

Available online at www.sciencedirect.com

Computer Communications 27 (2004) 1412-1423



www.elsevier.com/locate/comcom

Analysis of deadline assignment methods in distributed real-time systems

Dana Marinca^{a,1}, Pascale Minet^{*,2}, Laurent George^b

^aProject Hipercon, INRIA, BP105, 78153 Le Chesnay Cedex, Rocquencourt, France ^bLIIA, University of Paris 12, 120 rue Paul Armangot, 94400 Vitry sur Seine, France

Received 16 September 2003; revised 10 May 2004; accepted 12 May 2004

Available online 9 June 2004

Abstract

The deadline assignment problem arises in distributed systems where each subtask composing a distributed task must receive a local deadline such that the task end-to-end deadline is met. It also arises in multi-hop networks where the maximum sojourn time of a flow on each visited node must be bounded by a local deadline that allows the flow end-to-end deadline to be met. We first formalize the problem and identify the cases where the choice of a deadline assignment method has a strong impact on system performances. We then propose two deadline assignment methods: Fair Laxity Distribution (FLD) and Unfair Laxity Distribution (ULD). Both assign local deadlines to the flow. These deadlines are based on the flow minimum sojourn time that can be guaranteed on each visited node. FLD and ULD differ in the laxity distribution: fair between the visited nodes for FLD, and proportional to the minimum guaranteed sojourn time for ULD. Performances of FLD and ULD are compared with those of classical methods such as fair assignment and assignment proportional to the workload. Moreover, performance evaluation shows that FLD for NP-EDF scheduling and ULD for FIFO scheduling are good approximations of an optimal algorithm in the context of a video-on-demand system.

© 2004 Elsevier B.V. All rights reserved.

Keywords: Deadline assignment; Distributed real-time system; QoS network; Admission control; Real-time scheduling; EDF; FIFO

1. Introduction

In this paper, we address the problem of satisfying an end-to-end real-time constraint associated with a flow visiting a set of nodes in a network subject to traffic changes. For the sake of simplicity, we use the term flow even if the results given in this paper can be applied to distributed tasks. Usually, the end-to-end constraint associated with a flow is defined in terms of its maximum end-toend response time, called end-to-end deadline. Two basic approaches exist to deal with an end-to-end deadline.

• The first one consists in checking that the worst case endto-end response time of any flow is less than or equal to its end-to-end deadline. The worst case end-to-end response time can be bounded by the sum of the local worst case response times. In this approach a high

0140-3664/\$ - see front matter © 2004 Elsevier B.V. All rights reserved. doi:10.1016/j.comcom.2004.05.006

response time on a node can be compensated by low response times on other nodes. The main drawback of this approach is that a flow configuration change on a node can lead to recomputing the worst case response times on all nodes.

• The second approach consists in assigning local deadlines on each visited node so that the sum of the local deadlines is equal to the end-to-end deadline of the flow considered. The benefit of this approach is that a change in local flow configuration does not affect the other nodes as long as the local deadlines are still met. The main drawback of this solution is the possible rejection of a feasible flow configuration that meets the end-to-end deadline but not all the local deadlines.

In this paper, we adopt the second approach. Indeed, even if local deadline assignment is more restrictive, it is more suitable in a dynamic system where configuration changes occur frequently. The deadline assignment problem is a well-known problem arising in distributed real-time systems (e.g. real-time distributed databases, Video-on-Demand systems, production management and resource planing in an industrial process) or in multi-hop networks

^{*} Corresponding author. Tel.: +33-1-39-63-52-33; fax: +33-1-39-63-55-66.

E-mail addresses: pascale.minet@inria.fr (P. Minet), dana.marinca@ free.fr (D. Marinca), george@univ-paris12.fr (L. George).

¹ This work was realized during her PhD thesis.

² http://www.hipercom.inria.fr/~minet.

supporting quality of service, (QoS): e.g. LAN, Internet, Mobile Ad-hoc NETwork (MANET).

Our main contribution is the proposal of two deadline assignment algorithms: Fair Laxity Distribution (FLD) and Unfair Laxity Distribution (ULD). These algorithms can be applied with any scheduling algorithm able to determine the minimum local acceptable deadline for a new flow. This property is formally defined in Section 4 (see property 1). Both First In First Out (FIFO) and Non-Preemptive Earliest Deadline First (NP-EDF) meet this property. FIFO is a simple widespread algorithm and NP-EDF has been proved optimal in the uniprocessor case when flow arrival times are not known a priori Ref. [1]. We recommend the use of FLD when FIFO is the local node scheduling and ULD for NP-EDF. Both outperform well-known existing solutions (fair assignment and load proportional assignment): they allow more flows to be admitted.

This paper is organized as follows: in Section 2, we give a formal definition of the deadline assignment problem. Section 3 is a brief state of the art of existing solutions for deadline assignment. In Section 4, we present two new algorithms: FLD and ULD. Section 5 is a performance evaluation of FLD and ULD, comparatively with two classical deadline assignment methods: fair assignment and load proportional assignment. This performance evaluation shows the benefit of FLD and ULD algorithms. We then compare FLD and ULD to an optimal greedy algorithm that computes all valid deadline assignments when FIFO or NP-EDF is used. The intrinsic complexity of the greedy algorithm makes it inappropriate for an on-line deadline assignment. Being optimal, it is used as a reference for the evaluation of other algorithms. Finally, we conclude in Section 6.

2. The problem

On a node, the local deadline of a flow is the sojourn time that this flow can be guaranteed. This local deadline is obtained by establishing the following trade-off:

a small local deadline for a flow on a node makes it easier to meet the flow end-to-end deadline, but more difficult to accept subsequent flows;

• a high local deadline for a flow on a node makes it easier to accept this flow on the node, but more difficult to meet the end-to-end deadline of the flow.

As shown by the performance evaluation reported in Section 5, the benefit of a deadline assignment algorithm is maximum when the end-to-end deadlines are neither too small (i.e. the number of accepted flows will be low with any algorithm) nor too large (i.e. the processor utilization factor becomes the only limiting factor and the local deadlines are always met). In its generalized form, the deadline assignment problem, illustrated by Fig. 1, can be formalized as follows:

Deadline assignment problem: Consider a network characterized by a directed graph G(V, L), where V denotes a set of nodes and L a set of links. Let v_1 and v_{n+1} be any two nodes of V, let $P(v_1, v_{n+1})$ be a path composed of $\{l_1, l_2, \dots, l_{n-1}, l_n\}$ links and let D be a positive constant. Each link $l_i \in L, i = 1, ..., n$, binding node v_i to node v_{i+1} where v_i and v_{i+1} belong to V, can provide several QoS levels. Each QoS level $Q_{i,j}$ is characterized by two metrics: a delay $D_{i,j}$, and a cost $c_{i,j}$. The problem is to find for each link $l_i, i = 1, ...n$, the QoS level $Q_{i,j}$ meeting $\sum_{i=1}^n D_{i,j} \le D$ and minimizing $\max_{i=1,\ldots,n}(c_{i,j})$.

Deadline assignment problem applied to a multimedia system: The deadline assignment problem arises when a new multimedia flow F_i requests its admission. The QoS levels, associated with link l_i between nodes v_i and v_{i+1} , characterized by $(D_{i,j}, c_{i,j})$ can be defined as follows:

- the delay D_{i,j} denotes the deadline of flow F_i on node v_i;
 the cost c_{i,j} = -N^{clients}_{v_i}, where N^{clients}_{v_i} denotes the maximum number of clients that can be accepted on node v_i , after the acceptance of flow F_i with the local deadline $D_{i,j}$ on node v_i .

Arbitration between different solutions: If for a given problem, several solutions exist providing the same maximum cost, we propose the following arbitration technique: we replace in each solution the cost of links providing the maximum cost by $-\infty$. We then select among these solutions the ones that minimize the maximum cost of a link in the path. And so on...until we get either a unique solution or several



Fig. 1. Different QoS levels on a given path.

دريافت فورى 🛶 متن كامل مقاله

- امکان دانلود نسخه تمام متن مقالات انگلیسی
 امکان دانلود نسخه ترجمه شده مقالات
 پذیرش سفارش ترجمه تخصصی
 امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
 امکان دانلود رایگان ۲ صفحه اول هر مقاله
 امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
 دانلود فوری مقاله پس از پرداخت آنلاین
 پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات
- ISIArticles مرجع مقالات تخصصی ایران