12th IFAC/IEEE Workshop on Discrete Event Systems
Cachan, France. May 14-16, 2014

Performance Regulation via Integral Control in a Class of Stochastic Discrete Event Dynamic Systems

C. Seatzu * **Y. Wardi** **

* *Department of Electrical and Electronic Engineering, University of Cagliari, Italy (e-mail: seatzu @ diee.unica.it).*
** *School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, USA (e-mail: ywardi@ece.gatech.edu).*

Abstract: This paper presents a performance-regulation method for a class of Discrete Event Dynamic Systems (DEDS). The main idea is to use an integral controller with a variable gain, adaptively computed so as to guarantee effective output-tracking of a given reference value. The computation of the gain is based on the Infinitesimal Perturbation Analysis (IPA) gradient of the plant function with respect to the control variable, and the resultant tracking can be quite robust with respect to modeling inaccuracies and gradient-estimation errors. The proposed technique is tested on two examples concerning the regulation of the loss rate in a queue, and of inventory levels in a manufacturing-system represented by a Petri net. The results suggest its potential efficacy for a broader class of DEDS.

Keywords: Infinitesimal perturbation analysis, stochastic hybrid systems, Petri nets, performance regulation.

1. INTRODUCTION

Infinitesimal Perturbation Analysis (IPA) has been extensively applied to compute sample-path gradients of performance functions, defined on the state space of Discrete Event Dynamic Systems (DEDS), as functions of continuous parameters. Its main use so far has been in optimization of expected-value performance functions in conjunction with stochastic approximation algorithms. For extensive presentations of the IPA technique and its applications, please see Ho (1991); Glasserman (1991); Cassandras and Lafortune (1999) and references therein.

Recently IPA has been extended from DEDS to a class of stochastic hybrid systems, based on the Stochastic Flow Model (SFM) framework (see Cassandras et al. (2010); Wardi et al. (2010); Yao and Cassandras (2013); Wardi and Cassandras (2013) for recent surveys). The SFM paradigm, comprising a generalization of fluid queues, has several inherent features rendering it especially suitable to the application of IPA. In particular, IPA gradients in the setting of SFM are statistically unbiased in a far-larger class of systems than those in the setting of DEDS, and they often admit simple, model-free formulas and algorithms. Furthermore, in situations where an SFM acts as an abstraction of a DEDS, the SFM-based IPA gradients can provide sensitivity estimates of expected-value performance functions defined on the DEDS, which DEDS-related IPA gradients fail to yield. An additional feature of IPA in the SFM setting is that convergence to local minima of stochastic approximation algorithms exhibit considerable robustness to gradient-estimation errors; see Cassandras et al. (2002); Sun et al. (2004); Cassandras (2006); Panayiotou and Cassandras (2006); Yao and Cassandras (2013) for simulation results, and Wardi and Cassandras (2013a) for an initial analysis.

This paper leverages on the aforementioned results to propose a role for IPA in applications beyond optimization, namely in performance regulation of DEDS. In the scenario that we examine, a given reference performance value has to be maintained in the face of unpredictable factors such as system-modeling inaccuracies, input variations, changes in the system's characteristics, and the effects of noise. This can be achieved with a suitable feedback control law having an integrator in the loop. However, an integral control may result in inadequate stability margins or large output oscillations. Moreover, a controller with a fixed gain may yield inadequate performance under changing system's characteristics. To get around these difficulties we propose an integrator with a time-varying gain, adaptively computed in a way that (for a class of systems) broadens the stability margins and yields a faster tracking of the reference input than fixed-gain integral controllers.

To set the stage for our problem, consider a stochastic timed DEDS defined over a probability space $(\Omega, \mathcal{F}, P)$, suppose that its state variable evolves in a way that depends on a continuous parameter $\theta \in R$, and hence denoted by $x(\theta, t)$. We view $x(\theta, t)$ as a realization of a random process defined on the underlying probability

⋆ Research of the second author was supported in part by the NSF under Grant CNS-1239225.

space, and dependent on a sample $\omega \in \Omega$. Let us partition the positive-time axis into a sequence of contiguous intervals, or cycles, C_n, $n = 1, 2, \ldots$, and denote their respective lengths by T_n, $n = 1, 2, \ldots$. Thus, $C_1 = [0, T_1)$, $C_2 = [T_1, T_1 + T_2)$, etc. Furthermore, let $L_n(\theta)$ be a random function dependent on the state evolution during the nth cycle C_n. Given a performance reference r, we are concerned with regulating the process $L_n(\theta)$, $n = 1, 2, \ldots$, so that it approaches the value of r.

For example, consider a stable GI/G/1 queue whose service-time processes depend on a parameter θ. Given $T > 0$, let $L_n(\theta)$ be the mean delay of jobs that arrive during the cycle $C_n := [(n-1)T, nT)$. The queue's processes can be reset in various ways (or none at all) at the start of C_n, but that is not the point. We are concerned with regulating the functions L_n to a given reference value.

Generally, we consider successive realizations of the functions $L_n(\theta)$ as the output sequence of a nonlinear, time-varying system whose input is θ. Time is discrete and indicated by the counter n, and the time-variablity is due to the sample path and the boundary condition (state) at the starting time of C_n. Furthermore, notwithstanding the fact that the state evolution during C_n is inherently dynamic, we can view the system as memoryless by focusing merely on its input–output (θ–$L_n(\theta)$) relations. This setting is natural for the forthcoming discussion of the integral controller that we consider. We will set the gain of the integrator during C_n to the inverse of the IPA derivative $L'_{n-1}(\theta)$, computed during the previous cycle, C_{n-1}. As will be explained in the sequel, the control law defined in this way yields effective regulation.

The rest of the paper is structured as follows. Section 2 describes the control system in an abstract setting, Section 3 provides a queueing example, and Section 4 presents a simulation example of balancing inventory and backorder in a Petri-net model of a manufacturing system. Section 5 concludes the paper and discusses potential extensions of the results derived therein.

2. REGULATION FRAMEWORK

Consider the discrete-time feedback system shown in Figure 1, where the counter $n = 1, 2, \ldots$, represents time. Both plant and controller are assumed to be single-input-single-output subsystems so that all the signals indicated in the figure are one-dimensional. Note that we use the unusual notation θ_n for the input to the plant, but it is common in the setting of IPA, which will be used to compute the sample gradients of performance functions with respect to θ_n.

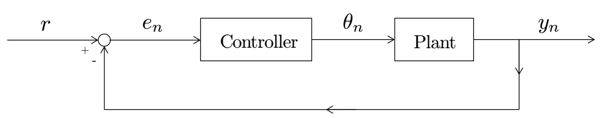

Fig. 1. Basic regulation system

Suppose that the plant is a nonlinear, memoryless, time-varying system represented by the functional relation $y_n = L_n(\theta_n)$, where the function $L_n : R \to R$ is assumed to be continuous in θ. The purpose of the feedback system is to have the output signal y_n asymptotically track a given reference value r.

To achieve such tracking, it is natural to have the controller contain an integrator. In its simplest form, an integral controller is a linear, time-invariant system with the transfer function

$$G_c(z) \;=\; A\frac{z^{-1}}{1 - z^{-1}}, \tag{1}$$

for a given gain $A > 0$, whose time-domain realization is defined by the equation $\theta_n = \theta_{n-1} + Ae_{n-1}$. However, a fixed-gain system may have the following two drawbacks. First, generally an integrator may result in a limited stability margin, and second, it may be impossible to determine a single gain which is adequate for every possible variation in the plant's characteristics. For these reasons we explore a variable gain, A_n, $n = 1, 2, \ldots$, so that the closed-loop system is defined by the following equations:

$$\theta_n = \theta_{n-1} + A_n e_{n-1}, \tag{2}$$

$$y_n = L_n(\theta_n), \tag{3}$$

and

$$e_n = r - y_n. \tag{4}$$

The gains A_n are computed adaptively according to the action of the system on the previous control variable, θ_{n-1}. An effective choice for A_n (for reasons that will become clear shortly) is $A_n = \big(L'_n(\theta_{n-1})\big)^{-1}$, where "prime" denotes derivative with respect to θ. However, this computation may not be exact, and only an approximation could be obtained. Thus, defining $K_n := \big(L'_n(\theta_{n-1})\big)^{-1}$, A_n has the form

$$A_n = K_n + \Delta K_n. \tag{5}$$

Note that the computation of the last four equations is recursive if it is made in the order (5)-(2)-(3)-(4).

Consider now the simple scenario where the plant is time invariant and the computation of K_n is exact, namely $L_n(\theta) = L(\theta)$ is independent of n, and $A_n = K_n$ in Equation (5). Then it can be seen that the system implements Newton's method for solving the equation $L(\theta) = r$, for which there are well-known results guaranteeing that $\lim_{n\to\infty} e_n = 0$ (and hence $\lim_{n\to\infty} y_n = r$). This limit also is satisfied under the time-invariance assumption with inexact computation of K_n in (5) as long as the relative error is less than γ for a $\gamma \in (0, 1)$, namely, $|\Delta K_n| \le \gamma |K_n|$ for every θ_{n-1}. Going a step further, suppose now that the plant is time varying in the sense that the functions $L_n(\theta)$ depend on n. In that case we cannot expect the error signal to converge to 0, but rather, to within a tolerance about 0 whose magnitude is a measure of the system's variability. In fact, Almoosa et al. (2012) showed that (under certain assumptions) for every $\epsilon > 0$ there exists $\delta > 0$ such that, if

$$|L_n(\theta_{n-1}) - L_{n-1}(\theta_{n-1})| < \delta \tag{6}$$

for all n, then

$$\limsup_{n\to\infty} |e_n| < \epsilon. \tag{7}$$

An important practical issue in a given control application is to compute the gain A_n in real time. This was addressed in Almoosa et al. (2012,a) for regulating power