

Schedulability-Driven Communication Synthesis for Time Triggered Embedded Systems

Paul Pop, Petru Eles, and Zebo Peng

*Dept. of Computer and Information Science, Linköping University, Sweden
{paupo, petel, zebpe}@ida.liu.se*

Abstract

We present an approach to static priority preemptive process scheduling for the synthesis of hard real-time distributed embedded systems where communication plays an important role. The communication model is based on a time-triggered protocol. We have developed an analysis for the communication delays proposing four different message scheduling policies over a time-triggered communication channel. Optimization strategies for the synthesis of communication are developed, and the four approaches to message scheduling are compared using extensive experiments.

1. Introduction

Depending on the particular application, an embedded system has certain requirements on performance, cost, dependability, size, etc. For hard real-time applications the timing requirements are extremely important. Thus, in order to function correctly, an embedded system implementing such an application has to meet its deadlines. One of the typical application areas for such systems is that of safety-critical automotive applications (e.g. drive-by-wire, brake-by-wire) [17].

In this paper we concentrate on certain aspects concerning the synthesis of embedded hard real-time systems which are implemented on distributed architectures consisting of multiple programmable processors and ASICs. Process scheduling is based on a static priority preemptive approach while the bus communication is statically scheduled.

Process scheduling for performance estimation and synthesis of embedded systems has been intensively researched in the last years. The existing approaches differ in the scheduling strategy adopted, system architectures considered, handling of the communication, and process interaction aspects.

Static non-preemptive scheduling of a set of processes on a multiprocessor system has been discussed in [2, 4, 5, 10]. Preemptive scheduling of independent processes with static priorities running on single processor architectures has its roots in [7]. The approach has been later extended to accommodate more general computational models and has been also applied to distributed systems [15]. The reader is referred to [1] for a survey on this topic.

Although different scheduling strategies have been adapted to accommodate distributed architectures, researchers have often ignored or very much simplified aspects concerning the communication infrastructure. One typical approach is to consider communication processes as processes with a given execution time (depending on the amount of information exchanged) and to schedule them as

any other process, without considering issues like communication protocol, bus arbitration, packaging of messages, clock synchronization, etc.

Currently, more and more real-time systems are used in physically distributed environments and have to be implemented on distributed architectures in order to meet reliability, functional, and performance constraints. Thus, in order to guarantee that real-time characteristics are fulfilled analysis has been done for different communication protocols. The CAN bus is analyzed in [13], a 802.5-style token ring and a simple TDMA protocol are analyzed in [14], in [16] P-NET networks are investigated, while in [3] the ATM protocol is considered.

In this paper we consider the time-triggered protocol (TTP) [6] as the communication infrastructure for a distributed real-time system. Processes are scheduled according to a static priority preemptive policy. We first perform the schedulability analysis considering four different approaches to message scheduling. After this, we go one step further by showing how the parameters of the communication protocol can be optimized in order to fit the communication particularities of a certain application.

TTP has been classically associated with non-preemptive static scheduling of processes, mainly because of fault tolerance reasons [5]. In [10] we have addressed the issue of non-preemptive static process scheduling and communication synthesis using TTP.

However, considering preemptive priority based scheduling at the process level, with time triggered static scheduling at the communication level can be the right solution under certain circumstances [8]. A communication protocol like TTP provides a global time base, improves fault-tolerance and predictability. At the same time, certain particularities of the application or of the underlying real-time operating system can impose a priority based scheduling policy at the process level.

The paper is divided into 7 sections. The next section presents the architectures considered for system implementation. The computational model assumed and formulation of the problem are presented in section 3, and section 4 presents the schedulability analysis for each of the four approaches considered for message scheduling. The optimization strategy is presented in section 5, and the four approaches are evaluated in section 6. The last section presents our conclusions.

2. System Architecture

2.1 Hardware Architecture

We consider architectures consisting of nodes connected




Figure 1. System Architecture
 by a broadcast communication channel (Figure 1). Every node consists of a TTP controller, a CPU, a RAM, a ROM and an I/O interface to sensors and actuators. A node can also have an ASIC in order to accelerate parts of its functionality.

Communication between nodes is based on the TTP [6]. TTP was designed for distributed real-time applications that require predictability and reliability (e.g. drive-by-wire). It integrates services necessary for fault-tolerant real-time systems.

The communication channel is a broadcast channel, so a message sent by a node is received by all the other nodes. The bus access scheme is time-division multiple-access (TDMA) (Figure 2). Each node N_i can transmit only during a predetermined time interval, the so called TDMA slot S_i . In such a slot, a node can send several messages packaged in a frame. A sequence of slots corresponding to all the nodes in the architecture is called a TDMA round. A node can have only one slot in a TDMA round. Several TDMA rounds can be combined together in a cycle that is repeated periodically. The sequence and length of the slots are the same for all the TDMA rounds. However, the length and contents of the frames may differ.

Every node has a TTP controller that implements the protocol services, and runs independently of the node's CPU. Communication with the CPU is performed through a so called message base interface (MBI) which is usually implemented as a dual ported RAM (see Figure 3).

The TDMA access scheme is imposed by a so called message descriptor list (MEDL) that is located in every TTP controller. The MEDL basically contains: the time when a frame has to be sent or received, the address of the frame in the MBI and the length of the frame. MEDL serves as a schedule table for the TTP controller which has to know when to send or receive a frame to or from the communication channel.

The TTP controller provides each CPU with a timer interrupt based on a local clock, synchronized with the local clocks of the other nodes. The clock synchronization is done by comparing the a priori known time of arrival of a frame with the observed arrival time. By applying a clock synchronization algorithm, TTP provides a global time-base of known precision, without any overhead on the communication.




Figure 2. Bus Access Scheme

2.2 Software Architecture

We have designed a software architecture which runs on the CPU in each node, and which has a real-time kernel as its main component. We consider a time-triggered system, which means that there are no other interrupts except for the timer interrupt. Each kernel has a so called tick scheduler that is activated periodically by the timer interrupts and decides on activation of processes, based on their priorities. Several activities, like polling of the I/O or diagnostics, take also place in the timer interrupt routine.

In order to run a predictable hard real-time application the overhead of the kernel and the worst case administrative overhead (WCAO) of every system call has to be determined. Our schedulability analysis takes into account these overheads, and also the overheads due to the message passing.

The message passing mechanism is illustrated in Figure 3, where we have three processes, P_1 to P_3 . P_1 and P_2 are mapped to node N_0 that transmits in slot S_0 , and P_3 is mapped to node N_1 that transmits in slot S_1 . Message m_1 is transmitted between P_1 and P_2 that are on the same node, while message m_2 is transmitted from P_1 to P_3 between the two nodes.

Messages between processes located on the same processor are passed through shared protected objects. The overhead for their communication is accounted for by the blocking factor, computed according to the priority ceiling protocol [9].

Message m_2 has to be sent from node N_0 to node N_1 . Thus, after m_2 is produced by P_1 , it will be placed into an outgoing message queue, called *Out*. The access to the queue is guarded by a priority-ceiling semaphore. A so called transfer process (denoted with *T* in Figure 3) moves the message from the *Out* queue into the MBI.

How the message queue is organized and how the message transfer process selects the particular messages and assembles them into a frame, depends on the particular approach chosen for message scheduling (see Section 4). The message transfer process is activated at certain a priori known moments, by the tick scheduler in order to perform the message transfer. These activation times are stored in a message handling time table (MHTT) available to the real-time kernel in each node. Both the MEDL and the MHTT are generated off-line as result of the schedulability analysis and optimization which will be discussed later. The MEDL




Figure 3. Message Passing Mechanism

imposes the times when the TTP controller of a certain node has to move frames from the MBI to the communication channel. The MHTT contains the times when messages have to be transferred by the message transfer process from the *Out* queue into the MBI, in order to further be broadcasted by the TTP controller. As result of this synchronization, the activation times in the MHTT are directly related to those in the MEDL and the first table results directly form the second one.

It is easy to observe that we have the most favorable situation when, at a certain activation, the message transfer process finds in the *Out* queue all the “expected” messages which then can be packed into the just following frame to be sent by the TTP controller. However, application processes are not statically scheduled and availability of messages in the *Out* queue can not be guaranteed at fixed times. Worst case situations have to be considered as will be shown in Section 4.

Let us come back to Figure 3. There we assumed a context in which the broadcasting of the frame containing message m_2 is done in the slot S_0 of Round 2. The TTP controller of node N_1 knows from its MEDL that it has to read a frame from slot S_0 of Round 2 and to transfer it into its MBI. In order to synchronize with the TTP controller and to read the frame from the MBI, the tick scheduler on node N_1 will activate, based its local MHTT, a so called delivery process, denoted with D in Figure 3. The delivery process takes the frame from the MBI, and extracts the messages from it. For the case when a message is split into several packets, sent over several TDMA rounds, we consider that a message has arrived at the destination node after all its corresponding packets have arrived. When m_2 has arrived, the delivery process copies it to process P_3 which will be activated. Activation times for the delivery process are fixed in the MHTT just as explained earlier for the message transfer process.

The number of activations of the message transfer and delivery processes depend on the number of frames transferred, and they are taken into account in our analysis, as well as the delay implied by the propagation on the communication bus.

3. Problem Formulation

We model an application as a set of processes. Each process p_i is allocated to a certain processor, has a known worst-case execution time C_i , a period T_i , a deadline D_i and a uniquely assigned priority. For aperiodic processes, T_i represents the minimum time between successive arrivals. We consider a preemptive execution environment, which means that higher priority processes can interrupt the execution of lower priority processes. A lower priority process can block a higher priority process (e.g., it is in its critical section), and the blocking time is computed according to the priority ceiling protocol. Processes exchange messages, and for each message m_i we know its size S_{m_i} . A message is sent once in every n_m invocations of the sending process, and has a unique destination process. Each process is allocated to a

node of our distributed architecture, and the messages are transmitted according to the TTP.

We are interested to synthesize the MEDL of the TTP controllers (and as a direct consequence, also the MHTTs) so that the process set is schedulable on an as cheap (slow) as possible processor set.

4. Schedulability Analysis

Under the assumptions presented in the previous section Tindell et al. [15] integrate processor and communication schedulability and provide a “holistic” schedulability analysis in the context of distributed real-time systems with communication based on a simple TDMA protocol. The basic idea is that the release jitter of a destination process depends on the communication delay between sending and receiving a message. The release jitter of a process is the worst case delay between the arrival of the process and its release (when it is placed in the run-queue for the processor). The communication delay is the worst case time spent between sending a message and the message arriving at the destination process.

Thus, for a process $d(m)$ that receives a message m from a sender process $s(m)$, the release jitter is: $J_{d(m)} = r_{s(m)} + a_m + r_{deliver} + T_{tick}$, where $r_{s(m)}$ is the response time of the process sending the message, a_m (worst case arrival time) is the worst case time needed for message m to arrive at the communication controller of the destination node, $r_{deliver}$ is the response time of the delivery process (see section 2.2), and T_{tick} is the jitter due to the operation of the tick scheduler. The communication delay for a message m is $C_m = a_m + r_{deliver}$. a_m itself is the sum of the access delay and the propagation delay. The access delay is the time a message queued at the sending processor spends waiting for the use of the communication channel. In a_m we also account for the execution time of the message transfer process (see section 2.2). The propagation delay is the time taken for the message to reach the destination processor once physically sent by the corresponding TTP controller.

The worst case time, message m takes to arrive at the communication controller of the destination node is determined in [15] using the arbitrary deadline analysis, and is given by:

$a_m = q = \max_{q=0,1,2,\dots} (w_m(q) + X_m(q) - qT_m)$, where the term $w_m(q) - qT_m$ is the access delay, $X_m(q)$ is the propagation delay, and T_m is the period of the message.

In [15] an analysis is given for the end-to-end delay of a message m in the case of a simple TDMA protocol. For this case, $w_m(q) = \left\lceil \frac{(q+1)P_m + I_m(w(q))}{S_p} \right\rceil T_{TDMA}$, where

P_m is the number of packets of message m , S_p is the size of the slot (in number of packets) corresponding to m , and I_m is the interference caused by packets belonging to messages of a higher priority than m . Although there are many similarities with the general TDMA protocol, the analysis in the case of TTP is different in several aspects and also

differs to a large degree depending on the policy chosen for message scheduling.

Before going into details for each of the message scheduling approaches, we analyze the propagation delay and the message transfer and delivery processes, as they do not depend on the particular message scheduling policy chosen. The propagation delay X_m of a message m sent as part of a slot S , with the TTP protocol, is equal to the time needed for the slot S to be transferred on the bus. This time depends on the slot size and on the features of the underlying buss.

The overhead produced by the communication activities must be accounted for not only as part of the access delay for a message, but also through its influence on the response time of processes running on the same processor. We consider this influence during the schedulability analysis of processes on each processor. We assume that the worst case computation time of the transfer process (T in Figure 3) is known, and that it is different for each of the four message scheduling approaches. Based on the respective MHTT, the transfer process is activated for each frame sent. Its worst case period is derived from the minimum time between successive frames.

The response time of the delivery process (D in Figure 3), r_{deliver} is part of the communication delay. The influence due to the delivery process must be also included when analyzing the response time of the processes running on the respective processor. We consider the delivery process during the schedulability analysis in the same way as the message transfer process.

The response times of the communication and delivery processes are calculated, as for all other processes, using the arbitrary deadline analysis from [15].

The four approaches we have considered for scheduling of messages using TTP differ in the way the messages are allocated to the communication channel (either statically or dynamically) and whether they are split or not into packets for transmission. The next subsections present an analysis for these approaches as well as the degrees of liberty a designer has, in each of the cases, when synthesizing the MEDL.

4.1 Static Single Message Allocation (SM)

The first approach to scheduling of messages using TTP is to statically (off-line) schedule each of the messages into a slot of the TDMA cycle, corresponding to the node sending the message. We also consider that the slots can hold each at maximum one single message. This approach is well suited for application areas (like automotive electronics) where the messages are typically short and the ability to easily diagnose the system is critical.

As each slot carries only one fixed, predetermined




Figure 4. Worst case arrival time for SM

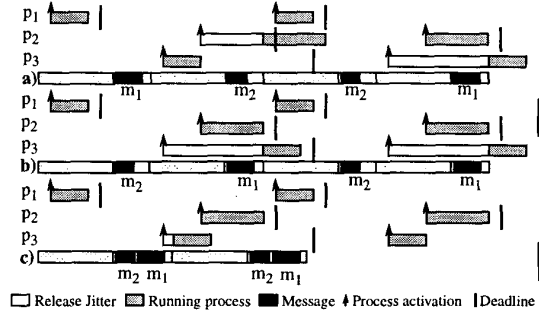


Figure 5. Optimizing the MEDL for SM and MM

message, there is no interference among messages. If a message m misses its slot it has to wait for the following slot assigned to m . The access delay for a message m in this approach is the maximum time between consecutive slots of the same node carrying the message m . We denote this time by $T_{m_{\text{max}}}$, illustrated in Figure 4.

In this case, the worst case arrival time a_m of a message m becomes $T_{m_{\text{max}}} + X_m$. Therefore, the main aspect influencing the schedulability analysis for the messages is the way the messages are statically allocated to slots, resulting different values for $T_{m_{\text{max}}}$. $T_{m_{\text{max}}}$, as well as X_m , depend on the slot sizes which in the case of SM are determined by the size of the largest message sent from the corresponding node, plus the bits for control and CRC, as imposed by the protocol.

During the synthesis of the MEDL, the designer has to allocate the messages to slots in such a way that the process set is schedulable. Since the schedulability of the process set can be influenced by the synthesis of the MEDL only through the $T_{m_{\text{max}}}$ parameters, these parameters have to be optimized.

Let us consider the simple example depicted in Figure 5, where we have three processes, p_1 , p_2 , and p_3 running each on different processors. When process p_1 finishes executing it sends message m_1 to process p_3 and message m_2 to process p_2 . In the TDMA configuration presented in Figure 5 a), only the slot for the CPU running p_1 is important for our discussion and the other slots are represented with light gray. With this configuration, where the message m_1 is allocated to the rounds 1 and 4 and the message m_2 is allocated to rounds 2 and 3, process p_2 misses its deadline because of the release jitter due to the message m_2 in round 2. However, if we have the TDMA configuration depicted in Figure 5 b), where m_1 is allocated to the rounds 2 and 4 and m_2 is allocated to the rounds 1 and 3, then all the processes meet their deadlines.

4.2 Static Multiple Message Allocation (MM)

This second approach is an extension of the first one. In this approach we allow more than one message to be statically assigned to a slot, and all the messages transmitted in the

same slot are packaged together in a frame. In this case there is also no interference, so the access delay for a message m is the same as for the first approach, namely, the maximum time between consecutive slots of the same node carrying the message m , $T_{m_{max}}$.

However, this approach offers more freedom during the synthesis of the MEDL. We have now to decide also on how many and which messages should be put in a slot. This allows more flexibility in optimizing the $T_{m_{max}}$ parameter. To illustrate this, let us consider the same example depicted in Figure 5. With the MM approach, the TDMA configuration can be arranged as depicted in Figure 5 c), where the messages m_1 and m_2 are put together in the same slot in the rounds 1 and 2. Thus, the deadline is met, and the release jitter is further reduced compared to the case presented in Figure 5 b) where the deadlines were also met but the process p_3 was experiencing large release jitter.

4.3 Dynamic Message Allocation (DM)

The previous two approaches have statically allocated one or more messages to their corresponding slots. This third approach considers that the messages are dynamically allocated to frames, as they are produced.

Thus, when a message is produced by a sender process it is placed in the *Out* queue ordered according to the priorities of the messages. At its activation, the message transfer process takes a certain number of messages from the head of the *Out* queue and constructs the frame. The number of messages accepted is decided so that their total size does not exceed the length of the data field of the frame. This length is limited by the size of the slot corresponding to the respective processor. Since the messages are sent dynamically, we have to identify them in a certain way so that they are recognized when the frame arrives at the delivery process. We consider that each message has several identifier bits appended at the beginning of the message.

Since we dynamically package the messages into frames in the order they are sorted in the queue, the access delay to the communication channel for a message m depends on the number of messages queued ahead of it.

The analysis in [15] bounds the number of queued ahead *packets* of messages of higher priority than message m , as in their case it is considered that a message can be split into packets before it is transmitted on the communication channel. We use the same analysis, but we have to apply it for the number of *messages* instead that of packets. We have to consider that messages can be of different sizes as opposed to packets which always are of the same size.

Therefore, the total size of higher priority messages queued ahead of a message m in a window w is:

$$I_m(w) = \sum_{\forall j \in hp(m)} \left\lceil \frac{w + r_{s(j)}}{T_j} \right\rceil S_j \text{ where } S_j \text{ is the size of the message } m_j, r_{s(j)} \text{ is the response time of the process sending message } m_j, \text{ and } T_j \text{ is the period of the message } m_j.$$

Further, we calculate the worst case time that a message m spends in the *Out* queue. The number of TDMA rounds needed, in the worst case, for a message m placed in the queue to be removed from the queue for transmission is $\left\lceil \frac{S_m + I_m}{S_s} \right\rceil$ where S_m is the size of the message m and S_s is the size of the slot transmitting m (we assume, in the case of DM, that for any message x , $S_x \leq S_s$). This means that the worst case time a message m spends in the *Out* queue is given by $\left\lceil \frac{S_m + I_m}{S_s} \right\rceil T_{TDMA}$, where T_{TDMA} is the time taken for a TDMA round.

To determine the term $w_m(q) - qT_m$ that gives the access delay (see Section 4), $w_m(q)$ is determined, using the arbitrary deadline analysis, as being:

$$w_m(q) = \left\lceil \frac{(q+1)S_m + I_m(w(q))}{S_s} \right\rceil T_{TDMA}. \text{ Since the size of}$$

the messages is given with the application, the parameter that will be optimized during the synthesis of the MEDL is the slot size. To illustrate how the slot size influences the schedulability, let us consider the example in Figure 6 a), where we have the same setting as for the example in Figure 5 a). The difference is that we consider message m_1 having a higher priority than message m_2 , and we schedule dynamically the messages as they are produced. With the TDMA configuration in Figure 6 a) message m_1 will be dynamically scheduled first in the slot of the first round, while message m_2 will wait in the *Out* queue until the next round comes, thus causing the process p_2 to miss its deadline. However, if we enlarge the slot so that it can accommodate both messages, message m_2 does not have to wait in the queue and it is transmitted in the same slot as m_1 . Therefore p_2 will meet its deadline, as presented in Figure 6 b). However, in general, increasing the length of slots does not necessarily improve the schedulability, as it delays the communication of messages generated by other nodes.

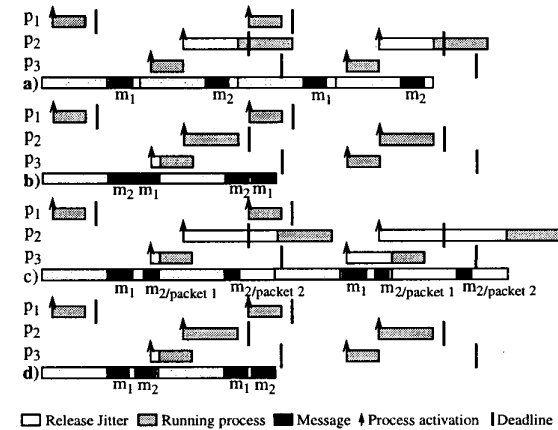


Figure 6. Optimizing the MEDL for DM and DP

4.4 Dynamic Packets Allocation (DP)

This approach is an extension of the previous one, as we allow the messages to be split into packets before they are transmitted on the communication channel. We consider that each slot has a size that accommodates a frame with the data field being a multiple of the packet size. This approach is well suited for the application areas that typically have large message sizes, and by splitting them into packets we can obtain a higher utilization of the buss and reduce the release jitter. However, since each packet has to be identified as belonging to a message, and messages have to be split at the sender and reconstructed at the destination, the overhead becomes higher than in the previous approaches.

For the analysis we use the formula from [15] which is based on similar assumptions as those for this approach:

$w_m(q) = \left\lceil \frac{(q+1)P_m + I_m(w(q))}{S_p} \right\rceil T_{TDMA}$, where P_m is the number of packets of message m , S_p is the size of the slot (in number of packets) corresponding to m , and

$I_m(w) = \sum_{\forall j \in hp(m)} \left\lceil \frac{w + r_{s(j)}}{T_j} \right\rceil P_j$, where P_j is the number of packets of a message m_j .

In the previous approach (DM) the optimization parameter for the synthesis of the MEDL was the size of the slots. Within this approach we can also decide on the packet size, which becomes another optimization parameter. Consider the example in Figure 6 c) where messages m_1 and m_2 have a size of 6 bytes each. The packet size is considered to be 4 bytes and the slot corresponding to the messages has a size of 12 bytes (3 packets) in the TDMA configuration. Since message m_1 has a higher priority than m_2 , it will be dynamically scheduled first in the slot of the first round, and it will need 2 packets. In the remaining packet, the first 4 bytes of m_2 are scheduled. Thus, the rest of 2 bytes from message m_2 have to wait for the next round, causing the process p_2 to miss its deadline. However, if we change the packet size to 3 bytes, and keep the same size of 12 bytes for the slot, we now have 4 packets in the slot corresponding to the CPU running p_1 (Figure 6 d). Message m_1 will be dynamically scheduled first, and will take 2 packets from the slot of the first round. This will allow us to send m_2 in the same round, therefore meeting the deadline for p_2 .

In this particular example, with one single sender processor and the particular message and slot sizes as given, the problem seems to be simple. This is, however, not the case in general. For example, the packet size which fits a particular node can be unsuitable in the context of the messages and slot size corresponding to another node. At the same time, reducing the packets size increases the overheads due to the transfer and delivery processes.

5. Optimization Strategy

Our problem is to analyze the schedulability of a given

process set and to synthesize the MEDL of the TTP controllers in a close to optimal way. The MEDL is synthesized according to the optimization parameters available for each of the four approaches to message scheduling discussed before. In order to guide the optimization process, we need a cost function that captures the “degree of schedulability” for a certain MEDL implementation. Our cost function is a modified version of that in [12]:

$$\text{cost function} = \begin{cases} f_1 = \sum_{i=1}^n \max(0, R_i - D_i), & \text{if } f_1 > 0 \\ f_2 = \sum_{i=1}^n R_i - D_i, & \text{if } f_1 = 0 \end{cases}$$

where n is the number of processes in the application, R_i is the response time of a process p_i , and D_i is the deadline of a process p_i . If the process set is not schedulable, there exists at least one R_i that is greater than the deadline D_i , therefore the term f_1 of the function will be positive. In this case the cost function is equal to f_1 . However, if the process set is schedulable, then all R_i are smaller than the corresponding deadlines D_i . In this case $f_1 = 0$ and we use f_2 as the cost function, as it is able to differentiate between two alternatives, both leading to a schedulable process set. For a given set of optimization parameters leading to a schedulable process set, a smaller f_2 means that we have improved the response times of the processes, so the application can be potentially implemented on a cheaper hardware architecture (with slower processors and/or buss). The release time R_i is calculated according to the arbitrary deadline analysis [15] based on the release jitter of the process (see section 4), its worst-case execution time, the blocking time, and the interference time due to higher priority processes. For a given application, we are interested to synthesize a MEDL such that the cost function is minimized. We are also interested to evaluate in different contexts the four approaches to message scheduling, thus offering the designer a decision support in choosing the right approach for his problem.

The synthesis of the MEDL is performed off-line, before implementing the application. This means that it is worth to explore the design alternatives and to try to derive near-optimal solutions. Further more, the comparison of the four approaches detailed in Section 4 is meaningful only if we take the near-optimal value for each of them.

Thus, we have developed an optimization procedure based on a simulated annealing (SA) strategy. The main characteristic of an SA strategy is that it tries to find the global optimum by randomly selecting a new solution from the neighbors of the current solution. The new solution is accepted if it is an improved one. However, a worse solution can also be accepted with a certain probability that depends on the deterioration of the cost function and on a control parameter called temperature [11].

In Figure 7 we give a short description of this algorithm. An essential component of the algorithm is the generation of a new solution x' starting from the current one x^{now} . The

```

simulated annealing
construct an initial TDMA round  $x^{now}$ 
temperature = initial temperature TI
repeat
  for i = 1 to temperature length TL
    generate randomly a neighboring solution  $x'$  of  $x^{now}$ 
    delta = cost function for  $x'$  - cost function for  $x^{now}$ 
    if delta < 0 then  $x^{now} = x'$ 
    else
      generate q = random (0, 1)
      if q <  $e^{-\text{delta} / \text{temperature}}$  then  $x^{now} = x'$  end if
    end if
  end for
  temperature =  $\alpha$  * temperature;
until stopping criterion is met
return solution corresponding to the best schedule

```

Figure 7. The Optimization Strategy

neighbors of the current solution x^{now} are obtained depending on the chosen approach. For SM, x' is obtained from x^{now} by inserting or removing a message in one of its corresponding slots. In the case of MM, we have to take additional care that the slots do not exceed the maximum allowed size (depends on the controller implementation), as we can allocate several messages to a slot. For these two static approaches, we also decide on the number of rounds in a cycle (e.g., 2, 4, 8, 16; limited by the size of the memory implementing the MEDL). The neighboring solution is obtained in the case of DM by increasing or decreasing the slots size within the bounds allowed by the particular TTP controller implementation, while in the DP approach we also increase or decrease the packet size.

For the implementation of this algorithm, the parameters TI (initial temperature), TL (temperature length), α (cooling ratio), and the stopping criterion have to be determined. They define the so called cooling schedule and have a strong impact on the quality of the solutions and the CPU time consumed. We were interested to obtain values for TI, TL and α that will guarantee the finding of good quality solutions in a short time. In order to tune the parameters we have first performed very long and expensive runs on selected large examples, and the best ever solution, for each example, has been considered as the near-optimum. Based on further experiments we have determined the parameters of the SA algorithm, for different sizes of examples, so that the optimization time is reduced as much as possible but the near-optimal result is still produced. These parameters have then been used by the large scale experiments presented in the following section. For example, for the graphs with 320 nodes, TI is 300, TL is 500 and α is 0.95. The algorithm stops if for three consecutive temperatures no new solution has been accepted.

6. Experimental Results

For evaluation of our scheduling approaches we first used sets of processes generated for experimental purpose. We considered architectures consisting of 2, 4, 6, 8 and 10 nodes. 40 processes were assigned to each node, resulting in sets of 80, 160, 240, 320 and 400 processes. 30 sets were generated for each dimension, thus a total of 150 sets of pro-

cesses were used for experimental evaluation. Worst case computation times, periods, deadlines, and message lengths were assigned randomly within certain intervals. For the communication channel we considered a transmission speed of 256 kbps. The maximum length of the data field in a slot was 32 bytes, and the frequency of the TTP controller was chosen to be 20 MHz. All experiments were run on a Sun Ultra 10 workstation.

For each of the 150 generated examples and each of the four scheduling approaches we have obtained, using our optimization strategy, the near-optimal values for the cost function. These values, for a given example, might differ from one approach to another, as they depend on the optimization parameters and the schedulability analysis determined for each of the approaches. We were interested to compare the four approaches to message scheduling based on the values obtained for the cost function.

Thus, Figure 8 a) presents the average percentage deviations of the cost function obtained by our optimization strategies in each of the four approaches, from the minimal value among them. The DP approach is generally the most performant, and the reason for this is that dynamic scheduling of messages is able to reduce release jitter because no space is wasted in the slots if the packet size is properly selected. However, by using the MM approach we can obtain almost the same result if the messages are carefully allocated to slots by our optimization strategy. Moreover, in the case of bigger sets of processes (e.g., 400) MM outperforms DP, as DP suffers from large overhead due to the handling of the packets. DM performs worse than DP because it does not split the messages into packets, and this results in a mismatch between the size of the messages dynamically queued and the slot size, leading to unused slot space that increases the jitter. SM performs the worst as its optimization strategy has not much room for improvement, leading to large amounts of unused slot space. Also, DP has produced a MEDL that resulted in schedulable process sets for 1.33 times more cases than the MM and DM. MM, in its turn, produced two times more schedulable results than the SM approach.

Together with the four approaches to message scheduling, a so called ad-hoc approach is presented. The ad-hoc approach performs scheduling of messages without trying to optimize the access to the communication channel. The ad-hoc solutions are based on the MM approach and consider a design with the TDMA configuration consisting of a simple, straightforward, allocation of messages to slots. The lengths of the slots were selected to accommodate the largest message sent from the respective node. Figure 8 a) shows that the ad-hoc alternative is constantly outperformed by any of the optimized solutions. This shows that by optimizing the access to the communication channel, significant improvements can be produced.

We were also interested to compare the four approaches




Figure 8: Comparison of the Four Approaches to Message Scheduling

with respect to the number of messages exchanged between different nodes and the maximum message size allowed. For the results depicted in Figure 8 b) and c) we have assumed sets of 80 processes allocated to 4 nodes. Figure 8 b) shows that as the number of messages increases, the differences between the approaches grow while the ranking among them remains the same. The same holds for the case when we increase the maximum allowed message size, with a notable exception. We can observe that for large message sizes MM becomes better than DP, since DP suffers from the overhead due to packet handling.

We have also considered a real-life example implementing an aircraft control system adapted from [15] where the ad-hoc solution and the SM approach failed to find a schedulable solution. However, the other approaches found schedulable solutions, DP having the smallest cost function followed in this order by MM and DM.

The above comparison between the four message scheduling alternatives is mainly based on the issue of schedulability. However, when choosing among the different policies, several other parameters can be of importance. Thus, a static allocation of messages can be beneficial from the point of view of testing and debugging and has the advantage of simplicity. Similar considerations can lead to the decision not to split messages. In any case, however, optimization of the communication structure is desirable.

7. Conclusions

We have presented an approach to static priority preemptive process scheduling for synthesis of hard real-time distributed embedded systems. The communication model was based on a time-triggered protocol. We have developed an analysis for the communication delays and optimization strategies for four different message scheduling policies. The four approaches to message scheduling were compared using extensive experiments.

References

[1] N. C. Audsley, A. Burns, R. I. Davis, K. Tindell, A. J. Wellings, "Fixed Priority Pre-emptive Scheduling: An Historical Perspective", *Real-Time Systems*, 8(2/3), 173-198, 1995.
 [2] P. Eles, K. Kuchcinski, Z. Peng, A. Doboli, P. Pop, "Scheduling of Conditional Process Graphs for the Synthesis of

Embedded Systems", *Proc. Design Aut. & Test in Europe*, 1998.
 [3] H. Ermedahl, H. Hansson, M. Sjödin, "Response-Time Guarantees in ATM Networks", *Proceedings of the 18th IEEE Real-Time Systems Symposium*, 274-284, 1997.
 [4] G. Fohler, "Joint Scheduling of Distributed Complex Periodic and Hard Aperiodic Tasks in Statically Scheduled Systems", *Proc. the 16th IEEE Real-Time Systems Symp.*, 1995.
 [5] H. Kopetz, *Real-Time Systems-Design Principles for Distributed Embedded Applications*, Kluwer Academic Publ. 1997.
 [6] H. Kopetz, G. Grünsteidl, "TTP-A Protocol for Fault-Tolerant Real-Time Systems", *IEEE Computer*, 27(1), 14-23, 1994.
 [7] C. L. Liu, J. W. Layland, "Scheduling Algorithms for Multiprogramming in a Hard-Real-Time Environment", *Journal of the ACM*, 20(1), 46-61, 1973.
 [8] H. Lonn, J. Axelsson, "A Comparison of Fixed-Priority and Static Cyclic Scheduling for Distributed Automotive Control Applications", *Proceedings of the 11th Euromicro Conference on Real-Time Systems*, 142-149, 1999.
 [9] L. Sha, R. Rajkumar, J. Lehoczky, "Priority Inheritance Protocols: An Approach to Real-Time Synchronization", *IEEE Transactions on Computers*, 39(9), 1175-1185, 1990.
 [10] P. Pop, P. Eles, Z. Peng, "Scheduling with Optimized Communication for Time-Triggered Embedded Systems", *Proc. Int. Workshop on Hardware-Software Co-design*, 78-82, 1999.
 [11] C. R. Reeves, *Modern Heuristic Techniques for Combinatorial Problems*, Blackwell Scientific Publications, 1993.
 [12] K. Tindell, A. Burns, A. J. Wellings, "Allocating Real-Time Tasks (An NP-Hard Problem made Easy)", *Real-Time Systems*, 4(2), 1992, 145-165.
 [13] K. Tindell, A. Burns, A. J. Wellings, "Calculating Controller Area Network (CAN) Message Response Times", *Control Eng. Practice*, 3(8), 1163-1169, 1995.
 [14] K. Tindell, A. Burns, A. J. Wellings, "Guaranteeing Hard Real Time End-to-End Communication Deadlines", *Dept. of Comp. Science, Univ. of York, Report Number RTRG/91/107*, 1991.
 [15] K. Tindell, J. Clark, "Holistic Schedulability Analysis for Distributed Hard Real-Time Systems", *Microprocessing and Microprogramming*, 40, 117-134, 1994.
 [16] E. Tovar, F. Vasques, A. Burns, "Adding Local Priority-Based Dispatching Mechanisms to P-NET Networks: a Fixed Priority Approach", *Proceedings of the 11th Euromicro Conference on Real-Time Systems*, 175-184, 1999.
 [17] X-by-Wire Consortium, "X-By-Wire: Safety Related Fault Tolerant Systems in Vehicles", URL: <http://www.vmars.tuwien.ac.at/projects/xbwire/>