# Analysis of Software- and Hardware-interrupt handling (with and without Cache)

Joakim Adomat, Lennart Lindh, E-mail: jat@mdh.se, llh@mdh.se,
Department of Computer Engineering (IDT), Mälardalens Real-Time Center (MRTC)
Mälardalens University, Sweden

**Extended Abstract**

*In this article, we present a comparative study of hardware and software handling of external interrupts. The main parameter of the study is the response time from the hardware interrupt point to the start of interrupt handler task. The test platform used is a stand-alone card (PowerPC - processor, MPC860) with a bus logger (100ns-time tick). The software interrupt system used is a commercial standard real-time kernel, VRTX [VRTX] and the hardware system is RTU (Real-Time Unit) [RTU] with interrupt and scheduling in specialized hardware devices. The total response time in a software solution is the interrupt plus the scheduling latency. In a hardware solution, this is concurrent with the CPU and is faster. The test results are presented for both cases, viz., with and without cache. The study shows that hardware interrupt handling is faster, more predictable, simpler, and safer than the software handler.*

## 1. Introduction

We are dealing with sporadic tasks, and accordingly the response time of concern is the time from the interrupt point to the start of interrupt task. We analyse, the following two main parameters:

- *Interrupt Task Response Time (ITRT) - the time from one high priority interrupt to the interrupt task start.*
- *Multiple Interrupt Task Response Time (MITRT) - the time from more than one high priority interrupts to the interrupt task start.*

The following two different solutions are tested and the response times are analysed

RTU: This is a hardware accelerator for Real-Time Kernels. It contains a scheduler, interrupt handler, semaphores etc, implemented in a hardware device outside the CPU.

VRTX: This is a commercial standard real-time kernel. The interrupt handler is implemented in software.

## 2. Methods of measurement and the physical test bench

The platform on which we made our measurements was a Power PC860 development board [MO860MBX]. A FPGA add-on board was designed and attached to the processor bus expansion slot. To facilitate accurate time measurement, a buslogger was implemented in the FPGA board. In this version, the resolution was 100ns and was considered sufficient. The buslogger asserts one or more interrupts to the system and starts a timer which is read by the interrupt task to measure the response time. There is also functionality for logging the time spent in specified areas in the address space. Address logging is useful for dynamic analysis of ISR WCET and kernel overhead.

## 3. Overall conclusions

The results show that the hardware solution has shorter response time than VRTX and is independent of the system load. The reason is that the RTU is implemented in hardware and works concurrently with the CPU, whereas in VRTX there are many functions in software and will utilize the CPU more.