

VLSI based memory peripheral for RTOS optimisation

S.Sheeba Rani Gnanamalar M.E., MBA.,(Ph.D)¹, A.Anie Selva Jothi M.E²

Assistant Professor, Department of Electrical and Electronic Engineering, Sri Krishna College of Engineering and Technology, Coimbatore, India¹

Assistant Professor, Department of Electronics and Communication Engineering, Sri Krishna College of Engineering and Technology, Coimbatore, India²

ABSTRACT: The growing technology in embedded application has resulted in use of increase in RTOS. Unfortunately RTOS can introduce significant performance degradation due to the following reasons: 1) when we implement real time task in software the processor utilization is more in order to schedule and manage the tasks. Hence lot of processor utilisation is wasted.2) The RTOS occupies considerable memory space to complete the task scheduling.3) Power consumption is more due to read/write operations of the memory and switching activities of the memory cycle.4) Speed of execution is dependent on the processor clock.This paper presents the Real-Time Task Manager (RTM)—a processor extension that minimizes the above performance drawbacks associated with RTOSes. The RTM accomplishes this by supporting, in hardware, a few of the common RTOS operations that are performance bottlenecks: task scheduling, time management, and event management. By exploiting the inherent parallelism of these operations, the RTM completes them in constant time, thereby significantly reducing RTOS overhead. It decreases both the processor time used by the RTOS and the maximum response time by an order of magnitude. This hardware support for RTOS is implemented in VHDL and the simulation and synthesis report are generated.

KEYWORDS: RTOS, Scheduling, Hardware-Software Codesign, Processor Utilisation

I. INTRODUCTION

RTOSes have become extremely important to the development of real-time systems as reflected by the growing market for RTOSes; half a billion dollars in shipments of RTOSes will be sold in 2002 [5]. RTOSes provide services for better hardware abstraction multitasking, task synchronization etc. However, the benefits that RTOSes provide do not come for free. The use of RTOSes has several important effects on various performance parameters of real-time systems.

Processor Utilization: The fraction of processing power that an application is able to use. In order to provide services like multitasking, pre-emption, and numerous others, the RTOS introduces an overhead which lowers the processing power available to the application

Response Time: The time it takes for a real-time system to respond to external stimuli. This delay is highly dependent on several factors, including whether or not the RTOS is pre-emptive and whether polling or interrupts are used to sense the stimulus. The effects of RTOSes on response time vary widely, but, in general, RTOSes increase response time to varying degrees.

Much of the performance degradation caused by RTOSes can be traced to their core operations, namely task scheduling, time management, and event management. Analysis of these functions reveals that they are executed frequently, that they perform inter-related actions, and that these actions exhibit parallelism. That parallelism cannot be exploited by software-based implementations; however, hardware-based solutions can exploit this parallelism and lower the performance degradation. Also, the rapidly dropping cost of logic [7] has made it possible to put custom hardware for enhancing frequently executing operations on embedded processors, without increasing costs significantly.

These factors suggest that RTOSes would greatly benefit from a custom hardware solution. This is the motivation for the Real-Time Task Manager (RTM), a memory-mapped on-chip peripheral designed to optimize task scheduling, time management, and event management that is compatible with a wide range of RTOSes. Measurements have been taken of the performance impact of the RTM on models of realistic real time systems. These measurements show that processor utilization and maximum response time are both reduced by an order of magnitude.

II. RTOS FUNCTIONS

To better understand that the commonly seen bottlenecks in RTOSes, we present some of the details regarding their nature— what they are, how they are implemented and how often they are invoked.

Task Scheduling: Process of determining which task should be running at any given time. The most commonly implemented approach to scheduling in commercial RTOSes is based on task priorities [8]. Priority-driven schedulers assign each task a priority and execute the task with the highest priority that is ready.

Time Management: The RTOS gets its sense of time by a periodic interrupt generated every clock tick by a hardware timer. The rate of the generation of this clock tick is the system clock frequency which is not to be confused with the CPU clock frequency. Typically the system clock frequency is on the order of kHz while the CPU clock is on the order of MHz. *Time management* refers to the RTOS's ability to allow tasks to use this mechanism to be scheduled at specific times i.e. have the task block for a specific time before becoming ready.

Event Management: Most real-time operating systems provide services for communication and synchronization between tasks—known as *Interprocess communication (IPC)*. Examples are semaphores, message queues etc. Event management involves keeping track of which tasks are blocked i.e. waiting for IPC and which tasks should be released i.e. have to receive IPC on resource availability.

Note that as shown in Table 1, the overhead of the RTOS due to these operations is proportional to the product of frequency and complexity. Since both of these components may increase linearly with the number of tasks, there may be a quadratic relationship between the number of tasks in the system and the RTOS overhead due to them. In addition the overhead of task scheduling and time management increases linearly with the system clock frequency.

The longest time that it takes to perform any critical section of code (including task scheduling, event or time management) adds to the maximum response time in pre-emptive systems. However, in nonpre-emptive systems where interrupts are polled, a response task is delayed more by executing workloads than by RTOS operations. This is because workloads execute for longer periods than RTOS operations and thus disable interrupt response for longer periods.

III. REAL-TIME TASK MANAGER

The RTM is a hardware module that implements a task database and thus supports the key RTOS operations: task scheduling, time management, and event management. It is an on-chip peripheral that communicates with the core via a memory-mapped interface. The RTM is pictured in Figure 1; where each record contains information about a single task. The RTM is accessed through the global address space and each record can be read from or written to as if it were another array of structures.

A record is composed of four individually addressable fields. Each record contains a *status field*, containing several bits that describe the status of the corresponding record. The *valid bit* is necessary to indicate if that the record is used by a task. The *delayed bit* indicates that the task is waiting for the amount of clock ticks specified by the *delay field* before being ready-to-run. The *event bit* indicates that the task is pending on the event with the identifier specified by the *event ID field*. The *suspended bit* indicates that the task has been suspended. If the valid bit is set, but the delayed, event, and suspended bits are clear, then the task is ready-to-run. Finally, the *priority field* indicates the task's priority. Also, the maximum number of records is some fixed constant, such as 64 or 256. The RTOS can issue instructions to retrieve certain information from the RTM via a set of memory-mapped registers. The RTM can be queried for the highest priority ready task and thus supports static-priority scheduling. The RTM does time management by decrementing the delay fields of all records by the number of clock ticks specified by the RTOS and updating the delay bit if this value reaches zero. Event

management is very similar to static-priority scheduling. The RTM queries for the highest priority task that is pending on a given event identifier, instead of querying its data structure for the highest priority ready task.

IV. ARCHITECTURE

A reference architecture is presented Above to illustrate the complexity and scalability of implementing the hardware architecture of the RTM. For a 64 record RTM with entries as shown in the figure the total number of flip-flops required is 2304, which is small enough to easily be implemented. These parameters are based on common limits and are sufficient for the majority of real-time applications.

Task scheduling is performed using 64 Ready Test cells and a binary tree of 63 Priority Test cells. Owing to space considerations we show this for a 32 entry case in Figure 2. The Ready Test cells simply determine which tasks are ready-to-run. Each Priority Test cell takes in as input two task priorities and ready bits. It then outputs which of these the highest priority ready task is. These priority values propagate down the comparator tree until the single sixth order cell where the index of the overall highest priority ready task is output. Each N^{th} order Priority Test cell contains an 8-bit comparator, an 8-bit 2:1 MUX, an $(N-1)$ -bit 2:1 MUX, an OR gate, an AND gate, and a NAND gate. The logic required task scheduling scales linearly with the number of records in the RTM (n RTM records requires n RT cells plus $n-1$ PT cells). The computational delay is proportional to the logarithm of the number of records. For a reasonable number of records i.e. the numbers typically used in production RTOSes, this implementation of task scheduling will be sufficiently fast and small.

In order to implement time management, 64 delay decrement cells are all that is required, as shown in Figure 3. Delay decrement cells consist of a simplified 16-bit adder and an AND gate. Delay cells decrement the delay field and perform a comparison to clear the delay bit when necessary. In this implementation, the logic scales linearly with the number of records; however, the computational delay is a constant. In designs where the CPU clock frequency is lower than the system clock frequency, the delay decrement cells can be reused and thus lower the number of delay cells required.

The event management is almost exactly the same as task scheduling. The only difference is that instead of Ready Test cells, it has Event Test cells, as seen in Figure 2. This allows the binary tree of Priority Test cells to be used for both scheduling and event management. The Event Test cells just check if the event bit is set and if the event ID of the resource being released matches the event ID in a specific record. They each use an 8-bit comparator and an AND gate. The combinational logic required for event management scales linearly, except that the majority can be shared with the task scheduler. Also, the computational delay is $O(\log(n))$, as it is with scheduling. The die area that the RTM needs is important in determining its feasibility. Based on existing area models for register files and caches [9], the RTM reference architecture requires approximately 2600 register-bit equivalents (RBEs). This is roughly equivalent to the die area used by a 32-bit by 64-word register file making it feasible to implement the RTM in hardware. This number is large compared to a simple core but is small compared to the whole chip. This is because a typical embedded processor consists of the core and other on-chip devices like memory and various I/O devices.




Fig. 1 Reference RTM Data Structure

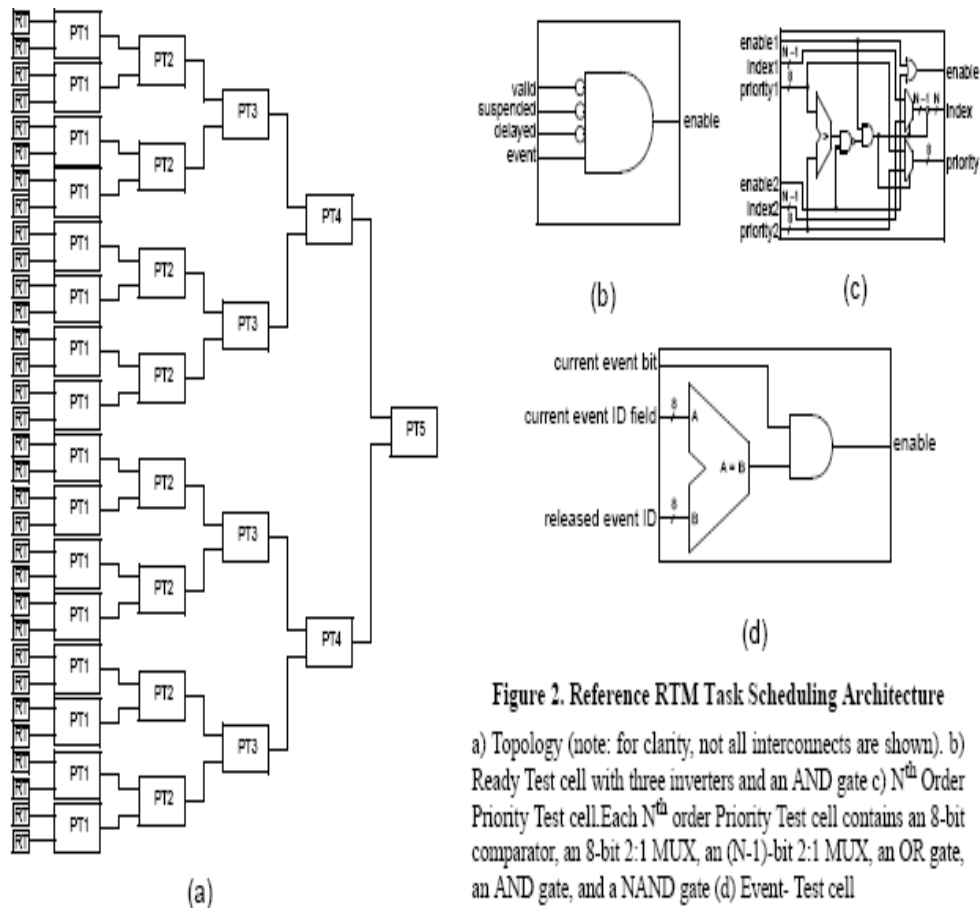


Figure 2. Reference RTM Task Scheduling Architecture

a) Topology (note: for clarity, not all interconnects are shown). b) Ready Test cell with three inverters and an AND gate c) N^{th} Order Priority Test cell. Each N^{th} order Priority Test cell contains an 8-bit comparator, an 8-bit 2:1 MUX, an $(N-1)$ -bit 2:1 MUX, an OR gate, an AND gate, and a NAND gate (d) Event- Test cell

Figure2.ReferenceRTMEventand TimeManagementArchitecture

To eliminate small objects, connected component labelling is applied to the resultant image.(c) represents text detection by applying second set of criteria which eliminates all the objects whose area is less than 300 and filled area is less than 500.

IV. EXPERIMENTS

In order to formally justify the use of the Real-Time Task Manager in actual real-time systems, an accurate quantification of the effects that it has on performance is done by analysing models of real-time systems that use RTOSes. All measurements are made on an in-house cycle-accurate C-based simulator of the Texas Instruments TMS320C6201. The

processor is a high-performance 32-bit VLIW fixed-point DSP with a 200 MHz clock that can issue eight 32-bit instructions per cycle [11]. The simulator is capable of loading the same executable binary files that run in actual systems and executing them exactly as they would on a real processor.

RTOS used in this study: □ C/OSII, a popular commercial pre-emptive RTOS [6] In □ C/OS, semaphores are used to synchronize the access of the data channel by the two tasks. An additional noise task (with a period of 32 ms) purpose

is to insert the type of background noise commonly present in real-time applications, e.g., many systems have LCD displays which have to be periodically refreshed. Aperiodic tasks which respond to external stimuli, such as I/O events are simulated as a by interrupts in MC/OS systems. The I/O events arrival times are geometrically distributed with an average arrival rate of 10 milliseconds.

□ **C/OS-II:** In Figure 4, The RTOS processor utilization has been divided into five categories for system configurations that use □ C/OS. These five categories include Scheduling, Interprocess Communication i.e. event Management, timer Interrupts, processing clock ticks i.e. time management or updating the timer queue and a miscellaneous category which includes task Creation etc.

As seen in the graph in fig. 4, the RTOS processor utilization without the RTM can be quite significant. The bulk of the reduction comes from converting the overhead of time management operations that ready tasks that reach their release times at every clock tick from a linear one to a constant one. The gains in event management and scheduling are lower, 14% and 30% respectively, because of the effectiveness of the software-based bit vector scheme employed in MC/OS. The timer interrupt overhead and miscellaneous category which are not optimized account for a constant less than 1% of the processor utilization, both with and without the RTM. The basic RTOS operations that the RTM implements result in the processing overhead required by □ C/OS to be reduced by 60% to 90%.

V. REAL-TIME TASK MANAGER IMPLEMENTATION

The RTM is a hardware module that implements a task database and thus supports the key RTOS operations: task scheduling, time management, and event management. It is an on-chip peripheral that communicates with the core via a memory-mapped interface. The RTM is pictured in Figure 1; where each record contains information about a single task. The RTM is accessed through the global address space and each record can be read from or written to as if it were another array of structures.

A record is composed of four individually addressable fields. Each record contains a *status field*, containing several bits that describe the status of the corresponding record. The *valid bit* is necessary to indicate if that the record is used by a task. The *delayed bit* indicates that the task is waiting for the amount of clock ticks specified by the *delay field* before being ready-to-run. The *event bit* indicates that the task is pending on the event with the identifier specified by the *event ID field*. The *suspended bit* indicates that the task has been suspended. If the valid bit is set, but the delayed, event, and suspended bits are clear, then the task is ready-to-run. Finally, the *priority field* indicates the task's priority. Also, the maximum number of records is some fixed constant, such as 64 or 256. The RTOS can issue instructions to retrieve certain information from the RTM via a set of memory-mapped registers. The RTM can be queried for the highest priority ready task and thus supports static-priority scheduling. The RTM does time management by decrementing the delay fields of all records by the number of clock ticks specified by the RTOS and updating the delay bit if this value reaches zero. Event management is very similar to static-priority scheduling. The RTM queries for the highest priority task that is pending on a given event identifier, instead of querying its data structure for the highest priority ready task. The software coding is in Turbo C for 1) scheduling 2) Event Management 3) Time Management and the timing analysis was calculated using Pentium Processor. The hardware design was designed and the VHDL coding was written for the hardware circuitry and downloaded using the XILINX-SPARTAN 2E FPGA kit. Results are as follows:

DESCRIPTION	HARWARE SIMULATION TIME	SOFTWARE SIMULATION TIME
Task scheduling	40.999 ns	5 micro s
Event Management	40.999ns	1.67 micro s
Time Management	5.977 ns	30 micro s

Table 1: Simulation result

REFERENCES

- [1] J. Adomat, J. Furunäs, L. Lindh, and J. Stårner. "Real-Time Kernel in Hardware RTU: A step towards deterministic and high performance real-time systems." In *Proceedings of Eighth Euromicro Workshop on Real-Time Systems*, L'Aquila, Italy, June 1996, pp. 164-168.
- [2] M. Bach. *The Design of the UNIX Operating System*. Prentice-Hall, Englewood Cliffs, NJ, 1986.
- [3] K. Baynes, C. Collins, E. Fiterman, B. Ganesh, P. Kohout, C. Smit, T. Zhang, and B. Jacob. "The Performance and Energy Consumption of Three Embedded Real-Time Operating Systems." In *Proceedings of the 2001 International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES '01)*, Atlanta, GA, November 2001, pp. 203-210.
- [4] R. Dick, G. Lakshminarayana, A. Raghunathan, and N. Jha. "Power Analysis of Embedded Operating Systems." In *Proceedings of the 37th Design Automation Conference*, Los Angeles, CA, June 2000.
- [5] International Technology Working Group. *International Technology Roadmap for Semiconductors 2001 Edition: Executive Summary*. Semiconductor Industry Association, 2001.
- [6] J. Labrosse. *MicroC/OS-II: The Real-Time Kernel*. R&D Books, Lawrence, KS, 1999.
- [7] C. Lee, M. Potkonjak, and W. Mangione-Smith. "MediaBench: A Tool for Evaluating and Synthesizing Multimedia and Communications Systems." In *Proceedings of the 30th Annual International Symposium on Microarchitecture (MICRO '97)*, Research Triangle Park, NC, December 1997.
- [8] J. Liu. *Real-Time Systems*. Prentice Hall, Upper Saddle River, NJ, 2000.
- [9] J. Mulder, N. Quach, and M. Flynn. "An Area Model for On-Chip Memories and its Application." *IEEE Journal of SOLID-STATE Circuits*, Vol. 26, No. 2, February 1991, pp. 98-106.
- [10] Texas Instruments. *Code Composer Studio User's Guide*. February 2000.
- [11] J. Turley and H. Hakkarainen. "TI's New C6x DSP Screams at 1,600 MIPS." *The Microprocessor Report*, Vol. 11, 1997, pp. 14-17.
- [12] V. J. Mooney and D. M. Blough. "A Hardware-Software Real-Time Operating System Framework for SOCs." *IEEE Design and Test of Computers*, pp. 44-51, November-December 2002.
- [13] B. E. S. Akgul and V. J. Mooney. "The System-on-a-Chip Lock Cache." *International Journal of Design Automation for Embedded Systems*, 7(1-2) pp. 139-174, September 2002.
- [14] V. Mooney and G. DeMicheli. "Hardware/Software Code Signoff of Run-Time Schedulers for Real-Time Systems." *Design Automation of Embedded Systems*, 6(1), pp. 89-144, September 2000.
- [15] T. Coope. "Embedded Intelligence." *InfoWorld*, November 17, 2000.