International Journal of Innovative
Computing, Information and Control
Volume **6**, Number **10**, October **2010**

 ISSN 1349-4198
pp. **4363–4377**

A NEW PERSPECTIVE PROPOSAL FOR PREEMPTIVE FEEDBACK SCHEDULING

Martina Maggio* and Alberto Leva

Dipartimento di Elettronica e Informazione
Politecnico di Milano
Piazza Leonardo da Vinci, 32 - 20133 Milano, Italy
{ maggio; leva }@elet.polimi.it
*PhD Student at the Dipartimento di Elettronica e Informazione

Received May 2009; revised September 2009

Abstract. *A novel way to formalise the problem of preemptive feedback scheduling is presented, based on discrete-time linear dynamic systems. Despite its extreme simplicity, the obtained formalism is suitable both for representing (and therefore dynamically analysing) existing scheduling methods, and for synthesizing new ones in a totally system-theoretical perspective. Also, by further specialising the class of systems employed, the modelling complexity can be scaled and tailored to the particular problem to be addressed. Examples are reported to explain and support the above claims, including some (simple) new scheduling policies.*
Keywords: Feedback scheduling, Preemptive scheduling, Discrete-time linear systems

1. **Introduction.** Feedback control of computing systems has been receiving much attention in recent years, see e.g. the book [6]. One of the main ideas behind such a huge research *corpus*, quoting again [6, p. xv], is that "by understanding the essential elements of control theory, computing practitioners can design systems that adapt in a more reliable manner".

Many attempts have been done to apply control techniques to computing systems by following the idea above, and a full review on the matter is impossible to give here. Applications range from QoS-oriented middleware [11] to flow control [17], from self managing software [4] to adaptive video streaming [14], and much more.

However, to the best of the authors' knowledge, one word in the statement quoted above has not yet been realised, and that word is "design". In fact, virtually the totality of the documented research starts from some computing system *already* designed, with limited (if any) control-theoretical foundation. Then, the said system is taken as "the plant" in the sense given to the term in the control domain, and (simplifying for brevity) some loops are closed around that plant to attain the desired specifications. Doing so is certainly a straightforward way to use control tools in computing systems, and the research presented here is in no sense a criticism to that *modus operandi*. However, strictly speaking, the mentioned *modus operandi* is not *designing* a computing system by means of the control theory.

Moreover, among the control application domains, the computing system one is unique. In fact, the "plant" and the "controller" are homogeneous (basically software) objects, making it possible and natural to design the overall system in a unitary manner right from the beginning. Also, computing systems are inherently immune from most of the problems encountered in other contexts. For example, measurement errors and state variables'