Automatica 48 (2012) 2801–2811

Contents lists available at SciVerse ScienceDirect

Automatica

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

Exploiting packet size in uncertain nonlinear networked control systems[☆]

Luca Greco[a,1], Antoine Chaillet[a], Antonio Bicchi[b]

[a] *L2S - EECI - University of Paris Sud 11 - Supélec. 3 rue Joliot-Curie. 91192 - Gif sur Yvette, France*

[b] *Interdep. Research Center "E. Piaggio", University of Pisa, Italy; IIT - Istituto Italiano di Tecnologia, Genova, Italy*

ARTICLE INFO

Article history:
Received 11 October 2010
Received in revised form
8 February 2012
Accepted 6 May 2012
Available online 23 September 2012

Keywords:
Networked control systems
Packet-switching networks
Model-based control
Time-varying delays

ABSTRACT

This paper addresses the problem of stabilizing uncertain nonlinear plants over a shared limited-bandwidth packet-switching network. While conventional control loops are designed to work with circuit-switching networks, where dedicated communication channels provide almost constant bit rate and delay, many networks, such as Ethernet, organize data transmission in packets, carrying larger amounts of information at less predictable rates. We adopt a model-based approach to remotely compute a predictive control signal on a suitable time horizon. By exploiting the inherent packets payload, this technique effectively reduces the bandwidth required to guarantee stability. Communications are assumed to be ruled by a rather general protocol model, which encompasses many protocols used in practice. An explicit bound on the combined effects of the maximum time between consecutive accesses to each node (MATI) and the transmission and processing delays (MAD), for both measurements and control packets, is provided as a function of the basin of attraction and the model accuracy. Our control strategy is shown to be robust with respect to sector-bounded uncertainties in the plant model. Sampling of the control signal is also explicitly taken into account. A case study is presented which enlightens the great improvements induced by the packet-based control strategy over methods that send a single control value in each packet.

1. Introduction

Industrial manufacturing is witnessing an ever more extensive use of communication networks to support automated scheduling, control and diagnostic activities (Moyne & Tilbury, 2007; Zurawski, 2005). The possibility offered by networks of replacing traditional point-to-point connections with more complex and dynamic schemes, opens unprecedented opportunities for factory control and management. Alongside allowing a pervasive adoption of decentralization and cooperation, networks convey many advantages in terms of flexibility, scalability and robustness. The adoption of a distributed networked architecture can induce a remarkable reduction of costs and delays for both installation and maintenance. These advantages justify the increasing interest in control over networks (see for instance Antsaklis & Baillieul, 2004, 2007; Bushnell, 2001; Di Benedetto, Johansson, Johansson, & Santucci, 2010; Wang & Liu, 2008).

In general terms, a Networked Control System (NCS) is a system in which sensors, actuators and controllers are spatially distributed and exchange information through a shared, digital, finite capacity channel. The use of the network as a communication medium and the distributed nature of the system make traditional control theory not always applicable. Issues such as quantization errors, data dropouts, variable transmission intervals, variable communication delays, and constrained access to the network, can no longer be ignored (Hespanha, Naghshtabrizi, & Xu, 2007). The NCS literature has separately addressed many of these problems, and sometimes the combinations thereof. An excellent discussion of the state-of-the-art is reported in Heemels, Teel, van de Wouw, and Nešić (2010). An essential aspect of NCS, not thoroughly analyzed in Heemels et al. (2010), is the packet-switching nature of many networks. As opposed to conventional control loops, which are designed to work with circuit-switching networks where dedicated communication channels provide almost constant bit rate and delay, networks such as Ethernet organize data transmission in packets, carrying larger amounts of information at less predictable rates.

[☆] This work was partially supported by the HYCON2 NoE, under grant agreement FP7-ICT-257462, and by the Contract IST 224428 (2008) (STREP) "CHAT — Control of Heterogeneous Automation Systems: Technologies for scalability, reconfigurability and security". The material in this paper was partially presented at the 18th IFAC World Congress, August 28–September 2, 2011, Milan, Italy. This paper was recommended for publication in revised form by Associate Editor Maurice Heemels under the direction of Editor Andrew R. Teel.

E-mail addresses: lgreco@ieee.org (L. Greco), chaillet@ieee.org (A. Chaillet), bicchi@ing.unipi.it (A. Bicchi).

[1] Tel.: +33 1 69 85 17 50; fax: +33 1 69 85 17 69.

0005-1098/$ – see front matter © 2012 Elsevier Ltd. All rights reserved.
doi:10.1016/j.automatica.2012.08.027

The organization of control information in data packets, which have relatively large transmission overhead, substantially alter the bandwidth/performance trade-off of traditional design. For instance, important data-rate theorems (Hespanha, Ortega, & Vasudevan, 2002; Nair & Evans, 2003, 2004) expressing a fundamental relationship between the degree of instability of a given physical system and the minimum bit rate required to stabilize it, do not account for the fact that data come in packets with a minimum size (e.g. 84 bytes in Ethernet). To simplify, transmitting a 16 bit record every millisecond requires as much bandwidth on average as sending a packet of 84 bytes every 42 ms; however, the implications on the effective sampling rate and feedback control performance are apparent. How to recover part of this performance is an objective of this study.

A second aspect inherent to packet-switching networks is transmission overhead. For instance, every Ethernet packet carries 38 bytes of headers and interframe separations, and useless information is necessarily padded into the payload to reach the minimum required packet length. As a consequence, transmitting a few bits per packet has essentially the same bandwidth cost as transmitting hundreds of them. A new, specific trade-off hence arises between packet rate and packet dimension for a given estimation/control task.

While the above aspects have been observed and described in the early literature on NCS (see e.g. the surveys Hespanha et al., 2007; Lian, Moyne, & Tilbury, 2001; Walsh & Ye, 2001), only recently have results appeared which address them explicitly in controller design. The goal can be succinctly described as to decrease the network utilization (in terms of bandwidth, or packets per unit of time) without compromising control performance. In Montestruque and Antsaklis (2004) a controller directly connected to the plant is considered and the number of measurement packets sent through the network is reduced by means of a state estimate provided by a model of the plant. In Bemporad (1998) the author pioneered the idea of sending feedforward control sequences, computed in advance on the basis of a model-based predictive (MBP) scheme, with the aim of compensating large delays in communication channels. A similar MBP scheme is exploited in Quevedo, Silva, and Goodwin (2007, 2008) and Quevedo and Nešić (2011) to counteract packet dropouts in the controller-to-plant channel. Compensation of delays and packet dropouts in nonlinear NCSs is the main concern also of the MBP controllers developed in Findeisen and Varutti (2009); Muñoz de la Peña and Christofides (2008). The following developments along these lines generalized the technique to address time-varying delays and transfer intervals (Polushin, Liu, & Lung, 2008), some robustness problems with respect to bounded perturbations (Pin & Parisini, 2009), as well as the constraints imposed by communication protocols on state measurement access (Chaillet & Bicchi, 2008).

In this paper we present a control strategy for packet-switching networks ensuring the stability of an uncertain nonlinear NCS affected by varying transmission intervals both on measurement and control sides, varying (and potentially large) delays, and constrained access to the network. Building upon our early results in Chaillet and Bicchi (2008), we adopt the feedforward approach to send in a packet not only the control value to be applied at a specific instant, but also a prediction of the control law valid on a given time horizon, so as to better exploit the payload. In the same spirit of other model-based approaches (e.g. Montestruque & Antsaklis, 2004; Pin & Parisini, 2009; Polushin et al., 2008; Quevedo et al., 2007), the control sequence is obtained by simulating an (imprecise) model of the closed-loop plant. The internal state of the model is asynchronously updated by means of the measurements of the plant state provided by sensors. Due to their spatial distribution, only portions of the model state can be updated in each instant. Therefore, we consider the constrained access to the network to be ruled by a protocol deciding which sensor can communicate at each instant. The large control packet, sent by the remote controller, is stored in an embedded memory on the plant side. Based on a local re-synchronization, made possible by a time-stamping of measurements, this strategy also allows us to compensate the effect of bounded communication delays in the control loop. Unlike the commonly assumed small-delays hypothesis (see for instance Heemels et al., 2010), we can compensate for delays larger than the transmission interval. We build our model upon the powerful hybrid formalism introduced in Nešić and Teel (2004), and we consider network imperfections affecting both sides of the control loop. We provide explicit bounds on the Maximum Allowable Delay (MAD (Heemels et al., 2010)) and on the Maximum Allowable Transfer Interval (MATI (Walsh, Ye, & Bushnell, 1999), i.e. the maximum duration between two successive communications) ensuring the exponential stability of the NCS over a prescribed basin of attraction. Finally, we clearly show, by means of a case study, the great improvement over existing methods that our feedforward control strategy induces on the aforementioned bounds.

A line of work close to ours is reported in Polushin et al. (2008), where the problem of stabilizing a nonlinear NCS with feedforward control sequences is addressed. Such sequences are computed by means of an approximate discrete-time plant model. The authors assume that the approximation algorithm is the only source of uncertainty in the model and that the inaccuracy of such a model can be reduced at will in order to achieve the desired MATI. In this paper, instead, we consider a robustness problem, where the plant uncertainty is given, and we provide a bound on the MATI in terms of the model inaccuracy (measured through its local Lipschitz constant).

2. Problem statement

Notation. Given a set $A \subset \mathbb{R}$ and $a \in A$, $A_{\geq a}$ denotes the set $\{s \in A \mid s \geq a\}$. Given a vector $x = (x_1, \ldots, x_n)^T \in \mathbb{R}^n$, $n \in \mathbb{N}_{\geq 1}$, $|x|$ denotes its Euclidean norm, i.e. $|x| \triangleq \left(\sum_{i=1}^n x_i^2\right)^{1/2}$. Given $R \geq 0$, B_R denotes the closed ball of radius R centered in zero: $B_R \triangleq \{x \in \mathbb{R}^n \mid |x| \leq R\}$. Given a locally essentially bounded signal $u : \mathbb{R}_{\geq 0} \to \mathbb{R}^n$, $\|u\|_{\mathcal{L}_\infty} \triangleq \operatorname{ess\,sup}_{t \geq 0} |u(t)|$. We use mod to denote the modulo operator, i.e. given $m, n \in \mathbb{N}$, $m \bmod n = p$ if and only if there exists $r \in \mathbb{N}$ such that $m = rn + p$ with $p < n$. We define the floor function $\lfloor \cdot \rfloor : \mathbb{R} \to \mathbb{Z}$ as $\lfloor x \rfloor \triangleq \max\{m \in \mathbb{Z} \mid m \leq x\}$.

2.1. Network model

We consider a NCS constituted of a remote controller receiving measurements from and sending commands to a physical plant through a shared communication channel (see Fig. 1). Control sequences are sent over the digital network as packets. An elementary embedded control device receives, decodes, and synchronizes these packets (see Buffer Synchronizer in Fig. 1) and applies control commands to the plant. Measurements are taken by physically distributed sensors and sent towards the controller as packets encoded with sufficient precision to neglect quantization effects. Sensors are assumed to be embedded with the plant and hence synchronized with it. Due to the distributed nature of the sensors, we also assume that the measurement part of the network is partitioned in ℓ nodes and only a *unique* node at a time can send its information (i.e. only partial knowledge of the plant state is available at each time instant). In other words, the state $x \in \mathbb{R}^n$, $n \in \mathbb{N}_{\geq 1}$, of the plant is decomposed as $x = (x_1^T, \ldots, x_\ell^T)^T$ with $x_i \in \mathbb{R}^{p_i}$ and $\sum_{i=1}^{\ell} p_i = n$.

We consider that measurements are taken and sent at instants $\{\tau_i^m\}_{i \in \mathbb{N}}$, and are received by the remote controller at instants