Design Optimization of Multi-Cluster Embedded Systems for Real-Time Applications

Paul Pop, Petru Eles and Zebo Peng
Linköping University, Sweden

Abstract

An increasing number of real-time applications are today implemented using distributed heterogeneous architectures composed of interconnected networks of processors. The systems are heterogeneous not only in terms of hardware components, but also in terms of communication protocols and scheduling policies. Each network has its own communication protocol, each processor in the architecture can have its own scheduling policy, and several scheduling policies can share a processor. In this context, the task of designing such systems is becoming increasingly important and difficult at the same time. The success of such new design methods depends on the availability of analysis and optimization techniques. In this paper, we present analysis and optimization techniques for heterogeneous real-time embedded systems. We address in more detail a particular class of such systems called multi-clusters, composed of several networks interconnected via gateways. We present a schedulability analysis for safety-critical applications distributed on multi-cluster systems and briefly highlight characteristic design optimization problems: the partitioning and mapping of functionality, and the packing of application messages to frames. Optimization heuristics for frame packing aiming at producing a schedulable system are presented. Extensive experiments and a real-life example show the efficiency of the frame-packing approach.

1 Introduction

Embedded real-time systems have to be designed such that they implement correctly the required functionality. In addition, they have to fulfill a wide range of competing constraints: development cost, unit cost, reliability, security, size, performance, power consumption, flexibility, time-to-market, maintainability, correctness, safety, etc. Very important for the correct functioning of such systems are their timing constraints: *"the correctness of the system behavior depends not only on the logical results of the computations, but also on the physical instant at which these results are produced"* [?]. Real-time systems have been classified as *hard* real-time and *soft* real-time systems [?]. Basically, hard real-time systems are systems where failing to meet a timing constraint can potentially have catastrophic consequences. For example, a brake-by-wire system in a car failing to react within a given time interval can result in a fatal accident. On the other hand, a multimedia system, which is a soft-real time system, can, under certain circumstances, tolerate a certain amount of delays resulting maybe in a patchier picture, without serious consequences besides some possible inconvenience to the user. Many real-time applications, following physical, modularity or safety constraints, are implemented using *distributed architectures*. Such systems are composed of several different types of hardware components, interconnected in a network. For such systems, the communication between the functions implemented on different nodes has an important impact on the overall system properties such as performance, cost, maintainability, etc. The analysis and optimization approaches presented in this paper are aimed towards heterogeneous distributed hard-real time systems that implement safety-critical applications where timing constraints are of utmost importance to the correct behavior of the application. The paper is organized in ten sections. Section2 presents the heterogeneous real-time embedded systems addressed in the paper,

and the type of applications considered. Sections 3 and 4 introduce the current state-of-the-art on the analysis and optimization of such systems. The rest of the paper focuses in more detail on some techniques for multi-cluster systems. The hardware and software architecture of multi-clusters, together with the application model, are outlined in Section5. Section6 identifies partitioning and mapping and frame packing as design optimization problems characteristic to multi-clusters. We present an analysis for multi-cluster systems in Section7, and show, in Section8, how this analysis can be used to drive the optimization of the packing of application messages to frames. The last two sections present the experimental results of the frame packing optimization and conclusions.

1.1 Automotive Electronics

Although the discussion in this paper is valid for several application areas, it is useful, for understanding the distributed embedded real-time systems evolution and design challenges, to exemplify the developments in a particular area. If we take the example of automotive manufacturers, they were reluctant, until recently, to use computer controlled functions onboard vehicles. Today, this attitude has changed for several reasons. First, there is a constant market demand for increased vehicle performance, more functionality, less fuel consumption and less exhausts, all of these at lower costs. Then, from the manufacturers side, there is a need for shorter time-to-market and reduced development and manufacturing costs. These, combined with the advancements of semiconductor technology, which is delivering ever increasing performance at lower and lower costs, has led to the rapid increase in the number of electronically controlled functions onboard a vehicle [?]. The amount of electronic content in an average car in 1977 had a cost of $110. Currently, that cost is$1341, and it is expected that this figure will reach \$1476 by the year 2005, continuing to increase because of the introduction of sophisticated electronics found until now only in high-end cars [?, ?]. It is estimated that in 2006 the electronics inside a car will amount to 25% of the total cost of the vehicle (35% for the high end models), a quarter of which will be due to semiconductors [?, ?]. High-end vehicles currently have up to 100 microprocessors implementing and controlling various parts of their functionality. The total market for semiconductors in vehicles is predicted to grow from \$8.9 billions in 1998 to \$21 billion in 2005, amounting to 10% of the total worldwide semiconductors market [?, ?]. At the same time with the increased complexity, the type of functions implemented by embedded automotive electronics systems has also evolved. Thanks to the semiconductors revolution, in the late 50s, electronic devices became small enough to be installed on board of vehicles. In the 60s the first analog fuel injection system appeared, and in the 70s analog devices for controlling transmission, carburetor, and spark advance timing were developed. The oil crisis of the 70s led to the demand of engine control devices that improved the efficiency of the engine, thus reducing fuel consumption. In this context, the first microprocessor based injection control system appeared in 1976 in the USA. During the 80s, more sophisticated systems began to appear, like electronically controlled braking systems, dashboards, information and navigation systems, air conditioning systems, etc. In the 90s, development and improvement have concentrated in the areas like safety and convenience. Today, it is not uncommon to have highly critical functions like steering or braking implemented through electronic functionality only, without any mechanical backup, like is the case in drive-by-wire and brake-by-wire systems [?, ?]. The complexity of electronics in modern vehicles is growing at a very high pace, and the constraintsin terms of functionality, performance, reliability, cost and time-to-marketare getting tighter. Therefore, the task of designing such systems is becoming increasingly important and difficult at the same time. New design techniques are needed, which are able to:

- successfully manage the complexity of embedded systems,

- meet the constraints imposed by the application domain,

- shorten the time-to-market, and

- reduce development and manufacturing costs.

Figure 1: Distributed Hard Real-Time Systems

The success of such new design methods depends on the availability of analysis and optimization techniques, beyond those corresponding to the state-of-the-art, which are presented in the next section.

2 Heterogeneous Real-Time Embedded Systems

2.1 Heterogeneous Hardware Architecture

Currently, distributed real-time systems are implemented using architectures where each node is dedicated to the implementation of a single function or class of functions. The complete system can be, in general, composed of several networks, interconnected with each other (see Figure reffig:fig1). Each network has its own communication protocol, and inter-network communication is via a *gateway* which is a node connected to both networks. The architecture can contain several such networks, having different types of topologies. A network is composed of several different types of hardware components, called *nodes*. Typically, every node, also called *electronic control unit* (ECU), has a communication controller, CPU, RAM, ROM and an I/O interface to sensors and actuators. Nodes can also have ASICs in order to accelerate parts of their functionality. The microcontrollers used in a node and the type of network protocol employed are influenced by the nature of the functionality and the imposed real-time, fault-tolerance and power constraints. In the automotive electronics area, the functionality is typically divided in two classes, depending on the level of criticalness:

- *Body electronics* refers to the functionality that controls simple devices such as the lights, the mirrors, the windows, the dashboard. The constraints of the body electronic functions are determined by the reaction time of the human operator that is in the range of 100 ms to 200 ms. A typical body electronics system within a vehicle consists of a network of ten to twenty nodes that are interconnected by a low bandwidth communication network like LIN. A node is usually implemented using a single-chip 8 bit micro-controller (e.g., Motorola 68HC05 or Motorola 68HC11) with some hundred bytes of RAM and Kilobytes of ROM, I/O points to connect sensors and to control actuators, and a simple network interface. Moreover, the memory size is growing by more than 25% each year [?, ?].

- *System Electronics* are concerned with the control of vehicle functions that are related to the movement of the vehicle. Examples of system electronics applications are engine control, braking, suspension, vehicle dynamics control. The timing constraints of system electronic functions are in the range of a couple of ms to 20 ms, requiring 16-bit or 32-bit micro-controllers (e.g., Motorola 68332) with about 16 Kilobytes of RAM and 256 Kilobytes of ROM. These microcontrollers have built-in communication controllers (e.g., the 68HC11 and 68HC12 automotive family of microcontrollers have an on-chip CAN controller), I/O to sensors and actuators, and are interconnected by high bandwidth networks [?, ?].

Section5 presents more details concerning the hardware and software architecture considered by our analysis and optimization techniques.

2.2 Heterogeneous Communication Protocols

As the communications become a critical component, new protocols are needed that can cope with the high bandwidth and predictability required. There are several communication protocols for real-time networks. Among the protocols that have been proposed for vehicle multiplexing, only the Controller Area Network (CAN) [?], the Local Interconnection Network (LIN) [?], and SAEs J1850 [?] are currently in use on a large scale. Moreover, only a few of them are suitable

for safety-critical applications where predictability is mandatory [**?**]. A survey and comparison of communication protocols for safety-critical embedded systems is available in [**?**]. Communication activities can be triggered either dynamically, in response to an event, or statically, at predetermined moments in time.

- Therefore, on one hand, there are protocols that schedule the messages statically based on the progression of time, for example, the SAFEbus [**?**] and SPIDER [**?**] protocols for the avionics industry, and the TTCAN [**?**] and Time-Triggered Protocol (TTP) [**?**] intended for the automotive industry.

- On the other hand, there are several communication protocols where message scheduling is performed dynamically, such as Controller Area Network (CAN) used in a large number of application areas including automotive electronics, LonWorks [**?**] and Profibus [**?**] for real-time systems in general, etc. Out of these, CAN is the most well known and widespread event-driven communication protocol in the area of distributed embedded real-time systems.

- However, there is also a hybrid type of communication protocols, such as Byteflight [**?**] introduced by BMW for automotive applications and the FlexRay protocol [**?**], that allows the sharing of the bus by event-driven and time-driven messages. The time-triggered protocols have the advantage of simplicity and predictability, while event-triggered protocols are flexible and have low cost. Moreover, protocols like TTP offer fault-tolerant services necessary in implementing safety-critical applications. However, it has been shown [**?**] that event-driven protocols like CAN are also predictable, and fault-tolerant services can also be offered on top of protocols like the TTCAN. A hybrid communication protocol like FlexRay offers some of the advantages of both worlds.

2.3   Heterogeneous Scheduling Policies

The automotive suppliers will select, based on their own requirements, the scheduling policy to be used in their ECU. The main approaches to scheduling are:

- *Static cyclic scheduling* algorithms are used to build, off-line, a schedule table with activation times for each process, such that the timing constraints of processes are satisfied.

- *Fixed priority scheduling* (FPS). In this scheduling approach each process has a fixed (static) priority which is computed off-line. The decision on which ready process to activate is taken on-line according to their priority.

- *Earliest deadline first* (EDF). In this case, that process will be activated which has the nearest deadline.

Typically, processes scheduled off-line using static cyclic scheduling are non pre-emptable, while processes scheduled using techniques such as FPS and EDF are pre-emptable. Another important distinction is between the event-triggered and time-triggered approaches.

- *Time-Triggered*: In the time-triggered approach activities are initiated at predetermined points in time. In a distributed time-triggered system it is assumed that the clocks of all nodes are synchronized to provide a global notion of time. Time-triggered systems are typically implemented using *non-preemptive static cyclic scheduling*, where the process activation or message communication is done based on a schedule table built off-line.

- *Event-Triggered*: In the event-triggered approach activities happen when a significant change of state occurs. Event-triggered systems are typically implemented using *preemptive fixed-priority based scheduling*, or *earliest deadline first*, where, as response to an event, the appropriate process is invoked to service it.

Figure 2: Distributed Safety-Critical Applications

There has been a long debate in the real-time and embedded systems communities concerning the advantages of each approach [**?**, **?**, **?**]. Several aspects have been considered in favour of one or the other approach, such as flexibility, predictability, jitter control, processor utilization, and testability. An interesting comparison of the ET and TT approaches, from a more industrial, in particular automotive perspective, can be found in [**?**]. The conclusion there is that one has to choose the right approach, depending on the particularities of the application. For certain applications, several scheduling approaches can be used together. Efficient implementation of new, highly sophisticated automotive applications, entails the use of time-triggered process sets together with event-triggered ones implemented on top of complex distributed architectures.

2.4 Distributed Safety-Critical Applications

Considering the automotive industry, the way functionality has been distributed on an architecture has evolved over time. Initially, distributed real-time systems were implemented using architectures where each node is dedicated to the implementation of a single function or class of functions, allowing the system integrators to purchase nodes implementing required functions from different vendors, and to integrate them into their system [**?**]. There are several problems related to this restricted mapping of functionality:

- The number of such nodes in the architecture has exploded, reaching, for example, more than 100 in a high-end car, incurring heavy cost and performance penalties.

- The resulting solutions are sub-optimal in many aspects, and do not use the available resources efficiently in order to reduce costs. For example, it is not possible to move a function from one node to another node where there are enough available resources (e.g., memory, computation power).

- Emerging functionality, such as brake-by-wire in the automotive industry, is inherently distributed, and achieving an efficient fault-tolerant implementation is very difficult in the current setting.

This has created a huge pressure to reduce the number of nodes by integrating several functions in one node and, at the same time, to distribute certain functionality over several nodes (see Figure reffig:fig2). Although an application is typically distributed over one single network, we begin to see applications that are distributed across several networks. For example, in Figure reffig:fig2, the third application, represented as black dots, is distributed over two networks. This trend is driven by the need to further reduce costs, improve resource usage, but also by application constraints like having to be physically close to particular sensors and actuators. Moreover, not only are these applications distributed across networks, but their functions can exchange critical information through the gateway nodes.

3 Schdulability Analysis

There is a large quantity of research [**?**, **?**, **?**] related to scheduling and schedulability analysis, with results having been incorporated in analysis tools such as TimeWiz [**?**], RapidRMA [**?**], RTA-OSEK Planner [**?**], and Aires [**?**]. The tools determine if the timing constraints of the functionality are met, and support the designer in exploring several design scenarios, and help to design optimized implementations. Typically, the timing analysis considers independent processes running on single processors. However, very often functionality consists of distributed processes that have data and control dependencies, exclusion constraints, etc. New schedulability analysis techniques are needed which can handle distributed applications, data and control dependencies,

and accurately take into account the details of the communication protocols that have an important influence on the timing properties. Moreover, highly complex and safety critical applications can in the future be distributed across several networks, and can use different, heterogeneous, scheduling policies. Preemptive scheduling of independent processes with static priorities running on single-processor architectures has its roots in the work of Liu and Layland [?]. The approach has been later extended to accommodate more general computational models and has also been applied to distributed systems [?]. The reader is referred to [?, ?, ?] for surveys on this topic. Static cyclic scheduling of a set of data dependent software processes on a multiprocessor architecture has also been intensively researched [?, ?]. In [?] an earlier deadline first strategy is used for non-preemptive scheduling of processes with possible data dependencies. Preemptive and non-preemptive static scheduling are combined in the cosynthesis environment described in [?, ?]. In many of the previous scheduling approaches researchers have assumed that processes are scheduled independently. However, processes can be sporadic or aperiodic, are seldom independent, and normally they exhibit precedence and exclusion constraints. Knowledge regarding these dependencies can be used in order to improve the accuracy of schedulability analyses and the quality of the produced schedules [?]. It has been claimed [?] that static cyclic scheduling is the only approach that can provide efficient solutions to applications that exhibit data dependencies. However, advances in the area of fixed priority preemptive scheduling show that such applications can also be handled with other scheduling strategies [?]. One way of dealing with data dependencies between processes in the context of static priority based scheduling has been indirectly addressed by the extensions proposed for the schedulability analysis of distributed systems through the use of the *release jitter* [?]. Release jitter is the worst case delay between the arrival of a process and its release (when it is placed in the ready-queue for the processor) and can include the communication delay due to the transmission of a message on the communication channel. In [?] and [?] time *offset* relationships and *phases*, respectively, are used in order to model data dependencies. Offset and phase are similar concepts that express the existence of a fixed interval in time between the arrivals of sets of processes. The authors show that by introducing such concepts into the computational model, the pessimism of the analysis is significantly reduced when bounding the time behavior of the system. The concept of *dynamic offsets* has been later introduced and used to model data dependencies [?]. 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. Researchers have often ignored or very much simplified the communication infrastructure. One typical approach is to consider communications 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. [?]. Tindell et al. [?] integrate processor and communication scheduling and provide a "holistic" schedulability analysis in the context of distributed real-time systems. The validity of the analysis has been later confirmed in [?]. In the case of a distributed system the response time of a process also depends on the communication delay due to messages. In [?] the analysis for messages is done in a similar way as for processes: a message is seen as an non pre-emptable process that is "running" on a bus. The response time analyses for processes and messages are combined by realizing that the *jitter* (the delay between the *arrival* of a processthe time when becomes ready for executionand the start of its execution) of a destination process depends on the *communication delay* (the time it takes for a message to reach the destination process, from the moment it has been produced by the sender process) between sending and receiving a message. Several researchers have provided analyses that bound the communication delay for a given communication protocol:

- Controller area network protocol [?];

- Time-division multiple access protocol [?];

- Asynchronous transfer mode protocol [?];

- Token ring protocol [?],

- Fiber distributed data interface protocol [**?**].

- Time-triggered protocol [**?**];

- FlexRay protocol [**?**].

Based on their own requirements, the suppliers choose one particular scheduling policy to be used. However, for certain applications, several scheduling approaches can be used together. One approach to the design of such systems, is to allow ET and TT processes to share the same processor as well as static (TT) and dynamic (ET) communications to share the same bus. Bus sharing of TT and ET messages is supported by protocols which support both static and dynamic communication [**?**]. we have addressed the problem of timing analysis for such systems [**?**]. A fundamentally different architectural approach to heterogeneous TT/ET systems is that of heterogeneous multi-clusters, where each cluster can be either TT or ET. In a *time-triggered cluster* processes and messages are scheduled according to a static cyclic policy, with the bus implementing a TDMA protocol such as, for example, the time-triggered protocol. On *event-triggered clusters* the processes are scheduled according to a priority based preemptive approach, while messages are transmitted using the priority-based CAN bus. In this context, we have proposed an approach to schedulability analysis for multi-cluster distributed embedded systems [**?**]. This analysis will be outlined in Section7. When several event-driven scheduling policies are used in a heterogeneous system, another approach to the verification of timing properties is to use the technique presented in [**?**] which couples the analysis of local scheduling strategies via an event interface model.

4 Design Optimization

4.1 Traditional Design Methodology

There are several methodologies for real-time embedded systems design. The aim of a design methodology is to coordinate the design tasks such that the time-to-market is minimized, the design constraints are satisfied, and various parameters are optimized. The main design tasks that have to be preformed are described in the following sections.

4.1.1 Functional Analysis and Design

The functionality of the host system, into which the electronic system is embedded, is normally described using a formalism from that particular domain of application. For example, if the host system is a vehicle, then its functionality is described in terms of control algorithms using differential equations, which are modeling the behavior of the vehicle and its environment. At the level of the embedded real-time system which controls the host system, the functionality is typically described as a set of functions, accepting certain inputs and producing some output values. The typical automotive application is a control application. The controller reads inputs from sensors, and uses the actuators to control the physical environment (the vehicle). A controller can have several modes of operation, and can interact with other electronic functions, or with the driver through switches and instruments. During the functional analysis and design stage, the desired functionality is specified, analyzed and decomposed into sub-functions based on the experience of the designer. Several suppliers and manufacturers have started to use tools like Statemate [**?**], Matlab/Simulink [**?**], ASCET/SD [**?**] and SystemBuild/ MatrixX [**?**] for describing the functionality, in order to eliminate the ambiguities and to avoid producing incomplete or incoherent specifications. At the level of functional analysis the exploration is currently limited to evaluating several alternative control algorithms for solving the control problem. Once the functionality has been captured using tools like Matlab/Simulink, useful explorations can involve simulations of executable specifications in order to determine the correctness of the behavior, and to assess certain properties of chosen solutions.

4.1.2 Architecture Selection and Mapping

The architecture selection task decides what components to include in the hardware architecture and how these components are connected. According to current practice, architecture selection is an *ad-hoc* process, based on the experience of the designer and previous product versions. The mapping task has to decide what part of the functionality should be implemented on which of the selected components. The manufacturers integrate components from suppliers, and thus the design space is severely restricted in current practice, by the fact that the mapping of functionality to an ECU is fixed.

4.1.3 Software Design and Implementation

This is the phase in which the software is designed and the code is written. The code for the functions is developed manually for efficiency reasons, and thus the exploration that would be allowed by automatic code generation is limited. At this stage the correctness of the software is analyzed through simulations, but there is no analysis of timing constraints, which is left for the scheduling and schedulability analysis stage.

4.1.4 Scheduling and Schedulability Analysis

Once the functions have been defined and the code has been written, the scheduling task is responsible for determining the execution strategy for the functions *inside an ECU*, such that the timing constraints are satisfied. Simulation is extensively used to determine if the timing constraints are satisfied. However, simulations are very time consuming and provide no guarantees that the timing constraints are met. In the context of static cyclic scheduling, deriving a schedule table is a complex design exploration problem. Static cyclic scheduling of a set of data dependent software processes on a multiprocessor architecture has been researched in [?, ?]. Such research has been used in commercial tools like TTP-Plan [?] which derives the static schedules for processes and messages in a time-triggered system using the time-triggered protocol for communication. If fixed priority preemptive scheduling is used, exploration is used to determine how to allocate priorities to a set of distributed processes [?]. Their priority assignment heuristic is based on the schedulability analysis from [?]. For earliest deadline first the issue of distributing the global deadlines to local deadlines has to be addressed [?].

5 Integration

In this phase the manufacturer has to integrate the ECUs from different suppliers. There is a lack of tools that can analyze the performance of the interacting functionality, and thus the manufacturer has to rely on simulation runs using the realistic environment of a prototype car. Detecting potential problems at such a late stage requires time-consuming extensive simulations. Moreover, once a problem is identified it takes a very long time to go through all the previous stages in order to fix it. This leads to large delays on the time-to-market. In order to reduce the large simulation times, and to guarantee that potential violations of timing constraints are detected, manufacturers have started to use in-house analysis tools and commercially available tools such as Volcano Network Architect (for the CAN and LIN buses) [?]. Volcano makes inter-ECU communication transparent for the programmer. The programmer only deals with *signals* that have to be sent and received, and the details of the network are hidden. Volcano provides basic API calls for manipulating signals. To achieve interoperability between ECUs from different suppliers, Volcano uses a *publish/subscribe* model for defining the signal requirements. Published signals are made available to the system integrator by the suppliers, while subscribed signals are required as inputs to the ECU. The system integrator makes the publish/subscribe connections by creating a set of CAN frames, and creating a mapping between the data in frames and signals [?]. Volcano uses the analysis in [?] for bounding the communication delay of messages transmitted using the CAN bus.

5.0.5 Calibration, Testing, Verification

These are the final stages of the design process. Because not enough analysis, testing and verification has been done in earlier stages of the design, these stages tend to be very time consuming, and problems identified here lead to large delays in product delivery.

5.1 Function Architecture Co-Design and Platform Based Design

New design methodologies are needed, which can handle the increasing complexity of heterogeneous systems, and their competing requirements in terms of performance, reliability, low power consumption, cost, time-to-market, etc. As the complexity of the systems continues to increase, the development time lengthens dramatically, and the manufacturing costs become prohibitively high. To cope with this complexity, it is necessary to reuse as much as possible at all levels of the design process, and to work at higher and higher abstraction levels. *Function/architecture co-design* is a design methodology proposed in [?, ?], which addresses the design process at higher abstraction levels. Function/architecture co-design uses a top-down synthesis approach, where trade-offs are evaluated at a high level of abstraction. The main characteristic of this methodology is the use, at the same time with the top-down synthesis, of a bottom-up evaluation of design alternatives, without the need to perform a full synthesis of the design. The approach to obtaining accurate evaluations is to use an accurate modeling of the behavior and architecture, and to develop analysis techniques that are able to derive estimates and to formally verify properties relative to a certain design alternative. The determined estimates and properties, together with user-specified constraints, are then used to drive the synthesis process. Thus, several architectures are evaluated to determine if they are suited for the specified system functionality. There are two extremes in the degrees of freedom available for choosing an architecture. At one end, the architecture is already given, and no modifications are possible. At the other end of the spectrum, no constraints are imposed on the architecture selection, and the synthesis task has to determine, from scratch, the best architecture for the required functionality. These two situations are, however, not common in practice. Often, a *hardware platform* is available, which can be *parameterized* (e.g., size of memory, speed of the buses, etc.). In this case, the synthesis task is to derive the parameters of the platform architecture such that the functionality of the system is successfully implemented. Once an architecture is determined and/or parameterized, the function/architecture co-design continues with the mapping of functionality onto the instantiated architecture. This methodology has been used in research tools like Polis [?] and Metropolis [?], and has also led to commercial tools such as the Virtual Component Co-design (VCC) [?]. In order to reduce costs, especially in the case of a mass market product, the system architecture is usually reused, with some modifications, for several product lines. Such a common architecture is denoted by the term *platform*, and consequently the design tasks related to such an approach are grouped under the term *platform-based design* [?]. The platform consists of a hardware infrastructure together with software components that will be used for several product versions, and will be shared with other product lines, in the hope to reduce costs and the time-to-market. The authors in [?] have proposed techniques for deriving such a platform for a given family of applications. Their approach can be used within any design methodology for determining a system platform that later on can be parameterized and instantiated to a desired system architecture. Considering a given application or family of applications, the system platform has to be instantiated, deciding on certain parameters, and lower level details, in order to suit that particular application(s). The search for an architecture instance starts from a certain platform, and a given application. The application is mapped and compiled on an architecture instance, and the performance numbers are derived, typically using simulation. If the designer is not satisfied with the performance of the instantiated architecture, the process is repeated. In the remainder of the paper we will consider a platform consisting of event- and time-triggered clusters, using the CAN and TTP protocols for communication, respectively. We will discuss analysis and optimization techniques for the configuration of the platform such that the given application is schedulable.

Figure 3: Time-Triggered Protocol

Figure 4: Controller Area Network Data Frame (CAN 2.0A)

6 Multi-Cluster Systems

One class of heterogeneous real-time embedded systems is that of *multi-cluster* systems. We consider architectures consisting of two clusters, one time-triggered, and the other event-triggered, interconnected by gateways (see Figure reffig:fig2):

- In a *time-triggered cluster* (TTC) processes and messages are scheduled according to a static cyclic policy, with the bus implementing a TDMA protocol such as, for example, the time-triggered protocol (TTP) [**?**].

- On *event-triggered clusters* (ETC) the processes are scheduled according to a priority based preemptive approach, while messages are transmitted using the priority-based CAN bus [**?**].

The next two sections present the hardware and software architecture of a two-cluster system, while Section5.3 presents the application model used. Section6 will introduce design problems characteristic for multi-cluster systems composed of time-triggered clusters interconnected with event-triggered clusters: the partitioning of functionality between the TT and ET clusters, the mapping of functionality to the nodes inside a cluster, and the packing of application message to frames on the TTP and CAN buses. Then, Section8 will present two optimization strategies for the frame packing problem.

6.1 Hardware Architecture

A *cluster* is composed of nodes which share a broadcast communication channel. Let N_T (N_E) be the set of nodes on the TTC (ETC). Every *node* $N_i \in N_T \cup N_E$ includes a communication controller and a CPU, along with other components. The gateways, connected to both types of clusters, have two communication controllers, for TTP and CAN. The communication controllers implement the protocol services, and run independently of the nodes CPU. Communication with the CPU is performed through a *Message Base Interface* (MBI); see Figure reffig:fig5. Communication between the nodes on a TTC is based on the TTP [**?**]. The TTP integrates all the services necessary for fault-tolerant real-time systems. The bus access scheme is time-division multiple-access (TDMA), meaning that each node N_i on the TTC, including the gateway node, can transmit only during a predetermined time interval, the TDMA slot S_i. In such a slot, a node can send several messages packed 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. The TDMA access scheme is imposed by a message descriptor list (MEDL) that is located in every TTP controller. The MEDL serves as a schedule table for the TTP controller which has to know when to send/receive a frame to/from the communication channel.

There are two types of frames in the TTP. The initialization frames, or I-frames, which are needed for the initialization of a node, and the normal frames, or N-frames, which are the data frames containing, in their data field, the application messages. A TTP data frame (Figure reffig:fig3) consists of the following fields: start of frame bit (SOF), control field, a data field of up to 16 bytes containing one or more messages, and a cyclic redundancy check (CRC) field. Frames are delimited by the inter-frame delimiter (IDF, 3 bits).

For example, the data efficiency of a frame that carries 8 bytes of application data, i.e., the percentage of transmitted bits which are the actual data bits needed by the application, is 69.5% (64 data bits transmitted in a 92-bit frame, without considering the details of a particular physical

Figure 5: Message Passing Example

layer). Note that no identifier bits are necessary, as the TTP controllers know from their MEDL what frame to expect at a given point in time. In general, the protocol efficiency is in the range of 6080% [**?**]. On an ETC, the CAN [**?**] protocol is used for communication. The CAN bus is a priority bus that employs a collision avoidance mechanism, whereby the node that transmits the frame with the highest priority wins the contention. Frame priorities are unique and are encoded in the frame identifiers, which are the first bits to be transmitted on the bus. In the case of CAN 2.0A, there are four frame types: data frame, remote frame, error frame, and overload frame. We are interested in the composition of the data frame, depicted in Figure reffig:fig3. A data frame contains seven fields: SOF, arbitration field that encodes the 11 bits frame identifier, a control field, a data field up to 8 bytes, a CRC field, an acknowledgement (ACK) field, and an end of frame field (EOF). In this case, for a frame that carries 8 bytes of application data, we will have an efficiency of 47.4% [**?**]. The typical CAN protocol efficiency is in the range of 2535% [**?**].

6.2   Software Architecture

A real-time kernel is responsible for activation of processes and transmission of messages on each node. On a TTC, the processes are activated based on the local schedule tables, and messages are transmitted according to the MEDL. On an ETC, we have a scheduler that decides on activation of ready processes and transmission of messages, based on their priorities. In Figure reffig:fig5 we illustrate our message passing mechanism. Here we concentrate on the communication between processes located on different clusters. For message passing within a TTC the reader is directed to [**?**], while the infrastructure needed for communications on an ETC has been detailed in [**?**]. Let us consider the example in Figure reffig:fig5, where we have an application consisting of four processes and four messages mapped on two clusters. Processes P_1 and P_4 are mapped on node N_1 of the TTC, while P_2 and P_3 are mapped on node N_2 of the ETC. Process P_1 sends messages m_1 and m_2 to processes P_2 and P_3, respectively, while P_2 and P_3 send messages m_3 and m_4 to P_4. All messages have a size of one byte.

The transmission of messages from the TTC to the ETC takes place in the following way (see Figure reffig:fig5). P_1, which is statically scheduled, is activated according to the schedule table, and when it finishes it calls the send kernel function in order to send m_1 and m_2, indicated in the figure by the number (1). Messages m_1 and m_2 have to be sent from node N_1 to node N_2. At a certain time, known from the schedule table, the kernel transfers m_1 and m_2 to the TTP controller by packing them into a frame in the MBI. Later on, the TTP controller knows from its MEDL when it has to take the frame from the MBI, in order to broadcast it on the bus. In our example, the timing information in the schedule table of the kernel and the MEDL is determined in such a way that the broadcasting of the frame is done in the slot S_1 of round 2 (2). The TTP controller of the gateway node N_G knows from its MEDL that it has to read a frame from slot S_1 of round 2 and to transfer it into its MBI (3). Invoked periodically, having the highest priority on node N_G, and with a period which guarantees that no messages are lost, the gateway process T copies messages m_1 and m_2 from the MBI to the TTP-to-CAN priority-ordered message queue Out_{CAN} (4). Let us assume that on the ETC messages m_1 and m_2 are sent independently, one per frame. The highest priority frame in the queue, in our case the frame f_1 containing m_1, will tentatively be broadcast on the CAN bus (5). Whenever f_1 will be the highest priority frame on the CAN bus, it will successfully be broadcast and will be received by the interested nodes, in our case node N_2 (6). The CAN communication controller of node N_2 receiving f_1 will copy it in the transfer buffer between the controller and the CPU, and raise an interrupt which will activate a delivery process, responsible to activate the corresponding receiving process, in our case P_2, and hand over message m_1 that finally arrives at the destination (7). Message m_3 (depicted in Figure reffig:fig5 as a grey rectangle labeled "m_3") sent by process P_2 from the ETC will be transmitted to process P_4 on the TTC. The transmission starts when P_2 calls its send function

and enqueues m_3 in the priority-ordered Out_{N_2} queue (8). When the frame f_3 containing m_3 has the highest priority on the bus, it will be removed from the queue (9) and broadcast on the CAN bus (10). Several messages can be packed into a frame in order to increase the efficiency of data transmission. For example, m_3 can wait in the queue until m_4 is produced by P_3, in order to be packed together with m_4 in a frame. When f_3 arrives at the gateways CAN controller it raises an interrupt. Based on this interrupt, the gateway transfer process T is activated, and m_3 is unpacked from f_3 and placed in the Out_{TTP} FIFO queue (11). The gateway node N_G is only able to broadcast on the TTC in the slot S_G of the TDMA rounds circulating on the TTP bus. According to the MEDL of the gateway, a set of messages not exceeding $size_{S_G}$ of the data field of the frame traveling in slot S_G will be removed from the front of the Out_{TTP} queue in every round, and packed in the S_G slot (12). Once the frame is broadcast (13) it will arrive at node N_1 (14), where all the messages in the frame will be copied in the input buffers of the destination processes (15). Process P_4 is activated according to the schedule table, which has to be constructed such that it accounts for the worst-case communication delay of message m_3, bounded by the analysis in Section 7.1, and, thus, when P_4 starts executing it will find m_3 in its input buffer. As part of our frame packing approach, we generate all the MEDLs on the TTC (i.e., the TT frames and the sequence of the TDMA slots), as well as the ET frames and their priorities on the ETC such that the global system is schedulable.

6.3   Application Model

We model an application Γ as a set of process graphs $G_i \in \Gamma$ (see Figure reffig:fig6). Nodes in the graph represent processes and arcs represent dependency between the connected processes. A *process* is a sequence of computations (corresponding to several building blocks in a programming language) which starts when all its inputs are available. When it finishes executing, the process produces its output values. Processes can be pre-emptable or non pre-emptable. *Non pre-emptable* processes are processes that cannot be interrupted during their execution, and are mapped on the TTC. *Pre-emptable* processes can be can be interrupted during their execution, and are mapped on the ETC. For example, a higher priority process has to be activated to service an event, in this case, the lower priority process will be temporary pre-empted until the higher priority process finishes its execution. A process graph is polar, which means that there are two nodes, called source and sink, that conventionally represent the first and last process. If needed, these nodes are introduced as dummy processes so that all other nodes in the graph are successors of the source and predecessors of the sink, respectively. The communication time between processes mapped on the same processor is considered to be part of the process worst-case execution time and is not modeled explicitly. Communication between processes mapped to different processors is performed by message passing over the buses and, if needed, through the gateway. Such message passing is modeled as a communication process inserted on the arc connecting the sender and the receiver process (the black dots in Figure reffig:fig6). Potential communication between processes in different applications is not part of the model. Technically, such a communication is implemented by the kernels based on asynchronous non-blocking send and receive primitives. Such messages are considered non-critical and are not affected by real-time constraints. Therefore, communications of this nature will not be addressed in this paper. Each process P_i is mapped on a processor $M(P_i)$ (mapping represented by hashing in Figure reffig:fig6), and has a worst case execution time C_i on that processor (depicted to the left of each node). The designer can provide manually such worst-case times, or tools can be used in order to determine the worst-case execution time of a piece of code on a given processor [?]. For each message we know its size (in bytes, indicated to its left), and its period, which is identical with that of the sender process. Processes and messages activated based on events also have a uniquely assigned priority, $priority_{P_i}$ for processes and $priority_{m_i}$ for messages. All processes and messages belonging to a process graph G_i have the

same period $T_i = T_{G_i}$ which is the period of the process graph. A deadline D_{G_i} is imposed on each process graph G_i. Deadlines can also be placed locally on processes. Release times of some processes as well as multiple deadlines can be easily modelled by inserting dummy nodes between certain processes and the source or the sink node, respectively. These dummy nodes represent processes with a certain execution time but which are not allocated to any processing element.

7 Multi-Cluster Optimization

Considering the type of applications and systems described in the previous section, and using the analysis outlined in Section7, several design optimization problems can be addressed. In this section, we present problems which are characteristic to applications distributed across multi-cluster systems consisting of heterogeneous TT and ET networks:

- Section6.1 briefly outlines the problem of partitioning the processes of an application into time-triggered and event-triggered domains, and their mapping to the nodes of the clusters.

- Section6.2 presents the problem of packing of messages to frames, which is of utmost importance in cost-sensitive embedded systems where resources, such as communication bandwidth, have to be fully utilized [?, ?, ?]. This problem will be discussed in more detail in Section8.

The goal of these optimization problems is to produce an implementation which meets all the timing constraints (i.e., the application is schedulable). In order to drive our optimization algorithms towards schedulable solutions, we characterize a given frame packing configuration using the degree of schedulability of the application. The *degree of schedulability* [?] is calculated as:

$$
\delta_\Gamma = \begin{cases} c_1 = \sum\limits_{i=1}^{n} \max(0, r_i - D_i), \text{ if } c_1 > 0 \\ \\ c_2 = \sum\limits_{i=1}^{n} (r_i - D_i), \text{ if } c_1 = 0 \end{cases}
\tag{1}
$$

where n is the number of processes in the application, r_i is the worst-case response time of a process P_i, and D_i its deadline. The worst-case response times are calculated by the `MultiClusterScheduling` algorithm using the response time analysis presented in Section7. If the application is not schedulable, the term c_1 will be positive, and, in this case, the cost function is equal to c_1. However, if the process set is schedulable, $c_1 = 0$ and we use c_2 as a 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 c_2 means that we have improved the worst-case response times of the processes, so the application can potentially be implemented on a cheaper hardware architecture (with slower processors and/or buses). Improving the degree of schedulability can also lead to an improvement in the quality of control for control applications.

7.1 Partitioning and Mapping

By partitioning we denote the decision whether a certain process should be assigned to the TT or the ET domain (and, implicitly, to a TTC or an ETC, respectively). Mapping a process means assigning it to a particular node inside a cluster. Very often, the partitioning decision is taken based on the experience and preferences of the designer, considering aspects like the functionality implemented by the process, the hardness of the constraints, sensitivity to jitter, legacy constraints, etc. Let P be the set of processes in the application Γ. We denote with $P_T \subseteq P$ the subset of processes which the designer has assigned to the TT cluster, while $P_E \subseteq P$ contains processes which are assigned to the ET cluster. Many processes, however, do not exhibit

Figure 7: Frame-Packing Optimization Example

certain particular features or requirements which obviously lead to their implementation as TT or ET activities. The subset $P^+ = P \setminus (P_T \cup P_E)$ of processes could be assigned to any of the TT or ET domains. Decisions concerning the partitioning of this set of activities can lead to various trade-offs concerning, for example, the schedulability properties of the system, the amount of communication exchanged through the gateway, the size of the schedule tables, etc. For part of the partitioned processes, the designer might have already decided their mapping. For example, certain processes, due to constraints like having to be close to sensors/actuators, have to be physically located in a particular hardware unit. They represent the sets $P_T^M \subseteq P_T$ and $P_E^M \subseteq P_E$ of already mapped TT and ET processes, respectively. Consequently, we denote with $P_T^* = P_T \setminus P_T^M$ the TT processes for which the mapping has not yet been decided, and similarly, with $P_E^* = P_E \setminus P_E^M$ the unmapped ET processes. The set $P^* = P_T^* \cup P_E^* \cup P^+$ then represents all the unmapped processes in the application. The mapping of messages is decided implicitly by the mapping of processes. Thus, a message exchanged between two processes on the TTC (ETC) will be mapped on the TTP bus (CAN bus) if these processes are allocated to different nodes. If the communication takes place between two clusters, two message instances will be created, one mapped on the TTP bus and one on the CAN bus. The first message is sent from the sender node to the gateway, while the second message is sent from the gateway to the receiving node. Using the notation introduced, the partitioning and mapping problem can be described more exactly as follows. As an input we have an application Γ given as a set of process graphs and a two-cluster system consisting of a TT and an ET cluster. As introduced previously, P_T and P_E are the sets of processes already partitioned into TT and ET, respectively. Also, $P_T^M \subseteq P_T$ and $P_E^M \subseteq P_E$ are the sets of already mapped TT and ET processes. We are interested to find a partitioning for processes in $P^+ = P \setminus (P_T \cup P_E)$ and decide a mapping for processes in $P^* = P_T^* \cup P_E^* \cup P^+$, where $P_T^* = P_T \setminus P_T^M$, and $P_E^* = P_E \setminus P_E^M$ such that imposed deadlines are guaranteed to be satisfied.

7.2 Frame Packing

In both the TTP and CAN protocols messages are not sent independently, but several messages having similar timing properties are usually packed into frames. In many application areas, like automotive electronics, messages range from one single bit (e.g., the state of a device) to a couple of bytes (e.g., vehicle speed, etc.). Transmitting such small messages one per frame would create a high communication overhead, which can cause long delays leading to an unschedulable system. For example, 65 bits have to be transmitted on CAN for delivering one single bit of application data. Moreover, a given frame configuration defines the exact behavior of a node on the network, which is very important when integrating nodes from different suppliers.

Let us consider the motivational example in Figure reffig:fig7, where we have the process graph from Figure reffig:fig7d mapped on the two-cluster system from Figure reffig:fig7e: P_1 and P_4 are mapped on node N_1 from the TTC, while P_2 and P_3 are mapped on N_2 from ETC. The data field of the frames is represented with a black rectangle, while the other frame fields are depicted with a grey color. We consider a physical implementation of the buses such that the frames will take the time indicated in the figure by the length of their rectangles. We are interested to find a frame configuration such that the application is schedulable. In the system configuration of Figure reffig:fig7a we consider that, on the TTP bus, the node N_1 transmits in the first slot (S_1) of the TDMA round, while the gateway transmits in the second slot (S_G). Process P_3 has a higher priority than process P_2, hence P_2 will be interrupted by P_3 when it receives message m_2. In such a setting, P_4 will miss its deadline, which is depicted as a thick vertical line in Figure reffig:fig7. Changing the frame configuration as in Figure reffig:fig7b, so that messages m_1 and m_2 are packed into frame f_1 and slot S_G of the gateway comes first, processes P_2 and P_3 will receive m_1 and m_2 sooner and thus reduce the worst-case response time of the process graph, which is still larger than the deadline. In Figure reffig:fig7c, we also pack m_3 and m_4 into f_2. In such a situation,

the sending of m_3 will have to be delayed until m_4 is queued by P_2. Nevertheless, the worst-case response time of the application is further reduced, which means that the deadline is met, thus the system is schedulable.

However, packing more messages will not necessarily reduce the worst-case response times further, as it might increase too much the worst-case response times of messages that have to wait for the frame to be assembled, like is the case with message m_3 in Figure reffig:fig7c. This design optimization problem can be formulated more exactly as follows. As input to the frame-packing problem we have an application Γ given as a set of process graphs mapped on an architecture consisting of a TTC and an ETC interconnected through a gateway. We consider that the partitioning and mapping of processes has been already decided. We are interested to find a mapping of messages to frames (a frame packing configuration) denoted by a 4-tuple $\Psi = <\alpha, \pi, \beta, \sigma>$ such that the application Γ is schedulable. Once a schedulable system is found, we are interested to further improve the "degree of schedulability" so the application can potentially be implemented on a cheaper hardware architecture (with slower buses and processors). Determining a frame configuration Ψ means deciding on:

- The mapping of application messages transmitted on the ETC to frames (the set of ETC frames α), and their relative priorities, π. Note that the ETC frames α have to include messages transmitted from an ETC node to a TTC node, messages transmitted inside the ETC cluster, and those messages transmitted from the TTC to the ETC.

- The mapping of messages transmitted on the TTC to frames, denoted by the set of TTC frames β, and the sequence σ of slots in a TDMA round. The slot sizes are determined based on the set β, and are calculated such that they can accommodate the largest frame sent in that particular slot. We consider that messages transmitted from the ETC to the TTC are not statically allocated to frames. Rather, we will dynamically pack messages originating from the ETC into the "gateway frame", for which we have to decide the data field length (see Section5.2).

Several details related to the schedulability analysis were omitted from the discussion of the example. These details will be discussed in the next section.

8   Multi-Cluster Analysis and Scheduling

Once a partitioning and a mapping is decided, and a frame packing configuration is fixed, the processes and messages have to be scheduled. For the TTC this means building the schedule tables, while for the ETC the priorities of the ET processes have to be determined and their schedulability has to be analyzed. The analysis presented in this section works under the following assumptions:

- All the processes belonging a process graph G have the same period T_G. However, process graphs can have different periods.

- The offsets are *static* (as opposed to *dynamic* [?]), and are smaller than the period.

- The deadlines are arbitrary, i.e., can be larger than the period.

The basic idea is that on the TTC an application is schedulable if it is possible to build a schedule table such that the timing requirements are satisfied. On the ETC, the answer whether or not a system is schedulable is given by a *schedulability analysis*. In this paper, for the ETC we use a *response time analysis*, where the schedulability test consists of the comparison between the worst-case response time r_i of a process P_i and its deadline D_i. Response time analysis of data dependent processes with static priority preemptive scheduling has been proposed in [43, 68, 80], and has been also extended to consider the CAN protocol [?]. The authors use the concept of *offset* in order to handle data dependencies. Thus, each process P_i is characterized by an

Figure 8: The MultiClusterScheduling Algorithm

offset O_i, measured from the start of the process graph, that indicates the earliest possible start time of P_i. Such an offset is, for example, O_2 in Figure reffig:fig7a, as process P_2 cannot start before receiving m_1. The same is true for messages, their offset indicating the earliest possible transmission time. The response time analysis employed is presented in Section7.1. However, determining the schedulability of an application mapped on a multi-cluster system cannot be addressed separately for each type of cluster, since the inter-cluster communication creates a circular dependency: the static schedules determined for the TTC influence through the offsets the worst-case response times of the processes on the ETC, which on their turn influence the schedule table construction on the TTC. In Figure reffig:fig7b packing m_1 and m_2 in the same frame leads to equal offsets for P_2 and P_3. Because of this, P_3 will delay P_2 (which would not be the case if m_2 sent to P_3 would be scheduled in round 3, for example) and thus the placement of P_4 in the schedule table has to be accordingly delayed to guarantee the arrivals of m_3 and m_4. In our analysis we consider the influence between the two clusters by making the following observations:

- The start time of process P_i in a schedule table on the TTC is its offset O_i.

- The worst-case response time r_i of a TT process is its worst case execution time, i.e. $r_i = C_i$ (TT processes are not preemptable).

- The worst-case response times of the messages exchanged between two clusters have to be calculated according to the schedulability analysis described in Section7.1.

- The offsets have to be set by a scheduling algorithm such that the precedence relationships are preserved. This means that, if process P_B depends on process P_A, the following condition must hold: $O_B \geq O_A + r_A$. Note that for the processes on a TTC which receive messages from the ETC this translates to setting the start times of the processes such that a process is not activated before the worst-case arrival time of the message from the ETC. In general, offsets on the TTC are set such that all the necessary messages are present at the process invocation.

The `MultiClusterScheduling` algorithm in Figure reffig:fig8 receives as input the application Γ, the frame configuration Ψ, and produces the offsets ϕ and worst-case response times ρ.

The algorithm sets initially all the offsets to 0 (line 1). Then, the worst-case response times are calculated using the `ResponseTimeAnalysis` function (line 4) using the analysis presented in Section7.1. The fixed-point iterations that calculate the response times at line 3 will converge if processor and bus loads are smaller than 100% [**?**]. Based on these worst-case response times, we determine new values $\phi^n ew$ for the offsets using a list scheduling algorithm (line 6).We now have a schedule table for the TTC and worst-case response times for the ETC, which are pessimistic. The following loop will reduce the pessimism of the worst-case response times. The multi-cluster scheduling algorithm loops until the degree of schedulability δ_Γ of the application Γ cannot be further reduced (lines 820). In each loop iteration, we select a new offset from the set of $\phi^n ew$ offsets (line 10), and run the response time analysis (line 11) to see if the degree of schedulability has improved (line 12). If δ_Γ has not improved, we continue with the next offset in $\phi^n ew$. When a new offset $O_i^n ew$ leads to an improved δ_Γ we exit the for-each loop 919 that examines offsets from $\phi^n ew$. The loop iteration 820 continues with a new set of offsets, determined by `ListScheduling` at line 15, based on the worst-case response times $\rho^n ew$ corresponding to the previously accepted offset. In the multi-cluster scheduling algorithm, the calculation of offsets is performed by the list scheduling algorithm presented in Figure reffig:fig9. In each iteration, the algorithm visits the processes and messages in the `ReadyList`. A process or a message in the application is placed in the `ReadyList` if all its predecessors have been already scheduled. The list is ordered based on the priorities presented in [**?**]. The algorithm terminates when all processes and messages have been visited. In each loop iteration, the algorithm calculates the earliest time moment (*offset*) when the process or message $node_i$ can start (lines 57). There are four situations:

Figure 9: ListScheduling Algorithm

Figure 10: Frame Scheduling on the TTC

1. The visited node is an ET message. The message m_i is packed into its frame f (line 9), and the offset O_f of the frame is updated. The frame can only be transmitted after all the sender processes that pack messages in this frame have finished executing. The offset of message m_i packed to frame f is equal to the frame offset O_f.

2. The node is a TT message. In this case, when the frame is ready for transmission, it is scheduled using the `ScheduleTTFrame` function (presented in Figure reffig:fig10), which returns the *round* and the *slot* where the frame has been placed (line 16 in Figure reffig:fig9). In Figure reffig:fig10, the round immediately following *offset* is the initial candidate to be considered (line 2). However, it can be too late to catch the allocated slot, in which case the next round is considered (line 4). For this candidate round, we have to check if the slot is not occupied by another frame. If so, the communication has to be delayed for another round (line 7). Once a frame has been scheduled, we can determine the offsets and worst-case response times (Figure reffig:fig9, line 18). For all the messages in the frame the offset is equal to the start of the slot in the TDMA round, and the worst-case response time is the slot length.

3. The algorithm visits a process P_i mapped on an ETC node. A process on the ETC can start as soon as its predecessors have finished and its inputs have arrived, hence $O_i = offset$ (line 22). However, P_i might experience, later on, interference from higher priority processes.

4. Process P_i is mapped on a TTC node. In this case, besides waiting for the predecessors to finish executing, P_i will also have to wait for its processor $M(P_i)$ to become available (line 25). The earliest time when the processor is available is returned by the `ProcessorAvailable` function.

Let us now turn the attention back to the multi-cluster scheduling algorithm in Figure reffig:fig8. The algorithm stops when the δ_Γ of the application Γ is no longer improved, or when a limit imposed on the number of iterations has been reached. Since in a loop iteration we do not accept a solution with a larger δ_Γ, the algorithm will terminate when in a loop iteration we are no longer able to improve δ_Γ by modifying the offsets.

8.1 Schedulability Analysis for the ETC

For the ETC we use a response time analysis. A *response time analysis* has two steps. In the first step, the analysis derives the worst-case response time of each process (the time it takes from the moment is ready for execution, until it has finished executing). The second step compares the worst case response time of each process to its deadline and, if the response times are smaller or equal to the deadlines, the system is schedulable. The analysis presented in this section is used in the `ResponseTimeAnalysis` function (line 4 of the algorithm in Figure reffig:fig8) . Thus, the response time analysis in [**?**] uses the following equation for determining the worst-case response time r_i of a process P_i: (2)

where C_i is the worst-case execution time of process P_i, T_j is the period of process P_j, and $hp(P_i)$ denotes the set of processes that have a priority higher than the priority of P_i. The summation term, representing the interference I_i of higher priority processes on P_i, increases monotonically in r_i, thus solutions can be found using a recurrence relation. Moreover, the recurrence relations that calculate the worst case response time are guaranteed to converge if the processor utilization is under 100%. The previously presented analysis assumes that the deadline of a process is smaller or equal to its period. This assumption has later been relaxed [**?**] to consider

arbitrary deadlines (i.e., deadlines can be larger than the period). Thus, the worst-case response time r_i of a process P_i becomes:

(2)

where J_i is the jitter of process P_i (the worst-case delay between the arrival of a process and the start of its execution), q is the number of busy periods being examined, and $w_i(q)$ is the width of the level-i busy period starting at time qT_i. The level-i busy period is defined as the maximum time a processor executes processes of priority greater than or equal to the priority of process P_i, and is calculated as [**?**]:

(3)

.

The pessimism of the previous analysis can be reduced by using the information related to the precedence relations between processes. The basic idea is to exclude certain worst case scenarios, from the critical instant analysis, which are impossible due to precedence constraints. Methods for schedulability analysis of data dependent processes with static priority preemptive scheduling have been proposed in [**?**, **?**, **?**, **?**]. They use the concept of *offset* (or *phase*), in order to handle data dependencies. In [**?**] Tindell shows that the pessimism of the analysis is reduced through the introduction of offsets. The offsets have to be determined by the designer. In their analysis [**?**], the response time of a process P_i is:

(4)

where T_G the period of the process graph G, O_i and O_j are offsets of processes P_i and P_j, respectively, and J_i and J_j are the release jitters of P_i and P_j. In Equation refeq:eq5, the level-i busy period starting at time qT_G is:

(5)

.

In the previous equation, the blocking term B_i represents interference from lower priority processes that are in their critical section and cannot be interrupted, and C_i represents the worst-case execution time of process P_i. The last term captures the interference I_i from higher priority processes in the application, including higher priority processes from other process graphs. The reader is directed to [**?**] for the details of the interference calculation. Although this analysis is exact (both necessary and sufficient), it is computationally infeasible to evaluate. Hence, [**?**] proposes a feasible but not exact analysis (sufficient but not necessary) for solving Equation refeq:eq5. Our implementations use the feasible analysis provided in [**?**] for deriving the worst-case response time of a process P_i. We are now interested to determine the worst-case response time of frames and the worst-case queuing delays experienced by a frame in a communication controller. Regarding

the worst-case response time of messages, we have extended the analysis from [**?**] and applied it for frames on the CAN bus:

(6)

.

In the previous equation J_f is the jitter of frame f which in the worst case is equal to the largest worst-case response time $r_{S(m)}$ of a sender process $S(m)$ which sends message m packed into frame f:

(7)

.

In Equation refeq:eq7, W_f is the *worst-case queuing delay* experienced by f at the communication controller, and is calculated as: (9)

where q is the number of busy periods being examined, and $w_f(q)$ is the width of the level-f busy period starting at time qT_f. Moreover, in Equation refeq:eq7, C_f is the worst-case time it takes for a frame f to reach the destination controller. On CAN, C_f depends on the frame configuration and the size of the data field, s_f, while on TTP it is equal to the slot size in which f is transmitted. The worst-case response time of message m packed into a frame f can be determined by observing that $r_m = r_f$. The worst-case queueing delay for a frame (W_f in Equation refeq:eq7) is calculated differently for each type of queue:

1. The output queue of an ETC node, in which case $W_f^{N_i}$ represents the worst-case time a frame f has to spend in the Out_{N_i} queue on ETC node N_i. An example of such a frame is the one containing message m_3 in Figure reffig:fig7a, which is sent by process P_2 from the ETC node N_2 to the gateway node N_G, and has to wait in the Out_{N_2} queue.

2. The TTP-to-CAN queue of the gateway node, in which case W_f^{CAN} is the worst-case time a frame f has to spend in the Out_{CAN} queue of node N_G. In Figure reffig:fig7a, the frame containing m_1 is sent from the TTC node N_1 to the ETC node N_2, and has to wait in the Out_{CAN} queue of gateway node N_G before it is transmitted on the CAN bus.

3. The CAN-to-TTP queue of the gateway node, where W_f^{TTP} captures the time f has to spend in the Out_{TTP} queue node N_G. Such a situation is present in Figure reffig:fig7a, where the frame with m_3 is sent from the ETC node N_2 to the TTC node N_1 through the gateway node N_G where it has to wait in the Out_{TTP} queue before it is transmitted on the TTP bus, in the S_G slot of node N_G.

On the TTC, the synchronization between processes and the TDMA bus configuration is solved through the proper synthesis of schedule tables, hence no output queues are needed. The frames sent from a TTC node to another TTC node are taken into account when determining the offsets, and are not involved directly in the ETC analysis. The next sections show how the worst queueing delays are calculated for each of the previous three cases.

8.1.1 Worst-case queuing delays in the Out_{Ni} and Out_{CAN} queues

The analyses for $W_f^{N_i}$ and W_f^{CAN} are similar. Once f is the highest priority frame in the Out_{CAN} queue, it will be sent by the gateways CAN controller as a regular CAN frame, therefore the same equation for w_f can be used:

(8)

.

The intuition is that f has to wait, in the worst case, first for the largest lower priority frame that is just being transmitted (B_f) as well as for the higher priority $f_j \in hp(f)$ frames that have to be transmitted ahead of f (the second term). In the worst case, the time it takes for the largest lower priority frame $f_k \in lp(f)$ to be transmitted to its destination is:

(9)

.

Note that in our case, $lp(f)$ and $hp(f)$ also include messages produced by the gateway node, transferred from the TTC to the ETC.

8.1.2 Worst-case queuing delay in the Out_{TTP} queue

The time a frame f has to spend in the Out_{TTP} queue in the worst case depends on the total size of messages queued ahead of f (Out_{TTP} is a FIFO queue), the size S_G of the data field of the frame fitting into the gateway slot responsible for carrying the CAN messages on the TTP bus, and the period $T_T DMA$ with which this slot S_G is circulating on the bus [?]:

(10)

where I_f is the total size of the frames queued ahead of f. Those frames $f_j \in hp(f)$ are ahead of f, which have been sent from the ETC to the TTC, and have higher priority than f:

(11)

where the frame jitter J_j is given by Equation refeq:eq8. The blocking term B_f is the time interval in which f cannot be transmitted because the slot S_G of the TDMA round has not arrived yet. In the worst case (i.e., the frame f has just missed the slot S_G), the frame has to wait an entire round $T_T DMA$ for the slot S_G in the next TDMA round.

Figure 11: The General Frame Packing Strategy

Figure 12: The Simulated Annealing Algorithm

9 Frame-Packing Optimization Strategy

The general multi-cluster optimization strategy is outlined in Figure reffig:fig11. The `MultiClusterConfiguration` strategy has two steps:

1. In the first step, line 3, the application is partitioned on the TTC and ETC clusters, and processes are mapped to the nodes of the architecture using the `PartitioningAndMapping` function. The partitioning and mapping can be done with an optimization heuristic like the one presented in [**?**]. As part of the partitioning and mapping process, an initial frame configuration $\Psi^0 = <\alpha^0, \pi^0, \beta^0, \sigma^0>$ is derived. Messages exchanged by processes partitioned to the TTC will be mapped to TTC frames, while messages exchanged on the ETC will be mapped to ETC frames. For each message sent from a TTC process to an ETC process, we create an additional message on the ETC, and we map this message to an ETC frame. The sequence σ^0 of slots for the TTC is decided by assigning in order nodes to the slots $(S_i = N_i)$. One message is assigned per frame in the initial set β^0 of TTC frames. For the ETC, the frames in the set α^0 initially hold each one single message, and we calculate the message priorities Π^0 based on the deadlines of the receiver processes.

2. The frame packing optimization, is performed as the second step (line 5 in Figure reffig:fig11). The `FramePackingOptimization` function receives as input the application Γ, the mapping M of processes to resources and the initial frame configuration Ψ^0, and it produces as output the optimized frame packing configuration Ψ. Such an optimization problem is NP complete [**?**], thus obtaining the optimal solution is not feasible. In this paper, we propose two frame packing optimization strategies, one based on a simulated annealing approach, presented in Section8.1, while the other, outlined in Section8.2, is based on a greedy heuristic that uses intelligently the problem-specific knowledge in order to explore the design space.

If after these steps the application is unschedulable, we conclude that no satisfactory implementation could be found with the available amount of resources. Testing if the application Γ is schedulable is done using the `MultiClusterScheduling` (MCS) algorithm (line 7 in Figure reffig:fig11). The multi-cluster scheduling algorithm, presented in Figure reffig:fig8, takes as input an application Γ, a mapping M and an initial frame configuration Ψ^0, builds the TT schedule tables, sets the ET priorities for processes, and provides the global analysis.

9.1 Frame Packing with Simulated Annealing

The first algorithm we have developed is based on a simulated annealing (SA) strategy [**?**], and is presented in Figure reffig:fig12. The algorithm takes as input the application Γ, a mapping M and an initial frame configuration Ψ^0, and determines the frame configuration Ψ which leads to the best degree of schedulability δ_Γ (the smaller the value, the more schedulable the system, see Section6). Determining a frame configuration Ψ means finding the set of ETC frames α and their relative priorities π, and the set of TTC frames β, including the sequence σ of slots in a TDMA round. The main feature of a SA strategy is that it tries to escape from a local optimum by randomly selecting a new solution from the neighbors of the current solution. The new solution is accepted if it is an improved solution (lines 910 of the algorithm in Figure reffig:fig12). 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 (lines 1213). In Figure reffig:fig12 we give a short description of this algorithm. An essential component of the algorithm is the generation of a new solution Ψ_new starting from the current one $\Psi_current$. The neighbors of

Figure 13: The OptimizeFramePacking Algorithm

the current solution $\Psi_c urrent$ are obtained by performing transformations (called moves) on the current frame configuration $\Psi_c urrent$ (line 8). We consider the following moves:

- moving a message m from a frame f_1 to another frame f_2 (or moving m into a separate single-message frame);

- swapping the priorities of two frames in α;

- swapping two slots in the sequence σ of slots in a TDMA round.

For the implementation of this algorithm, the parameters TI (initial temperature), TL (temperature length), ϵ(cooling ratio), and the stopping criterion have to be determined. They define the "cooling schedule" and have a decisive impact on the quality of the solutions and the CPU time consumed. We are interested to obtain values for TI, TL and ϵthat will guarantee the finding of good quality solutions in a short time. We performed long runs of up to 48 hours with the SA algorithm, for ten synthetic process graphs (two for each graph dimension of 80, 160, 240 320, 400, see Section9) and the best ever solution produced has been considered as the optimum. Based on further experiments we have determined the parameters of the SA algorithm so that the optimization time is reduced as much as possible but the near-optimal result is still produced. For example, for the graphs with 320 nodes, TI is 700, TL is 500 and ϵ is 0.98. The algorithm stops if for three consecutive temperatures no new solution has been accepted.

9.2 Frame Packing Greedy Heuristic

The `OptimizeFramePacking` greedy heuristic (Figure reffig:fig13) constructs the solution by progressively selecting the best candidate in terms of the degree of schedulability.

We start by observing that all activities taking place in a multi-cluster system are ordered in time using the offset information, determined in the `StaticScheduling` function based on the worst-case response times known so far and the application structure (i.e., the dependencies in the process graph). Thus, our greedy heuristic outlined in Figure reffig:fig13, starts with building two lists of messages ordered according to the ascending value of their offsets, one for the TTC, $messages_\beta$, and one for ETC, $messages_\alpha$. Our heuristic is to consider for packing in the same frame messages which are adjacent in the ordered lists. For example, let us consider that we have three messages, m_1 of 1 byte, m_2 of 2 bytes and m_3 of 3 bytes, and that messages are ordered as m_3, m_1, m_2 based on the offset information. Also, assume that our heuristic has suggested two frames, frame f_1 with a data field of 4 bytes, and f_2 with a data field of 2 bytes. The `PackMessages` function will start with m_3 and pack it in frame f_1. It continues with m_2, which is also packed into f_1, since there is space left for it. Finally, m_3 is packed in f_2, since there is no space left for it in f_1. The algorithm tries to determine, using the `for-each` loops in Figure reffig:fig13 the best frame configuration. The algorithm starts from the initial frame configuration Ψ^0, and progressively determines the best change to the current configuration. The quality of a frame configuration is measured using the `MultiClusterScheduling` algorithm, which calculates the degree of schedulability δ_Γ (line 13). Once a configuration parameter has been fixed in the outer loops it is used by the inner loops:

- Lines 1015: The innermost loops determine the best size S_α for the currently investigated frame f_α in the ETC frame configuration $\alpha_c urrent$. Thus, several frame sizes are tried (line 11), each with a size returned by `RecomendedSizes` to see if it improves the current configuration. The `RecomendedSizes`($messages_\alpha$) list is built recognizing that only messages adjacent in the $messages_\alpha$ list will be packed into the same frame. Sizes of frames are determined as a sum resulted from adding the sizes of combinations of adjacent messages, not exceeding 8 bytes. For the previous example, with m_1, m_2 and m_3, of 1, 2 and 3 bytes,

respectively, the frame sizes recommended will be of 1, 2, 3, 4, and 6 bytes. A size of 5 bytes will not be recommended since there are no adjacent messages that can be summed together to obtain 5 bytes of data.

- Lines 916: This loop determines the best frame configuration α. This means deciding on how many frames to include in α (line 9), and which are the best sizes for them. In α there can be any number of frames, from one single frame to n_α frames (in which case each frame carries one single message). Once a configuration $\alpha_b est$ for the ETC, minimizing δ_Γ, has been determined (saved in line 16), the algorithm looks for the frame configuration β which will further improve δ_Γ.

- Lines 717: The best size for a frame f_β is determined similarly to the size for a frame f_α.

- Lines 618: The best frame configuration $\beta_b est$ is determined. For each frame configuration β tried, the algorithm loops again through the innermost loops to see if there are better frame configurations α in the context of the current frame configuration $\beta_c urrent$.

- Lines 419: After a $\beta_b est$ has been decided, the algorithm looks for a slot sequence σ, starting with the first slot and tries to find the node which, when transmitting in this slot, will reduce δ_Γ. Different slot sequences are tried by swapping two slots within the TDMA round (line 5).

For the initial message priorities π^0 (initially, there is one message per frame) we use the "heuristic optimized priority assignment" (HOPA) approach in [?], where priorities in a distributed real-time system are determined, using knowledge of the factors that influence the timing behavior, such that the degree of schedulability of the system is improved (line 1). The ETC message priorities set at the beginning of the algorithm are not changed by our greedy optimization loops. The priority of a frame $f_\alpha \in \alpha$ is given by the message $m \in f_\alpha$ with the highest priority. The algorithm continues in this fashion, recording the best ever $\Psi_b est$ configurations obtained, in terms of δ_Γ, and thus the best solution ever is reported when the algorithm finishes.

10   Experimental Results

For the evaluation of our frame-packing optimization algorithms we first used process graphs generated for experimental purpose. We considered two-cluster architectures consisting of 2, 4, 6, 8 and 10 nodes, half on the TTC and the other half on the ETC, interconnected by a gateway. Forty processes were assigned to each node, resulting in applications of 80, 160, 240, 320 and 400 processes. We generated both graphs with random structure and graphs based on more regular structures like trees and groups of chains. We generated a random structure graph deciding for each pair of two processes if they should be connected or not. Two processes in the graph were connected with a certain probability (between 0.05 and 0.15, depending on the graph dimension) on the condition that the dependency would not introduce a loop in the graph. The width of the tree-like structures was controlled by the maximum number of direct successors a process can have in the tree (from 2 to 6), while the graphs consisting of groups of chains had 2 to 12 parallel chains of processes. Furthermore, the regular structures were modified by adding a number of 3 to 30 random cross-connections. The mapping of the applications to the architecture has been done using a simple heuristic that tries to balance the utilization of processors while minimizing communication. Execution times and message lengths were assigned randomly using both uniform and exponential distribution within the 10 to 100 ms, and 1 bit to 2 bytes ranges, respectively. For the communication channels we considered a transmission speed of 256 Kbps and a length below 20 meters. All experiments were run on a SUN Ultra 10. The first result concerns the ability of our heuristics to produce schedulable solutions. We have compared the degree of schedulability $\delta_\Gamma obtained from our \texttt{OptimizeFramePacking}(OFP) heuristic (Figure ref fig : fig13) with the near-optimal values obtained by SA (Figure ref fig : fig12). Obtaining solutions that have a better degree of schedulability means o$

Figure 14: Hardware Architecture for the Cruise Controller

caseresponsetimes, increasingthechancesofmeetingthedeadlines.Table1presentstheaveragepercentagedeviationoftheo
optimalvaluesobtainedwithSA.TogetherwithOFP, astraightforwardapproach(SF)ispresented.TheSFapproachdoesno

In Table1 we have one row for each application dimension of 80 to 400 processes, and a header for each optimization algorithm considered. For each of the SF and OFP algorithms we have three columns in the table. In the first column, we present the average percentage deviation of the algorithm from the results obtained by SA. The percentage deviation is calculated according to the formula:

$$(12)$$

.

The second column presents the maximum percentage deviation from the SA result, and the third column presents the average execution time of the algorithm, in seconds. For the SA algorithm we present only its average execution times. Table1 shows that when packing messages to frames, the degree of schedulability improves dramatically compared to the straightforward approach. The greedy heuristic `OptimizeFramePacking` performs well for all the graph dimensions, having, for example, run-times which are on average under 50 for applications with 240 processes. When deciding on which heuristic to use for design space exploration or system synthesis, an important issue is the execution time. In average, our optimization heuristics needed a couple of minutes to produce results, while the simulated annealing approach had an execution time of up to 6 hours.

10.1 The Vehicle Cruise Controller

A typical safety critical application with hard real-time constraints, is a vehicle cruise controller (CC). We have considered a CC system derived from a requirement specification provided by the industry. The CC delivers the following functionality: it maintains a constant speed for speeds over 35 Km/h and under 200 Km/h, offers an interface (buttons) to increase or decrease the reference speed, and is able to resume its operation at the previous reference speed. The CC operation is suspended when the driver presses the brake pedal. The specification assumes that the CC will operate in an environment consisting of two clusters. There are four nodes which functionally interact with the CC system: the Anti-lock Braking System (ABS), the Transmission Control Module (TCM), the Engine Control Module (ECM), and the Electronic Throttle Module (ETM) (see Figure reffig:fig14). It has been decided to map the functionality (processes) of the CC over these four nodes. The ECM and ETM nodes have an 8-bit Motorola M68HC11 family CPU with 128 Kbytes of memory, while the ABS and TCM are equipped with a 16-bit Motorola M68HC12 CPU and 256 Kbytes of memory. The 16-bit CPUs are twice as fast than the 8-bit ones. The transmission speed of the communication channel is 256 Kbps and the frequency of the TTP controller was chosen to be 20 MHz. We have modeled the specification of the CC system using a set of 32 processes and 17 messages as described in [**?**], where the mapping of processes to the nodes is also given. The period was chosen 250 ms, equal to the deadline. In this setting, the straightforward approach SF produced an end-to-end worst-case response time of 320 ms, greater than the deadline, while both the OFP and SA heuristics produced a schedulable system with a worst-case response time of 172 ms. This shows that the optimization heuristic proposed, driven by our schedulability analysis, is able to identify that frame packing configuration which increases the schedulability degree of an application, allowing the developers to reduce the implementation cost of a system.

11 Conclusions

Heterogeneous distributed real-time systems are used in several application areas to implement increasingly complex applications that have tight timing constraints. The heterogeneity is manifested not only at the hardware and communication protocol levels, but also at the level of the scheduling policies used. In order to reduce costs and use the available resources more efficiently, the applications are distributed across several networks. We have introduced the current state-of-the-art analysis and optimization techniques available for such systems, and addressed in more detail a special class of heterogeneous distributed real-time embedded systems called multi-cluster systems. We have presented an analysis for multi-cluster systems and outlined several characteristic design problems, related to the partitioning and mapping of functionality, and the optimization of the access to the communication infrastructure. An approach to schedulability-driven frame packing for the synthesis of multi-cluster systems was presented as an example of solving such a design optimization problem. We have developed two optimization heuristics for frame configuration synthesis which are able to determine frame configurations that lead to a schedulable system. We have shown that by considering the frame packing problem, we are able to synthesize schedulable hard-real time systems and to potentially reduce the overall cost of the architecture. The main message of this paper is that efficient analysis and optimization methods are needed and can be developed for the efficient implementation of applications distributed over interconnected heterogeneous networks.

12 References

References

[1] Aires, http://kabru.eecs.umich.edu/aires/

[2] Ascet/SD, $http://en.etasgroup.com/products/ascet_sd/G.Agrawal, B.Chen, W.Zhao, S.Davari, "GuaranteeingSy$ $-339, March 1994.$

[3] N. C. Audsley, A. Burns, M. F. Richardson, A. J. Wellings, "Hard Real-Time Scheduling: The Deadline Monotonic Approach, in Proceedings of the 8th IEEE Workshop on Real-Time Operating Systems and Software, pages 127–132, 1991.

[4] N. Audsley, A. Burns, R. Davis, K. Tindell, A. Wellings, "Fixed Priority Preemptive Scheduling: An Historical Perspective, in Real-Time Systems, 8(2/3), 173-198, 1995.

[5] N. Audsley, K. Tindell, A. Burns, "The End of Line for Static Cyclic Scheduling?, in Proceedings of the Euromicro Workshop on Real-Time Systems, 36-41, 1993.

[6] F. Balarin et al., Hardware-Software Co-Design of Embedded Systems: The POLIS Approach, Kluwer Academic Publishers, Boston, 1997.

[7] F. Balarin, L. Lavagno, P. Murthy, A. Sangiovanni-Vincentelli, "Scheduling for Embedded Real-Time Systems, in IEEE Design and Test of Computers, January-March, 71-82, 1998.

[8] F. Balarin, Y. Watanabe, H. Hsieh, L. Lavagno, C. Paserone, A. Sangiovanni-Vincentelli, "Metropolis: An Integrated Electronic System Design Environment, Computer , volume 36, issue 4, pages 45–52, 2003.

[9] J. Berwanger, M. Peller, R. Griessbach, A New High Performance Data Bus System for Safety-Related Applications, http://www.byteflight.de, 2000.

[10] Robert Bosch GmbH, CAN Specification, Version 2.0, http://www.can.bosch.com/, 1991.

[11] A. Burns, A. Wellings, Real-Time Systems and Programming Languages, Addison Wesley, 2001.

[12] M. Chiodo, "Automotive Electronics: A Major Application Field for Hardware-Software Co-Design, in Hardware/ Software Co-Design, Kluwer Academic Publishers, pages 295– 310, 1996.

[13] B. P. Dave, N. K. Jha, "COHRA: Hardware-Software Cosynthesis of Hierarchical Heterogeneous Distributed Systems, in IEEE Transactions on CAD, 17(10), 900-919, 1998.

[14] B. P. Dave, G. Lakshminarayana, N. J. Jha, "COSYN: Hardware-Software Co-Synthesis of Heterogeneous Distributed Embedded Systems, in IEEE Transactions on VLSI Systems, 7(1), 92-104, 1999.

[15] EAST-EEA project, ITEA Full Project Proposal, http://www.itea-office.org, 2002.

[16] Echelon, LonWorks: The LonTalk Protocol Specification, http://www.echelon.com, 2003.

[17] P. Eles, A. Doboli, P. Pop, Z. Peng, "Scheduling with Bus Access Optimization for Distributed Embedded Systems, in IEEE Transactions on VLSI Systems, 472-491, 2000.

[18] H. Ermedahl, H. Hansson, M. Sjdin, "Response-Time Guarantees in ATM Networks, in Proceedings of the IEEE Real-Time Systems Symposium, pages 274–284, 1997.

[19] R. Ernst, "Codesign of Embedded Systems: Status and Trends, in IEEE Design Test of Computers, April-June, 1998.

[20] The FlexRay Group, FlexRay Requirements Specification, Version 2.0.2, http://www.flexray-group.com/, 2002.

[21] J. J. Gutirrez Garca, M. Gonzlez Harbour, "Optimized Priority Assignment for Tasks and Messages in Distributed Hard Real-Time Systems, in Proceedings of the Workshop on Parallel and Distributed Real-Time Systems, 124-132, 1995.

[22] P. Hansen, The Hansen Report on Automotive Electronics, http://www.hansenreport.com/, July–August, 2002.

[23] K. Hoyme, K. Driscoll, "SAFEbus, in IEEE Aerospace and Electronic Systems Magazine, volume 8, number 3, pages 34– 39, 1992.

[24] International Organization for Standardization, "Road vehiclesController area network (CAN)Part 4: Time-triggered communication, ISO/DIS 11898-4, 2002.

[25] J. Jonsson, K. G. Shin, "Robust Adaptive Metrics for Deadline Assignment in Distributed Hard Real-Time Systems, Real-Time Systems: The International Journal of Time-Critical Computing Systems, Vol. 23, No. 3, pages 239–271, 2002.

[26] K. Jost, "From Fly-by-Wire to Drive-by-Wire, Automotive Engineering International, 2001.

[27] K. Keutzer, S. Malik, A. R. Newton, "System-Level Design: Orthogonalization of Concerns and Platform-Based Design, in IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, volume 19, number 12, December 2000.

[28] B. Kienhuis, E. Deprettere, K. Vissers, P. Van Der Wolf, "An Approach for Quantitative Analysis of Application-Specific Dataflow Architectures, in Proceedings of the IEEE International Conference on Application-Specific Systems, Architectures and Processors, pages 338 –349, 1997.

[29] H. Kopetz: "Automotive Electronics–Present State and Future Prospects, 25th International Symposium on Fault-Tolerant Computing, 1995.

[30] H. Kopetz, Real-Time Systems – Design Principles for Distributed Embedded Applications, Kluwer Academic Publishers, 1997.

[31] H. Kopetz, "Automotive Electronics, in Proceedings of the 11th Euromicro Conference on Real-Time Systems, pages 132– 140, 1999.

[32] H. Kopez, R. Nossal, "The Cluster-Compiler – A Tool for the Design of Time Triggered Real-Time Systems, in Proceedings of the ACM SIGPLAN Workshop on Languages, Compilers, and Tools for Real-Time Systems, 108-116, 1995.

[33] C. Lee, M. Potkonjak, W. Wolf, "Synthesis of Hard Real-Time Application Specific Systems, in Design Automation for Embedded Systems, 4(4), 215-241, 1999.

[34] G. Leen, D. Heffernan, "Expanding automotive electronic systems, in Computer, Volume: 35, Issue: 1, Pages 88-93, 2002.

[35] Local Interconnect Network Protocol Specification, http://www.lin-subbus.org, 2003.

[36] C. L. Liu, J. W. Layland, "Scheduling Algorithms for Multiprogramming in a Hard-Real-Time Environment, in Journal of the ACM, 20(1), 46-61, 1973.

[37] H. Lnn, J. Axelsson, "A Comparison of Fixed-Priority and Static Cyclic Scheduling for Distributed Automotive Control Applications, in Proceedings of the Euromicro Conference on Real-Time Systems, 142-149, 1999.

[38] Matlab/Simulink, http://www.mathworks.com

[39] K. Melin, Volvo S80: Electrical System of the Future, Volvo Technology Report, 1998.

[40] P. S. Miner, "Analysis of the SPIDER Fault-Tolerance Protocols, in Proceedings of the 5th NASA Langley Formal Methods Workshop, 2000.

[41] T. Nolte, H. Hansson, C. Norstrm, S. Punnekkat, "Using bit-stuffing distributions in CAN analysis, in Proceedings of the IEEE/IEE Real-Time Embedded Systems Workshop, 2001

[42] J. C. Palencia, M. Gonzlez Harbour, "Schedulability Analysis for Tasks with Static and Dynamic Offsets, in Proceedings of the 19th IEEE Real-Time Systems Symposium, 26-37, 1998.

[43] J. C. Palencia, M. Gonzlez Harbour, "Exploiting Precedence Relations in the Schedulability Analysis of Distributed Real-Time Systems, in Proceedings of the 20th IEEE Real-Time Systems Symposium, pages 328–339, 1999.

[44] T. Pop, P. Eles, Z. Peng, "Holistic Scheduling and Analysis of Mixed Time/Event-Triggered Distributed Embedded Systems, in Proceedings of the International Symposium on Hardware/Software Codesign, 187-192, 2002.

[45] P. Pop, P. Eles, Z. Peng, "Scheduling with Optimized Communication for Time Triggered Embedded Systems, in Proceedings of the International Workshop on Hardware-Software Codesign, 178-182, 1999.

[46] P. Pop, P. Eles, Z. Peng, "Bus Access Optimization for Distributed Embedded Systems Based on Schedulability Analysis, in Proceedings of the Design Automation and Test in Europe Conference, 567-574, 2000.

[47] P. Pop, P. Eles, Z. Peng, "Schedulability Analysis and Optimization for the Synthesis of Multi-Cluster Distributed Embedded Systems, in Proceedings of Design Automation and Test in Europe Conference, 184-189, 2003.

[48] P. Pop, Analysis and Synthesis of Communication-Intensive Heterogeneous Real-Time Systems, Linkping Studies in Science and Technology, Ph.D. Dissertation No. 833, 2003, available at http://www.ida.liu.se/ paupo/thesis

[49] P. Pop, P. Eles, Z. Peng, "Schedulability-Driven Communication Synthesis for Time-Triggered Embedded Systems, in Real-Time Systems Journal, No. 24, 297–325, 2004.

[50] P. Pop, P. Eles, Z. Peng, V. Izosimov, M. Hellring, O. Bridal, "Design Optimization of Multi-Cluster Embedded Systems for Real-Time Applications, in Proceedings of Design, Automation and Test in Europe Conference, 1028–1033, 2004.

[51] Profibus International, PROFIBUS DP Specification, http://www.profibus.com/, 2003.

[52] P. Puschner, A. Burns, "A Review of Worst-Case Execution-Time Analyses, in Real-Time Systems Journal, Vol. 18, No. 2/3, May 2000.

[53] A. Rajnak, K. Tindell, L. Casparsson, Volcano Communications Concept, Volcano Communication Technologies AB, 1998.

[54] RapidRMA, http://www.tripac.com

[55] C. R. Reevs, Modern Heuristic Techniques for Combinatorial Problems, Blackwell Scientific Publications, 1993.

[56] K. Richter, M. Jersak, R. Ernst, "A Formal Approach to MpSoC Performance Verification, in Computer, volume 36, issue 4, pages 60–67, 2003.

[57] RTA-OSEK Planner, http://www.livedevices.com

[58] J. Rushby, "Bus Architectures for Safety-Critical Embedded Systems, Springer–Verlag Lecture Notes in Computer Science, volume 2211, pages 306–323, 2001.

[59] SAE Vehicle Network for Multiplexing and Data Communications Standards Committee, SAE J1850 Standard, 1994.

[60] K. Sandstrm, C. Norstrm, "Frame Packing in Real-Time Communication, in Proceedings of the International Conference on Real-Time Computing Systems and Applications, 399-403, 2000.

[61] Statemate, http://www.ilogix.com

[62] J. A. Stankovic, K. Ramamritham, Advances in Real-Time Systems, IEEE Computer Society Press, 1993.

[63] J. K. Strosnider, T. E. Marchok, "Responsive, Deterministic IEEE 802.5 Token Ring Scheduling, in Journal of Real-Time Systems, volume 1, issue 2, pages 133–158, 1989.

[64] SystemBuild/MatrixX, http://www.ni.com/matrixx

[65] B. Tabbara, A. Tabbara, A. Sangiovanni-Vincentelli, Function/Architecture Optimization and Co-Design of Embedded Systems, Kluwer Academic Publishers, 2000.

[66] TimeWiz, http://www.timesys.com

[67] K. Tindell, Adding Time-Offsets to Schedulability Analysis, Department of Computer Science, University of York, Report No. YCS-94-221, 1994.

[68] K. Tindell, A. Burns, A. Wellings, "Calculating CAN Message Response Times, in Control Engineering Practice, 3(8), 1163-1169, 1995.

[69] K. Tindell, J. Clark, "Holistic Schedulability Analysis for Distributed Hard Real-Time Systems, in Microprocessing Microprogramming, Vol. 50, No. 2-3, 1994.

[70] TTP-Plan, http://www.tttech.com/

[71] TTTech, TTP/C Specification Version 0.5, 1999, availabe at http://www.tttech.com/

[72] TTTech, Comparison CAN–Byteflight–FlexRay–TTP/C, Technical Report, availabe at http://www.tttech.com/

[73] Virtual Component Co-design, http://www.cadence.com/

[74] Volcano Network Analyzer, http://www.volcanoautomotive.com/

[75] W. Wolf, "A Decade of Hardware/Software Codesign, in Computer, 36/4, 38–43, 2003.

[76] X-by-Wire Consortium, X-By-Wire: Safety Related Fault Tolerant Systems in Vehicles, http://www.vmars.tuwien.ac.at/projects/xbywire/, 1998.

[77] J. Xu, D. L. Parnas, "On satisfying timing constraints in hard-real-time systems, in IEEE Transactions on Software Engineering, 19(1), 1993.

[78] J. Xu, D. L. Parnas, "Priority Scheduling Versus Pre-Run-Time Scheduling, in Journal of Real Time Systems, volume 18, issue 1, pages 7–24, 2000.

[79] T. Y. Yen, W. Wolf, Hardware-Software Co-Synthesis of Distributed Embedded Systems, Kluwer Academic Publishers, 1997.