Impact Case Study: The world's smallest automotive real-time operating system

Robert I. Davis
Real-Time Systems Research Group,
Department of Computer Science,
University of York, York, UK.
rob.davis@york.ac.uk

Nigel Tracey
ETAS Ltd.
York, UK
nigel.tracey@etas.com

Impact Summary

Research from the Real-Time Systems Group at the University of York was exploited to design an exceptionally efficient Real-Time Operating System (RTOS), used in automotive Electronic Control Units (ECUs), and its associated schedulability analysis tools. The RTOS has been deployed in over 1 billion ECUs. It has been standardised upon by many of the world's leading automotive powertrain systems and chassis electronics suppliers, and is used in cars produced by nearly all of the major car manufacturers world-wide.

Background

In real-time embedded systems, such as the ECUs used in vehicles, system functionality is decomposed into multiple software tasks running on a microprocessor. The system requirements place time constraints on these tasks. Hence a task may be required to execute every 10 milliseconds, read and process data from sensors, and output its results within a specific time constraint or deadline. When there are multiple tasks with different periods and deadlines running on the same microprocessor, an RTOS is needed to schedule when each task should execute. It is essential that all of the tasks are guaranteed to meet their deadlines during operation; otherwise the system may suffer from intermittent timing faults that compromise its functionality and reliability.

Given the complex behaviour of these systems, it is impossible to obtain a 100% guarantee that tasks will always meet their deadlines via testing. Instead, a rigorous scientific and systematic solution to this problem is s*chedulability analysis*; a set of techniques used to determine off-line if each task can be guaranteed to meet its deadline under a specific scheduling policy. Schedulability analysis is used to compute the *worst-case response time*, the longest time that can elapse from a task being released to it outputting its results and completing execution. If this is less than the deadline, then the task can be guaranteed to always meet its time constraints.

Research

In the 1990's seminal research into schedulability analysis [1], [2], [3], [4], and [5] for fixed priority pre-emptive scheduling, originally called *Deadline Monotonic Schedulability Analysis* but now widely referred to as *Response Time Analysis*, was introduced by the Real-Time Systems Research Group (RTSRG) at the University of York.

This analysis is applicable to fixed priority scheduling, and a task model that accurately accounts for the detailed timing behaviours of tasks in automotive systems. These timing behaviours include: tasks that are invoked sporadically (i.e. with minimum inter-arrival times, but not necessarily strictly periodically in time – for example tasks that are triggered off of a crank angle sensor measuring engine rotation); tasks with deadlines that are less than their periods and prior to completion [1], [2]– - accounting for tasks that need to make a response prior to their next invocation to avoid buffer overruns, and to carry out further computations after a response has been made, in preparation for the next cycle; tasks with offset release times [4] – used as a means of avoiding peak load in short time intervals; tasks with jittered released times [5] – that are triggered by the arrival of messages that can take a variable amount of time to be transmitted, and tasks that share resources [1], [2] – such as data structures and peripheral devices used for communication. The analysis also accounts for the overheads of a well-designed RTOS [3].

This research therefore introduced for the first time, schedulability analysis that could be applied in practice to commercial real-time systems, providing a rigorous approach to obtaining timing correctness. This was recognised in the EPSRC International Review of Computer Science undertaken in 2002:

The techniques developed built upon other important research contributions such as the Stack Resource Policy [11] for resource locking; however, without the work of the researchers in the RTSRG, the impact would not have been possible due to the fact that the underlying models used by prior schedulability analyses were too limited to be used in practice.

Route to Impact

In 1997, Robert Davis and Ken Tindell (both previously members of the RTSRG) co-founded a company called Northern Real-Time Applications (NRTA) Ltd., with the aim of developing an RTOS and schedulability analysis tools specifically tailored to automotive applications that use low cost microcontrollers. In doing so, they utilised the research that they had been involved in and heavily exposed to while at the University of York

There were two fundamental design goals:

1. The real-time behaviour of systems built using the RTOS must be fully analysable using schedulability analysis tools. In other words the behaviour of the RTOS must match the assumptions of the underpinning schedulability analysis techniques.
2. The memory and execution time overheads of the RTOS must be significantly less than those of any other RTOS available for use in automotive applications.

Robert Davis led the team that developed the SSX5 RTOS and associated schedulability analysis tools (originally called the "Time Compiler", later "Real-Time Architect (RTA)"). The schedulability analysis tools implemented Response Time Analysis as introduced in [1], [2], [3], [4], and [5]. The SSX5 RTOS was developed precisely to meet the assumptions of this analysis. The execution time overheads were minimised and made constant, independent of the number of tasks, allowing them to be accurately measured and this data used in the schedulability analysis. The memory overheads of applications built on SSX5 were radically reduced by comparison with other automotive RTOS. This was achieved via the use of single-stack execution and compile time, i.e. off-line, configuration of the RTOS data structures to minimise RAM usage.

NRTA attracted significant venture capital funding in 1998 (£1 million from 3i) and again in 2000 (£9.2 million from 3i and TecCapital). In 2001, the company changed its name to LiveDevices Ltd.

In March 2003 LiveDevices was sold to ETAS GmbH, a wholly owned subsidiary of Robert Bosch GmbH. The reason for the trade sale was that Robert Bosch had benchmarked RTA-OSEK and found it to be significantly more efficient than its subsidiary's Ercos RTOS. Rather than attempt to write a new OSEK RTOS from scratch and compete with LiveDevices, ETAS chose to buy the company, bringing the RTA-OSEK technology and the 20+ LiveDevices engineering team in-house.

Standards

During the development of the SSX5 RTOS, the automotive industry was working on standards via the OSEK organisation. As a Technical Committee Member of OSEK [7], NRTA influenced the OSEK OS standard [8] ensuring that the basic conformance classes (BCCx) could be achieved with a single-stack RTOS, leveraging the execution time and memory savings which that approach facilitates [12]. NRTA modified the SSX5 RTOS to comply with the OSEK standard, in the process renaming the product RTA-OSEK.

Subsequently, ETAS, as a premium partner [9] of the AUTOSAR (AUTomotive Open System ARchitecture) partnership, have been heavily involved in specifying the AUTOSAR operating system standard [10], which extends the OSEK operating system standard. ETAS derived an AUTOSAR compliant RTOS called RTA-OS from RTA-OSEK [6]. (Note in [6] RTA-OSEK 'Planner' is the new name for the schedulability analysis tools, while 'Builder' is the name for the off-line configuration tool).

Impact

ETAS currently sell two versions of the RTOS, RTA-OSEK and RTA-OS compliant with the OSEK (***O****ffene* ***S****ysteme und deren Schnittstellen für die* ***E****lektronik in* ***K****raftfahrzeugen*; in English: "Open Systems and their Interfaces for the Electronics in Motor Vehicles") and AUTOSAR (AUTomotive Open System ARchitecture) operating system standards respectively.

The RTOS is currently available for more than 50 different ECU microcontrollers [6] including: Renesas: V850E, SH2, SH2A, H8S, H8SX, M16C; Xilinx Microblaze, PPC405 Core; Texas Instruments TMS470P, TMS570P; Infineon Tricore TC17x6, C166, XC2000; Freescale Star12, MPC555, MPC55xx, S12X, MPC56x, HC12X16, HC08, HCS12; Fujitsu 16LX; Analog Devices Blackfin, STMicroelectronics ST30, ST7, ST10.

ETAS customers for the RTOS cover a wide range of application areas within Automotive Electronics. It has been standardised upon (used by default in all ECUs) by many of the world's leading automotive powertrain systems and chassis electronics suppliers, and is used in cars produced by nearly all of the world's major car manufacturers. By 2015, the RTOS had been deployed in over 1 billion ECUs. This number is increasing at the quite astonishing rate of between 1 and 2 million new ECUs *per week*.

Beneficiaries

Use of the RTOS and its associated schedulability analysis tools has benefitted automotive manufacturers and their Tier 1 suppliers in the following ways: (i) A reduced memory footprint means that cheaper microcontroller variants with smaller on-chip RAM / Flash memory can be used. (The code size of RTA-OS is typically in the range 1 Kbytes to 1.5 Kbytes depending on the processor – making it the world's smallest AUTOSAR OS – so small that the hexadecimal machine code can be reproduced in the graphic below). This has reduced unit costs in production. (ii) The very low execution time overheads of the RTOS mean that more functionality can be included on a given low cost microprocessor reducing costs by avoiding the need for hardware upgrades to more capable but expensive devices. (iii) A reduction in the time spent debugging intermittent timing issues. Schedulability analysis and appropriate use of proven real-time mechanisms have enabled off-line analysis of task response times, reducing system integration time and testing effort, and improving reliability. For these reasons the world's major ECU suppliers and car manufacturers have adopted this technology. In a competitive market, some of these benefits will have been passed on to their customers in the form of cheaper, more reliable vehicles.

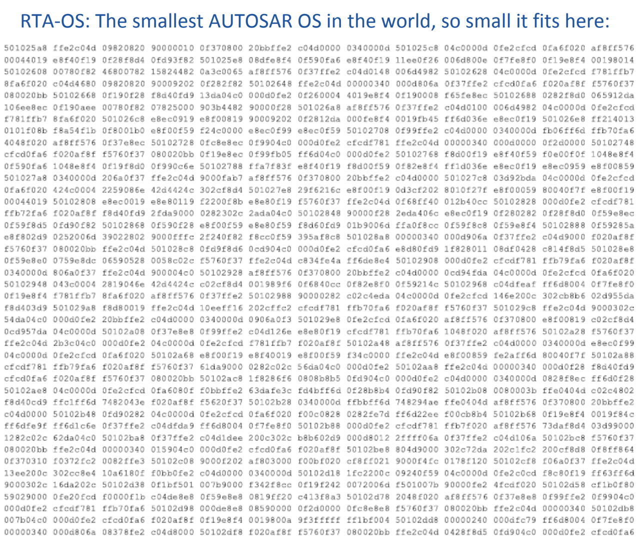

The Automotive Electronics market is both huge and highly competitive, with electronics now contributing 15-30% of overall vehicle production costs. For the reasons given above, the world's leading Automotive OEMs and Tier-1 suppliers have adopted the RTA-OSEK and RTA-OS operating systems. They have done so for the substantial benefits it brings to them and to their customers.

The technology has led directly to the creation and sustaining, over a period of more than 15 years of a large number of high technology jobs in York. The fact that ETAS has offices in York is a direct consequence of the research conducted at the University of York. (ETAS is head-quartered in Germany and has offices in 12 other countries).

Future Challenges

Automotive systems are now moving towards implementations on multicore hardware. The availability of such high performance multiprocessors means that it is now cost effective to integrate different applications that would otherwise have run on independent ECUs onto the same hardware platform. These different applications have different criticality levels which leads to a host of interesting problems. Mixed Criticality Systems [13] are currently a hot topic in real-time systems research. The Real-Time Systems Research Group is a world leader in this area with a number of projects funded by both the UK EPSRC and the EU. The group continues to have close ties with ETAS who act as industrial advisors on these projects as well as sponsoring research students.

References

[1] N.C. Audsley, A. Burns, M. Richardson, K.W. Tindell, A.J. Wellings, "Applying New Scheduling Theory to Static Priority Pre-emptive Scheduling". Software Engineering Journal, Volume 8, Issue 5, pages 284-292, 1993. ISSN :0268-6961

[2] N.C. Audsley, A. Burns, A.J. Wellings, "Deadline Monotonic Scheduling: Theory and Application". Control Engineering Practice, Volume 1, No. 1, pages 71-78, 1993. DOI:10.1016/0967-0661(93)92105-D.

[3] A. Burns, A. J. Wellings, "Engineering a Hard Real-Time System: From Theory to Practice" Software Practice and Experience, Volume 25, Issue 7, pages 705-726, July 1995. DOI: 10.1002/spe.4380250702.

[4] N.C. Audsley, K.W. Tindell, A. Burns, "The End Of The Line For Static Cyclic Scheduling?" In proceedings, 5th Euromicro Workshop on Real-Time Systems, pages 36-41, June 1993. DOI: 10.1109/EMWRT.1993.639042

[5] K.W. Tindell, A. Burns, A.J. Wellings, "An extendible approach for analyzing fixed priority hard real-time tasks". Real-Time Systems, Volume 6, Issue 2, pages 133-151, 1994. DOI: 10.1007/BF01088593.

[6] http://www.etas.com/en/products/rta_software_products.php, http://www.etas.com/en/products/rta_osek.php, http://www.etas.com/en/products/rta_os.php

[7] http://portal.osek-vdx.org/index.php?option=com_content&task=view&id=8&Itemid=1 1

[8] http://portal.osek-vdx.org/files/pdf/specs/os223.pdf

[9] http://www.autosar.org/index.php?p=2&up=1&uup=2&uuup=0

[10] http://www.autosar.org/download/R4.1/AUTOSAR_SWS_OS.pdf

[11] T.P. Baker, "Stack-based Scheduling of Real-Time Processes." Real-Time Systems Journal (3)1, pp. 67-100. 1991.

[12] R.I. Davis, N. Merriam, N.J. Tracey, "How Embedded Applications Using an RTOS can stay within On-chip Memory Limits". In proceedings of the Work in Progress and Industrial Experience Sessions, 12th EuroMicro Conference on Real-Time Systems. Stockholm, Sweden. June 2000.

[13] A. Burns, R.I.Davis, "Mixed Criticality Systems: A Review", 2014 available from http://www-users.cs.york.ac.uk/~burns/review.pdf