



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/167644/>

Version: Accepted Version

Proceedings Paper:

Akesson, Benny, Nasri, Mitra, Nelissen, Geoffrey et al. (2020) An Empirical Survey-based Study into Industry Practice in Real-time Systems. In: 2020 IEEE Real-Time Systems Symposium (Proceedings). 2020 IEEE Real-Time Systems Symposium, 01-04 Dec 2020

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

An Empirical Survey-based Study into Industry Practice in Real-time Systems

Benny Akesson¹, Mitra Nasri², Geoffrey Nelissen², Sebastian Altmeyer³, Robert I. Davis⁴

¹ESI (TNO) & University of Amsterdam, Netherlands.

²Eindhoven University of Technology, Netherlands.

³University of Augsburg, Germany.

⁴University of York, United Kingdom.

Abstract—This paper presents results and observations from a survey of 120 industry practitioners in the field of real-time embedded systems. The survey provides insights into the characteristics of the systems being developed today and identifies important trends for the future. The survey aims to inform both academics and practitioners, helping to avoid divergence between industry practice and fundamental academic research.

Index Terms—Real-time systems, Survey, Industry Practice.

I. INTRODUCTION

The real-time embedded systems field covers a broad range of systems, from simple control loops on micro-controllers to complex interconnected distributed systems. These systems span many different application domains, including avionics, automotive, consumer electronics, industrial automation, and medical systems, each with its own requirements, standards, and practices. This diversity makes industrial real-time systems and their associated design methods difficult to characterize.

Some fields, such as software engineering [1]–[3] and systems engineering [4], have a history of systematically researching industry practice using surveys, interviews, and literature reviews [5]. This provides a view of the perceived relevance, benefits, and drawbacks of different technologies and methods; identifies trends and opportunities for future research; and tracks the adoption of existing research results. By contrast, there is no tradition of empirical studies into industry practice in the real-time systems field. This omission contributes to a gap, and potentially a divergence, between industry practice and academic research. This paper addresses that gap via an empirical survey-based study into industry practice. The five objectives of the study were to:

- O1** Establish whether timing predictability is of concern to the real-time embedded systems industry,
- O2** Identify relevant industrial problem contexts, including hardware platforms, middleware, and software,
- O3** Determine which methods and tools are used to achieve timing predictability,
- O4** Establish which techniques and tools are used to satisfy real-time requirements,
- O5** Determine trends for future real-time systems.

A survey targeting industry practitioners in the area of real-time embedded systems was developed and distributed. The survey comprised 32 questions related to the five objectives.

The two main contributions of this survey are:

- 1) Insights into the characteristics of real-time systems based on responses from 120 industry practitioners from a variety of organizations, countries, and application domains.

- 2) Discovery of statistically significant differences between the three largest application domains: avionics, automotive, and consumer electronics.

The remainder of the paper is organized as follows: Section II outlines related work. Section III describes the methodology used and the design of the survey. The survey questions and results are elaborated in Section IV, along with key observations and a discussion of statistically significant differences between domains. Section V concludes.

II. RELATED WORK

Research methods used for understanding industry practice can be broadly divided into three categories:

- 1) *Survey-based research* targeting industry practitioners in one or more application domains [6]–[13].
- 2) *Interviews* with industry practitioners, either open or based on a framework of questions, with the answers subsequently analyzed [14,15].
- 3) *Literature surveys reviewing case studies* with the goal of categorizing industry experiences [5].

Some works combine both survey-based research and interviews, exploiting their complementary nature to improve overall quality [8,11]. There are also replication studies investigating how results generalize to other populations [1,12].

In contrast to the fields of software and systems engineering, there has been little if any research undertaken into industry practice in the real-time systems field. Instead, the academic community tends to look inwards, surveying and classifying its own work rather than studying industry practice, contexts, and needs. Examples of well-known literature review surveys include those on uniprocessor scheduling [16,17], multiprocessor scheduling [18], limited preemptive scheduling [19], mixed-criticality scheduling [20], resource allocation and mapping [21], timing analysis [22], and multi-core timing analysis [23]. Recent literature surveys [23]–[25] and other initiatives [26,27] take this one step further and include diagrams illustrating how the number of publications on different research topics has varied over time. This allows hot-topic areas to be identified. While these works may be useful to identify trends in academic real-time systems research, these trends may not be reflected in industry. In conclusion, there is no existing work that systematically surveys industry practice in the real-time systems field. The aim of this paper is to address that omission and to help close the gap between industry practice and academic research.

III. METHODOLOGY

The study described in this paper has five objectives **O1** to **O5** (listed in Section I) that focus on industry practice. To meet these objectives, we chose, as the research method, a *survey* asking industry practitioners a set of predetermined questions. We found no existing surveys in the relevant area and it was therefore necessary to develop a new survey instrument. In designing the survey, we first considered appropriate validity criteria and their relevant threats, before making design choices to mitigate those threats. Due to space limitations, threats to validity and their mitigation are discussed in an extended version of this survey published as a technical report [28]. Note, in the description below, we follow the structure, classifications, and terminology proposed by Kitchenham et al. [3] for survey-based research.

Survey design: The survey was designed as a *cross-sectional study*, i.e. a snapshot taken over a particular period of time, December 2019 to April 2020 in this case. We used a *self-administered questionnaire* on SurveyMonkey [29] that could be answered without the need for intervention, and without providing respondent or company identification. Further, we did not automatically collect any data relating to respondents identities (IP addresses etc.), thus preserving anonymity. The aim here was to reduce the risk of self-exclusion by enabling those who work on confidential projects to still answer the survey questions. As an additional guarantee of anonymity, we promised respondents to only release summarized and aggregated results [30].

Survey questions: We focused on *closed questions*, where respondents are asked to select one or more answers from a list of predefined options. Closed questions are typically faster to answer and easier to analyze than open ones, thus they reduce the likelihood that a participant abandons the survey before reaching the end. The drawback of closed questions is that they limit the range of possible responses.

The questions were carefully formulated to be neutrally worded, as precise as possible, and avoid unnecessary jargon. Key terms were explained where necessary, helping to ensure that the questions and the list of predefined options were as unambiguous and easy to comprehend as possible. We did not, however, manage to completely eliminate ambiguity caused by jargon. For example, the term “distributed system” used in Questions 10 and 11 may have been interpreted in two different ways. (This is discussed further in Section IV).

With closed questions, care was taken to ensure that the predefined options were unbiased, mutually exclusive, and as exhaustive as possible. We included “Other”, where appropriate, to give respondents the opportunity to go beyond the predefined options when necessary. Where appropriate, we also allowed multiple options to be selected to prevent arbitrary choices between equally valid answers. This was particularly important in the context of real-time systems comprising different sub-systems to which different answers could apply. Both of these techniques are listed as best-practices [3]. The category “I do not know” was added to a number of questions where we did not expect all respondents to be able to give an answer, despite belonging to the target

population. The “I do not know” category has the benefit of making the lack of knowledge explicit, as distinct from skipped questions and arbitrary answers.

We had a choice of asking respondents to consider either one or several real-time systems that were being developed in their organization. We chose the former, since although this approach gathers data about fewer systems, it enables conclusions to be drawn about individual systems, and answers to different questions to be related to one another.

Instructions for participants: The welcome page of the survey explained the purpose of the survey, defined the target population, and suggested that it would take 10-15 minutes to answer the 32 questions. It also explained that the survey was anonymous and that the output of the survey would be an academic paper. The incentive for the respondents was the opportunity to shape future research in the area of real-time systems, and align it with industry practice and needs.

Survey validation: A draft of the survey was validated by a test group, consisting of 13 people from industry, research institutes, and universities, representing domain experts from the target population and users of the survey results. The test group filled out a draft version of the survey and provided feedback. Based on this feedback, we improved the instructions, questions, and predefined options.

Sampling method: Since there is no list that identifies all industry practitioners who work on real-time embedded systems, it was not possible to perform a *random sampling* to invite the participants. Further, the target population is highly specific and has limited availability, which prevents the use of probabilistic sampling methods. Instead, we used a combination of *convenience sampling* and *snowball sampling*. Convenience sampling means that we reached out to the target population via the authors’ combined networks using emails and personal messages on LinkedIn [31]. We sent them a personalized invitation, followed by a reminder a few weeks later, as suggested in [3].

To increase the reach of the survey beyond the authors’ networks, we applied snowball sampling in two different ways. First, by encouraging those who we invited to take the survey to forward it to other practitioners. However, we instructed them to only forward the invitation to people working on different real-time systems. Second, we used snowball sampling to mitigate a geographical bias towards contacts from the region where most of the authors’ networks reside. We asked 20 academics, primarily based in other regions to forward the invitation to members of the target population in their networks. In total, 120 respondents completed the survey, 90 via the authors’ network, and 30 via the authors’ academic contacts on other continents. Due to snowball sampling, we do not know the exact response rate of the survey.

IV. RESULTS

This section lists all of the survey questions, in the order in which they appeared, along with graphs of the results, and our observations. The survey was divided into a number of topics, which are separated by horizontal lines in the text below.


Where results are given as percentages, unless otherwise stated, these correspond to the proportion of respondents who

selected that specific option out of all of the respondents who answered that particular question. The graphs presenting the results are color-coded. Red bars are used for questions with distinct alternatives, and hence the total sums to 100%. Blue bars are used where respondents were asked to “select all options that apply”, and hence the percentages sum to more than 100%. Multi-colored bars (e.g. Question 6) indicate the percentage of respondents who selected the corresponding scores or rankings. Where the answers have an ordering (e.g. Question 1) then the results are presented in that order, otherwise they have been re-ordered with the most popular answer first for ease of reference, nevertheless, “I do not know” and “Other” are always placed last.

Our observations include a commentary on the results, and a more in-depth look at the data. In some cases, we comment on the results for sub-groups that have been identified via the answers to Question 4 (Avionics, Automotive, and Consumer Electronics) and Question 5 (safety-critical components, and no safety-critical components). We *only* comment on the difference in results between these sub-groups where these differences are statistically significant at the $p < 0.05$ level¹. Finally, the number X of respondents answering each question is given as ($n=X$) at the right hand side of each question box.


Demographics: This part of the survey asked questions about the respondent’s organization and professional experience.

Question 1: How many employees does your organization have? ($n=120$)



Observations: Approximately two thirds of the respondents were from large companies (> 1000 employees), with around one third from small and medium sized enterprises (SMEs).


Question 2: Which position best describes your current role in your organization? ($n=120$)



Observations: Approximately 60% of the respondents were directly involved in system development (software, system, or hardware), while approximately 27% were involved in industrial research. Note that the category “Academic Researcher” includes staff on secondment to industry.

¹According to the analysis provided by the SurveyMonkey toolset, see https://help.surveymonkey.com/articles/en_US/kb/Significant-Differences.


Question 3: How many years industrial experience do you have? ($n=120$)



Observations: The majority of respondents had many years of industrial experience, with 41% having more than 10 years, and only 23% having five years or less.


System Context: This part of the survey asked questions about hardware, software and the execution of the system. Respondents were asked to think about a particular system where they were familiar with these aspects, and to consider the same system for all questions to ensure consistent responses.

Question 4: To what domain(s) does the considered system belong? ($n=107$)



Observations: The survey has broad coverage of the different application domains. Note, that multiple domains could be selected. The largest overlaps were between: Avionics and Defense 9.4%, Automotive and Industrial Automation 6.6%, Automotive and Consumer Electronics 5.7%, Automotive and Avionics 5.7%, and Space and Defense 4.7%. Automotive alone was indicated by 65% of those selecting that domain, similarly, Avionics alone by 60% of those selecting that domain, and Consumer Electronics alone by 56% of those selecting that domain. Of the 11 respondents who indicated “Other domain”, five specified “Telecomms.” i.e. 4.7%.


Question 5: Is (parts of) the considered system safety-critical? ($n=107$)



Observations: Even though the response to Question 4 indicates broad domain coverage, a large majority (75%) of the systems considered had some part that was safety critical.

Of those respondents who selected Avionics in Question 4, 100% answered “Yes” to this question, compared to 91% of those who selected Automotive, and just 52% of those who selected Consumer Electronics.

Question 6: Give a score to the importance of different system aspects for the considered system. ($n=107$)




Observations: Timing predictability, although viewed as less important than functional correctness, reliability and safety, was seen as more important than security, computing power, cost, and thermal considerations. (This is perhaps unsurprising, since the survey targeted those working on real-time systems).

Of the respondents who selected Avionics in Question 4, 87% thought that timing predictability was very important, compared to 48% of those who selected Automotive, and just 26% of those who selected Consumer Electronics. In contrast, unit cost of the execution platform was rated as very important by 45% of those respondents who selected Automotive in Question 4, 32% of those who selected Consumer Electronics, and just 7% of those who selected Avionics.

Hardware Platform: This part of the survey asked questions about the hardware and software configurations of the considered system. Here, most questions allowed the selection of multiple options to capture the characteristics of complex systems with many components. Respondents were asked to select all options that apply to the system they were considering.

Question 7: What Operating Systems are running on the considered system? (n=103)




Observations: While 78% of respondents indicated that some parts of their system use an RTOS or Micro kernel, a significant minority (37%) had parts that use no Operating System (OS) at all. RTOS alone was selected by 22.5%, Linux alone by 7.8%, and Bare metal alone by 4.9%. None of the respondents used Windows alone. There were many systems that used more than one OS (62.7%). The largest overlaps were between RTOS and Linux (42.2%), Bare metal and RTOS (28.4%), and Bare metal and Linux (17.6%). The combination of Bare metal, RTOS, and Linux was used by 14.7% of respondents.

Of the respondents who indicated in Question 5 that their system contained some safety-critical components, 87% used an RTOS. This figure reduced to 50% of those who indicated no safety-critical components. By contrast, the corresponding figures for the use of Windows were 3% and 25%, respectively.


As an optional addition to this question, respondents were asked to name the operating systems they were using. 32 respondents did so, many citing multiple operating systems. The following lists the operating systems named and the number of times they were cited: Autosar (8), QNX (8), VXWorks (4), OSEK (3), Redhat Linux (3), Free RTOS (2), Linux (2), PikeOS (2), Ubuntu (2), Yocto Linux (2), Arinc-653 (1), DEOS (1), EmbOS (1), Erika (1), Integrity (1), LynxOS (1), RTEMS (1), SafeRTOS (1), ThreadX (1), Windows (1), and Zephyr (1).

Question 8: Select the options that describe the processing hardware of the considered system. (n=103)




Observations: The majority of systems (81%) include multi-core components, while just under 40% include single core components. Similarly, 35% to 42% of systems include FPGAs, GPUs, and other hardware accelerators. Of the 10 respondents who indicated "Other", three specified DSP (i.e. 2.9%) and two specified System-on-Chip (i.e. 2.0%).

Question 9: Select the options that describe the memory hierarchy of the considered system. (n=103)




Observations: The majority of systems (over 63%) have elements of a complex memory hierarchy including mass storage devices, DRAM, and multiple levels of cache. Core local memory and single level caches are also prevalent.

Question 10: How many distributed nodes (e.g. ECUs) are there in the considered system? (n=102)



Observations: The majority of systems are distributed (73%), with only 17% identified as having a single node (ECU).

Question 11: Which of the following options describe the connectivity within the (distributed) system? (n=103)




Observations: Wireless networks were used in around 25% of systems, with Ethernet (64%) and CAN (41%) the most popular forms of wired network. Many systems (48%) used multiple types of wired network, with 34.3% using Ethernet and CAN, 27.5% Ethernet and Serial, 19.6% CAN and Serial, and 14.7% Ethernet, CAN, and Flexray. 9.8% of systems used Ethernet as the only wired network, while less than 3% used CAN, Flexray, or Serial alone. Wireless was used as the only network by 8.8% of respondents, about one third of those using that technology.

Of the respondents who selected Automotive in Question 4, 74% used CAN, and 34% used Flexray, this reduced to 21% and 3% for those who selected Avionics. Flexray was only used by one respondent (1%) who did not select Automotive.

There was some inconsistency in what was understood by one node or ECU in Question 10 (17%) and what was understood by “not distributed” (11%) in Question 11. This could be because respondents were considering “Nodes” or “ECUs” in Question 10, and “connectivity” in Question 11. (A single node or ECU may contain multiple connected processing units).


Timing Characteristics: This part of the survey asked questions about the timing characteristics of the considered system.

Question 12: Which of the following sentences are true about task activations in your system? (n=101)



Observations: While periodic activation is the most common at 82%, over 60% of systems included aperiodic activations. 22% of responses indicated highly predictable behaviors (utilizing either periodic or time triggered activation) with no sporadic or aperiodic tasks while 4% (and 2%) of respondents indicated to have purely sporadic (aperiodic) activations with no time-triggered or periodic tasks. Interestingly, 74% of respondents indicated at least two, and 25% all four types of activations.

Question 13: Which of the following timing constraints exist(s) in your system? (n=101)




Note, a more detailed explanation of the terms used was provided in the survey².

Observations: Given the scope of the survey, it is unsurprising that just under 90% of respondents indicated that their system had some form of timing constraints. Many systems (62%) had a combination of two or more different constraints: Hard and Firm 38%, Hard and Soft 36%, Firm and Soft 42%, and all three 27%. In contrast, far fewer systems had only one type of timing constraint: Hard 5%, Firm 10%, and Soft 15%.

Of the respondents who selected Avionics in Question 4, 79% indicated Hard constraints, compared to 56% of those who selected Automotive, and only 27% of those who selected Consumer Electronics. Of the respondents who indicating in Question 5 that their system contained some safety-critical components, 64% indicated Hard constraints. This reduced to 21% of those who indicated no safety-critical components.


Question 14: For the most time-critical functions in the system, roughly how frequently can the deadline of a function be missed without causing a system failure. (n=101)



Observations: A substantial number of respondents (35%) were unable to give a specific answer to this question, and answered “I do not know”. Only a small proportion (15%) of systems were considered strictly hard real-time, with deadlines that must never be missed. By contrast, 45% of respondents indicated that the most time critical functions in the system could miss some deadlines, and 20% indicated that deadline misses more often than 1 in 100 could be tolerated.


Question 15: What is the largest number of consecutive deadline misses that could be tolerated, assuming that such a blackout does not reoccur for a very long time. (n=101)

²Hard implies that violating the timing constraint is considered a failure of the system. Firm implies that violating the timing constraint is highly undesirable. Soft means that occasionally violating the timing constraint is acceptable, but negatively impacts the perceived quality of the system.




Observations: The responses to this question follow a similar pattern to those of Question 14, with 34% of respondents indicating that the system can tolerate black-out periods in the range of 1 to more than 10 deadline misses. Here, only about 60% of respondents were able to give a specific answer, with 40% answering “I do not know”.

Question 16: What are relevant timing constraints to your system? (n=99)



Observations: End-to-end response time was considered the most important timing constraint, with the largest percentage of “very important” scores and the highest average score of 4.3. However, task running time (3.78), response jitter (3.64) and activation jitter (3.42) also need to be considered. 72.7% of respondents rated end-to-end response time highest or equal highest. For task running time, response jitter, and activation jitter, this was the case for 45.5%, 35.4%, and 32.3% of respondents, respectively.

Question 17: How does the considered system react if tasks miss deadlines? (n=102)




Observations: The most common (45%) reaction to a missed deadline is to report the issue and continue, with 10% of systems also doing nothing. Other systems take actions on a deadline miss, including 30% rebooting, and 30% restarting tasks. Further, although 15% said that a deadline miss may never occur (Question 14) only 7% trust their system enough to state that “This case never happens”.

Of the respondents who indicated in Question 5 that their system contained some safety-critical components, 36% indicated “Reboots the system”. This reduced to just 8% of those

who indicated no safety-critical components. By contrast, the figures for “Does nothing” were 6% and 21%, respectively.


Managing Timing Behavior: This part of the survey asked questions about the methods used to analyze and influence the timing behavior of the system.

Question 18: Which methods are used for Worst-Case Execution Time (WCET) estimation in the considered system? (n=99)



Observations: Measurement-Based Timing Analysis (MBTA) tools are used by substantially more respondents than Static Timing Analysis (STA) tools, with more than 50% using in-house MBTA tools compared to 15% for in-house STA tools. This distinction is less stark when it comes to third party solutions with 34% using third party MBTA tools and 21% using third party STA tools. Overall, 67.4% of respondents used some form of MBTA tool, 33.7% used some form of STA tool, and 23.5% used both.

Question 19: What steps are taken to help increase timing predictability? (n=97)




Observations: While more than 50% of respondents use watchdog timers, static scheduling, and appropriate hardware selection, it is clear that there is no “silver bullet” to improving timing predictability. Each of the wide range of different techniques is used by at least 20% of respondents, and 46%

of respondents answered “Yes” to at least 5 of the techniques listed. Some of the techniques that are least frequently employed are, however, those that have the largest impact on average-case execution times (e.g. disabling caching, and turning off all but one core).

There was considerable uncertainty in answering parts of this question, reflected in approximately 30% of respondents answering “I don’t know” with respect to the use of scratchpads, cache locking, memory bandwidth regulation, and refactoring code into memory and computation phases.


Question 20: Which task scheduling policy/policies are used in the considered system? (n=97)



Observations: The most popular scheduling policies were fixed priority and static cyclic / table driven, with each used by more than half of the respondents. Round-robin and FIFO, which are not traditionally viewed as real-time policies, were employed in around 30% of systems, while EDF was employed in less than 17% of systems, less than one third as often as fixed priority scheduling.

Of the respondents who selected Automotive in Question 4, 27% used EDF scheduling, compared to just 3% of those who selected Avionics, and 11% of those who selected Consumer Electronics.

Question 21: Indicate the types of preemption that are supported in the considered system. (n=97)




Note, a more detailed explanation of the terms used was provided in the survey³.

Observations: While preemptive scheduling is the most popular choice, used in two thirds of systems, both non-preemptive and co-operative scheduling are used in more than one third of systems.

Question 22: Indicate how task migration can take place between different cores in the considered system. (n=98)


³Preemptive implies that task execution can be preempted by other tasks at any time, non-preemptive implies that task execution cannot be preempted by other tasks before completion, and cooperative means that task execution can be preempted by other tasks, but only at predefined preemption points.



Observations: Although timing predictability is typically easier to achieve without task migration, the proportion of systems permitting migration (37%) is similar to the proportion that do not permit it (40%).

Of the respondents who selected Avionics in Question 4, only 7% indicated that task migration is permitted while the task is executing. By comparison, this figure was 27% for those who selected Automotive, and 30% for those who selected Consumer Electronics.

Question 23: How do you ensure that the functions in the considered system respect their deadlines? (n=97)



Observations: Less than 10% of respondents are using commercial schedulability analysis tools, while more than 30% use in-house solutions. The main off-line approach is schedule correctness by construction, using a static schedule and checking that execution time budgets hold. However, the most common approach overall is to run tests and check for overruns (61%), with a similar proportion to those that use watchdogs timers / run-time monitors (see Question 19).


None of the respondents who selected Avionics in Question 4, indicated “no specific action undertaken”, compared to 12% of those who selected Automotive, and 16% of those who selected Consumer Electronics. Of the respondents who indicating in Question 5 that their system contained some safety-critical components, 7% answered “no specific action undertaken”. The corresponding figure was 29% of those who indicated no safety-critical components.

Timelines for Hardware Adoption: This part of the survey asked questions about timelines for hardware adoption.

Question 24: By which year did or do you expect development projects for real-time embedded systems in your department to begin using multi-core embedded processors (i.e. processors with 2 to 16 cores)? (n=97)


Question 25: By which year did or do you expect development projects for real-time embedded systems in your department to begin using heterogeneous multi-cores with different types of CPUs, GPUs, and other accelerators? (n=97)

Question 26: By which year did or do you expect development projects for real-time embedded systems in your department to begin using many-core embedded processors (i.e. processors with more than 16 cores)? (n=97)



Observations: Multi-core systems (Question 24) are already widely used in current developments, with 80% of respondents indicating their use by 2021, and only 10% answering “I do not know”. The uptake of heterogeneous multi-core systems (Question 25) lags behind simpler multi-core systems, but nevertheless just under 60% of respondents indicate their use by 2021, with 20% answering “I do not know”. Finally, the uptake of many-core systems (Question 26) is less certain, with 36% of respondents answering “I do not know”, 33% indicating take up by 2021, and 48% take up by 2029.

Question 27: By which year did or do you expect new development projects for real-time embedded systems in your department to stop using single-core embedded processors (i.e. processors with one core)? (n=97)




Observations: Although the proportion of respondents expecting to use single-core devices drops in future years, a substantial minority (31%) still expect to use single cores after 2029. Interestingly, this is the case for respondents who indicated each of the Automotive, Avionics, and Consumer Electronics domains in Question 4, with 30%, 34.5%, and 30%, respectively, expecting to use single-cores after 2029.

Familiarity with Real-time Systems Research: This part of the survey asked questions about familiarity with the real-time systems research community and its results.


Question 28: How many research publications (e.g. conference or journal papers) in the real-time systems field have you read in the last year? (n=96)

Question 29: How many real-time systems research publications (e.g. conference or journal papers) have you published as a (co-)author in the last 5 years? (n=96)



Observations: Around 79% of respondents read at least one research publication in the past year, and around 55% contributed to research publications in the past 5 years.


Question 30: How do you interact with the real-time research community? (n=97)



Observations: Only 16% of respondents have no interactions with the real-time research community. “Other interaction” (6%) included: research internships, co-supervisions, and interacting with researchers directly.

Follow Up: The final part of the survey asked questions about following up on this survey and general remarks.

Question 31: Indicate the purposes for which we may contact you again, if any. (n=60)



Observations: 48 respondents provided their email addresses for subsequent follow up.

Question 32: Enter feedback or remarks (n=23)

Observations: 23 respondents provided feedback. The most common comments were complementary remarks about the survey, and a desire to see the results.

V. CONCLUSIONS

An absence of any systematic studies into industry practice increases the risk that academic research will diverge from areas that are crucial to the development of future industrial systems. This may lead to research that is less relevant, less likely to be adopted, and has a lower potential for impact.

While empirical survey-based research is well-established in software and system engineering, there were previously no such studies in the real-time systems field. This paper addresses that omission by presenting the results of a survey,

containing 32 questions related to methods, tools, and trends in industrial real-time systems development. The survey was completed by 120 industry practitioners from a variety of different organizations, countries, and application domains.

The survey results show that industry recognizes the importance of timing predictability, but that other design aspects are of equal or greater importance, such as functional correctness and reliability/availability. Hence, it is important for real-time systems research to be cognizant of its impact on these aspects.

Many real-time systems today are distributed systems that use multi-core processors, and have complex memory hierarchies. Further, multiple different operating systems and networking technologies are typically utilized within the same system, as are different types of timing constraints and task activation patterns. Many respondents did not consider timing constraints to be sacrosanct, with even the most time critical functions allowed to miss some deadlines.

There is no silver bullet to manage timing behavior in complex real-time systems. Instead, the survey reveals that a wide range of different tools, techniques, and policies are used for timing analysis, scheduling, and to increase timing predictability. There is no one size fits all solution.

The trends suggest that single-core systems are still widely used today, and are expected to remain relevant for new developments for at least the next ten years. However, more complex (heterogeneous) multi- and many-core systems are already prevalent and their adoption is expected to increase significantly during the 2020s.

Finally, we would like to end by encouraging the real-time systems research community to read this survey carefully and reflect on how they can address the variety and complexity of future industrial real-time systems in their research.

ACKNOWLEDGMENTS

This research was funded, in part, by the Innovate UK HI-CLASS project and by the Netherlands Organisation for Applied Scientific Research TNO.

REFERENCES

- [1] A. Höfer and W. F. Tichy, "Status of empirical research in software engineering," in *Empirical Software Engineering Issues. Critical Assessment and Future Directions*. Springer, 2007, pp. 10–19.
- [2] C. Wohlin, P. Runeson, M. Höst, M. C. Ohlsson, B. Regnell, and A. Wesslén, *Experimentation in software engineering*. Springer Science & Business Media, 2012.
- [3] B. A. Kitchenham and S. L. Pfleeger, "Personal opinion surveys," in *Guide to advanced empirical software engineering*. Springer, 2008, pp. 63–92.
- [4] G. Muller, "Systems engineering research methods," *Procedia Computer Science*, vol. 16, pp. 1092–1101, 2013.
- [5] P. Mohagheghi and V. Dehlen, "Where is the proof? – a review of experiences from applying MDE in industry," *Lecture Notes in Computer Science*, vol. 5095, pp. 432–443, 2008.
- [6] M. Torchiano, F. Tomassetti, F. Ricca, A. Tiso, and G. Reggio, "Relevance, benefits, and problems of software modelling and model driven techniques – a survey in the Italian industry," *Journal of Systems and Software*, vol. 86, no. 8, pp. 2110–2126, 2013.
- [7] G. Liebel, N. Marko, M. Tichy, A. Leitner, and J. Hansson, "Model-based engineering in the embedded systems domain: an industrial survey on the state-of-practice," *Software & Systems Modeling*, vol. 17, no. 1, pp. 91–113, 2018.
- [8] J. Hutchinson, J. Whittle, and M. Rouncefield, "Model-driven engineering practices in industry: Social, organizational and managerial factors that lead to success or failure," *Science of Computer Programming*, vol. 89, pp. 144–161, 2014.

- [9] A. Forward and T. C. Lethbridge, "Problems and opportunities for model-centric versus code-centric software development: a survey of software professionals," in *Proceedings of the 2008 international workshop on Models in software engineering*. ACM, May 2008, pp. 27–32.
- [10] M. Broy, S. Kirstan, H. Krčmar, B. Schätz, and J. Zimmermann, "What is the benefit of a model-based design of embedded software systems in the car industry?" *Emerging Technologies for the Evolution and Maintenance of Software Models*, pp. 343–369, 2012.
- [11] J. Whittle, J. Hutchinson, and M. Rouncefield, "The state of practice in model-driven engineering," *IEEE software*, vol. 31, no. 3, pp. 79–85, 2014.
- [12] A. Vetro, W. Bohm, and M. Torchiano, "On the benefits and barriers when adopting software modelling and model driven techniques – an external, differentiated replication," in *Empirical Software Engineering and Measurement (ESEM), 2015 ACM/IEEE International Symposium on*. IEEE, 2015, pp. 1–4.
- [13] F. Hermans, M. Pinzger, and A. Van Deursen, "Domain-specific languages in practice: A user study on the success factors," in *International Conference on Model Driven Engineering Languages and Systems*. Springer, 2009, pp. 423–437.
- [14] J. Whittle, J. Hutchinson, M. Rouncefield, H. Burden, and R. Heldal, "A taxonomy of tool-related issues affecting the adoption of model-driven engineering," *Software & Systems Modeling*, vol. 16, no. 2, pp. 313–331, 2017.
- [15] A. Kuhn, G. C. Murphy, and C. A. Thompson, "An exploratory study of forces and frictions affecting large-scale model-driven development," in *International Conference on Model Driven Engineering Languages and Systems*. Springer, 2012, pp. 352–367.
- [16] N. C. Audsley, A. Burns, R. I. Davis, K. W. Tindell, and A. J. Wellings, "Fixed priority pre-emptive scheduling: An historical perspective," *Real-Time Syst.*, vol. 8, no. 2–3, p. 173–198, Mar. 1995.
- [17] L. Sha, T. Abdelzaher, K.-E. Årzén, A. Cervin, T. Baker, A. Burns, G. Buttazzo, M. Caccamo, J. Lehoczky, and A. K. Mok, "Real time scheduling theory: A historical perspective," *Real-Time Syst.*, vol. 28, no. 2–3, p. 101–155, Nov. 2004.
- [18] R. I. Davis and A. Burns, "A survey of hard real-time scheduling for multiprocessor systems," *ACM Comput. Surv.*, vol. 43, no. 4, Oct. 2011.
- [19] G. C. Buttazzo, M. Bertogna, and G. Yao, "Limited preemptive scheduling for real-time systems. a survey," *IEEE Transactions on Industrial Informatics*, vol. 9, no. 1, pp. 3–15, 2013.
- [20] A. Burns and R. I. Davis, "A survey of research into mixed criticality systems," *ACM Computing Surveys (CSUR)*, vol. 50, no. 6, pp. 1–37, 2017.
- [21] A. K. Singh, P. Dziurzanski, H. R. Mendis, and L. S. Indrusiak, "A survey and comparative study of hard and soft real-time dynamic resource allocation strategies for multi-/many-core systems," *ACM Comput. Surv.*, vol. 50, no. 2, Apr. 2017.
- [22] R. Wilhelm, J. Engblom, A. Ermedahl, N. Holsti, S. Thesing, D. Whalley, G. Bernat, C. Ferdinand, R. Heckmann, T. Mitra *et al.*, "The worst-case execution-time problem—overview of methods and survey of tools," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 7, no. 3, p. 36, 2008.
- [23] C. Maiza, H. Rihani, J. M. Rivas, J. Goossens, S. Altmeyer, and R. I. Davis, "A survey of timing verification techniques for multi-core real-time systems," *ACM Computing Surveys (CSUR)*, vol. 52, no. 3, pp. 1–38, 2019.
- [24] R. I. Davis and L. Cucu-Grosjean, "A survey of probabilistic schedulability analysis techniques for real-time systems," *LITES: Leibniz Transactions on Embedded Systems*, pp. 1–53, 2019.
- [25] —, "A survey of probabilistic timing analysis techniques for real-time systems," *Leibniz Transactions on Embedded Systems*, vol. 6, no. 1, pp. 03–1, 2019.
- [26] [Online]. Available: <https://github.com/ebni/classify-rt>
- [27] [Online]. Available: <https://site.ieee.org/tcrts/classification-of-rtss-papers/>
- [28] B. Akesson, M. Nasri, G. Nelissen, S. Altmeyer, and R. I. Davis, "A comprehensive survey of industry practice in real-time systems," *Tech. Rep.*, 2020. [Online]. Available: <https://akesson.nl/publications/>
- [29] [Online]. Available: <https://www.surveymonkey.com>
- [30] B. Akesson, M. Nasri, G. Nelissen, S. Altmeyer, and R. I. Davis, "Data Set Real-Time Systems Survey," Oct. 2020. [Online]. Available: https://uvaauas.figshare.com/articles/dataset/Data_Set_Real-Time_Systems_Survey/13117709
- [31] [Online]. Available: <https://www.linkedin.com/>