FACULTY
OF MATHEMATICS
AND PHYSICS
Charles University

Prague, September 7, 2020

Doctoral Thesis Review

Thesis: Scheduling Algorithms For Time-Triggered Communication Protocols

Author: Jan Dvořák

Reviewer: Lubomír Bulej

For some time, the automotive and aviation industries have been moving away from traditional (mechanical) systems for controlling vehicular functions, replacing direct mechanical linkages between the control and the controlled elements with electrical signals transmitted between sensors (including human-machine interface elements), control units, and actuators. In addition to saving weight, the resulting x-by-wire systems enable more precise control and tuning, as well as computer-controlled interpretation of human input (and emergency intervention) leading to improved safety.

To preserve flexibility (and to eliminate excessive wiring), the components of the control system are attached to a shared, fault-tolerant interconnect, resulting in a network of interacting components responsible for sensing, control, and actuation. To reduce cost, the interconnect is used to carry both time-critical and best-effort traffic, but reserves bandwidth for time-critical communication, which typically follows a deterministic schedule. As the sophistication of the control systems increases, so does the number components and the amount of data they need to exchange, stretching the bandwidth requirements for time-critical and best-effort traffic. Further exacerbating the bandwidth problem is the need to maintain backward (between past and present variants) and lateral (between multiple current variants) compatibility of the time-critical communication schedules to allow sharing of components within product lines.

In this context, the thesis deals with scheduling of time-triggered communication on the FlexRay and Time-Triggered Ethernet interconnects, typically used in the automotive and aviation domains. Specifically, the thesis aims to advance the state of the art in scheduling of time-triggered communication by addressing the lack of support for incremental and/or multi-variant scheduling. This is a major pain point for manufacturers with a large portfolio of products with long, overlapping lifespans, and requires algorithms that can minimize the bandwidth requirements of time-triggered communication while respecting the trade-offs needed to provide backward and lateral compatibility. I therefore consider *the topic and the objectives of the thesis to be both sufficiently relevant and sufficiently challenging.*

While the thesis provides an explicit and detailed list of various major and minor contributions, the key contributions correspond to the three main chapters of the thesis. The first key contribution is an algorithm for scheduling time-triggered communication in the static segment of the FlexRay interconnect over two independent channels, which can double the available bandwidth for signals that need not be fault-tolerant. The second key contribution is a heuristic algorithm for incremental and multi-variant scheduling, also on

Department of Distributed and Dependable Systems
Malostranské nám. 25, 118 00 Praha 1, Czech Republic
phone +420-951554245, fax +420-951554323
https://d3s.mff.cuni.cz

the FlexRay interconnect.. And finally the third key contribution is an algorithm for incremental scheduling of time-triggered traffic on the Time-Triggered Ethernet interconnect. *The contributions of the thesis advance the state of the art and address an important problem hindering network communication design in industrial applications.*

The author follows an established research methodology, in which he first surveys the state of the art in a particular area, showing that the problem being addressed is an actual problem, provides a problem statement and a break-down into sub-problems that need to be addressed separately, presents solutions to the sub-problems and the integrated solution, and provides an extensive evaluation of the solution on benchmark problem instances. Due to their nature, many of the sub-problems can be formulated (and the author formulates them) as graph theory, constraint programming, or (mixed) integer linear programming problems. In addition to using established terminology, this also provides a sense of problem complexity and motivation for heuristic solutions. The evaluation uses both synthetic and real-world problem sets, and explores sensitivity to various problem parameters in addition to key performance metrics. I therefore conclude that *the approach and the research methodology are sound and sufficient.*

The high-level goals and objectives of the thesis are formulated in a very generic fashion—almost at a meta-level—prescribing a general approach rather than specific goals. This is apparently meant to provide a common umbrella for the content of Chapters 2, 3, and 4, which are based on published journal papers. In light of the contributions of the thesis and due to alignment of the overall approach and methodology with the stated objectives of the thesis, I conclude that *the thesis has achieved the stated goals and objectives.*

The presentation of the thesis exposes a few rough edges, apparently caused by putting together previously published material. These manifest as minor inconsistencies in the symbols used (despite the table of symbols at the end of the thesis), and presentation issues, e.g., some terms (such as release time) are never defined, while some other terms or symbols are explained well after first use. Each of the constraints in the formulation in various problems would deserve a commentary, instead of an overblown equation full of symbols, which is completely unreadable without a symbol lookup table (especially Eq. 4.1). When evaluating the FlexRay scheduling algorithms, the numbers of allocated slots are presented with a fractional part, even though only whole slots can be allocated, therefore the numbers should be rounded up to the next integer. The text (and figure caption, e.g. Figure 4.5) talks about differences, but the plot shows absolute values. The thesis lists a comparison of the ECU-to-channel assignment methods to an alternative based on a genetic algorithm as a contribution, but the genetic algorithm is never really described in the thesis. These issues do add unnecessary friction to the presentation flow, but in the end I do not consider them critical.

While I am generally satisfied with the thesis, I decided to list some questions that might motivate discussion during the defense:

- In Chapter 2, schedule comparison only takes into account the total number of allocated slots. Is there no value in balancing the utilization of the two channels?
- In Chapter 2, during hyper-graph construction, why the algorithm does not distinguish between senders and receivers? For two hyper-edges with the same

Department of Distributed and Dependable Systems
Malostranské nám. 25, 118 00 Praha 1, Czech Republic
phone +420-951554245, fax +420-951554323
https://d3s.mff.cuni.cz

number of ECUs, a hyper-edge with more senders will need more slots to be allocated.

- In Chapter 2, what is the real-world benefit of the CAH heuristic? The ILP approach can—apparently—handle the real-world benchmark instances just fine, and it is only the synthetic benchmarks that benefit from the heuristic.
- In Chapter 3, does the extensibility optimization have a real-world impact? Figure 3.15 does show a measurable difference, but the optimization saves ~7 slots out ~122 in the first iteration, and ~3 slots out of ~139 in the last iteration.
- In Chapter 3, why the algorithm does not take advantage of the dual-channel scheduling presented in Chapter 2?
- What are the industrial requirements? What is unacceptable algorithm execution time and in what context? Is there any use case for periods that are not harmonic? Discussing these requirements explicitly would provide a better frame of reference for various optimizations.
- Would it be possible to provide a common framework for the scheduling of time-triggered communication over different interconnects? It appears that despite differences in terminology, the concepts, execution, and requirements related to time-triggered communication are very similar, while some of the constraints appear to be due to topology and/or technology (bus vs. point-to-point). Perhaps it is not such a stretch to expect a common framework that could be specialized to FlexRay and Time-Triggered Ethernet.

In closing, I am fully convinced that the thesis presents original and relevant contributions to the state of the art and therefore meets the established dissertation standards. The author has demonstrated his ability to conduct independent research and **I recommend the thesis for presentation and defense.**

Lubomír Bulej

Department of Distributed and Dependable Systems
Malostranské nám. 25, 118 00 Praha 1, Czech Republic
phone +420-951554245, fax +420-951554323
https://d3s.mff.cuni.cz