

Dual-Priority versus Background Scheduling: A Path-Wise Comparison

Bruno Gaujal, Nicolas Navet, Jörn Migge

► **To cite this version:**

Bruno Gaujal, Nicolas Navet, Jörn Migge. Dual-Priority versus Background Scheduling: A Path-Wise Comparison. [Research Report] RR-3734, INRIA. 1999. <inria-00072931>

HAL Id: inria-00072931

<https://hal.inria.fr/inria-00072931>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

***Dual-Priority versus Background Scheduling :
a Path-wise Comparison***

Bruno Gaujal , Nicolas Navet , Jörn Migge

No 3734

July 13, 1999

————— THÈME 1 —————



*Rapport
de recherche*

Dual-Priority versus Background Scheduling : a Path-wise Comparison

Bruno Gaujal ^{*} , Nicolas Navet [†] , Jörn Migge [‡]

Thème 1 — Réseaux et systèmes
Projet TRIO

Rapport de recherche n° 3734 — July 13, 1999 — 30 pages

Abstract: In this paper two well-known scheduling policies for Real Time Systems, namely Background Scheduling and Dual-Priority are compared in terms of response times for Soft Real Time traffic (SRT). It is proved that, when the SRT traffic is FIFO, the Dual-Priority policy always behaves better, in the preemptive case as well as in the non-preemptive case. The proof is based on a trajectorial method. As a complementary result, some non-FIFO examples where the Background Scheduling outperforms the Dual-Priority, are given.

Key-words: Real-Time Systems, Scheduling Algorithm, Local Area Network

(Résumé : tsvp)

^{*} Loria, 615, rue du Jardin botanique, 54602 Villers-les-Nancy, France. Email: Bruno.Gaujal@loria.fr

[†] Loria-Inpl, Ensem, 2 avenue de la Forêt de la Haye, 54516 Vandoeuvre, France. Email: Nicolas.Navet@loria.fr

[‡] Loria-Inpl, Ensem, 2 avenue de la Forêt de la Haye, 54516 Vandoeuvre, France. Email: Jorn.Migge@loria.fr

Unité de recherche INRIA Lorraine
Technopôle de Nancy-Brabois, Campus scientifique,
615 rue de Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY (France)
Téléphone : 03 83 59 30 30 - International : +33 3 3 83 59 30 30
Télécopie : 03 83 27 83 19 - International : +33 3 83 27 83 19
Antenne de Metz, technopôle de Metz 2000, 4 rue Marconi, 55070 METZ
Téléphone : 03 87 20 35 00 - International: +33 3 87 20 35 00
Télécopie : 03 87 76 39 77 - International : +33 3 87 76 39 77

Dual-Priority contre Background Scheduling : une comparaison trajectorielle

Résumé : Dans cet article, nous comparons, en termes de temps de réponse du trafic à contraintes souples (SRT), deux politiques d'ordonnancement classiques dans les systèmes temps réel, qui sont l'ordonnancement en arrière-plan ("Background Scheduling") et la politique dite à changement de priorité ("Dual-Priority"). Nous prouvons que lorsque le trafic SRT est FIFO, la politique à changement de priorité se comporte toujours mieux que l'ordonnancement en arrière-plan dans le cas préemptif comme dans le cas non-préemptif. La preuve utilise une méthode trajectorielle. Le résultat complémentaire est qu'il existe des exemples non-FIFO sur lesquels l'ordonnancement en arrière-plan est meilleur que la politique à changement de priorité.

Mots-clé : Systèmes temps réel, algorithmes d'ordonnancement, réseaux locaux

1 Introduction

Context of the paper In real-time systems, hard and soft timing constraints generally coexist. The problem of jointly scheduling Hard Real-Time (HRT) traffic and Soft Real-Time (SRT) traffic is an important issue in real-time computing and it arises both for tasks (scheduling on the CPU) and messages (scheduling on the network). The problem that we consider is to study some schedules that reduce as much as possible the average response time of SRT traffic while ensuring the timing requirements of HRT traffic to be met. In this study, it is assumed that the HRT traffic is periodic or sporadic, while no assumptions are placed on SRT traffic for most of the paper (a deterministic stability condition under the form of a (σ, ρ) bound, is used for implementation issues).

Existing work The simplest strategy for scheduling both SRT traffic and HRT traffic is to schedule the SRT traffic in the "background" (that will be called the background scheduling policy or BS for short), *i.e.* with a lower priority than any HRT tasks or messages. Experiments have shown that it leads to poor performances in terms of responsiveness of SRT traffic (see [3, 10] for the scheduling of tasks and for message scheduling [21, 12]). Several approaches performing better have been developed for the scheduling of tasks. Existing scheduling schemes include Earliest Deadline Last (EDL, [7]), Deferrable Server (DS, [18]), Priority Exchange (PE, [18]), Extended Priority Exchange (EPE, [27]), Static Slack Stealing algorithm (STS, [17]) and Dynamic Slack Stealing (DSS, [9]). As pointed out in [3] and [8], they all have disadvantages : DS, PE and EPE do not make use of all the available slack time, EDL and DSS are computationally expensive while SSS may require, depending on the task set, to store huge amount of information.

In [8, 10] an elegant and simple alternative, termed the Dual-Priority (DP) policy, has been proposed. In essence, this policy facilitates the responsive execution of SRT tasks by executing all HRT tasks immediately, when there are no SRT tasks ready, or as late as possible where SRT tasks are ready to be ran.

The DP scheme is applicable over a wide range of problems with non-significant overheads and no restrictive hypotheses except the knowledge of the worst-case response time for HRT tasks. Experiments published in [8, 10] and later in [3, 4] have shown through a series of simulations that the DP policy is highly effective in terms of responsiveness of SRT tasks in the context of the preemptive scheduling of tasks. From a practical point of view, the implementation of DP is quite straightforward and can be either done at the kernel level such as in the *Monstre* real-time OS [24]

(quoted in [3]) or at the application level for instance using ADA constructs as detailed in [6] and [3].

The dual-priority scheme is also usable for message scheduling on a network with bounded access time to the medium and this, even when nodes are not synchronized because the knowledge of the first release time of a periodic source is not mandatory. Its utilization has been proposed in [28] and it has shown to be very efficient [21, 12] in the context of a CAN (Controller Area Network [14]) based in-vehicle multiplexing system. Recently, the dual-priority strategy has been enhanced in various directions. For the scheduling of tasks, a method for computing tighter bounds on the response time for purely periodic task sets has been published in [3, 4]. In the context of message scheduling, a mechanism that provides probabilistic guarantees to prevent hard real-time frames from missing their deadlines when transmission errors can occur is proposed in [21].

Goal of the paper As previously mentioned, numerous experiments [8, 10, 3, 4, 21, 12] have shown the DP strategy to be very efficient both for the scheduling of tasks and the scheduling of messages. However, the problem of proving, in a precise manner, the efficiency of DP has not been addressed yet. A first important question is whether DP always behaves better than the classical strategy, termed Background Scheduling (BS for short), with which SRT tasks are assigned the lowest priority levels under Fixed Priority Scheduling.

In this paper, we will prove that the response times of the SRT traffic are always better under the DP policy if and only if the whole SRT traffic is FIFO. This result will be shown in the preemptive and the non-preemptive case.

The comparison between DP and BS will be done in a path-wise manner where a path (also called a trajectory) of the system is determined by the tasks (resp. messages) activation dates and by their execution (resp. transmission) times. These path-wise comparisons have several advantages. Indeed, they will hold for arbitrary distributions of the arrival process of the SRT traffic and for any increasing functional of the response time, such as the expectation and moments of any order or even logarithm, exponential and Laplace transforms of the response times.

This result is two-fold. On one hand, it reinforces the results hinted by previous experimental studies showing substantial gain of DP over BS by providing a theoretical basis. On the other hand, it suggests that when the FIFO ordering of SRT tasks is not satisfied, one may experience a surprising phenomenon where some SRT tasks have a lower response time under BS. Indeed, in some cases where the FIFO condition does not hold, we exhibit examples where the response times of some SRT

tasks (messages) are better under BS than under DP. Also, several simulations were run under a non-FIFO case for SRT traffic in the context of the CAN network. They also have the same kind of behavior where the response times of some SRT messages is smaller under BS than under DP.

Organization of the paper In Section 2, we introduce the framework and the formulation of the problem as well as the notations used in the following. Section 3 deals with the non-preemptive case while Section 4 treats the preemptive case. Section 5 reports some simulations of a CAN priority bus that illustrate the different behaviors coming up in the theoretical parts. Finally, in Section 6, we investigate how to ensure the FIFO ordering of SRT traffic.

2 Framework

The context of this study is the preemptive and the non-preemptive scheduling of mixed SRT and HRT traffic on a shared resource that can be either a processor or a network. Up to Section 4, we will focus on the non-preemptive case. In order to keep a unified vocabulary, we will talk about tasks even though everything, in the non-preemptive case, is transposable to messages.

Throughout this paper, we will use some concepts and ideas introduced in [16] and refined later in [19]. The system under study can be modeled by a finite set of m *recurrent tasks* and one resource that executes the successive *instances* of these tasks. Regarding the timing constraints, the set of tasks can be split into two subsets :

- the HRT tasks, $\mathcal{H} = \{\tau_1, \dots, \tau_p\}$,
- and the set of SRT tasks $\mathcal{S} = \{\tau_{p+1}, \dots, \tau_m\}$.

For any task τ_i , $\tau_{k,n}$ denotes the n -th instance of task τ_k , $A_{k,n}$ is the release time of instance $\tau_{k,n}$, and $C_{k,n}$ is the load brought by instance $\tau_{k,n}$. The HRT tasks are assumed to be periodic (resp. sporadic) with a period (resp. minimal inter-arrival time) denoted by T_k for task τ_k . As for the SRT tasks, no assumption are placed in most of the paper, and SRT traffic can be completely arbitrary. Only in Section 6, a deterministic stability condition will be introduced in order to compute upper bounds on busy periods.

The resource (which capacity is fixed to 1) is shared by all the tasks according to a scheduling policy (S) which assigns the resource to the instances. Under S, $B_{k,n}^{(S)}$ is the time when execution of $\tau_{k,n}$ begins and $E_{k,n}^{(S)}$ is the time when execution of $\tau_{k,n}$

ends. $\mathbf{R}_{k,n}^{(S)}$ is the response time of $\tau_{k,n}$ and by definition $\mathbf{R}_{k,n}^{(S)} = E_{k,n}^{(S)} - A_{k,n}^{(S)}$. If an instance $\tau_{k,n}$ has been released before time t but has not been completed at time t (i.e. $A_{k,n} \leq t < E_{k,n}^S$), then $\tau_{k,n}$ is said to be *pending* at time t . The set of all instances pending at time t is denoted by $\mathcal{H}^S(t)$. Each instance of an HRT task τ_k has a *deadline* D_k relatively to its arrival time. The system is said *feasible* under S if each instance of all HRT tasks meets its deadline. Formally,

$$\forall k, \quad \forall n, \quad \mathbf{R}_{k,n}^S \leq D_k. \quad (1)$$

The resource, scheduled under the policy S, is *busy* at time t if a task is being executed at time t . The busy indicator is the right-continuous function $\beta^S(t)$ equal to 1 whenever the resource is busy at time t and 0 otherwise.

For the sake of simplicity, in the rest of the paper, the following restrictions are placed :

1. Tasks have no jitter in their release date.
2. Context switch latencies are neglected.
3. HRT tasks have deadlines that are less than or equal to their periods.

Jitter in task availability dates (assumption 1) and context switch latencies (assumption 2) can be taken into account in the schedulability analysis as in [5]. The third assumption (deadlines must not be greater than periods) can be relaxed as in [10].

2.1 The Background Scheduling BS policy

Under fixed priority scheduling, the priority assignment with which all SRT tasks are given lower priorities than HRT tasks is called the Background Scheduling BS policy. Although experiments [3, 10, 21] have shown that BS performances are poor in terms of responsiveness of SRT tasks under heavy load, BS has 2 key advantages; it is straightforward to implement and the feasibility of HRT tasks is easily ensured.

Priorities BS is a Fixed Priority Scheduling, thus all instances $\tau_{k,n}$ are given a fixed priority level, denoted $\pi^{BS}(k, n)$. We will assume with no loss of generality that all instances are ordered according to their numbering. In other words, $\pi^{BS}(k, n) = (k, n)$, ordered using the alpha-numerical ordering. The priority numbering scheme adopted is "the smaller the number, the higher the priority" and the priority rule says that if $\pi^{BS}(k, n) < \pi^{BS}(k', n')$, then instance $\tau_{k,n}$ has a *higher priority* than instance $\tau_{k',n'}$ (and $\tau_{k',n'}$ has a lower priority than $\tau_{k,n}$). If several instances of the

same task are pending at the same instant, then, the earliest release has priority over all other instances of the same task (the system is FIFO within one task). As previously mentioned, under the BS scheme, all HRT tasks have higher priority than all SRT tasks.

Allocation rule In the non-preemptive case, the BS policy behaves according to the following rule:

*as soon as the resource is not busy, the pending instance
with the highest priority starts being executed.*

Note that under the non-preemptive allocation rule, the instance being executed can change only at completion times.

Feasibility is equivalent to $\mathbf{R}_k^{BS} \leq D_k \quad \forall \tau_k \in \mathcal{H}$, where \mathbf{R}_k^{BS} is the worst-case response time of task τ_k . The value of \mathbf{R}_k^{BS} in the non-preemptive case, is the maximum time needed by the task to gain the resource (denoted by I_k) plus C_k . When $D_j < T_j$ for all HRT tasks (the case $D_j > T_j$ is slightly more complicated, see [1]), then τ_k can be delayed by higher priority tasks and by one lower priority task that has already obtained the resource (in the worst-case, it is the execution time of the biggest task with priority lower than τ_k). From [13], we have :

$$\mathbf{R}_k^{BS} = C_k + I_k \quad (2)$$

where I_k is the longest time all higher priority tasks can occupy the bus plus the execution time of the biggest lower priority task. I_k is defined as the limit, when n goes to infinity, of the sequence

$$I_k^0 = 0, \quad I_k^n = \max_{i>k} (C_i) + \sum_{j<k} \left(\left\lfloor \frac{I_k^{n-1}}{T_j} \right\rfloor + 1 \right) C_j. \quad (3)$$

The quantity I_k is computed starting with $I_k^0 = 0$, until convergence or until $I_k^n > D_k - C_k$. In the latter case, τ_k is not guaranteed to respect its deadline and the set of tasks is non-feasible.

2.2 The Dual-Priority DP policy

The main difference with BS is that the priority level of the HRT tasks may change dynamically over time. This policy has been proposed in [8, 10]. Under DP, the priority range must be partitioned into three bands: "low hard", "soft", "high hard"

in increasing level of priority. All the priorities in the "low hard" range are lower than all "soft" which are lower than "high hard". An HRT task is first queued with a priority within the "low hard" band and later, when it becomes urgent, it will be promoted to the "high hard" range. Instead of executing HRT tasks as soon as they are available, they can be deferred in favor of SRT tasks until they become urgent. We define a *critical interval* for a HRT task τ_k as any time interval of the form $]n\tau_k + D_k - \mathbf{R}_k^{BS}, n\tau_k + D_k]$, for all $n \in \mathbb{N}$. The priorities of HRT tasks change over time. For all $k \leq p$,

$$\pi^{DP}(k, n, t) = \begin{cases} (k, n) & \text{if } t \text{ is in a critical interval for task } \tau_k \\ & \text{"high hard" priority),} \\ (m + k, n) & \text{otherwise ("low hard" priority).} \end{cases} \quad (4)$$

The priority of SRT tasks τ_k remains fixed over time in the "soft" priority range : $\forall k > p, \forall t, \pi^{DP}(k, n, t) = (k, n)$.

Once the priorities are defined, the policy follows the same rule as BS: *as soon as an execution is completed, the pending instance with the current highest priority starts being executed.*

An illustration of a trajectory under BS and DP policies is represented in Figure 1. The rectangles represent the instances of the tasks, with a grey area when the instance is being executed and a white area when it is pending and waiting to be executed. In this example, four tasks are competing for the resource : τ_1 and τ_2 are HRT tasks while τ_3 and τ_4 are SRT tasks (with the notations adopted in Section 2, we have $p = 2$ and $m = 4$). The release times are $A_{1,1} = 0, A_{2,1} = 2, A_{3,1} = 4, A_{4,1} = 6$. The execution times are respectively : $C_1 = 5, C_2 = 3, C_3 = 5, C_4 = 4$. At time 0, the only pending instance is $\tau_{1,1}$, therefore, it starts being executed immediately, under DP as well as under BS: $B_{1,1}^{BS} = B_{1,1}^{DP} = 0$. When instance $\tau_{1,1}$ is completed (at time 5), then $\tau_{2,1}$ and $\tau_{3,1}$ are pending. Under BS, their respective priorities are: $\pi^{BS}(2, 1) = (2, 1)$ and $\pi^{BS}(3, 1) = (3, 1)$. Therefore, $\tau_{2,1}$ gets the resource: $B_{2,1}^{DP} = 5$. Under DP, their respective priorities are: $\pi^{DP}(2, 1, 5) = (2 + 4, 1)$ (we assume that instance $\tau_{2,1}$ is not in its critical interval at time 5) and $\pi^{DP}(3, 1, 5) = (3, 1)$. Therefore, $\tau_{3,1}$ gets the resource and $B_{3,1}^{DP} = 5$. Under DP, when $\tau_{3,1}$ is completed at time 10, $\tau_{2,1}$ and $\tau_{4,1}$ are pending. Their respective priorities are: $\pi^{DP}(2, 1, 10) = (2, 1)$ (it is in its critical interval) and $\pi^{DP}(4, 1, 10) = (4, 1)$. Therefore, $\tau_{2,1}$ gets the resource: $B_{2,1}^{DP} = 10$. Finally, under DP the response time of $\tau_{3,1}$ and $\tau_{4,1}$ is respectively 6 and 11 versus 9 and 11 under BS.

Lemma 1. *The policy DP is feasible if and only if BS is feasible.*




Figure 1: A trajectory under DP and BS policies.

This lemma is a direct consequence of the computation of a bound on the worst response time of HRT tasks under DP. This computation has been proposed in [10]. Note however that under DP, even if the response time of a HRT task is guaranteed to be smaller than the deadline, it may be increased.

3 Dual-Priority VS Background Scheduling : the non-preemptive case

Numerous simulations suggest that DP improves substantially the average response time of SRT tasks over BS. For the non-preemptive, simulations were performed in [21, 12] for message scheduling on a CAN (Controller Area Network) priority bus. These simulations use exponential inter-arrival time for SRT messages and measure the average response time and the variance of response times. For both criteria, the gain provided by DP over BS is very important. For instance, the observed gain for average SRT response time ranges from a factor of 2.4 for a 60% total load to a factor of 3.8 for a 95% total load. It is also noteworthy that the dual-priority scheme offers very good resistance to an increase of the network load up to 90%.

However, up to the authors' knowledge, no formal proof of the fact that policy DP actually improves the response times of SRT tasks over BS has been published so far. In this section, we will provide the necessary and sufficient condition under which DP is better than BS on any trajectory. The path-wise method has the advantage that one will be able to compare any increasing functional of the response time for any arrival process of the SRT tasks.

3.1 Busy periods

We define the *total workload* at time t as the left continuous function

$$W(t) = \sum_{(k,n)} C_{k,n} \mathbf{1}_{\{A_{k,n} < t\}} - \int_0^t \beta(t) dt. \quad (5)$$

The function $W(t)$ can be seen as the total amount of work which has arrived before time t and which is still waiting to be done. The first term in Equation (5) is the amount of work arrived before t while the second term is the work done between 0 and t , with time 0 being the availability date of the resource.

Lemma 2. *The total workload is the same under BS and DP.*

Proof. In order to prove this lemma, we will use the fact that both policies are *non-idling* which means that whenever there is something ready to be executed, the resource is busy.

By definition of the resource busy functions, we have:

$$W^{BS}(t) = \sum_{(k,n)} C_{k,n} \mathbf{1}_{\{A_{k,n} < t\}} - \int_0^t \beta^{BS}(t) dt, \quad (6)$$

$$W^{DP}(t) = \sum_{(k,n)} C_{k,n} \mathbf{1}_{\{A_{k,n} < t\}} - \int_0^t \beta^{DP}(t) dt. \quad (7)$$

Therefore, if $\beta^{BS}(t) = \beta^{DP}(t)$ for all t , then the workloads will coincide. Let us assume that the busy functions are not the same under both policies. By right continuity of the functions $\beta^{DP}(t)$ and $\beta^{BS}(t)$, then there exists a time $t_0 > 0$ such that:

$$\begin{aligned} \beta^{DP}(t) &= \beta^{BS}(t), \quad \forall t \text{ s.t. } 0 \leq t < t_0, \\ \beta^{DP}(t_0) &\neq \beta^{BS}(t_0). \end{aligned}$$

Let us assume that we have $\beta^{DP}(t_0) = 0$ and $\beta^{BS}(t_0) = 1$. Since the release times of all tasks as well as their execution durations are the same under both policies, and using the fact that the busy functions are the same in both cases up to time t_0 , Equations (6) and 7 yield

$$W^{BS}(t_0) = W^{DP}(t_0).$$

Now, using the non-idleness of policy DP, $\beta^{DP}(t_0) = 0$ means that $W^{DP}(t_0) = 0$. As for BS, we have $\beta^{BS}(t_0) = 1$ and $W^{BS}(t_0) = 0$ which is only possible if there is an arrival at time t_0 in BS. But since the arrival times coincide under BS and DP, this implies that there is an arrival under DP at time t_0 as well, contradicting the fact that $\beta^{DP}(t_0) = 0$. The same argument holds for proving that $\beta^{DP}(t_0) = 1$ and $\beta^{BS}(t_0) = 0$ is impossible. \square

Definition 1 (busy period, cluster). *A busy period is a time interval $[t_1, t_2[$ such that $\beta(t) = 1$ for all $t \in [t_1, t_2[$ and $W(t_1) = 0$, $W(t_2^+) = 0$. The set of all instances arriving in a busy period, form a cluster, and will be denoted by \mathcal{C} .*

An immediate corollary of Lemma 2 is that busy periods and clusters coincide under BS and DP.

3.2 A counter-example of $DP > BS$

In this section, we present a case where the response times of some SRT tasks are smaller under BS than under DP. Let us consider the trajectory illustrated in Figure 2 with 2 instances of HRT tasks ($A_{1,1} = 0$ with $C_1 = 5$, $A_{2,1} = 3$ with $C_2 = 3$) and 3 instances of SRT tasks ($A_{3,1} = 6$ with $C_3 = 5$, $A_{4,1} = 4$ with $C_4 = 7$). Considering that neither instance $\tau_{1,1}$ nor instance $\tau_{2,1}$ are in a critical interval, we obtain on this trajectory $E_{1,1}^{BS} = 5$, $E_{2,1}^{BS} = 8$, $E_{3,1}^{BS} = 13$, $E_{4,1}^{BS} = 20$ for BS and $E_{1,1}^{DP} = 5$, $E_{2,1}^{DP} = 20$, $E_{3,1}^{DP} = 17$, $E_{4,1}^{DP} = 12$ for DP.

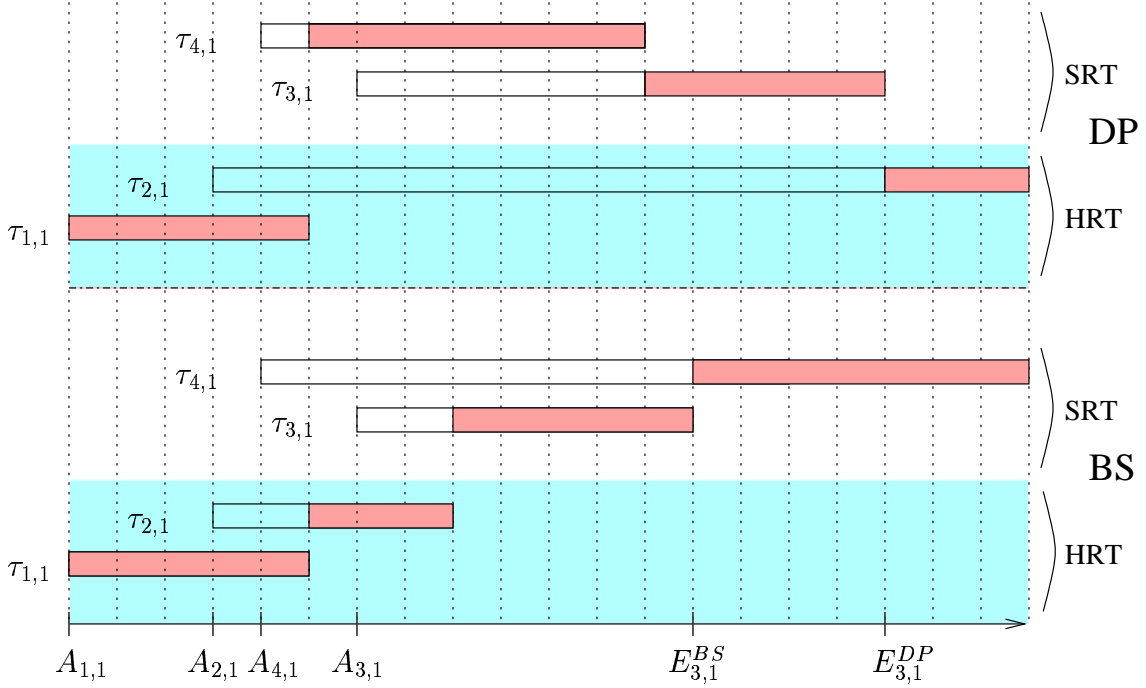


Figure 2: A trajectory under DP and BS in the non-preemptive case, where BS performs better than DP for $\tau_{3,1}$.

Note that for instance $\tau_{3,1}$, BS performs better than DP because $E_{3,1}^{BS} = 13 < E_{3,1}^{DP} = 17$. This inversion is due to the fact that SRT tasks arrive in a non FIFO order, as it is proved in the next section.

3.3 The FIFO Case

In this section, all instances of SRT tasks are executed in the FIFO order on the bus. With no loss of generality, we can assume that there exists a single SRT task τ_m , with $m = p + 1$, that gathers all the instances of SRT tasks.

The first step is to compare DP and BS over a single busy period involving a cluster \mathcal{C} . The global comparison will be derived from the individual comparisons over each cluster of the trajectory.

Over one busy period, we transform DP into a fixed non-preemptive priority policy, called EP in the following, which behaves exactly as DP by choosing appropriate priorities. Consider the completion times of all the instances in the cluster \mathcal{C} under DP and assign priorities to all the instances according to the order of their completion time. Namely, $\pi^{EP}(\tau_{k,n}) = \#\{(i, j) \in \mathcal{C} : E_{i,j}^{DP} < E_{k,n}^{DP}\}$.

Lemma 3. *For all instances $\tau_{k,n}$ in \mathcal{C} , $E_{k,n}^{DP} = E_{k,n}^{EP}$.*

Proof. First note that EP is non-idling and has the same clusters and busy periods as DP. The proof holds by induction on the size of the cluster considered. If the cluster is made by only one task, then clearly, $E_{k,n}^{DP} = A_{k,n} + C_{k,n} = E_{k,n}^{EP}$. Now assume that the cluster is made of i instances. We remove the task executed last under DP, say instance $\tau_{a,b}$. We get a cluster made of $i - 1$ tasks. By induction, $E_{k,n}^{DP} = E_{k,n}^{EP}$ for all tasks in this reduced cluster.

Under DP, the completions of all tasks in the original cluster, are the same as in the reduced cluster since $\tau_{a,b}$ is executed last. As for EP, task $\tau_{a,b}$ has the lowest priority by construction of EP. Since the remaining $i - 1$ tasks form a cluster, the additional task $\tau_{a,b}$ with the lowest priority is transmitted last. \square

Lemma 4. *Consider an arbitrary instance of a SRT instance, say $\tau_{m,n}$. Then $E_{m,n}^{EP} \leq E_{m,n}^{BS}$.*

Proof. We first show that under EP the priorities of all SRT instances are ordered according to their arrival. Let $\tau_{m,n}$ be an arbitrary SRT instance. Under policy DP, the priorities of SRT tasks are static and FIFO (by hypothesis). Since $A_{m,n} < A_{m,n+1}$ then, $\pi^{DP}(m, n, t) < \pi^{DP}(m, n + 1, t)$ for all t . Using the execution rule of policy DP, if instance $\tau_{m,n}$ has not yet been transmitted at time $A_{m,n+1}$, then both tasks are pending. Whenever there is a execution opportunity, $\tau_{m,n}$ has priority over $\tau_{m,n+1}$ by the FIFO rule, therefore, it will be transmitted earlier. This means that $E_{m,n} < E_{m,n+1}$. As for EP, by definition, $\pi^{EP}(m, n) < \pi^{EP}(m, n + 1)$.

The priority order among SRT instances is the same under EP and BS. As for real-time instances, their priority in EP is a modification with respect to those in BS. By applying Lemma 10, given in Appendix 7, we obtain $E_{m,n}^{EP} \leq E_{m,n}^{BS}$ for all n . \square

Theorem 1. *For all SRT instances $\tau_{m,n}$, $E_{m,n}^{DP} \leq E_{m,n}^{BS}$. In other words, the response times of all SRT tasks are smaller under DP than BS.*

Proof. The proof is a direct consequence of Lemmas 3 and 4. \square

4 Dual-Priority VS Background Scheduling : the preemptive case

Several simulations published in [8, 10, 3, 4] suggest to us that DP improves greatly the response time of SRT tasks over BS in the preemptive case as well. However, as for the non-preemptive case, to the best of our knowledge, no formal proof, of the efficiency of DP has been given.

In the preemptive case, the priorities defined for BS and DP respectively are the same as in Section 3 (non-preemptive case). The only difference lies in the resource allocation rule which is now, for both policies :

the instance of the task being executed at time t is the pending task of highest priority.

The difference is that the task being executed can change suddenly, (for example as soon as a high priority task is released) and not only at completion times as in the non-preemptive case. Thus, the execution of a task $\tau_{k,n}$ can be split in several *pieces*, $T_{k,n,i}$ being the i th piece with execution beginning and completion respectively denoted by $B_{k,n,i}$ and $E_{k,n,i}$. Each interruption in the processing of a task is caused by an instance of a higher priority task that takes the resource. Note that under DP, the change can occur when an instance entering a critical interval, has its priority promoted.

As for the non-preemptive case, we will compare DP and BS over each trajectory. Here again, DP performs always better than BS when all SRT tasks are FIFO. Surprisingly, the proof does not work in the general case for completely different reasons. In the non-preemptive case, changing the priorities only improves the lowest priority task (Lemma 10), while in the preemptive case, all tasks with an improved priority benefit from it (Lemma 11). The problem for the preemptive case comes from the comparison of BS with a fixed priority policy that behaves as DP, where the FIFO property is needed to have compatible priorities.

As for the non preemptive case, feasibility under BS is always equivalent to $\mathbf{R}_k^{BS} \leq D_k \quad \forall \tau_k \in \mathcal{P}$. The computation of \mathbf{R}_k^{BS} in the preemptive case is classical, it has been for instance studied by Joseph and Pandya [15] and later by Audsley et al. [2] :

$$\mathbf{R}_k^{BS} = I_k \quad (8)$$

where I_k , the longest time that all higher priority tasks can occupy the resource, is the limit of the sequence

$$I_k^0 = 0, \quad I_k^n = C_k + \sum_{\forall j < k} \left\lceil \frac{I_k^{n-1}}{T_j} \right\rceil C_j \quad (9)$$

I_k is computed starting with $I_k^0 = 0$ until convergence or until $I_k^n > D_k - C_k$.

4.1 Busy periods

Lemma 5. *The total workload is the same under BS and DP.*

Proof. The proof is exactly the same as in the non-preemptive case, see Lemma 2. \square

An immediate corollary is that busy periods as well as clusters are the same under DP and under BS. One can also verify that the workload (as well as busy periods and clusters) are the same under preemptive and non-preemptive systems because of the non-idleness.

4.2 Completion time ordering

In a similar way as in the non-preemptive case, we construct an intermediate policy EP with fixed priorities that behaves exactly as policy DP. However, the construction is more involved.

To construct EP, we consider each piece $\tau_{k,n,i}$ of execution of task $\tau_{k,n}$ under DP as an instance of task τ_k , called $\tau_{k,n,i}$, with arrival time $A_{k,n}$. Therefore, under EP, several instances (as many as there are pieces under DP for $\tau_{k,n}$) are released at the same instant. The priorities under EP are given according to the end of execution of all the pieces:

$$\pi^{EP}(k, n, i) = \#\{(x, y, z) | E_{x,y,z}^{DP} < E_{k,n,i}^{DP}\}. \quad (10)$$

Lemma 6. *For all k, n, i , $E_{k,n,i}^{DP} = E_{k,n,i}^{EP}$.*

Proof. The proof holds by induction on the size of the cluster that includes $\tau_{k,n}$. If $\tau_{k,n}$ is alone in the cluster, then $\tau_{k,n}$ is not interrupted during its execution and it is formed by a single piece $\tau_{k,n,1} = \tau_{k,n}$. By Lemma 5, we have $E_{k,n}^{DP} = E_{k,n,1}^{EP}$. Now, let us assume that the cluster is formed of n pieces. The reduced set composed of the same pieces except the last one to be executed under DP forms a cluster with $n - 1$ pieces. By induction, the completion times are the same under both policies. As for the last task, it is executed last under DP, therefore, according to Equation (10) it has the lowest priority under EP and will be thus executed last among the considered cluster. \square

4.3 A counter-example of DP > BS

We now show an example of a trajectory for which the preemptive scheduling under DP is not better than under BS for all SRT tasks. As for the non-preemptive case of Section 3.2, this counter example uses SRT tasks arriving in a non FIFO order. Here, we also need a HRT task which has its priority promoted.

Let us consider a trajectory (see Figure 3) with a single HRT task ($A_{1,1} = 5$ and $C_1 = 5$) and two SRT tasks ($A_{2,1} = 8$ and $C_2 = 3$, $A_{3,1} = 0$ and $C_4 = 7$), the HRT instance $\tau_{2,1}$ entering its critical interval at time 9. Under this trajectory, we obtain $E_{1,1}^{DP} = 13$, $E_{2,1}^{DP} = 15$, $E_{3,1}^{DP} = 7$ and $E_{1,1}^{BS} = 10$, $E_{2,1}^{BS} = 13$, $E_{3,1}^{BS} = 15$.

Figure 3 represents this trajectory, one notes that $E_{2,1}^{BS} < E_{2,1}^{DP}$, thus for $\tau_{2,1}$ BS performs better than DP. Again, this phenomenon is due to the non-FIFO feature of the SRT traffic. Note that it is possible to build examples for which the difference $E_{i,j}^{DP} - E_{i,j}^{BS}$ is arbitrarily large, for some SRT instances $\tau_{i,j}$.

4.4 The FIFO case

In this section, we consider that all SRT instances are ordered in the FIFO order under BS. As for the non-preemptive case, we can assume that there exists a single SRT task τ_m , with $m = p + 1$, that gathers all the instances of SRT tasks.

We first need to modify the tasks under BS by considering each piece of execution $\tau_{k,n,i}$ (defined in Section 4.2) to be a new task with arrival time $A_{k,n}$ and priority $\pi^{BS} = (k, n, i)$.

Lemma 7. *Under EP, the SRT tasks verify $\pi_{m,n,i}^{EP} < \pi_{m,n+1,j}^{EP}$ and $\pi_{m,n,i}^{EP} < \pi_{m,n,i+1}^{EP}$ for all n, j, i .*

Proof. By Lemma 6, we have $E_{m,n,i}^{DP} = E_{m,n,i}^{EP}$. By construction of the tasks $\tau_{m,n,i}$, we have $E_{m,n,i}^{DP} < E_{m,n+1,j}^{DP}$ and $E_{m,n,i}^{DP} < E_{m,n,i+1}^{DP}$ for all n, j, i . By definition of the




Figure 3: A trajectory under DP and BS in the preemptive case, where BS performs better than DP for $\tau_{2,1}$.

priorities under EP, see Equation (10), we obtain $\pi_{m,n,i}^{EP} < \pi_{m,n+1,j}^{EP}$ and $\pi_{m,n,i}^{EP} < \pi_{m,n,i+1}^{EP}$ for all n, j, i . \square

Lemma 8. *Under the FIFO assumption on SRT tasks $E_{m,n,i}^{DP} \leq E_{m,n,i}^{EP}$ for all n and i .*

Proof. First note that a task $\tau_{k,n,i}$ may not be executed in a single piece under BS, unlike under EP. If a single cluster is considered, then all the SRT tasks have compatible priorities under BS and EP using Lemma 7. As for the HRT tasks, their priorities can be very different under both policies. However, for each SRT task $\tau_{m,n,i}$, all HRT tasks have higher priority than $\tau_{m,n,i}$ under BS. Combining the two previous properties, the set of tasks with higher priority under EP is included in the set of tasks with higher priority under BS. Therefore we can apply Lemma 11, given in Appendix 7, which implies that $E_{m,n,i}^{EP} \leq E_{m,n,i}^{BS}$. \square

Theorem 2. *For any task $\tau_{k,n}$, $E_{k,n}^{DP} \leq E_{k,n}^{BS}$.*

Proof. First, note that $E_{k,n}^{DP} = E_{k,n,h}^{DP}$, where piece $\tau_{k,n,h}$ is the last piece of task $\tau_{k,n}$ under DP. Using Lemma 6, we obtain $E_{k,n,h}^{EP} = E_{k,n,h}^{DP}$. Now, using Lemma 8, $E_{k,n,h}^{EP} \leq E_{k,n,h}^{BS}$. Finally, using Lemma 7, $E_{k,n}^{BS} = E_{k,n,h}^{BS}$, which ends the proof. \square

5 Experiments

Many simulation results have been published in the literature that show a substantial gain in terms of response time for SRT traffic when DP was used instead of BS. Under heavy load ($> 70\%$), the gains usually encountered were above 60% (see experiments 5.3 and 5.4 in [10], Figure 4 in [4], Figure 2 in [21] and Figure 5 in [12]). In this section, in order to exhibit the practical implications of Theorem 2, we present new simulations done in the non-preemptive case.

In the first experiment, we consider a realistic CAN-based in-vehicle application provided by the car industry (see [22] for a detailed description) where 6 devices (e.g. engine controller, automatic gear box, ...) exchange messages. The traffic consists of a set of 12 HRT messages (e.g. speed and torque from the engine controller) with periods $T_1 = 10$, $T_2 = 14$, $T_3 = 20$, $T_4 = 15$, $T_5 = 20$, $T_6 = 40$, $T_7 = 15$, $T_8 = 50$, $T_9 = 20$, $T_{10} = 100$, $T_{11} = 50$ and $T_{12} = 100$ ms, all having a size of 125 bits. The periodic sources are assumed to begin transmitting from the start of the simulation. In addition, there exists a stream of SRT frames whose inter-arrival times are exponentially distributed. Among the SRT traffic, we distinguish

15 different messages all of size 100 bits. The transmission rate of the CAN bus is 125kbit/s and the total network load induced by the 27 messages is 90% with 53.46% for the periodic part.

In order to show the importance of the priority order among the SRT messages, we have imposed the SRT messages to be emitted in a quasi-LIFO order, that is in the order :

$$\tau_{27,1}, \tau_{26,1}, \tau_{25,1}, \dots, \tau_{15,1}, \tau_{14,1}, \tau_{13,1}, \tau_{27,2}, \dots$$

The simulations using DP and BS policies respectively were run on more than 10,000 instances of HRT messages.

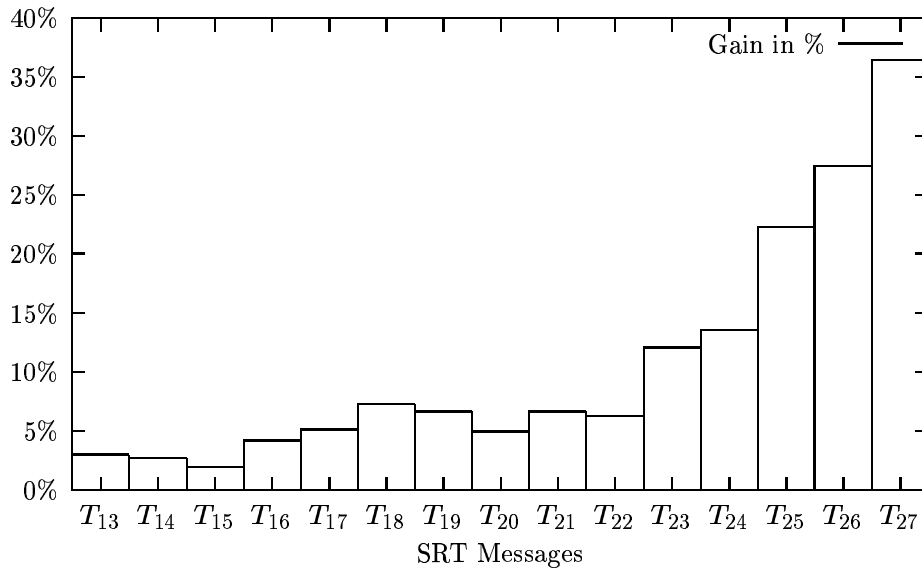


Figure 4: Average gain in response time of DP over BS for each SRT message with 12 HRT messages.

The results, represented in Figure 4, show the gain (given in percentage) of DP over BS for each SRT message in terms of the average response time. It is worth noting that while the gain remains substantial for message τ_{27} (around 40%), this is not the case for messages in the high and mid-priority ranges (the gain for τ_{15} is for instance around 2%). This is very small compared to other simulations which were published in the literature.

In the second set of experiments, we have chosen an environment which was likely to exhibit the behavior illustrated in Figure 2, where some SRT tasks (here messages) have a better response time under BS than under DP. We simulate a CAN priority bus with a stream composed of 15 different SRT messages of size 100 bits and with a single HRT message of length 115bits and period $\tau_1 = 1\text{ms}$. The transmission rate is 250kbit/s, the HRT message induces a load of 46% while the total load is 95%. Again, the SRT messages are released in a quasi LIFO order

$$\tau_{16,1}, \tau_{15,1}, \tau_{14,1}, \dots, \tau_{4,1}, \tau_{3,1}, \tau_{2,1}, \tau_{16,2}, \dots$$

The results of these simulations are shown in Figure 5. Here, it is remarkable to

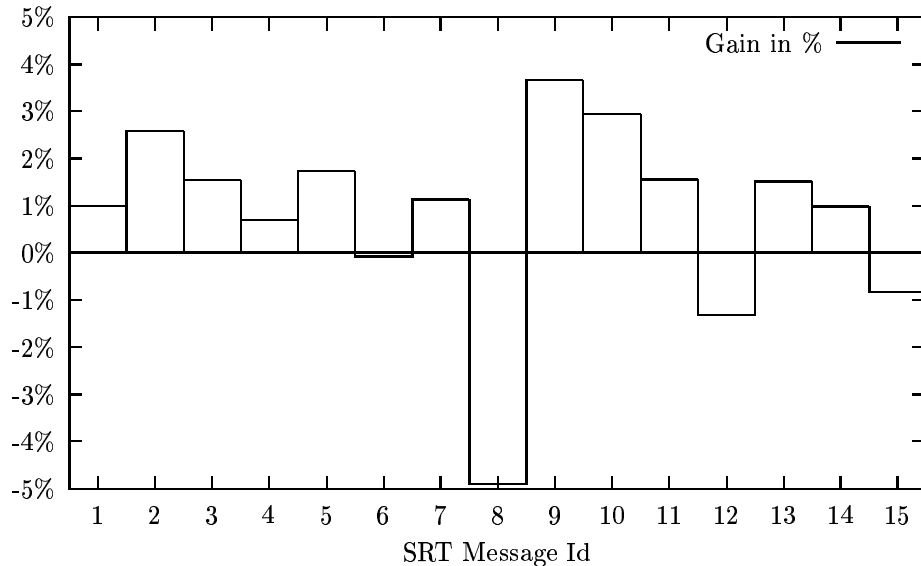


Figure 5: Average gain in response time of DP over BS for each SRT message with 1 HRT message.

notice that 4 out of 15 SRT messages have an average response time better under BS than under DP, namely tasks $\tau_6, \tau_8, \tau_{12}$ and τ_{15} . Note also that the gain obtained for task τ_8 under BS is about 5%, and is the largest absolute gain. The latter simulation shows that the behavior exhibited on a single well-chosen trajectory may also be observed in average over long typical runs.

6 Implementation Issues

Throughout the paper we have seen that the FIFO ordering for SRT tasks is critical for the efficiency of DP. In this section, we are concerned with the practical problem of ensuring the FIFO ordering for SRT traffic.

In a centralized framework, the FIFO ordering for SRT tasks is easily achieved through prioritization, the problem that will be addressed in this section is to find a way of imposing the FIFO ordering for the scheduling of messages in a distributed environment. We will focus on the CAN priority bus but the principles of the analysis remains valid for other priority buses such as the J1850 [26] or VAN [11]. The use of DP for message scheduling on CAN has been proposed by Tindell and Hansson in [28]. The mechanisms described in the following are compatible with the use of DP on CAN.

Each station may emit SRT messages as well as HRT messages. Each message is composed of K bits for priority encoding and M bits for control and data. In the following, we will propose a way to encode the FIFO feature on those K bits, provided that all stations have a synchronous clock. Note that the encoding of information using some bits of the identifier is rather classical in CAN, see for instance [29] and [23]. To synchronize the stations, we have to add an initialization phase, which consists in sending a special frame that serves as global starting signal. This signal will serve as the origin of time. The granularity of time being the bit-time. Since the CAN protocol ensures that all nodes are synchronized on the bit-time, all stations will have coherent clocks.

For HRT messages, the K priority bits are used in the following way : the first 2 bits are set to 00 when the message is in a critical interval and 10 otherwise (see Figure 6). The following $K - 2$ bits are used to encode the priority level. This is possible as long as $K - 2 \geq \log_2(p)$ where p is the number of HRT frames. For SRT messages, the first 2 bits of the K priority bits are set to 01. The following $K - S - 2$ bits are used to encode the current time t , common to all stations by assumption, which is called the *time-stamp* of the frame. The last $S = \lceil \log_2(\# \text{ stations}) \rceil$ bits are used to encode the identifier of the sender station (in order to break ties in case of identical time-stamps). There is enough space as long as $K - S - 2 \geq \log_2(t)$. Note that distinguishing between the 2 types of HRT frames and SRT frames using the first 2 bits of the CAN frame has already been proposed in [28].

With this method the priority for an instance of a SRT message is its release instant. Therefore, with the assumption that all stations are synchronized, it makes sure that all the SRT traffic satisfies the FIFO ordering. However, the drawback of this method is the fact that it is impossible to encode SRT messages after time 2^{K-S-2} . In order

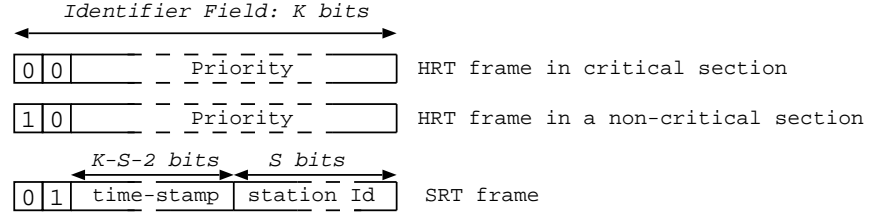


Figure 6: Identifier field of the CAN frame for HRT frames and SRT frames.

to deal with this problem, we propose to add one particular SRT message (called the *reset message*, denoted τ_{m+1}) that will be released every H units of time, with a priority level defined as for the SRT messages (i.e. function of its time-stamp and station identifier). On reception of this particular message, each station decreases its time clock by H units of time, as well as the time-stamp of all SRT messages which are waiting to be sent.

The question now is how to choose H and K such that no overflow of the time clock ever occurs. First, let us compute an upper bound L on the length of all busy periods in the system, this will give us a bound on the response time of the reset message. This calculus can be done under some assumptions on the SRT traffic, namely, a deterministic load arrival bound such as a (σ, ρ) condition (see [19]). We assume that for any message, τ_k , the work released in any interval of size t is bounded by the affine formula $\sigma_k + \rho_k t$. In particular, periodic and sporadic messages satisfy a (σ, ρ) condition. The deterministic stability of the system says: $\sum_{k=1}^{m+1} \rho_k < 1$. Now, let the interval $[t_1, t_2]$ be a busy period of the system. This means that the total load is null at time t_1 and at time t_2 . Therefore, the length of the busy period equals the load which has arrived during the interval $[t_1, t_2]$. This yields

$$L \leq \frac{\sum_{k=1}^{m+1} \sigma_k}{1 - \sum_{k=1}^{m+1} \rho_k}.$$

Lemma 9. *With*

$$H = 2^{K-S-2} - L, \quad (11)$$

there is no time overflow.

Proof. The k^{th} instance of the reset message is released at kH and completes its execution at e_k , when the clock is set to zero, see Figure 7.

The initial reset message being released just after the synchronization and before all other SRT instances, no overflow occurs until e_0 . Suppose that with the chosen H no time overflow has occurred until e_{k-1} . Because of the FIFO order, SRT instances pending at e_{k-1} have been released after $(k-1)H$. Due to the reset at e_{k-1} , their time-stamps and those of the instances released in $[e_{k-1}, e_k[$ vary between 0 and $e_k - e_{k-1}$. To avoid overflow, $e_k - e_{k-1}$ must be shorter than the longest time that can be encoded, i.e. shorter than 2^{K-S-2} . Since all response times are shorter than the longest busy period, we have $e_k \leq kH + L$. Furthermore $e_{k-1} \geq (k-1)H$. Thus,

$$e_k - e_{k-1} \leq kH + L - (k-1)H = H + L = 2^{K-S-2}.$$

□

Actually any choice $H \leq 2^{K-S-2} - L$ is an acceptable solution, but to limit the overhead induced by the reset message, H should be chosen as long as possible.

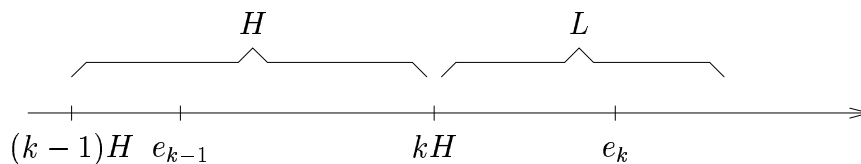


Figure 7: Choosing the period H of the reset message.

Note that the reset message contributes to the total load and also to the bound L of the busy periods, with σ_{m+1} being the transmission time of the reset message and $\rho_{m+1} = \sigma_{m+1}/H$. When K and the characteristics of the traffic are fixed, then Equation (11) can be solved in H (quadratic form).

In the context of CAN, the number K of bits allocated for priority encoding is fixed to either 11 bits (CAN 2.0A) or 29 bits (CAN 2.0B). A typical case (as for in-vehicle networks, see [25, 22]) would involve less than 64 stations, so that $S = 6$ is enough. With $K = 11$, this leaves only 3 bits for time-stamp encoding. Considering the 12 HRT messages described in Section 5 and a SRT traffic inducing a load of 75% made of 20 messages of 100 bits, the total load at 500kbit/s is 86%. When solving (11), we obtain two negative values for the period H of the reset message. Thus the proposed method is not applicable on CAN2.0A.

With CAN2.0B, $K = 29$ thus 21 bits are left for the time-stamp encoding which enables to encode a time-stamp range of 4194 ms. Using the same traffic as previously, the bound L on the longest busy period is 76 ms and the solution of Equation (11) is

$H = 4118$ ms. In these conditions, the load overhead of the reset message is 4.10^{-5} which is quasi neglectable.

7 Conclusion

In this paper, it has been proven for the preemptive as well as for the non-preemptive scheduling that the response times of the SRT traffic are always better under the Dual-Priority than under Background-Scheduling if and only if the whole SRT traffic is FIFO. This was confirmed by simulations of a CAN bus which have shown the response times of some SRT messages to be smaller under BS than under DP when the FIFO condition is not satisfied. The result of the study reinforces the results hinted by previous experimental studies showing substantial gain of DP over BS by providing a theoretical basis which gives practical guidelines for application designers willing to implement DP. Finally, such a path-wise method of comparing scheduling policies may also be used for several other scheduling policies such as EDL [7], PE [18] or EPE [27]. This is currently under investigation.

Acknowledgment The authors would like to thank Gerald Cabus for some helpful discussions about the implementation issues of the FIFO feature on priority buses.

Appendix A: Changing priorities (non-preemptive case)

This appendix is devoted to the proof of a general technical lemma about the effect on the change of priorities in the BS policy in the non-preemptive case.

A formula for the response time for an instance $\tau_{k,n}$ can be derived from [20] (p.22). It satisfies the following equality.

$$R_{k,n} = C_{k,n} + \min \{t \geq 0 \mid \Omega_{k,n}(A_{k,n}) + \Gamma_{k,n}(t) + \rho_{k,n}(A_{k,n}) = t\}, \quad (12)$$

where

- $\Omega_{k,n}(x)$ is the workload present at time x contributed by all instances with priority higher than $\pi(k, n)$,
- $\Gamma_{k,n}(x) = \sum_{(i,j)} \mathbf{1}_{\{\pi(i,j) < \pi(k,n)\}} \mathbf{1}_{\{A_{k,n} \leq A_{i,j} \leq A_{k,n} + x\}} C_{i,j}$ is the high priority work arrived between the release time, $A_{k,n}$ and time $A_{k,n} + x$.
- $\rho_{k,n}(x)$ is the remaining execution time of an instance with a priority lower than $\pi(k, n)$, which has started its execution before time x and is not completed yet.

As for the completion time,

$$\begin{aligned} E_{k,n} &= A_{k,n} + R_{k,n} \\ &= A_{k,n} + C_{k,n} + \min \{t \geq 0 \mid \Omega_{k,n}(A_{k,n}) + \Gamma_{k,n}(t) + \rho_{k,n}(A_{k,n}) = t\}. \end{aligned}$$

We construct a new fixed priority function π' in the following way: $\pi'(m, n) = \pi(m, n)$ for all instances of the lower priority task τ_m . As for any other task τ_k , $\pi'(k, i)$ is arbitrary (possibly larger than $\pi'(m, n)$).

All quantities related with the new priority will be denoted with a prime sign.

Lemma 10. *For any instance of the lowest priority class, $\tau_{m,n}$, $E'_{m,n} \leq E_{m,n}$.*

Proof. Consider an instance of the lower priority task $\tau_{m,n}$. This task belongs to a cluster, say \mathcal{C} of all instances involved in a busy period $[t_1, t_2[$. In the following of the proof, we will only consider the instances of tasks belonging to \mathcal{C} , since no other task will influence the completion time of task $\tau_{m,n}$ under both priorities.

Under the original priorities, since $\tau_{m,n}$ is the instance with the lowest priority released so far, $\rho_{m,n}(A_{m,n}) = 0$ and

$$E_{m,n} = A_{m,n} + C_{m,n} + \min \{t \geq 0 \mid \Omega_{m,n}(A_{m,n}) + \Gamma_{m,n}(t) = t\}. \quad (13)$$

As for the value of $E'_{m,n}$,

$$E'_{m,n} = A_{m,n} + C_{m,n} + \min \{t \geq 0 \mid \Omega'_{m,n}(A_{m,n}) + \Gamma'_{m,n}(t) + \rho'_{m,n}(A_{m,n}) = t\}. \quad (14)$$

In order to compare both values, we will examine closely the values of the different terms involved in Equations (13) and (14).

Since $\tau_{m,n}$ is the lowest priority tasks which has ever been released by time $A_{m,n}$, then

$$\Omega_{m,n}(A_{m,n}) = W(A_{m,n}) \quad (15)$$

$$= W'(A_{m,n}) \quad (16)$$

$$\geq \Omega'_{m,n}(A_{m,n}) + \rho'_{m,n}(A_{m,n}), \quad (17)$$

where, Equation (15) comes from the fact that when instance $\tau_{m,n}$ is released, it has the lowest priority so far, Equation (16) from the fact that the workload of all non-idling policies are equal (see Lemma 2) and Equation (17) from the definition of Γ' and ρ' which are both distinct part of the workload, $\Gamma'_{m,n}(A_{m,n})$ is the fraction

of the workload at time $A_{m,n}$ due to instances of high priority and $\rho'_{m,n}(A_{m,n})$ is the fraction of the workload due to an instance with lower priority than $\tau_{m,n}$ being transmitted at time $A_{m,n}$.

In addition, for all time t ,

$$\Gamma_{m,n}(t) = \sum_{(i,j)} C_{i,j} \mathbf{1}_{\{\pi(i,j) < \pi(m,n)\}} \mathbf{1}_{\{A_{m,n} \leq A_{i,j} \leq A_{m,n} + t\}}, \quad (18)$$

$$\geq \sum_{(i,j)} C_{i,j} \mathbf{1}_{\{\pi'(i,j) < \pi'(m,n)\}} \mathbf{1}_{\{A_{m,n} \leq A_{i,j} \leq A_{m,n} + t\}} \quad (19)$$

$$= \Gamma'_{m,n}(t). \quad (20)$$

Inequality (19) comes from the fact that the set of all instances $(\tau_{i,j})$ such that $\pi(i,j) > \pi(m,n)$ arriving after time $A_{m,n}$ is exactly the set $\{\tau_{m,i}, i > n\}$. It is included in the set of instances such that $\pi'(i,j) > \pi'(m,n)$ and arriving after time $A_{m,n}$, since π' does not modify the priorities among the task τ_m . The complementary sets are included in the reversed direction. For all (i,j) ,

$$\mathbf{1}_{\{\pi(i,j) < \pi(m,n)\}} \mathbf{1}_{\{A_{m,n} \leq A_{i,j} \leq A_{m,n} + t\}} \leq \mathbf{1}_{\{\pi'(i,j) < \pi'(m,n)\}} \mathbf{1}_{\{A_{m,n} \leq A_{i,j} \leq A_{m,n} + t\}}.$$

Combining both Equation (17) and (20) yields $E_{m,n} \geq E'_{m,n}$. \square

Appendix B: Changing priorities (preemptive case)

In the preemptive case, one can prove a more powerful result than Lemma 10. Here the response time of any instance is improved whenever its priority is improved, while Lemma 10 only considers the lowest priority task.

The completion time of an instance $\tau_{k,n}$ is given by the formula:

$$E_{k,n} = A_{k,n} + \min \{t \geq 0 \mid \Omega_{k,n}(A_{k,n}) + \theta_{k,n}(t) + C_{k,n} = t\},$$

where

- $\Omega_{k,n}(x)$ is the workload present at time x contributed by all instances with priority higher than $\pi(k,n)$,
- $\theta_{k,n}(x) = \sum_{(i,j)} C_{i,j} \mathbf{1}_{\{\pi(i,j) < \pi(k,n)\}} \mathbf{1}_{\{A_{k,n} \leq A_{i,j} < x + A_{k,n}\}}$ is the high priority load arrived between the release time $A_{k,n}$ and time $A_{k,n} + x$.

We modify some priorities such that the priority of instance $\tau_{k,n}$ is improved: for all (i, j) , $\pi'(i, j) < \pi'(k, n) \Rightarrow \pi(i, j) < \pi(k, n)$. Under this assumption on π' , the following lemma can be established.

Lemma 11. $E_{k,n} \geq E'_{k,n}$

Proof. The proof is similar to the proof of Lemma 10.

Again, using the non-idling conservation Lemma 5, the busy periods as well as the clusters are identical under both priorities. Task $\tau_{k,n}$ belongs to one cluster involved in a busy period, $[t_1, t_2[$. In the following only tasks in this cluster will be considered.

$$E_{k,n} = A_{k,n} + \min \{t \geq 0 \mid \Omega_{k,n}(A_{k,n}) + \theta_{k,n}(t) + C_{k,n} = t\},$$

and

$$E'_{k,n} = A_{k,n} + \min \{t \geq 0 \mid \Omega'_{k,n}(A_{k,n}) + \theta'_{k,n}(t) + C_{k,n} = t\}.$$

A first sequence of inequalities gives

$$\Omega_{k,n}(A_{k,n}) = W_{k,n}(A_{k,n}) \quad (21)$$

$$\geq W'_{k,n}(A_{k,n}) \quad (22)$$

$$= \Omega'_{k,n}(A_{k,n}), \quad (23)$$

where $W_{k,n}(t)$ (resp. $W'_{k,n}(t)$) is the workload at time t of a system where all tasks with priority lower than $\pi(k, n)$ (resp. $\pi'(k, n)$) have been removed.

Equality (21) comes from the fact that in preemptive systems, low priority tasks have no influence on the execution of higher priority tasks. Equality (22) comes from the fact that the set of the instances with higher priority than $\tau_{k,n}$ under π includes the the set of the instances with higher priority than $\tau_{k,n}$ under π' . Equality (23) is the same as Equality (21) but with π' .

The second series of inequalities is

$$\theta_{m,n}(t) = \sum_{(i,j)} C_{i,j} \mathbf{1}_{\{\pi(i,j) < \pi(m,n)\}} \mathbf{1}_{\{A_{m,n} \leq A_{i,j} < A_{m,n} + t\}} \quad (24)$$

$$\geq \sum_{(i,j)} C_{i,j} \mathbf{1}_{\{\pi'(i,j) < \pi'(m,n)\}} \mathbf{1}_{\{A_{m,n} \leq A_{i,j} < A_{m,n} + t\}} \quad (25)$$

$$= \theta'_{m,n}(t). \quad (26)$$

Inequality (25) comes from the definition of π' which implies that for all (i, j) ,

$$\mathbf{1}_{\{\pi(i,j) < \pi(m,n)\}} \mathbf{1}_{\{A_{m,n} \leq A_{i,j} < A_{m,n+t}\}} \leq \mathbf{1}_{\{\pi'(i,j) < \pi'(m,n)\}} \mathbf{1}_{\{A_{i,j} A_{m,n} \leq A_{i,j} < A_{m,n+t}\}}.$$

Combining both Equation (23) and (26) yields $E_{m,n} \geq E'_{m,n}$. \square

References

- [1] N. Audsley, A. Burns, M. Richardson, K. Tindell, and A. Wellings. Applying New Scheduling Theory to Static Priority Pre-emptive Scheduling. *Software Engineering Journal*, 8(5):284–292, September 1993.
- [2] N.C. Audsley, A. Burns, M.F. Richardson, and A.J. Wellings. Hard real-time scheduling: The deadline monotonic approach. In *Proceedings 8th IEEE Workshop on Real-Time Operating Systems and Software*, 1991.
- [3] G. Bernat. *Specification and Analysis of Weakly Hard Real-Time Systems*. Phd thesis, Universitat de les Illes Balears, Departament de Ciències Matemàtiques i Informàtica, January 1998.
- [4] G. Bernat and A. Burns. Combining (n, m) -hard deadlines and dual priority scheduling. In *Proceedings of the 18th IEEE Real-Time Systems Symposium*, December 1997.
- [5] A. Burns, K. Tindell, and A.J. Wellings. Effective analysis for engineering real-time fixed priority schedulers. *IEEE Transactions on Software Engineering*, 21(5):475–480, May 1995.
- [6] A. Burns and A.J. Wellings. Dual priority scheduling in ada95 and real-time posix. In *Proceedings of the 21st IFAC/IFIP Workshop on Real-Time Programming*, pages 45–50, 1996.
- [7] H. Chetto and M. Chetto. Some Results of the Earliest Deadline Scheduling Algorithm. *IEEE Transactions on Software Engineering*, 15(10):1261–1269, October 1989.
- [8] R. Davis. Dual Priority scheduling : A means of Providing Flexibility in Hard Real-Time Systems Systems. Technical report, Department of Computer Science, Univ. of York (UK), May 1994. Technical Report YCS230.

-
- [9] R. Davis, K. Tindell, and A. Burns. Scheduling slack time in fixed priority pre-emptive systems. In *Proceedings of the 14th IEEE Real-Time Systems Symposium*, pages 222–231, December 1993.
 - [10] R. Davis and A.J. Wellings. Dual priority scheduling. In *Proceedings of the 16th IEEE Real-Time Systems Symposium*, pages 100–109, December 1995.
 - [11] Association Française de Normalisation (AFNOR). Véhicules routiers - transmission de données, 1990. document R13-708.
 - [12] B. Gaujal and N. Navet. Traffic shaping in real-time distributed systems: a low-complexity approach. *Computer Communications*, accepted for publication, to appear in 1999. Also available as INRIA Research Report RR-3719.
 - [13] L. George, N. Rivierre, and M. Spuri. Preemptive and non-reemptive real-time uni-processor scheduling. Technical Report 2966, INRIA, september 1996.
 - [14] International Standard Organization ISO. *Road Vehicles - Low Speed serial data communication - Part 2: Low Speed Controller Area Network*. ISO, 1994. ISO 11519-2.
 - [15] M. Joseph and P. Pandya. Finding response times in a real-time system. *The Computer Journal*, 29(5):390–395, 1986.
 - [16] S. Lefebvre-Barbaroux, A. Jean-Marie, and C. Chaouiya. Problèmes d'ordonnancement d'une superposition de flux periodiques sous contraintes temps-réel. Rapport de Recherche 1576, INRIA, January 1992.
 - [17] J. Lehoczky and S. Ramos-Thuel. Aperiodic Task Scheduling for Hard-Real-Time Systems. In *Proceedings IEEE Real-Time Systems*, pages 110–123, December 1992.
 - [18] J. Lehoczky, L. Sha, and Y. Ding. Enhancing aperiodic responsiveness in hard real-time environment. In *Proceedings IEEE Real-Time Systems*, pages 261–270, December 1987.
 - [19] J.M. Migge and A. Jean-Marie. Timing analysis of real-time scheduling policies: a trajectory based model. Technical Report 3561, INRIA, 1998.
 - [20] J.M. Migge and A. Jean-Marie. Non-preemption, critical sections and round robin. Technical Report 3678, INRIA, 1999.

-
- [21] N. Navet and Y.-Q. Song. Reliability improvement of the dual-priority protocol under unreliable transmission. *Control Engineering Practice*, accepted for publication, to appear in 1999.
 - [22] N. Navet, Y.-Q. Song, and F. Simonot. Worst-case deadline failure probability in real-time applications distributed over CAN (Controller Area Network). *Journal of Systems Architecture*, accepted for publication, to appear in 1999.
 - [23] M.-A. Peraldi and J.-D. Decotignie. Combining real-time features of local area network fip and can. Technical report, École Polytechnique de Lausanne, 1996.
 - [24] F. Riera. *Monstre: A Microkernel orientat a les necessitats de sistemes de temps real empotrats*. Final year project, Universitat de les Illes Balears, Departament de Ciències Matemàtiques i Informàtica, 1997.
 - [25] Society of Automotive Engineers SAE. Class c application requirement considerations. Technical report, Society of Automotive Engineers, Jun 1993. J2056/1.
 - [26] SAE Vehicle Network for Multiplexing and Data Communications Standards Committee. Class b data communications network interface - sae j1850 standard, 1996. Rev. NOV96.
 - [27] S.P. Sprunt, L. Sha, and J. Lehoczky. Aperiodic Task Scheduling for Hard-Real-Time Systems. *Real-Time Systems*, 1(1):27–60, 1989.
 - [28] K. Tindell and H. Hansson. Babbling Idiots, The Dual Priority Protocol, and Smart CAN Controllers. In *2nd international CAN Conference, ICC'95*, pages 7.22–7–28, Septembre 1995.
 - [29] K. Zuberi and K.G. Shin. Scheduling on controller area network for real-time cim applications. *IEEE Transactions on Robotics and Automation*, 13(2):310–314, April 1997.



Unit e de recherche INRIA Lorraine, Technop ole de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS L ES NANCY
Unit e de recherche INRIA Rennes, Irisa, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unit e de recherche INRIA Rh one-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unit e de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unit e de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

 diteur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399