

Behavior and performance evaluation of Windows Embedded Compact 7 on ATOM

Copyright

© Copyright Dedicated Systems Experts NV. All rights reserved, no part of the contents of this document may be reproduced or transmitted in any form or by any means without the written permission of Dedicated Systems Experts NV, Diepenbeemd 5, B-1650 Beersel, Belgium.

Disclaimer

Although all care has been taken to obtain correct information and accurate test results, Dedicated Systems Experts, VUB-Brussels, RMA-Brussels and the authors cannot be liable for any incidental or consequential damages (including damages for loss of business, profits or the like) arising out of the use of the information provided in this report, even if these organizations and authors have been advised of the possibility of such damages.

Authors

Luc Perneel (1, 2), Hasan Fayyad-Kazan(2) and Martin Timmerman (1, 2, 3)
1: Dedicated Systems Experts, 2: VUB-Brussels, 3: RMA-Brussels

<http://download.dedicated-systems.com>

E-mail: info@dedicated-systems.com

EVALUATION REPORT LICENSE

This is a legal agreement between you (the downloader of this document) and/or your company and the company DEDICATED SYSTEMS EXPERTS NV, Diepenbeemd 5, B-1650 Beersel, Belgium.
It is not possible to download this document without registering and accepting this agreement on-line.

1. **GRANT.** Subject to the provisions contained herein, Dedicated Systems Experts hereby grants you a non-exclusive license to use its accompanying proprietary evaluation report for projects where you or your company are involved as major contractor or subcontractor. You are not entitled to support or telephone assistance in connection with this license.
2. **PRODUCT.** Dedicated Systems Experts shall furnish the evaluation report to you electronically via Internet. This license does not grant you any right to any enhancement or update to the document.
3. **TITLE.** Title, ownership rights, and intellectual property rights in and to the document shall remain in Dedicated Systems Experts and/or its suppliers or evaluated product manufacturers. The copyright laws of Belgium and all international copyright treaties protect the documents.
4. **CONTENT.** Title, ownership rights, and an intellectual property right in and to the content accessed through the document is the property of the applicable content owner and may be protected by applicable copyright or other law. This License gives you no rights to such content.
5. **YOU CANNOT:**
 - You cannot, make (or allow anyone else make) copies, whether digital, printed, photographic or others, except for backup reasons. The number of copies should be limited to 2. The copies should be exact replicates of the original (in paper or electronic format) with all copyright notices and logos.
 - You cannot, place (or allow anyone else place) the evaluation report on an electronic board or other form of on line service without authorisation.
6. **INDEMNIFICATION.** You agree to indemnify and hold harmless Dedicated Systems Experts against any damages or liability of any kind arising from any use of this product other than the permitted uses specified in this agreement.
7. **DISCLAIMER OF WARRANTY.** All documents published by Dedicated Systems Experts on the World Wide Web Server or by any other means are provided "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. This disclaimer of warranty constitutes an essential part of the agreement.
8. **LIMITATION OF LIABILITY.** Neither Dedicated Systems Experts nor any of its directors, employees, partners or agents shall, under any circumstances, be liable to any person for any special, incidental, indirect or consequential damages, including, without limitation, damages resulting from use of OR RELIANCE ON the INFORMATION presented, loss of profits or revenues or costs of replacement goods, even if informed in advance of the possibility of such damages.
9. **ACCURACY OF INFORMATION.** Every effort has been made to ensure the accuracy of the information presented herein. However Dedicated Systems Experts assumes no responsibility for the accuracy of the information. Product information is subject to change without notice. Changes, if any, will be incorporated in new editions of these publications. Dedicated Systems Experts may make improvements and/or changes in the products and/or the programs described in these publications at any time without notice. Mention of non-Dedicated Systems Experts products or services is for information purposes only and constitutes neither an endorsement nor a recommendation.
10. **JURISDICTION.** In case of any problems, the court of BRUSSELS-BELGIUM will have exclusive jurisdiction.

Agreed by downloading the document via the internet.

| | | |
|-------|--|----|
| 1 | Document Intention..... | 6 |
| 1.1 | Purpose and scope | 6 |
| 1.2 | Test framework used: 2.9..... | 6 |
| 1.3 | Conventions | 6 |
| 1.4 | Related documents | 7 |
| 2 | Introduction | 8 |
| 2.1 | Overview | 8 |
| 2.2 | Evaluated (RTOS) Product | 8 |
| 2.2.1 | Software | 8 |
| 2.2.2 | Hardware | 8 |
| 3 | Evaluation results summary..... | 9 |
| 3.1 | Positive points | 9 |
| 3.2 | Negative points | 9 |
| 3.3 | Ratings | 9 |
| 4 | Test Results | 10 |
| 4.1 | Calibration system test (CAL) | 10 |
| 4.1.1 | Tracing overhead (CAL-P-TRC)..... | 10 |
| 4.1.2 | CPU power (CAL-P-CPU) | 11 |
| 4.2 | Clock tests (CLK) | 13 |
| 4.2.1 | Operating system clock setting (CLK-B-CFG) | 13 |
| 4.2.2 | Clock tick processing duration (CLK-P-DUR)..... | 14 |
| 4.3 | Thread tests (THR) | 16 |
| 4.3.1 | Thread creation behaviour (THR-B-NEW) | 16 |
| 4.3.2 | Round robin behaviour (THR-B-RR) | 17 |
| 4.3.3 | Thread switch latency between same priority threads (THR-P-SLS)..... | 17 |
| 4.3.4 | Thread creation and deletion time (THR-P-NEW)..... | 20 |
| 4.4 | Semaphore tests (SEM)..... | 24 |
| 4.4.1 | Semaphore locking test mechanism (SEM-B-LCK) | 24 |
| 4.4.2 | Semaphore releasing mechanism (SEM-B-REL)..... | 25 |
| 4.4.3 | Time needed to create and delete a semaphore (SEM-P-NEW)..... | 25 |
| 4.4.4 | Test acquire-release timings: non-contention case (SEM-P-ARN)..... | 28 |
| 4.4.5 | Test acquire-release timings: contention case (SEM-P-ARC) | 29 |
| 4.5 | Mutex tests (MUT)..... | 32 |
| 4.5.1 | Priority inversion avoidance mechanism (MUT-B-ARC) | 32 |
| 4.5.2 | Mutex acquire-release timings: contention case (MUT-P-ARC) | 33 |
| 4.5.3 | Mutex acquire-release timings: non-contention case (MUT-P-ARN) | 35 |
| 4.6 | Interrupt tests (IRQ) | 37 |
| 4.6.1 | Interrupt latency (IRQ_P_LAT)..... | 37 |
| 4.6.2 | IST to interrupted thread latency (IRQ_P_DLT)..... | 39 |
| 4.6.3 | Interrupt to thread latency (IRQ_P_TLT)..... | 40 |
| 4.6.4 | Maximum sustained interrupt frequency (IRQ_S_SUS)..... | 41 |

Doc: **EVA-2.9-TST-CE7-ATOM**Issue: **v2 on 1-May-2012**Tests Date: **Mar-Apr 2012**

| | | |
|-------|------------------------------------|----|
| 4.7 | Memory tests | 43 |
| 4.7.1 | Memory leak test (MEM_B_LEK) | 44 |
| 5 | Appendix A: Vendor comments | 45 |
| 6 | Appendix B: Acronyms | 46 |

SAMPLE



Doc: **EVA-2.9-TST-CE7-ATOM**

Issue: **v2 on 1-May-2012**

Tests Date: **Mar-Apr 2012**

DOCUMENT CHANGE LOG

| Issue No. | Revised Issue Date | Para's / Pages Affected | Reason for Change |
|-----------|--------------------|-------------------------|---|
| 0.0 | 21-March-2012 | All | Initial draft |
| 1.0 | 18 April-2012 | All | Final draft |
| 2.0 | 1 May-2012 | IRQ | Changing sustained values due to an error |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |

SAMPLE

1 Document Intention

1.1 Purpose and scope

This document presents the quantitative evaluation results of the **Windows Embedded Compact7** OS on ATOM-based platform.

The layout of this report follows the one depicted in “The OS evaluation template” [Doc. 4]. The test specifications can be found in “The evaluation test report definition” [Doc. 3]. For more detailed references, See section “Related documents” of this document. These documents have to be seen as an integral part of this report!

Due to the tightly coupling between these documents, the framework version of “The evaluation test report definition” has to match the framework version of this evaluation report (which is 2.9). More information about the documents and tests versions together with their corresponding relation between both can be found in “The evaluation framework”, see [Doc. 1] in section “Related documents” of this document.

The generic test code used to perform these tests can be downloaded on our website by using the link in the “related documents” section.

1.2 Test framework used: 2.9

This document shows the test results in the scope of the evaluation framework 2.9. More details about this framework are found in Doc 1 (see section “Related documents”).

1.3 Conventions

Throughout this document, we use certain typographical conventions to distinguish technical terms. Our used conventions are the following:

- ❖ ***Bold Italic*** for OS Objects
- ❖ **Bold** for Libraries, packets, directories, software, OSs...
- ❖ `Courier New` for system calls (APIs...)

2 Introduction

This chapter talks about the OS that we are going to test and evaluate, and the hardware on which the under testing OS will be employed to be tested.

2.1 Overview

Releasing a new OS with a different name (changed from **Windows CE** to **Windows Embedded Compact 7**) does not mean that we are up with a new OS! Such naming change was mainly done for marketing purposes, as there were no fundamental changes in the OS itself!

Further in the document, the full name “**Windows Embedded Compact 7**” or the short names “**Compact 7**” and “**CE7**” will be used.

2.2 Evaluated (RTOS) Product

This section describes the OS that Dedicated Systems tested using their Evaluation Testing Suite, and the hardware on which this OS was running during the testing.

2.2.1 Software

The RTOS that will be evaluated and tested is **Windows Embedded Compact 7**. This OS was launched by Microsoft Corporation at the beginning of 2011. In fact, this OS “**Windows Embedded Compact**” is the successor of **Windows CE6R3**.

The tests for evaluating this OS were done in March 2012.

2.2.2 Hardware

We tested the **Windows Embedded Compact 7** on an ATOM-based platform. This platform has the following characteristics:

- Motherboard: Advantech SOM-6760, PCI bus at 33MHz, using the System Controller Hub US15W.
- CPU: Intel Atom Z530 1.6GHz 133MHz Front Side Bus.
 - 32KByte L1 Instruction Cache,
 - 24KByte L1 Write Back Data Cache,
 - 512KByte 8-way L2 Cache (which can be reduced up to zero in some processor sleep states)
 - 1 core with hyper-threading support (however hyper threading was disabled during this test).
- RAM: 512MB DDR2
- VMETRO PCI exerciser in PCI slot 3 (PCI interrupt level D, local bus interrupt level 10)
- VMETRO PBT-315 PCI analyser in PCI slot 4.
- External and CPU internal cache was enabled during the tests.

3 Evaluation results summary

Following is a summary of the results of evaluating **Windows Embedded Compact 7**, released by Microsoft Corporation, Inc.

3.1 Positive points

- All protection primitives use priority inheritance, which is a major plus for achieving real-time behavior
- Good debugging tools: Available also for kernel/driver debugging.
- Very easy to install and to set-up a target (from templates).
- provides the same flexibility as a 32-bit general purpose OS

3.2 Negative points

- The operating system documentation has taken a step backwards compared with the previous versions. A lot of background information is removed.
- Customizing the kernel and adding custom drivers (BSP) stays a daunting task once you go away from the default configurations.
- The remote tool has been changed since last version. We noticed two issues, the more important of which is that there is no officially-supported method to include the remote tools within a device image using Platform Builder. Additionally, we noticed during our testing that establishing a connection between the tools and the target took in excess of a minute, which was longer than our expectation.

3.3 Ratings

For a description of the ratings, see [Doc. 3].

| | | | |
|----------------------|---|---|----|
| RTOS Architecture | 0 | 8 | 10 |
| OS Documentation | 0 | 6 | 10 |
| OS Configuration | 0 | 7 | 10 |
| Internet Components | 0 | 9 | 10 |
| Development Tools | 0 | 8 | 10 |
| Installation and BSP | 0 | 8 | 10 |
| Support | 0 | 8 | 10 |

4 Test Results

Test Results 0  10

As usual, Compact 7 real-time behavior is excellent.

4.1 Calibration system test (CAL)

“Calibration tests” are performed to calibrate the tracing overhead compared with the processing power of the platform. Such tests are important to understand the accuracy of the measurements done in scope of this report, and for measuring the processing power of the platform. This calibration permits comparison with the results on other platforms.

4.1.1 Tracing overhead (CAL-P-TRC)

On this board we used the PCI analyzer as a mechanism to measure event durations.

“Tracing overhead test” calibrates the tracing system overhead. It is more related to the hardware than the OS because its aim is to correct the measured time values.

In the rest of the document, the tracing overhead is subtracted from the obtained results.

In general, the results in this report are correct to +/- 0.2 μseconds. Therefore the results shown in the tables are rounded to 0.1 microseconds.

4.1.1.1 Test results

| Test | result |
|--------------------------|----------|
| Average tracing overhead | 271 nsec |
| minimum tracing overhead | 271 nsec |
| maximum tracing overhead | 271 nsec |

4.1.2 CPU power (CAL-P-CPU)

The “CPU power” test calibrates the CPU performance and the memory bandwidth of the used platform. This test is measured in different situations, starting from the situation where code and data are cached, until the situation where neither code nor data are cached. With such different situation tests, the effects of the cache can be calculated.

We have been seriously reworking this test lately. The CPU test uses only one data address; The non-cached version is about 128KB in size (instructions), while the cached version uses a loop (a bit unrolled to have a small loop overhead but so it fits in the L1 I-cache and it uses only two data words). The instruction cache test is done twice:

- The instructions have not been mapped yet (leading to TLB exceptions and page faults)
- There will not be any page faults (TLB exceptions will still happen).

This gives us some indication about the impact of page faults.

For this specific ATOM platform, we used a factor 5 for the test, so that $5 \times 128\text{KB} = 640\text{KB}$ is larger than the L1/L2 caches. After the test, we divide the results by 5, in order to be capable to compare the results with other platforms

Further, we divided the data cache tests into a read test (reading content of a large array in non-cached case, and read a small array in a loop in the cached case) and a write test. Remark that we flush the caches in between the tests.

This rework shows that a worst-case / best-case scenario can cause significant performance impacts, something that in reality will almost surely never be that large (or you should be able to run everything using only L1 caches).

The impact of either having the code in the I-Cache or not, has serious effect on the results of the tests.

Remark that the results of such tests will depend also, to a high extent, on the cache organization:

- Number of ways
- Line size
- Number of address bits used for index
- Virtual or physical addresses used as index.

4.1.2.1 Test results

The results for our Pentium II 233 MHz platform running Compact 7, averaged over 10 tests, are shown below as a reference:

| Test | no cache | cached | cache effect |
|--------------------------------------|----------|----------|--------------|
| CPU test: first load. | 1.450 ms | | |
| CPU test: ICache effect | 500.5 us | 216.8 us | 2.4 |
| MEM write test | 333.6 us | 309.2 us | 1.1 |
| MEM read test | 253.2 us | 166.4 us | 1.5 |
| Average caching effect (CPU and MEM) | | | 1.7 |

The results for the same code for **Compact 7** on this ATOM platform are shown below:

| Test | no cache | cached | cache effect |
|--------------------------------------|----------|---------|--------------|
| CPU test: first load. | 62.5 us | | |
| CPU test: ICache effect | 57.8 us | 16.7 us | 3.5 |
| MEM write test | 35.8 us | 22.6 us | 1.6 |
| MEM read test | 33.1 us | 27.7 us | 1.2 |
| Average caching effect (CPU and MEM) | | | 2.1 |

Comparing both platforms show the enormous increase in speed! ATOM seems to be almost 10 times faster than MMX platform.

Here are some conclusions regarding the ATOM:

- Caching of instructions has the most impact! This is logical because for each instruction, memory has to be fetched containing this instruction. When handling data, you will always have some instructions without data access (register manipulations and operations) which are not impacted by the data cache.
- A strange behavior was noticed on this processor: caching has a major impact on data writes compared with data reads. On the other processor, caching impact is less (as writes can be postponed).

- In general, this platform is less impacted by cache misses compared with, for instance, the ARM Beagle platform. It seems that the RAM speed is still adequate compared with the caches.

The results cannot be completely compared with other tests that we did on the same platform. Even if the same code is used, these figures can be different depending on compiler optimizations and compiler versions.

4.2 Clock tests (CLK)

“Clock tests” measure the time needed by the operating system to handle its clock interrupt. On the tested platform, the clock tick interrupt is set on the highest hardware interrupt level, interrupting any other thread or interrupt handler.

4.2.1 Operating system clock setting (CLK-B-CFG)

The “OS clock setting” test examines the setting of the clock tick period in the operating system. This test shows the default clock timing as they are set by the BSP and/ or the kernel.

For this test, the `SleepTillTick()` function call is used. Following the manual, the delay should be based on the clock tick and wait until the next clock tick.

4.2.1.1 Test results

| Test | result |
|------------------------|--------|
| Test succeeded | Yes |
| Tested clock period | 1ms |
| Clock period adaptable | NO |

4.2.2 Clock tick processing duration (CLK-P-DUR)

The “clock tick processing duration” test examines the clock tick processing duration in the kernel. The test results are extremely important, as the clock interrupt will disturb all the other performed measurements.

The bottom line of the figures in section 4.2.2.2 represents the normal loop time of the test if no clock interrupt occurs during the test loop. The upper line is generated by the samples when a clock interrupt occurred during the loop. The difference between the two lines is the clock tick processing duration.

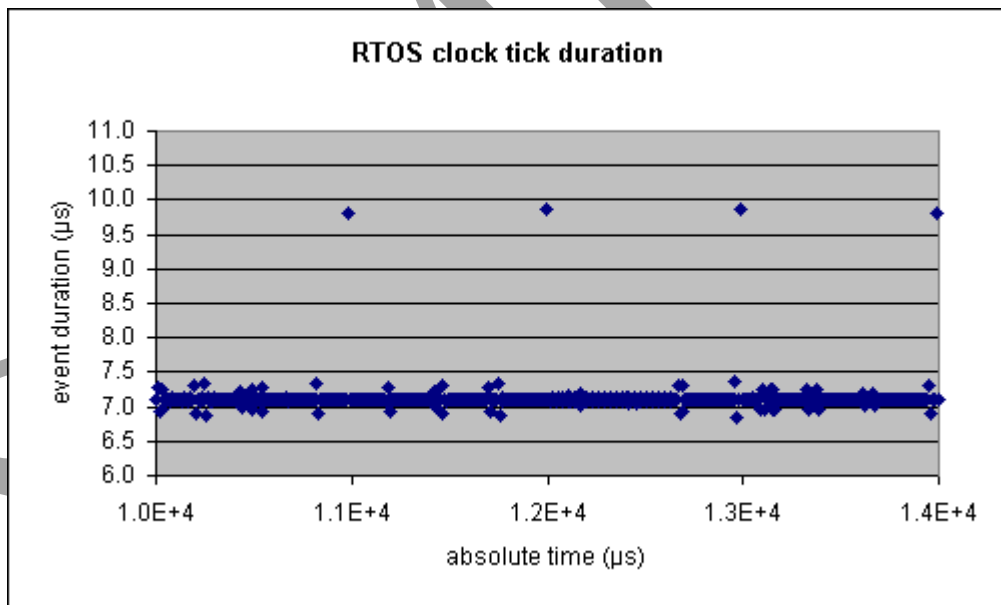

Be careful: the OAL layer may stop the clock tick interrupt if not needed (tickles kernel). However, this will happen only if the CPU load is minimal, in which case the used memory bandwidth will be small.

Even that an ATOM is much faster than our standard Pentium MMX 200MHz evaluation platform, interrupt latency is not that much better. A part of the interrupt latency is also hardware dependent and thus not only dependent on CPU speed.

4.2.2.1 Test results

| Test | result |
|-------------------------------------|-----------------------------------|
| CLOCK_LOOP_COUNTER | 1000 |
| Normal busy loop time | 7.0 μ s |
| Busy loop time with clock interrupt | 10 μ s, worst case 19 μ s |
| Clock interrupt duration | 3 μ s to 12 us |

4.2.2.2 Diagrams



Zoomed in extract from diagram above.

4.3 Thread tests (THR)

“Thread tests” measure the scheduler performance.

4.3.1 Thread creation behaviour (THR-B-NEW)

The “thread creation behavior” test examines the OS behavior when it creates threads. This test attempts to answer the question: Does the OS behave as it should in order to be considered a real-time operating system? Following scenarios are tested:

- If a thread is created with a lower priority than the creating thread, then are we sure that it is not activated until the creating thread is finished?
- If a thread is created with the same priority as the creating thread, will it be placed at the ready tail?
- When yielding after the creation in the above test, does the newly created thread becomes active?
- If a thread is created with a higher priority than the creating thread, is it then immediately activated?

This test succeeded without any problems.

4.3.1.1 Test results

| Test | result |
|-------------------------------|--------|
| Test succeeded | YES |
| Lower priority not activated? | YES |
| Same priority at tail? | YES |
| Yielding works? | YES |
| Higher priority activated? | YES |

4.3.2 Round robin behaviour (THR-B-RR)

The “round robin behavior” test checks if the scheduler uses a fair round robin mechanism to schedule threads that use the round robin scheduling policy, are of the same priority, and are in the ready-to-run state!

☹ A problem was detected here: the first time a thread becomes active, it takes a longer time slice (100ms) than in the other cases (1ms = clock tick).

We took a look back to the test results of **CE 6.0** and discovered that the same problem was indeed present there as well. To be sure that we didn't do anything wrong in the test, we compared the code for this test with the test code of other RTOS which did not have this behavior and no differences were found.

The problem is caused by the fact that the default round robin time slice for a thread used in CE 7.0 is 100ms. When adapting the time slice to 1ms, then this will only have impact on the next scheduling round. Even if you set the time slice before starting the thread (suspended mode) it will only be applied on the next end of the time slice.

Although it is strange behavior, creating dynamically threads in a real-time system is a bad practice which is normally never done. As such, this problem will not have any impacts in real use cases.

4.3.2.1 Test results

| Test | result |
|-----------------------------------|---|
| Test succeeded | No (first time slice after thread creation is longer) |
| RR Time slice following this test | 1 clock tick normally (first slice takes 100 clock ticks) |

4.3.3 Thread switch latency between same priority threads (THR-P-SLS)

The “thread switch latency between same priority threads” test measures the time needed to switch between threads of the same priority. For this test, threads must voluntarily yield the processor for other threads.

This test was performed in order to generate the worst-case behavior. We performed the test with an increasing number of threads, starting with two (2) and going up to 1000 in order to observe the behavior in a worst-case scenario. As we increase the number of active threads, the caching effect becomes evident since the thread context will no longer be able to reside in the cache (on this platform the L1 caches are 32KB, for the instruction cache and 28KB for the data cache).


As the L2 cache is much larger, 512KB, it is not likely that the clock tick will not be located in this cache anymore.



4.3.3.1 Test results

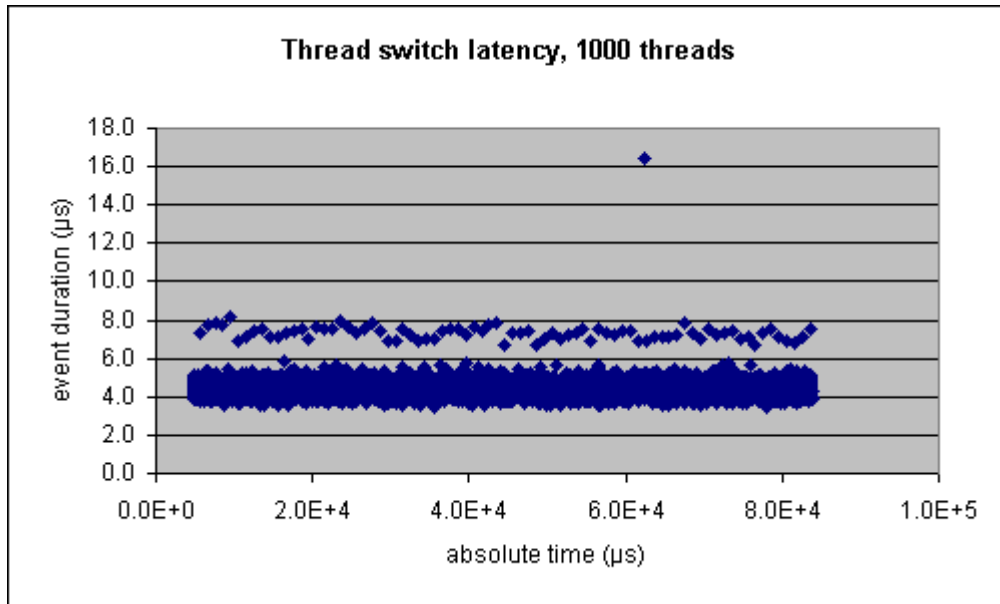
| Test | result |
|----------------|--------|
| Test succeeded | YES |

| Test | Sample qty | Avg | Max | Min |
|-------------------------------------|------------|-------------|--------------|-------------|
| Thread switch latency, 2 threads | 16383 | 2.1 μ s | 5.9 μ s | 1.8 μ s |
| Thread switch latency, 10 threads | 16383 | 2.3 μ s | 6.7 μ s | 1.9 μ s |
| Thread switch latency, 128 threads | 16320 | 3.7 μ s | 14.6 μ s | 2.6 μ s |
| Thread switch latency, 1000 threads | 16383 | 4.4 μ s | 16.4 μ s | 3.5 μ s |

4.3.3.2 Diagrams







4.3.4 Thread creation and deletion time (THR-P-NEW)

The “thread creation and deletion time” test examines the time required to create a thread, and the time required to delete a thread in the following different scenarios:

- Scenario 1 “never run”: The created thread has a lower priority than the creating thread and is deleted before it has any chance to run. No thread switch occurs in this test.
- Scenario 2 “run and terminate”: The created thread has a higher priority than the creating thread and will be activated. The created thread immediately terminates itself (thread does nothing).
- Scenario 3 “run and block”: The same as the previous scenario (scenario 2: run and terminate), but the created thread does not terminate (it lowers its priority when it is activated).

In the scenarios where the thread actually runs (2 and 3), the creation time is measured as the duration from the system call creating the thread until the time when the created thread is activated. For the “never run” scenario, the creation time is measured as the duration of the system call.


In this test we see clearly the impact of the caches: the first time everything has to be retrieved from RAM (large number of instructions which are not cached, this combined with some data) and as a result, the first measured event takes more than twice the time than the subsequent events.


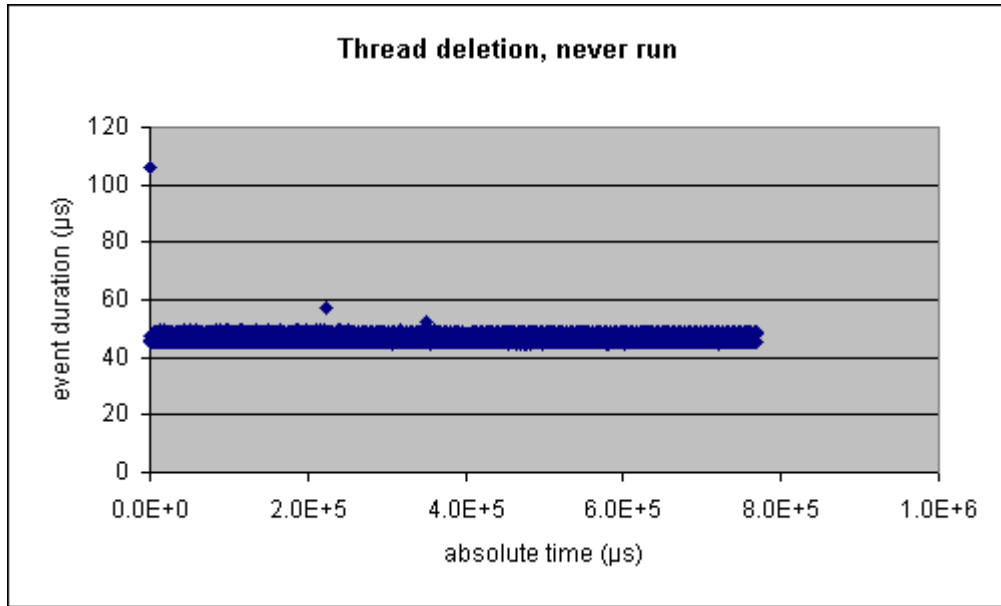
4.3.4.1 Test results



| Test | result |
|----------------|--------|
| Test succeeded | YES |

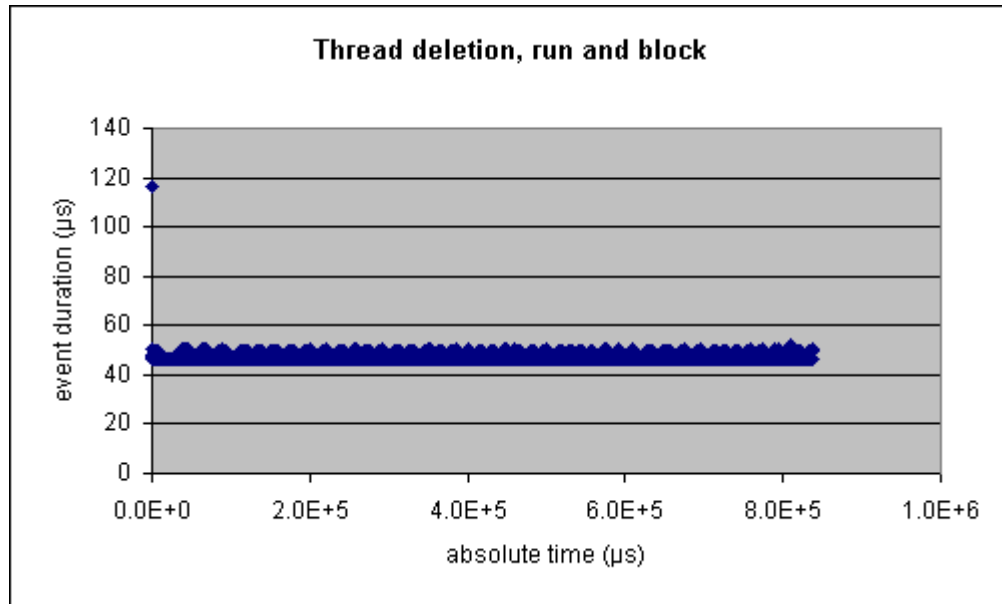
| Test | Sample qty | Avg | Max | Min |
|------------------------------------|------------|--------------|-------------|--------------|
| Thread creation, never run | 7500 | 51.6 μ s | 136 μ s | 50.7 μ s |
| Thread deletion, never run | 7500 | 45.6 μ s | 106 μ s | 44.9 μ s |
| Thread creation, run and terminate | 7500 | 56.2 μ s | 149 μ s | 54.9 μ s |
| Thread deletion, run and terminate | 7500 | 2.1 μ s | 5.4 μ s | 1.6 μ s |
| Thread creation, run and block | 7500 | 56.1 μ s | 149 μ s | 54.7 μ s |
| Thread deletion, run and block | 7500 | 46.6 μ s | 117 μ s | 46.0 μ s |

4.3.4.2 Diagrams









4.4 Semaphore tests (SEM)

“Semaphore tests” examine the behavior and performance of the OS counting semaphore. The counting semaphore is a system object that can be used to synchronize threads.

4.4.1 Semaphore locking test mechanism (SEM-B-LCK)

In this test, we verify if the counting semaphore locking mechanism works as it is expected to work. If this mechanism works as expected, then:

- The P () call will block only when the count is zero.
- The V () call will increment the semaphore counter.
- In the case where the semaphore counter is zero, the V () call will cause a rescheduling by the OS, and blocked threads may become active.

The semaphore behaves correctly as a protection mechanism.

4.4.1.1 Test results

| Test | result |
|--------------------------|---------------------------|
| Test succeeded | YES |
| Maximum semaphore value? | Limited by the “int” type |
| Rescheduling on free? | OK |

4.4.2 Semaphore releasing mechanism (SEM-B-REL)

The “semaphore releasing mechanism” test verifies that the highest priority thread being blocked on a semaphore will be released by the release operation. This action should be independent of the order of the acquisitions taking place.

Windows Embedded Compact 7 passed this test.

4.4.2.1 Test results

| Test | result |
|----------------|--------|
| Test succeeded | YES |

4.4.3 Time needed to create and delete a semaphore (SEM-P-NEW)

The “time needed to create and delete a semaphore” test is performed to gain an insight about the time needed to create a semaphore and the time needed to delete it. The deletion time is checked in two cases:

- The *semaphore* is used between the creation and deletion.
- The *semaphore* is NOT used between the creation and deletion.

Remark that although we do not use “named” *semaphores*, there seems to be a system call required to create/delete a semaphore.



The clock tick line is clearly visible.



4.4.3.1 Test results

| Test | result |
|----------------|--------|
| Test succeeded | YES |

| Test | Sample qty | Avg | Max | Min |
|-------------------------------------|------------|-------------|--------------|-------------|
| Semaphore creation time, used | 7500 | 1.8 μ s | 9.8 μ s | 1.4 μ s |
| Semaphore deletion time, used | 7500 | 2.3 μ s | 15.5 μ s | 1.8 μ s |
| Semaphore creation time, never used | 7500 | 1.7 μ s | 9.5 μ s | 1.3 μ s |
| Semaphore deletion time, never used | 7500 | 2.0 μ s | 15.4 μ s | 1.6 μ s |

4.4.3.2 Diagrams





4.4.4 Test acquire-release timings: non-contention case (SEM-P-ARN)


The “acquire-release timings: non-contention case” test measures the acquisition and release time in the non-contention case. Since in this test the semaphore does not neither block nor causes any rescheduling (thread switching), the duration of the call should be short.

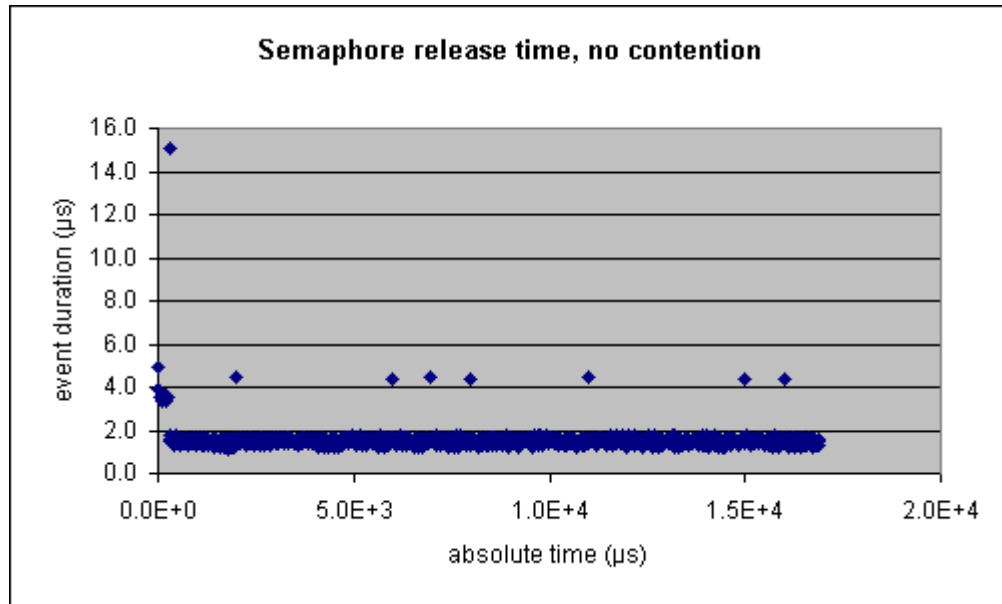
4.4.4.1 Test results

| Test | result |
|----------------|--------|
| Test succeeded | YES |

| Test | Sample qty | Avg | Max | Min |
|---|------------|-------------|--------------|-------------|
| Semaphore acquisition time, no contention | 3850 | 1.8 μ s | 6.1 μ s | 1.4 μ s |
| Semaphore release time, no contention | 3849 | 1.5 μ s | 15.1 μ s | 1.1 μ s |

4.4.4.2 Diagrams





4.4.5 Test acquire-release timings: contention case (SEM-P-ARC)

The “acquire release timings: contention case” test is performed to test the time needed to acquire and release a semaphore, depending on the number of threads blocked on the semaphore. It measures the time in the contention case when the acquisition and release system call causes a rescheduling to occur.

The purpose of this test is to see if the number of blocked threads has an impact on the times needed to acquire and release a semaphore. It attempts to answer the question: “How much time does the OS needs to find out which thread should be scheduled first?”

In this test, since each thread has a different priority, the question is how the OS handles these pending thread priorities on a semaphore. To have a more clear view on our test, you can take a look on the expanded diagrams during a small time frame (e.g. one test loop):

- We create 128 threads with different priorities. The creating thread has a lower priority than the threads being created.
- When the thread starts execution, it tries to acquire the *semaphore*; but as it is taken, the thread stops and the kernel switch back to the creating thread. The time from the acquisition attempt (which fails) to the moment the creating thread is activated again is called here the “acquisition time”. Thus, this time includes the thread switch time.
- Thread creation takes some time, so the time between each measurement point is large compared with most other tests.

- After the last thread is created and is blocked on the *semaphore*, the creating thread starts to release the *semaphore* repeating this action the same number of times as the number of blocked threads on the semaphore.
- We start timing at the moment the *semaphore* is released which in turn will activate the pending thread with the highest priority, which will stop the timing (thus again the thread switch time is included).

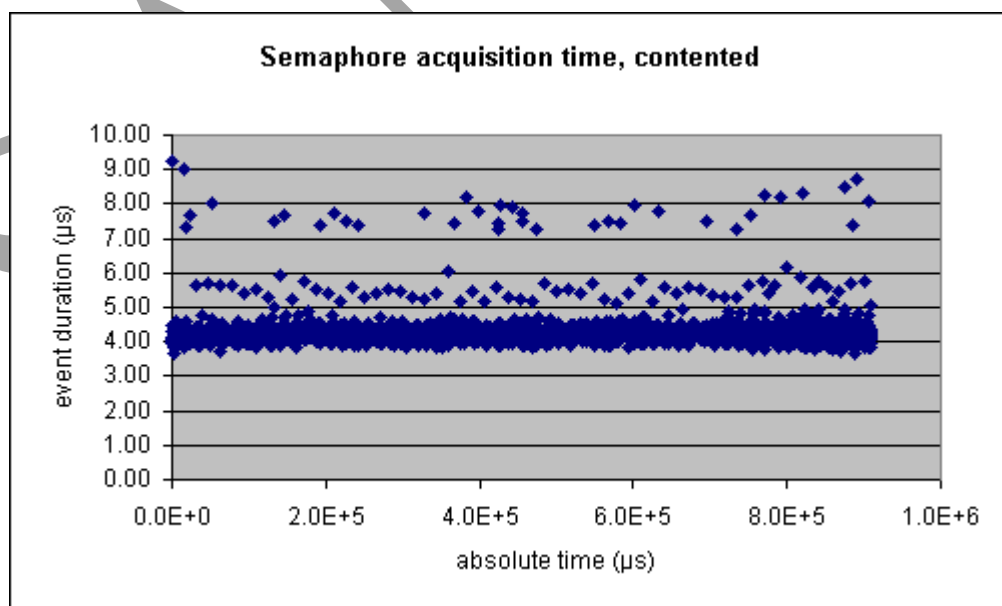
Now, the most important part of this test is to see if the number of threads pending on a *semaphore* has an impact on release times. Clearly, it doesn't, so this is good.

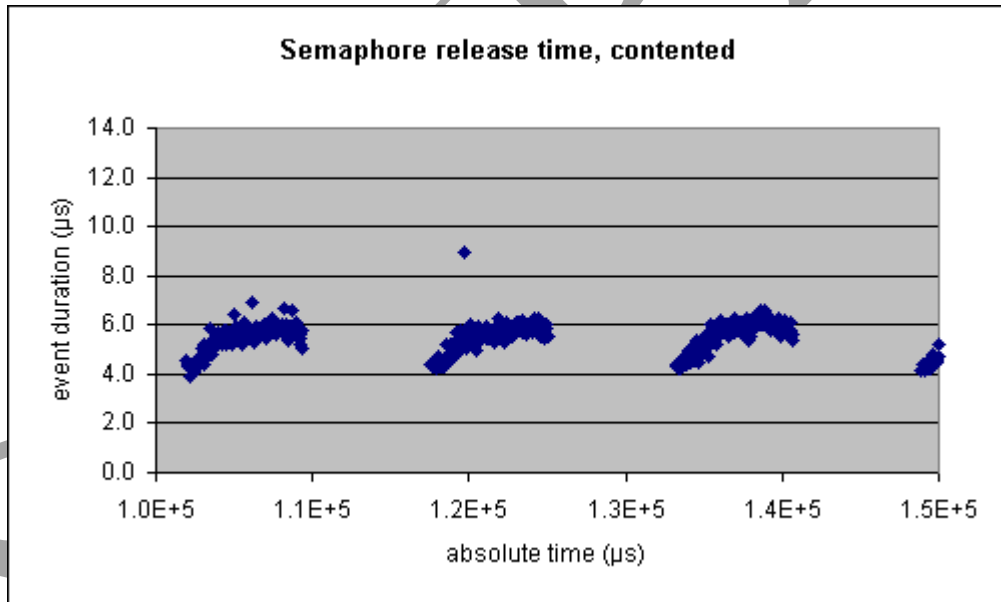
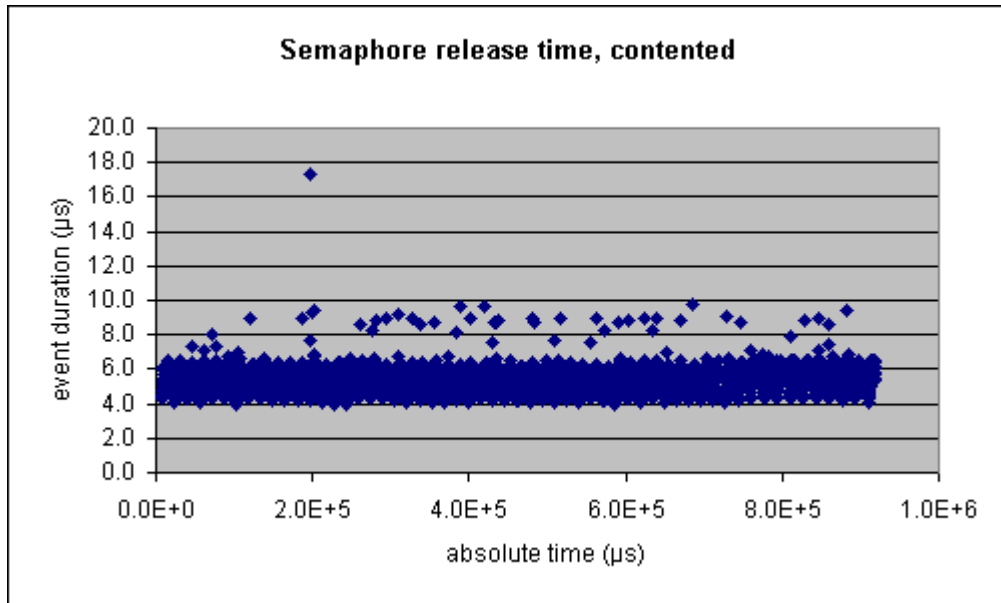
4.4.5.1 Test results

| Test | result |
|-------------------------------|--------|
| Test succeeded | YES |
| Max number of threads pending | 128 |

| Test | Sample qty | Avg | Max | Min |
|---------------------------------------|------------|-------------|--------------|-------------|
| Semaphore acquisition time, contented | 7424 | 4.2 μ s | 9.3 μ s | 3.6 μ s |
| Semaphore release time, contented | 7424 | 5.5 μ s | 17.3 μ s | 3.9 μ s |

4.4.5.2 Diagrams





Zoomed in version of previous diagram.

4.5 Mutex tests (MUT)

Our “mutex tests” help us evaluate the behavior and performance of the mutual exclusive semaphore.

Although the mutual exclusive semaphore (further called mutex) is usually described as being the same as a counting semaphore where the count is one, this is not true. The behavior of a mutex is completely different than the behavior of a semaphore. Unlike semaphores, mutexes use the concept of a “lock owner”, and can thus be used to prevent priority inversions. Semaphores cannot do this, and it goes without saying that mutexes (and not semaphores) should not be used semaphores for critical section protection mechanisms. In scope of the framework, this test will look into detail of a mutex system object that avoids priority inversion. The one used is the critical section one.

Remark that, Windows Embedded Compact 7 has as well a *Mutex* system object; but this should be used only between processes as it always requires a long round-trip to the kernel even if the lock is not contented.

Remark as well that there exists InterlockedXXX functions, which use the available CPU instruction set to provide atomic behavior and as a result, these are fast.

4.5.1 Priority inversion avoidance mechanism (MUT-B-ARC)

The “priority inversion avoidance mechanism” test determines if the system call being tested prevents the priority inversion case. To check this possibility, the test artificially creates a priority inversion.

The behavior test was performed using the Windows Embedded Compact 7 Critical section object. The mutex object can be used between processes but therefore it requires of course each time a call to the kernel.

Therefore it is better to use the Critical Section object for protection between threads within a process. In this case, no kernel calls are needed if the lock is not contented, which increases a lot the speed of the lock under normal usage scenarios.

Priority inversion behaves as expected for both lock objects.

4.5.1.1 Test results

| Test | result |
|--|--|
| Priority inversion avoidance system call present | Yes |
| System call used | InitializeCriticalSection, EnterCriticalSection, LeaveCriticalSection. |
| Test succeeded | YES |
| Priority inversion avoided | YES |
| Mechanism used if any? | Priority inversion cannot be disabled in Compact 7 , which is a plus! |

4.5.2 Mutex acquire-release timings: contention case (MUT-P-ARC)

The “mutex acquire-release timings: contention case” test is the same test as the “priority inversion avoidance mechanism” test described above, but performed in a loop. In this case, we measure the time needed to acquire and release the mutex in the priority inversion case.

Our test is designed so that the acquisition enforces a thread switch:

- The acquiring thread is blocked
- The thread with the lock is released.

We measured the acquisition time from the request for the mutex acquisition to the activation of the lower priority thread with the lock.

Note that before the release, an intermediate priority level thread is activated (between the low priority one having the lock and the high priority one asking the lock). Due to the priority inheritance, this thread does not start to run (the low priority thread having the lock inherited the high priority of the thread asking the lock).

We measured the release time from the release call to the moment the thread requesting the mutex was activated; so this measurement also includes a thread switch.

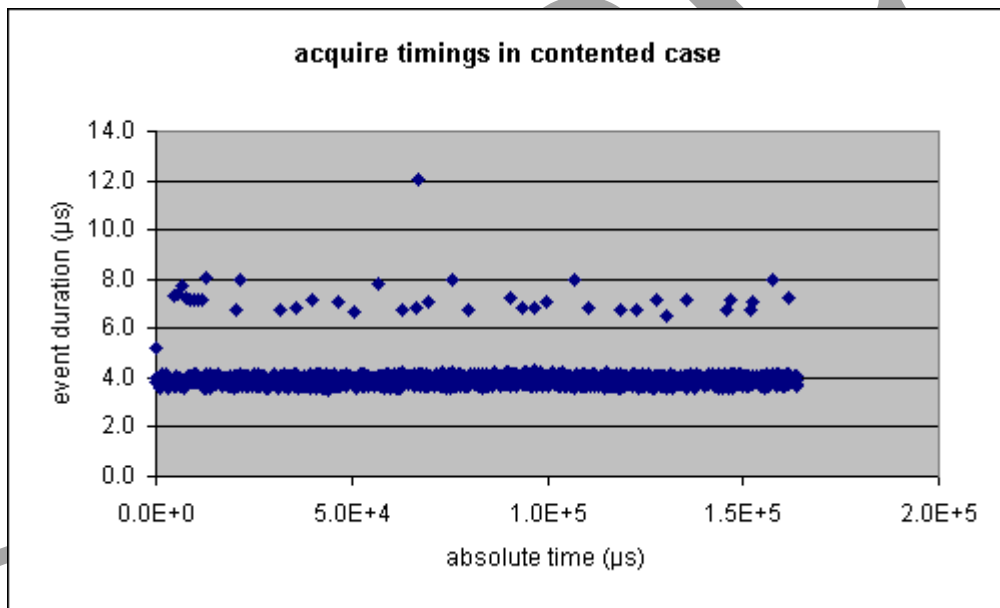
The clock tick interrupt can be clearly seen (as usual) (figures of section 4.5.2.2).

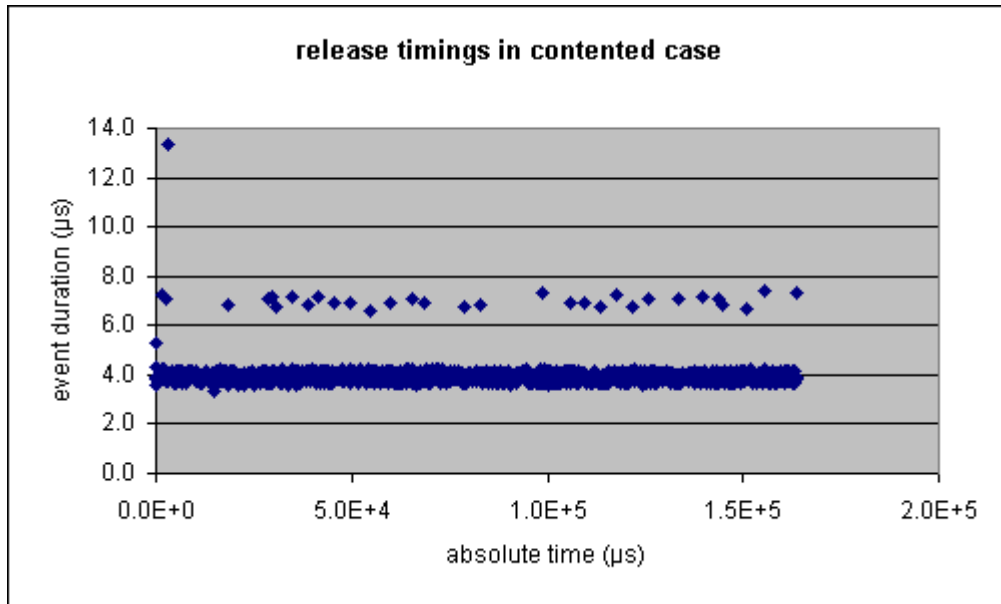
4.5.2.1 Test results

| Test | result |
|----------------|--------|
| Test succeeded | Yes |

| Test | Sample qty | Avg | Max | Min |
|------------------------------------|------------|-------------|--------------|-------------|
| Mutex acquisition time, contention | 7500 | 3.9 μ s | 12.0 μ s | 3.5 μ s |
| Mutex release time, contention | 16383 | 3.9 μ s | 13.3 μ s | 3.3 μ s |

4.5.2.2 Diagrams:





4.5.3 Mutex acquire-release timings: non-contention case (MUT-P-ARN)

The “mutex acquire-release timings: no contention case” test measures the overhead incurred by using a lock when this lock is not owned by any other thread. Well-designed software will use non-contended locks most of the time, and only in some rare cases the lock will be taken by another thread.

Therefore, it is important that the non-contention case should be fast. Remark that this is only possible if:

- The CPU supports some type of atomic instruction, so that no system call is needed when no contention is detected.
- A lock is not shared between processes.

The last requirement is valid only when the Critical Section object in Compact 7 is used, because the Mutex object is an object meant to be used in a shared process scenario and does not contain this optimization.


Performance is excellent and too small to measure!


4.5.3.1 Test results

| Test | result |
|----------------|--------|
| Test succeeded | Yes |

| Test | Sample qty | Avg | Max | Min |
|---------------------------------------|------------|--------------|--------------|--------------|
| Mutex acquisition time, no contention | 7500 | <0.1 μ s | <0.1 μ s | <0.1 μ s |
| Mutex release time, no contention | 7500 | <0.1 μ s | 1.6 μ s | <0.1 μ s |

4.5.3.2 Diagrams:





4.6 Interrupt tests (IRQ)

“Interrupt tests” evaluate how the operating system performs when handling interrupts.

Interrupt handling is a key system capability of real-time operating systems. Indeed, RTOSs are typically event driven.

For our interrupt tests, we use a PCI board to generate interrupts. The PCI board we used has an independent programmable interrupt generator, which protects it from influence by the RTOS clock. This protection allows us to guarantee that an independent interrupt source is not synchronized in any way with the platform being tested.

4.6.1 Interrupt latency (IRQ_P_LAT)

The “interrupt latency” test measures the time it takes to switch from the hardware interrupt to an interrupt handler. This time is measured from the moment the interrupt line on the PCI is toggled up to the instruction in the interrupt handler which clears the interrupt. Therefore the measurement does take into account the hardware interrupt latency as well.

Remark that in Windows Embedded Compact 7, the interrupt handler is normally already a thread. However on this platform we added our own interrupt handler which is run in interrupt context and activates then the Interrupt Servicing Thread (IST).

The 2us between the two lines was caused by outstanding PCI transaction which has to be finished. This transaction takes 1.9µs, so if the CPU wanted to jump to the interrupt context with an outstanding transaction then the context switch is still blocked for 1.9us which gives the extra


delay. However this transaction gives us a good way to detect when the interrupt line was toggled. Anyhow, this impacts more the hardware latency than the latency introduced by the OS.

The clock time is easily detected again (it has the highest interrupt level).

4.6.1.1 Test results

| Test | Sample qty | Avg | Max | Min |
|----------------------------|------------|--------|---------|--------|
| Latency from HW IRQ to ISR | 8173 | 6.7 us | 16.7 us | 5.6 us |

4.6.1.2 Diagrams



4.6.2 IST to interrupted thread latency (IRQ_P_DLT)


The “interrupt dispatch latency” test measures the time the OS takes to switch from the interrupt servicing thread (standard mechanism to handle interrupts in Windows CE) back to the interrupted thread.

4.6.2.1 Test results

| Test | Sample qty | Avg | Max | Min |
|--|------------|--------|--------|--------|
| Latency from IST to interrupted thread | 8172 | 3.0 us | 8.5 us | 2.8 us |

4.6.2.2 Diagrams

SAMPLE




4.6.3 Interrupt to thread latency (IRQ_P_TLT)

The “interrupt to thread latency” test measures the time the OS takes to switch from the interrupt handler (running in interrupt context) to the Interrupt Servicing Thread.

4.6.3.1 Test results

| Test | Sample qty | Avg | Max | Min |
|--|------------|--------|--------|--------|
| Latency from ISR to Interrupt Servicing Thread | 8173 | 2.7 us | 8.8 us | 2.6 us |

4.6.3.2 Diagrams



4.6.4 Maximum sustained interrupt frequency (IRQ_S_SUS)

The “maximum sustained interrupts frequency” test measures the probability that an interrupt might be missed. It attempts to answer the question: Is the interrupt handling duration stable and predictable?

Remark first that for this test, an interrupt service thread is used; so compared with other RTOS, extra context switch latency should be added (needed to be added twice in this case: interrupt start and interrupt leave).

We note that the rather long duration to handle the clock tick interrupt is due to the implementation by the BSP vendor, which in this case came from a 3rd party. We expect that improvements to the BSP would decrease the long duration.

This test is done in 3 phases:

- 1000 interrupts as an initial phase: a fast test just to see where we have to start searching.
- 1 000 000 interrupts as a second phase based on the results from the first phase. This test still takes less than a minute and gives already accurate results.
- 1 billion interrupts as a last phase, which takes few hours, and sometimes more than 24 hours, depending on the used platform and OS. This phase is done to verify stability; therefore, we cannot run this phase many times, especially when it comes to large interrupt latencies.

As one can observe in the above interrupt test results, the interrupt latency is 6.7 μ s in the best case. But due to the clock interrupt, the latency increases. For a long run, we need about 44 μ s not to miss one interrupt, as shown in the test results below.

4.6.4.1 Test results

| Interrupt period | #interrupts generated | #interrupts lost |
|------------------|-----------------------|------------------|
| 10.1 μ s | 1000 | 7 |
| 10.6 μ s | 1000 | 0 |
| 16 μ s | 10 000 | 1 |
| 16.4 μ s | 10 000 | 0 |
| | | |

| Interrupt period | #interrupts generated | #interrupts lost |
|------------------|-----------------------|------------------|
| 18.3 μ s | 1 000 000 | 4 |
| 21 μ s | 1 000 000 | 0 |
| | | |
| 24 μ s | 10 000 000 | 0 |
| | | |
| 24 μ s | 100 000 000 | 0 |
| | | |
| 26 μ s | 1 billion | 204 |
| 31 μ s | 1 billion | 0 |
| | | |

SAMPLE

4.7 Memory tests

This examines the memory leaks of OS.

4.7.1 Memory leak test (MEM_B_LEK)

This test continuously create/remove objects in the operating system (threads, *semaphores*, *mutexes* ...).

| Test | result |
|--|----------|
| Test succeeded | YES |
| Test duration (how long we let the endless loop run) | >10h |
| Number of main test loops done | > 50 000 |

SAMPLE

5 Appendix A: Vendor comments

All vendor comments were integrated within the document as there were no disagreements.

SAMPLE

6 Appendix B: Acronyms

| Acronym | Explanation |
|---------|---|
| API | Application Programmers Interface: calls used to call code from a library or system. |
| BSP | Board Support Package: all code and device drivers to get the OS running on a certain board |
| DSP | Digital Signal Processor |
| FIFO | First In First Out: a queuing rule |
| GPOS | General Purpose Operating System |
| GUI | Graphical User Interface |
| IDE | Integrated Development Environment (GUI tool used to develop and debug applications) |
| IRQ | Interrupt Request |
| ISR | Interrupt Servicing Routine |
| MMU | Memory Management Unit |
| OS | Operating System |
| PCI | Peripheral Component Interconnect: bus to connect devices, used in all PCs! |
| PIC | Programmable Interrupt Controller |
| PMC | PCI Mezzanine Card |
| PrPMC | Processor PMC: a PMC with the processor |
| RTOS | Real-Time Operating System |
| SDK | Software Development Kit |
| SoC | System on a Chip |
| | |