

Design Process and Tools for the HW/SW-Codesign and Rapid-Prototyping of Parallel and Heterogeneous Real-Time Communication Systems¹

Frank Slomka, Matthias Dörfel, Simone Spitz², Richard Hofmann
University of Erlangen-Nuremberg
Department of Computer Architecture and Performance Evaluation
Martensstr. 3, 91058 Erlangen, Germany

Abstract

Based on the methodology for the development of communication systems, a framework for early performance evaluation and hardware/software codesign of such systems is presented. We describe how performance requirements can be formulated in a formal way during the analysis phase of a project and how these requirements are used by the synthesis tools for hardware/software codesign during the design and implementation phase of the development cycle. Additionally, the concepts of our codesign tools and a heterogeneous rapid-prototyping environment are described. It is shown, how our prototyping board can be used for validating the results of the early performance evaluation. We demonstrate the approach on an example of a real-time communication system.

1 Introduction

The development of heterogeneous communication systems for performance critical applications as real-time video-conference and multimedia requires to bridge the gap between specification-level analysis techniques for communication protocols and the synthesis of embedded systems. This leads to an integrated, tool-supported development cycle from the analysis to the design phase [6][11] and the automatic implementation of formal specifications on configurable rapid-prototyping platforms. To validate as many design alternatives of the real-time communication system as possible, fast synthesis and integrated measurement techniques are needed. This includes the automatic formulation of constraints from the early analysis phase in the specification language used in the design phase. The resulting specification then will automatically be synthesized to the hardware description and software implementation of the system.

In the remainder, we describe a development process and tools for supporting the rapid development

of communication systems. To support the complete development cycle from beginning to the end, we use the formal specification languages SDL (Specification and Description Language [13]) and MSC (Message Sequence Chart [14]). We extended these standardized languages with annotations that support performance evaluation in early design phases and integrate the hardware/software codesign of embedded systems [19] in the methodology. Our extension of MSC is called PMSC [7] and our extension of SDL is called SDL* [28]. The specification of design constraints in PMSC and SDL* are compatible, so it is possible to transfer constraints from the PMSC specification to the corresponding SDL* specification during the development of the system. The support of the complete design cycle is the difference to the SDL based approaches for codesign described in [2] and in [10]. The main focus of this papers was the semiautomatic [2] or automatic implementation [10] of hardware/software systems described with SDL. Early design phases, the specification of performance constraints with MSC, and the integration of the measurement of constraints in a rapid-prototyping environment are not considered in the papers described above.

This paper continues with a short introduction to the design methodology. In section 3 we describe the different codesign tools used in our design process and section 4 regards the rapid-prototyping hardware and the measurement tools. Finally, a short overview of a case study of a multimedia system is presented in section 5.

2 Design Process

2.1 Designing Communication Systems

As introduced in section 1, the design process of a communication system starts with the analysis

1. This work has been funded by the Deutsche Forschungsgemeinschaft (DFG) under grant He1408/4-2.
2. Simone Spitz is now with Siemens AG, Switzerland.

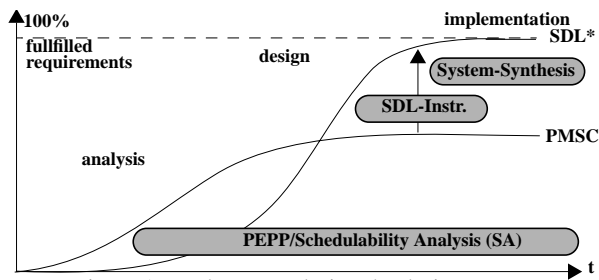


Figure 1: Tool support during the design process

phase. In this phase, the functionality and the requirements of the system are determined. Using the language MSC, the interaction between different system components is described formally. This description is augmented by performance requirements, stated in the extension PMSC. Therefore, PMSC allows to describe the estimated timing behavior of system components to perform an early performance analysis of the overall system [7]. This can be a stochastic graph analysis using the analysis tool PEPP [3] or a real time schedulability analysis as described in [15] and [27].

Figure 1 shows the fraction of fulfilled requirements over the time of a complete design process within our framework. In addition it is shown in figure 1 which tool of figure 2 is used at the different phases of the design cycle. As discussed above, the designer first specifies the systems interaction and the performance requirements using PMSC. After the system components and its communication relationships are specified, the engineer starts to develop an SDL specification. During the analysis and the design phase of the system the designer often refines the PMSC and the SDL specification [19], e.g. if the current PMSC does not keep the systems requirements and the system components or communication mechanisms in the specification were changed.

This phase results in a complete PMSC specification and an SDL specification that contains the block and processes structure. Next the designer specifies the complete functional behavior of the different SDL processes. At his time, the SDL specification contains more detailed information than the PMSC specification: In the PMSC specification, only requirements of the system and the communication interactions are described. In contrast to that, the SDL specification determines the complete functional behavior.

As can be seen in figure 1, it is not possible to specify the complete functionality of the system with PMSC. In addition to the dynamic behavior of the system, the functional behavior must be specified with SDL.

To use the system synthesis tools described in section 3, the performance requirements specified with PMSC must be transferred to the SDL specifi-

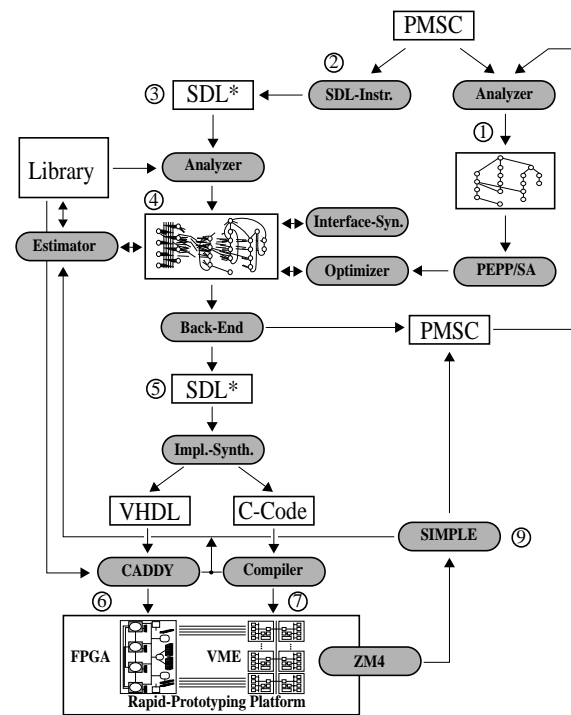


Figure 2: Tool flow

cation. To describe such performance and other system requirements we defined a few extensions to SDL [19][28]. The new language is called SDL*.

Using the SDL instrumentation tool described in [4] all performance requirements given in the PMSC specification are transferred automatically to the SDL* specification. Additionally it is possible in SDL* to specify the mapping of SDL functionality to hardware components to support a semiautomatic hardware/software codesign. Now the SDL* specification contains all informations needed by the design tools described in section 3 to implement the specification on the rapid-prototyping platform described in section 4.

2.2 Tool Flow for Codesign and Prototyping

In figure 2, a complete overview of the SDL*/PMSC framework is given. To support the early performance evaluation of the system, the framework starts with the PMSC analysis (1). Based on PMSC, different analysis techniques are possible: A stochastic graph analysis for communication systems [3], a real-time schedulability analysis based on the event stream model for hard real-time systems [15] and a schedulability analysis for heterogeneous systems [27]. To support all the different analysis techniques, the PMSC specification is translated to an internal graph representation [17].

On the other hand a complete synthesis path is supported by the framework. As described in sec-

tion 2.1, it is possible to transfer performance requirements specified in PMSC to the corresponding SDL* specification (see also figure 4 and figure 5). This is done by a SDL/MSC instrumentation tool (2). This tool also allows to use PMSC for the definition of measurement points in the implementation. This technique is called *specification driven monitoring* [4]. The data generated by the measurement tools (9) are used to validate the real-time constraints.

After all, performance requirements are transferred to the SDL* specification (3), which in turn is translated into an internal graph representation. The functional behavior of the SDL* specification is translated to a cyclic control-flow (CFG) graph. This graph is mapped on an architecture graph of the rapid-prototyping hardware, which is loaded from a library (4). Furthermore, the SDL* specification may contain instructions about the mapping of functional units to architecture modules. In the model, this is described by the mapping of nodes of the CFG to nodes of the architecture graph.

This internal representation of the SDL* specification is used by different design tools: an estimation tool (see section 3.1), which estimates the costs and the timing behavior of graph nodes on different architectures, a communication synthesis tool (see section 3.2), which decides how communication links of the SDL* specification will be implemented and an optimization tool based on tabu search [22] to decide the mapping of the nodes of the CFG to hardware- or software components (see section 3.3).

The internal graph representation can be retranslated to a SDL* (5) or PMSC specification. After the optimization process, the SDL* specification contains all mapping information needed by the following implementation synthesis tool (see section 3.4). Generating a PMSC specification with estimated execution times for the allocated architecture components, it is possible to use the PMSC performance analysis tools again.

The SDL* specification is translated into conventional implementation languages as VHDL for the hardware modules and C for the software parts of the system. This is done by the implementation synthesis tool.

The VHDL description is synthesized by the architecture synthesis system CADDY [1] and special FPGA synthesis tools (6). The software part of the system is translated by conventional compilers (7).

After the SDL* specification is implemented on the rapid-prototyping hardware (see section 4) the performance of the system can be measured by the hardware monitor ZM4 [12]. The results of the measurement will be evaluated by the tool SIMPLE [12] (9). Using this data, it is possible to back-

annotate the estimated values of the CFG and to iterate the optimization loop.

3 Design Tools

3.1 Estimation of execution times and costs

For each node of the CFG, the software and hardware execution time and the costs of the allocated architecture components are estimated. The internal graph representation contains two different types of nodes: Nodes with an underlying data flow graph (DFG) and nodes which must be mapped on library components, e.g. signal queues, timer modules and communication interfaces. The software and hardware execution times as well as the costs of elements are loaded from the component library shown on the left hand side of figure 2.

For each DFG, the estimation tool calculates the hardware execution time and costs by using the LEFT-EDGE [18] algorithm. For a given DFG, the number of resources needed is calculated to guarantee the shortest latency. Based on the calculated hardware resources, the area of the application specific circuit is estimated. To map more than one node of the control flow graph on one specific hardware implementation, the estimation algorithm described in [29] is used. For estimating software execution times, a different algorithm calculates the latency of the given data flow graph based on a restricted number of resources (registers/ALUs) with fixed costs.

After synthesis and implementation on the prototyping board, the execution times calculated by CADDY or measured by the SIMPLE monitoring software are used to refine the estimation.

3.2 Interface Synthesis

In heterogeneous architectures with parallel processors and application specific hardware components, communication synthesis is an important task. The semantics of SDL allows different communication models for the implementation of specifications: The processes can communicate via signals, which will be supported by the run-time environment. Another approach is to communicate via common memory, e.g. sharing variables. This model is called the data referencing model. For each SDL channel, the communication synthesis tool selects the cheapest and fastest implementation model. This is normally the shared variable model, because the sending of a signal via the run-time system needs many copy operations.

To select an implementation model for an SDL channel, the communication synthesis tool works on the internal control flow graph described in section 2.2. Based on the CFG, the communication synthesis tool selects all SDL channels, which can

be implemented by data referencing [26]. All other channels must be implemented using the signal mechanism of the run-time system. The results of the interface synthesis are used to define a starting solution for the optimizer.

3.3 Optimization

The optimizer searches for a system architecture with minimal cost that meets the given constraints. To perform this the optimizer calculates the mapping of nodes of the CFG to nodes of the architecture graph. Thus, the optimizer subsequently modifies the system architecture to improve its quality. The search for the optimal implementation of the system is based on a tabu search method [22]. If a solution does not meet the constraints, the optimizer allocates new hardware components of the prototyping environment and/or changes the mapping of the SDL processes. If a 'good' solution is found, the search algorithm exits and the system architecture provides the guidelines to generate the implementation of the system. To support a semi-automatic codesign approach, it is possible for the designer to specify a starting solution for the optimization using the mapping construct of SDL*.

3.4 Implementation Synthesis

To implement the optimized specification on the rapid prototyping board, the SDL* specification that contains all hardware and software mappings is translated to the implementation languages VHDL for the hardware components and C for the software modules.

We developed a tool to translate SDL* specifications to VHDL [21]. It should be easy to modify the back-end of the tool to generate the C description for the software modules, too. Because the tool generates a unique process identification for each process (hardware or software) used in the SDL specification, the synthesis of software-software, hardware-hardware, hardware-software and software-hardware communication is supported.

This tool generates a behavioral description of the SDL process and a structural description in VHDL based on the mapping information of SDL*. The structure description includes the communication links between the different modules of the resulting hardware. From the resulting VHDL code the architecture synthesis system CADDY [1] synthesizes a register-transfer description of the hardware modules. For each SDL process an application specific processor consisting of a datapath and a hardwired controller is generated. Using this description, RTL and logic synthesis tools generate the programming description for the FPGAs.

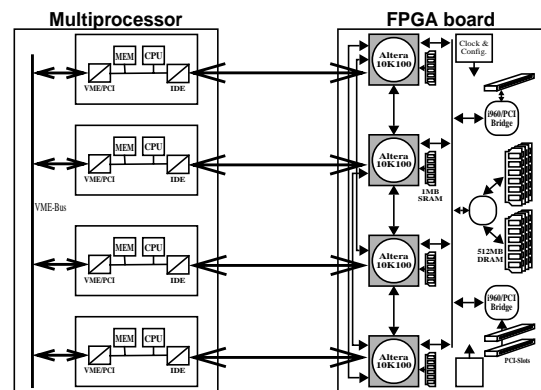


Figure 3: The prototyping platform

4 Rapid Prototyping Environment

4.1 FPGA Board and Multiprocessor System

The prototyping environment [5] consists of two parts: a VME based multiprocessor system and an FPGA board. Both parts are interconnected with several parallel connections.

The usage of a VME system [30] enables the use of standard processor boards which lowers the cost and also increases the flexibility by using several types of processors for specialized functions. Currently, we have installed two Pentium-133 boards with 32 MB of DRAM on each, an ethernet interface, a video adapter and several serial and parallel interfaces for debugging and monitoring purposes. They are compatible to the PC/AT standard, diskless and booted via network.

Each FPGA on the FPGA board is connected to an processor board through an own IDE interface. This simple harddisk interface offers up to 16 MB/s communication bandwidth and is flexible enough to implement different communication methods ranging from simple port accesses to high speed DMA transfers without much processor interaction. While the software side of the interface is integrated in the chipset of the VME processor boards the hardware side is implemented in the FPGAs by the use of matching library components.

The FPGA board shown in figure 3 contains 4 Altera Flex10k100 FPGAs [9]. They are interconnected with 70 signals between every two chips, which can be used to communicate between processes in different FPGAs but also to split processes among multiple FPGAs. Additionally, each FPGA is connected to a shared global bus. Via this bus, a dedicated memory controller [31] for 512 MB DRAM and two PCI controllers [32] can be reached. One of the PCI controllers connects the board to a host PC with a bus extender, the other one implements a local PCI bus with two slots for

network interface cards. We have chosen this configuration, because the implementation of the shared bus with its protocol takes less FPGA gates than the implementation of the memory controller and the PCI interfaces in the FPGAs.

For the targeted application area, we have integrated additional components to implement useful functionality: a dedicated SRAM module and a separate monitoring interface for each FPGA, a programmable clock generator for up to 6 different clocks with low-jitter buffers. The whole configuration of the FPGAs and the additional components is done by a microcontroller to support the remote control from an integrated development platform.

4.2 Runtime Support System

In addition to the prototyping hardware we have developed a set of library functions which together form the so-called runtime support system.

The runtime library of the parts of the SDL system implemented in software is based on the real-time operating system RTEMS [25]. All necessary functions for the implementation of code generated by the tools described in section 3.4 are resolved to the corresponding functions of the operating system. Also, the communication between hardware and software parts is handled by this runtime system. We have implemented a set of functions which offer different communication methods with different bandwidth and varying resource usage. The optimization phase selects the appropriate functions together with the corresponding hardware functions.

The same way, the software parts are linked with the library, the VHDL output of the tools is merged with library functions. This happens at different stages of the synthesis. The highest level contains functions which are integrated in the architecture generation of CADDY. The body of the function is directly inserted into the high level description. One level below are functions from which only the interface is considered by CADDY while generating the appropriate behavior. At the lowest level, between the high level synthesis of CADDY and the RT synthesis, complete library entities are connected to the entities generated by CADDY.

The functions used for hardware synthesis range from computational elements optimized for the target architecture over components needed to fulfil the semantic model of SDL like FIFO buffers to simple interfaces connecting the SDL processes to fixed components like the interface to the DRAM or a PCI based network card.

4.3 Measurement Tools

After code generation for the system, the application is executed in the real environment under aug-

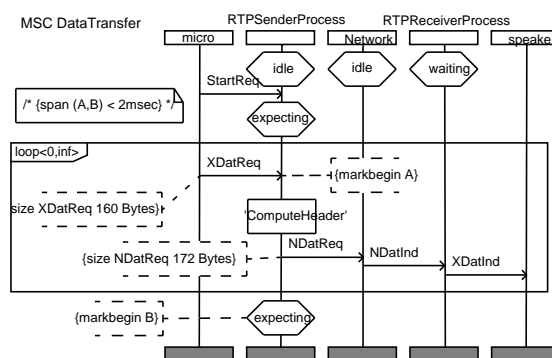


Figure 4: PMSC Specification of RTP Data Transfer

mentation of a performance monitor. For this, we use the ZM4/SIMPLE environment [12]. With the hardware monitor ZM4, a trace of timestamped events is recorded. Single events are for example the receipt of a signal at SDL level or the occurrence of an interrupt. Events in the runtime system itself like scheduling or buffer overflow can be recorded, too.

The interface between the system under test and the hardware monitor is implemented with functions resp. components from the runtime support system. After selecting interesting events in the high level specification or deducing them from nonfunctional constraints, the correct insertion points of these instrumentation commands have to be found in the functional specification. This is done automatically using a technique called specification driven instrumentation developed by our group [4]. The instrumentation is inserted in the SDL specification before compilation and resolved to functions in the runtime system.

The resulting event traces are analyzed with different methods implemented in our toolset SIMPLE. The necessary trace descriptions and query commands are also generated automatically from the specification. With the monitoring results, the system under test can be judged, and measured durations for each part of the system can be fed back in analysis tools like PEPP and the optimization step by using PMSC.

5 Case Study: Implementation of RTP

The methodology and tools mentioned above are applied to implement a multimedia communication system with stringent real time requirements also known as Quality of Service (QoS) requirements. One part of the system will be explained in the next section to illustrate how the tools are applied.

The example illustrated below concentrates on the transmission of voice data in the Internet using the protocol RTP (Real Time Transport Protocol) [24].

RTP is a light-weight application protocol which is used to transfer time-sensitive data. It does not guarantee the timely delivery of the data but provides mechanisms that allow the receiver to reconstruct the ordering, to synchronize different media streams and detect the loss of packets. Depending on the application, appropriate recovery mechanisms are executed. In addition to the data transfer protocol, RTP contains a control protocol (RTCP). Each participant of a session sends and receives reports in order to estimate the quality of the transmission and to perform adaption if the quality is too low. RTP performs on top of UDP [20] or ATM AAL5.

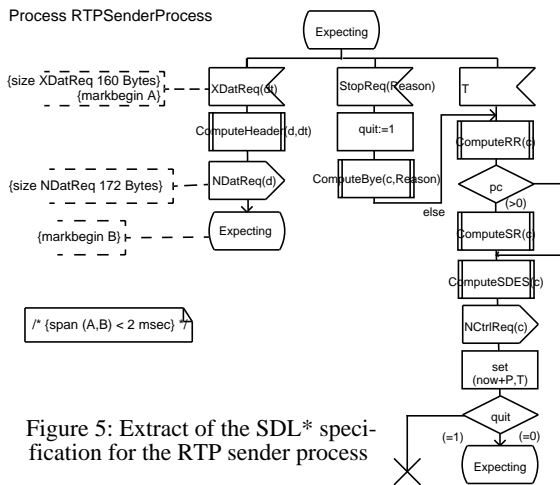


Figure 5: Extract of the SDL* specification for the RTP sender process

In the first phase of the software engineering process, the global behavior of the system is specified using PMSC in order to integrate performance aspects e.g. QoS requirements. In our example, we use RTP for transmitting voice data. Therefore, the PMSC specification in figure 4 consists of an instance RTPSenderProcess which describes the behavior of the sender. The sender receives periodically raw audio data from an audio process, which is not described here in detail. The audio data is packed into RTP packets and sent into the network using the signal NdatReq. For this case, we specify the following performance and QoS requirements. The application generates audio data with a rate of 64 kbps. The size of the payload is 160 bytes. The RTP protocol layer adds its own header to the payload, yielding to a total packet size of 172 bytes. The time requirement for processing the packets is 2 msec. (s. figure 4).

The PMSC specification describes the dynamic behavior of our system, e.g. the exchange of signals. It does not describe the functional behavior. For a complete functional specification we use SDL* (figure 5). Since the SDL* specification is the main input of the codesign tools all the dynamic requirements in the PMSC specification (figure 4) are transferred and matched.




Figure 6: Measurements of SW implementation of RTP sender process

In figure 6, some measurements for a software implementation are displayed, e.g. the processing time for the creation and sending of each RTP packet. The timestamps were taken when the audio data was sent to the SDL system and when the RTP packets were sent to the UDP sockets. The processing times of the RTP packets range from approximately 1.6 to 2.7 msec due to other concurrent processes operating in the system at the same time.

Using the implementation synthesis tool discussed in section 3.4 a behavioral VHDL description of the RTP specification is generated. With the architecture synthesis tool CADDY this description is transformed to a RT level implementation (figure 7). Performing the sender process of RTP in hardware the processing time to send a packet is smaller than 100 ns. The RTP specification contains one sender and one receiver process and CADDY generates a architecture with 7400 gates to implement one process.

The result of the experiment is that it is possible to implement a fast solution for the RTP protocol in hardware. A software solution is also possible if only the RTP processes are running on the hardware.




Figure 7: Synthesizing hardware from the RTP specification

6 Conclusion

In the paper we discussed the complete development cycle of a communication system. We have

extended MSC to include early performance analysis in the design process and SDL to support HW/SW codesign. The design phase of the development process is supported by several codesign tools and a rapid-prototyping environment. Presenting a short case study of the development of the RTP protocol we have shown how PMSC and SDL* are used to specify non functional requirements and how they are measured.

7 Acknowledgments

The authors thanks Lennard Lambert for the idea of figure 1 and a lot of discussions about the use of PMSC. We also thank Frank Lemmen for the discussion about the specification-driven monitoring approach.

8 References

- [1] R. Camposano, W. Rosenstiel. Synthesizing circuits from behavioral descriptions. IEEE Transactions Computer-Aided Design, Vol. 8, February 1989
- [2] J.M. Daveau, G. Marchioro, A.J. Jerraya. Hardware/Software Co-Design of an ATM Network Interface Card: a Case Study. 6th International Workshop on Hardware/Software Codesign, IEEE Computer Society Press, Seattle, 1998
- [3] P. Dauphin and A. Quick, "PEPP: Performance Evaluation of Parallel Programs," in 7. ITG/GI-Fachtagung Messung, Modellierung und Bewertung von Rechen- und Kommunikationssystemen, RWTH Aachen, Short Papers and Tool Presentation (B. Walke, ed.), Aachen, September 1993.
- [4] Dauphin, W. Dulz und F. Lemmen, "Specification-driven Performance Monitoring of SDL/MSC-specified Protocols", in Proc. of 8th Int. Workshop on Protocol Test Systems (A. Cavalli und S. Budkowski, eds.), September 1995.
- [5] M. Dörfel, R. Hofmann. A Prototyping System for High Performance Communication Systems. 9th IEEE International Workshop on Rapid System Prototyping. June 1998.
- [6] W. Dulz. A Framework for the Performance Evaluation of SDL/MSC-specified Systems. Proceedings ECM'96, 1996
- [7] N. Faltin, L. Lambert, A. Mitschele-Thiel, F. Slomka. An Annotational Extension of Message Sequence Charts to Support Performance Engineering. In SDL'97, Time for Testing, SDL, MSC and Trends, Proceedings of the eight SDL-Forum, Elsevier Science Publishers, 1997.
- [8] F. Fischer, A. Muth, G. Färber. Towards Interprocess Communication and Interface Synthesis for a Heterogeneous Real-Time Rapid Prototyping Environment. 6th International Workshop on Hardware/Software Codesign, IEEE Computer Society Press, Seattle, 1998
- [9] Flex10K Embedded Programmable Logic Family. Data Sheet. Altera, 1996
- [10] T. Hadlich, T. Szczepanski. The ODE-System - A SDL based Approach to Hardware-/Software-Codesign. Embedded Systems, OMI, 1995
- [11] U. Herzog. Performance Evaluation and Formal Description. Advanced Computer Technology, Proc. Reliable Systems and Applications, IEEE Computer Society Press, May 1991
- [12] R. Hofmann, R. Klar, B. Mohr, A. Quick, M. Siegle. Distributed Performance Monitoring: Methods, Tools, and Applications. IEEE Transactions on Parallel and Distributed Systems, Vol. 5, No. 6, June 1994
- [13] ITU-T. Z.100, Appendix I. ITU, SDL Methodology Guidelines. ITU, 1993
- [14] ITU-T. Z.120, Message Sequence Chart. ITU, 1996
- [15] T. Kolloch. SDL/MSC Design Methodology for Hard Real-Time Systems. Workshop on Performance and Time in SDL and MSC. Technical Report 1/98, IMMD VII, University of Erlangen-Nuremberg, Erlangen, February 1998.
- [16] T. Kolloch, G. Färber. Mapping an Embedded Hard Real-Time Systems SDL Specification to an Analyzable Task Network - A Case Study. Proceedings of the ACM SIGPLAN Workshop on Languages, Compilers and Tools for Embedded Systems (LCTES'98), Lecture Notes in Computer Science, Montreal, 1998.
- [17] L. Lambert. Bewertung von MSC-Spezifikationen mit Task-Graphen. Formale Beschreibungstechniken für verteilte Systeme, 7. GI/ITG-Fachgespräch, Berlin, Juni 1997.
- [18] G. De Micheli. Synthesis and Optimization of Digital Circuits. McGraw Hill, 1994
- [19] A. Mitschele-Thiel, F. Slomka. A Methodology for Hardware/Software Codesign of Real-Time Systems with SDL/MSC. CONSYSE '97, International Workshop on Conjoint Systems Engineering, K. Buchenrieder, A. Sedlmeier (editors), ITpress Verlag, Bruchsal / Chicago, to appear 1999
- [20] J. Postel. RFC 768: User Datagram Protocol, August 1980
- [21] D. Reichelt. Konzeption und Implementierung eines Werkzeugs zur automatischen Generierung einer VHDL Hardwarebeschreibung aus einer annotierten SDL-Systembeschreibung. Diplomarbeit am Lehrstuhl für Rechnerarchitektur und Verkehrstheorie, Universität Erlangen-Nürnberg, 1998
- [22] C.R. Reeves. Modern Heuristic Techniques for Combinatorial Problems. Blackwell Scientific Publications, Oxford 1993
- [23] J.Rozenblit, K. Buchenrieder. Codesign - Computer-Aided Software/Hardware Engineering. IEEE Press, 1995
- [24] H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson. RFC 1889: RTP: A Transport Protocol for Real-Time Applications, January 1996
- [25] RTEMS. The RTEMS Homepage. <http://www.oarcorp.com/rtems/rtems.html>
- [26] F. Slomka. Communication Synthesis for Parallel Implementations of SDL Specifications. submitted.
- [27] F. Slomka, J. Zant, L. Lambert. Schedulability Analysis of Heterogeneous Systems for Performance Message Sequence Chart. 6th International Workshop on Hardware/Software Codesign, IEEE Computer Society Press, Seattle, 1998
- [28] S. Spitz, F. Slomka, M. Dörfel. SDL* - An Annotated Specification Language for Engineering Multimedia Communication Systems. 6th Open Workshop On High Speed Networks, Institut für Nachrichtenvermittlung und Datenverarbeitung, Universität Stuttgart, 1997
- [29] F.Vahid, D.D. Gajski. Incremental Hardware Estimation during Hardware/Software Functional Partitioning. IEEE Transactions on VLSI-Systems, Vol. 3, No. 3, Sept. 1995
- [30] VMEbus International Trade Association, <http://www.vita.com>
- [31] V96BMC Rev. D. High Performance DRAM Controller. Data Sheet. V3 Semiconductor, 1996
- [32] V961PBC Rev. B1. Local Bus to PCI Bridge. Data Sheet. V3 Semiconductor, 1996