

# Optimized Deadline Assignment and Schedulability Analysis for Distributed Real-Time Systems with Local EDF Scheduling

Juan M. Rivas, J. Javier Gutiérrez, J. Carlos Palencia, and Michael González Harbour

*Computers and Real-Time Group, Universidad de Cantabria, 39005-Santander, SPAIN  
{rivasjm, gutierjj, palencij, mgh}@unican.es*

## Abstract<sup>1</sup>

*The assignment of scheduling parameters under the Earliest Deadline First (EDF) scheduling policy is trivial in single processor systems because deadlines are used directly. However, it is still difficult to find a feasible deadline assignment for EDF distributed systems when the utilization levels of the CPUs and communication networks are pushed near to their limits. Most distributed applications specify end-to-end deadlines for each transaction and there are no individual deadlines assigned to their tasks or messages. This paper presents a new heuristic algorithm, called HOS-DA (Heuristic Optimized Scheduling Deadline Assignment), for optimizing the assignment of deadlines to tasks and messages in distributed hard real-time systems. The algorithm is based on HOPA (Heuristic Optimized Priority Assignment), a previous method for the assignment of priorities in fixed priority distributed systems. The results of the proposed algorithm are compared with two other algorithms that exist for solving the same problem, and show that a utilization increase of up to 18% is possible. The paper also proposes a new schedulability analysis technique for EDF distributed systems with local scheduling deadlines.*

## 1. Introduction

Distributed real-time systems are increasingly important in today's embedded systems, since low-cost networking facilities allow the interconnection of multiple devices and their controllers into a single large system. The system's software is composed of concurrent tasks that are often statically allocated to processing nodes where each task may exchange messages with other tasks in the same processing node or in a different one.

Although fixed priority scheduling is the most popular on-line scheduling policy, usage of the Earliest deadline

first (EDF) policy is starting to get more attention in industrial environments, given its benefits in terms of increased resource usage. Analysis techniques are available to determine whether a given system, either single processor or distributed, will meet all of its timing requirements. EDF is now available in real-time languages like Ada 2005 [21] or RTSJ [23]. It is available in real-time operating systems such as SHaRK [18] or ERIKA [4], and has been implemented at the application level in OSEK/VDX embedded operating systems [2]. There are also real-time networks using EDF for scheduling messages, for instance in general purpose networks [3], or in the CAN Bus [14]. Given its maturity, it is expected that the number of industrial applications using EDF will increase in the near future.

One of the problems for scheduling distributed real-time systems is finding an assignment of scheduling parameters that leads to a feasible scheduling. This problem is fully solved for single processor systems. However, most timing requirements for distributed transactions are in the form of end-to-end deadlines, and finding scheduling parameters for the tasks and messages that constitute the transaction has no known optimum solution other than an intractable brute-force mechanism.

Different heuristics that can provide acceptable solutions to the assignment of scheduling parameters in a reasonable time have been studied for fixed priorities: simulated annealing or genetic algorithms can be used, as general-purpose optimization techniques; in [6] an algorithm called HOPA (Heuristic Optimized Priority Assignment), based on iteratively applying response time analysis (RTA), is shown to usually find better solutions than simulated annealing, in less time. There are also some algorithms for EDF that distribute end-to-end deadlines into individual deadlines for tasks and messages. The technique presented in [10] tries to reduce the number of missed deadlines in soft real-time systems. In [8][9] the assignment of deadlines is solved together with the allocation of tasks to processors; only a subset of the tasks are constrained by predetermined assignments to specific processors while the others can be allocated by the algorithm in order to make the system schedulable with the deadlines

---

1. This work has been funded in part by the Spanish Ministry of Science and Technology under grant number TIN2008-06766-C03-03 (RT-MODEL), and by the IST Programme of the European Commission under project FP6/2005/IST/5-034026 (FRESCOR). This work reflects only the author's views; the EU is not liable for any use that may be made of the information contained herein.

previously assigned. Other works, like the one proposed in [7], deal with strategies to analyze the schedulability of the distributed system based on pipelines, which is a restrictive model.

In [11], Liu proposes four basic strategies for assigning deadlines in EDF distributed systems. However these are not iterative algorithms incapable of improving on the initial assignment. In this paper we explore a new heuristic algorithm that tries to reuse some of the ideas of HOPA [6]. We call the new algorithm HOSDA (Heuristic Optimized Scheduling Deadline Assignment), and its goal is to find a feasible assignment of deadlines in a real-time distributed system by distributing end-to-end deadlines into intermediate deadlines, and by iterating over response time analysis to improve on the solution found.

In distributed EDF scheduling it is possible to find two kinds of schedulers: *global-deadline* schedulers have their deadlines referenced to the arrival of the event that releases the transaction, possibly in a different resource; *local-deadline* schedulers have deadlines referenced to the release time of each task in its own processing resource. Global-deadline schedulers require clock synchronization among all the processing resources involved, and the precision of the deadlines depends on the precision of the clock synchronization mechanism. While some systems do have precise clock synchronization this is not always the case. Local-deadline schedulers just use the local clock of each processor and are more general and easier to implement. The techniques proposed in [8][9][11] for assigning deadlines are based on local-deadline schedulers, and [10] used global deadlines.

It is interesting to notice that current response-time analysis techniques for EDF [20][13] are developed for global-deadline schedulers. Since local-deadline schedulers are more useful to general distributed platforms, in Section 3 we show how to adapt RTA techniques to support local-deadline schedulers. These RTA techniques are the basis for the HOSDA deadline distribution algorithm that we propose in this paper. As we will see, HOSDA outperforms the basic algorithms.

The paper is organized as follows. In Section 2 we provide a quick review of the model that we use for the distributed system, and we briefly describe the state of the art for the analysis techniques. Section 3 describes the proposed analysis algorithm for EDF with local scheduling deadlines. In Section 4 we show the details of the heuristic algorithm for optimizing the assignment of deadlines in distributed real-time systems. Section 5 evaluates the results of our algorithm by comparing them with those obtained by current deadline distribution strategies, showing that in most cases our algorithm finds better solutions

in a reasonable time. Finally, in Section 6 we give our conclusions.

## 2. System Model and Current Analysis Techniques

For our system model we will use the one described in [6], that is, we'll consider a task model with periodic distributed transactions. Every transaction instance (*transaction job*) is released by the arrival of an external event, and each *task job* is released once the previous task in the transaction finishes its execution, so a precedence relation inside the transaction is imposed. We denote the  $i$ -th transaction as  $\Gamma_i$ . We assume that all event sequences that arrive to the system are known in advance, and that tasks are statically assigned to processors (similarly messages to communication networks). Since the analysis of message traffic on the networks can be carried out using techniques that are used in the CPU (plus a small amount of non preemption), for simplifying purposes, we will treat messages as if they were tasks executing in a processor.

The timing requirements that we consider are end-to-end deadlines ( $D_i$ ) that start at the transaction job's period, and must be met by the final task in the transaction. Each task also has an associated local deadline,  $d_{ij}$ , which is relative to the release time of the task job, and a global deadline,  $D_{ij}$ , which is relative to the start of the transaction job's period.

Figure 1 shows an example of one of these transactions, with just three tasks, each executing in a different resource (two CPUs and one network in this case). The external event that releases the transaction is labeled  $e_i$ .




Figure 1. System model

For each task  $\tau_{ij}$  job we define its response time as the difference between its completion time and the instant at which the period of the transaction job started,  $t_n$ . The worst-case response time will be called  $R_{ij}$ .

Tasks after the first one are affected by the variability of the response times of the preceding tasks in the transaction. This implies that every task may have release jitter, which we will call  $J_{ij}$ . If we conservatively assume that best case execution times are zero, the best case response time of any task is zero and therefore the release time of the  $n$ -th job of

task  $\tau_{ij}$ ,  $j \neq 1$ , is in  $[t_n, t_n + J_{ij}]$  with  $J_{ij}$  being equal to the worst-case response time of the previous task,  $R_{i,j-1}$ .

We will assume that task synchronization is achieved using a hard real-time synchronization protocol (such as Baker's protocol [1]). The effects of this synchronization is modelled by a blocking term,  $B_{ij}$ .

This model described here is well-suited to represent a large number of real-time architectures that are found in practice. For example, if the system is using a client-server approach, each portion of a task that invokes a service from a remote server may be decomposed into the following activities: the activity before invoking the service, the message sent to the server, the server's activity, the reply message, and the activity after invoking the server.

There are different analysis techniques that can be used to analyze a system model like the one proposed here [20][13][15][16]. However, all these techniques rely on global-deadline schedulers, and for that reason, here we develop a new analysis to address local-deadline schedulers.

### 3. EDF Response-Time analysis for Local Deadlines

In this section we will extend the holistic analysis [20][22] for EDF distributed systems to support local deadline schedulers. Even though the new analysis is based on similar principles, the underlying model and the resulting equations are different.

In the holistic analysis we assume for each analysis step that every task is independent of the other tasks, even from those belonging to the same transaction. After each analysis step, the dependencies are captured into the jitter terms of each task.

Response time analysis is based on the creation of the longest busy period, found from a critical instant. The following theorem helps us to find the critical instant for a task in the context of the independent tasks assumption. It was proven by Spuri [19] for global deadlines but we adapt it to local deadlines.

**Theorem 1.** The worst-case response time of a task  $\tau_{ab}$  is found in a busy period in which each task  $\tau_{ij}$  in the same processor (different from  $\tau_{ab}$ ) is scheduled such that its first job that is inside the busy period is released at the beginning of that busy period, after having experienced its maximum jitter (i.e., the start of the corresponding task job's period was  $J_{ij}$  time units before the start of the busy period), and the remainder of the jobs are released with the minimum jitter that makes the job start inside the busy period.

**Proof:** The proof can be found in [17] and is omitted here because it is similar to the one provided by Spuri in [19].  $\square$

Note that, contrary to the other tasks, releasing the analyzed task at the start of the busy period may not lead to its worst-case response time. So, the critical instant for a task is found in a busy period that is started by the simultaneous release of all tasks except perhaps the one under analysis.

Under the conditions of theorem 1, the worst-case contribution of a task  $\tau_{ij}$  to a busy period of length  $l$  when the deadline of  $\tau_{ab}$  occurs at instant  $D$  is [17]:

$$W_{ij}(l, D) = \min(p_l, p_D) \cdot C_{ij} \quad (1)$$

where  $p_l$  is the number of releases of  $\tau_{ij}$  in the busy period:

$$p_l = \left\lceil \frac{l + J_{ij}}{T_i} \right\rceil \quad (2)$$

and  $p_D$  the number of releases with deadline before or at  $D$ :

$$p_D = \begin{cases} 0 & \text{if } D < d_{ij} \\ \left\lfloor \frac{J_{ij} + D - d_{ij}}{T_i} \right\rfloor + 1 & \text{otherwise} \end{cases} \quad (3)$$

Using this expression we can calculate the worst-case response time of task  $\tau_{ab}$ . Unfortunately, we don't know how to phase the release time of  $\tau_{ab}$  in relation to the busy period, but it is easy to see that the worst case situation must be found when the release time is placed at the beginning of the busy period, or at an instant such that the deadline of the analyzed job of  $\tau_{ab}$  coincides with the deadline of a task  $\tau_{ij}$ 's job. The set of instants,  $\Psi_{ij}$ , at which the deadline of  $\tau_{ab}$ 's job coincides with the deadlines of one of the task jobs in the busy period is:

$$\Psi_{ij} = \bigcup \{(p-1)T_i - J_{ij} + d_{ij}\} \quad \forall p = 1 \dots \left\lceil \frac{L + J_{ij}}{T_i} \right\rceil \quad (4)$$

where  $L$  corresponds to longest busy period, calculated as:

$$L = \sum_{\forall \tau_{ij}} \left\lceil \frac{L + J_{ij}}{T_i} \right\rceil \cdot C_{ij} \quad (5)$$

This set  $\Psi_{ij}$  must be augmented with the deadlines corresponding to task  $\tau_{ab}$

$$\Psi_{ab} = \bigcup \{(p-1)T_a + d_{ab}\} \quad \forall p = 1 \dots \left\lceil \frac{L}{T_a} \right\rceil \quad (6)$$

And so the full set of situations for which  $\tau_{ab}$  has to be analyzed corresponds to those releases whose deadline is in the set

$$\Psi = \Psi_{ij} \cup \Psi_{ab} \quad (7)$$

Each potential release time for  $\tau_{ab}$  is obtained by subtracting  $d_{ab}$  from each value in  $\Psi$ . Checking the response times under all these release times we can find the one that causes the worst-case response time of the task. Given that there may be several releases of  $\tau_{ab}$  in the busy period, we must analyze them all. For each value  $\psi \in \Psi$ , the completion time of release  $p$  of  $\tau_a$ ,  $w_{ab}^\psi(p)$ , can be calculated by adding the worst-case contribution of all tasks in the same processor, and the blocking term:

$$w_{ab}^\psi(p) = B_{ab} + pC_{ab} + \sum_{\forall ij \neq ab} W_{ij}(w_{ab}^\psi(p), \psi) \quad (8)$$

The worst-case response time is calculated by subtracting the start of the job's period from the resulting completion time:

$$R_{ab}^\psi(p) = w_{ab}^\psi(p) - (\psi - d_{ab} - J_{ab}) \quad (9)$$

For each value of  $p$ , we only need to check the values of  $\psi$  in one period, because if the release time corresponding to  $\psi$  was greater than the corresponding period, then we would be analyzing another job with a different value of  $p$ . This allows us to restrict the set of values to be checked:

$$\Psi^* = \{\psi_x \in \Psi \mid (p-1)T_a + d_{ab} \leq \psi_x < pT_a + d_{ab}\} \quad (10)$$

Finally, to calculate the worst-case response time of task  $\tau_{ab}$  we must determine the maximum response times within all the potential release times examined:

$$R_{ab} = \max(R_{ab}^\psi(p)) \quad \forall p = 1 \dots \left\lceil \frac{L}{T_a} \right\rceil, \quad \forall \psi \in \Psi^* \quad (11)$$

We can now feed these response times into the holistic analysis loop like in [20], obtaining new jitter values from the response times and repeating the analysis until a stable solution is obtained. Since the dependencies of response times on jitters are monotonically increasing, the algorithm is known to converge to the final solution, except when the utilization is close to 100% and in special cases that experience shows that are uncommon.

#### 4. Heuristic Algorithm for Optimized Deadline Assignment

Paper [6] proposes the HOPA algorithm that, unlike general-purpose optimization algorithms such as simulated annealing, uses knowledge of the factors that influence the

timing behavior to find an optimized solution to the priority assignment problem in fixed-priority distributed systems.

HOPA is based on the distribution of the global deadlines of each transaction among the different tasks that compose it. Once each task is assigned an artificial local deadline, deadline monotonic priorities are assigned in each processing resource and an analysis of the whole system is carried out. As a result of the analysis, new local deadlines are calculated. The iteration proceeds until a schedulable solution is found or some stopping condition is reached.

In this section we propose a new heuristic algorithm called HOSDA which is the adaptation of HOPA to systems scheduled by EDF. The objective is to check if the basic method to calculate the artificial deadlines that lead to schedulable solutions for fixed priorities can also be used to assign local scheduling deadlines for EDF distributed systems. From a high-level point of view, the algorithm is:

```

algorithm HOSDA is
begin
  assign initial scheduling deadlines;
  loop
    calculate worst-case response times;
    exit when some stopping criterion;
    calculate new scheduling deadlines;
  end loop;
end HOSDA;

```

The initial scheduling deadlines should be assigned in some way that preserves the global deadlines:

$$D_{ij} = \sum_{k \in pr_i(j)} d_{ik} \quad (12)$$

where:

$d_{ij}$  is the local scheduling deadline of task  $\tau_{ij}$ .

$pr_i(j)$  is the set of tasks preceding task  $\tau_{ij}$  in the transaction  $i$  to which it belongs, including itself.

$D_{ij}$  is the global deadline (intermediate or end-to-end) of task  $\tau_{ij}$ .

The proportional deadline assignment algorithm proposed in [11] can be used as the initial deadline assignment.

After all the local scheduling deadlines have been assigned, the system is analyzed using the technique described in Section 3. To avoid convergence problems or very long calculations that the holistic analysis may have when utilizations are very high, we have added a termination condition to the holistic analysis that makes it stop when the response time of a task exceeds the imposed deadline by a configurable factor. In this way we bound and shorten the analysis time and we assure that the deadline assignment algorithm can continue working. If the

solution is not schedulable, new scheduling deadlines are calculated by redistributing the global deadlines among the tasks of each transaction.

The redistribution of local deadlines uses the concept of “excess” of each task which, intuitively, and in the same way than in HOPA, measures the distance that separates each task from schedulability.

The original algorithm for fixed priorities [6] had two different definitions for the excess of a particular task  $\tau_{ij}$ , “excess of response time”, which is based on the difference between the local response time and the local deadline, and led to faster solutions; and “excess of compute time”, based on the calculation of the slack time, and led to more schedulable solutions. Currently we do not use the excess of computation time, since it requires large amounts of computational time, except for a very small range of examples.

In [6], 2 definitions for the Excess of Response Time were presented:

$$exc(\tau_{ij}) = \begin{cases} (R_{ij} - J_{ij} - d_{ij}) \frac{\Delta R_{ij}}{\Delta D_{ij}} & (1) \\ or & (13) \\ (R_{ij} - d_{ij}) \frac{\Delta R_{ij}}{\Delta D_{ij}} & (2) \end{cases}$$

Where:

$\Delta R_{ij}$  is the difference between the worst case global response times of the task before  $\tau_{ij}$  that has a global deadline and the task after  $\tau_{ij}$  that also has a global deadline.

$\Delta D_{ij}$  is the difference between the global deadlines of the tasks before and after  $\tau_{ij}$  that have a global deadline.

If there is no previous task with a global deadline, then 0 is used to find the difference in  $\Delta R_{ij}$  and  $\Delta D_{ij}$ . In the common case when there is only a single end-to-end deadline for the transaction all the  $\Delta R_{ij}$  terms of the same transaction have the same value related to the response time of the last task in the transaction,  $\Delta R_{ij} = R_{im_i}$ , and the same applies to  $\Delta D_{ij} = D_i$ .

Both definitions of excess of response time can lead separately to feasible deadline assignments. In our preliminary tests, usage of definition (2) led to slightly better results (around 3% more schedulable resource utilization achieved on average), so this is the recommended option to try first.

We use the same definitions as in [6] for the excess of each processing resource,  $exc(PR_k)$ ; the maximum excess of all the processing resources,  $Mex(PR)$ ; and the maximum of the excesses of all the tasks belonging to a particular transaction responding to external event  $e_i$ , which we call  $Mex(e_i)$ .

Given these definitions of excess times, we calculate the new scheduling deadline for each task  $d_{ij}(new)$  (14) as a function of the old scheduling deadline for that task, the excess for that task, the excess for the resource to which that task belongs, and the values of two empirical constants  $k_R$  and  $k_a$ . Like in HOPA, once the new local deadlines  $d_{ij}$  have been obtained for all the tasks in the transaction, we adjust them proportionally to make them fit into the global deadlines.

$$d_{ij}(new) = d_{ij}(old) \left(1 + \frac{exc(PR_k)}{k_R \cdot Mex(PR)}\right) \left(1 + \frac{exc(\tau_{ij})}{k_a \cdot Mex(e_i)}\right) \quad (14)$$

The  $k_R$  and  $k_a$  constants, like in HOPA, control the relative influence of the processing resource and task components, respectively, in the calculation of the new deadline. In order to empirically determine which values give the best results, extensive preliminary experiments were made. These experiments consisted of applying the HOSDA algorithm over a wide set of examples. These examples are comprised of systems with different number of processors [2,15], transaction lengths (number of tasks between 1 and the number of processors), periods (differences of up to 3 orders of magnitude), and end-to-end deadlines (between the period,  $T$ , and  $2 \cdot T \cdot \text{number of processors}$ ). For each example generated, different sets of values for  $k_R$  and  $k_a$  were applied. We found out that higher utilizations could be reached if the values for this parameters were between 1.0 and 3.0. Normally, in the HOSDA algorithm we start with values of  $k_R = k_a = 1.5$ , and then, after a given number of iterations, we change both values to 2.0, 2.5, 3.0, etc., until one stopping condition is reached. The stopping conditions are:

- A schedulable solution has been found.
- Two consecutive deadline assignments are identical (in which case the algorithm would continue providing the same solution).
- Since the algorithm is not guaranteed to converge to a solution, a Maximum Number of Iterations is defined. This number is configurable and, for a given size of the system, it sets a limit to the time that the user is willing to wait for obtaining a solution.

As in HOPA, after the algorithm finds a deadline assignment that makes the system schedulable, it is capable of finding a better optimized solution by executing more iterations.

## 5. Evaluation of the Heuristic Algorithm

In this section we compare the performance of the heuristic algorithm proposed with existing techniques for deadline distribution described by J. Liu in [11]. We choose

the best two of them: Proportional Deadline (PD) and Normalized Proportional Deadline (NPD).

In order to carry out the comparison between the HOSDA algorithm and the two selected techniques (PD, NPD), we have implemented and integrated the deadline assignment algorithms and the analysis technique into the MAST suite of tools [12]. This allows us to use both a precise model [5] to describe the systems under test and a uniform way to check the results. Since MAST is free software this also makes the implementation of the techniques available to the community.

A generator of examples has also been developed to test the execution of the algorithms over a wide spectrum of cases. This generator starts with a base system, characterized by the number of processors and transactions, and for each generated example, it assigns a random number of tasks and periods to each transaction. The WCET of each task is sequentially increased from a very low utilization until the utilization of the system reaches 100%. Systems generated by this tool are stored to be processed later.

In order to illustrate the performance of the HOSDA algorithm, we have chosen three base systems with different levels of complexity: a Small, Intermediate, and Big Size Example (SSE, ISE, and BSE respectively). Table 1 shows the number of processors and transactions for each base system. The number of tasks in each transaction varies randomly from 1 to the number of processors, and can be different in each example. We have generated 100 examples for each base system, and we apply five types of end-to-end deadlines to the transactions: deadline equal to period ( $D=T$ ); deadline proportional to the period and the number of processors ( $N$ ) traversed by the transaction multiplied by 0.5 ( $D = 0.5 \cdot N \cdot T$ ), one ( $D = N \cdot T$ ), or two ( $D = 2 \cdot N \cdot T$ ); and finally a random deadline between  $T$  and  $2 \cdot N \cdot T$ .

**Table 1. Base systems for the examples generator**

	SSE	ISE	BSE
<b>Number of Processors</b>	3	5	8
<b>Number of Transactions</b>	6	8	12

Table 2 shows the average maximum schedulable utilization reached by the algorithms for all the case studies. In this table we can see that, on average, HOSDA can schedule systems with higher utilizations than PD and NPD. However, there is a small amount of examples in which HOSDA, at its first attempt with the constants selected, has not been better than NPD. Of course, HOSDA is always better or equal than PD, as this is the starting assignment.

HOSDA has been applied to many more examples with very similar results and was able to increase the processor

**Table 2. Average Maximum Schedulable Utilization (%)**

		T	NT/2	NT	2NT	Random
<b>SSE</b>	<b>HOSDA</b>	60.2	60.8	92.7	98.8	93.9
	<b>PD</b>	59.0	57.7	84.3	97.6	88.9
	<b>NPD</b>	58.0	56.7	83.9	97.7	88.0
<b>ISE</b>	<b>HOSDA</b>	43.2	57.5	84.3	98.3	86.8
	<b>PD</b>	42.0	52.5	73.2	89.9	76.4
	<b>NPD</b>	41.4	52.4	72.8	90.0	76.4
<b>BSE</b>	<b>HOSDA</b>	34.6	59.5	75.8	88.1	76.1
	<b>PD</b>	33.8	53.1	67.5	78.0	67.5
	<b>NPD</b>	32.8	53.0	66.8	77.3	66.6

utilization by up to 18% in some of these examples. The best results are obtained in systems with deadlines larger than periods, usually in those in which deadlines are proportional to the period multiplied by the number of tasks in the transaction.

The CPU time spent by each algorithm to find a feasible solution was measured for all examples, and it was found that HOSDA took, on average, an order of magnitude more time to reach a solution than PD and NPD. For the PD and NPD algorithms this CPU time is composed of the time needed to make the assignment of the deadlines plus the time to run the schedulability analysis over the solution proposed. The time spent by HOSDA depends on the number of iterations needed to find the solution, which could lead to repeat the analysis potentially many times.

In some cases, the  $k_R$  and  $k_a$  constants used in the calculation of the excess (14) can be rebalanced to minimize the influence of the excess in the resource (e.g., by assigning values of  $k_R=10k_a$ , or  $k_R=100k_a$ ) or the excess in the transactions (e.g., by assigning values of  $k_a=10k_R$ , or  $k_a=100k_R$ ), which increase the possibility of reaching a feasible solution in some examples. As future work we plan to make a broad study of the impact that the  $k$  pair of values have in the schedulability of distributed systems, in order to be able to automate the selection and application of different sets of  $k_a$  and  $k_R$ , for the purpose of seeking better deadline assignments.

## 6. Conclusions

In this paper we propose a heuristic algorithm for assigning scheduling deadlines in distributed hard real-time systems. It is an adaptation to EDF of a previous algorithm called HOPA, designed for fixed priorities. We have shown that this method can find feasible optimized solutions in a reasonable amount of time, even in situations where the utilization of the resources is high and thus the number of solutions is small.

Since HOSDA is based on iteratively applying response time analysis and previous techniques for EDF scheduling

were based on global deadlines, in this paper we have extended the analysis techniques to support local deadlines, which are more useful in distributed environments because they do not require global clock synchronization.

We have compared HOSDA with two other reference algorithms previously proposed to address the same problem but which are not able to optimize. The results of the comparison are that our method finds feasible solutions in many cases where these other methods fail. This is especially noticeable when utilizations are high or when the deadlines of the transactions are larger than the periods of the events that release them.

The quality of the results obtained by HOSDA are quite similar to those obtained by HOPA. We are currently working on comparing whether HOPA for fixed priorities or HOSDA for EDF can find better solutions to schedule the same set of transactions. We are also working on the assignment of scheduling parameters in mixed systems (systems with EDF and FP scheduled processing resources). Finally, we also plan to test HOSDA with more advanced RTA techniques, such as those based on offsets, and to give HOSDA the capability to manage deadlines that cannot be changed or that must be less than or equal to a specific value (preassigned scheduling deadlines).

## REFERENCES

- [1] T. P. Baker, "Stack-based scheduling for realtime processes," *Real-Time Systems*, v.3 n.1, pp.67-99, March 1991
- [2] Claas Diederichs, Ulrich Margull, Frank Slomka, and Gerhard Wirrer, "An application-based EDF scheduler for OSEK/VDX," *Proceedings of the conference on Design, Automation and Test in Europe*, Munich (Germany), pp. 1045-1050, 2008.
- [3] M. Di Natale, and A. Meschi, "Scheduling Messages with Earliest Deadline Techniques," *Real-Time Systems*, 20, pp. 255-285, Kluwer Academic Publishers, 2001.
- [4] ERIKA (Embedded Real time Kernel Architecture) home page, <http://erika.sssup.it/>
- [5] M. González Harbour, J.J. Gutiérrez, J.C. Palencia, and J.M. Drake, "MAST: Modeling and Analysis Suite for Real Time Applications," *Proceedings of 13th Euromicro Conference on Real-Time Systems*, Delft (The Netherlands), pp. 125-134, 2001.
- [6] J.J. Gutiérrez, and M. González Harbour, "Optimized Priority Assignment for Tasks and Messages in Distributed Real-Time Systems," *Proceedings of 3rd Workshop on Parallel and Distributed Real-Time Systems*, Santa Barbara (California), pp. 124-132, 1995.
- [7] Praveen Jayachandran, Tarek Abdelzaher, "A Delay Composition Theorem for Real-Time Pipelines," *Proceedings of 19th Euromicro Conference on Real-Time Systems (ECRTS'07)*, pp. 29-38, IEEE, 2007.
- [8] Jan Jonsson, and Kang G. Shin, "Robust Adaptive Metric for Deadline Assignment in Distributed Hard Real-Time Systems," *Real-Time Systems*, 23, pp. 239-271, Kluwer Academic Publishers, 2002.
- [9] Jan Jonsson, and Kang G. Shin, "Deadline Assignment in Distributed Hard Real-Time Systems with Relaxed Locality Constraints," *Proc. of the IEEE Int'l Conf. on Distributed Computing Systems (ICDCS'97)*, Baltimore (Maryland) pp. 432-440, 1997.
- [10] B. Kao, and H. García-Molina, "Deadline Assignment in a Distributed Soft Real-Time System," *IEEE Transactions on Parallel and Distributed Systems*, Vol. 8, No. 12, pp. 1268-1274, 1997.
- [11] J. Liu, "Real-Time Systems," Prentice Hall, 2000.
- [12] MAST home page, <http://mast.unican.es/>
- [13] J.C. Palencia, and M. González Harbour, "Offset-Based Response Time Analysis of Distributed Systems Scheduled under EDF," *Proceedings of the 15th Euromicro Conference on Real-Time Systems, ECRTS, Porto (Portugal)*, 2003.
- [14] P. Pedreiras, and L. Almeida, "EDF Message Scheduling on Controller Area Network," *Computing & Control Engineering Journal*, 13(4), pp. 163-170, 2002.
- [15] R. Pellizzoni, and G. Lipari, "Improved Schedulability Analysis of Real-Time Transactions with Earliest Deadline Scheduling," *Proceedings of the 11th IEEE Real Time on Embedded Technology and Applications Symposium, RTAS*, pp. 66 - 75, 2005.
- [16] R. Pellizzoni, and G. Lipari, "Holistic analysis of asynchronous real-time transactions with earliest deadline scheduling," *Journal of Computer and System Sciences*, Volume 73, Issue 2, pp. 186-206, 2007.
- [17] Juan M. Rivas, et al. "Optimized Deadline Assignment for Tasks and Messages in Distributed Real-Time Systems". Technical report, University of Cantabria. URL: <http://www.ctr.unican.es/publications/jmr-jjg-jcp-mgh-2009a.pdf>
- [18] S.Ha.R.K. (Soft Hard Real-Time Kernel) home page, <http://shark.sssup.it/>
- [19] M. Spuri, "Analysis of Deadline Scheduled Real-Time Systems," RR-2772, INRIA, France, 1996.
- [20] M. Spuri, "Holistic Analysis for Deadline Scheduled Real-Time Distributed Systems," Tech. Rep. RR-2873, INRIA, France, April 1996
- [21] S. Tucker Taft, Robert A. Duff, Randall L. Brukardt, Erhard Ploedereder, and Pascal Leroy (Eds.), "Ada 2005 Reference Manual. Language and Standard Libraries. International Standard ISO/IEC 8652:1995(E) with Technical Corrigendum 1 and Amendment 1," LNCS 4348, Springer, 2006.
- [22] K. Tindell, and J. Clark, "Holistic Schedulability Analysis for Distributed Hard Real-Time Systems," *Microprocessing & Microprogramming*, Vol. 50, Nos.2-3, pp. 117-134, 1994.
- [23] RTSJ (Real-Time Specification for Java) home page, <http://www.rtsj.org/>