Almost Sure Stability of Anytime Controllers via Stochastic Scheduling

Luca Greco, Daniele Fontanelli and Antonio Bicchi

***Abstract*— In this paper we consider scheduling the execution of a number of different software tasks implementing a hierarchy of real-time controllers for a given plant. Controllers are considered to be given and ordered according to a measure of the performance they provide. Software tasks implementing higher-performance controllers are assumed to require a larger worst-case execution time. Control tasks are to be scheduled within given availability limits on execution time, which change stochastically in time. The ensuing switching system is prone to instability, unless a very conservative policy is enforced of always using the simplest, least performant controller. The presented method allows to condition the stochastic scheduling of control tasks so as to obtain a better exploitation of computing capabilities, while guaranteeing almost sure stability of the resulting switching system.**

I. Introduction

In the implementation of controllers on an embedded programmable processor, it is often the case that control tasks have to share computational resources with several other tasks. A multitasking Real-Time Operating System (RTOS) is often used to schedule the execution of tasks within stringent temporal constraints, typically using priorities. Real-time preemptive algorithms can suspend the execution of a task in the presence of tasks having a higher priority. Representative examples of such scheduling algorithms are Rate Monotonic (RM) and Earliest Deadline First (EDF) [1], [2], [3]. As a consequence, time allocated to execute a control task can change unpredictably in time.

Assuming that the RTOS guarantees a minimum time τ_{min} for the execution of the control task at hand, a conventional approach to cope with this problem is to design controllers simple enough to be executable within τ_{min}. A less conservative approach consists in designing "scalable" control algorithms, which can be interrupted at any time, while still providing a valid result whose performance increases with the execution time actually allowed. The idea, stemming from the theory of imprecise computation [4], is aptly described by the term "anytime control", used already in [5]. With respect to most anytime algorithms, however, the fact that anytime controllers interact in feedback with dynamic systems introduces severe difficulties in their synthesis and in the analysis of the resulting closed-loop performance. Indeed, the stochastic switching system ensuing from executing different controllers at different times is even prone to instability.

In this paper, we focus on the analysis of stability of linear anytime controllers, whose scheduling is decided stochastically. Being the scheduler modelled as a Markov chain, an anytime-controlled linear system turns out to be a Markov Jump Linear System (MJLS). We present a method to condition the stochastic properties of the scheduler so as to obtain a better exploitation of computational power, while guaranteeing stability of the resulting switching system in a suitable probabilistic sense, i.e. Almost Sure (AS) stability [6], [7].

Work supported by E.C. under contracts IST-2004-004536 "RUNES"and IST-2004-511368 "HYCON"

Authors are with the Interdept. Research Center "E. Piaggio", University of Pisa, Italy

II. Problem Formulation and Background

Let Σ be a given linear, discrete time, invariant plant described as

$$\begin{array}{rcl} w_{t+1} & = & Aw_t + Bu_t \\ y_t & = & Cw_t. \end{array}$$

Let also Γ_i, $i \in I \triangleq \{1, 2, \ldots, n\}$, be a family of feedback controllers for Σ, described by

$$\begin{array}{rcl} z^i_{t+1} & = & F_i z^i_t + G_i y_t \\ u^i_t & = & H_i z^i_t + L_i y_t. \end{array}$$

Assume that the feedback connection of Σ with Γ_i, obtained by setting $u(\cdot) = u_i(\cdot)$, is asymptotically stable for all i. Let the closed-loop system thus obtained be denoted Σ_i and described by

$$x_{t+1} = \widehat{A}_i x_t, \tag{1}$$

where

$$\widehat{A}_i = \begin{bmatrix} A + BL_iC & BH_i \\ G_iC & F_i \end{bmatrix}.$$

The set of controllers is assumed to be given and to proceed from a synthesis technique (such as e.g. cascade design) providing a hierarchy of performance and complexity. In other terms, we assume that application of controller i provides better closed-loop performance than controller j if $i > j$. On the other hand, the worst-case execution time (WCET) of the software code implementing more performant controllers is larger ($WCET_i > WCET_j$, $i > j$).

A. Scheduling Description Using Markov Chains

Neglecting for simplicity delays and jitter, we assume that measurements are acquired and control inputs are released at every sampling instant tT_g, $t \in \mathbb{N}$, where T_g is a fixed sampling time.

Let $\gamma_t \in [\tau_{min}, \tau_{max}]$, $\tau_{max} < T_g$, denote the time allotted to the control task during the t-th sampling interval. By hypothesis, $WCET_1 \leq \tau_{min}$ and $WCET_n \leq \tau_{max}$.

Define an event set $L_\tau \triangleq \{\tau_1, \ldots, \tau_n\}$, and a map

$$\begin{array}{rcccl} \mathcal{T} & : & [\tau_{min}, \tau_{max}] & \to & L_\tau \\ & & \gamma_t & \mapsto & \tau(t) \end{array}$$

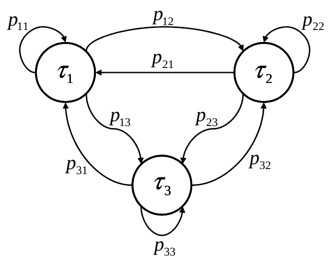

Fig. 1. Markov chain model of a scheduler.

where

$$\tau(t) = \begin{cases} \tau_1, & \text{if} \quad \gamma_t \in [\tau_{min}, WCET_2) \\ \tau_2, & \text{if} \quad \gamma_t \in [WCET_2, WCET_3) \\ \vdots & \text{if} \quad \vdots \\ \tau_n, & \text{if} \quad \gamma_t \in [WCET_n, \tau_{max}] \end{cases}$$

A simple stochastic description of the random sequence $\{\tau(t)\}_{t\in\mathbb{N}}$, inherited from the probability distribution of γ_t, can be given in terms of an independently and identically distributed (i.i.d.) process, with probability distribution $\Pr\{\tau(t) = \tau_i\} = p_i$, with $0 \le p_i \le 1$ and $\sum_{i\in I} p_i = 1$ associated to the event that "at time t the time slot γ_t is such that all controllers $\Gamma_j, j \le i$, but no controller $\Gamma_k, k > i$, can be executed."

A slightly more complex, but more general, model allowing for non-stationary probability distributions is provided by a finite state discrete-time homogeneous irreducible aperiodic Markov chain. In this case $\Pr\{\tau(t) = \tau_i\} = \pi_i(t)$ is time-dependent, while transition probabilities $p_{ij} \triangleq \Pr\{\tau(t+1) = \tau_j \mid \tau(t) = \tau_i\}$ are time-independent. A Markov-chain description of the scheduler for $n = 3$ is depicted in figure 1.

B. Almost Sure Stability

In this section a brief review of results on Almost Sure stability (AS–stability) for discrete–time systems is reported following [7].

Consider the discrete-time Markov Jump Linear System (MJLS)

$$x_{t+1} = A_{\sigma_t} x_t \tag{2}$$

where $\sigma = \{\sigma_t\}_{t\in\mathbb{N}}$ is a finite state, discrete–time and time homogeneous Markov chain taking values in $I \triangleq \{1, 2, \ldots, n\}$. The statistics of the Markov process are identified by the transition probability matrix $P = (p_{ij})_{n\times n}$, where $p_{ij} \triangleq \Pr\{\sigma_{t+1} = j \mid \sigma_t = i\}$, and by an initial probability measure π_0 defined on I, with $\pi_{0i} \triangleq \Pr\{\sigma_0 = i\}$. The evolution of the probability distribution $\pi(t)$ of the process at time t is given by

$$\pi(t+1) = P^T \pi(t). \tag{3}$$

If the Markov chain is irreducible and aperiodic (see for instance [8]), then there exists a unique *invariant probability distribution* $\overline{\pi}$ such that $\overline{\pi}_i > 0 \; \forall\, i \in I$ and $\lim_{t\to\infty} \pi(t) = \overline{\pi}$ for any π_0. Such a distribution provides the steady-state probability distribution for the process σ.

Definition 1: The MJLS (2) is said (exponentially) almost surely stable (AS-stable) if there exists $\mu > 0$ such that, for any $x_0 \in \mathbb{R}^N$ and any initial distribution π_0, the following condition holds

$$\Pr\left\{\limsup_{t\to\infty} \frac{1}{t} \ln \|x_t\| \le -\mu \right\} = 1.$$

Let $\|\cdot\|$ be a matrix norm induced by some vector norm. The following sufficient condition for AS-stability was proved in [6]:

Theorem 1 (1–step average contractivity): [6] If

$$\xi_1 = \prod_{i\in I} \|A_i\|^{\overline{\pi}_i} < 1 \tag{4}$$

then the MJLS (2) is AS-stable.

Inequality (4) can be interpreted as an average contractivity of the state norm over a one-step horizon. In this vein, a less restrictive condition involving the average contractivity over a multi–step horizon has been presented in [7]. Namely, a new MJLS, called "lifting of period m", is associated to the MJLS (2). Such a system represents the sampling of the original one at time instants hm, $h \in \mathbb{N}$, and its stability properties mirror those of the original system. More precisely, the lifted version of period m of system (2) is defined by

$$\widetilde{x}_{h+1} = \widetilde{A}_{\widetilde{\sigma}_h} \widetilde{x}_h$$

with

$$\begin{aligned} \widetilde{x}_h &= x_{mh} \\ \widetilde{\sigma}_h &= \begin{bmatrix} \sigma_{mh} & \cdots & \sigma_{mh+m-1} \end{bmatrix}^T \\ \widetilde{A}_{\widetilde{\sigma}_h} &= A_{\sigma_{mh+m-1}} A_{\sigma_{mh+m-2}} \cdots A_{\sigma_{mh}}. \end{aligned}$$

Moreover, $\widetilde{\sigma} = \{\widetilde{\sigma}_h\}_{h\in\mathbb{N}}$ is a stationary Markov process taking values in $\widetilde{I} = I^m$ and characterized as follows: for $\widetilde{\imath} = (i_1, i_2, \ldots, i_m) \in \widetilde{I}$ and $\widetilde{\jmath} = (j_1, j_2, \ldots, j_m) \in \widetilde{I}$, the transition probability $\widetilde{p}_{\widetilde{\imath}\widetilde{\jmath}} \triangleq \Pr\{\widetilde{\sigma}_{h+1} = \widetilde{\jmath} \mid \widetilde{\sigma}_h = \widetilde{\imath}\}$ is given by $\widetilde{p}_{\widetilde{\imath}\widetilde{\jmath}} = p_{i_m j_1} \prod_{k=1}^{m-1} p_{j_k j_{k+1}}$. This process has the following invariant probability distribution $\widetilde{\pi}_{\widetilde{\imath}} = \prod_{k=1}^{m-1} p_{i_k i_{k+1}} \overline{\pi}_{i_1}$.

The 1–step average contractivity condition applied to the lifted system yields

Theorem 2 (m–step average contractivity): [7] If

$$\xi_m = \prod_{\widetilde{\imath}\in\widetilde{I}} \|\widetilde{A}_{\widetilde{\imath}}\|^{\widetilde{\pi}_{\widetilde{\imath}}} < 1 \tag{5}$$

then the MJLS (2) is AS-stable.

The importance of condition (5) is related to the fact that for increasing values of m it provides a sequence of sufficient conditions and, most importantly, to the following result

Theorem 3: [7] System (2) is exponentially AS-stable if and only if $\exists m \in \mathbb{N}$ such that condition (5) holds.

III. Stochastic Scheduler and AS-stability

Once a set of controllers Γ_i and the associated closed loop dynamics $\widehat{A}_i$ are given, and a Markov chain description of the scheduler is provided, in order to actually execute an anytime controller it is necessary to determine a switching policy for choosing among possible controllers. In our context, a *switching policy* is defined as a map $s : \mathbb{N} \to I$, $t \mapsto s(t)$, and determines an upper bound to the index i of the controller to be executed at time t, $i \leq s(t)$. In other terms, at time tT_g, the system starts computing the controller algorithm until it can provide the output of $\Gamma_{s(t)}$, unless a preemption event occurs forcing it to provide only $\Gamma_{\tau(t)}$, the highest controller computed before preemption. Application of a switching policy s to a set of feedback systems Σ_i, $i \in I$ under a scheduler τ generates a switching linear system (Σ_i, τ, s) which, under suitable hypotheses, is also a MJLS.

As an example, the most conservative policy is to set $s(t) \equiv 1$, i.e. forcing always the execution of the simplest controller Γ_1, regardless of the probable availability of more computational time. By assumption, this (non-switching) policy guarantees stability of the resulting closed loop system.

On the opposite, a "greedy" strategy would set $s(t) \equiv n$, which leads to providing $\Gamma_{\tau(t)}$ for all t. Although this policy attempts at maximizing the utilization of the most performing controller, it is well known that switching arbitrarily among asymptotically stable systems Σ_i may easily result in an unstable behavior [9].

A sufficient condition for the greedy switching policy to provide an AS-stable system is provided by Theorem 1, where $A_i = \widehat{A}_i$ and $\overline{\pi} = \overline{\pi}_\tau$. If this condition is not verified, it is important to investigate whether there exist intermediate switching policies which can ensure AS-stability. To do so, consider first the following

Problem 1: Given a set of matrices A_i, $i \in I$, find the set

$$\overline{\Pi} = \{\overline{\pi}_\tau \mid \prod_{i \in I} \|A_i\|^{\overline{\pi}_{\tau_i}} < 1\}$$

of invariant probability distributions of homogeneous, irreducible, aperiodic Markov chains, such that the associated MJLS is 1-step contractive.
Assume that Problem 1, applied to the set of closed loop matrices (i.e., with $A_i = \widehat{A}_i$), admits a solution $\overline{\Pi} \neq \emptyset$. Then, the following synthesis problem is of relevance:

Problem 2: Given a set of closed loop matrices $\widehat{A}_i$, $i \in I$ and an invariant probability scheduler distribution $\overline{\pi}_\tau$, find a (stabilizing) switching law s such that the resulting switching system (Σ_i, τ, s) is a MJLS with invariant probability distribution $\overline{\pi}_d \in \overline{\Pi}$.

It is worth noting that the existence of a solution to Problem $P1$ is not obvious in general. Indeed, for a given set of Schur closed-loop matrices, it is not always possible to find an invariant probability distribution $\overline{\pi}_\tau$ ensuring one-step contractivity. Indeed, for the 1-step contractivity to hold in a given norm, it is necessary that $\exists i \in I$ such that $\left\|\widehat{A}_i\right\| < 1$, which condition may well not be verified. Furthermore, solution of the synthesis Problem 2 may result very difficult if the set $\overline{\Pi}$ is too small.

To ease the synthesis, Problem 1 could be extended to lifted systems by exploiting the m–step contractivity condition (5). A characterization of viable schedulers in this case should be given in terms of a set of Markov chain transition matrices P. Unfortunately, the synthesis of a switching policy solving the analogous to Problem 2 in this case turns out to be very complex.

Our approach consists in exploiting the possibility of supervising the controller choice so that some control patterns, i.e. substrings of symbols in I, are preferentially used with respect to others. For a substring of length m $\widetilde{\imath} \in \widetilde{I} = I^m$, $\widetilde{\imath} = (i_1, i_2, \ldots, i_m)$, let $\widehat{A}_{\widetilde{\imath}} = \widehat{A}_{i_1}\widehat{A}_{i_2}\cdots\widehat{A}_{i_m}$. The following result will prove useful in the sequel.

Proposition 1: If $\exists i \in I$ such that $\widehat{A}_i$ is Schur, then $\exists m \in \mathbb{N}$ such that Problem 1 applied to $\widehat{A}_{\widetilde{\imath}}$ admits a solution, i.e.

$$\overline{\Pi} = \{\overline{\pi}_{\widetilde{\tau}} \mid \prod_{\widetilde{\imath} \in \widetilde{I}} \|\widehat{A}_{\widetilde{\imath}}\|^{\overline{\pi}_{\widetilde{\tau}_{\widetilde{\imath}}}} < 1\} \neq \emptyset$$

Proof. If $\widehat{A}_i$ is Schur, then, for any given matrix norm, $\exists m \in \mathbb{N}$ such that $\left\|\widehat{A}_i^m\right\| < 1$. Let $\imath_i = \{i, \ldots, i\} \in \widetilde{I}$ and $\widehat{A}_{\widetilde{\imath}} = \widehat{A}_i^m$. Hence $\exists\, \overline{\epsilon} > 0$ such that, for all probability distributions with $\overline{\pi}_{\tau_{\imath_i}} = 1 - \epsilon$, $\sum_{\jmath \in \widetilde{I}}^{\jmath \neq \imath_i} \overline{\pi}_{\tau_\jmath} = \epsilon$, $\epsilon \leq \overline{\epsilon}$, the solution $\overline{\Pi} \neq \emptyset$. ∎

Given a solution set to Problem 1 $\overline{\Pi}$, further constraints have to be satisfied in order for Problem 2 to be solvable. Indeed, let $\overline{\pi}_d \in \overline{\Pi}$. In order for a switching policy $s(t)$ to exist which can alter a given scheduler probability distribution $\overline{\pi}_\tau$ into $\overline{\pi}_d$, the following conditions must hold:

$$\begin{aligned}
\overline{\pi}_{d_n} &\leq \overline{\pi}_{\tau_n} && (C.1)\\
\overline{\pi}_{d_{n-1}} &\leq \overline{\pi}_{\tau_{n-1}} + (\overline{\pi}_{\tau_n} - \overline{\pi}_{d_n}) && (C.2)\\
&\vdots \\
\overline{\pi}_{d_1} &\leq \overline{\pi}_{\tau_1} + (\overline{\pi}_{\tau_2} - \overline{\pi}_{d_2}) + \cdots + (\overline{\pi}_{\tau_n} - \overline{\pi}_{d_n}) && (C.n)
\end{aligned}$$

where $\overline{\pi}_\tau = \begin{bmatrix}\overline{\pi}_{\tau_1} & \overline{\pi}_{\tau_2} & \cdots & \overline{\pi}_{\tau_n}\end{bmatrix}^T$,
and $\overline{\pi}_d = \begin{bmatrix}\overline{\pi}_{d_1} & \overline{\pi}_{d_2} & \cdots & \overline{\pi}_{d_n}\end{bmatrix}^T$.

Inequalities $(C.1)$–$(C.n)$ take into account the fact that no switching law can alter the scheduler so as to give more computational time to control tasks than it is made available by the scheduler. Furthermore, constraints $(C.2)$–$(C.n)$ model the fact that the probability of the i-th controller can be increased only at the expenses of a reduction of the probabilities $\overline{\pi}_{d_j}$, $j > i$ of more complex controllers.

To take into account the above constraints explicitly, Problem 1 can be reformulated as follows:

Problem 3: Given a set of matrices $\widehat{A}_i$, $i \in I$, find the set $\overline{\Pi}_d = \{\overline{\pi}_d\}$ with

$$\begin{aligned}
&P3.1) && \prod_{i=1}^{n} \left\|\widehat{A}_i\right\|^{\overline{\pi}_{d_i}} < 1\\
&P3.2) && 0 < \overline{\pi}_{d_i} < 1\\
&P3.3) && \sum_{i=1}^{n} \overline{\pi}_{d_i} = 1\\
&P3.4) && \overline{\pi}_{d_i} \leq \overline{\pi}_{\tau_i} + \sum_{j=i+1}^{n} \overline{\pi}_{\tau_j} - \sum_{j=i+1}^{n} \overline{\pi}_{d_j}.
\end{aligned}$$

IV. Conditioning Chain

Assume, for a given anytime-controlled system and a scheduler described by a homogeneous irreducible aperiodic Markov chain τ with transition probability P_τ and steady-state probability distribution $\overline{\pi}_\tau$, that $\overline{\Pi}_d \neq \emptyset$. In this section,

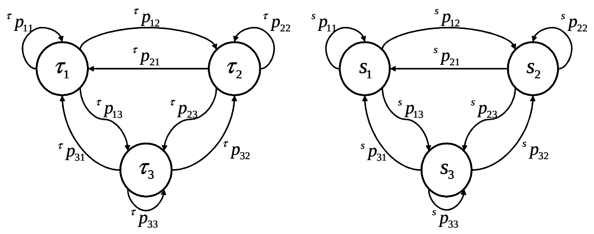

Fig. 2. State transition diagram for the scheduler (with states τ_i) and for the switching law (with states σ_i) with $n = 3$.

we tackle the synthesis Problem 2 by developing a *stochastic* switching law.

Consider a homogeneous irreducible aperiodic Markov chain σ with the same number n of states as the chain τ. The states are labelled as σ_i, with the meaning that if the associated process form $\sigma(t)$ is equal to σ_i, then $s(t) = i$, i.e. in the next sampling interval tT_g at most the i–th controller is computed (if no preemption occurs). We will refer to σ as the conditioning Markov chain.

It should be noted that the choice of determining the switching law by a Markov chain σ not dependent on the scheduler chain τ (see figure 2) renounces to full generality. However, it has the advantage of not requiring on-line computations, and simplifies considerably the analysis.

In the following paragraphs, we study how the stochastic properties of the scheduler and conditioning chain interact to produce a resulting switching system. The purpose is to synthetize a conditioning chain which can produce a MJLS with a steady-state probability distribution $\overline{\pi}_d \in \overline{\Pi}$ ensuring AS-stability.

A. Merging Markov Chains: Mixing

Consider two independent finite-state homogeneous irreducible aperiodic Markov chains α and β such that $\alpha(t) : \Omega \to L_\alpha \triangleq \{\alpha_i \mid i \in I\}$ and $\beta(t) : \Omega \to L_\beta \triangleq \{\beta_j \mid j \in I\}$, $t \in \mathbb{N}$. Let the statistics of α and β be given by the transition probability matrices $P_\alpha = \left({}^{\alpha}p_{ij}\right)_{n\times n}$ and $P_\beta = \left({}^{\beta}p_{ij}\right)_{n\times n}$ with $i, j \in I$, and by the initial probability distributions $\pi_\alpha(0)$ and $\pi_\beta(0)$. Denote with $\overline{\pi}_\alpha$ and $\overline{\pi}_\beta$ the (unique) steady-state probability distributions of α and β respectively. Define the stochastic process $\alpha\beta : \Omega \to L_{\alpha\beta} \triangleq L_\alpha \times L_\beta$ such that $\alpha\beta(t) = (\alpha(t), \beta(t))$.

Theorem 4: The following assertions are true:

i) $\alpha\beta$ is a finite-state homogeneous irreducible aperiodic Markov chain whose statistics are given by the transition probability matrix $P_{\alpha\beta} = \left({}^{\alpha\beta}p_{ij}\right)_{n\times n} = P_\alpha \otimes P_\beta$ and by the initial probability distribution $\pi_{\alpha\beta}(0) = \pi_\alpha(0) \otimes \pi_\beta(0)$;

ii) the evolution of the chain $\alpha\beta$ is given by

$$\pi_{\alpha\beta}(t) = \pi_\alpha(t)\otimes\pi_\beta(t) = \left[P^t_{\alpha\beta}\right]^T (\pi_\alpha(0) \otimes \pi_\beta(0)) \tag{6}$$

with $t \in \mathbb{N}$. $\pi_{\alpha\beta}(t)$ converges to the unique invariant probability distribution

$$\overline{\pi}_{\alpha\beta} = \overline{\pi}_\alpha \otimes \overline{\pi}_\beta \tag{7}$$

for any initial distribution $\pi_{\alpha\beta}(0)$.

The proof is reported in the Appendix.

B. Merging Markov Chains: Aggregating

The Markov chain $\tau\sigma$ obtained mixing the chains τ and σ having both n states, has n^2 states ($L_{\tau\sigma}$ is the Cartesian product of L_τ and L_σ). Our goal is to produce a form process assuming values in the set of the controller indices, that is we must build a process with a desired stationary probability $\overline{\pi}_d \in \overline{\Pi}$, whose cardinality is n. Hence, after mixing the two chains we make use of an *aggregation function* to reduce the number of states. This function is based on the constraints imposed by the scheduler and groups together all the states of $\tau\sigma$ actually producing the execution of the same controller. The k-th controller is hence executed in the next interval if:

1) $\tau(t) \geq \tau_k$ and $\sigma(t) = \sigma_k$ (i.e. the suggested controller is $s(t) = k$);
2) $\sigma(t) \geq \sigma_k$ (i.e. the suggested controller $s(t) \geq k$), and $\tau(t) = \tau_k$ (i.e. the scheduler makes preemption at τ_k).

Therefore, assuming without loss of generality that the indices are chosen such that $L_\tau \equiv L_\sigma$, we define an aggregated process $(\tau\sigma)^\star$, taking values in the set $L_{(\tau\sigma)^\star} \equiv L_\tau \equiv L_\sigma$ with cardinality n, as

$$(\tau\sigma)^\star(t) = \min\{\tau(t), \sigma(t)\}.$$

The characterization of the aggregated process $(\tau\sigma)^\star$ is rather easy, as shown below.

Proposition 2: The evolution of the probability distribution $\pi^\star(t) = [\pi^\star_1(t), \ldots, \pi^\star_n(t)]$ of the process $(\tau\sigma)^\star$ aggregated by means of the *minimum* function is given by

$$\pi^\star(t) = H(\pi_\tau(t) \otimes \pi_\sigma(t)) \tag{8}$$

with $H \in \{0,1\}^{n\times n^2}$ such that $H = [H_1, H_2, \ldots, H_n]$ and

$$H_i = \left[\begin{array}{c|c} I_{i,i} & \begin{array}{c} 0_{i-1,n-i} \\ 1_{1,n-i} \end{array} \\ \hline 0_{n-i,i} & 0_{n-i,n-i} \end{array}\right].$$

Remark 1: The previous proposition asserts that the evolution of the aggregated process $(\tau\sigma)^\star$ is related to the evolutions of the chains τ and σ by a linear time-invariant mapping. Therefore, the process $(\tau\sigma)^\star$ admits an invariant distribution to which (at least) each initial distribution of type $H(\pi_\tau(0) \otimes \pi_\sigma(0))$ converges, if the chains τ and σ have their steady-state distributions. In particular, from (8), we have

$$\overline{\pi}^\star = H(\overline{\pi}_\tau \otimes \overline{\pi}_\sigma). \tag{9}$$

V. Markov Based Switching Law

A. One-Step Solution

In this section we address the solution of Problem 2 by investigating the existence of a Markov chain σ conditioning the scheduler chain τ so as to produce a MJLS with an invariant probability distribution $\overline{\pi}_d \in \overline{\Pi}_d$ solving Problem 3.

Based on results of the previous section on merging Markov chains, the desired solution must have a structure as in (9), or, more explicitly,

$$\overline{\pi}_{d_i} = \sum_{(\tau_h,\sigma_k)\in\chi_{d_i}} \overline{\pi}_{\tau_h}\overline{\pi}_{\sigma_k} \tag{10}$$

where

$$\chi_{d_k} = \{(\tau_i,\sigma_j) \in L_{\tau\sigma} \mid \min(\tau_i,\sigma_j) = d_k \in L_{(\tau\sigma)^\star}\}.$$

It is worth noting that, even if the process $(\tau\sigma)^\star$ may not be a homogeneous irreducible aperiodic Markov chain, the problem of AS-stability still makes sense and the contractivity conditions can be used. Indeed, the state evolution of the JLS driven by the aggregated process $(\tau\sigma)^\star$ is the same as the one produced by an equivalent MJLS with n^2 states and driven by the Markov chain $\tau\sigma$. The equivalent MJLS is constructed by associating to $(\tau_i,\sigma_j) \in L_{\tau\sigma}$ the index $\mu(\tau_i,\sigma_j) \triangleq \min\{i,j\}$, hence the controlled system $\hat{A}_{\mu(\tau_i,\sigma_j)}$.

Indeed condition $P3.1)$ of Problem 3 can be rewritten using (10) as

$$\prod_{i=1}^{n} \left\|\widehat{A}_i\right\|^{\sum_{(\tau_h,\sigma_k)\in\chi_{d_i}} \overline{\pi}_{\tau_h}\overline{\pi}_{\sigma_k}} = \prod_{(\tau_h,\sigma_k)\in L_{\tau\sigma}} \left\|\widehat{A}_{\mu(\tau_h,\sigma_k)}\right\|^{\overline{\pi}_{\tau_h}\overline{\pi}_{\sigma_k}}.$$

Our synthesis Problem 2 is therefore reduced to finding a vector $\overline{\pi}_\sigma = \begin{bmatrix}\overline{\pi}_{\sigma_1} & \cdots & \overline{\pi}_{\sigma_n}\end{bmatrix}^T$ such that the resulting $\overline{\pi}_d$ given by (10) is a solution to Problem 3, i.e. $\overline{\pi}_d \in \overline{\Pi}_d$.

It actually turns out that the choice of the structure of $\overline{\pi}_d$ described in (10), resulting from the choice of an independent conditioning chain, simplifies the formulation of the synthesis problem substantially. Indeed, the following lemma can be proved by simple if lenghty arguments, which are omitted for brevity.

Lemma 1: Constraints $P3.2), P3.3)$, and $P3.4)$ in Problem 3 are satisfied by any $\overline{\pi}_d = \begin{bmatrix}\overline{\pi}_{d_1} & \cdots & \overline{\pi}_{d_n}\end{bmatrix}^T$ with $\overline{\pi}_{d_i} = \sum_{(\tau_h,\sigma_k)\in\chi_{d_i}} \overline{\pi}_{\tau_h}\overline{\pi}_{\sigma_k}$, $0 < \overline{\pi}_{\sigma_i} < 1$, and $\sum_{i=1}^n \overline{\pi}_{\sigma_i} = 1$.

Furthermore, for $\overline{\pi}_d$ as in (10), constraint $P3.1)$ can be rewritten (proviso $\left\|\widehat{A}_i\right\| \neq 0\ \forall i$) as

$$\ln\left(\prod_{i=1}^{n}\left\|\widehat{A}_i\right\|^{\overline{\pi}_{d_i}}\right) = \sum_{i=1}^{n}\overline{\pi}_{d_i}\ln\left(\left\|\widehat{A}_i\right\|\right)$$
$$= \sum_{i=1}^{n}\overline{\pi}_{\sigma_i}\sum_{h=1}^{n}\overline{\pi}_{\tau_h}\ln\left(\left\|\widehat{A}_{\mu(\tau_h,\sigma_i)}\right\|\right) < 0.$$

In the light of previous analysis, the synthesis Problem 2 can be written as the Linear Programming problem:

$$\begin{aligned}
&\text{Find a vector } \overline{\pi}_\sigma = \begin{bmatrix}\overline{\pi}_{\sigma_1} & \cdots & \overline{\pi}_{\sigma_n}\end{bmatrix}^T \text{ such that}\\
&1)\qquad \sum_{i=1}^{n} c_i\overline{\pi}_{\sigma_i} < 0\\
&2)\qquad 0 < \overline{\pi}_{\sigma_i} < 1\\
&3)\qquad \sum_{i=1}^{n}\overline{\pi}_{\sigma_i} = 1,
\end{aligned} \tag{11}$$

where

$$c_i = \sum_{h=1}^{n}\overline{\pi}_{\tau_h}\ln\left(\left\|\widehat{A}_{\mu(\tau_h,\sigma_i)}\right\|\right).$$

B. Multi-Step Solution

So far, the design of a stochastic switching law based on a conditioning Markov chain has been formulated using the one-step average contractivity condition (4). As already pointed out, this condition might well be not satisfiable for a given set of controllers. To tackle this problem, we will use a multi-step lifting technique, as described in section II-B. However, as suggested in section III, instead of using the lifted version of a chain σ for conditioning, we employ an unconstrained chain σ^m. This has the consequence of associating to substrings of matrices a steady-state probability of occurrence that is in general different from the product of the probabilities of each single matrix. For instance, the probability of occurrence of the string $\widehat{A}_1\widehat{A}_2$ can be different from the probability of $\widehat{A}_2\widehat{A}_1$. It should be pointed out that, while in the one-step solution it is attempted to execute the controller suggested by the particular realization of the process $\sigma(t)$, in the multi-step solution the switching policy suggests the sequence of controllers to be executed in the next m steps.

In an m-step lifting, the scheduler Markov chain states become strings of the original symbols τ_i, taking values in the new state space L_τ^m with cardinality n^m. Let τ^m denote the lifted chain and $\widetilde{\tau}_{\widetilde{\imath}} \in L_\tau^m$ its states ($\widetilde{\imath} \in I^m$). To the aim of designing a switching law, consider a Markov chain σ^m with n^m states taking values in a finite state space L_{σ^m}. To simplify the description of the switching policy, assume that any $\widetilde{\sigma}_{\widetilde{\jmath}} \in L_{\sigma^m}$ is a string of symbols $\sigma_j \in L_\sigma$, hence we have $L_{\sigma^m} \triangleq L_\sigma^m$. Notice however that, notwithstanding the choice of using the same set of symbols, σ^m is *not* in general the lifted version of a chain of n nodes, rather it is a chain with the same number of states as τ^m.

Suppose now that a set of steady-state probability distributions $\widetilde{\overline{\Pi}}_d$ exists solving Problem 3 for the lifted system. The synthesis problem is then again to find a steady-state probability distribution $\widetilde{\overline{\pi}}_\sigma$ for the chain σ^m such that the aggregated process $(\tau^m\sigma^m)^*$ has steady-state distribution $\widetilde{\overline{\pi}}^\star_{\tau^m\sigma^m} \in \widetilde{\overline{\Pi}}_d$.

If we set $L_\sigma = L_\tau$, the aggregation function is again the *minimum* function, applied element-wise. With these assumptions, the overall problem can be formulated in the parameters $\widetilde{\overline{\pi}}_{\sigma_{\widetilde{\imath}}}$ as in (11).

VI. Conclusions

We considered the problem of scheduling the execution of different, hierarchically ordered tasks designed for anytime control of a linear plant. Given a stochastic model of the scheduler, and the set of controllers, we formulated a linear program whose solutions provide a switching law that conditions the scheduler so that the resulting switching system is stable in a probabilistic sense. Although solvability for this problem is not guranteed for one-step switching laws, we have shown that for any set of stabilizing anytime controllers it is possible to find a long enough step horizon m, such that a m-step switching law exists providing almost sure stability. Further work will be devoted to provide constructive methods to synthesize anytime controllers for which the above results can be applied with small m, and to study the performance of the controlled system under switching.

VII. Acknowledgement

The authors gratefully acknowledge Bruno Picasso for suggesting the use of Almost Sure stability tools to approach the problem of anytime control, and for generously providing his insights and technical comments in the writing of this paper.

Appendix

Before proving Theorem 4 we need some preliminary results on primitive matrices. Recall that the transition probability matrix of a finite-state homogeneous irreducible aperiodic Markov chain is a time-invariant stochastic irreducible aperiodic matrix of finite dimension.

Definition 2 ([10] p.127): A matrix is *primitive* if it is irreducible and aperiodic.

Theorem 5 ([10] p.128): Let A be a nonnegative matrix[1]. The following are equivalent.

1) A is primitive.
2) $A^m > 0$ for some $m \geq 1$.
3) $A^m > 0$ for all sufficiently large m.

Primitivity is preserved by the Kronecker product.

Lemma 2: Given $A \geq 0$ and $B \geq 0$ primitive matrices, then $A \otimes B$ is a primitive matrix.

Proof. From Theorem 5 we know that there exist $m_1, m_2 \geq 1$ such that $A^{m_1} > 0$ and $B^{m_2} > 0$. From point 3 of the same theorem we know that there exists $m \geq \max(m_1, m_2)$ such that $A^m > 0$ and $B^m > 0$. Recalling the definition of Kronecker product, it is apparent that

$$A^m \otimes B^m > 0$$

and using the 'mixed product rule'

$$A^m \otimes B^m = (A \otimes B)^m > 0,$$

hence $A \otimes B$ is primitive. ■

Proof. [of Theorem 4]

i) Let us prove first that $P_{\alpha\beta}$ is a time-invariant stochastic primitive matrix. The first two properties follow directly by the same properties of P_α and P_β and by the definition of Kronecker product. The other property is proved by the Lemma 2. We show now that the statistics of the process $\alpha\beta$ are given by $P_{\alpha\beta}$ and $\pi_{\alpha\beta}(0) = \pi_\alpha(0) \otimes \pi_\beta(0)$. To this aim, let us compute the transition probability $\Pr\{\alpha\beta(t+1) = (\alpha_k, \beta_h) \mid \alpha\beta(t) = (\alpha_i, \beta_i)\}$ for any $i, j, h, k \in I$ and any $t \in \mathbb{N}$. Recalling that $\alpha\beta(t) = (\alpha_i, \beta_i)$ can be considered as the joint event $(\alpha(t) = \alpha_i) \cap (\beta(t) = \beta_i)$, we can omit the dependence of t and write $\Pr\{\alpha\beta(t+1) = (\alpha_k, \beta_h) \mid \alpha\beta(t) = (\alpha_i, \beta_i)\} = \Pr\{\alpha_k \cap \beta_h \mid \alpha_i \cap \beta_j\}$. Due to the independence of the events α_k and β_h, we can write

$$\Pr\{\alpha_k \cap \beta_h \mid \alpha_i \cap \beta_j\} = \Pr\{\alpha_k \mid \alpha_i \cap \beta_j\} \Pr\{\beta_h \mid \alpha_i \cap \beta_j\}.$$

Let us consider $\Pr\{\alpha_k \mid \alpha_i \cap \beta_j\}$ (the case of $\Pr\{\beta_h \mid \alpha_i \cap \beta_j\}$ is similar) and apply the Bayes rule

$$\begin{aligned}
\Pr\{\alpha_k \mid \alpha_i \cap \beta_j\} &= \frac{\Pr\{\alpha_i \cap \beta_j \mid \alpha_k\} \Pr\{\alpha_k\}}{\Pr\{\alpha_i \cap \beta_j\}} \\
&= \frac{\Pr\{\alpha_i \mid \alpha_k\} \Pr\{\beta_j \mid \alpha_k\} \Pr\{\alpha_k\}}{\Pr\{\alpha_i\} \Pr\{\beta_j\}} \\
&= \left[\frac{\Pr\{\alpha_k \mid \alpha_i\} \Pr\{\alpha_i\}}{\Pr\{\alpha_k\}} \frac{\Pr\{\alpha_k \mid \beta_j\} \Pr\{\beta_j\}}{\Pr\{\alpha_k\}}\right] \cdot \\
&\quad \frac{\Pr\{\alpha_k\}}{\Pr\{\alpha_i\} \Pr\{\beta_j\}} \\
&= \frac{\Pr\{\alpha_k \mid \alpha_i\} \Pr\{\alpha_k \mid \beta_j\}}{\Pr\{\alpha_k\}} \\
&= \Pr\{\alpha_k \mid \alpha_i\},
\end{aligned}$$

where the last term is obtained noting again that the events α_k and β_j are independent (hence $\Pr\{\alpha_k \mid \beta_j\} = \Pr\{\alpha_k\}$). Therefore

$$\Pr\{\alpha_k \cap \beta_h \mid \alpha_i \cap \beta_j\} = \Pr\{\alpha_k \mid \alpha_i\} \Pr\{\beta_h \mid \beta_j\} = {}^{\alpha}p_{ik} {}^{\beta}p_{jh}.$$

Keeping the transition $\alpha_i \longrightarrow \alpha_k$ and considering all the transitions $\beta_h \longrightarrow \beta_j$ $\forall \beta_h, \beta_j \in L_\beta$, we find that all the transition probabilities are given by ${}^{\alpha}p_{ik} P_\beta$. Using the same argument for each transition $\alpha_i \longrightarrow \alpha_k$ $\forall \alpha_i, \alpha_k \in L_\alpha$ and defining the indices $l = (i-1)n + j$ and $r = (k-1)n + h$, one can easily find that

$${}^{\alpha\beta}p_{lr} = \Pr\{\alpha_k \cap \beta_h \mid \alpha_i \cap \beta_j\} = {}^{\alpha}p_{ik} {}^{\beta}p_{jh},$$

or in matrix form

$$P_{\alpha\beta} = P_\alpha \otimes P_\beta.$$

Moreover, from the properties of independent random variables, we have that

$$\pi_{\alpha\beta}(0) = \pi_\alpha(0) \otimes \pi_\beta(0).$$

ii) From the independence of the random variables $\alpha(t)$ and $\beta(t)$ $\forall t \in \mathbb{N}$, we have that

$$\pi_{\alpha\beta}(t) = \pi_\alpha(t) \otimes \pi_\beta(t) \tag{12}$$

and from the previous point

$$\pi_{\alpha\beta}(t) = \left[P_{\alpha\beta}^t\right]^T (\pi_\alpha(0) \otimes \pi_\beta(0)).$$

From (12) it is apparent that $\lim_{t\to\infty} \pi_{\alpha\beta}(t) = \overline{\pi}_\alpha \otimes \overline{\pi}_\beta$ for any $\pi_{\alpha\beta}(0) = \pi_\alpha(0) \otimes \pi_\beta(0)$. To extend this property to any initial distribution $\pi_{\alpha\beta}(0)$, it is sufficient to recall that the steady-state probability distribution of a homogeneous irreducible aperiodic Markov chain is unique. ■

[1]With $M \geq 0$ ($M > 0$) we mean nonnegative (positive) matrices. Stochastic matrices are a subset of nonnegative matrices.

REFERENCES

[1] C. L. Liu and J. W. Layland, “Scheduling algorithms for multiprogramming in a hard-real-time environment,” *Journal of the Association for Computing Machinery*, vol. 20, no. 1, 1973.

[2] J. W. S. Liu, *Real–Time Systems*, T. Holm, Ed. Upper Saddle River, NJ: Prentice Hall Inc., 2000.

[3] G. Buttazzo, *Hard Real-Time Computing Systems: Predictable Scheduling Algorithms and Applications*. Boston: Kluwer Academic Publishers, 1997.

[4] J. W. S. Liu, W.-K. Shih, K.-J. Lin, R. Bettati, and J.-Y. Chung, “Imprecise computation,” *Proceedings of the IEEE*, vol. 82, no. 1, pp. 83–93, January 1994.

[5] R. Bhattacharya and G. J. Balas, “Anytime control algorithm: Model reduction approach,” *Journal of Guidance, Control, and Dynamics*, vol. 27, no. 5, pp. 767–776, 2004.

[6] Y. Fang, K. Loparo, and X. Feng, “Almost sure and δ-moment stability of jump linear systems,” *Int. J. Control*, vol. 59, no. 5, pp. 1281–1307, 1994.

[7] P. Bolzern, P. Colaneri, and G. D. Nicolao, “On almost sure stability of discrete-time Markov jump linear systems,” in *Proc. 43rd IEEE Conf. On Decision and Control*, vol. 3, 2004, pp. 3204–3208.

[8] C. G. Cassandras and S. Lafortune, *Introduction to Discrete Event Systems*. Kluwer Academic Publishers, 1999.

[9] D. Liberzon, J. P. Hespanha, and A. S. Morse, “Stability of switched systems: A Lie-algebraic condition,” *Systems & Control Letters*, vol. 37, no. 3, pp. 117–122, 1999.

[10] D. A. Lind and B. Marcus, *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press, 1995.