Available online at www.sciencedirect.com
SCIENCE DIRECT®
automatica
Automatica 41 (2005) 1439–1445
www.elsevier.com/locate/automatica

Brief paper

Robust optimal control of regular languages☆

Constantino M. Lagoa[a,*], Jinbo Fu[b], Asok Ray[b]

[a] *Department of Electrical Engineering, Pennsylvania State University, University Park, PA 16802, USA*
[b] *Department of Mechanical and Nuclear Engineering, Pennsylvania State University, University Park, PA 16802, USA*

Received 18 November 2003; received in revised form 9 November 2004; accepted 15 March 2005

Abstract

This paper presents an algorithm for robust optimal control of regular languages under specified uncertainty bounds on the event cost parameters of the language measure that has been recently reported in literature. The performance index for the proposed robust optimal policy is obtained by combining the measure of the supervised plant language with uncertainty. The performance of a controller is represented by the language measure of the supervised plant and is minimized over the given range of event cost uncertainties. Synthesis of the robust optimal supervisory control policy requires at most n iterations, where n is the number of states of the deterministic finite-state automaton (DFSA) model, generated from the regular language of the unsupervised plant behavior. The computational complexity of the control synthesis method is polynomial in n.

Keywords: Discrete event supervisory control; Optimal control; Robust control; Language measure

1. Introduction

This paper addresses the problem of robust supervisory control of regular languages, representing deterministic finite-state automata (DFSA) models of the physical plants under decision and control. Specifically, algorithms are formulated for both robust analysis and optimal robust supervisor synthesis for regular languages, or equivalently, for their DFSA models. Also, mathematical foundations of these algorithms are rigorously established.

Recent results on the analysis of DFSA and on the synthesis of optimal supervisory control policies motivate the work presented in this paper. More precisely, a novel way of accessing the performance of DFSA has been proposed in Wang and Ray (2004) and Ray and Phoha (2003), where a signed real measure of regular languages has been developed. This measure is computed using an event cost matrix and a characteristic vector and it provides a quantitative tool for evaluating the performance of regular languages, representing discrete-event dynamic behavior of DFSA plant models. This work was followed by Fu, Ray, and Lagoa (2004) where, based on this performance measure, an algorithm is developed for the design of optimal supervisors. The optimal controller is obtained by selectively disabling controllable events in order to maximize the overall performance. However, uncertainty was not addressed. This paper considers structural uncertainty in the DFSA model: (i) uncertainty in the presence of some of the state transitions and (ii) uncertainty in the entries of the event cost matrix (and hence the state transition matrix) used to compute the performance of the system. The first source of uncertainty results from inaccuracies in modelling of the discrete-event dynamic behavior of the plant under supervision. The second source of uncertainty is inherent to the process of parameter

☆ This paper was not presented at any IFAC meeting. This paper is recommended for publication in revised form by Associate Editor Xiren Cao under the direction of Editor Ian Petersen. This work has been supported in part by the Army Research Office under Grant No. DAAD19-01-1-0646, National Science Foundation under Grants Nos. ECS-9984260 and ANI-0125653, and NASA Glenn Research Center under Grant Nos. NAG3-2448 and NNC04GA49G.
* Corresponding author.
E-mail addresses: lagoa@engr.psu.edu (C.M. Lagoa), jxf293@psu.edu (J. Fu), axr2@psu.edu (A. Ray).

0005-1098/$ - see front matter © 2005 Elsevier Ltd. All rights reserved.
doi:10.1016/j.automatica.2005.03.016

identification of the event cost matrix. Since the entries of the event cost matrix can be often determined using a scaled version of the probabilities obtained by Monte Carlo simulations, their values are only known as bounded intervals with a specified level of statistical confidence. The robust supervisory control algorithm, presented in this paper, has a complexity that is polynomial in the number of states of the DFSA model.

1.1. Previous work

The problem of robust control of discrete-event dynamical systems (DEDS) has been addressed by several researchers. Park and Lim (2002) have studied the problem of robust control of nondeterministic DEDS. The performance measure used was nonblocking property of the supervised plant. Necessary and sufficient conditions for existence of a robust nonblocking controller for a given finite set of plants are provided. However, no algorithm for controller design is provided. The problem of designing nonblocking robust controllers was also addressed by Cury and Krogh (1999) with the additional constraint that the infinite behavior belongs to a given set of allowable behaviors. In this work, the authors concentrated on the problem of designing a controller that maximizes the set of plants for which their supervised behavior belong to the admissible set of behaviors. Takai (1999) addresses a similar problem. However, it considers the whole behavior (not just the infinite behavior) and it does not consider nonblockingness. Lin (1993) adopted a different approach, where both the set of admissible plants and the performance are defined in terms of the marked language. Taking the set of admissible plants as the plants whose marked language is in between two given behaviors, the authors provided conditions for solvability of the problem of designing a discrete event supervisory controller such that the supervised behavior of any of the admissible plants contains a desired behavior K.

To address a subject related to that of this paper several researchers have proposed optimal control algorithms for DFSA based on different assumptions. Some of these researchers have attempted to quantify the controller performance using different types of cost assigned to the individual events. Passino and Antsaklis (1989) proposed path costs associated with state transitions and hence optimal control of a discrete event system is equivalent to following the shortest path on the graph representing the uncontrolled system. Kumar and Garg (1995) introduced the concept of payoff and control costs that are incurred only once regardless of the number of times the system visits the state associated with the cost. Consequently, the resulting cost is not a function of the dynamic behavior of the plant. Brave and Heyman (1993) introduced the concept of optimal attractors in discrete-event control. Sengupta and Lafortune (1998) used control cost in addition to the path cost in optimization of the performance index for trade-off between finding the shortest path and reducing the control cost.

A limitation of the work mentioned above is that the controllers are designed so that the closed loop system has certain specified characteristics. No performance measure is given that can compare the performance of different controllers. To address this issue, Wang and Ray (2004) and Ray and Phoha (2003) have proposed a signed real measure for regular languages. This novel tool of addressing the performance of DFSAs enables the developing of a new approach to supervisor design. The design of optimal supervisor has been reported by Fu, Ray, and Lagoa (2003) and Fu et al. (2004) for without and with event disabling cost, respectively. This paper extends the previous work on optimal control by including robustness in the problem of designing supervisors in the presence of uncertainty.

2. Brief review of the language measure

This section briefly reviews the concept of signed real measure of regular languages introduced in Wang and Ray (2004). Let the plant behavior be modelled as a DFSA $G_i \equiv (Q, \Sigma, \delta, q_i, Q_m)$ where Q is the finite set of states with $|Q| = n$ excluding the dump state (Ramadge & Wonham, 1987) if any, and $q_i \in Q$ is the initial state; Σ is the (finite) alphabet of events; Σ^* is the set of all finite-length strings of events including the empty string ε; the (possibly partial) function $\delta : Q \times \Sigma \to Q$ represents state transitions and $\hat{\delta}^* : Q \times \Sigma^* \to Q$ is an extension of δ; and $Q_m \subseteq Q$ is the set of marked states. The set Q_m is partitioned into Q_m^+ and Q_m^-, where Q_m^+ contains all *good* marked states that are desired to terminate on and Q_m^- contains all *bad* marked states that are not desired to terminate on, although it may not always be possible to avoid terminating on the bad states while attempting to reach the good states. The marked language associated with DFSA G_i $L_m(G_i)$ is partitioned into $L_m^+(G_i)$ and $L_m^-(G_i)$ consisting of good and bad strings that, starting from q_i, terminate on Q_m^+ and Q_m^-, respectively.

The language of all strings that starts at a state $q_i \in Q$ and terminates on a state $q_j \in Q$, is denoted as $L(q_i, q_j)$. That is, $L(q_i, q_j) \equiv \{s \in L(G_i) : \hat{\delta}^*(q_i, s) = q_j\}$. Furthermore, consider a characteristic function $\chi : Q \to [-1,\ 1]$ satisfying

$$\chi(q_j) \in \begin{cases} (0, 1] & \text{if } q_j \in Q_m^+, \\ \{0\} & \text{if } q_j \notin Q_m, \\ [-1, 0) & \text{if } q_j \in Q_m^-. \end{cases}$$

Now, the event cost $\tilde{\pi} : \Sigma^* \times Q \to [0, 1]$ is defined as

- $\tilde{\pi}[\sigma_k|q_j] = 0$ if $\delta(q_j, \sigma_k)$ is undefined; $\tilde{\pi}[\varepsilon|q_j] = 1$;
- $\tilde{\pi}[\sigma_k|q_j] \equiv \tilde{\pi}_{jk} \in [0, 1)$; $\sum_k \tilde{\pi}_{jk} < 1$;
- $\tilde{\pi}[\sigma_k\ s|q_j] = \tilde{\pi}[\sigma_k|q_j]\tilde{\pi}[s|\delta(q_j, \sigma_k)]$.

Given this, the signed real measure μ of a singleton string set $\{s\} \subset L(q_i, q_j) \subseteq L(G_i) \in 2^{\Sigma^*}$ is defined as

$$\mu(\{s\}) \equiv \chi(q_j)\, \tilde{\pi}(s|q_i) \quad \forall s \in L(q_i, q_j).$$

The signed real measure of $L(q_i, q_j)$ is defined as

$$\mu(L(q_i, q_j)) \equiv \sum_{s \in L(q_i, q_j)} \mu(\{s\})$$

and the signed real measure of a DFSA G_i, initialized at the state $q_i \in Q$, is denoted as

$$\mu_i \equiv \mu(L(G_i)) = \sum_j \mu(L(q_i, q_j)).$$

Taking $\mu \equiv [\mu_1\ \mu_2 \cdots \mu_n]^{\mathrm{T}}$, it was proven in Wang and Ray (2004) that

$$\mu = \Pi\mu + X,$$

where Π is an $n \times n$ matrix whose (i, j) entry is

$$\pi_{jk} \equiv \pi(q_k|q_j) = \sum_{\sigma \in \Sigma : \delta(q_j, \sigma) = q_k} \tilde{\pi}(\sigma|q_j)$$

and

$$\pi_{jk} = 0 \quad \text{if } \{\sigma \in \Sigma : \delta(q_j, \sigma) = q_k\} = \emptyset$$

and $X \equiv [\chi_1\ \chi_2 \cdots \chi_n]^{\mathrm{T}}$. So that the vector μ is well defined, it is assumed that there exist a $\theta > 0$ such that, for all i

$$\sum_j \pi_{ij} \leqslant (1 - \theta).$$

Remark 1. It should be noted that if one defines the function (or operator) $T : \Re^n \rightarrow \Re^n$

$$T(x) \equiv \Pi x + X,$$

finding the language measure μ is equivalent to finding the fixed point of T. This observation plays an important role in the proofs of the results in this paper.

3. Modelling uncertainty

This section introduces the model of uncertainty of a DFSA along with the concept of robust performance. Uncertainty in the event cost matrix and in the existence of state transitions are the types of uncertainty studied in this paper. Other types of uncertainties (e.g., those due to the number of states and controllability of the events) that are not addressed here are topics of future work.

3.1. Uncertainty

As mentioned above, two types of uncertainties are considered in the model used in this paper: Uncertainty in the presence of state transitions and uncertainty in the event cost matrix. These two sources of uncertainty can be modelled by modifying the definition of the measure used to compute the performance. More precisely, define the *uncertain event cost* as

$$\tilde{\pi}^{\Delta}[\sigma_j|q_i] = (\tilde{\pi}^0[\sigma_j|q_i] + \Delta_{cost}[\sigma_j|q_i])\Delta_{model}[\sigma_j|q_i],$$

where $\tilde{\pi}^0[\sigma_j|q_i]$ is the event cost of the nominal model, $\Delta_{cost}[\sigma_j|q_i]$ represents the uncertainty associated with the determination of the event cost and belongs to an interval that is a proper subset of [0, 1]. Finally, $\Delta_{model}[\sigma_j|q_i]$ represents uncertainty in the existence of this specific transition of the automaton, i.e., if this transition is always present then

$$\Delta_{model}[\sigma_j|q_i] = 1.$$

If there is uncertainty in the presence of this transition, then

$$\Delta_{model}[\sigma_j|q_i] \in \{0, 1\}.$$

Furthermore, it is assumed that the uncertain event costs vary independently, i.e., each uncertain parameter only enters in one event cost. Finally, it is also assumed that there exists a constant $\theta > 0$ such that for all admissible uncertainty values and all i

$$\sum_j \sum_{\sigma \in \Sigma : \delta(q_i, \sigma) = q_j} \tilde{\pi}^{\Delta}[\sigma|q_i] \leqslant 1 - \theta.$$

The set of admissible values for the uncertainty Δ is denoted by $\mathbf{\Delta}$ and is a compact subset of $[0, 1]^{|\Sigma| \times n} \times \{0, 1\}^{|\Sigma| \times n}$. Now, given any $\Delta \in \mathbf{\Delta}$, the uncertain state transition matrix is given by

$$\Pi(\Delta) = \begin{bmatrix} \pi_{11}(\Delta) & \pi_{12}(\Delta) & \dots & \pi_{1n}(\Delta) \\ \pi_{21}(\Delta) & \pi_{22}(\Delta) & \dots & \pi_{2n}(\Delta) \\ \vdots & \vdots & \ddots & \vdots \\ \pi_{n1}(\Delta) & \pi_{n2}(\Delta) & \dots & \pi_{nn}(\Delta) \end{bmatrix},$$

where

$$\pi_{ij}(\Delta) = \sum_{\sigma \in \Sigma : \delta(q_i, \sigma) = q_j} \tilde{\pi}^{\Delta}[\sigma|q_i].$$

3.2. Additional notation

Given a supervisor S, let $\Pi(S, \Delta)$ be the uncertain state transition matrix under supervisor S, i.e., $\Pi(S, \Delta)$ has entries

$$\pi_{ij}(S, \Delta) = \sum_{\sigma \in \Sigma : \delta(q_i, \sigma) = q_j} \tilde{\pi}^{S, \Delta}[\sigma|q_i],$$

where

$$\tilde{\pi}^{S, \Delta}[\sigma|q_i] = \begin{cases} 0 & \text{if } \sigma \text{ controllable and disabled by } S, \\ \tilde{\pi}^{\Delta}[\sigma|q_i] & \text{otherwise.} \end{cases}$$

For a given admissible value for the uncertainty Δ, the performance of the plant under the supervisor S, denoted by $\mu(S, \Delta)$, is the solution of

$$\mu(S, \Delta) = \Pi(S, \Delta)\mu(S, \Delta) + X.$$

4. Robust performance

In this section, a precise definition of robust performance is provided and an algorithm for computing is presented.

4.1. Definition of robust performance

Consider an uncertain automaton controlled by a supervisor S. The *robust performance of supervisor S*, denoted by $\underline{\mu}(S)$, is defined as the worst-case performance, i.e.,

$$\underline{\mu}(S) = \min_{\Delta \in \mathbf{\Delta}} \mu(S, \Delta),$$

where the above minimum is taken elementwise. Even though the minimization is done element by element, this performance is achieved for some $\Delta^* \in \mathbf{\Delta}$. The precise statement of this result is given below and its proof is provided in Section 5.2.

Lemma 2. *Let S be a supervisor. Then, there exists a* $\Delta^* \in \mathbf{\Delta}$ *such that, for all admissible* $\Delta \in \mathbf{\Delta}$,

$$\underline{\mu}(S) = \mu(S, \Delta^*) \leqslant \mu(S, \Delta),$$

where the above inequality is implied elementwise.

An algorithm for computing $\underline{\mu}(S)$ is presented below.

Algorithm 1. *Computation of the worst-case performance of supervisor S.*

Step 0. *Let* $k = 0$ *and select* $\Delta^0 \in \mathbf{\Delta}$.

Step 1. *Let* Δ^{k+1} *be such that*[1]

$$\pi_{ij}(S, \Delta^{k+1}) = \begin{cases} \max_{\Delta \in \mathbf{\Delta}} \pi_{ij}(S, \Delta) & \text{if } \mu_j(S, \Delta^k) < 0, \\ \min_{\Delta \in \mathbf{\Delta}} \pi_{ij}(S, \Delta) & \text{if } \mu_j(S, \Delta^k) \geqslant 0. \end{cases}$$

Step 2. *If* $\mu(S, \Delta^{k+1}) = \mu(S, \Delta^k)$ *then* $\underline{\mu}(S) = \mu(S, \Delta^{k+1})$ *and stop. Else let* $k \leftarrow k+1$ *and go to Step* 1.

The theorem below presents the formal result that indicates that the above algorithm converges in a finite number of steps.

Theorem 3. *Given a supervisor S, the above algorithm converges to its robust performance, i.e.,*

$$\mu(S, \Delta^k) \to \underline{\mu}(S).$$

Furthermore, it stops after n steps, where n is the number of states of the automaton.

[1] Note that a Δ^{k+1} can always be found since the uncertainty in each entry of the matrix $\Pi(S, \Delta)$ is independent of the uncertainty in the other entries.

5. Proofs of Lemma 2 and Theorem 3

The results proven in this section could be derived using a similar reasoning to the one in Fu et al. (2004). However, we present here a new approach that is needed to prove the later results on robust controller design.

5.1. Additional notation

Given a supervisor S and uncertainty value $\Delta \in \mathbf{\Delta}$, let $T_\Delta^S : \mathfrak{R}^n \to \mathfrak{R}^n$ be defined as

$$T_\Delta^S(\mu) \doteq \Pi(S, \Delta)\mu + X.$$

Furthermore, let $T^S : \mathfrak{R}^n \to \mathfrak{R}^n$ be given by

$$T^S(\mu) = \min_{\Delta \in \mathbf{\Delta}} T_\Delta^S(\mu),$$

where the above minimum is taken entry by entry. Note that $T_\Delta^S(\cdot)$ is well-defined since, as mentioned in Section 3.1, the uncertainty in each entry of $\Pi(S, \Delta)$ is independent of the uncertainties in all other entries. Finally, given $x \in \mathfrak{R}^n$, define the max-norm $\|x\| = \max_i |x_i|$. Given $x, y \in \mathfrak{R}^n$, it follows that $x \leqslant y$ if $x_i \leqslant y_i$ for all $i = 1, 2, \ldots, n$. It also follows that $x < y$ if $x \leqslant y$ and $x_i < y_i$ for some i.

Before providing the proofs of Lemma 2 and Theorem 3, a number of relevant properties of the functions $T_\Delta^S(\cdot)$ and $T^S(\cdot)$ are established as supporting lemmas.

Lemma 4. *Let S be a supervisor and* $\Delta \in \mathbf{\Delta}$ *be given, then* T_Δ^S *is a contraction.*

Proof. Recall that there exists a $\theta \in (0, 1)$ such that $\sum_{j=1}^n \pi_{ij}(S, \Delta) \leqslant 1 - \theta$ for all control policies. Now, let $x, y \in \mathfrak{R}^n$ be two vectors, then the ith coordinate of $T_\Delta^S(x) - T_\Delta^S(y)$ satisfies the following inequality:

$$|(T_\Delta^S(x) - T_\Delta^S(y))_i| = |[\Pi(S, \Delta)(x - y)]_i| \leqslant (1 - \theta)\|x - y\|.$$

Hence

$$\|T_\Delta^S(x) - T_\Delta^S(y)\| \leqslant (1 - \theta)\|x - y\|.$$

The proof is completed by noting that $0 < \theta < 1$. □

Lemma 5. *Let S be a supervisor and let* $\Delta, \Delta' \in \mathbf{\Delta}$ *be given. If* $T_\Delta^S(\mu(S, \Delta')) \leqslant \mu(S, \Delta')$ *then*

$$\mu(S, \Delta) \leqslant \mu(S, \Delta').$$

Proof. Note that, since all entries of $\Pi(S, \Delta)$ are nonnegative then

$$x \geqslant y \Rightarrow T_\Delta^S(x) = \Pi(S, \Delta)x + X \geqslant \Pi(S, \Delta)y + X = T_\Delta^S(y).$$

Therefore, it follows that

$$(T_\Delta^S)^2(\mu(S, \Delta')) \doteq T_\Delta^S[T_\Delta^S(\mu(S, \Delta')] \leqslant T_\Delta^S(\mu(S, \Delta')) \leqslant \mu(S, \Delta').$$

Repeating the above reasoning, for any k,

$$(T_\Delta^S)^k(\mu(S,\Delta')) \leqslant \mu(S,\Delta').$$

Now, by Lemma 4 and using the contraction mapping theorem (Naylor & Sell, 1982), it follows that

$$\mu(S,\Delta) = \lim_{k\to\infty}(T_\Delta^S)^k(\mu(S,\Delta')) \leqslant \mu(S,\Delta').$$

Corollary 6. *Let S be a supervisor and let $\Delta, \Delta' \in \mathbf{\Delta}$ be given. Then, $T_\Delta^S(\mu(S,\Delta')) < \mu(S,\Delta')$ implies that*

$$\mu(S,\Delta) < \mu(S,\Delta').$$

Proof. By Lemma 5 it is known that $\mu(S,\Delta) \leqslant \mu(S,\Delta')$. If it is assumed that $\mu(S,\Delta) = \mu(S,\Delta')$, then

$$T_\Delta^S(\mu(S,\Delta')) = T_\Delta^S(\mu(S,\Delta)) = \mu(S,\Delta) = \mu(S,\Delta')$$

and a contradiction is reached. □

Corollary 7. *Let S be a controllable supervisor. Then, the operator T^S is a contraction.*

Proof. Let $x, y \in \Re^n$ be two vectors and let Δ_y be the control policy satisfying $T^S(y) = T_{\Delta_y}^S(y)$. Then,

$$\begin{aligned} T^S(x) - T^S(y) &= \min_{\Delta\in\mathbf{\Delta}} T_\Delta^S(x) - \min_{\Delta\in\mathbf{\Delta}} T_\Delta^S(y) \\ &= \min_{\Delta\in\mathbf{\Delta}} T_\Delta^S(x) - T_{\Delta_y}^S(y) \\ &\leqslant T_{\Delta_y}^S(x) - T_{\Delta_y}^S(y). \end{aligned}$$

Hence,

$$\begin{aligned} (T^S(x) - T^S(y))_i &\leqslant [\Pi(S,\Delta_y)(x-y)]_i \\ &\leqslant |[\Pi(S,\Delta_y)(x-y)]_i| \\ &\leqslant (1-\theta)\|x-y\|. \end{aligned}$$

Exchanging the roles of x and y, it follows that

$$(T^S(y) - T^S(x))_i \leqslant (1-\theta)\|x-y\|.$$

Hence,

$$|(T^S(x) - T^S(y))_i| \leqslant (1-\theta)\|x-y\|$$

and

$$\|(T^S(x) - T^S(y))\| \leqslant (1-\theta)\|x-y\|.$$

The proof is completed by noting that $0 < \theta < 1$. □

Lemma 8. *Let S and S' be two supervisors, if $T^S(\underline{\mu}(S')) \geqslant \underline{\mu}(S')$ then $\underline{\mu}(S) \geqslant \underline{\mu}(S')$. In addition, if $T^S(\underline{\mu}(S')) > \underline{\mu}(S')$ then $\underline{\mu}(S) > \underline{\mu}(S')$.*

Proof. Similar to the proofs of Lemma 5 and Corollary 6. □

Having these preliminary results, one can now proceed with the proofs of Lemma 2 and Theorem 3.

5.2. Proof of Lemma 2

First, note that, by Corollary 7, T^S is a contraction. Hence, there exists a $\underline{\mu}(S)$ such that

$$\underline{\mu}(S) = T^S(\underline{\mu}(S)) = \min_{\Delta\in\mathbf{\Delta}} T_\Delta^S(\underline{\mu}(S)).$$

Since $T_\Delta^S(\underline{\mu}(S))$ depends continuously on Δ and the minimization is done over the compact set $\mathbf{\Delta}$, then there exists a $\Delta^* \in \mathbf{\Delta}$ such that $\underline{\mu}(S) = T_{\Delta^*}^S(\underline{\mu}(S))$ and, therefore, $\underline{\mu}(S) = \mu(S,\Delta^*)$. Furthermore, for any supervisor S and any $\Delta \in \mathbf{\Delta}$, T_Δ^S is a monotone operator, i.e.,

$$x \geqslant y \Rightarrow T_\Delta^S(x) \geqslant T_\Delta^S(y).$$

This is a direct consequence of the fact that the entries of $\Pi(S,\Delta)$ are nonnegative. This implies that T^S is a monotone operator, i.e.,

$$x \geqslant y \Rightarrow T^S(x) \geqslant T^S(y).$$

Now, given any Δ and the associated performance $\mu(S,\Delta)$, the definition of T^S implies that

$$T^S(\mu(S,\Delta)) \leqslant T_\Delta^S(\mu(S,\Delta)) = \mu(S,\Delta).$$

Hence

$$\underline{\mu}(S) = \lim_{k\to\infty}(T^S)^k(\mu(S,\Delta)) \leqslant \mu(S,\Delta).$$

Therefore, $\underline{\mu}(S) = \mu(S,\Delta^*)$ and the proof is complete.

5.3. Proof of Theorem 3

Since π_{ij} multiplies μ_j in the computation of $T(\mu)$, it follows that Δ^{k+1} in Step 1 of Algorithm 1 must satisfy the condition

$$\begin{aligned} T_{\Delta^{k+1}}^S(\mu(S,\Delta^k)) &= \min_{\Delta\in\mathbf{\Delta}} T_\Delta^S(\mu(S,\Delta^k)) \\ &\leqslant T_{\Delta^k}^S(\mu(S,\Delta^k)) = \mu(S,\Delta^k). \end{aligned}$$

Hence, by Lemma 5,

$$\mu(S,\Delta^{k+1}) \leqslant \mu(S,\Delta^k).$$

Moreover, the algorithm stops only if $\mu(S,\Delta^{k+1}) = \mu(S,\Delta^k)$. Therefore, the stopping rule is equivalent to having

$$\begin{aligned} \mu(S,\Delta^{k+1}) &= T_{\Delta^{k+1}}^S(\mu(S,\Delta^{k+1})) \\ &= T_{\Delta^{k+1}}^S(\mu(S,\Delta^k)) = \min_{\Delta\in\mathbf{\Delta}} T_\Delta^S(\mu(S,\Delta^k)). \end{aligned}$$

Hence,

$$\mu(S,\Delta^k) = \min_{\Delta\in\mathbf{\Delta}} T_\Delta^S(\mu(S,\Delta^k)) = \underline{\mu}(S).$$

To prove convergence in n steps, note that $\mu(S,\Delta^{k+1}) \neq \mu(S,\Delta^k)$ if and only if there exists a j such that the sign of $\mu_j(S, \Delta^{k+1})$ is different from the sign of $\mu_j(S, \Delta^k)$. Moreover, $\mu(S, \Delta^{k+1}) \leqslant \mu(S, \Delta^k)$. The above monotonicity property implies that each entry of $\mu_j(S, \Delta^k)$ changes sign at most one time and, since $\mu(S, \Delta^k)$ is a vector of dimension n, the algorithm will stop after n steps.

6. Optimal robust supervisor

We now provide an algorithm that generates a sequence of supervisors which converges to the most permissive robust optimal supervisor in a finite number of steps.

Algorithm 2. *Optimal robust supervisor design.*

Step 0. *Let S^0 be a controllable supervisor and let $k = 0$.*

Step 1. *Given S^k, determine $\underline{\mu}(S^k)$ using Algorithm* 1.

Step 2. *Determine S^{k+1} by disabling the controllable events leading to states with $\underline{\mu}_j(S^k) < 0$ and enabling the events leading to states with $\underline{\mu}_j(S^k) \geqslant 0$.*

Step 3. *If $S^{k+1} = S^k$ then set $S^* = S^k$, $\underline{\mu}(S^*) = \underline{\mu}(S^k)$ and stop. Else return to Step* 1 *with $k \leftarrow (k+1)$.*

Theorem 9. *The control policy S^* obtained by Algorithm* 2 *is an optimal control policy over the uncertainty range and the robust performance of the closed-loop system is given by $\underline{\mu}(S^*)$. Moreover, S^* is the maximally permissive controller among all robust optimal controllers. Furthermore, the algorithm terminates in at most n steps, where n is the number of states of the* (*deterministic*) *automaton of the plant model*.

Proof. We start by introducing an additional function: Let $T : \Re^n \to \Re^n$ be defined as

$$T(\mu) \doteq \max_S T^S(\mu).$$

Some relevant properties of $T(\cdot)$ are established in the following two lemmas. The proofs are omitted since they are similar to the proofs of the lemmas in Section 5.

Lemma 10. *The transformation T is a contraction*.

Lemma 11. *There exists a S^* such that*

$$\underline{\mu}^* = \underline{\mu}(S^*) = T(\underline{\mu}(S^*)).$$

Furthermore, for all supervisor S, $\underline{\mu}^* \geqslant \underline{\mu}(S)$.

The proof of Theorem 9 is started by noting that S^{k+1} in Step 1 of Algorithm 2 is chosen in such a way as to maximize π_{ij} when $\mu_j(S^k) \geqslant 0$ and minimize π_{ij} when $\mu_j(S^k) < 0$. Then, it follows that

$$T^{S^{k+1}}(\underline{\mu}(S^k)) = \max_S T^S(\underline{\mu}(S^k)) \geqslant T^{S^k}(\underline{\mu}(S^k)) = \underline{\mu}(S^k).$$

Hence, by Lemma 8, $\underline{\mu}(S^{k+1}) \geqslant \underline{\mu}(S^k)$. Moreover, the algorithm stops only if $\underline{\mu}(S^{k+1}) = \underline{\mu}(S^k)$ yielding

$$\underline{\mu}(S^{k+1}) = T^{S^{k+1}}(\underline{\mu}(S^{k+1})) \\ = T^{S^{k+1}}(\underline{\mu}(S^k)) = \max_S T^S(\underline{\mu}(S^k)).$$

Therefore,

$$\underline{\mu}(S^k) = \max_S T^S(\underline{\mu}(S^k)) = \underline{\mu}(S^*),$$

where S^* is the optimal controller. To prove convergence in n steps, note that $\underline{\mu}(S^{k+1}) \neq \underline{\mu}(S^k)$ if and only if there exists a j such that the sign of $\underline{\mu}_j(S^{k+1})$ is different from the sign of $\underline{\mu}_j(S^k)$. Because of the fact that $\underline{\mu}(S^{k+1}) \geqslant \underline{\mu}(S^k)$, this monotonicity property implies that each entry of $\underline{\mu}_j(S^k)$ changes sign at most one time and, since $\underline{\mu}_j(S^k)$ is a vector of dimension n, the algorithm will stop after at most n steps.

To prove that the controller S^* obtained is the most permissive one among all the optimal controllers, let $\underline{\mu}^*$ be the optimal performance and let $\mathscr{S}^*$ be the set of all optimal supervisors. Then,

$$\mathscr{S}^* = \{S : T(\underline{\mu}^*) = T^S(\underline{\mu}^*) = \underline{\mu}^*\}.$$

Given this, all supervisors in $\mathscr{S}^*$ only differ in the enabling or disabling or events leading to states q_j with optimal performance $\underline{\mu}_j^* = 0$. Since the supervisor $S^* \in \mathscr{S}^*$ obtained by the algorithm enables all events leading to states q_i with $\underline{\mu}_i^* \geqslant 0$, then it is maximally permissive among all optimal supervisors in $\mathscr{S}^*$.

7. Summary and conclusions

This paper presents the theory of a state-based robust optimal control policy of regular languages for finite state automata that may have already been subjected to constraints such as control specifications (Ramadge & Wonham, 1987). The synthesis procedure is quantitative and relies on a recently developed signed real measure of formal languages (Wang & Ray, 2004). The objective is to maximize the worst-case performance vector over the event cost uncertainty range without any further constraints. The robust optimal control policy maximizes performance by selectively disabling controllable events that may terminate on “bad” marked states and, simultaneously, ensuring that the remaining controllable events are kept enabled. The worst-case performance is guaranteed within the specified event cost uncertainty bounds. The control policy induced by the updated state transition cost matrix yields maximal performance and is unique in the sense that the controlled language is the most permissive (i.e., least restrictive) among all controller(s) having robust optimal performance. The computational complexity of the robust optimal control synthesis procedure is polynomial in the number of states of the automaton.

Future areas of research in robust optimal control include: (i) incorporation of the cost of disabling controllable events (Cury & Krogh, 1999) and (ii) robustness of the control policy relative to other uncertainties in the plant model including loss of controllability and observability.

References

Brave, Y., & Heyman, M. (1993). On optimal attraction of discrete-event processes. *Information Science*, *67*(3), 245–276.

Cury, J. E. R., & Krogh, B. H. (1999). Robustness of supervisors for discrete-event systems. *IEEE Transactions on Automatic Control*, *44*(2), 376–379.

Fu, J., Ray, A., & Lagoa, C. M. (2003). Optimal control of regular languages with event disabling cost. *Proceedings of American control conference* (pp. 1691–1695). Denver, Colorado.

Fu, J., Ray, A., & Lagoa, C. M. (2004). Unconstrained optimal control of regular languages. *Automatica*, *40*(4), 639–648.

Kumar, R., & Garg, V. (1995). Optimal supervisory control of discrete event dynamical systems. *SIAM Journal of Control and Optimization*, *33*(2), 419–439.

Lin, F. (1993). Robust and adaptive supervisory control of discrete event systems. *IEEE Transactions on Automatic Control*, *38*(12), 1848–1852.

Naylor, A. W., & Sell, G. R. (1982). *Linear operator theory in engineering and science*. Berlin: Springer.

Park, S.-J., & Lim, J.-T. (2002). Robust and nonblocking supervisory control of nondeterministic discrete event systems using trajectory models. *IEEE Transactions on Automatic Control*, *47*(4), 655–658.

Passino, K., & Antsaklis, P. (1989). On the optimal control of discrete event systems. *Proceedings of 28th IEEE decision and control conference* (pp. 2713–2718). Tampa, Florida.

Ramadge, P. J., & Wonham, W. M. (1987). Supervisory control of a class of discrete event processes. *SIAM Journal of Control and Optimization*, *25*(1), 206–230.

Ray, A., & Phoha, S. (2003). Signed real measure of regular languages for discrete-event automata. *International Journal of Control*, *76*(18), 1800–1808.

Sengupta, R., & Lafortune, S. (1998). An optimal control theory for discrete event systems. *SIAM Journal of Control and Optimization*, *36*(2), 488–541.

Takai, S. (1999). Design of robust supervisors for prefix-closed language specifications. *Proceedings of the 38th IEEE conference on decision and control* (pp. 1725–1730).

Wang, X., & Ray, A. (2004). A language measure for performance evaluation of discrete-event supervisory control systems. *Applied Mathematical Modelling*, *28*(9), 817–833.

Constantino Lagoa got his B.S. and M.Sc. degrees from the Instituto Superior Técnico, Technical University of Lisbon, Portugal in 1991 and 1994, respectively and his Ph.D. degree from the University of Wisconsin at Madison in 1998. He joined the Electrical Engineering Department of the Pennsylvania State University in August 1998, where he currently holds the position of Associate Professor. He has a wide range of research interests including robust control, controller design under risk specifications, control of computer networks and discrete event dynamical systems. In 2000 he received the NSF CAREER award for his proposal on system design under risk constraints.

Jinbo Fu received the Bachelor and Master degrees, both in Precision Instruments and Mechanology, from Tsinghua University, Beijing, China, in 1996 and 1999, respectively. He also received Master degrees in Electrical Engineering and Mechanical Engineering from Pennsylvania State University, University Park in 2002 and 2003, respectively, and the Ph.D. degree in Mechanical Engineering from Pennsylvania State University, University Park in 2003.

Dr. Fu has been working on the theories and applications of discrete event supervisory (DES) control since 2000 in several projects. His research interests include DES control theories and application, intelligent systems, engine propulsion system command and control, damage mitigating control, fault detection and information fusion.

Asok Ray earned the Ph.D. degree in Mechanical Engineering from Northeastern University, Boston, MA, and also graduate degrees in each discipline of Electrical Engineering, Mathematics, and Computer Science. Dr. Ray joined the Pennsylvania State University in July 1985, and is currently a Distinguished Professor of Mechanical Engineering. Prior to joining Penn State, Dr. Ray held research and academic positions at Massachusetts Institute of Technology and Carnegie-Mellon University as well as research and management positions at GTE Strategic Systems Division, Charles Stark Draper Laboratory, and MITRE Corporation. Dr. Ray has been a Senior Research Fellow at NASA Glenn Research Center under a National Academy of Sciences award. Dr. Ray's research experience and interests include: Control and optimization of continuously varying and discrete-event dynamical systems, intelligent instrumentation for real-time distributed systems and modeling and analysis of complex dynamical systems from thermodynamic perspectives in both deterministic and stochastic settings, as applied to Aeronautics and Astronautics, Undersea Vehicles and Surface Ships, Power and Processing plants, and Robotics. Dr. Ray has authored or co-authored three hundred seventy research publications including over one hundred and fifty scholarly articles in refereed journals such as transactions of ASME, IEEE and AIAA, and research monographs. Dr. Ray is a Fellow of IEEE and ASME, and an Associate Fellow of AIAA.