Automatica 46 (2010) 528–536

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

On-line control of the threshold policy parameter for multiclass systems☆

Michael M. Markou, Christos G. Panayiotou *

KIOS Research Center for Intelligent Systems and Networks and Department of Electrical and Computer Engineering, University of Cyprus, Nicosia, Cyprus

ARTICLE INFO

Article history:
Received 21 May 2007
Received in revised form
11 June 2009
Accepted 20 November 2009
Available online 4 January 2010

Keywords:
Stochastic fluid models
Infinitesimal perturbation analysis
Control
Optimization
Quality of service
Buffer management
Discrete event systems
Queueing systems
Communication systems

ABSTRACT

Motivated by the problem of Quality of Service (QoS) provisioning, this paper investigates a system with two classes of traffic and addresses the problem of dynamically controlling the buffer threshold used with the threshold policy, such that an objective function is optimized. For the analysis, a Stochastic Fluid Model (SFM) is used to derive unbiased sensitivity estimators of various metrics (workload, loss, throughput and packet expiration) with respect to the buffer threshold. These estimators are computed using information obtained from the sample path of the "real" discrete event system. Thus, one can use these estimators together with stochastic approximation techniques in order to maintain the system at a near-optimum point despite any changes in the incoming traffic or the transmission capacity of the network.

1. Introduction

In this paper we consider a node that serves two classes of traffic, with different Quality of Service (QoS) requirements. An example is a network that is required to support multimedia applications with real-time data such as MPEG video streams. In the MPEG protocol, three different types of picture (video frame) are used for video coding: I frames, which are complete images, and P and B frames, which have additional information that improves the image quality. Dropping or delaying an I frame has a significantly greater impact on the quality of the received video than dropping or delaying a B or P frame. As a result, an application has an incentive to mark its I frames as high priority and its B and P frames as low priority traffic; if any node experiences congestion, it will start dropping the low priority packets (B and P frames) and thus I frames have a better chance to get through. In this paper, we assume that the source will voluntarily mark its packets; however, other marking approaches are also possible (see for example Heinanen & Guérin, 1999).

In this paper, we assume that the buffer of any intermediate node (router) uses the Threshold Policy (TP) to accept or reject packets; the TP is described in detail in Section 2. Since in multimedia applications it is often required that the order of the packets received is preserved, only First-In First-Out (FIFO) policies are considered. Among this family of policies we adopted the TP, since it better protects high priority packets than other FIFO policies (e.g., pushout and limited RED policies) and it is very easy to implement (Cidon, Guérin, & Khamisy, 1994). In addition, we assume that packets have a time-to-live (TTL) field and thus a node does not forward any expired packets.

In our approach, we adopt a Stochastic Fluid Model (SFM) and derive Infinitesimal Perturbation Analysis (IPA) estimators Cassandras and Lafortune (1999) (gradient estimators) of various performance metrics of interest. These SFM-based estimators are evaluated based on data observed from the "real" Discrete Event System (DES) and are then used together with stochastic approximation techniques Kushner and Yin (1997) to perform on-line control of the threshold parameter of the TP. We emphasize that we use SFMs only for *control and optimization* rather than performance evaluation, and various studies have shown that such approximate models are able to effectively determine the optimal solutions; for example, see Cassandras, Wardi, Melamed, Sun, and Panayiotou (2002) and references therein. Motivated by the approaches in

☆ Supported in part by the Cyprus Research Promotion Foundation under contract ENIΣX/0504/21. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Giuseppe De Nicolao under the direction of Editor Ian R. Petersen.

* Corresponding address: University of Cyprus, KIOS Research Center for Intelligent Systems and Networks and Department of Electrical and Computer Engineering, 75 Kallipoleos St., PO Box 20537, 1678 Nicosia, Cyprus. Tel.: +357 22 892298; fax: +357 22 892260.

E-mail addresses: markou@ucy.ac.cy (M.M. Markou), christosp@ucy.ac.cy (C.G. Panayiotou).

0005-1098/$ – see front matter
doi:10.1016/j.automatica.2009.12.002

Panayiotou and Cassandras (2001) and Panayiotou, Cassandras, Sun, and Wardi (2004), our goal is to develop an algorithm to *dynamically* control the buffer of network nodes that support multiple classes of traffic with different QoS requirements. An important advantage of the proposed methodology is that the gradient estimation process does not require any knowledge of the system's underlying stochastic processes; in other words, it is model free.

The contribution of this paper is that it derives the IPA sensitivity estimates of certain performance measures of interest with respect to the threshold parameter for a network buffer that involves two priority classes of traffic, and which limits network traffic by discarding expired packets. A similar problem has been investigated in Panayiotou et al. (2004) and Cassandras, Sun, Panayiotou, and Wardi (2003), but only for nodes that forward all packets. We point out that the introduction of packet expiration introduces a significant complication when comparing the model of this paper and the models in Panayiotou et al. (2004) and Cassandras et al. (2003). Specifically, in this paper the system dynamics depend on the state of the system $x(\cdot;\cdot)$, while Panayiotou et al. (2004) and Cassandras et al. (2003) assume no such dependency. Another important distinction is that the derived sensitivities are no longer piecewise constant functions. The authors of Yu and Cassandras (2003) and Markou and Panayiotou (2005) consider a problem of deriving IPA estimates for systems where the system dynamics depend on the state of the system; however, they only consider the single class case. Finally, this paper extends the results of Markou and Panayiotou (2006) by providing all lemma and theorem proofs as well as additional simulation results.

2. System model

We investigate a single node and assume that each source marks its packets as high priority (HP) or low priority (LP). HP packets are always accepted. On the other hand, LP packets are accepted only if the state of the buffer is below a threshold. All accepted packets are served using a FIFO policy, though some packets may be removed from the buffer and discarded before service if they violate the predefined delay QoS requirements (e.g., they have expired). It is worth mentioning that this system is a Discrete Event System (DES), in the sense that its operation is defined over a set of discrete events (e.g., packet arrival/departure, packet expiration). Moreover, the buffer state (the number of packets in the buffer) can take only non-negative integer values.

2.1. The stochastic fluid modeling framework

Fig. 1 shows the SFM queuing equivalent of the node with two distinct traffic streams of HP and LP packets (Cassandras et al., 2003). In this context, we assume that $\alpha_h(t)$ and $\alpha_l(t)$ correspond to the time-varying inflow processes of HP and LP traffic, respectively. The process $\beta(t)$ corresponds to the maximum service rate (successful transmission rate). The buffer is assumed to have infinite capacity and admits all HP flows; thus no HP fluid will be dropped due to buffer overflow. However, HP fluid can be dropped due to expiration. On the other hand, LP flows are accepted only when the buffer content $x(t;\theta)$ (queue length) is below a threshold θ; otherwise, the LP fluid is dropped. Using the SFM framework, the TP works as follows:

- If $0 \le x(t;\theta) < \theta$, all incoming flows are accepted.
- If $x(t;\theta) = \theta$, HP and *some* LP flows are accepted.[1]
- If $x(t;\theta) > \theta$, only HP flows are accepted.

[1] Note that in a discrete event modeling framework if an LP packet arrives when $x(t;\theta) = \theta$ it will always be dropped.

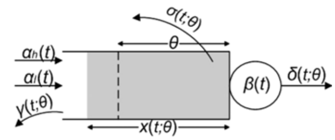

Fig. 1. The SFM of a node's buffer.

Based on the input processes $\alpha_h(t)$, $\alpha_l(t)$ and $\beta(t)$ and the threshold policy (TP), the following processes are derived: the buffer workload $x(t;\theta)$, the node's outflow $\delta(t;\theta)$, the LP fluid overflow $\gamma(t;\theta)$ and the fluid expiration $\sigma(x;t;\theta)$ due to unfulfilled QoS requirements. For the remainder of this paper, the buffer threshold θ is the control parameter of interest. Next, we make the following assumption regarding the input processes.

Assumption 1. With probability 1, $\alpha_h(t) < \infty, \alpha_l(t) < \infty$ and $\beta(t) < \infty$ are piecewise constant functions.

Assumption 1 can be extended to arbitrary piecewise differentiable functions but the piecewise constant assumption is used for simplicity. Finally, all processes evolve in a given time horizon $[0, T]$ for some fixed $0 < T < \infty$.

2.2. System dynamics

To ease the notation we also define $A(t) = \alpha_h(t) + \alpha_l(t) - \beta(t)$ and $B(t) = \alpha_h(t) - \beta(t)$. The state dynamics of $x(t;\theta)$ are given by

$$\frac{\mathrm{d}x(t;\theta)}{\mathrm{d}t^+} = \begin{cases} 0, & \text{if } x(t;\theta) = 0 \text{ and } A(t) \le 0 \\ 0, & \text{if } x(t;\theta) = \theta \text{ and } A(t) - \sigma(x;t;\theta) > 0 \\ & \text{and } B(t) - \sigma(x;t;\theta) < 0 \\ B(t) - \sigma(x;t;\theta), & \text{if } x(t;\theta) > \theta \\ A(t) - \sigma(x;t;\theta), & \text{otherwise.} \end{cases} \tag{1}$$

It is worth pointing out that the system dynamics differ from those in Cassandras et al. (2003) since they also depend on $\sigma(x;t;\theta)$ (which is a function of $x(t;\theta)$). In the real system, the packet expiration will be determined by a time stamp on each packet and will depend on the application's requirements, the number of hops a packet has to travel and the network conditions. For the fluid model, we assume that $\sigma(x;t;\theta) = c \cdot x(t;\theta)$. (Note that this is a good approximation when the packet arrivals are Poisson and the expiration deadline of each packet is exponentially distributed. In this case, since any packet can expire irrespective of its place in the queue, the expiration process is equivalent to the birth–death process of an $M/M/m$ system and the rate which takes the system from x to $x-1$ is μx, where μ is the service rate.)

2.3. Notation and preliminaries

A typical sample path can be decomposed into two types of alternating intervals: *empty periods* (EPs), during which the system is empty ($x(t;\theta) = 0$) and *non-empty periods* (NEPs), during which the system is non-empty ($x(t;\theta) > 0$). Fig. 2 shows a typical NEP.

On the system sample path we identify the events where (i) the buffer becomes or ceases to be empty, (ii) $x(t;\theta)$ hits the threshold θ (either from above or below) and (iii) $x(t;\theta)$ becomes (and stays) equal to θ or ceases being equal to θ. These events are classified as *exogenous* in the sense that they do not depend on the parameters of the system, or *endogenous* otherwise. For example, the event buffer ceases to be empty (point $v_0 = \xi$ in Fig. 2) is exogenous because it is caused by a sign change of $A(t)$. Similarly, the event buffer ceases to be equal to θ (points v_2, v_6, and v_8) is also exogenous. All other events are endogenous. Endogenous and exogenous events are important when determining the event time derivatives. Since exogenous events do not depend on the parameters of the system, their event time derivative with respect to θ, $\frac{\mathrm{d}v_i}{\mathrm{d}\theta} = 0$. On the other hand, in general, for endogenous events $\frac{\mathrm{d}v_i}{\mathrm{d}\theta} \neq 0$.