THE UNIVERSITY of York

Transferring Real-Time Systems Research into Industrial Practice Four Impact Case Studies

Robert I. Davis[1], Iain Bate[1], Guillem Bernat[2], Ian Broster[2], Alan Burns[1], Antoine Colin[2], Stuart Hutchesson[3], and Nigel Tracey[4]

[1]University of York, [2]Rapita Systems Ltd., [3]Rolls-Royce PLC, [4]ETAS Ltd.

This talk is different!

- Most Presentations:
  - Are about research done in the previous year
  - Give technical details about real-time systems research and its evaluation
  - Look forward to the research having some impact in the future
- This Presentation:
  - Is about research done in the previous century
  - Explains how real-time systems research was transferred into industrial use
  - Looks back at the impact of the technology over the last 20 years
  - Discusses some key success factors and roadblocks along the way

Four Impact Case Studies

- Where real-time systems research has been successfully transferred into industrial practice

- 1. Volcano: Guaranteeing the real-time performance of in-vehicle networks
- 2. RTA-OSEK and RTA-OS: Automotive real-time operating systems

- 3. RapiTime: A tool suite for analysing the timing behaviour of real-time software
- 4. Visual FPS: The first CAA certified use of a fixed priority scheduler in a high criticality avionics system

Science and Engineering

- Edward Lee - RTSS 2017 Award speech: A Personal View of Real-Time Computing:
  "An engineer seeks a physical system to match a model, whereas a scientist seeks a model to match a physical system."

- Impact case studies each involved elements of both science and engineering
  - Science – derivation of models and analysis for (idealised) real-time systems
  - Engineering – development of middleware enabling systems to be built that closely matched the assumptions of the models
  - Further science – to refine the models and analysis to match the detailed behaviour of the engineered systems

Let’s go back in time…from today …to the mid 1990s

Impact Case Study 1: Volcano: Guaranteeing the real-time performance of in-vehicle networks

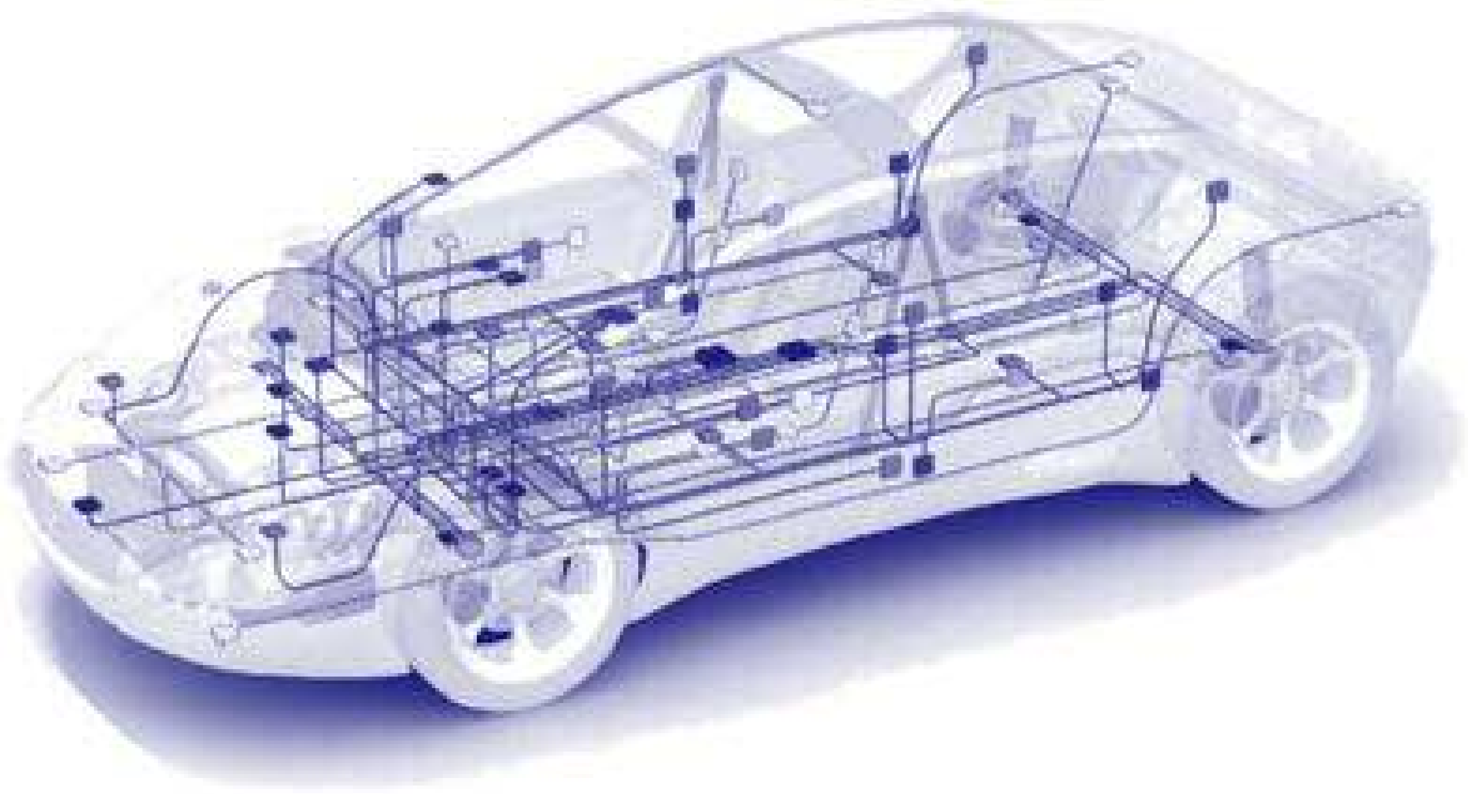

Impact Case Study 1: Volcano: Guaranteeing the real-time performance of in-vehicle networks

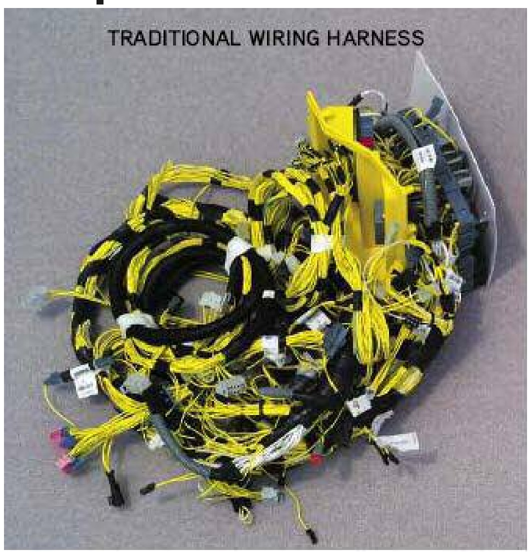

- Background: Early 1990s cars used point-to-point wiring
  - A typical luxury car had:
    - > 1000m of copper wire (30Kg)
    - > 300 connectors, 2000 terminals, 1500 wires
  - Expensive to manufacture, install and maintain
  - Unreliable due to very large number of connections
- Controller Area Network (CAN)
  - Simple, robust, reliable in-vehicle digital communications network
  - Small extra cost of CAN controllers and transceivers offset by massive reduction in wiring costs
  - Signals packed into messages which are broadcast on the network connecting ECUs
  - End-to-end deadline on signals lead to real-time constraints on message transmission (5ms to 1 sec)

Underpinning Research

- Schedulability analysis for CAN
  - Calculates the longest time that each message can take before it is transmitted over Controller Area Network (CAN)
  - Can be used to prove if all messages are guaranteed to meet their deadlines
  - Systematic approach was much better than testing and hoping the worst-case has been seen

- Analysis
  - Message Length

$$C_m = \left(g + 8s_m + 13 + \left\lfloor \frac{g + 8s_m - 1}{4} \right\rfloor \right) \tau_{bit}$$

  - Queuing delay

$$w_m^{n+1} = B^{MAX} + \sum_{\forall k \in hp(m)} \left\lceil \frac{w_m^n + J_k + \tau_{bit}}{T_k} \right\rceil C_k$$

  - Response time

$$R_m = J_m + w_m + C_m \leq D_m$$

Start-up Company #1: Northern Real-Time Technologies Ltd

- Origins
  - Start-up company NRTT founded in 1995 to develop "Volcano" technology for Volvo Car Corporation
- Objectives for Volcano
  - Ensure that systems built using the technology could be analysed using network schedulability analysis tools
  - Achieve very low execution time overheads and memory footprint for the on-target software
  - Support reconfiguration of signal to message mapping and message IDs post production
- Products developed
  - **Volcano Target Package:** API software, CAN device drivers, and configuration tools
  - **Volcano Network Architect:** Network schedulability analysis tools (in conjunction with Swedish company Kimble AB)

Advantages of using Volcano

- Guaranteed real-time network performance
  - Reduces the time and cost spent testing
  - Eliminates intermittent timing faults on the network reducing warranty costs and no fault found replacement of ECUs
- High network utilisation
  - Possible to configure networks to use 70-80% of the bandwidth compared with circa 30% with ad-hoc methods reliant on testing
  - Enables more ECUs to be connected to the same network thus supporting more functionality at lower cost and with higher reliability
- Post production re-configuration
  - Changing signal to message mappings and message IDs enables upgrades and lucrative 'software-only' options

Volcano in production

- First used in Volvo S80 in 1997

- Subsequently in Volvo XC90, S80, S/V/XC70, S60, S40, and V50

Impact Case Study 2: RTA-OSEK and RTA-OS: Automotive real-time operating systems

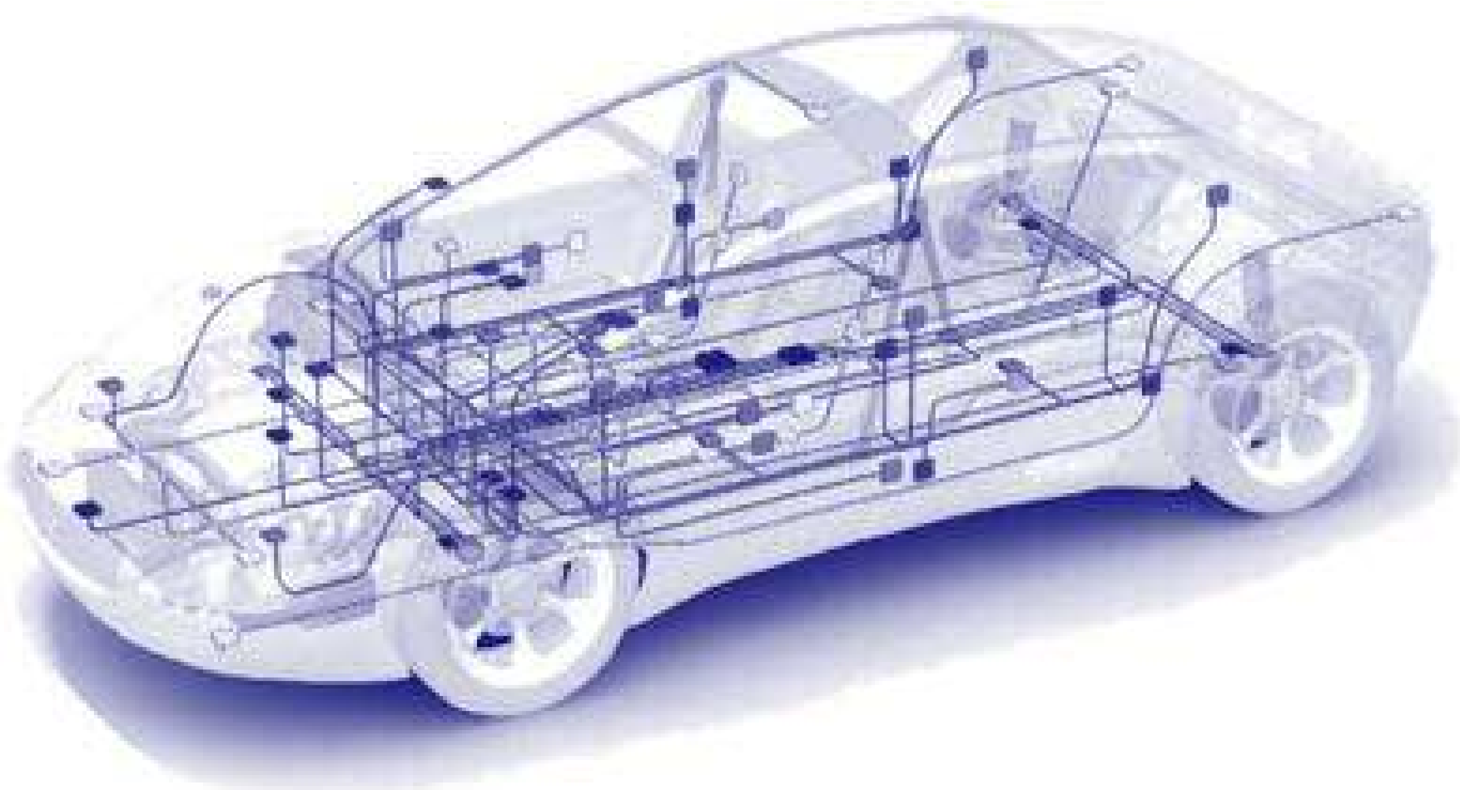

Impact Case Study 2: RTA-OSEK and RTA-OS: Automotive real-time operating systems

- Background: Automotive Electronics circa late 1990s
  - 15-25 ECUs connected via two or more communications networks (CAN)
  - Relatively simple low cost microprocessors (single-core)
- System functionality
  - Multiple software tasks running on each ECU
- Real-Time Operating System (RTOS)
  - Needed to schedule when each task could run so that all tasks meet their timing constraints
  - Essential otherwise the system could suffer intermittent timing faults and poor reliability
  - RTOS of the time were arguably not fit for purpose – large memory footprints, high overheads, and didn't meet assumptions of theory (e.g. issues with priority inversion)

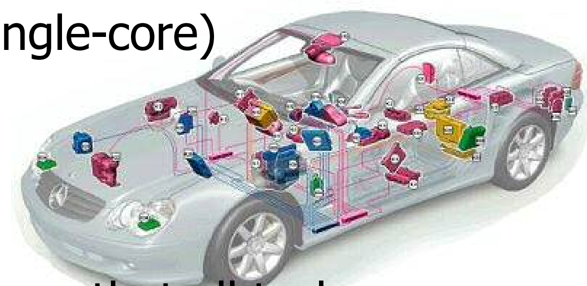

Underpinning Research

- Schedulability Analysis for Processors
  - Response Time Analysis for Fixed Priority Scheduling
  - Accounts for resource sharing, non-preemptive execution, periodic/sporadic arrivals, deadlines prior to completion, and other aspects needed for tasks in automotive systems
  - Accounts for the overheads of a well designed RTOS

$$w_{i,q}^{m+1} = B_i + (q+1)C_i + \sum_{\forall j \in hp(i)} \left\lceil \frac{w_{i,q}^{m} + J_j}{T_j} \right\rceil C_j$$

$$R_i = \max_{\forall q = 0,1,2 \ldots Q_i - 1} (w_{i,q} - qT_i + J_i)$$

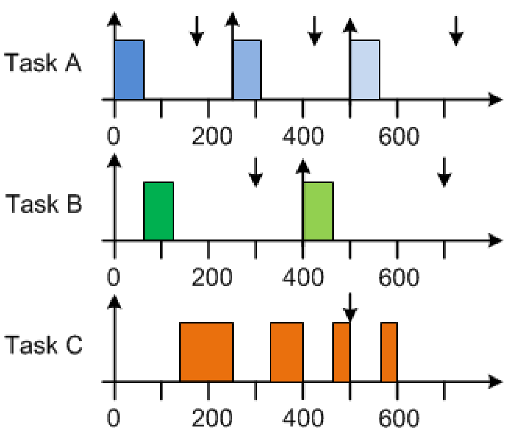

Start-up Company #2: Northern Real-Time Applications Ltd

- Origins
  - Start-up company NRTA founded in 1997 specifically to develop a RTOS for automotive applications
- Objectives for the RTOS
  - Ensure that systems built using the RTOS could be analysed using schedulability analysis tools
  - Execution time overheads and memory footprint must be much smaller than any other automotive RTOS
  - Sell the RTOS to **many** car manufacturers and suppliers
- Products developed
  - **Real-Time Architect** schedulability analysis tools
  - **RTA-OSEK** real-time operating system

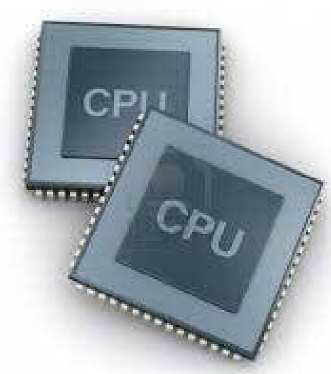

Advantages of using the RTOS

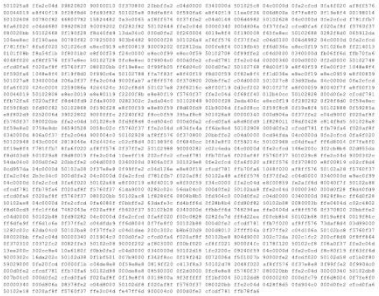

- Low memory footprint
  - Enables use of cheaper microprocessor variants which reduce unit costs in production
  - 1 KBytes to 1.5 Kbytes (data publically available for all variants)
- Low and bounded execution time overheads
  - Allows more useful functionality to be added without the need to upgrade to more expensive processors (Data publically available for all variants)
- Analysable behaviour
  - Guaranteed timing behaviour leads to more reliable systems
  - Reduces time spent debugging intermittent timing problems

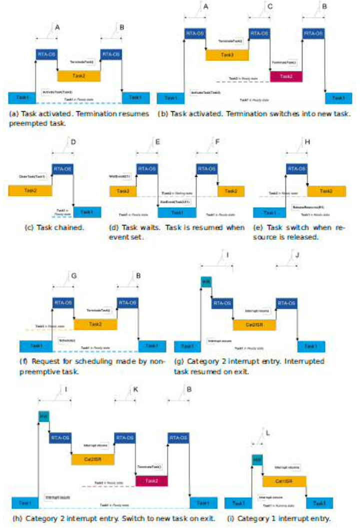

Start-up Company #2: Northern Real-Time Applications Ltd*

- Benchmarking
  - Ford benchmarked the RTOS and found it to be much more efficient than 10 other competitors
  - ETAS (a subsidiary of Bosch) also benchmarked the RTOS against their in-house offering and found it was much more efficient
- Trade sale
  - Faced with the option to start from scratch and build a new RTOS or buy the company, ETAS bought the company in 2003
  - ETAS adapted the operating system to meet the AUTOSAR standard (RTA-OS)

*Note in 2001, Northern Real-Time Applications changed its name to LiveDevices as it was also exploring products in the IoT domain

Success Factors: Common to the impact case studies

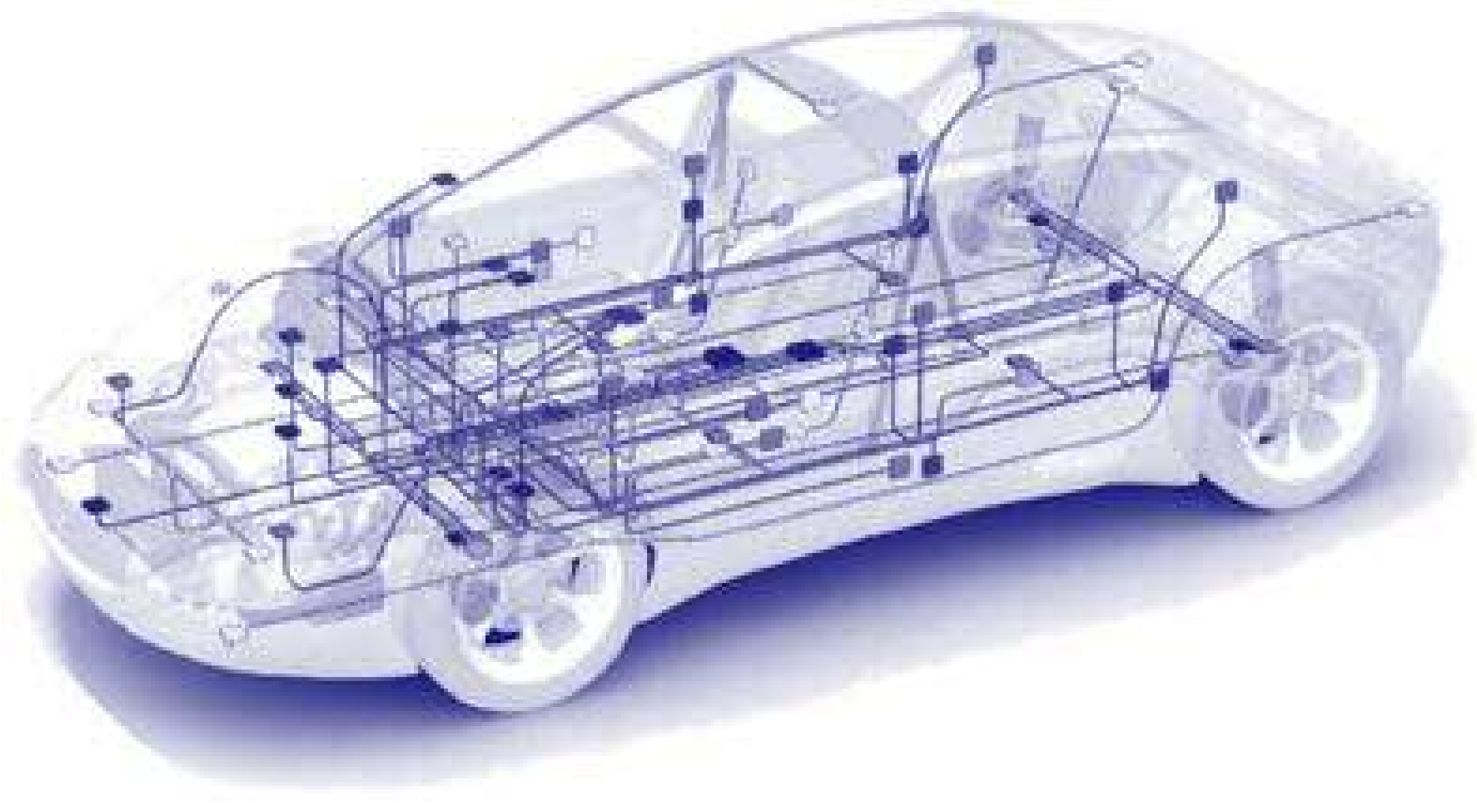

Success Factors

- 1. Having an idea and then a product that made a step change for customers - providing a return on their investment
  - Volcano: increased network utilisation from 30 to 80% with improved reliability – reduced development, production and warranty costs
  - RTOS: reduced memory footprint and overheads result in lower production costs. Improved reliability gave lower warranty costs

- 2. Core team of excellent people
  - Typically the founders of the company and the first few employees who worked very hard over long periods of time (years) to ensure the company was a success

THE UNIVERSITY of York

Success Factors

- 3. A product that was not easy to replicate – barrier to competition
  - Important in obtaining funding and getting a foothold in the market
  - Very evident with the RTOS since the company was bought by one of its competitors
- 4. A high product quality and outstanding customer support
  - When a company is small and has only been around for a year or two it needs to build an excellent reputation
  - Quality is absolutely essential - make or break in terms of winning the trust of major companies who are considering adopting the technology

Success Factors

- 5. Balanced team of people
  - Not just the technologists and software engineers, but also field application engineers and support staff who can do an exceptional job in handling customer issues
  - Marketing and sales staff who understand the technology and can therefore talk effectively to both engineers and managers at customer sites
- 6. Previous experience
  - Having someone on board who has previous experience in a successful start-up company in the same field is hugely advantageous
  - They will understand what is needed to grow a company successfully and help avoid all manner of pitfalls

Success Factors

- 7. Attracting an acquisition
  - In each case, trade sale of the company / technology led to a large acceleration in the rate of adoption
  - Structuring the company not only for standalone success but also for acquisition was an important success factor

Major Roadblocks

- 1. Funding the initial development from academic ideas and prototypes to saleable product
  - A high quality industry ready product is far removed from typical academic prototypes
  - Effort is needed when the company first starts and has few sales
  - Self-funding can work if the founders can afford not to be paid for a while, and or can get one or two early contracts
  - Business angel or venture capital funding is effective but comes at a cost of giving up some proportion of the shares in the company
  - Assistance from the host institution in terms of providing time to cover initial development efforts can be greatly beneficial

Major Roadblocks

- 2. Finding the right sales staff
  - It proved remarkably difficult to find people who were both good at sales and understood the technology
  - In each company, sales were led by someone with a strong technical background who had the right personality and turned themselves into an excellent salesman via appropriate training
  - Bringing in ‘high flying’ sales staff without a strong technology background can be an expensive mistake!

Major Roadblocks

- 3. Convincing major companies to fully adopt a new technology
  - Problematic due to the conservative approach often taken to purchasing from small companies
  - Concerns:
    - Will the company be around in a year's time?
    - Can it handle the volume of support that may be needed?
    - Is the product really of a high enough quality to rely upon?
  - The main factors in addressing this were product and customer service quality, and time – it becomes easier to make these larger sales once the company has been established for a few years
  - In each case, a much higher level of sales was achieved once the company was acquired by a larger organisation (trade sale)

Back to today…
Volcano

- Volcano Technology now owned and marketed by Mentor Graphics
- Available for more than 30 different microprocessors used in automotive systems
  - Fujitsu 16LX, FR Series; Hitachi H8S, SH7055, SH7058; Infineon C16x, TC179x, TC176x, XC800, XC2000; Renesas M16C, R32C/M32C; Freescale HC08, HC12, MC683xx, MPC5xx, MAC71xx; S12, S12X, MPC55xx, MPC 56xx; Mitsubishi M32R, MC32C; PowerPC; National CR16; NEC V85x, 78K0; ST Microelectronics ST9, ST10; Texas Instruments TMS470; Toshiba TMP92/TMP94.
- Used by Volvo in more than 5 million cars since 1997
- Also used by Ford, Jaguar, Land Rover, Aston Martin, Mazda, SAIC…

THE UNIVERSITY of York

Back to today…
RTA-OSEK / RTA-OS

- RTA-OSEK / RTA-OS part of ETAS' product line
- Now available for over 50 different microprocessor families 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
- Standardized upon by many of the world's leading automotive suppliers
- Used by almost all of the world's leading car companies

Barcelona: Camp Nou

- At current production rates how long to produce one ECU containing the RTOS, for every one of the 100,000 seats in the stadium?
  Answer: 10 hours (10,000 units per hour)
- If we put all the ECUs containing the RTOS ever produced evenly on all 100,000 seats, how high would the pile of boxes be on every seat? Answer: 1.25 Km

Photo by Johannes Falke - Own work, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=21538527

An apples vs. oranges comparison

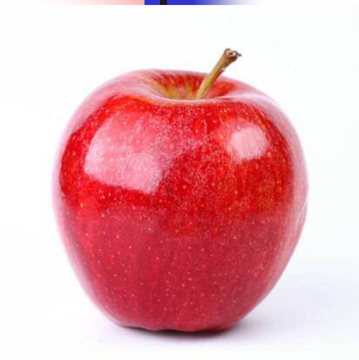

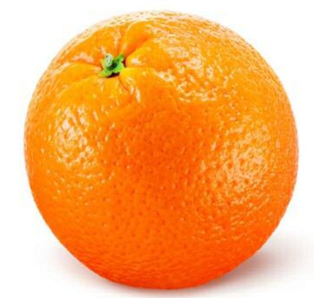

- Total worldwide production volumes to 2017
- Apple iPhones (all models) 1.16 billion
  (Source: https://www.lifewire.com/how-many-iphones-have-been-sold-1999500)
- ECUs containing RTA-OS/RTA-OSEK 1.25 billion
  (Source: ETAS UK)

Note iPhone production rate is higher so will overtake soon

iPhones

RTA-OS
RTA-OSEK
ECUs

Summary: One common thread

- Excellent research
  - Many of the underpinning research papers were later recognised as seminal ones in the field (some cited over 500 times)
- Exploitation via a start-up company
  - Huge amount of hard work over many years by those involved
  - …and a reasonable amount of luck!
- World-wide impact
  - Products have been adopted and standardised upon by many large companies in the automotive and aerospace industries
  - Huge benefits to society from more efficient, more reliable, lower cost vehicles

Take home message

- Real-Time Systems research matters and has a huge impact (even if it often goes unseen)
  - If you have some excellent ideas and the commitment and persistence to see them through then you could make a real success of transferring your research to industry
  - Starting-up a company is not difficult
  - It could lead wide-ranging impact benefitting many people

  - So what are you waiting for?

    …maybe this year could be the start of a great opportunity

Questions?

- There’s a lot more info in the paper