

RTOS-Centric Cosimulator for Embedded System Design

Shinya HONDA[†], Takayuki WAKABAYASHI[†], *Nonmembers*, Hiroyuki TOMIYAMA^{††a)},
and Hiroaki TAKADA^{††}, *Members*

SUMMARY With the growing design complexity of contemporary embedded systems, real-time operating systems (RTOSs) have become one of important components of such complex embedded systems. This paper presents an RTOS-centric hardware/software cosimulator which we have developed for embedded system design. One of the most remarkable features in our cosimulator is that it has a complete simulation model of an RTOS which is widely used in industry, so that application tasks including RTOS service calls are natively executed on a host computer. Our cosimulator also features cosimulation with functional simulation models of hardware written in C/C++ and cosimulation with HDL simulators. A case study with a JPEG decoder application demonstrates the effectiveness of our cosimulator.

key words: RTOS, cosimulation, embedded systems

1. Introduction

Nowadays, real-time operating systems (RTOSs) have become one of the most important components of embedded systems due to the growing complexity of the system functionality as well as the increasing time-to-market pressures. With this trend, a demand for fast cosimulation of hardware and embedded software including an RTOS is also becoming stronger in order to validate the functionality of the overall systems. For this purpose, we have developed a fast hardware/software cosimulator which fully supports services of an RTOS. Since the RTOS model plays a central role in our cosimulator, we call our cosimulator an *RTOS-centric cosimulator*.

The RTOS model in our cosimulator is provided to a system designer in the form of C/C++ source code. Application software tasks written by the designer are compiled and linked with the RTOS model to generate an object code which is directly executable on the host computer. Due to the native execution, the speed of our cosimulator is much higher than that of traditional cosimulation which uses an instruction-set simulator (ISS) of the target processor for software execution. It should be noted that simulation using our cosimulator is basically untimed, but timed simulation is also possible by inserting a delay function into the software code.

Our cosimulator supports two types of hardware simulation. One is cosimulation with functional hardware models written in C or C++. This enables fast functional validation of the overall systems. The other is cosimulation with HDL simulators, e.g., *ModelSim* [1]. Using this ability, register-transfer level (RTL) design implementation of the hardware can be simulated together with embedded software (i.e., RTOS and application tasks). In this case, our cosimulator works as a testbench which generates input data to and receives output data from the hardware in an interactive manner. Thus, the hardware design can be validated efficiently.

Our cosimulator supports all the services of the *μITRON 4.0* standard [2], [3]. *μITRON* is one of the most popular RTOSs in Japan for small- to middle-scale embedded systems. *μITRON* is not a specific RTOS product, but is an application programming interface (API) standard. It only defines a set of API functions, and implementation of the function bodies may differ among *μITRON*-based RTOSs. It is reported in [3] that over 40% of RTOSs used in Japan are based on the *μITRON* standard. Our simulator implements all of 91 API functions which are defined by *μITRON 4.0 Standard Profile*. Therefore, one can use the same application code for both simulation and final implementation if a *μITRON*-based RTOS used in the implementation. Due to this, the design productivity can be significantly improved.

In summary, our RTOS-centric cosimulator features

- native (hence, fast) software execution on a host computer,
- cosimulation with functional hardware models in C/C++,
- cosimulation with HDL simulators, and
- complete support of RTOS services.

To our knowledge, no other cosimulator satisfies all of the above features.

This paper is organized as follows. Section 2 surveys related work on hardware/software cosimulation with RTOS supports. Section 3 describes principles and implementation of the RTOS-centric cosimulator which we have developed. A case study using a JPEG decoder application is presented in Sect. 4. Section 5 concludes this paper with a summary.

Manuscript received March 19, 2004.

Final manuscript received August 5, 2004.

[†]The authors are with the Department of Information and Computer Science, Toyohashi University of Technology, Toyohashi-shi, 441-8580 Japan.

^{††}The authors are with the Department of Information Engineering, the Graduate School of Information Science, Nagoya University, Nagoya-shi, 464-8603 Japan.

a) E-mail: hiroyuki@acm.org

2. Related Work

Hardware/software cosimulation has been studied around the world for more than a decade, and a number of commercial cosimulators have come onto the market. The cosimulators are currently one of indispensable CAD tools in the design of embedded systems and systems-on-chip. However, since the traditional cosimulators do not feature explicit supports for RTOSs, a designer needs to use an ISS in many cases to run embedded software including an RTOS. Due to the slow execution speed, extensive simulation of large software is impossible.

In the recent years, several research efforts have been made to model RTOSs for system-level design and cosimulation. *SoCOS* presented in [4] is a system-level design environment where the *OSAPI* library provides generic RTOS system calls to application software. *OSAPI* is a virtual RTOS to enable native execution of embedded software. After simulation-based validation, the *OSAPI* calls are replaced with the system calls of the actual RTOS used in the final implementation to obtain the final software code. In [5], a similar approach is presented. A main difference is that the RTOS model in [5] is build upon an existing system-level design language, i.e., SpecC [6], so that existing CAD tools such as simulators can be used. Techniques presented in [7] also use the SpecC language to model the preemptive behavior. In [8], an OS model is proposed for fast and time-accurate cosimulation. The model focuses on accurately modeling the RTOS overhead during task execution as well as the preemptive behavior. In [9], a method which automatically generates RTOS-dependent software from SystemC description is presented. The method replaces SystemC's constructs for concurrency and communication with corresponding RTOS service calls. In this sense, it can be considered that SystemC involves a simple RTOS model in itself.

A common weakness of these OS models is that they support only a limited set of RTOS services in order to make the models generic and independent of specific RTOSs. μ ITRON-based RTOSs, for example, have more than 90 service calls. These services may need to be fully utilized in order to write high-quality software. However, the RTOS model in [5], for example, supports only 16 service calls. Therefore, it is easily imagined that the quality of software automatically generated by these previous methods is lower than that of hand-crafted RTOS-dependent software. Since the RTOS-centric cosimulator which we have developed fully supports RTOS services, such high-quality software can be designed and simulated. Instead, our simulator has a different weakness that it can be used only for μ ITRON-based platforms.

In fact, our RTOS-centric cosimulator is not a replacement of the cosimulators with generic RTOS modeling, but is complementary to them. The cosimulators with generic RTOS modeling can be used in the high-level design phases where the RTOS to be used in the implementation is not

determined yet. After RTOS selection, the software is refined into RTOS-dependent code. Of course, the RTOS-dependent software needs to be cosimulated to verify the correctness of its functionality. So far, the RTOS-dependent software has been executed using an ISS, so the cosimulation speed is slow. With our RTOS-centric cosimulator, however, the RTOS-dependent software can be cosimulated at a very high speed. After the cosimulation-based functional validation, cycle-accurate cosimulation will be done using an ISS. Thus, our cosimulator fills the gap of abstraction between system-level RTOS-independent cosimulation and cycle-accurate cosimulation.

Some RTOSs have their simulation models for native execution on a host computer. For example, *VxWorks* from WindRiver Systems [10] has its simulation model named *VxSim*. Our cosimulator is similar to the RTOS simulation models. To our knowledge, however, the RTOS simulation models do not explicitly support cosimulation with hardware. On the other hand, our cosimulator supports cosimulation with standard HDL simulators as well as untimed hardware models in C or C++.

A different approach to embedded software design is presented in [11], where an application-specific RTOS and its simulation model is automatically generated. In their approach, application software is analyzed first, and then only the RTOS services used in the application are included in the final RTOS. The work is similar to ours in the sense that a set of services which can be used in application software is predefined. The major difference is that the work in [11] puts a special focus on customizing an RTOS while our cosimulator is based on a standard RTOS.

3. RTOS-Centric Cosimulator

The RTOS-centric cosimulator which we have developed implements all the services defined in *μ ITRON 4.0 Standard Profile*. In this section, an overview of *μ ITRON* is presented, and then our cosimulator is described in detail.

3.1 μ ITRON

μ ITRON is a set of APIs which has been standardized mainly in Japan. *μ ITRON* is designed for small- to middle-scale embedded systems such as consumer electronic products, mobile devices, automotive electronic control systems, and so on. *μ ITRON*-based RTOSs have been widely used in industry, especially in Asian countries. It is reported in [3] that over 40% of RTOSs used in Japan are based on the *μ ITRON* standard. The latest formal release at present is *μ ITRON 4.0*. The minimal set of *μ ITRON* APIs is called *Standard Profile*, and *μ ITRON 4.0 Standard Profile* includes 91 API functions in total[†].

μ ITRON employs a priority-based preemptive scheduling policy, and the priority of tasks can be changed at runtime. *μ ITRON 4.0 Standard Profile* states that the number

[†]In this paper, the term *μ ITRON* often refers to *μ ITRON 4.0 Standard Profile* depending on the context.

of priority levels must be at least 16. μ ITRON supports several basic synchronization and communication mechanisms, such as semaphores, events, data queues, and mailboxes. Dynamic memory allocation using so-called memory pools is supported in μ ITRON. Since μ ITRON is designed for small embedded systems, virtual memory or dynamic module loading is not supported. Interrupt handling is one of important mechanisms in RTOSs. μ ITRON provides a number of API functions for defining interrupt handlers, allowing and prohibiting interrupts, changing interrupt masks, and so on. The number of interrupt levels is not determined by μ ITRON but implementation-dependent. μ ITRON has three types of time event handlers, i.e., cyclic handlers, alarm handlers, and overrun handlers. Cyclic handlers are invoked periodically, alarm handlers are invoked at a specified time, and overrun handlers are invoked when the execution time of a task exceeds a specified time.

In addition, μ ITRON provides a number of services which are necessary for industrial use. More detailed documents on μ ITRON are found in [3].

3.2 Cosimulator Overview

In the rest of this section, we describe our RTOS-centric cosimulator which is based on the μ ITRON standard. The cosimulator implements all the services defined in μ ITRON 4.0 Standard Profile. Figure 1 shows the overall organization of our RTOS-centric cosimulator. The shaded parts in the figure are included in the cosimulator.

The cosimulator is developed in the C and C++ languages and runs on Microsoft Windows-based computers. The cosimulator is freely available as a part of the *TOPPERS/JSP Kernel* package [12][†]. The TOPPERS/JSP Kernel is a μ ITRON-based RTOS kernel which we have developed earlier, and it is freely available for both research and commercial purposes.

The RTOS model, consisting of the RTOS kernel and *HAL* (*hardware abstraction layer*) in Fig. 1, is provided to a system designer in the form of C/C++ source code. Ap-

plication tasks which are written in the C language by the designer are compiled and linked with the RTOS model, in order to generate an object code which is directly executable on the Windows-based host computer. Our RTOS-centric cosimulator supports cosimulation with functional hardware models written in C or C++, which enables fast validation of the overall system functionality. Cosimulation with standard HDL simulators is also supported so that RTL hardware designs can be validated efficiently.

3.3 Task Management

Application tasks and the RTOS model are compiled together using a native C/C++ compiler on a host computer, and an object file is generated. The object file is executable directly on the host computer. Thus, from a viewpoint of the host computer, the object code (i.e., application tasks and the cosimulator) is no more than an application process, and is scheduled in a time-sharing manner with the other Windows applications.

Within the cosimulator, each of the application tasks is implemented as a thread on the host computer, and the RTOS kernel schedules the threads in a priority-based preemptive manner. This thread-based implementation facilitates debugging of the application tasks since the context of each task, such as mode (running, ready, waiting, suspended, etc.), program counter, contents of stack memory, and so on, can be easily monitored with normal C/C++ debuggers. If all the tasks are linked into a thread, it is not easy to retrieve and monitor the contexts of individual tasks without RTOS-specific support tools.

3.4 Timer

Unlike the reference simulator of SystemC [13] or SpecC [6], our RTOS-centric cosimulator does not have a global timer for managing simulation cycles. Therefore, our cosimulator does not support cycle-accurate execution of software. Instead of the global clock for simulation, our cosimulator uses the timer of the host computer running on MS-Windows. In Fig. 1, the arrow labeled (a) denotes the system call to the host computer for getting the time.

In our cosimulator, the period of timer ticks is customizable, but it must be equal to or longer than 1 millisecond. If an application task executes a system call, *dly_task(100)*, the RTOS kernel suspends the execution of the task for 100 milliseconds in the real time on the host computer. It should be noted that the system call does not mean that the task waits for 100 simulation cycles.

3.5 Cosimulation Mechanism

As mentioned earlier, our cosimulator supports two types of hardware models, i.e., functional simulation models in

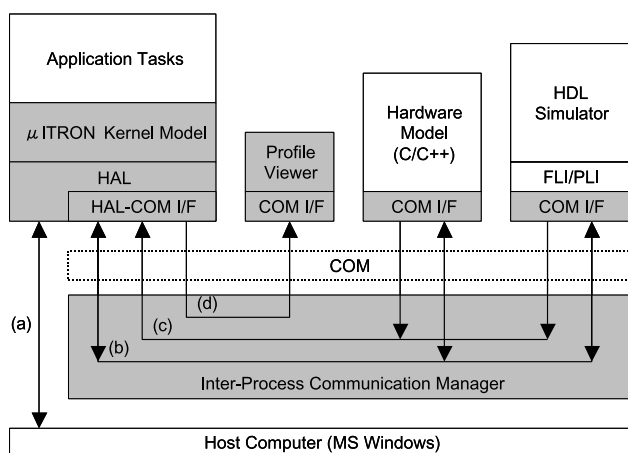


Fig. 1 Organization of the RTOS-centric cosimulator.

[†]The ability to communicate with HDL simulators is not included in the current release.

C/C++ and hardware designs in HDL. The functional simulation models are compiled with a native C/C++ compiler to generate object code which is executable on the host computer. The HDL designs are simulated on HDL simulators which support PLI (Programming Language Interface) or FLI (Foreign Language Interface)[†]. It should be noted that both functional simulation models and HDL simulators can be executed simultaneously, and the number of the functional models and that of HDL simulators can be more than one. For example, it is possible to cosimulate a software model (consisting of the RTOS model and application tasks), two functional hardware models, and three HDL simulators simultaneously.

In our cosimulator, as shown in Fig. 1, *Inter-Process Communication Manager (IPCM)* manages communication between the software model and the hardware models. IPCM is based on the *Component Object Model (COM)* technology developed by Microsoft Corporation [14]. COM is an object-oriented interface technology which enables binary software components to interact with other software components in order to realize higher-level services. In our cosimulator, the software model, functional hardware models, and HDL simulators run as different processes on the host computer. Each of the software model, the functional hardware models and the HDL simulators interacts with IPCM through COM. Then, communication between the software and the hardware is performed by way of IPCM.

μ ITRON defines a set of guidelines for designing hardware devices and their drivers. In the guidelines, it is recommended that all devices (except a timer and a primary serial I/O port) be accessed through memory-mapped I/O. The guidelines also define a hardware abstraction layer (HAL) which includes a set of functions for memory-mapped I/O accesses. For example, a function, *sil_reb_mem(mem)*, reads 1-byte data from the device which is mapped to the address pointed by *mem*. Our RTOS-centric cosimulator assumes that hardware devices and the drivers are designed based on the guidelines.

Let us assume that *sil_reb_mem(mem)* is executed in the software. First, the RTOS model sends a request message to IPCM through COM. The message includes the access type (read or write), the size of required data, and the address. IPCM has a memory-map table, and based on the table and the received address, IPCM sends the request message to the corresponding functional hardware model or HDL simulator. The requested data is then sent from the hardware to the software by way of IPCM. Thus, IPCM serves as a router of the messages and data. In Fig. 1, communication based on memory-mapped I/O is depicted by arrow (b). Interrupts from the hardware to the software are also handled by IPCM, and are depicted by arrow (c) in Fig. 1.

The cosimulation mechanism with IPCM is flexible in several points. First, our cosimulator can simulate any number of hardware models simultaneously as long as there is no conflict in the address map. Also, it is possible to simulate both C/C++ models and HDL models simultaneously. Second, a C/C++ model can be easily replaced with

an HDL model or vice versa without modifying application tasks, provided the address map of the hardware remains unchanged. This ability facilitates validation of hardware designs. Finally, the RTOS model can be easily replaced with an ISS running on a μ ITRON-based RTOS by adding the HAL-COM interface to the ISS, without modifying application tasks or hardware models. Thus, our cosimulator enables plug-and-play of software models as well as hardware models. This is due to the flexible communication mechanism of IPCM as well as the complete supports of the μ ITRON standard.

3.6 Profiler

Another remarkable feature in our cosimulator is the profiler to collect various data during execution. The data includes start/finish times and return code (if any) of application tasks, interrupt handlers, time event handlers, service calls, and so on. Such data is useful to debug or optimize the designs.

The obtained data can be displayed at runtime on a window (i.e., named *Profile Viewer* in Fig. 1). The profiler is implemented within HAL, and the profiled data is sent to the viewer through IPCM. Of course, profiling requires some amount of overhead on simulation speed and memory requirement. In order to reduce the overhead, the profiler can be enabled and disabled during simulation. Therefore, only the necessary data can be obtained with the minimal overhead.

It should be noted that much more information can be obtained by using normal C/C++ debuggers if the cosimulator is compiled with the debugging option enabled, i.e., the “-g” option for many C compilers. It is also possible to embed any C code for profiling into the cosimulator, since the source code of the cosimulator is available.

4. Design Example

In order to evaluate the effectiveness of our RTOS-centric cosimulator, we designed a JPEG decoder system and simulated it with our cosimulator. The simulation was done on dual Xeon 2.4 GHz processors with hyper-threading technology, running on MS-Windows XP. Figure 2 shows the flow of the JPEG decoder where there are mainly four functions: VLD, Dequantization, IDCT, and YUV2RGB. In addition, there is a function, Display, to display the decoded image on a window of the host computer.




Fig. 2 A JPEG decoder example.

[†]At present, only the ModelSim HDL simulator [1] is supported in our cosimulator.

Table 1 Cosimulation time of the JPEG decoder. All of the functions are implemented in software.

CPU Utilization for the JPEG Task	100%	50%	25%
Our Cosimulator	0.009 sec	0.008 sec	0.008 sec
ARMulator	23.8 sec	59.3 sec	125 sec

Table 2 Cosimulation time of the JPEG decoder. IDCT is implemented in hardware, and is described as a functional simulation model in C.

CPU Utilization for the JPEG Task	100%	50%	25%
Our Cosimulator	46 sec	49 sec	97 sec
ARMulator	71 sec	94 sec	141 sec

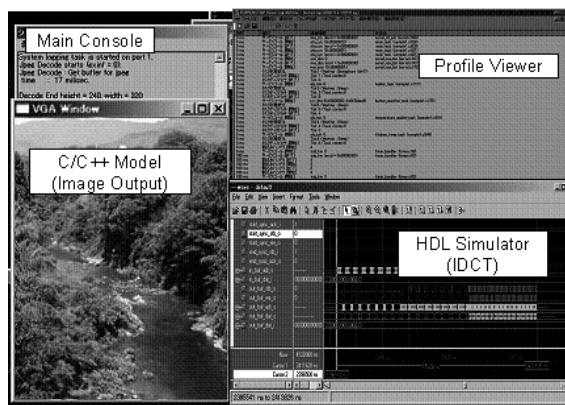
First, we implemented all of the functions in software except the display function. The display function was designed only for the simulation purpose, so it was written as a functional simulation model in C/C++. A JPEG image is decoded by software, and then the decoded image is sent from the software to the display window through IPCM. The overall JPEG decoder was implemented as an application task. We compared cosimulation speed of our cosimulator with that of an instruction-set simulator, ARMulator [15]. On ARMulator, the JPEG task was executed with the TOPPERS/JSP Kernel [12]. Note that we used the same JPEG code for both our cosimulator and ARMulator without modification. Table 1 summarizes the results on the simulation time. In this table, the first row labeled “*CPU Utilization for the JPEG Task*” indicates the percentage of CPU time allocated to the JPEG task on the target system. 100% utilization means that there is no other application task. Even in this case, there exist a few system tasks such as the timer and the profiler. 50% means that another task exists while 25% means that three tasks exist in the target system, in addition to the JPEG task. The application tasks running concurrently with the JPEG application include an infinite loop whose body is empty. All of these application tasks had the same priority level as the JPEG application on the RTOS. Due to this, all of the application tasks (including the JPEG task) were scheduled in a time-sharing manner. The timer tick was set to 10 milliseconds, so that the application tasks switch every 10 milliseconds. Table 1 demonstrates that our cosimulator is three or four orders of magnitude faster than the ISS.

Next, we changed hardware/software partitioning. IDCT was moved from software to hardware, and we described a functional simulation model of the IDCT function in C. The cosimulation time is shown in Table 2. With our cosimulator, most of the cosimulation time was spent for communication between the software and the hardware when the CPU utilization was 100%. When the CPU utilization was 50% and 25%, the simulation time presented in Table 2 includes the time for simulating the other tasks.

Finally, we designed an IDCT circuit at the register-transfer level in VHDL. The IDCT simulation model was replaced with an HDL simulator, ModelSim [1], to simulate the RT-level IDCT design. Note that the JPEG appli-

Table 3 Cosimulation time of the JPEG decoder. IDCT is implemented in VHDL, and is executed with an HDL simulator.

CPU Utilization for the JPEG Task	100%	50%	25%
Our Cosimulator	182 sec	327 sec	333 sec
ARMulator	248 sec	961 sec	2,802 sec

**Fig. 3** A snapshot of JPEG simulation with the RTOS-centric cosimulator.

cation code did not require any modification from the one cosimulated with the C-level hardware model. The simulation speed is shown in Table 3.

A snapshot of this cosimulation is presented in Fig. 3, where the software model (i.e., the RTOS model and the application tasks), the HDL simulator, the C/C++ simulation model of the display function, and the profiler were running simultaneously, with communicating with each other through IPCM.

5. Conclusions

RTOSs have become one of necessities in the design of complex embedded systems. In order to validate the overall functionality of such embedded systems, RTOSs need to be simulated together with application software and hardware. We have developed a fast, flexible RTOS-centric cosimulator for embedded system design. Our cosimulator features complete supports of RTOS services, native execution of software, cosimulation with functional hardware models in C or C++, and cosimulation with HDL simulators. In this paper we have described principles and implementation of our cosimulator. We have also shown a case study using a JPEG decoder example to demonstrate the effectiveness of the cosimulator.

At present, the timing of communication and synchronization is inaccurate in our cosimulator. Improvement of timing accuracy remains as one of our future work.

Acknowledgment

This work was partially supported by JSPS Grant-in-Aid for Young Scientists (B) #16700058.

References

- [1] Mentor Graphics Corporation, <http://www.mentor.com/>
- [2] H. Takada and K. Sakamura, "μITRON for small-scale embedded systems," *IEEE Micro*, vol.15, no.6, pp.46–54, 1995.
- [3] ITRON, <http://www.assoc.tron.org/itron/>
- [4] D. Desmet, D. Verkest, and H. De Man, "Operating system based software generation for systems-on-chip," *Proc. Design Automation Conference (DAC)*, pp.396–401, 2000.
- [5] A. Gerstlauer, H. Yu, and D. Gajski, "RTOS modeling for system level design," *Proc. Design Automation and Test in Europe (DATE), Embedded Software Forum*, pp.130–135, 2003.
- [6] SpecC Technology Open Consortium, <http://www.specc.org/>
- [7] H. Tomiyama, Y. Cao, and K. Murakami, "Modeling fixed-priority preemptive multi-task systems in SpecC," *Proc. Workshop on Synthesis and System Integration of Mixed Information Technologies (SASIMI)*, pp.93–100, 2001.
- [8] Y. Yi, D. Kim, and S. Ha, "Virtual synchronization technique with OS modeling for fast and time-accurate cosimulation," *Proc. Int'l Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS)*, pp.1–6, 2003.
- [9] F. Herrera, H. Posadas, P. Sanchez, and E. Villar, "Systematic embedded software generation from SystemC," *Proc. Design Automation and Test in Europe (DATE), Embedded Software Forum*, pp.142–147, 2003.
- [10] WindRiver Systems Inc., <http://www.wrs.com/>
- [11] S. Yoo, G. Nicolescu, L. Gauthier, and A.A. Jerraya, "Automatic generation of fast timed simulation models for operating systems in SoC design," *Proc. Design Automation and Test in Europe (DATE)*, pp.620–627, 2002.
- [12] TOPPERS Project, <http://www.toppers.jp/>
- [13] SystemC Open Initiative, <http://www.systemc.org/>
- [14] Microsoft Corporation, <http://www.microsoft.com/>
- [15] ARM Corporation, <http://www.arm.com/>



Hiroyuki Tomiyama received his Ph.D. degree in computer science from Kyushu University in 1999. From 1999 to 2001, he was a visiting postdoctoral researcher with the Center of Embedded Computer Systems, University of California, Irvine. From 2001 to 2003, he was a researcher at the Institute of Systems & Information Technologies/KYUSHU. In 2003, he joined the Graduate School of Information Science, Nagoya University, as an assistant professor, where he is now an associate professor. His research interests include system-level design automation, architectures and compilers for embedded systems and systems-on-chip. He is currently serves as an editorial board member of *International Journal on Embedded Systems*. He has also served on the organizing and program committees of several premier conferences including ICCAD, ASP-DAC, CODES+ISSS, HLDVT, SCOPEs and so on. He is a member of ACM, IEEE, and IPSJ.



Hiroaki Takada is a Professor at the Department of Information Engineering, the Graduate School of Information Science, Nagoya University. He received his Ph.D. degree in Information Science from University of Tokyo in 1996. He was a Research Associate at University of Tokyo from 1989 to 1997, and was an Assistant Professor and then an Associate Professor at Toyohashi University of Technology from 1997 to 2003. His research interests include real-time operating systems, real-time scheduling theory, and embedded system design. He is a member of ACM, IEEE, IPSJ, and JSSST.



Shinya Honda received the B.E. and M.E. degrees in information and computer sciences from Toyohashi University of Technology, Aichi, Japan, in 2000 and 2002, respectively. Currently he is a Ph.D. candidate at the Department of Electronic and Information Engineering, Toyohashi University of Technology. His research interests include system-level design automation and real-time operating systems. He received the best paper award from IPSJ in 2003. He is a member of IPSJ.



Takayuki Wakabayashi received the B.E., M.E. and Ph.D. degrees through his research life in information and computer sciences from Toyohashi University of Technology. His research interests include development environments for embedded systems, codesign, cosimulation, PC-simulation and rapid prototyping.