

Schedulability Analysis for CAN-based Networked Control Systems with Dynamic Bandwidth Management

Manel Velasco, Pau Martí, José Yépez, Ricard Villà and Josep M. Fuertes
Automatic Control Department
Technical University of Catalonia

{manel.velasco,pau.marti,jose.yepeze,ricard.villa,josep.m.fuertes}@upc.edu

Abstract

This paper presents the schedulability analysis for control messages when networked control loops, built on top of the Controller Area Network (CAN), are dynamically allocating bandwidth in terms of their controlled plants' dynamics. The bandwidth allocation policy is theoretically described by an optimization problem and practically solved by the distributed bitwise arbitration of CAN messages when message identifiers, i.e., priorities, reflect control applications demands. This poses the problem of assessing whether the set of real-time messages will meet their deadlines regardless of run-time priority changes. This is solved by a schedulability analysis based on recent results on worst-case response time techniques for real-time CAN applications. The analysis ends up with the schedulability test for this type of applications.

1. Introduction

For control applications where several control loops share limited computing resources, the problem of efficiently managing resources with respect to control performance has been recently achieved by resource allocation (RA) policies. Most of the existing work on RA concentrates on performance optimization in processor-based embedded control systems, e.g. [1–6]. In short, most of these results suggest methods to efficiently select the controllers' sampling periods. Formally, many of them are based on specifying and solving at run-time an optimization problem that relates control performance and CPU utilization. In practice, a resource manager, that is periodically executed, knowing the current plant states, assigns new sampling periods such that a given cost function is optimized.

Looking at maximizing control performance under communication bandwidth constraints when several control loops are closed through a real-time network, i.e., networked control system (NCS, [7]), different approaches have been presented [8–14]. However, in NCS, a key aspect has to be considered due to the inherent distributed architecture. Different than CPU-based RA approaches,

in NCS, the information required such as plants' current states is physically distributed. As a consequence, a resource manager can not easily perform the run-time optimization and assign new sampling periods.

The previous observation explains why many of the approaches trading off control performance and bandwidth for NCS provide off-line solutions which may be enhanced by heuristic run-time adaptations (e.g., [9] or [11]). Other solutions, where run-time centralized optimization is carried out, skip the implementation problems posed by the distributed architecture, such as in [10]. On the contrary, the run-time RA approaches presented in [8], [12], [13] and [14] take into account the distributed nature of NCS. However, none of them addresses schedulability analysis aspects.

In particular, [14] presented a distributed resource allocation approach targeting control performance optimization for several control loops built on top of the Controller Area Network (CAN) [15]. The approach is theoretically described by an optimization problem and practically solved by the distributed bitwise arbitration of CAN messages when message identifiers, i.e., priorities, reflect run-time control applications demands. The solution, that applies to single-input/single-output networked control systems, requires using an application profile that mandates to specify types of nodes (sensors, controllers and actuators) and defines how messages identifiers should be specified with respect to the controlled plant dynamics, which ends up with the definition of four message classes.

To structure messages in classes is not a new approach in CAN. For example, it was done in the mixed traffic scheduling scheme [16] or more recently in the hybrid priority scheduling schemes [17] and [18]. In [16] message classes were defined to improve schedulable utilization while [14] targets to efficiently use the available schedulable utilization in order to improve control applications' performance. In [17] and [18] different messages classes were defined by hybrid priority schemes with the objective of improving control loops stability performances or minimizing bandwidth occupancy. Their approach can be considered as complementary to the profile defined in [14].

The application of the profile defined by [14] provokes

run-time changes in message priorities because control demands evolve at run-time. This poses the problem of assessing whether the set of real-time messages will meet their deadlines regardless of run-time priority changes. The contribution of this paper, which extends [14], is to present the schedulability analysis for these systems. The schedulability analysis is based on recent results on worst-case response time techniques for real-time CAN applications [19], but with the property that messages are divided into classes, and that message priorities (within a class) can change at run-time. The analysis ends up with the schedulability test for this type of applications. The reader is referred to [20] for an extended research report.

The rest of this paper is organized as follows. Section 2 summarizes previous results [14] in order to provide the appropriated background. Section 3 sets the preliminaries for the schedulability analysis, which is developed in Section 4. Section 5 presents the schedulability test and extends its applicability. Section 6 concludes the paper.

2. Background

The networked architecture considered in this work is the following: each controller and plant constitute a control loop closed over the network, and each controller is assumed to be implemented in separate nodes. That is, sampling, control algorithm computation and actuation are performed in separated nodes. These nodes exchange sample and control signal informations through the network shared by all nodes of all control loops. Therefore, network bandwidth is the scarce resource that must be allocated among the control loops.

2.1. RA problem, solution and practical considerations

The problem to be solved is how to assign the scarce network bandwidth to the set of control loops such that all messages are schedulable and overall control performance is maximized, knowing that the controllers will provide better performance when given more network bandwidth. The key assumption is that all controllers can not run, simultaneously, at their highest possible sampling frequency due to bandwidth limitations, so they cannot provide the best possible control performance. In [14] it is shown that the problem to be solved can be formulated as a constrained optimization problem whose solution dictates that, after a minimum bandwidth share is guaranteed to each control loop, the remaining available bandwidth should be assigned to the control loop with largest error (possibly weighted), where the error is defined as a function of the plant state. Note that a uniform error model is required to use this solution.

Therefore, the optimal solution requires to compute a function, *largest error*, that takes parameters, *errors*, from different nodes. This function has to be computed efficiently and with a small (and bounded) number of messages. In fact, since control systems are extremely sensitive to timing variations, it is required that the time spent

on computing and executing the bandwidth reallocation solution should be negligible with respect the dynamics of the set of controlled plants. Therefore, it would be desirable to have a solution with a time-complexity that does not depend on the number of nodes.

Depending on the type of communication protocol considered in the implementation, the previous requirement may not be met. In [14] it is discussed that the optimal policy could be achieved using centralized protocols. However, their primary limitation is that they assume the existence of a master node that knows the current state of all controlled plants. Maintaining a global state can be problematic in practice, increasing network traffic and latencies, and introducing a central point of failure.

Without assuming that all the information required to solve the allocation problem is localized in a single node, a different solution is required. The key idea in [14] is to observe that CAN permits to schedule messages on a priority based semantics, and that a feasible implementation of the optimal bandwidth allocation policy could take advantage of this property. Intuitively, for the distributed implementation of the optimal solution, [14] proposes to encode at run-time plants' errors into each message identifier in order to have the solution implemented at the CAN bitwise arbitration resolution. Therefore, structural properties of CAN permit to efficiently implement the optimal policy because calculating the largest error of all plants reduces to resolve the arbitration and to observe its winner.

2.2. Application profile

The previous intuitive solution is achieved in [14] by defining an application profile for CAN extended frames (known as CAN 2.0 B) but the extension to CAN base frames (known as CAN 2.0 A) is straightforward.

The application profile allows the following functionalities. First, control signal messages are given the highest priority in order to be guaranteed and thus minimize latencies in each control loop operation. Also, sampler nodes have to be assigned a minimum bandwidth share (periodicity) to guarantee stability. Secondly, non-control messages behave as expected in terms of being granted bus access once control messages and periodic sampler messages have been guaranteed. Finally, the sampler node with highest error is assigned the bandwidth left over. These functionalities are achieved by deploying the set of networked control loops using different type of nodes and messages classes. Four types of nodes are defined:

- **Sampler nodes:** they sample the plant at the maximum rate (given by the sensor hardware limitations), compute the plant error, and try to send this information (plant error and plant state) over the network.
- **Controller nodes:** each controller node, after receiving a sampler message, computes (using the plant state) and sends the control signal to the corresponding actuator.

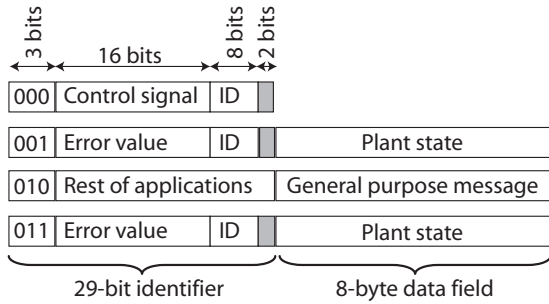


Figure 1. Bit encoding for message classes.

- **Actuator nodes:** they apply the control signal to the plant upon reception of the control signal message.
- **Other nodes:** they do other non-control activities.

Four messages classes are defined, characterized by the first 3 bits of the 29 bits message identifier (Figure 1):

- **000 (control messages):** This class contains the highest priority messages, which are always granted to win bus access. These messages will be used by controller nodes to send control signal messages to actuator nodes. They are only send upon reception of a sampler message.
- **001 (granted sensor messages):** This class contains the second higher priority messages, which have to be also always granted in the bus. These messages, with a guaranteed minimum periodicity, will be used by sensor nodes to send plant states to controller nodes. Apart from the first 3 bits, the following 16 bits will encode the 1's complement of the plant error in such a way that the sample node with highest error will always win access to the bus in front of the other sampler nodes in the case of collisions.
- **010 (general purpose messages):** This class contains all non-control messages and their identifier encoding should be as thought originally, but with 010 in the first 3 bits.
- **011 (best effort sensor messages):** This class contains the lowest priority messages, and they are not guaranteed. These messages will be used by all sampler nodes to send as many messages as possible. Apart from the first 3 bits, the following 16 bits will code the plant error like messages of class 001. By doing so, the bandwidth left over will be re-assigned at run-time to the node with highest error.

The four traffic classes are distinguished by the first 3 bits of the message identifier. Although 2 bits would suffice (as originally presented in [14]), practical considerations prevent using only 2 bits. CAN does not allow to

have the most significant 7 bits recessive, i.e., 1. Therefore, if only 2 bits were used to code the class types, the last class would be distinguished by 11. This code together with the error value (see Figure 1) could result in 7 recessive bits. To avoid this problem, 3 bits are used.

By analyzing the CAN arbitration applied to all possible combinations of collisions of all messages of all classes, it is analytically shown and experimentally validated in [14] that the desired functionality mandated by the optimal policy is fully achieved [20].

In [14] it is observed that after sampler messages (classes 001 and 011) there will always be in the bus control messages (class 000). This precedence relation is said to hold because it is assumed that the time spent by the controller node to both compute the control signal and place the message in a transmitter buffer is less than the inter-frame time (3 bits time), which is a realistic assumption taking into account standard CAN bit rates and current microprocessor's speed. The schedulability analysis presented next takes this assumption into account. However, later on, this assumption is removed to extend the analysis to all possible cases.

3. Schedulability analysis: preliminaries

This section motivates the need for the schedulability analysis, shortly reviews the latest results on worst-case response time analysis, and presents convenient notation.

3.1. Motivating example

Figure 2 shows an example of how the CAN traffic is affected when the application profile is applied. The top subfigure presents the traffic of CAN in standard operation. The bottom subfigure shows the traffic of CAN when using the profile. In both subfigures, general purpose traffic (dashed boxes, class 010) share the network with those messages required in the operation of two control loops, where black boxes represent control messages (class 000), white boxes represent granted sensor messages (class 001) and gray boxes represent best effort sensor messages (class 011) that only appear when the profile is used (bottom subfigure).

In the top subfigure, pairs of white and black boxes in the bottom correspond to the sampling and control messages of one control loop. Similarly, pairs of white and black boxes in the middle of the top subfigure correspond to the sampling and control messages of the second control loop. As it can be seen, they are periodically triggered, and they have the same periodicity. When using the application profile, bottom subfigure, these pairs still appear. But in addition, more control messages (black boxes) are sent because they are triggered by best effort sensor messages (gray boxes). In fact, the new pairs of gray and black boxes apply to either control loop depending on which loop is experiencing the largest error. As it can be seen, the largest error is mainly being detected in the control loop whose messaging is in the middle of the




Figure 2. CAN traffic with (bottom subfigure) or without (top subfigure) the application profile.

bottom subfigure.

The use of the application profile slightly alters the transmission timing and ordering of the set of messages. Two main alterations can be observed:

- **Permutations:** By comparing the top and bottom subfigures of Figure 2, it can be observed that two pairs of granted sensor/control messages (white/black boxes) of the top subfigure right after time 0.02s affecting both control loops permute in time in the bottom subfigure due to a priority change. This is due to the fact that, using the application profile, granted sensor messages competing for the network win access to CAN depending on the errors that the sensed plants are experiencing.
- **Delayed transmissions:** In addition, the use of the application profile may delay pairs of granted sensor/control messages or general purpose messages. In particular, as it can be seen in Figure 2, the pair of granted sensor/control messages sent at time 0.04s in the top subfigure are delayed in the bottom subfigure due to insertion in the bus of pairs of best effort sensor/control messages. For the same reason, the third general purpose message (dashed box) after time 0.02s starts later in the bottom subfigure.

Hence, comparing the standard CAN traffic and the resulting traffic when using the application profile, some anomalies can be detected and must be formally analyzed.

3.2. CAN worst-case response time analysis revisited

Response time analysis for CAN aims to provide a method of calculating the worst-case response time of each message. This value can be compared to the message deadlines to determine if the system is schedulable.

From [19], the worst-case response time of a message i on CAN without using the application profile is given by

$$R_i = J_i + W_i + C_i \quad (1)$$

where J_i is the queuing jitter, W_i is the queuing delay and C_i is the transmission time. Parameter J_i corresponds to the longest time between the initiating event and the message being queued and ready to be transmitted. Henceforth, and without losing generality, we will omit J_i because it is not affected by the application of the profile.

Parameter W_i corresponds to the longest time that the message can remain in the queue before commencing successful transmission. It comprises two elements: blocking time B_i due to messages which may be in the process of being transmitted and interference time due to higher priority messages which may win arbitration and be transmitted in preference to message i . In [19] it is shown that

$$W_i = B_i + \sum_{j \in hp(i)} \left\lceil \frac{W_i + t_{bit}}{T_j} \right\rceil C_j \quad (2)$$

where $hp(i)$ is the set of messages with priorities higher than i , $\lceil a/b \rceil$ is notation for the *ceiling* function which returns the smallest integer greater than or equal to a/b , t_{bit} is the bit time, and T_j is the period of the j^{th} message. Equation (2) is solved using the recurrence relation

$$W_i^{k+1} = B_i + \sum_{j \in hp(i)} \left\lceil \frac{W_i^k + t_{bit}}{T_j} \right\rceil C_j \quad (3)$$

with $W_i^0 = B_i$ as a suitable initial condition. Denoting by D_i the i^{th} -message deadline, (3) iterates until, either

$$R_i^{k+1} = W_i^{k+1} + C_i > D_i \quad (4)$$

in which case the message is not schedulable, or

$$W_i^{k+1} = W_i^k \quad (5)$$

in which case the message i is schedulable provided

$$B_i = B^{MAX} \quad \text{and} \quad R_i^{k+1} = W_i^{k+1} + C_i \leq D_i. \quad (6)$$

The later condition has been shown (through equation (17) in [19]) to be a sufficient condition for schedulability test of the message i , where B^{MAX} is the blocking time that is equal to the transmission time of the longest possible CAN message (8 bytes), irrespective of the characteristics and priorities of the messages in the system.

The worst-case response time of the i^{th} message is given at the end of the recurrence iteration by

$$R_i^{k+1} = W_i^{k+1} + C_i = B^{MAX} + \sum_{j \in hp(i)} \left\lceil \frac{W_i^k + t_{bit}}{T_j} \right\rceil C_j + C_i. \quad (7)$$

Eq. (7) holds for deadlines no greater than periods [19].

3.3 Notation

The following notation is introduced:

- Ω denotes the set of messages to be scheduled.
- Ω_0 denotes the subset of control messages: $\Omega_0 = \{\tau_i \in \Omega : class(\tau_i) = 000\}$
- Ω_1 denotes the subset of granted sampler messages: $\Omega_1 = \{\tau_i \in \Omega : class(\tau_i) = 001\}$
- Ω_2 denotes the subset of general purpose messages: $\Omega_2 = \{\tau_i \in \Omega : class(\tau_i) = 010\}$
- Ω_3 denotes the subset of best effort sampler messages: $\Omega_3 = \{\tau_i \in \Omega : class(\tau_i) = 011\}$

If each granted sampler message or each best effort sampler message is followed by a control message, we can simplify the notation by joining Ω_1 or Ω_3 with Ω_0 as

- Ω_1^+ denotes the subset of join messages of Ω_0 and Ω_1 : $\Omega_1^+ = \{\tau_i \in \Omega : class(\tau_i) = 001 \text{ with } C_i = C_i^{001} + t_{interframe} + C_i^{000}\}$. That is, each i^{th} -message in Ω_1^+ has a transmission time that includes the transmission time of the i^{th} -granted sampler message, the interframe time, and the transmission time of the i^{th} -control message.
- Ω_3^+ denotes the subset of join messages of Ω_0 and Ω_3 : $\Omega_3^+ = \{\tau_i \in \Omega : class(\tau_i) = 011 \text{ with } C_i = C_i^{011} + t_{interframe} + C_i^{000}\}$. That is, each i^{th} -message in Ω_3^+ has a transmission time that includes the transmission time of the i^{th} -best effort sampler message, the interframe time and the transmission time of the i^{th} -control message.

4. Schedulability analysis: case study

The schedulability analysis explores all the relevant cases that may happen due to the dynamic priority changes. The controlled plant dynamics causes priority changes in messages belonging to Ω_1^+ . Therefore, we have to analyze how these changes affect the worst case response time of messages belonging to Ω_1^+ (*Case 1*), to Ω_2 (*Case 2*), and to Ω_3 (*Case 3*). For the sake of completeness of the analysis, we also have to analyze how the worst case response time of messages belonging to $\Omega_1^+ \cup \Omega_2$ is affected when best effort messages are introduced. This analysis demands a slight modification of (7), announced as *Case 0* and it is fully covered in *Case 4*.

4.1. Case 0

The worst case response time of a message was given in (7) where B^{MAX} is the blocking time corresponding to the transmission time of the longest possible CAN message. The application profile introduces new traffic: new pairs of best effort sensor/control messages appear in the network. Hence, (7) should also include the case that any

message willing to enter the bus may be blocked by one sequence of these new pairs of messages. Formally, the recurrence when using the application profile should be

$$R_i^{k+1} = W_i^{k+1} + C_i = B^{MAX^+} + \sum_{j \in hp(i)} \left\lceil \frac{W_i^k + t_{bit}}{T_j} \right\rceil C_j + C_i \quad (8)$$

where $B^{MAX^+} = \max(B^{MAX}, \text{longest}(\tau_i) \in \Omega_3^+)$.

The situation analyzed in this case provokes delayed transmissions, as previously illustrated in Figure 2.

Observation 1 *When using the application profile, the worst case response time analysis of any message has to be performed using (8).*

4.2. Case 1

This case studies the worst case response time of message $\tau_k \in \Omega_1^+$ when a priority change occurs between two messages of Ω_1^+ . This case can be alternatively announced as follows. If R_k^{n+1} is the solution of (8) for message $\tau_k \in \Omega_1^+$ when the priority of $\tau_i \in \Omega_1^+$ is higher than the priority of $\tau_j \in \Omega_1^+$, then R_k^{n+1} may not be the solution of (8) for message τ_k when the priority of τ_j is higher than the priority of τ_i .

Proposition 1 *The priority change between any two messages of Ω_1^+ can alter the worst case response time of a message of Ω_1^+ .*

Proof. The possible priority changes with respect to a message $\tau_k \in \Omega_1^+$ are:

1. $\tau_i \in hp(k)$ and $\tau_j \in hp(k)$ (the two messages that permute their priorities have higher priorities than τ_k): For this case, the worst case response time R_k^{n+1} of τ_k is not altered by the priority changes. R_k^{n+1} is

$$R_k^{n+1} = B^{MAX^+} + \sum_{j \in hp(k)} \left\lceil \frac{W_k^n + t_{bit}}{T_j} \right\rceil C_j + C_k. \quad (9)$$

Knowing that $hp(k)$ is the set of higher priority messages than τ_k , it holds that $hp(k)_{\text{original}} = hp(k)_{\text{permuted}}$. Therefore, to alter the order of two factors of the summand in (9) do not affect the result.

2. $\tau_i \notin hp(k)$ and $\tau_j \notin hp(k)$ (the two messages that permute their priorities have lower priorities than τ_k): Since there are no modifications in $hp(k)$, the worst case response time of τ_k remains the same.
3. $\tau_i \notin hp(k)$ and $\tau_j \in hp(k)$ (from the two messages that permute their priorities, one has higher priority than τ_k): We have to consider two scenarios. The first scenario consists in considering that all messages of Ω_1^+ always send 8 data bytes. This implies that all messages are equal in terms of transmission time if $T_i = T_j$ (their periods are the same), which implies that we can not make a difference between τ_i

and τ_j . In this case, the worst case response time is not altered because

$$\left\lceil \frac{W_k + t_{bit}}{T_i} \right\rceil C_i = \left\lceil \frac{W_k + t_{bit}}{T_j} \right\rceil C_j \quad (10)$$

The second scenario consists in considering that messages of Ω_1^+ do not send the same number of bytes or that $T_i = T_j$. Therefore, condition (10) does not hold anymore and the worst case response time of τ_k can be altered by the permutation.

4. $\tau_i \notin hp(k)$ and $\tau_j = \tau_k$ (one of the two messages that permute their priorities is τ_k , and the other one has a lower priority): In this case the number of summands of (9) before the permutation is less than the number that we can have after the permutation, which implies that the worst case response time of τ_k can drastically increase. The only situation in which we have guarantees that the permutation will benefit the worst case response time of τ_k is when τ_k is the lowest priority message of Ω_1^+ right before the permutation. In this case, any permutation provokes a reduction of its worst case response time.
5. $\tau_i \in hp(k)$ and $\tau_j = \tau_k$ (one of the two messages that permute their priorities is τ_k , and the other one has a higher priority): In this case, then number of summands of (9) before the permutation is bigger than the number that we can have after the permutation, which implies that the worst case response time of τ_k always decreases.

□

The situation given by Proposition 1 was illustrated in Figure 2, where it can be seen that pairs of granted sensor and control messages are permuted at run time. And this actually affects their worst case response time.

Observation 2 *Proposition 1 implies that schedulability of Ω is affected by messages of Ω_1^+ when priority changes occurs between messages of Ω_1^+ . Then, in order to guarantee schedulability while accounting for priority changes, the worst case response time analysis of all k^{th} messages of Ω_1^+ has to be performed considering that $hp(k) = \Omega_1^+ - \{\tau_k\}$, that is, considering that all messages of Ω_1^+ are of higher priority than the message under analysis, τ_k .*

4.3. Case 2 and Case 3

These cases share a common fact: both look at how the worst case response time of messages belonging to lower priority classes (Ω_2 or Ω_3) than messages of Ω_1^+ would be affected with respect to permutations of messages of Ω_1^+ . For this reason the rest of this subsection focuses on Case 2. The study of Case 3 would follow exactly the same argumentation.

Hence, this case studies the change on the worst case response time of message $\tau_k \in \Omega_2$ when a priority change occurs between two messages of Ω_1^+ , which can be alternatively announced as follows. If R_k^{n+1} is the solution of (8) for message $\tau_k \in \Omega_2$ when the priority of $\tau_i \in \Omega_1^+$ is higher than the priority of $\tau_j \in \Omega_1^+$, then R_k^{n+1} is also the solution of (8) for message τ_k when the priority of τ_j is higher than the priority of τ_i .

Proposition 2 *The priority change between any two messages of Ω_1^+ does not alter the worst case response time of any message of Ω_2 .*

Proof. The worst case response time R_k^{n+1} of $\tau_k \in \Omega_2$ is

$$R_k^{n+1} = B^{MAX^+} + \sum_{j \in hp(k)} \left\lceil \frac{W_k^n + t_{bit}}{T_j} \right\rceil C_j + C_k. \quad (11)$$

Knowing that $hp(k)$ is the set of higher priority messages of τ_k , it holds that $\Omega_1^+ \subseteq hp(k)$. Therefore, to alter the order of two factors of the summand in (11) do not affect the result. □

As it is illustrated in the example shown in Figure 2, the first general purpose message after 0.002 is not altered by the permutation occurring right before.

Observation 3 *Proposition 2 implies that schedulability of Ω is not affected by messages of Ω_2 when priority changes occurs between messages of Ω_1^+ . Therefore, to check schedulability we only require to evaluate the worst case response time of all messages of Ω_2 in only one scenario of priorities.*

4.4. Case 4

This case studies the change on the worst case response time of message $\tau_k \in \Omega_1^+ \cup \Omega_2$ when messages of Ω_3 are introduced in the bus. This case can be alternatively announced as follows. If R_k^{n+1} is the solution of (8) for message $\tau_k \in \Omega_1^+ \cup \Omega_2$ when considering that $\Omega = \Omega_1^+ \cup \Omega_2$, then R_k^{n+1} is also the solution of (8) for message τ_k when $\Omega = \Omega_1^+ \cup \Omega_2 \cup \Omega_3$.

Proposition 3 *The priority change between any two messages of Ω_3 does not alter the worst case response time of any message of $\Omega_1^+ \cup \Omega_2$. The same happens if new messages of Ω_3 appear in the bus or when messages of Ω_3 terminate.*

Proof. The worst case response time R_k^{n+1} of any message $\tau_k \in \Omega_1^+ \cup \Omega_2$ is given by

$$R_k^{n+1} = B^{MAX^+} + \sum_{j \in hp(k)} \left\lceil \frac{W_k^n + t_{bit}}{T_j} \right\rceil C_j + C_k. \quad (12)$$

Knowing that $hp(k)$ is the set of higher priority messages of τ_k , R_k^{n+1} is not altered if two messages $\tau_i, \tau_j \in$

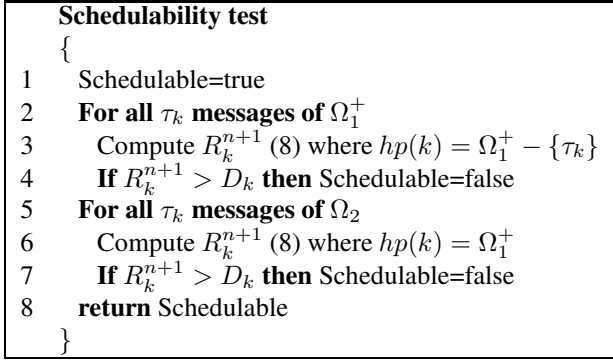


Figure 3. Schedulability test.

Ω_3 change their priorities because $\Omega_3 \notin hp(k)$. The same holds regardless of whether the number of messages in Ω_3 increases or decreases. \square

As it was illustrated in the example shown in Figure 2, the insertion of best effort messages could delay the transmission of general purpose messages. But their worst-case response time remains the same.

Observation 4 Messages of Ω_3 are not guaranteed in the bus by definition of the application profile because they are considered best effort messages, with the lowest set of priorities. Therefore they have to be discarded in the schedulability analysis.

5. Schedulability test

To guarantee that all messages of CAN-based control applications enabled with the presented profile will meet their deadlines, the schedulability test shown in Figure 3 has to be applied. It is built upon observations 1, 2, 3 and 4. Observation 1 is included in lines 3 and 6. Observation 2 is reflected in lines 2-4. Observation 3 is reflected in lines 5-7. The last observation is also reflected in the test, in the sense that since best effort messages (Ω_3) are not guaranteed, they do not need to be analyzed for their worst case response time.

Note that the schedulability test (Figure 3) is based on the previous case study (Section 4) that addressed the schedulability in the final states of the priority changes, not “during” the priority changes. Since priority changes give rise to mode changes, the schedulability analysis should also include mode changes. However, observation 2 already imposed a pessimistic rule in the schedulability test that includes the effect that mode changes would have in the worst case response time of messages of class Ω_1^+ .

The new schedulability test restricts message set schedulability with respect to the standard schedulability test presented in [19]. In particular, the restriction affects only granted sensor and control messages when accounting for the priority changes. That is, a given specification of timing constraints for these messages may result in a feasible schedule according to [19] but unfeasible ac-

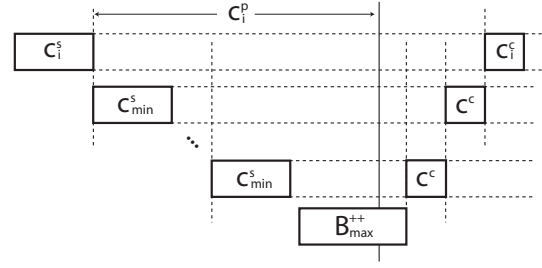


Figure 4. Analysis of control node execution time being longer than the interframe time.

ording to the presented schedulability test. In any case, priority changes do not affect general purpose message feasibility (recall observation 3).

The apparent limitation given by the feasibility restriction is overcome by the improvement on control performance that is achieved when using the application profile (see [14]). It is left for future work to assess the trade off between the achievable benefits in terms of control performance vs. the schedulability restrictions that are imposed.

5.1. Extension

The schedulability analysis and the derived test presented in Sections 4 and 5 assumed that granted sampler messages or best effort sampler messages are always followed by a control message (messages of classes Ω_1^+ and Ω_3^+). The reason for this assumption is that the time spent by the i^{th} controller node upon reception of any sampler message to both compute the control signal and place the message in a transmitter buffer, namely C_i^P , is less than the inter-frame time, t_{if} . Therefore, any other message of class Ω_1 , Ω_2 or Ω_3 ready to be transmitted, will not win bus access until the control message has been sent.

This assumption is realistic because at the faster CAN bit rate, the inter-frame time is 3ms. And for a standard microcontroller with a standard clock, e.g. for a PIC18F458 with a 20MHz oscillator, the execution time of a control algorithm is of the order of 10 μ s. Hence, the controller computation time C_i^P is by far shorter than the interframe time. However, in some implementations, the assumption may not hold and further analysis is required. In this case, when $C_i^P > t_{if}$, the transmission time of the messages belonging to classes Ω_1^+ and Ω_3^+ can no be computed as defined in Section 3.3.

The transmission time C_i of these messages can be longer. First, C_i has to include the time spent for transmitting the sensor message ($C_i^S = C_i^{001}$ or C_i^{011} for granted or best effort messages) and the control message (C_i^{000}). Note that according to the message classes definition, all control messages take the same transmission time, namely C^C , while C_i^S vary. Let C_{min}^S denote the the transmission time of a the shortest sensor message in Ω .

Second, C_i has to include the controller node processing time C_i^P , and all time spent by other messages that can be transmitted within the node processing time and

that may delay the transmission of the control message. Figure 4 illustrates this case in the worst scenario. The top boxes in the left and right side represent the i^{th} sensor message and corresponding control message. The controller node processing time C_i^P is also represented. Although the i^{th} control message, C_i^C , is ready for transmission after C_i^P , it is blocked by B^{MAX++} (in the bottom of the subfigure) and by several control messages C^C generated by sensor messages transmitted during C_i^P . The more sensor message that enter during C_i^P , the more control messages C^C that can delay C_i^C . The worst case scenario is given when all the control messages have C_{min}^S . And B^{MAX++} is the maximum among the transmission time of the longest general purpose message or the longest time spent by the messaging of a j^{th} closed loop operation that includes sensor message, processing time and control message. Formally, being $C_j = C_j^S + C_j^P + C_j^C$,

$$B^{MAX++} = \max(B^{MAX}, \text{longest}(\tau_j) \in \Omega_1^+ \cup \Omega_3^+). \quad (13)$$

By putting all together, the transmission time C_i of any message in Ω_1^+ and Ω_3^+ is given by

$$C_i = C_i^S + C_i^P + B^{MAX++} + \left\lfloor \frac{C_i^P}{C_{min}^S} \right\rfloor C^C + C_i^C \quad (14)$$

where $\lfloor a/b \rfloor$ is the notation for the *floor* function which returns the biggest integer smaller than or equal to a/b .

By updating the schedulability analysis and schedulability test in Sections 4 and 5 using (14), the case of having control nodes with a longer processing time than the interframe time is also covered.

6. Conclusions

This paper has presented a complete schedulability analysis of an application profile for networked embedded control systems that allows to implement an optimal bandwidth allocation policy among networked control loops using the CAN arbitration. From the analysis, a pessimistic schedulability test has been derived.

Acknowledgments

This work was partially supported by C3DE DPI2007-61527, and by ArtistDesign NoE IST-2008-214373.

References

[1] D. Seto, J. P. Lehoczky, L. Sha, and K. Shin, "On task schedulability in real-time control systems," *17th IEEE Real-Time Systems Symposium*, pp. 13-21, Dec. 1996.

[2] P. Martí, C. Lin, S. Brandt, M. Velasco and J.M. Fuertes, "Optimal State Feedback Based Resource Allocation for Resource-Constrained Control Tasks", *25th IEEE Real-Time Systems Symposium*, Dec. 2004.

[3] D. Henriksson and A. Cervin, "Optimal On-line Sampling Period Assignment for Real-Time Control Tasks Based on

Plant State Information.", *44th IEEE Conference on Decision and Control and European Control Conference ECC 2005*, Dec. 2005.

[4] E. Bini and A. Cervin, "Delay-Aware Period Assignment in Control Systems.", *29th IEEE Real-Time Systems Symposium*, December 2008.

[5] P. Martí, C. Lin, S.A. Brandt, M. Velasco and Josep M. Fuertes, "Draco: Efficient resource Management for Resource-Constrained Control tasks.", *IEEE Transactions on Computers*, vol. 58, n. 1, Jan. 2009.

[6] M. El Mongi Ben Gaid, A. Çela and Y. Hamam, "Optimal Real-Time Scheduling of Control Tasks With State-Feedback Resource Allocation," *IEEE Transactions on Control Systems Technology*, 2009.

[7] Y. Tipsuwan and M. Y. Chow, "Control methodologies in networked control systems", *Control Engineering Practice*, Vol. 11, pp. 1099-1111, 2003

[8] G.C. Walsh, and Y. Hong, "Scheduling of networked control systems," *IEEE Control Systems Magazine*, vol. 21, n. 1, pp. 57-65, Feb. 2001

[9] D. Hristu-Varsakelis, "Feedback control systems as users of a shared network:communication sequences that guarantee stability", *IEEE Conf. on Decision and Control*, 2001

[10] M.S. Branicky, S.M. Phillips, and W. Zhang "Scheduling and Feedback Co-Design for Networked Control Systems", *IEEE Conference on Decision and Control*, 2002.

[11] H. Rehbinder, and M. Sanfridson, "Scheduling of a limited communication channel for optimal control", *Automatica*, Vol. 30, No. 3, pp. 491-500, March 2004.

[12] M. Velasco, J.M. Fuertes, C. Lin, P. Martí, and S. Brandt, "A Control Approach to Bandwidth Management in Networked Control Systems.", *30th Annual Conf. of the IEEE Industrial Electronics Society*, Nov. 2004.

[13] Al-Hammouri, A.T.; Branicky, M.S.; Liberatore, V.; Phillips, S.M., "Decentralized and dynamic bandwidth allocation in networked control systems," In 20th Int. Parallel and Distributed Processing Symposium, April 2006.

[14] M. Velasco, P. Martí, R. Castañé, J. Guardia and J.M. Fuertes, "A CAN Application Profile for Control Optimization in Networked Embedded Systems", *32th Annual Conf. of the IEEE Industrial Electronics Society*, Nov. 2006.

[15] CAN Specification version 2.0. Robert Bosch GmbH, Postfach 30 02 40, D-70442 Stuttgart, 1991.

[16] K.M. Zuberi, and K.G. Shin, "Design and Implementation of Efficient Message Scheduling for Controller Area Network", *IEEE Transactions on Computers*, vol. 49(2), pp. 182-188, Feb. 2000.

[17] G. Juanole and G. Mouney, "Networked Control Systems: Definition and Analysis of a Hybrid Priority Scheme for the Message Scheduling," *13th IEEE Int. Conf. on Embedded and Real-Time Computing Systems and Applications*, 2007.

[18] T. Bai and Z. Wu, "Hybrid Bandwidth Scheduling for CAN-based Networked Control Systems," *Acta Automatica Sinica*, vol. 33, n. 9, pp. 963-967, Sept. 2007.

[19] R.I. Davis, A. Burns, R. J. Bril, and J. J. Lukkien, "Controller Area Network (CAN) schedulability analysis: Refined, revisited and revised," *Real-Time Systems*, vol. 35, n. 3, pp. 239-272, Apr. 2007

[20] M. Velasco, P. Martí and Josep M. Fuertes, "Schedulability Analysis for CAN-based Control Applications with Dynamic Bandwidth Management," Research report ESAIL-RR-08-04, Nov. 2008. <http://www.upcnet.es/~pmc16/08RRCAN.pdf>