L(G)}(s)$  and defined as

$$\Sigma_{L(G)}(s) := \{\sigma \in \Sigma \mid s\sigma \in L(G)\}.$$

**Example 1** (Active Set). At any point in time where there are no aircraft airborne, the active set for the ATC problem is restricted to the alphabet of uncontrollable events. This is because the controllable events are only applicable when there is an aircraft in the airspace to be directed.

$$\Sigma_{L(G)}(\epsilon) = \Sigma_{uc}$$

**Definition 2** (Post-language). The post-language of a language  $L$  after the trace  $s$  is denoted  $L/s$  and defined as

$$L/s := \{t \in \Sigma^* \mid st \in L\}.$$

**Example 2** (Post-language). If the aircraft from Figure 3 is the only aircraft airborne and

$$s = \tilde{\delta}_1 \delta_1 v_1 \alpha_1 \beta_1 \gamma_1$$

<sup>4</sup> Chung et al., 1992b

The active set allows the supervisor to reason about what actions are possible in the plant from any given state.

The post-language operation provides the "lookahead" portion of LLP by allowing the future language of the plant to be determined on the basis of the history of the system. Since the possibility of some events is dependent on the state of the plant, this kind of operation cannot be done ahead of time and must be computed on-the-fly.

occurs, then the post-language contains strings such as

$$\begin{aligned} L(G)/s = \{ & \theta_1, \theta_1\chi_1^3, \tilde{\rho}_2^1, \tilde{\rho}_2^1\theta_1, \tilde{\rho}_2^1\theta_1\rho_2^1, \tilde{\rho}_2^1\theta_1\rho_2^1\chi_1^3, \\ & \tilde{\rho}_2^1\theta_1\rho_2^1\chi_1^3\alpha_2, \tilde{\rho}_2^1\theta_1\rho_2^1\chi_1^3\alpha_2\beta_2, \\ & \tilde{\rho}_2^1\theta_1\rho_2^1\chi_1^3\alpha_2\beta_2\gamma_2, \tilde{\rho}_2^1\theta_1\rho_2^1\chi_1^3\alpha_2\beta_2\gamma_2\eta_2, \\ & \tilde{\rho}_2^2, \tilde{\rho}_2^2\rho_2^2, \tilde{\rho}_2^2\rho_2^2\theta_1, \tilde{\rho}_2^2\rho_2^2\theta_1\gamma_2, \\ & \tilde{\rho}_2^2\rho_2^2\theta_1\gamma_2\chi_1^3, \tilde{\rho}_2^2\rho_2^2\theta_1\gamma_2\chi_1^3\eta_2, \dots \} \end{aligned}$$

**Definition 3** (Truncation). The truncation of a language  $L$  to the strings of length  $N$  is denoted  $L|_N$  and defined as

$$L|_N := \{t \in L \mid |t| \leq N\}$$


where  $|t|$  is the length of the string  $t$ .

**Example 3** (Truncation). Extending Example 2 the truncated post-language  $L(G)/s|_3$  contains strings such as

$$\begin{aligned} \{ & \theta_1, \theta_1\chi_1^3, \tilde{\rho}_2^1, \tilde{\rho}_2^1\theta_1, \tilde{\rho}_2^1\theta_1\rho_2^1, \\ & \tilde{\rho}_2^2, \tilde{\rho}_2^2\rho_2^2, \tilde{\rho}_2^2\rho_2^2\theta_1, \dots \} \end{aligned}$$

## 4 CHARACTERIZING LLP SUPERVISORS

WITH THESE USEFUL DEFINITIONS INTRODUCED, we now turn our attention to formalizing the ideas behind LLP supervisors. An LLP supervisor (Figure 4) is a general supervisory framework and can be modelled in many different ways, nevertheless we formalize the operations in terms of automata, languages and strings as done by the original authors.<sup>5</sup>



The LLP control scheme (Algorithm 1) aligns with traditional DES approaches: the biggest departure is in the third step where the controller must make assumptions about plant behaviour that is neither definitely legal nor definitely illegal. In practice the controller does this by adopting an attitude towards behaviour of uncertain legality: the controller can optimistically believe that all uncertain strings

The truncation operation provides the “limited” portion of LLP by allowing the language of the plant to be trimmed to a tractable set of strings. As seen from Example 2 it can be difficult to write strings out in full, let alone enumerate all possible strings. Instead, truncation restricts the language to those of its members whose length is no more than  $N$ .

<sup>5</sup> Chung et al., 1992b

Figure 4: An LLP supervisor projects the plant’s behaviour for the next  $N$  steps (adapted from Chung et al., 1992b)

can be steered towards legality or the controller can conservatively restrict the plant to behaviour that is known to be legal.

---

Algorithm 1: LLP Supervisor

- 1: **while** events occur in the plant **do**
- 2:   Project the plant's behaviour over the next  $N$  events

$$L(G)/s|_N \text{ and } L_m(G)/s|_N$$

- 3:   Determine legality of all strings

$$\overline{K}/s|_N \text{ and } K/s|_N$$

- 4:   Label pending strings according to attitude

$$f_a^N(s) = \begin{cases} \text{cons.} & K/s|_{N-1} \\ \text{optm.} & K/s|_N \cup (\overline{K}/s|_N \setminus \overline{K}/s|_{N-1}) \end{cases}$$

- 5:   Find the modified tree's supremal controllable sublanguage

$$f^N(s) = [f_a^N(s)]^{\uparrow/s|_N}$$

- 6:   Output control action

$$\gamma^N(s) = \overline{f^N(s)|_1} \cup \overline{\Sigma_{uc} \cap \Sigma_{L(G)}(s)}$$


---

For the ATC problem this algorithm leads the LLP supervisor through the following determinations: whether aircraft might arrive and/or take off and where all aircraft in the airspace can move to; which of these behaviours respect the specifications and which do not; taking either a conservative or optimistic attitude towards behaviours that are not clearly legal or illegal; calculating the supremal controllable sublanguage of this pruned tree; and then outputting the control action for all aircraft under its control.

In order to be convinced that this formalization of an LLP supervisor leads to correct behaviour we must answer three questions:<sup>6</sup>

1. What happens if the supervisor says that there is no policy that can keep the plant's behaviour legal?

From a purely practical point of view, we are very interested in knowing how our supervisor will act when its calculations tell it that illegal behaviour cannot be prevented.

2. Is there a reason to prefer one attitude over the other? How differently do the optimistic and conservative attitudes behave?

The two proposed attitudes seem to be suited to different scenarios.

Algorithm 1 is adapted from Chung et al., 1992b. The LLP control scheme is a five step process, with the supervisor performing the following calculations after each event: predict the possible future behaviour of the plant for the next  $N$  steps; eliminate any traces the controller knows to be illegal; decide how to label the pending traces; calculate the pruned tree's supremal controllable sublanguage; and produce the control action.

<sup>6</sup> Chung et al., 1992b

If this is true we would like to know how we can expect each attitude to behave and in which scenarios we should pick one attitude over the other.

3. Is it worth using an LLP supervisor instead of simply calculating an offline supervisor in the traditional manner?

Given the weight of the literature showing that traditional supervisors are correct-by-construction and provably correct, we would like to know if we are trading away any of these guarantees when we decide to use an LLP supervisor. We will return to these three questions throughout the following presentation of Chung et al.'s technical results.

#### 4.1 Characteristic Error of the LLP Supervisor

In answering the first question we get right to the heart of what distinguishes LLP supervisors. In contrast with a traditional offline supervisor, the characteristic and defining error of the LLP supervisor is the run-time error.

**Definition 4** (Run-time error). If  $s \in L(G, \gamma^N)$  and  $f^N(s) = \emptyset$ , then we say that there is a **run-time error** happening in  $L(G, \gamma^N)$  at trace  $s$ .

A run time error exists when the LLP does not have any guaranteed safe action, in the sense that there is no trajectory from the current state of the system to a legal marked state that isn't at risk of being driven uncontrollably into illegal behaviour. The answer to the first question is therefore that if the supervisor says that no policy can prevent illegal behaviour in the plant then something has gone wrong and the supervisor should have prevented the plant from getting into this state in the first place.

A special case of the run-time error is when it occurs at the beginning of the plant's behaviour. This is called a starting error and it indicates that no supervisor can guarantee safe behaviour in the plant.

**Definition 5** (Starting error). If a run-time error occurs for  $s = \epsilon$ , we call this a **starting error**.

With this in mind, we should begin by asking whether the ATC problem has a starting error? If so, then it is not suitable for supervisory control and we must either reshape the problem or make use of other techniques.



## 4.2 Validity of the LLP Supervisor

The question remains whether a valid supervisor can be synthesized for the ATC problem with either the optimistic or conservative attitude. This is the question that we will focus on next, beginning with Chung et al.'s definition of supervisor validity.<sup>7</sup>

**Definition 6** (Valid supervisor). An LLP supervisor with control policy  $\gamma^N$  is called **valid** if

$$L(G, \gamma^N) = \overline{K^\dagger}.$$

A valid supervisor for the ATC problem would therefore enforce flight safety, but only by restricting aircraft when necessary. The takeaway is that a valid LLP supervisor necessarily balances correctness with permissiveness.

<sup>7</sup> Chung et al., 1992b

It is worth noting that the terminology of an LLP supervisor's *validity* is potentially confusing given what it expresses. Clearly, a valid supervisor avoids blocking states as well as states which lead uncontrollably to illegal strings. Definition 6 is stronger than this, however, and also requires that the supervisor be minimally restrictive.

## 5 ATTITUDES FOR THE LLP SUPERVISOR

CHUNG ET AL. PRESENTED TWO POSSIBLE ATTITUDES that LLP supervisors could take towards pending strings in their  $N$ -step lookahead windows: optimistic and conservative.<sup>8</sup> We will consider the size of language produced by each attitude, the effect of larger windows on these language and whether these attitudes result in valid supervisors.

To begin, we state a result that will prove useful in our analysis. Put plainly, it states that if an attitude is at least as permissive for every possible string as another attitude, then the language allowed by the former attitude will be at least as big as the language allowed by the latter. This is stated more formally in Lemma 1 and is proved by induction.<sup>9</sup>

<sup>8</sup> Chung et al., 1992b

**Lemma 1.** If  $\gamma_i(s) \supseteq \gamma_j(s)$  for all  $s \in \Sigma^*$ , then  $L(G, \gamma_i) \supseteq L(G, \gamma_j)$ .

<sup>9</sup> Chung et al., 1992b

### 5.1 The Optimistic Attitude

An LLP supervisor with the optimistic attitude marks all pending strings as legal, which means that it will presume that it can steer any trace to a string in  $K$  until there is proof that it cannot.

**Example 4** (Uncontrollable illegality). In the ATC problem, an optimistic supervisor could allow an aircraft in both the top-left and top-middle sections of airspace (Figure 5).

This trace isn't illegal itself and is the prefix of a marked string, but it could uncontrollably lead to an illegal string: the event  $\rho_i^1$  would require two controlled actions to avoid having two aircraft

This attitude will naturally allow every string in  $\overline{K^\dagger}$  to occur, since there will never be reason to believe that they will uncontrollably lead to an illegal state. It might also allow additional strings which do lead uncontrollably to illegal strings, however, since the information required to avoid these strings may lie outside the lookahead window.




Figure 5: Although strings that lead to this state are prefixes of legal, marked strings, the plant can lead uncontrollably to illegal behaviour if the event  $\rho_i^1$  occurs.

in the same section of airspace but that would contravene the specification that only one controlled event can happen before  $\rho_i^1$ .

**The language produced by the optimistic attitude** The optimistic attitude prioritizes maximal permissiveness over preventing illegal behaviour. The size of the language it produces,  $L(G, \gamma_{optm}^N)$ , is provably bounded between the prefix-closure of the supremal controllable sublanguage of  $K$  and the infimal controllable superlanguage of  $K$ .<sup>10</sup>

<sup>10</sup> Chung et al., 1992b

**Theorem 1.**  $\overline{K^\uparrow} \subseteq L(G, \gamma_{optm}^N) \subseteq K^\downarrow$

Intuitively, then, we hope that if we increase the size of the lookahead window then the LLP supervisor will have more of the information it needs to avoid illegal strings and that  $L(G, \gamma_{optm}^N) = \overline{K^\uparrow}$  for some large enough  $N$ . Chung et al. proved this as Theorem 2 which states that increasing the size of the lookahead window will never increase the size of the language allowed by an optimistic supervisor.<sup>11</sup>

<sup>11</sup> Chung et al., 1992b

**Theorem 2.**  $L(G, \gamma_{optm}^N) \supseteq L(G, \gamma_{optm}^{N+1})$

**Run-time errors with the optimistic attitude** An LLP supervisor with the optimistic attitude will never prohibit a string that is contained in  $\overline{K^\uparrow}$ , so we know that it will be minimally restrictive. We do have to verify, however, whether an illegal string might be allowed to occur.

**Theorem 3.** If a supervisor with the optimistic attitude is valid, then there will be no run-time errors in  $L(G, \gamma_{optm}^N)$ .

*Proof.* Chung et al.'s proof<sup>12</sup> of Theorem 3 is by contradiction and is broken down into two cases: when the current state's active set contains at least one uncontrollable action and, conversely, when the current state's active set is entirely controllable.

<sup>12</sup> Chung et al., 1992b

1. Consider when  $\Sigma_u \cap \Sigma_{L(G)}(s) \neq \emptyset$ .

Because we assume the supervisor is valid and that a run-time error has occurred, we have that  $f_{optm}^N(s) = K^\dagger/s|_1 = \emptyset$ . But,  $\gamma_{optm}^N(s)$  must include the uncontrollable events that can occur in the plant, which makes it non-empty and a strict superset of  $\overline{K^\dagger}/s|_1$ . This contradicts the definition of a valid supervisor, therefore  $\gamma_{optm}^N$  is not valid.

2. Consider when  $\Sigma_u \cap \Sigma_{L(G)}(s) = \emptyset$ .

Because we assume that a run-time error has occurred we have that  $s \notin K$ , otherwise the supervisor could permit  $\epsilon$  to occur. If  $\gamma_{optm}^N(s) = \emptyset$  and  $s \notin K$  then our supervisor  $\gamma_{optm}^N$  is blocking at  $s$ , which means that  $s$  is not a prefix of any string in  $K$ . From this it follows that  $L(G, \gamma_{optm}^N) \neq K^\dagger$  and therefore  $\gamma_{optm}^N$  is not valid.

For both cases, if a run-time error occurs then the supervisor was not valid, which is a contradiction. Therefore we conclude that if a supervisor with the optimistic attitude is valid, then there will be no run-time errors in  $L(G, \gamma_{optm}^N)$ .  $\square$

This result is intuitive because by definition a valid supervisor is one that produces the prefix-closure of the supremal controllable sublanguage of the legal language. Unfortunately, the converse of Theorem 3 is not true. Chung et al. give a counterexample to show this, with the key being the potential for an unlimited number of uncontrollable events occurring.<sup>13</sup>

<sup>13</sup> Chung et al., 1992b

## 5.2 Conservative Attitude

An LLP supervisor with the conservative attitude marks all pending strings as illegal, which means that it presumes a trace leads uncontrollably to an illegal string unless there is proof that it does not.

**Example 5** (Overly restrictive). In the ATC problem, this means that the supervisor could take more actions than necessary in order to ensure the legality of the produced strings.

For example, consider a supervisor responsible for controlling the configuration shown in Figure 6. If this supervisor is notified that an aircraft is ready to depart, the event  $\tilde{\delta}_i$  occurs, then the supervisor might unnecessarily delay that aircraft's departure until all airborne aircraft have departed the airspace.

This attitude will never permit a string that is not part of the plant's supremal controllable sublanguage, but it may also prevent strings in  $\overline{K^\dagger}$  from occurring, since it may not be clear that their prefixes can be controlled all the way to a marked string.

**The language produced by the conservative attitude** The conservative attitude results in prioritizing the prevention of illegal behaviour over maximal permissiveness. The size of the language it produces,  $L(G, \gamma_{cons}^N)$ , is provably bounded such that it is no larger




Figure 6: Having received notification that an aircraft is ready to take off, the conservative attitude may lead a supervisor to prevent this until all the airborne aircraft have departed the airspace.

than the prefix-closure of the supremal controllable sublanguage of  $K$ : as long as  $\overline{K^\dagger} \neq \emptyset$ , a conservative supervisor composed with the plant is guaranteed to produce a subset of  $\overline{K^\dagger}$ .<sup>14</sup>

<sup>14</sup> Chung et al., 1992b

**Theorem 4.**  $K^\dagger \neq \emptyset \iff L(G, \gamma_{cons}^N) \subseteq \overline{K^\dagger}$

We would like to think, then, that by increasing the size of the lookahead window the LLP supervisor will be able to see that more strings in  $\overline{K^\dagger}$  can be permitted and we will have  $L(G, \gamma_{cons}^N) = \overline{K^\dagger}$  for some large enough  $N$ . Chung et al. formalize that increasing the lookahead window's size will never decrease the size of the language allowed by a conservative supervisor.<sup>15</sup>

<sup>15</sup> Chung et al., 1992b

**Theorem 5.**  $L(G, \gamma_{cons}^{N+1}) \supseteq L(G, \gamma_{cons}^N)$

*Proof.* Our proof follows Chung et al.'s from their addendum of technical results.<sup>16</sup> Due to Lemma 1 the result in Theorem 5 will follow if we show

<sup>16</sup> Chung et al., 1992a

$$f_{cons}^N \subseteq f_{cons}^{N+1} \quad \forall s \in \Sigma^*.$$

By the definition of the  $f_u^N$  block (Algorithm 1, Line 6) we have

$$f_{cons}^N := (K/s|_{N-1})^{\uparrow/s|_N} = (K/s|_{N-1})^{\uparrow/s|_{N+1}}.$$

By the definition of truncation we have

$$(K/s|_{N-1}) \subseteq (K/s|_N).$$

Since the post-language and supremal controllable sublanguage operations don't alter inclusion relationships

$$\begin{aligned} (K/s|_{N-1}) &\subseteq (K/s|_N) \\ \implies (K/s|_{N-1})^{\uparrow/s|_{N+1}} &\subseteq (K/s|_N)^{\uparrow/s|_{N+1}} \\ \implies f_{cons}^N &\subseteq f_{cons}^{N+1} \end{aligned}$$

□

**Run-time errors with the conservative attitude** From Theorem 4 we can see that as long as the plant's supremal controllable sublanguage is nonempty then the language produced by a supervisor with the conservative attitude will be nonempty as well.

**Theorem 6.** If there is no starting error in  $L(G, \gamma_{cons}^N)$ , then there will be no run-time error in  $L(G, \gamma_{cons}^N)$ .

*Proof.* Chung et al.'s proof<sup>17</sup> proceeds by induction on the length of the trace  $s$ .

<sup>17</sup> Chung et al., 1992a

The base case is  $s = \epsilon$ , where the length of  $s$  is 0. There is no run-time error at  $s$  because we have assumed that there is no starting error in  $L(G, \gamma_{cons}^N)$ .

Our hypothesis is that no run-time error has occurred for  $s$ , a trace of length  $i$ , and we will show that this implies there is no run-time error for  $s\sigma$ , a trace of length  $i + 1$  where  $\sigma \in \gamma_{cons}^N$ .

Since there was no run-time error at  $s$  we know that  $f_{cons}^N(s) \neq \emptyset$ ; as a minimum  $\sigma \in f_{cons}^N(s)$  since the conservative attitude allowed it to occur. With this in mind, as well as Theorem 5's result, if we consider the post-language  $f_{cons}^N(s)/\sigma$  then we have

$$\{\epsilon\} \subseteq f_{cons}^N(s)/\sigma = f_{cons}^{N-1}(s\sigma) \subseteq f_{cons}^N(s\sigma)$$

which means  $f_{cons}^N(s\sigma) \neq \emptyset$  and no run-time error occurs at  $s\sigma$ . □

Given its emphasis on enforcing legal behaviour, it makes sense that as long as the plant begins in a legal state, a supervisor with the conservative attitude will never allow the plant enter an illegal state.

### 5.3 Comparing the two attitudes

Having described the optimistic and conservative attitudes, we can ask ourselves how do these attitudes relate to each other, i.e. Chung et al.'s second question. We first consider the differences in the control actions they produce for one step ahead in the plant, formulated as Chung et al.'s first comparison result.<sup>18</sup>

<sup>18</sup> Chung et al., 1992b

**Theorem 7** (Comparing One Step Ahead).  $\overline{f_{cons}^M(s)}|_1 \subseteq \overline{f_{optm}^N(s)}|_1 \quad \forall s \in \Sigma^*, \quad \forall N, M \in \mathbb{N}$

Although this result is obvious when  $N = M$ , it is less clear that it should be true otherwise. With our previous results, however, the proofs are straightforward.

*Proof.* Consider three cases.

1.  $N = M$ . As stated above, this follows directly from the attitudes' respective definitions.
2.  $N < M$ . We have

$$\overline{f_{cons}^M(s)}|_1 \subseteq \overline{f_{optm}^M(s)}|_1 \subseteq \overline{f_{optm}^N(s)}|_1$$

because increasing the lookahead window size does not increase the size of the language permitted by the optimistic attitude (Theorem 2)

3.  $N > M$ . Similarly, we have

$$\overline{f_{cons}^M(s)}|_1 \subseteq \overline{f_{cons}^N(s)}|_1 \subseteq \overline{f_{optm}^N(s)}|_1$$

because increasing the the lookahead window size does not decrease the size of the language permitted by the conservative attitude (Theorem 5)

□

This is a powerful relationship between the two attitudes and with Lemma 1 it extends directly to comparing the languages produced by the two attitudes, formalized in Chung et al.'s second comparison result.<sup>19</sup>

<sup>19</sup> Chung et al., 1992b

**Theorem 8** (Comparing Complete Languages).  $L(G, \gamma_{cons}^M) \subseteq L(G, \gamma_{optm}^N) \quad \forall N, M \in \mathbb{N}$

The key takeaways from Theorem 8 are that the choice of attitude has a definite effect on  $L(G, \gamma^N)$  and that, unintuitively, this effect is independent of the size of lookahead window used. The power of this result is that we can be sure that the conservative attitude will never produce behaviour that is more permissive than the behaviour produced by the optimistic attitude, regardless of their respective lookahead window sizes.

## 6 COMPARING ONLINE AND OFFLINE SUPERVISORS

WE WILL NOW ADDRESS Chung et al.'s third question: do we give up any guarantees or abilities when we forego an offline supervisor and we use an LLP supervisor instead?

The results covered in Section 5 allow us to note that the sequences of languages

$$\{L(G, \gamma_{cons}^N)\}_{N=1}^{\infty} \text{ and } \{L(G, \gamma_{optm}^N)\}_{N=1}^{\infty}$$

are both bounded and monotone, and we can therefore conclude that their limits,  $L(G, \gamma^{\infty})$ , exist. As a final answer for the second question we would like to know if these limits are the same language and, to answer the third question, are these limits the same as the language produced by an offline controller? Formally, we are asking whether or not it is true that:

$$\lim_{N \rightarrow \infty} L(G, \gamma^N) = \overline{K^\dagger} \quad (3)$$

Unfortunately, Equation 3 is not true in the general case, even under the assumption that  $L_m(G)$  and  $K$  are regular languages.<sup>20</sup> All is not lost, however, since we can place conditions on the problem and on the value of  $N$  that do guarantee that  $L(G, \gamma^N) = \overline{K^\dagger}$ . This can be done whether or not  $K$  is prefix-closed.

<sup>20</sup> Chung et al., 1992b

### 6.1 Case I: $K$ is prefix-closed

We start with the easier of the two cases, where the supervisor never has to worry about blocking, because the proofs are correspondingly simpler to follow.

To facilitate the analysis, Chung et al. introduce the measure  $N_u(L)$  which is the length of the longest finite subtrace of uncontrollable events that occurs in language  $L$ .<sup>21</sup>

<sup>21</sup> Chung et al., 1992b

$$N_u(L) := \begin{cases} \max\{|s| \mid s \in \Sigma_{uc} \wedge \exists u, v \in \Sigma^* \\ (usv \in L)\} \text{ if it exists} \\ \text{undefined otherwise} \end{cases} \quad (4)$$

**The optimistic attitude** Although we know that the absence of run-time errors does not imply that an optimistic supervisor is valid, we can make this claim if  $K = \overline{K}$ .<sup>22</sup> Chung et al. proved that if  $N$  can be made large enough to ensure no run-time errors in  $L(G, \gamma_{optm}^N)$  then it also ensures that  $\gamma_{optm}^N$  is a valid supervisor.<sup>23</sup>

<sup>22</sup> Chung et al., 1992b

<sup>23</sup> Chung et al., 1992b

**Theorem 9** (Conditions on  $N$  – Case I, Optimistic Attitude). If  $K = \overline{K}$  then

$$(N \geq N_u(L(G)) + 1) \vee (N \geq N_u(K) + 2) \\ \implies L(G, \gamma_{optm}^N) = \overline{K}^\dagger.$$

Of note, for the first condition in Theorem 9's antecedent we need only to look one event further than the longest finite uncontrollable subtrace that can be generated by the plant. As shown in Figure 7 if  $t$  is the longest finite subtrace of uncontrollable events in  $L(G)$ , then an LLP supervisor with the optimistic attitude only needs to look  $|t| + 1$  events ahead in order to allow  $t$  to occur.

By contrast, in the second condition we must look two events further because the plant may be capable of generating longer uncontrollable subtraces than are included in the legal language. Consider in Figure 7 the case where  $t$  is the longest finite subtrace of uncontrollable events in  $K$ . Since it is still possible that the plant could produce additional uncontrollable events after  $t$ , the optimistic LLP supervisor must look  $|t| + 2$  events ahead in order to see that the plant can be controlled after  $t$  occurs.




Figure 7: Given that the trace  $s$  has occurred in the plant, how big a lookahead window is needed to decide whether  $\sigma_{c1}$  should be enabled or disabled?

**The conservative attitude** Taking a different approach for the conservative attitude, we can show that if  $L(G, \gamma_{cons}^N)$  has no starting error then the conservative supervisor is valid.<sup>24</sup> The key is the language

$$K_{pruned}^N = K \setminus (K / \Sigma_u^{N-1}) \Sigma^* \quad (5)$$

where if there is no starting error in  $L(G, \gamma_{cons}^N)$  and if  $K \cap \Sigma_u^{(N-1)} = \emptyset$  then

$$L(G, \gamma_{cons}^N) = (K_{pruned}^N)^\dagger \quad (6)$$

Chung et al. prove that if the plant begins in a legal state and  $N$  can be made large enough to see the longest finite uncontrollable subtrace in  $K$  (i.e. if  $K_{pruned}^N = K$ ) then it also ensures that  $\gamma_{cons}^N$  is a valid supervisor (Theorem 10).<sup>25</sup>

**Theorem 10** (Conditions on  $N$  – Case I, Conservative Attitude). If  $K = \overline{K}$  and there is no starting error in  $L(G, \gamma_{cons}^N)$  then

$$N \geq N_u(K) + 2 \implies L(G, \gamma_{cons}^N) = \overline{K}^\dagger.$$

<sup>24</sup> Chung et al., 1992b

<sup>25</sup> Chung et al., 1992b

Because  $K \subseteq L(G)$  we can use  $N_u(L(G))$  to approximate  $N_u(K)$ . This transforms Theorem 10 into

$$N \geq N_u(L(G)) + 2 \implies L(G, \gamma_{cons}^N) = \overline{K}^\dagger$$

and is useful since it may be more practical to calculate  $L(G)$  than  $K$  for some problems.(Chung et al., 1992b).



It is notable that, unlike for the optimistic attitude, it is not sufficient that  $N \geq N_u(L(G)) + 1$  to guarantee that  $\gamma_{cons}^N$  is a valid supervisor. This is because if the longest uncontrollable subtrace in  $L(G)$  appears in a string in  $K$ , then it is possible that  $K \cap \Sigma_u^{N-1} \neq \emptyset$  and therefore  $K_{pruned}^N \neq K$ .

This is illustrated in Figure 7 by having a conservative LLP supervisor look  $|t| + 1$  events ahead. Even if  $t$  is the longest uncontrollable subtrace in  $L(G)$  and even though the supervisor can see to the end of the subtrace in the plant's behaviour, the supervisor cannot see if any uncontrollable actions are possible after  $t$  and will therefore conservatively prevent  $\sigma_{c1}$  from occurring.

**Characterizing the ATC Problem** Theorems 9 and 10 make it clear that, at a minimum, a problem only admits valid LLP supervisors if the plant cannot generate an infinitely long subtrace of uncontrollable actions. This is true of the ATC problem by construction, since the supervisors for adjacent spaces (Figure 2) have been restricted from putting our supervisor in an uncontrollably illegal state. The longest uncontrollable subtrace that can occur in the plant is of length two, e.g.  $\delta_i \bar{\rho}_j^1$ , and after this subtrace occurs, the plant's active set is guaranteed to be restricted to controllable actions.

## 6.2 Case II: $K$ is not prefix-closed

The case where  $K \neq \bar{K}$  is more difficult because our LLP supervisor must avoid blocking states as well as illegal states. Unfortunately, this is the relevant case for many real-world problems, including the ATC problem.

To analyze this case, Chung et al. introduce the idea of a "frontier" of legal states beyond which the plant can uncontrollably end up in an illegal state. This leads to the definition of two sets of traces: legal, marked, controllable traces; and uncontrollably crossing traces.<sup>26</sup>

<sup>26</sup> Chung et al., 1992b

**Definition 7** (Legal, marked, controllable traces). The set of legal, marked, controllable traces is denoted  $K_{mc}$  and is defined as

$$K_{mc} := \{s \in K \cap L_c(G)\}$$

where the language  $L_c(G)$  contains all the traces that can be produced by the plant after which the active set contains only controllable actions.

$$L_c(G) := \{s \in L(G) \mid s\sigma \notin L(G) \quad \forall \sigma \in \Sigma_{uc}\}$$

**Definition 8** (Uncontrollably crossing traces). The set of uncontrollably crossing traces is denoted  $K_{f\bar{c}}$  and contains all the traces that lead from  $\bar{K}$  to  $L(G) \setminus \bar{K}$  on account of an uncontrollable event.

$$K_{f\bar{c}} := ((L(G) \setminus \bar{K}) / \Sigma_{uc}) \cap \bar{K}$$

**The optimistic attitude** The challenge for an optimistic supervisor is that it may allow traces to occur that lead to run-time errors. The key, then, is to ensure that the supervisor knows which strings in  $K_{mc}$  are prefixes of strings in  $K_{f\bar{c}}$ , i.e. the supervisor must know which legal strings must be disallowed because they may uncontrollably lead the plant into illegal behaviour.

This is achieved by defining the measure  $N_{mcf\bar{c}}$ , which is the length of the longest subtrace  $t$  that can occur after a legal, marked, controllable string  $s$  such that  $st$  is the only illegal string in  $\overline{st}/s$ .<sup>27</sup>

<sup>27</sup> Chung et al., 1992b

$$N_{mcf\bar{c}} := \begin{cases} \max\{|t| \mid \exists s \in K_{mc} \cup \{\epsilon\} \left( st \in K_{f\bar{c}} \wedge \right. \\ \quad \left. \forall \epsilon < v < t \left( sv \notin K_{f\bar{c}} \cup K_{mc} \right) \right)\} \\ \text{if it exists} \\ \text{undefined otherwise} \end{cases} \quad (7)$$

Chung et al. then prove that an LLP supervisor with the optimistic attitude is valid as long as  $N$  is greater than  $N_{mcf\bar{c}}$ .<sup>28</sup>

<sup>28</sup> Chung et al., 1992b

In Figure 7 consider the case where  $s\sigma_{c1}$  is a legal string and  $t$  is the longest subtrace that can occur such that  $s\sigma_{c1}t$  is the only illegal string in the plant's post-language after  $s\sigma_{c1}$ . If  $N < N_{mcf\bar{c}} + 1$  then the LLP supervisor will not see the illegal state that occurs after  $s\sigma_{c1}t$  and will optimistically allow the subtrace  $t$  to occur. If  $N \geq N_{mcf\bar{c}} + 1$ , on the other hand, the supervisor will have enough information to disable the event  $\sigma_{c1}$  and prevent the subtrace  $t$  from occurring.

**Theorem 11** (Conditions on  $N$  – Case II, Optimistic Attitude). If  $K^\uparrow \neq \emptyset$  then

$$N \geq N_{mcf\bar{c}} + 1 \implies L(G, \gamma_{optm}^N) = \overline{K^\uparrow}.$$

**The conservative attitude** In a similar fashion, the challenge for a conservative supervisor is that it may disallow traces that controllably lead to legal, marked states. The answer here is to ensure that the supervisor can see definitively whether or not every branch leads to a string in  $K_{mc}$  or to an illegal string.

The measure to accomplish this is  $N_{mcmc}$ , which is defined as the length of the longest string  $t$  that can occur after a legal, marked, controllable string  $s$  such that  $st$  is itself is the only legal, marked and controllable string in  $\overline{st}/s$ .<sup>29</sup>

<sup>29</sup> Chung et al., 1992b

$$N_{mcmc} := \begin{cases} \max\{|t| \mid \exists s \in K_{mc} \cup \{\epsilon\} \left( st \in K_{mc} \wedge \right. \\ \quad \left. \forall \epsilon < v < t \left( sv \notin K_{mc} \right) \right)\} \\ \text{if it exists} \\ \text{undefined otherwise} \end{cases} \quad (8)$$

This measure is used to formulate the guarantee in Theorem 12 namely that if  $N$  is greater than  $N_{mcmc}$  then a supervisor with the conservative attitude will have enough information to be a valid supervisor.<sup>30</sup>

<sup>30</sup> Chung et al., 1992b

The reasoning that supports  $N_{mcmc}$  as a measure is similar to the reasoning in support of  $N_{mcf\bar{c}}$ . In Figure 7 consider the case where  $s\sigma_{c1}$  is a legal string and  $t$  is the longest subtrace that can occur such that  $s\sigma_{c1}t$  is the only legal string in the plant's post-language after  $s\sigma_{c1}$ . If  $N < N_{mcmc} + 1$  then the LLP supervisor will not see the legal state that occurs after  $s\sigma_{c1}t$  and will conservatively prevent the subtrace  $t$  from occurring. If  $N \geq N_{mcmc} + 1$ , on the other hand, the supervisor will have enough information to allow the event  $\sigma_{c1}$  and it will not needlessly restrict the legal string  $s\sigma_{c1}t$  from occurring.

**Theorem 12** (Conditions on  $N$  – Case II, Conservative Attitude). If  $\bar{K} = \overline{K_{mc}}$  and there is no starting error in  $L(G, \gamma_{cons}^N)$  then

$$N \geq N_{mcmc} + 1 \implies L(G, \gamma_{cons}^N) = \bar{K}^\dagger.$$

**Characterizing the ATC problem** We recognize that the ATC problem's legal language,  $K$ , is not prefix-closed and that therefore Case II applies. To begin, we must check if the ATC problem meets the conditions for finding a value of  $N$  where  $L(G, \gamma^N) = \bar{K}^\dagger$ .

1. Does  $K^\dagger \neq \emptyset$ ?

Unfortunately, no; the supremal controllable sublanguage of the ATC problem is empty. Without synthesizing the whole legal language  $K$ , we can consider the following example.

**Example 6** (Showing that  $K^\dagger = \emptyset$ ). Our supervisor receives a notification that an aircraft will arrive from  $\rho^1$  whenever possible; these notifications are uncontrolled events, since they are the result of other supervisors, and therefore can't be disabled. But if the trace

$$s = \bar{\rho}_1^1 \rho_1^1 \bar{\rho}_2^1 \alpha_1 \rho_2^1 \bar{\rho}_3^1$$

occurs, then we will have aircraft in the top-left and top-middle sections of airspace and a third that must arrive to the top-left section after a maximum of one controllable action.

As seen in Example 4 this cannot be done while respecting the specification that two aircraft cannot occupy the same section of airspace. The only controllable action that can occur after  $s$  is  $\beta_1$ , but then both  $\alpha_2$  and  $\rho_3^1$  need to occur at the same time to prevent an illegal string.

With this, we conclude that the ATC problem as formulated does not admit a supervisor that can guarantee all the requirements are




Figure 8: The scenario in Example 4 can be made to occur uncontrollably by the supervisor of the neighbouring airspace.

met, let alone such an LLP supervisor. If we make adjustments to the plant, however, we can ensure that adjacent supervisors do not force our supervisor into an illegal state. We formulate the following additional specifications and incorporate them into our plant:

- Departing aircraft must be separated by ten actions;
- Aircraft arriving through the same gate must be separated by five actions; and
- Only five aircraft can be in the airspace at a given time.

The result is a new plant that has on the order of  $10^7$  states but ensures that  $K^\uparrow \neq \emptyset$ .

### 6.3 Summary

We have seen that, although in the general case  $L(G, \gamma^\infty) \neq \overline{K^\uparrow}$ , it is possible for some problems to choose a finite value for  $N$  such that an LLP supervisor allows behaviour in the plant that is identical to that allowed by an offline supervisor.

To start, we must ensure some simple characteristics for the problem, namely that  $K^\uparrow \neq \emptyset$  and that there is no starting error in  $L(G, \gamma^N)$ . The key is to calculate the appropriate measure,  $N_u(K)$ ,  $N_u(L(G))$ ,  $N_{mfc\bar{c}}$  or  $N_{mcmc}$ , and to use this to set a lower bound on the size of the lookahead window. We have seen that these measure are undefined for some problems, however, which forces us to conclude that these problems do not admit a valid LLP supervisor with the given attitude.

## 7 SYNTHESIZING AN LLP SUPERVISOR

GIVEN THAT OUR APPLICATION IS SAFETY-CRITICAL, it is tempting to think that our LLP supervisor should implement the conservative attitude. Theorems 11 and 12 make it clear, however, that both the optimistic attitude and conservative attitude can be used by a valid supervisor and that, in the limit, these result in the same language.

### 7.1 An Optimistic LLP Supervisor

In the first case, Theorem 11 tells us that we having  $N \geq N_{mcf\bar{c}} + 1$  will guarantee that an optimistic supervisor for the ATC problem will be valid. As has been discussed previously, the most dangerous situations that appear are when we allow aircraft to be arranged such that an arriving aircraft cannot be accommodated (see Example 4).

**Example 7** ( $N_{mcf\bar{c}}$  for the ATC problem).

$$\begin{aligned} s &= \epsilon \\ t &= \tilde{\rho}_1^3 \rho_1^3 \kappa_1 \lambda_1 \mu_1 \nu_1 \tilde{\rho}_2^3 \alpha_1 \rho_2^3 \kappa_2 \lambda_2 \mu_2 \nu_2 \\ &\quad \tilde{\rho}_3^3 \beta_1 \rho_3^3 \kappa_3 \lambda_3 \mu_3 \alpha_2 \tilde{\rho}_4^3 \rho_4^3 \gamma_1 \kappa_4 \lambda_4 \\ &\quad \tilde{\rho}_5^4 \nu_3 \rho_5^4 \end{aligned}$$

We are required to take  $s = \epsilon$  because  $K \cap L_c(G) = \emptyset$  for the ATC problem, which implies that  $K_{mc} = \emptyset$ .

The string  $t$  is the longest string that takes us from  $s$  to an uncontrollably crossing trace; this is achieved by having four aircraft arrive from the gate with the furthest possible travel time and array themselves in a “cross pattern” (Figure 9)). Now any action except for  $\eta_1$  runs into the problem that an uncontrollable notification regarding a fifth arriving aircraft can force us uncontrollably into illegal behaviour.

Based on our string  $t$ , we have  $N_{mcf\bar{c}} = |t| = 28$  and can therefore say that an optimistic LLP supervisor for the ATC problem should have a lookahead window of size

$$N \geq 29.$$

### 7.2 A Conservative LLP Supervisor

Theorem 12 tells us that  $N \geq N_{mcmc} + 1$  guarantees that a conservative LLP supervisor for the ATC problem will be valid. Unfortunately, as previously noted,  $K_{mc} = \emptyset$  for the ATC problem which means that  $N_{mcmc}$  is undefined. This reflects that at any point in the system, before we can reach a marked state we may uncontrollably




Figure 9: The cross pattern.

have another aircraft turn up in the airspace. Borrowing terminology from other control theory domains, the ATC problem exhibits orbital stability and not asymptotic stability.

Interestingly, this means that we cannot find a valid conservative LLP supervisor with the current problem formulation. An obvious way to adjust the problem would be to mark all states representing safe behaviour throughout the plant in order to allow for the calculation of  $N_{mcmc}$ .

**Example 8** ( $N_{mcmc}$  for modified ATC problem).

$$\begin{aligned}
 s &= \epsilon \\
 t &= \tilde{\rho}_1^3 \rho_1^3 \kappa_1 \lambda_1 \mu_1 \nu_1 \tilde{\rho}_2^3 \alpha_1 \rho_2^3 \kappa_2 \lambda_2 \mu_2 \nu_2 \\
 &\quad \tilde{\rho}_3^3 \beta_1 \rho_3^3 \kappa_3 \lambda_3 \mu_3 \alpha_2 \tilde{\rho}_4^3 \rho_4^3 \gamma_1 \kappa_4 \nu_3 \beta_2 \lambda_4 \mu_4 \alpha_3 \\
 &\quad \tilde{\rho}_5^3 \nu_3 \rho_5^3
 \end{aligned}$$

Although  $K_{mc} \neq \emptyset$  for the modified ATC problem,  $s = \epsilon$  still gives us the trace  $t$  of maximum length.

The string  $t$  is the longest string that takes us from  $s$  to a string in  $K_{mc}$ ; this is achieved by having four aircraft arrive and take as many actions as possible before a fifth aircraft arrives and prevents any uncontrollable actions from happening.

Based on  $t$ , we have  $N_{mcmc} = |t| = 32$  and can therefore say that a conservative LLP supervisor for the modified ATC problem should have a lookahead window of size

$$N \geq 33.$$

### 7.3 Summary

Of particular note in the last two sections is that, although the optimistic and conservative attitudes can both be used to produce  $L(G, \gamma^N) = \overline{K}^\dagger$ , the particular value of  $N$  required for each may be different. This results from the fact that the two attitudes need different kinds of information to produce valid supervisors; depending on the structure of the plant it may take more steps into the future for a given attitude to see what it needs to see.

## 8 EXTENSIONS TO LIMITED LOOKAHEAD SUPERVISORS

IN ADDITION TO THE MATERIAL PRESENTED IN THEIR seminal work,<sup>31</sup> the authors also published four notable extensions of LLP supervisors. These works introduced a technique to efficiently recompute the lookahead tree at every step, a more nuanced attitude towards pending strings, the ability to incorporate state information, and the ability to deal with partially observed plants.

<sup>31</sup> Chung et al., 1992a; and Chung et al., 1992b

### 8.1 Efficient Computation of the Lookahead Tree

The original authors presented a method for efficiently calculating the lookahead tree in a recursive manner, addressing a practical aspect of implementing the LLP control scheme.<sup>32</sup>

<sup>32</sup> Chung et al., 1993

---

Algorithm 2: Recursive calculation for an LLP supervisor

**Input:** An  $N$ -step lookahead tree where  $X^t(j)$  is the set of states at the  $j$ th level of the tree that was calculated after the  $t$ th event.  $j = 1, \dots, N$  and  $t = 0, 1, 2, \dots$

- 1: **for all** States in  $X^t(N)$  **do**
- 2:     **if** Conservative attitude **then**
- 3:         The cost of this state is  $\infty$ .
- 4:     **else if** Optimistic attitude **then**
- 5:         **if**  $x$  is an illegal state **then**
- 6:             The cost of this state is  $\infty$ .
- 7:         **else**
- 8:             The cost of this state is 0.
- 9: **for** Each layer,  $j$ , recursively back through the lookahead tree **do**
- 10:     **for all** States in  $X^t(j)$  **do**
- 11:         **if** Conservative attitude and this state's cost was 0 in the last computed tree **then**
- 12:             The cost of this state is 0.
- 13:         **else if** Optimistic attitude and this state's cost was  $\infty$  in the last computed tree **then**
- 14:             The cost of this state is  $\infty$ .

---

Algorithm 2 is adapted from Chung et al., 1993. To begin, the states in the  $N$ th layer of the tree have not been seen and are labelled according to the supervisor's attitude (Lines 1 to 8).

Finally, if the cost of the initial state is 0, then the least restrictive control policy can be found using the "cost-to-go" values; if the cost of the initial state is  $\infty$ , then a run-time error has occurred (Lines 23 to 26).

---

```

15:     else if This state is marked or if this state is transient and
      there is a controlled action that leads to a state with cost 0; and if
      no uncontrollable action leads to a state with cost  $\infty$  then
16:         The cost of this state is 0
17:     else
18:         The cost of this state is  $\infty$ .
19:     if The cost of all states in this layer are 0 then
20:         Set the cost of all states in preceding layers to 0 and termi-
      nate.
21:     else if The cost of all states in this layer match their costs in the
      last computed tree then
22:         Set the cost of all states in preceding layers to their costs in
      the last computed tree.
23: if The cost of the initial state is 0 then
24:     return The least restrictive control policy.
25: else
26:     return Run-time error

```

---

This recursive calculation (Algorithm 2) is based on converting the LLP control problem to a  $0/\infty$  optimal control problem. In this problem, the cost of a policy  $g(s)$  after trace  $s$  has occurred in the plant is 0 if  $s$  is legal, if  $g(s)$  permits all uncontrollable actions that can occur in the plant after the trace  $s$ , and if  $g(s)$  is non-empty whenever  $s$  is a prefix of a marked trace but not marked itself. All other policy/trace combinations are assigned a cost of  $\infty$ . This setup allows a dynamic programming approach to be used and makes it simple to apply when costs other than 0 and  $\infty$  are desired.

## 8.2 Variable Lookahead Policies

Chung et al. also extended their work by introducing variable lookahead policies (VLP) and the undecided cost,  $U$ .<sup>33</sup> This cost is added to the optimal control framework that was introduced for recursively computing the lookahead tree,<sup>34</sup> and is defined so that the following ordering exists amongst costs:

$$0 < U < \infty.$$

The power of the “undecided” cost is that pending strings can be assessed truthfully as opposed to arbitrarily labelling them based on supervisor attitude. This cost and its place in the ordering of costs reflect that an undecided string is not as desirable as a string with cost 0, but it is certainly more desirable than a string with cost  $\infty$ .

Additional refinements presented included allowing the size of the lookahead window to be unbounded and again making intelli-

<sup>33</sup> Chung et al., 1994

<sup>34</sup> Chung et al., 1993



---

Algorithm 3: VLP supervisor labels pending string,  $s$

**Input:** A pending string  $s$  that leads from the plant's current state to another state in the  $N$ -step lookahead tree.

```

1: if  $|s| = N$  then
2:   if  $s$  is not a prefix of a legal string then
3:     Cost-to-go of  $s$  is  $\infty$ 
4:   else
5:     Cost-to-go of  $s$  is  $U$ 
6: else if  $s$  leads to a marked, controllable state then
7:   Cost-to-go of  $s$  is 0
8: else if  $s$  is not a prefix of a legal string then
9:   Cost-to-go of  $s$  is  $\infty$ 
10: else
11:   if An uncontrollable event can occur then
12:     Cost-to-go of  $s$  is the maximum cost-to-go that occurs if an uncontrollable event follows  $s$ 
13:   else
14:     Cost-to-go of  $s$  is the minimum cost-to-go that occurs if a controllable event follows  $s$ 

```

---

gent reuse of previously computed trees unless the underlying plant has varied in the meantime. The authors presented experimental results from simulations and concluded that VLP supervisors expand significantly less of the lookahead tree than traditional LLP supervisors, that even a bounded VLP supervisor reduces the uncertainty in control policies, and that the unbounded VLP supervisor is the best supervisor on average.<sup>35</sup>

### 8.3 Using State Information

Work up to date had assumed that little to no information was available regarding the underlying plant, although this allows easy generalization, it also ignores useful information.<sup>36</sup>

---

Algorithm 4: Online computation of control policy by VLP-S

**Input:** Initial state  $x$  and underlying plant  $G$

```

1: Add all uncontrollable events that can occur in state  $x$  to the control policy
2: if We don't know the cost-to-go for  $x$  then
3:   Calculate the cost-to-go of  $x$  using Algorithm 5
4: if The cost-to-go of  $x$  is not  $\infty$  then
5:   for all Controllable actions  $\sigma$  that can occur in state  $x$  do
6:     if We don't know the cost-to-go for  $\delta(x, \sigma)$  then
7:       Calculate the cost-to-go of  $\delta(x, \sigma)$  using Algorithm 5
8:     if The cost-to-go of  $\delta(x, \sigma)$  is 0 then
9:       Add  $\sigma$  to the control policy
10: return the control policy

```

---

Algorithm 3 is adapted from Chung et al., 1994. Similar to the calculations in Algorithm 2, a cost is assigned directly to  $s$  if it leads to the edge of the lookahead tree (Lines 1 to 5). If the string leads to a state earlier on in the lookahead tree, though, then the costs-to-go of future states are used to determine the cost-to-go of the current state (Lines 6 to 14).

<sup>35</sup> Chung et al., 1994

<sup>36</sup> Hadj-Alouane et al., 1994

Algorithm 4 is adapted from Hadj-Alouane et al., 1994. To calculate a state's cost-to-go, the lookahead window is first constructed by recursively expanding the plant  $G$  from the current state  $x$ . This can be done with customizable stopping conditions to ensure that the procedure halts and that states are only expanded if necessary (Line 1). Once this is done, all of the illegal states, blocking states, and states that lead uncontrollably to these are assigned the cost-to-go  $\infty$ . Any remaining states are assigned the cost-to-go 0, which allows the cost-to-go of the initial state to be returned.

To address this, Hadj-Alouane et al. presented variable lookahead policies with state information (VLP-S), which allows reduced computation whenever states appear in multiple locations in the tree (Algorithm 4).<sup>37</sup> The VLP-S supervisor constructs the control policy for the current state by allowing all uncontrollable events that can occur in the plant and then any controllable actions that lead to states with a cost-to-go of 0.

---

Algorithm 5: Calculating the cost-to-go of a state with VLP-S

**Input:** State  $x$  and underlying plant  $G$

- 1: Recursively expand the subplant of  $G$  from the state  $x$
  - 2: Assign all illegal states the cost-to-go  $\infty$
  - 3: **do**
  - 4:   Assign all states that lead uncontrollably to a state with cost-to-go  $\infty$  the same cost-to-go
  - 5:   Assign all blocking states the cost-to-go  $\infty$
  - 6: **while** Any states were assigned a cost at Line 5
  - 7: Assign all uncosted states the cost-to-go 0
  - 8: **return** the cost-to-go for state  $x$
- 

<sup>37</sup> Hadj-Alouane et al., 1994

Algorithm 5 is adapted from Hadj-Alouane et al., 1994.

Algorithm 4, Lines 2 and 6 are of particular note. The supervisor checks to see if the cost-to-go of a given state can be reused from a previous computation, which allows previously computed costs-to-go to be reused for future iterations of Algorithm 4.

## 8.4 Dealing with Partial Observation

The last direct extension that we discuss is Hadj-Alouane et al.'s work to enable variable lookahead policies to consider plants under partial observation (VLP-PO).<sup>38</sup> The authors present this as an extension of their previous work, VLP-S, to the specific case where the legal language is prefix-closed.

Since DES problems with partial observation do not have supremal controllable and observable sublanguages, these policies generate instead a maximal controllable and observable sublanguage of the legal language on the basis of an argument ordering of the controllable events.<sup>39</sup>

Only minimal changes are necessary: first, the set of states that the plant could be in is computed based on the last observable event; then, the control action is computed in the manner of VLP-S. These computations are adapted to the partially observed nature of the plant and with respect to a global ordering provided for all events. Finally, the set of states that the plant could be in until the next observable event is computed, which will be an argument the next time that VLP-PO is called.

<sup>38</sup> Hadj-Alouane et al., 1996

Varying this ordering can lead to a different maximal sublanguage; the authors highlight that if this ordering is done with care then a desirable maximal sublanguage can be obtained depending on the domain-specific meaning of "desirable."

<sup>39</sup> Hadj-Alouane et al., 1996

The paper also presented a distributed version of VLP-PO and showed that it produced the same policies produced by the sequential VLP-PO. Both versions were demonstrated on a resource allocation problem, which allowed the authors to estimate that the distributed version would require a tenth of the run-time required by the sequential VLP-PO.

## 9 ADDITIONAL APPLICATIONS OF LIMITED LOOKAHEAD

MANY OTHER AUTHORS have contributed to the literature on limited lookahead policies, building on Chung et al.'s broad conceptual base. We organize these into four categories (Table 1) and cover them in the following sections.

---

### Different Underlying Plants

Modifying an optimal controller to provide control for a partially observed system.

Decentralized and modular plants.

Probabilistic DES.

Hybrid systems.

Finite State Machines with Variables.

Partially observed Petri Nets with forbidden states.

Fault-tolerant supervisors.

---

### Different Online Supervisors

No explicit calculation of the supremal controllable sublanguage.

Near-optimal control of dynamic DES .

Enacting robust control.

Learning optimal LLP control using reinforcement learning.

---

### Different Calculations

Extending rather than truncating traces.

Estimating the lookahead tree's state-space size.

---

### Applications

Robots — Navigation and task allocation.

Software – Fault detection and enforcing concurrency restraints.

---

Table 1: An overview of the literature investigating and applying LLP supervisors.

### 9.1 Different Underlying Plants

Heymann and Lin noted that synthesizing supervisors for partially observed plants is an NP-hard problem and presented a process for modifying an optimal controller of the original, fully-observed system in an online fashion to provide control for a partially observed system. This process runs in  $\mathcal{O}(n)$  time, where  $n$  is the number of states contained in the the plant composed with the legal specification.<sup>40</sup> This was improved upon by Ushio, whose method always produces maximal controllable and observable sublanguages.<sup>41</sup>

For complex plants resulting from  $n$  disjoint subsystems, Minhas and Wonham showed how modular calculations could be used to determine conditions on the size of the lookahead window.<sup>42</sup> Online decentralized supervisors can be synthesized for these plants even if

<sup>40</sup> Heymann and Lin, 1994

<sup>41</sup> Ushio, 1999

<sup>42</sup> Minhas and Wonham, 2003

their total structure is unknown beforehand.<sup>43</sup>

The underlying plant might also be probabilistic. Winacott and Rudie compared the performance of an LLP supervisory control approach named Recursive Utility Based Limited Lookahead (RUBLL) in this setting against the optimal solution derived using Markov Chain analysis. They concluded that RUBLL's performance was consistent with the optimal solution but was more readily applied to scenarios where exhaustive approaches are infeasible.<sup>44</sup>

For hybrid systems, such as a switching circuit, Dupuis and Fan compared an LLP supervisor's performance with that of an evolved finite state controller. They concluded that the LLP supervisor was better adapted to systems with slow dynamics and a vector field that flows towards the target, although their study only considered a 1-step lookahead policy for the LLP supervisor.<sup>45</sup>

Finally, for partially observed Petri Nets Ru et al. synthesized supervisors that allow for control to be effected for any finite set of forbidden states.<sup>46</sup>

<sup>43</sup> Dai and Lin, 2014

<sup>44</sup> Winacott and Rudie, 2009

<sup>45</sup> Dupuis and Fan, 2010

<sup>46</sup> Ru et al., 2014

## 9.2 Different Online Supervisors

Gu et al. propose a method for online synthesis of DES supervisors that divides the strings in a supervisor's lookahead tree into three types. Type I strings are those that include forbidden states and/or forbidden event strings; Type II strings include neither forbidden states nor forbidden event strings; and Type III strings include forbidden event strings when the past behaviour of the plant is considered. Obviously the strings of Types I and III are illegal, while those of Type II need to be tested for legality. The authors propose such a test based on the controllability of the related trace and claim that their method "realizes optimal supervisory control without calculating the supremal controllable sublanguage."<sup>47</sup>

An optimal control framework is natural to consider if the supervisor must maximize the reward gained instead of guaranteeing legal plant behaviour. Grigorov and Rudie proposed a method to effect (near-)optimal control of dynamic DES, considering aspects such as system reliability and normalizing for risk. A key takeaway was that looking too far into the future and not far enough can both cause problems when planning.<sup>48</sup>

Zhao et al. applied limited lookahead to deterministic finite state machines with variables and showed how they could enforce safety for a power grid in an online manner.<sup>49</sup>

Boroomand and Hashtrudi-Zad demonstrated that their Robust Limited Lookahead (RLL) supervisor could enact nonblocking robust control by taking a conservative attitude towards pending strings in

<sup>47</sup> Gu Tianlong et al., 1996

A dynamic DES is one where the structure of the underlying plant is itself time-varying.

<sup>48</sup> Grigorov and Rudie, 2006

<sup>49</sup> Zhao et al., 2012

the lookahead tree. They also showed that RLL resulted in maximally permissive policies if the lookahead window was of size  $N_{nnf}$ , the length of the longest neighbouring nonblocking frontier trace.<sup>50</sup>

<sup>50</sup> Boroomand and Hashtrudi-Zad, 2013

In order to allow for supervisors to adapt as they enact control, Umemoto and Yamasaki showed that a supervisor could learn an optimal LLP by using reinforcement learning. Specifically, for a non-stationary environment the LLP supervisor is able to learn the state transition probabilities, expected rewards and expected costs of enacting control. The authors conducted simulations that showed how different learning rates trade convergence rate for asymptotic performance.<sup>51</sup>

<sup>51</sup> Umemoto and Yamasaki, 2015

For situations where uncontrolled faults can occur in the plant, Dai et al. considered whether limited lookahead could be used to achieve fault-tolerance in supervisors. Their proposed method contained three parts: a learned nominal supervisor; a learned fault detector; and a post-fault supervisor. Necessary and sufficient conditions are presented for the existence of a satisfactory post-fault supervisor.<sup>52</sup>

<sup>52</sup> Dai et al., 2016

### 9.3 Different Calculations

Kumar et al. approached the idea of LLP from a notion of extending rather than truncating, as Chung et al. and most other surveyed works had done. The result of this approach is an extension based Limited Lookahead (ELL) supervisor, which estimates the plant behaviour by adding all finite-length event sequences to the N-step projection of plant behaviour. The key benefit of this approach is that it obviates the need for the supervisor to take an attitude by ensuring that no pending strings occur while requiring the same order of computation as the traditional LLP supervisor.<sup>53</sup>

<sup>53</sup> Kumar et al., 1998

Addressing a topic of practical concern, Winacott et al. presented and analysed a method for estimating the state-space size for an LLP supervisor's lookahead tree. This method is based on the underlying plant's adjacency matrix and depends on estimating a parameter  $\tau$  that captures the state space's size.<sup>54</sup>

<sup>54</sup> Winacott et al., 2013

### 9.4 Applications

Besides theoretical developments, there have also been a number of works applying LLP theory to specific problem domains.

For robotics, Kobayashi and Ushio applied LLP to a Petri net model of mobile robots navigating through a building,<sup>55</sup> while Tsalatsanis et al. used LLP to dynamically allocate tasks to a team of robots.<sup>56</sup>

<sup>55</sup> Kobayashi and Ushio, 2000

<sup>56</sup> Tsalatsanis et al., 2012

For software applications, Zhao et al. used limited lookahead DES to actively monitor software in order to predict faults and avoid them

before they occurred.<sup>57</sup> Auer et al. considered the problem of multithreaded applications and applied limited lookahead to guarantee that concurrency restraints were enforced.<sup>58</sup>

<sup>57</sup> Zhao et al., 2010

<sup>58</sup> Auer et al., 2014

## 10 CONCLUSION

IN THIS TUTORIAL we introduced the basic terminology required when discussing LLP supervisors. Although the underlying ideas for LLP supervisors are intuitive, we presented Chung et al.'s formal proofs to rigorously confirm their correctness. We demonstrated that both the optimistic and conservative attitudes can be used to produce valid supervisors as well as the conditions for which this is the case.

The ATC problem illustrated Chung et al.'s ideas throughout the tutorial, including: characterizing the underlying problem, recognizing when the problem did not allow for LLP control and highlighting that the required window size may vary based on the supervisor's attitude. The ATC problem also illuminated some possible routes for developing the theory of LLP supervision, including the idea that some plants may exhibit a kind of orbital stability. In such cases it would be useful to bound the window size  $N$  which guarantees that the supervisor can prevent illegal behaviour and ensure that each of its subplants reaches their respective marked states in finite time without requiring that the plant as a whole reaches its marked state.

Finally, we reviewed extensions to Chung et al.'s original paper, many by the original authors, as well as applications of this theory to a wide variety of specific problems and problem classes. We believe that this tutorial will enable other DES practitioners to make use of LLP supervisors when appropriate in their own work.

## ACKNOWLEDGMENTS

The authors acknowledge that Queen's University is situated on traditional Anishinaabe and Haudenosaunee Territory.

All DFA figures were produced using the Integrated Discrete Event Systems (IDES) software.<sup>59</sup>

IDES is available under the AGPL-3.0 open source license at <https://github.com/krudie/IDES>.

<sup>59</sup> Rudie, 2006

This research was supported by the Natural Sciences and Engineering Research Council of Canada as well as the Faculty of Engineering and Applied Science at Queen's University.

## REFERENCES

- Anthony Auer, Juergen Dingel, and Karen Rudie. Concurrency control generation for dynamic threads using discrete-event systems. *Science of Computer Programming*, 82:22–43, 2014. DOI: 10.1016/j.scico.2013.01.007.
- Farzam Boroomand and Shahin Hashtrudi-Zad. A Limited Lookahead Policy in Robust Nonblocking Supervisory Control of Discrete Event Systems. In *Proceedings of the American Control Conference*, pages 935–939. IEEE, Jun 2013. DOI: 10.1109/ACC.2013.6579956.
- Sheng-Luen Chung, Stéphane Lafortune, and Feng Lin. Addendum to "Limited Lookahead Policies in Supervisory Control of Discrete Event Systems": Proofs of Technical Results (Report No. CGR-92-6). Technical report, College of Engineering, University of Michigan, Ann Arbor, Michigan, 1992a.
- Sheng-Luen Chung, Stéphane Lafortune, and Feng Lin. Limited lookahead policies in supervisory control of discrete event systems. *IEEE Transactions on Automatic Control*, 37(12):1921–1935, 1992b. DOI: 10.1109/9.182478.
- Sheng-Luen Chung, Stéphane Lafortune, and Feng Lin. Recursive computation of limited lookahead supervisory controls for discrete event systems. *Discrete Event Dynamic Systems: Theory and Applications*, 3(1):71–100, Mar 1993. DOI: 10.1007/BF01439177.
- Sheng-Luen Chung, Stéphane Lafortune, and Feng Lin. Supervisory control using variable lookahead policies. *Discrete Event Dynamic Systems: Theory and Applications*, 4(3):237–268, 1994. DOI: 10.1007/BF01438709.
- Jin Dai and Hai Lin. A learning-based synthesis approach to decentralized supervisory control of discrete event systems with unknown plants. *Control Theory and Technology*, 12(3):218–233, 2014. DOI: 10.1007/s11768-014-4082-2.
- Jin Dai, Ali Karimodini, and Hai Lin. Achieving fault-tolerance and safety of discrete-event systems through learning. *Proceedings of the American Control Conference*, 2016-July:4835–4840, 2016. DOI: 10.1109/ACC.2016.7526118.
- Jean François Dupuis and Zhun Fan. Comparing an evolved finite state controller for hybrid system to a lookahead design. *IEEE Congress on Evolutionary Computation*, 2010. DOI: 10.1109/CEC.2010.5586402.
- Lenko Grigorov and Karen Rudie. Near-Optimal Online Control of Dynamic Discrete-Event Systems. *Discrete Event Dynamic Systems*, 16(4):419–449, Dec 2006. DOI: 10.1007/s10626-006-0020-x.
- Gu Tianlong, Gao Jinchang, and Zhou Chunhui. On-line synthesis of supervisors for discrete events in automated manufacturing systems. In *Proceedings of the IEEE International Conference on Industrial Technology (ICIT'96)*, pages 176–179, Shanghai, China, 1996. IEEE. DOI: 10.1109/ICIT.1996.601566.
- Nejib Ben Hadj-Alouane, Stéphane Lafortune, and Feng Lin. Variable lookahead supervisory control with state information. *IEEE Transactions on Automatic Control*, 39(12):2398–2410, 1994. DOI: 10.1109/9.362854.

- Nejib Ben Hadj-Alouane, Stéphane Lafortune, and Feng Lin. Centralized and distributed algorithms for on-line synthesis of maximal control policies under partial observation. *Discrete Event Dynamic Systems*, 6(4):379–427, Oct 1996. DOI: 10.1007/BF01797138.
- Michael Heymann and Feng Lin. On-line control of partially observed discrete event systems. *Discrete Event Dynamic Systems: Theory and Applications*, 4(3):221–236, Jul 1994. DOI: 10.1007/BF01438708.
- Keigo Kobayashi and Toshimistu Ushio. An application of LLP supervisory control with Petri net models in mobile robots. In *IEEE International Conference on Systems, Man and Cybernetics*, volume 4, pages 3015–3020, Nashville, USA, 2000. IEEE. DOI: 10.1109/ICSMC.2000.884460.
- Ratnesh Kumar, Hok M. Cheung, and Steven I. Marcus. Extension based limited lookahead supervision of discrete event systems. *Automatica*, 34(11):1327–1344, Nov 1998. DOI: 10.1016/S0005-1098(98)00077-6.
- Rajinderjeet Minhas and W.M. Wonham. Online supervision of discrete event systems. In *Proceedings of the American Control Conference*, volume 2, pages 1685–1690, Denver, USA, 2003. IEEE. DOI: 10.1109/ACC.2003.1239836.
- Yu Ru, Maria Paola Cabasino, Alessandro Giua, and Christoforos N. Hadjicostis. Supervisor synthesis for discrete event systems under partial observation and arbitrary forbidden state specifications. *Discrete Event Dynamic Systems: Theory and Applications*, 24(3):275–307, 2014. DOI: 10.1007/s10626-012-0152-0.
- Karen Rudie. The Integrated Discrete-Event Systems Tool. *8th International Workshop on Discrete Event Systems*, pages 394–395, 2006. DOI: 10.1109/WODES.2006.382403.
- Athanasios Tsalatsanis, Ali Yalcin, and Kimon P. Valavanis. Dynamic task allocation in cooperative robot teams. *Robotica*, 30(5):721–730, 2012. DOI: 10.1017/S0263574711000920.
- Hijiri Umemoto and Tatsushi Yamasaki. Optimal LLP Supervisor for Discrete Event Systems Based on Reinforcement Learning. In *IEEE International Conference on Systems, Man, and Cybernetics*, pages 545–550, Kowloon, China, Oct 2015. IEEE. DOI: 10.1109/SMC.2015.106.
- Toshimitsu Ushio. On-Line Control of Discrete Event Systems with a Maximally Controllable and Observable Sublanguage. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, E82-A(9):1965–1970, 1999.
- Creag Winacott and Karen Rudie. Limited lookahead supervisory control of probabilistic discrete-event systems. In *47th Annual Allerton Conference on Communication, Control, and Computing*, pages 660–667, Monticello, USA, 2009. IEEE. DOI: 10.1109/ALLERTON.2009.5394925.
- Creag Winacott, Behnam Behinaein, and Karen Rudie. Methods for the estimation of the size of lookahead tree state-space. *Discrete Event Dynamic Systems*, 23(2):135–155, Jun 2013. DOI: 10.1007/s10626-012-0138-y.



Changzhi Zhao, Wei Dong, and Zhichang Qi. Active Monitoring for Control Systems under Anticipatory Semantics. In *10th International Conference on Quality Software*, pages 318–325, Zhangjiajie, China, Jul 2010. IEEE. DOI: 10.1109/QSIC.2010.82.

Junhui Zhao, Yi Liang Chen, Zhong Chen, Feng Lin, Caisheng Wang, and Hongwei Zhang. Modeling and control of discrete event systems using finite state machines with variables and their applications in power grids. *Systems and Control Letters*, 61(1):212–222, 2012. DOI: 10.1016/j.sysconle.2011.10.010.