

EVENT-TRIGGERED REAL-TIME SCHEDULING OF STABILIZING CONTROL TASKS

PAULO TABUADA

ABSTRACT. In this note we revisit the problem of scheduling stabilizing control tasks on embedded processors. We start from the paradigm that a real-time scheduler should be regarded as a feedback controller that decides which task is executed at any given instant. This controller has for objective guaranteeing that software tasks meet its deadlines and that stabilizing control tasks asymptotically stabilize the plant. According to this feedback paradigm, the decision of executing control tasks should not be based on release times and deadlines but rather on the state of the plant. We investigate the feasibility of a simple event-triggered scheduler based on the state norm and provide some schedulability results.

1. INTRODUCTION

Small embedded microprocessors are quickly becoming an essential part of the most diverse applications. A particularly interesting example are physically distributed sensor/actuator networks responsible for collecting and processing information, and to react to this information through actuation. The embedded microprocessors forming the computational core of these networks are required to execute a variety of tasks comprising the relay of information packets in multi-hop communication schemes, monitoring physical quantities, and computation of feedback control laws. Since we are dealing with resource limited microprocessors it becomes important to assess to what extent we can increase the functionality of these embedded devices through novel real-time scheduling algorithms based on event-triggered rather than time-triggered execution of control tasks.

We investigate in this note a very simple event-triggered scheduling algorithm that preempts running tasks to execute the control task whenever a certain error becomes large when compared with the state norm. This idea is an adaptation to the scheduling context of several techniques used to study problems of control under communication constraints [NE00, BL00, EM01]. We take explicitly into account the execution time of the control task and show that the proposed scheduling policy guarantees global asymptotical stability. We also provide sufficient conditions for co-schedulability of the control task with other tasks competing for processor time. The proposed approach is illustrated with simulation results.

Real-time scheduling of control tasks has received renewed interest from the academic community in the past years [SLCB00, CE00, ACR⁺00, CEBA02, BA02, CLS03, LHLQ06]. Common to all these approaches is the underlying principle that better control performance is achieved by providing more CPU time to control tasks. This can be accomplished in two different ways: letting control tasks run for longer amounts of time using anytime implementations or model predictive controllers; or by scheduling control tasks more frequently. All these approaches assume the existence of a performance criterion for the control task such as a cost function used to design an optimal linear quadratic regulator. Scheduling strategies are then obtained through optimization algorithms seeking to determine schedules maximizing the performance criterion. The work presented in this note does not resort to optimization and does not require a performance criterion. Instead, the decision to execute the control task is determined by a feedback mechanism based on the state of the plant.

Closer to the results presented in this note is the work described in [PPV⁺02, PPBSV05], where resource allocation and feedback control are designed in an integrated fashion. Several concurrent controllers described

This research was partially supported by the National Science Foundation EHS award 0509313.

by scalar gains and activation rates of the corresponding processes are designed so as to ensure stability of the controlled processes as well as real-time schedulability. Close at the technical level, although addressing very different problems, is the recent work on stabilization under communication constraints [NE00, BL00, EM01, BPZ02, Lib03, NT04]. All these approaches are concerned with the stabilization of continuous systems under reduced communication and the employed techniques share with some of the techniques described in this note a common ancestor: the perturbation approach to stability analysis of control systems, described for example in [Kha96]. Similar techniques have also been used in [LNT02] to show how sample-and-hold implementations of stabilizing controllers guarantee stability under sufficiently fast time-triggered executions. Finally, we would like to refer the reader to [AB02] where some advantages of event-driven control over time-driven control are presented in a stochastic setting. A preliminary version of the results presented in this note was reported in [TW06].

2. NOTATION AND PROBLEM STATEMENT

2.1. Notation. We shall use the notation $|x|$ to denote the Euclidean norm of an element $x \in \mathbb{R}^n$. Given matrices A and B , $[A|B]$ denotes the matrix formed by the columns of matrix A followed by the columns of matrix B . A function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is said to be Lipschitz continuous on compacts if for every compact set $S \subset \mathbb{R}^n$ there exists a constant $L > 0$ such that:

$$|f(x) - f(y)| \leq L|x - y|$$

for every $x, y \in S$. A continuous function $\alpha : [0, a[\rightarrow \mathbb{R}_0^+$, $a > 0$, is said to be of class \mathcal{K} if it is strictly increasing and $\alpha(0) = 0$. It is said to be of class \mathcal{K}_∞ if $a = \infty$ and $\alpha(r) \rightarrow \infty$ as $r \rightarrow \infty$.

2.2. Problem statement. We consider a control system:

$$(2.1) \quad \dot{x} = f(x, u), \quad x \in \mathbb{R}^n, u \in \mathbb{R}^m$$

for which a feedback controller:

$$(2.2) \quad u = k(x)$$

has been designed rendering the closed loop system:

$$(2.3) \quad \dot{x} = f(x, k(x + e))$$

Input-to-State Stable (ISS) with respect to measurement errors $e \in \mathbb{R}^n$. We shall not need the definition¹ of ISS in this note but rather the following characterization.

Definition 2.1. A smooth function $V : \mathbb{R}^n \rightarrow \mathbb{R}_0^+$ is said to be an ISS Lyapunov function for the closed loop system (2.3) if there exist class \mathcal{K} functions $\underline{\alpha}, \bar{\alpha}, \alpha$ and γ satisfying:

$$(2.4) \quad \underline{\alpha}(|x|) \leq V(x) \leq \bar{\alpha}(|x|)$$

$$(2.5) \quad \frac{\partial V}{\partial x} f(x, k(x + e)) \leq -\alpha(|x|) + \gamma(|e|)$$

Closed loop system (2.3) is said to be ISS with respect to measurement errors $e \in \mathbb{R}^n$ if there exists an ISS Lyapunov function for (2.3).

The implementation of the feedback law (2.2) on an embedded processor is typically done by sampling the state at time instants:

$$t_0, t_1, t_2, t_3, t_4 \dots$$

computing $u(t_i) = k(x(t_i))$ and updating the actuator values at time instants:

$$t_0 + \Delta, t_1 + \Delta, t_2 + \Delta, t_3 + \Delta, t_4 + \Delta, \dots$$

¹See, for example, [Son04] for an introduction to ISS and related notions.

where $\Delta \geq 0$ represents the time required to read the state from the sensors, compute the control law and update the actuators. This means that to a sequence of measurements:

$$x(t_0), x(t_1), x(t_2), x(t_3), x(t_4), \dots$$

there corresponds a sequence of actuation updates:

$$u(t_0 + \Delta), u(t_1 + \Delta), u(t_2 + \Delta), u(t_3 + \Delta), u(t_4 + \Delta), \dots$$

Between actuator updates the input value u is held constant according to:

$$(2.6) \quad t \in [t_i + \Delta, t_{i+1} + \Delta[\implies u(t) = u(t_i + \Delta)$$

Furthermore, the sequence of times $t_0, t_1, t_2, t_3, t_4, \dots$ is typically periodic meaning that $t_{i+1} - t_i = T$ where $T > 0$ is the period. We can thus regard the execution of the control task implementing the control law (2.2) as being time-triggered. In this note we consider instead event-triggered executions where the sequence $t_0, t_1, t_2, t_3, t_4, \dots$ of execution times is no longer periodic neither specified in advance but rather implicitly defined by an execution rule based on the state of the plant. To introduce this execution rule we define the measurement error e to be:

$$(2.7) \quad t \in [t_i + \Delta, t_{i+1} + \Delta[\implies e(t) = x(t_i) - x(t)$$

We can thus describe the evolution of (2.1) under the implementation (2.6) of control law (2.2) by:

$$\dot{x} = f(x, k(x + e))$$

with $e \in \mathbb{R}^n$ as defined in (2.7).

Let us first consider the hypothetical case $\Delta = 0$ with the single purpose of explaining the execution rule. From (2.5) we see that if we restrict the error to satisfy:

$$(2.8) \quad \gamma(|e|) \leq \sigma \alpha(|x|), \quad \sigma > 0$$

the dynamics of V is bounded by:

$$\frac{\partial V}{\partial x} f(x, k(x + e)) \leq (\sigma - 1) \alpha(|x|)$$

thus guaranteeing that V decreases provided that $\sigma < 1$. Inequality (2.8) can be enforced by executing the control task when:

$$(2.9) \quad \gamma(|e|) \geq \sigma \alpha(|x|)$$

since $\Delta = 0$ implies that if the control task is executed at time t_i we will have $e(t_i) = x(t_i) - x(t_i) = 0$ and $\gamma(|e(t_i)|) = 0$ thus enforcing (2.8). When $\Delta > 0$, the control task needs to be executed *before* the inequality $\gamma(|e|) \geq \sigma \alpha(|x|)$ is satisfied in order to account for the delay Δ between measuring the state and updating the actuators.

Although the simple execution rule (2.9) guarantees global asymptotical stability by construction, there are three important questions that need to be answered in order to assess the feasibility of this scheduling policy:

- (1) Since the execution times are only implicitly defined, can we guarantee that they will not become arbitrarily close resulting in an accumulation² point?
- (2) In the absence of accumulation points, can we compute an estimate of the time elapsed between consecutive executions of the control task?
- (3) How can we use the execution rule (2.9) when there are more tasks competing for processor time and still guarantee that no deadlines are missed?

Answering the above questions is the objective of the following sections.

²In the context of hybrid systems this corresponds to another example of the infamous Zeno behavior [ATS06].

3. EXISTENCE OF A LOWER BOUND FOR INTER-EXECUTION TIMES

We start immediately with one of the main contributions of this note.

Theorem 3.1. *Let $\dot{x} = f(x, u)$ be a control system and let $u = k(x)$ be a control law rendering the closed loop system ISS with respect to measurement errors. If the following assumptions are satisfied:*

- (1) $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ is Lipschitz continuous on compacts;
- (2) $k : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is Lipschitz continuous on compacts;
- (3) There exists an ISS Lyapunov function V for the closed loop system satisfying (2.5) with α^{-1} and γ Lipschitz continuous on compacts,

then, for any compact set $S \subset \mathbb{R}^n$ containing the origin, there exists an $\varepsilon > 0$ such that for all $\Delta \in [0, \varepsilon]$ there exists a time $\tau \in \mathbb{R}^+$ such that for any initial condition in S the inter-execution times $\{t_{i+1} - t_i\}_{i \in \mathbb{N}}$ implicitly defined by the execution rule (2.9) are lower bounded by τ , that is, $t_{i+1} - t_i \geq \tau$ for any $i \in \mathbb{N}$.

Proof. Let R be the compact set defined by all the points $x \in \mathbb{R}^n$ satisfying $V(x) \leq \mu$ where $\mu > 0$ is large enough so that $S \subset R$. Such μ always exists since by (2.4) V is proper or radially unbounded. Set R is forward invariant for the closed loop system since the execution rule (2.9) guarantees $\dot{V} \leq 0$. We now define another compact set E by all the points $e \in \mathbb{R}^n$ satisfying $|e| \leq \gamma^{-1}(\sigma\alpha(|x|))$ for all $x \in R$. Since α^{-1} and γ are Lipschitz continuous on compacts, then so is $\alpha^{-1}(\gamma(|r|)/\sigma)$. Let P be the Lipschitz constant for the compact set E so that $|\alpha^{-1}(\gamma(|r|)/\sigma) - \alpha^{-1}(\gamma(|s|)/\sigma)| \leq P|r - s|$. If $r = e$ and $s = 0$ we obtain $\alpha^{-1}(\gamma(|e|)/\sigma) \leq P|e|$. Note that by enforcing $P|e| \leq |x|$ we guarantee $\alpha^{-1}(\gamma(|e|)/\sigma) \leq |x|$ (which is (2.8)) so that it suffices to show that the inter-execution times are bounded for the execution rule $P|e| \geq |x|$. As a first step towards showing boundedness we note that it follows from Lipschitz continuity on compacts of $f(x, u)$ and $k(x)$ that $f(x, k(x + e))$ is also Lipschitz continuous on compacts, that is:

$$|f(r, k(r + s)) - f(r', k(r' + s'))| \leq L|(r, s) - (r', s')|$$

by taking $r = x$, $s = e$ and $r' = 0 = e'$ we obtain:

$$(3.1) \quad |f(x, k(x + e))| \leq L|(x, e)| \leq L|x| + L|e|$$

when $(x, e) \in R \times E$. Note that $R \times E$ is forward invariant for the state as well as the error dynamics. We can now bound the inter-event times by looking at the dynamics of $|e|/|x|$.

$$\begin{aligned}
 \frac{d}{dt} \frac{|e|}{|x|} &= \frac{d}{dt} \frac{(e^T e)^{1/2}}{(x^T x)^{1/2}} \\
 &= \frac{(e^T e)^{-1/2} e^T \dot{e} (x^T x)^{1/2} - (x^T x)^{-1/2} x^T \dot{x} (e^T e)^{1/2}}{x^T x} \\
 &= \frac{e^T \dot{x}}{|e||x|} - \frac{x^T \dot{x}}{|x||x|} \frac{|e|}{|x|} \\
 &\leq \frac{|e||\dot{x}|}{|e||x|} + \frac{|x||\dot{x}|}{|x||x|} \frac{|e|}{|x|} \\
 &= \left(1 + \frac{|e|}{|x|}\right) \frac{|\dot{x}|}{|x|} \\
 &\leq \left(1 + \frac{|e|}{|x|}\right) \frac{L|x| + L|e|}{|x|} \quad \text{by (3.1)} \\
 (3.2) \quad &= L \left(1 + \frac{|e|}{|x|}\right) \left(1 + \frac{|e|}{|x|}\right)
 \end{aligned}$$

If we denote $|e|/|x|$ by y we have the estimate $\dot{y} \leq L(1 + y)^2$ and we conclude that $y(t) \leq \phi(t, \phi_0)$ where $\phi(t, \phi_0)$ is the solution of $\dot{\phi} = L(1 + \phi)^2$ satisfying $\phi(t, \phi_0) = \phi_0$.

Assume now that $\Delta = 0$. Then, the inter-execution times are bounded by the time it takes for ϕ to evolve from 0 to $1/P$, that is, the inter-execution times are bounded by the solution $\tau \in \mathbb{R}^+$ of $\phi(\tau, 0) = 1/P$. Since $\phi(\tau, 0) = \tau L / (\tau L - 1)$ we obtain $\tau = 1/(L + LP)$.

For $\Delta > 0$ we need a more detailed analysis. First, pick σ' satisfying $\sigma < \sigma' < 1$ (take for example $\sigma' = \sigma + (1 - \sigma)/2$) and let P' be the Lipschitz constant for $\alpha^{-1}(\gamma(|e|)/\sigma')$. Let now $\varepsilon_1 \in \mathbb{R}^+$ satisfy $\phi(\varepsilon_1, 1/P) = 1/P'$. Such ε_1 always exists since ϕ is continuous, $\dot{\phi} > 0$ and $1/P < 1/P'$. Then, by executing the control task at time t_i , defined by $P|e| = |x|$, we guarantee that for $t \in [t_i, t_i + \varepsilon_1[$ we have $|e| \leq |x|/P'$ and thus also $\gamma(|e|) \leq \sigma'\alpha(|x|)$. Since $\sigma' < 1$ asymptotic stability is still guaranteed. The inter-execution times are now bounded by $\Delta + \tau$ where τ is the time it takes for ϕ to evolve from $|e(t_i + \Delta)|/|x(t_i + \Delta)| = |x(t_i) - x(t_i + \Delta)|/|x(t_i + \Delta)|$ to $1/P$. We thus need to pick Δ small enough so that $|e(t_i + \Delta)|/|x(t_i + \Delta)| < 1/P$ since $\dot{\phi} > 0$. It now follows from continuity³ of $|x(t_i) - x(t_i + \Delta)|/|x(t_i + \Delta)|$ with respect to Δ the existence of $\varepsilon_2 > 0$ such that for any $0 \leq \Delta \leq \varepsilon_2$ we have $|x(t_i) - x(t_i + \Delta)|/|x(t_i + \Delta)| < 1/P$. The proof is now finished by taking $\varepsilon = \min\{\varepsilon_1, \varepsilon_2\}$. \square

Theorem 3.1 shows that the simple execution rule (2.9) results in a sequence of inter-execution times for the control task that is guaranteed to be lower bounded provided that Δ is sufficiently small. The techniques used in the proof rely on Lipschitz continuity and are necessarily conservative for general nonlinear systems. However, for linear systems they provide reasonable estimates and we can even provide computable bounds for Δ as discussed in the next section.

4. THE LINEAR CASE

In this section we discuss the linear case in some detail since the arguments in the proof of Theorem can be refined in order to provide a computational estimate for τ . We thus assume the control system to be of the form:

$$\dot{x} = Ax + Bu, \quad x \in \mathbb{R}^n, u \in \mathbb{R}^m$$

with A and B matrices of appropriate dimensions. We also assume the existence of a linear feedback:

$$u = Kx$$

rendering the closed loop system globally asymptotically stable and where K is a matrix of appropriate dimensions. Note that in the linear case any such K renders the closed loop system ISS with respect to measurement errors. We thus have a Lyapunov function $V : \mathbb{R}^n \rightarrow \mathbb{R}_0^+$ satisfying:

$$(4.1) \quad \underline{a}|x|^2 \leq V(x) \leq \bar{a}|x|^2$$

$$(4.2) \quad \frac{\partial V}{\partial x}(Ax + BKx + BKe) \leq -a|x|^2 + g|e||x|$$

with $\underline{a}, \bar{a}, a, g \in \mathbb{R}^+$.

Corollary 4.1. *Let $\dot{x} = Ax + Bu$ be a linear control system, let $u = Kx$ be a linear control law rendering the closed loop system globally asymptotically stable and assume that $\Delta = 0$. For any initial condition in \mathbb{R}^n the inter-event times $\{t_{i+1} - t_i\}_{i \in \mathbb{N}}$ implicitly defined by the execution rule:*

$$|e| \geq \sigma|x|$$

are lower bounded by the time τ satisfying:

$$(4.3) \quad \phi(\tau, 0) = \sigma$$

where $\phi(t, \phi_0)$ is the solution of:

$$(4.4) \quad \dot{\phi} = |A + BK| + (|A + BK| + |BK|)\phi + |BK|\phi^2$$

³Note that $x(t_i + \Delta)$ is never zero since the closed loop system converges asymptotically to zero and thus never reaches zero in finite time.

satisfying $\phi(0, \phi_0) = \phi_0$. Furthermore, for $\Delta > 0$ and for any desired $\sigma > 0$, the execution rule:

$$|e| \geq \sigma'|x|$$

with:

$$(4.5) \quad \frac{\Delta |A + BK|BK|(\sigma + 1)}{1 - \Delta |A + BK|BK|(\sigma + 1)} \leq \sigma' \leq \phi(-\Delta, \sigma)$$

enforces for any $i \in \mathbb{N}$ and for any $t \in [t_i + \Delta, t_{i+1} + \Delta[$ the following inequality:

$$|e(t)| \leq \sigma|x(t)|$$

with inter-execution times bounded by $\Delta + \tau$ where time τ satisfies:

$$(4.6) \quad \phi\left(\tau, \frac{\Delta |A + BK|BK|(\sigma + 1)}{1 - \Delta |A + BK|BK|(\sigma + 1)}\right) = \sigma'$$

Proof. Estimate (4.4) follows directly from (3.2) by using $|\dot{x}| \leq |A+BK||x| + |BK||e|$. Inequality $\sigma' \leq \phi(-\Delta, \sigma)$ guarantees that if $|e| = \sigma'|x|$ holds at time t_i , triggering an execution of the control task, then the inequality $|e| \leq \sigma|x|$ will hold for $t \in [t_i, t_i + \Delta[$. To conclude boundedness of inter-execution times we need to ensure that $|e(t_i + \Delta)|/|x(t_i + \Delta)| \leq \sigma'$ since if this inequality fails to hold, an execution of the control task is requested before the termination of the previous execution. Let e' denote the error defined by $e'(t) = x(t_i) - x(t)$ for $t \in [t_i, t_i + \Delta]$ and note that $|e(t_i + \Delta)|/|x(t_i + \Delta)|$ is by definition equal to $|e'(t_i + \Delta)|/|x(t_i + \Delta)|$. To estimate $|e'(t)|/|x(t)|$ we compute $\frac{d}{dt}|e'(t)|/|x(t)|$ following the same argument as in the proof of Theorem 3.1 to obtain:

$$(4.7) \quad \frac{d}{dt} \frac{|e'|}{|x|} \leq L \left(1 + \frac{|e'|}{|x|}\right) \left(1 + \frac{|e|}{|x|}\right)$$

with $L = |A + BK|BK|$. The evolution of $|e'|/|x|$ is thus bounded by $\psi(t)$ which is the solution satisfying $\psi(t_i) = 0$ of:

$$\dot{\phi} = L(1 + \phi)^2 \quad \dot{\psi} = L(1 + \psi)(1 + \phi)$$

since, as shown in the proof of Theorem (3.1), $|e|/|x|$ is bounded by $\phi(t)$. Since:

$$\psi(t) = \frac{tL \left(\frac{|e(t_i)|}{|x(t_i)|} + 1\right)}{1 - tL \left(\frac{|e(t_i)|}{|x(t_i)|} + 1\right)} = \frac{tL(\sigma' + 1)}{1 - tL(\sigma' + 1)} \leq \frac{tL(\sigma + 1)}{1 - tL(\sigma + 1)}$$

where the inequality follows from $\frac{\partial \psi}{\partial \sigma'} > 0$, we have:

$$\psi(\Delta) < \frac{\Delta L(\sigma + 1)}{1 - \Delta L(\sigma + 1)}$$

and thus assumption (4.5) guarantees that $|e(t_i + \Delta)|/|x(t_i + \Delta)| < \sigma'$. The inter-execution times are now bounded by $\Delta + \tau$ where τ is the time necessary for ϕ to evolve from $\Delta L(\sigma + 1)/(1 - \Delta L(\sigma + 1))$ to σ' , that is, the solution τ of $\phi(\tau, \Delta L(\sigma + 1)/(1 - \Delta L(\sigma + 1))) = \sigma'$. □

Note that for $\Delta = 0$ we can chose $\sigma' = \sigma$ in (4.5) and (4.6) becomes (4.3). Although the formulas in Corollary 4.1 only provide estimates, they are sufficiently accurate to be useful in practical situations as described in the next section.

5. AN ACADEMIC EXAMPLE

We now illustrate the previous results on the linear control system:

$$(5.1) \quad \begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -2 & 3 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u$$

stabilized by the following linear feedback:

$$(5.2) \quad u = x_1 - 4x_2$$

Using $V = x^T P x$ as a Lyapunov function with P defined by:

$$(5.3) \quad P = \begin{bmatrix} 1 & \frac{1}{4} \\ \frac{1}{4} & 1 \end{bmatrix}$$

we obtain $\frac{\partial V}{\partial x}(Ax + BKx) = -x^T Q x$ where:

$$(5.4) \quad Q = \begin{bmatrix} \frac{1}{2} & \frac{1}{4} \\ \frac{1}{4} & \frac{3}{2} \end{bmatrix}$$

Taking the measurement error into account we have:

$$\frac{\partial V}{\partial x}(Ax + BKx + BK\epsilon) \leq -a|x|^2 + g|e||x|$$

where:

$$a = \lambda_m(Q) > 0.44 \quad g = |K^T B^T P + PBK| = 8$$

and $\lambda_m(Q)$ is the smallest eigenvalue of Q . Since σg needs to be smaller than 0.44 we select $\sigma = 0.05$ and obtain the execution rule $|e| \geq 0.05|x|$. For $\Delta = 0.005s$ and according to Corollary 4.1 we can select any σ' satisfying:

$$0.0235 \leq \sigma' \leq 0.0405$$

The theoretical value for the inter-execution time corresponding to $\sigma' = 0.04$ is:

$$\Delta + \tau = 0.005s + 0.0091s = 0.0141s$$

while simulation results for the initial conditions:

$$(x_1(0), x_2(0)) = (10 \cos(\frac{2\pi}{30}\kappa), 10 \sin(\frac{2\pi}{30}\kappa)) \quad \kappa = 1, \dots, 30$$

provided a lower bound of $0.0237s$.

In Figure 1 we can see how the error norm never reaches $\sigma|x|$ even though it goes beyond $\sigma'|x|$ which is used as execution rule. Although for small values of time there is a large gap between the maximum value reached by $|e|$ and $\sigma|x|$, this gap decreases as the state approaches the origin.

For different values of σ' we obtained the inter-execution time estimates reported in Table 1. We can see that the estimated values, although conservative, do not overestimate the values obtained through simulation by more than a factor of 3.

6. CO-SCHEDULABILITY OF STABILIZING CONTROL TASKS

The lower bound for inter-execution times of the control task, presented in Theorem 3.1, can be used to analyze schedulability of a set of tasks $T = \{T_i\}_{i \in I}$. We shall assume a preemptive scheduler in which the control task has the highest priority and thus cannot be preempted by any other task and is executed without delays when $\gamma(|e|) \geq \sigma\alpha(|x|)$. Note that timing overheads associated with context switching can be captured in the proposed framework by suitably enlarging Δ . We regard execution of the control task as a timing overhead imposed on the tasks T_i by the scheduler in the sense that T_i may need to be interrupted to execute the control task. When a set of control tasks T can be scheduled despite the overhead associated with the control

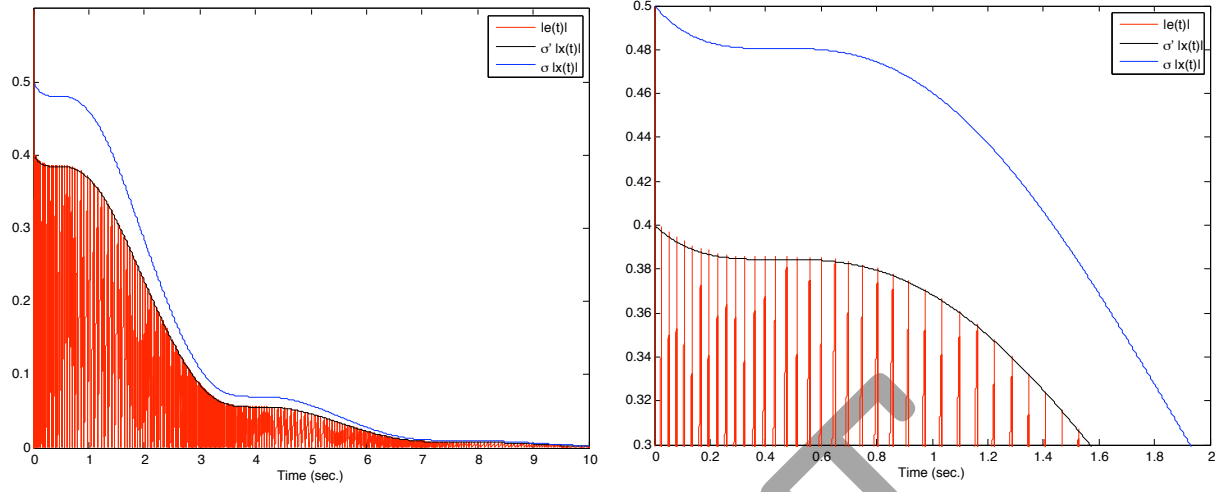


FIGURE 1. Evolution of $|e|$, $\sigma'|x|$ and $\sigma|x|$ for $\Delta = 0.005$, $\sigma = 0.05$ and $\sigma' = 0.04$.

| σ' | $\Delta + \tau$ (estimated) | $\Delta + \tau$ (simulated) | simulated/estimated |
|-----------|-----------------------------|-----------------------------|---------------------|
| 0.0250 | 0.0058s | 0.0151s | 2.60 |
| 0.0275 | 0.0072s | 0.0166s | 2.31 |
| 0.0300 | 0.0086s | 0.0180s | 2.09 |
| 0.0325 | 0.0100s | 0.0194s | 1.94 |
| 0.0350 | 0.0114s | 0.0209s | 1.83 |
| 0.0375 | 0.0127s | 0.0223s | 1.76 |
| 0.0400 | 0.0141s | 0.0237s | 1.68 |

TABLE 1. Simulated and estimated inter-execution times for several values of σ' .

task we say that T is co-schedulable with the control task. Co-schedulability is now ensured by the following sufficient condition where $\lceil r \rceil$ denotes the smallest integer greater than $r \in \mathbb{R}$.

Theorem 6.1. *Let $\dot{x} = f(x, u)$ be a control system, let $u = k(x)$ be a control law rendering the closed loop system ISS with respect to measurement errors and let $T = \{T_i\}_{i \in I}$ be a set of tasks with execution times $\{\Delta_i\}_{i \in I}$. If the set of tasks T is schedulable with new execution times Δ'_i given by:*

$$\Delta'_i = \Delta_i + \lceil \Delta_i / \tau \rceil \Delta$$

then T is co-schedulable with the control task under the execution rule (2.9).

Proof. Since the time between executions of the control task is lower bounded by $\Delta + \tau$, each task T_i can be interrupted at most $\lceil \Delta_i / \tau \rceil$ times. Each interruption will delay the execution of T_i by Δ units of time resulting in a total execution time of $\Delta_i + \lceil \Delta_i / \tau \rceil \Delta$. \square

Co-schedulability is ensured by the existence of a schedule where the duration Δ_i of each task T_i has been inflated to Δ'_i in order to accommodate the timing overhead imposed by the control task. Theorem 6.1 is just one example of the use of the event-triggered scheduling rule discussed in this note. Other possibilities are discussed in the next section.

7. DISCUSSION

7.1. Real-time scheduling policies for non-preemptible tasks. The simple preemptive scheduling strategy presented in Section 6 relied on the possibility to preempt all but the control task. In many situations this assumption may not hold and more elaborate scheduling strategies are necessary. The results presented in this note are also relevant in this more general context since the lower bound on the inter-execution times can be used to construct a timed-automaton model for the control task. This model can then be composed with models of the remaining tasks and a scheduler, regarded as a supervisor, can be synthesized by resorting to control and/or game theoretic techniques for timed automata [AGS02, FMPY06].

7.2. ISS with respect to actuation errors and networked control systems. The results presented in this note are based on the assumption that the controller to be implemented renders the closed loop system ISS with respect to measurement errors. From a theoretical point of view it is not difficult to see that Theorem 3.1 still holds, *mutatis mutantis*, if ISS with respect to actuation errors is assumed in place of ISS with respect to measurement errors. This is an assumption that can be made without loss of generality for control affine systems. In this case, the execution rule would be based on the error $e(t) = k(x(t_i)) - k(x(t))$ which makes this approach uninteresting from the practical point of view as $k(x(t))$ needs to be computed to determine if the control task computing $k(x(t))$ should be executed! Under the ISS with respect to measurement errors assumption used in this note we only need to compute $|x(t_i) - x(t)|$ and $|x(t)|$ to determine if the control task should be executed and this can be done in hardware without requiring processor time. However, the ISS with respect to actuation errors does make sense in a networked control setting where the shared resource to be scheduled is not processor time but rather the transmission medium. Similar ideas have been more or less explicitly explored in [YTS02, MA04, NT04].

8. ACKNOWLEDGMENTS

The author would like to thank his colleague Mike Lemmon for insightful discussions about some of the issues addressed in this note.

REFERENCES

- [AB02] K. J. Astrom and B.M. Bernhardsson. Comparison of riemann and lebesgue sampling for first order stochastic systems. In *Proceedings of the 41st IEEE Conference on Decision and Control*, December 2002.
- [ACR⁺00] P. Albertos, A. Crespo, I. Ripoll, M. Valles, and P. Balbastre. Rt control scheduling to reduce control performance degrading. In *Proceedings of the 39th IEEE Conference on Decision and Control*, Sydney, Australia, 2000.
- [AGS02] K. Altisen, G. Gssler, and J. Sifakis. Scheduler modeling based on the controller synthesis paradigm. *Real-Time Systems: The International Journal of Time-Critical Computing Systems*, 23(1):55–84, 2002.
- [ATS06] Aaron D. Ames, Paulo Tabuada, and Shankar Sastry. On the stability of zeno equilibria. In Joao P. Hespanha and Ashish Tiwari, editors, *Hybrid Systems: Computation and Control 2006*, volume 3927 of *Lecture Notes in Computer Science*. Springer-Verlag, Santa Barbara, CA, USA, 2006.
- [BA02] G. Buttazzo and L. Abeni. Adaptive workload management through elastic scheduling. *Real-Time Systems*, 23(1):7–24, 2002.
- [BL00] R.W. Brockett and D. Liberzon. Quantized feedback stabilization of linear systems. *IEEE Transactions on Automatic Control*, 45(7):1279–1289, 2000.
- [BPZ02] M. Branicky, S. Phillips, and W. Zhang. Scheduling and feedback co-design for networked control systems. In *Proceedings of the 41st IEEE Conference on Decision and Control*, Las Vegas, NV, December 2002.
- [CE00] Anton Cervin and Johan Eker. Feedback scheduling of control tasks. In *Proceedings of the 39th IEEE Conference on Decision and Control*, Sydney, Australia, 2000.
- [CEBA02] A. Cervin, J. Eker, B. Bernhardsson, and K.-E. Arzen. Feedbackfeedforward scheduling of control tasks. *Real-Time Systems*, 23(1):25–53, 2002.
- [CLS03] R. Chandra, X. Liu, and L. Sha. On the scheduling of flexible and reliable real-time control systems. *Real-Time Systems*, 24(2):153–169, 2003.
- [EM01] N. Elia and S.K. Mitter. Stabilization of linear systems with limited information. *IEEE Transactions on Automatic Control*, 46(9):1384–1400, 2001.

- [FMPY06] E. Fersman, L. Mokrushin, P. Pettersson, and . Yi. Schedulability analysis of fixed priority systems using timed automata. *Theoretical Computer Science*, 354(2):301 – 317, 2006.
- [Kha96] Hassan K. Khalil. *Nonlinear Systems*. Prentice Hall, New Jersey, second edition, 1996.
- [LHLQ06] D. Liu, S. Hu, M. Lemmon, and L. Qiang. Firm real-time system scheduling based on a novel qos constraint. *IEEE Transactions on Computers*, 55(3):320–333, 2006.
- [Lib03] Daniel. Liberzon. Hybrid feedback stabilization of systems with quantized signals. *Automatica*, 39(9):1543–1554, 2003.
- [LNT02] D.S. Laila, D. Netic, and A.R. Teel. Open and closed loop dissipation inequalities under sampling and controller emulation. *European Journal of Control*, 18:109–125, 2002.
- [MA04] L.A. Montestruque and P.J. Antsaklis. Stability of model-based networked control systems with time-varying transmission times. *IEEE Transactions on Automatic Control*, 49(9):1562–1572, 2004.
- [NE00] Girish N. Nair and Robin J. Evans. Stabilization with data-rate-limited feedback: Tightest attainable bounds. *Systems and Control Letters*, 41(1):49–56, 2000.
- [NT04] D. Netic and A.R. Teel. Input to state stability of networked control systems. *Automatica*, 40:2121–2128, 2004.
- [PPBSV05] L. Palopoli, C. Pinello, A. Bicchi, and A. Sangiovanni-Vincentelli. Maximizing the stability radius of a set of systems under real-time scheduling constraints. *IEEE Transactions on Automatic Control*, 50(11):1790–1795, 2005.
- [PPV+02] Luigi Palopoli, Claudio Pinello, Alberto Sangiovanni Vincentelli, Laurent Elghaoui, and Antonio Bicchi. Synthesis of robust control systems under resource constraints. In Claire Tomlin and Mark R. Greenstreet, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science, pages 337–350. Springer-Verlag, 2002.
- [SLCB00] Lui Sha, Xue Liu, Marco Caccamo, and Giorgio Buttazzo. Online control optimization using load driven scheduling. In *Proceedings of the 39th IEEE Conference on Decision and Control*, Sydney, Australia, 2000.
- [Son04] E.D. Sontag. Input to State Stability: Basic concepts and results. In *CIME Summer Course on Nonlinear and Optimal Control Theory*, Lecture Notes in Mathematics, pages 462–488. Springer-Verlag, 2004.
- [TW06] Paulo Tabuada and Xiaofeng Wang. Preliminary results on state-triggered scheduling of stabilizing control tasks. In *Proceedings of the 45th IEEE Conference on Decision and Control*, San Diego, CA, 2006. Submitted.
- [YTS02] J.K. Yook, D.M. Tilbury, and N.R. Soparkar. Trading computation for bandwidth: Reducing communication in distributed control systems using state estimators. *IEEE Transactions on Control Systems Technology*, 10(4):503–518, July 2002.

UCLA ELECTRICAL ENGINEERING DEPARTMENT, 66-147F ENGINEERING IV BUILDING, LOS ANGELES, CA 90095-1594

E-mail address: tabuada@ee.ucla.edu