Timing verification of automotive communication architectures using quantile estimation

Nicolas NAVET – University of Luxembourg – L1359 Luxembourg
Shehnaz LOUVART – Renault SAS Technocentre – F78288 Guyancourt
Jose VILLANUEVA – Renault SAS Technocentre – F78288 Guyancourt
Sergio CAMPOY-MARTINEZ – Renault SAS Technocentre – F78288 Guyancourt
Jörn MIGGE – RealTime-at-Work (RTaW) – F54600 Villers-lès-Nancy

Keywords : Timing verification, automotive networks, Controller Area Network, simulation, statistics.

Abstract: Early stage timing verification on CAN traditionally relies on simulation and schedulability analysis, also known as worst-case response time (WCRT) analysis. Despite recent progresses, the latter technique remains pessimistic especially in complex networking architectures with gateways and heterogeneous communication stacks. Indeed, there are practical cases where no exact WCRT analysis is available, and merely upper bounds on the response times can be derived, on the basis of which unnecessary conservative design choices may be made. Simulation, on the other hand, does not provide any guarantees per se and, in the context of critical networks, should only be used along with an adequate methodology. In this paper, we argue for the use of quantiles of the response time distribution as performance metrics providing an adjustable trade-off between safety and resource usage optimization. We discuss how the exact value of the quantile to consider should be chosen with regard to the criticality of the frames, and illustrate the approach on two typical automotive use-cases.

1 Introduction

1.1 Increased bandwidth requirements and more complex architectures

Automotive multiplexing technologies, and specifically CAN, helped to keep the wiring harness complexity under control, improve existing Electrical and Electronic (EE) functions and introduce a wide range of new ones. In today's cars, there are thousands of signals exchanged by several tens of ECUs, with some signals having timing constraints below 5ms. The trend towards increased bandwidth requirements has never weakened, and along with topology and functional domain constraints, has led to the use of several CAN clusters within a car, more than 4 or 5 in some cases [1,6]. Also, the data rates of the CAN buses are now higher (e.g., 250kbit/s or 500kbit/s for a body network when it used to be 125kbit/s) and the bus load level has increased (e.g., sometimes greater than 50%). The architectures are becoming complex because of gateways between the CAN buses or sometimes between a CAN bus and other networking technologies (typically MOST and FlexRay today, see [6], and probably Ethernet in the future).

1.2 Optimized CAN networks

Optimizing CAN networks, in order to reach higher load levels, has now become an industrial requirement. Techniques and tools serving that purpose have been developed in the industry over the last 10 years (see [1] for a recap), and they are now mature and well supported by a number of COTS tools. For instance, one of the most efficient techniques, that is routinely used in the design of automotive networks, is the use of offsets to desynchronize the streams of frames [2,7]. Also, as far as possible, communication stacks should be free from implementation choices and hardware limitations that cause departures from the ideal priority-based behavior of CAN (e.g., the use of a FIFO waiting queue for frames waiting for transmission, insufficient number of transmission buffers in the communication controller, etc - see §3.1).

1.3 The need for more accurate timing verification model

In the early design phases of a vehicle, verification has to be done on models of the system. The two main verification techniques are simulation and schedulability analysis (also often called worst-case response time analysis), which basically consists in building a mathematical model of the worst possible situations that can be encountered at run-time. The most important performance metrics is the communication latency, or frame response time, that is the time it takes for a message to be received by its consumer nodes.

Optimized CAN networks mean higher network loads, and indeed they may now easily exceed 50% whereas 30%-35% were considered for a long time as a threshold not to be exceeded. Because there is less error margin, there is a need for a more rigorous development process in conjunction with timing verification models that are more fine-grained than they were in the past. For instance, a single station with a FIFO queue can create bursts of high priority frames that will impact the latencies of the frames sent by all the other stations (see [1]), possibly it may even propagate to other networks through the increased jitters of the frames that are forwarded through the gateways. Such behaviors cannot be overlooked in the timing analysis.

1.4 Simulation versus schedulability analysis

Schedulability analysis allows to calculate upper bounds on frame response times, and thus in principle it provides firm and deterministic guarantees. However, in many practical cases, schedulability analyses are not able to capture the complexity of the real systems (see Section 3), and if they do, results are often pessimistic (see [17]), which can lead to unnecessarily conservative design choices. Besides, a mathematical analysis with deterministic assumptions is intrinsically not well suited to model random phenomena such as transmission errors, or the transmission of aperiodic frames. Our view is that if schedulability analysis is absolutely required, the system has to be conceived accordingly from the beginning. For instance, all departures from the temporal behaviors that can be analyzed must be prevented. This imposes overhead in the design process, for instance, more constraining specifications to the suppliers.

Simulation models on the other hand are easier to develop and validate (see [17]), and they can be more realistic. For instance, the simulator RTaW-Sim [14] captures all timing-relevant characteristics of a CAN-based communication architecture, including behaviors in the gateways. Furthermore, the information obtained from simulation runs is much more fine-grained than schedulability analysis, and basically any statistics of interest can be collected. For instance, knowing the extent to which successive frame deadline misses might jeopardize the stability of a control law can be answered with the knowledge of the distribution of successive deadline misses.

1.5 Contribution of the paper

Quantiles are used as risk measures in several application domains. For instance, in finance, it is the basis of Value-at-Risk (VaR) which is a cornerstone of the risk assessment framework. This paper advocates the use of quantiles of the frame response time distribution derived by simulation as a practical performance evaluation technique for automotive communication architectures. We highlight the limitations of worst-case schedulability analysis for CAN-based automotive systems, discuss the use-cases of quantiles and provide guidelines for the use of quantile-based performance evaluation.

2 Metrics for the evaluation of frame latencies

The main performance metric for real-time communication networks is the communication latency, also called frame response time, that is the time from the production of a message until the reception by the stations that consume the message. The latency constraint, or deadline constraint, is the maximum allowed value for the response time. This deadline is typically inherited from applicative level constraints or regulatory constraints (e.g. time to answer a diagnosis request). The aim of timing verification is to make sure that deadline constraints are met. Timing verification on models, by simulation or schedulability analysis, allows deriving a number of metrics on the frame response times. Those metrics, along with the corresponding timing verification techniques are shown in Figure 1.

2.1 Bound, worst-case response time and simulation maximum

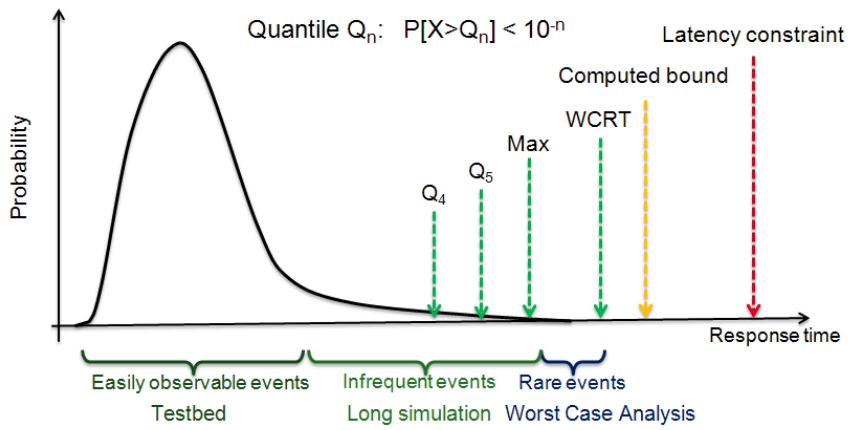

Figure 1: Metrics of the frame latencies and techniques to verify them. The black curve shows an idealized distribution of a frame response times.

The bound on the response time, which is the outcome of a schedulability analysis, is usually larger than the true worst-case possible response time (denoted by WCRT). For instance, as soon as there is a gateway or non-ideal communication stacks, schedulability analysis is pessimistic to an extent that cannot be predicted. The maximum value seen during a simulation is less than the WCRT, here again the distance between both values is unknown.

Figure 2 shows, for a typical 500kbit/s automotive network, the values of the quantiles and the upper bound computed with the schedulability analysis published in [3]. In the specific setup of this experimentation, the schedulability analysis is almost exact (see footnote 3 in [3]), thus the upper bounds on the response times, denoted by WCRT on Figure 2, will most likely be very close to the true worst-case possible response times.

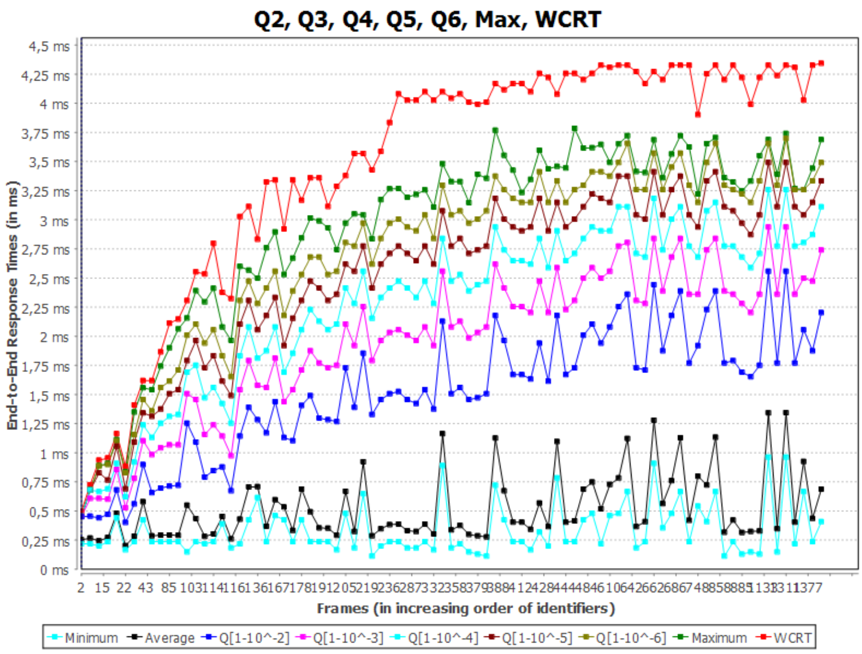

Figure 2: Illustration of the values of the possible performances metrics for the set of frames of a realistic 500kbit/s automotive CAN network with offsets. The frames are sorted by decreasing priority. The following inequations always hold: Min ≤ Average ≤ Q2 ≤ Q3 ≤ Q4 ≤ Q5 ≤ Q6 ≤ Max ≤ WCRT. The quantiles are evaluated by simulation using RTaW-Sim[14] and the worst-case response time using NETCAR-Analyzer plug-in for RTaW-Sim[15].

2.2 Quantiles of the response times

Formally, for a random variable X, a p-quantile is the largest value x such that

$$P[X \leq x] \geq p \quad \text{or equivalently,} \quad P[X>x] < 1 - p$$

In other words, it is a threshold L such that for any response time,

- the probability to be smaller than L is larger than p,
- the probability to be larger than L is smaller than 1 – p.

For example, the probability that a response-time is larger than the $(1\text{-}10^{-3})$-quantile, also denoted by Q3 quantile or Q3 for short, is lower than 10^{-3}. For a frame with a period of 10ms, the Q3 will be exceeded on average once every $10^{3}\cdot 10\text{ms}=10^{4}\text{ms}$, that is 10s. Table 1 shows how quantiles translate to deadline miss frequency and average time between deadline misses, for frames with a period equal to 10ms and 500ms and deadlines assumed to be equal to quantiles.

Quantile	One frame every	Mean time to deadline miss if frame period is 10ms	Mean time to deadline miss if frame period is 500ms
Q3	1000	10 s	8mn 20s
Q4	10 000	1mn 40s	≈ 1h 23mn
Q5	100 000	≈ 17mn	≈ 13h 53mn
Q6	1000 000	≈ 2h 46mn	≈ 5d 19h

Table 1: Quantiles and corresponding frame deadline miss frequencies for frame periods equal to 10ms and 500ms, and frame deadlines assumed to be equal to quantiles values.

It should be noted that quantile values derived by simulation do not say anything about possible successive quantile overshoots, which may pose a threat for the stability of control laws for instance. This can happen if large response times are temporally correlated, as it typically happens in overload conditions. The distribution of the number of successive response times above a certain quantile, obtained with RTaW-Sim in our case, helps the designer to answer this question. It is possible to go beyond the verification of strictly successive quantile overshoots by testing linear and non-linear temporal dependences among the response time values above a certain quantile. Classical statistical tests that can serve that purpose are the auto-correlation and BDS test (Brock, Dechert, Scheinkman - see [20]).

2.3 Defining target quantile and simulation length

The first step in timing verification is to identify the actual timing constraints of the frames. Then, depending on the criticality of the frames, the deadline miss frequency that can be tolerated is defined (see Section 4 for examples). This value directly determines the target quantile Qn (see Table 1), which may be different for each frame. Finally the communication architecture is simulated and if, for each frame, the value of the target quantile is below the timing constraint, the system meets the performance objective. If not, the system has to be at least partly redesigned, for instance by increasing the priority for some frames, or optimizing further the offsets values.

The simulation length is crucial in order to obtain robust statistics, because the higher the quantile the longer the simulation time needed. The evaluation of how long is enough can be done for instance on the basis of the confidence intervals of the quantiles, that can be derived using computational methods such as boostrapping or other statistical methods (see [21] for a review and comparative assessment). It should be paid attention here that the statistical method used for computing confidence interval is suited to higher quantiles, and that the system does not depart significantly from the underlying assumptions of the method (typically, the sequence of response times is independent and identically distributed – i.i.d). Basically, with a high-performance simulation engine such as RTaW-Sim, it is possible to obtain reasonable values for Q5 and Q6 (for frames with periods <500ms) in a few hours of simulation. It should be noted that good tool support can significantly help the designer here, for instance by suggesting the right simulation length for a given quantile, or identifying values computed on too small samples. In RTaW-Sim for instance, this is done by displaying these values with several intensities of gray, whereas values that are computed on sufficiently large samples are displayed in plain black text.

3 Limitations of schedulability analyses – illustration on CAN

Worst-case response time (WCRT) analysis is generally considered as the technique that is the best suited to provide the guarantees that are needed in critical networks. Here we highlight – in the context of automotive CAN networks - that this approach may suffer from several shortcomings.

3.1 Pessimism due to coarse-grained models of the communication stack

If the first deterministic analytical model of CAN communication, proposed in the 90s in [10], was a milestone in timing verification, it was a rather simplified abstraction with regard to the traffic models and the characteristics of the communication stacks. Over the years, these limitations were identified and partially lifted. This includes:

- Non-abortable transmit requests [5], as some communication stacks/controllers may not offer the possibility to cancel lower-priority transmission requests when a higher priority frame is released,
- Limited number of transmit buffers, and delays in refilling the transmit buffers [5,6],
- The use of a FIFO waiting queue for frames, or any other policy than the Highest Priority First [2,9],
- Internal CAN controller message arbitration based on transmit buffer number rather than frame ID,
- Frame queuing not done in priority order (but for example by PDU index in Autosar),
- Possible asynchronisms between the applicative level tasks that produce the signals and the communication task responsible for building the frames from the signals, and issuing the transmission requests. In some cases, this task may suffer delays caused by higher priority activities,
- The additional jitters due to gateways [13].

However, most often the mathematical models developed are conservative in the sense that they are only able to compute safe upper bounds on the response time and not the exact value. This is in particular the case in [2,9] for FIFO waiting queues, where conservative assumptions are placed in order to be able to handle the overall complexity of the problem. Also, to the best of our knowledge, there is no exact analysis for the common case today, where CAN buses are interconnected by gateways. Moreover, a global framework that would make it possible to combine the analyses is still missing (though there is a first step in that direction in [16]).

This pessimism of the analyses is obviously an issue when the hardware resources have to be used at high-load levels, and the degree of pessimism can hardly be quantified. Basically it can go up to largest busy period on the bus, which, in most cases, is equal to the worst-case response time of the lowest priority frame computed under the assumption of an ideal CAN behavior ("whenever message arbitration starts on the bus, the highest priority message queued on each node is entered into arbitration"). This upper-bound on the pessimism is not helpful in practice.

3.2 Complexity of schedulability analysis

In [5] the authors explore another direction than coarse-grained models and develop an almost exact analysis for non-abortable transmission requests. However, the complexity of such an accurate analysis makes it error-prone should it be extended to integrate other non-ideal behaviors. The flaw in the original CAN analysis corrected in [18] after 13 years of extensive use, suggests to us that simple analyses, even if coarse-grained, should be preferred.

The algorithmic complexity of the analysis is also an issue that must be dealt with. For instance, as soon as frame offsets are considered, the complexity of an exact schedulability analysis becomes exponential in time. Ignoring frame offsets is not an option in most setups because the analysis becomes then so pessimistic that is hardly usable. If, with some optimizations[3] and taking advantage of the limited number of stations in an automotive CAN network, it is often possible to compute the exact value of the response times in the ideal case, to the best of our knowledge, this is not feasible for most departures from the ideal priority-based behavior of CAN. In our view, in many cases, an exact analysis that can be used on industrial size applications will remain out of reach because of the algorithmic complexity of the problem.

3.3 The gap between analytical traffic models and real automotive traffic

Besides the drawback of the pessimism, deterministic schedulability analyses are not able to model all kinds of traffic in a fully satisfactory manner. This is the case for non periodic traffic (fully event-triggered transmissions[8] or Autosar mixed transmission model[22]), diagnosis transactions (e.g., OBD2 requests) and the transmission of segmented data (e.g., code upload).

This is also the case for transmission errors which are random by nature and thus cannot be well captured by a deterministic model.
When the expressive power of a deterministic model is not sufficient, there are classically two options: either making pessimistic assumptions (e.g., modeling aperiodic frames like periodic ones), which is not always possible because it may result in an overloaded bus, or ignoring what cannot be modeled (e.g., ignoring transmission errors and aperiodic traffic). Obviously, both options are unsatisfactory because they are either inefficient in terms of resource usage or potentially unsafe. In our view, ideal network timing verification models should:

- Account for transmission errors, and possibly other errors states (e.g., ECU reboots, bus-off states in CAN),
- Provide realistic models of traffic, in particular the non-periodic part of the traffic [8], multi-packet messages such as diagnosis frames or code upload, and the transmission jitters (especially for frames that are forwarded from one network to another [13]). The non-periodic traffic is generally difficult to characterize, but if overlooked, one will tend to underestimate the frame latencies as shown in [8].

4 Use-cases of quantile-based performance evaluation

In this section, we illustrate how quantiles, collected by simulation, help to verify timing constraints and make design choices on two typical automotive use-cases. The execution time of a simulation mostly depends on the number of frames exchanged in a certain scenario. With typical automotive CAN-based communication architectures, a speedup of 200 compared to the simulated functioning time can be achieved with RTaW-Sim[14], which means that statistics with sample size larger than 100 for Q5 can be obtained in hours for frames with a period lower than or equal to 500ms.

4.1 OBD-II requests

OBD-II is standard diagnostics protocol to query information on the status of components; it is used during driving and maintenance. By regulation, the time between the request frame and the first frame of the response must be less than 50ms. The OBD frames, which are of low priority, will be delayed by all higher-priority frames exchanged on the bus. The ECU that is queried will need some time to produce an answer, and the delay from the reception of the OBD request until the availability of the first frame of the response is here assumed to be 30ms. Besides, when the nodes are not located on the same network, request and answer frames must be transferred through a gateway which increases the communication latencies. This is the case in the OBD transaction under study, with the two interconnected CAN buses respectively loaded at around 50% and 40%.
The aim is to assess if the communication architecture is able to meet the 50ms timing constraint. Precisely, a deadline miss must not occur more often than once every 1000 OBD requests. The evaluation is done with conservative assumptions on the communication stacks (e.g., FIFO waiting queues for the ECUs and the gateway), and transmission errors that can take place on the bus. Offsets are implemented for some of the frames to desynchronize the streams of frames and reduce the network latency, see [2,7]. Simulations were performed with RTaW-Sim which includes a dedicated traffic model for OBD transactions. The simulation length was set to 10 days, and with a time between two OBD requests equal to 100ms, a sample of $8.64 \cdot 10^6$ response times was collected.

Metrics	OBD response times
Min	31.94
Average	34.29
Q3	46.55
Q4	**49.31**
Q5	53.45
Q6	55.32
Max	56.57

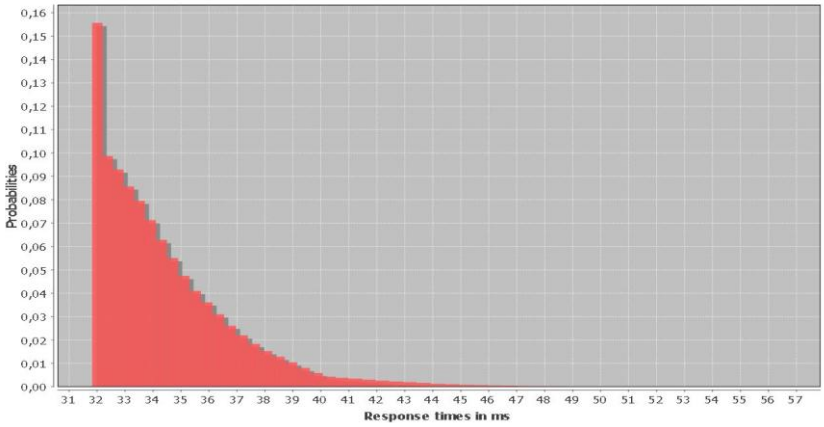

Figure 4: Quantiles and distribution of the OBD request response times.

One sees in Figure 4 that the target objective regarding the deadline miss frequency is met since the Q3, and even the Q4 quantile, is below 50ms. This means that, on average, less than one OBD request out of 10000 will have a response time larger than 50ms.

4.2 End-to-end latencies

The second use-case is to evaluate the probability that the freshness of the signals contained in a frame is not sufficient when received by consumer nodes located on another network. The frame under study, denoted by T13, is a 10ms frame containing signals used for the control of the vehicle. The experimental conditions are identical to the previous use-case. The end-to-end latency is made up of the delay on the source network, the gatewaying delay, and the latency on the destination network.

In Figure 5, one sees that T13 has a maximum end-to-end response time of 12.13ms which is larger than its period. However, latencies up to the value of Q6, that is 8.87ms, are acceptable, on the basis of the evaluation of their impact at the functional level. This impact evaluation involves converting the latencies into physical values such as speed, acceleration, and distances. Possible values above Q6 are considered rare enough to not jeopardize by any means the safety with regard to the dynamics of a vehicle.

ShortName	Payload	TxMode	Period	TxOffset	Min	Average	Q2	Q3	Q4	Q5	Q6	Max
T1	8	P	10	0	0,236	0,284	0,491	0,597	0,614	0,614	0,614	1,47
T2	8	P	10	0	0,478	0,534	0,786	0,848	0,857	0,857	0,857	1,712
T3	4	B			0,382	0,594	1,71	3,614	5,376	7,086	8,278	9,241
T4	6	P	10	0	0,2	0,262	0,705	0,973	1,15	1,305	1,414	1,421
T5	8	P	10	0	0,442	0,53	1,085	1,29	1,456	1,596	1,662	1,663
T6	8	P	10	0	0,684	0,803	1,398	1,605	1,73	1,886	1,905	1,905
T7	8	P	10	0	0,926	1,074	1,72	1,908	2,044	2,14	2,147	2,147
T8	8	P	10	0	0,236	0,357	1,374	1,811	2,068	2,228	2,374	2,389
T9	8	P	10	0	0,478	0,658	1,854	2,228	2,415	2,573	2,631	2,631
T10	6	P	10	0	0,684	0,924	2,241	2,472	2,648	2,805	2,837	2,837
T11	4	P	10	0	0,166	0,341	1,681	2,32	2,648	2,821	2,971	3,009
T12	8	P	10	0	0,424	0,658	2,153	2,741	3,023	3,184	3,251	3,251
T13	8	B			0,522	0,866	2,573	4,149	6,244	7,593	8,87	12,129
T14	8	P	20	0	0,72	1,058	2,726	3,258	3,511	3,614	3,719	3,735
T15	8	P	20	0	1,168	1,588	3,094	3,511	3,741	3,784	3,962	3,977
T16	8	P	20	20	0,236	0,483	2,307	3,129	3,593	3,895	4,102	4,215
T17	8	P	20	10	0,72	1,102	2,971	3,511	3,828	3,985	4,102	4,172
T18	8	P	20	10	0,962	1,442	3,412	3,828	4,125	4,246	4,37	4,443
T19	3	P	100	0	0,148	0,447	2,529	3,451	3,917	4,221	4,37	4,374
T20	8	P	100	0	0,39	0,779	3,041	3,895	4,345	4,602	4,616	4,616
T21	8	P	100	0	0,236	0,534	2,664	3,491	4,031	4,295	4,395	4,439

Figure 5: Statistics of the response times obtained by simulation with RTaW-Sim for frames of use-case 2.

5 Conclusion

Networking technologies such as AFDX have been conceived with the requirement that the temporal behavior of the network must be predictable, and this is why AFDX is amenable to worst-case timing verification by schedulability analysis with limited pessimism (see [19]). Other technologies, based on the Time-Triggered paradigm, such as TTP or TTEthernet, offer full determinism in the temporal domain and thus an easy timing verification. CAN-based automotive architectures are not as easily analyzable from a timing point of view, because of heterogeneous hardware and software components, and because the temporal behaviors of the ECUs and gateways are less constrained. As a result, CAN schedulability analyses are often not able to capture the entire complexity of the system with the risk to be pessimistic and possibly unsafe. Besides, it is acceptable for many functions to tolerate occasional deadline misses, provided that the risk is well assessed and acceptable with regard to the dynamics and the criticality of the functions.

This paper explores the use of quantile-based performance evaluation by simulation as a practical alternative to schedulability analysis for automotive systems. We discuss how quantiles can be used in the timing verification process, how to determine the target quantile for a frame with regard to its criticality, and how to set the simulation length accordingly. We emphasize that the extent to which successive deadline misses can occur can be assessed. Then, two case-studies illustrate how this approach based on quantiles provides the designer with fine-grained quantitative information supporting design choices.

6 References

[1] N. Navet, H. Perrault, "CAN in Automotive Applications: a Look Forward", 13th International CAN Conference, Hambach, Germany, March 5-6, 2012.

[2] R. Davis, N. Navet, "Controller Area Network (CAN) Schedulability Analysis for Messages with Arbitrary Deadlines in FIFO and Work-Conserving Queues", 9th IEEE International Workshop on Factory Communication System (WFCS 2012), May 21-24, 2012, Lemgo/Detmold, Germany.

[3] P. Meumeu-Yomsi, D. Bertrand, N. Navet, R. Davis, "Controller Area Network (CAN): Response Time Analysis with Offsets", 9th IEEE International Workshop on Factory Communication System (WFCS 2012), May 21-24, 2012, Lemgo/Detmold, Germany.

[4] A. Monot, N. Navet, B. Bavoux, C. Maxim, "Fine-grained Simulation in the Design of Automotive Communication Systems", Embedded Real-Time Software and Systems (ERTS 2012), Toulouse, France, February 1-3, 2012.

[5] D. Khan, R. Davis, N. Navet, "Schedulability analysis of CAN with non-abortable transmission requests", 16th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA 2011), Toulouse, France, September 2011.

[6] H. Kellerman, G. Nemeth, J. Kostelezky, K. Barbehön, F. El-Dwaik, L. Hochmuth, "BMW 7 Series architecture", ATZextra, November 2008.

[7] M. Grenier, L. Havet, N. Navet, "Pushing the limits of CAN – Scheduling frames with offsets provides a major performance boost", Proc. of the 4th European Congress Embedded Real Time Software (ERTS 2008), Toulouse, France, January 29 – February 1, 2008.

[8] D. Khan, N. Navet, B. Bavoux, J. Migge, "Aperiodic Traffic in Response Time Analyses with Adjustable Safety Level", IEEE ETFA2009, Mallorca, Spain, September 22-26, 2009.

[9] R.I. Davis, S. Kollmann, V. Pollex, F. Slomka, "Controller Area Network (CAN) Schedulability Analysis with FIFO queues". In proceedings 23rd Euromicro Conference on Real-Time Systems (ECRTS), pages 45-56, July 2011.

[10] K.W. Tindell, A. Burns. "Guaranteeing message latencies on Controller Area Network (CAN)", In Proceedings of 1st International CAN Conference, pp.1-11, September 1994.

[11] M. Di Natale, "Evaluating message transmission times in Controller Area Networks without buffer preemption", In 8th Brazilian Workshop on Real-Time Systems, 2006.

[12] D.A. Khan, R.J. Bril, N. Navet, "Integrating hardware limitations in CAN schedulability analysis," IEEE International Workshop on Factory Communication Systems (WFCS) pp.207-210, 18-21 May 2010.

[13] R. Davis, N. Navet, "Traffic Shaping to Reduce Jitter in Controller Area Network (CAN)", WiP at the 24th Euromicro Conference on Real-Time Systems (ECRTS12), Pisa, Italy, July 11-13, 2012.

[14] RealTime-at-Work, "RTAW-Sim: Timing-accurate simulation of Controller Area Network with fault-injection capabilities," available at http://www.realtimeatwork.com/software/rtaw-sim/, 2009.

[15] RealTime-at-Work, "NETCAR-Analyzer: Timing analysis and resource usage optimization for Controller Area Network," available at http://www.realtimeatwork.com/software/netcar-analyzer/, 2009.

[16] R. Davis, N. Navet, "Controller Area Network (CAN) Schedulability Analysis for Messages with Arbitrary Deadlines in FIFO and Work-Conserving Queues", 9th IEEE International Workshop on Factory Communication System (WFCS 2012), May 21-24, 2012, Lemgo/Detmold, Germany.

[17] N. Navet, "Quantile-based performance evaluation on CAN", presentation at the 14th International CAN Conference, Paris, November 12-13, 2013.

[18] R. Davis, A. Burn, R. Bril, and J. Lukkien, "Controller Area Network (CAN) schedulability analysis: refuted, revisited and revised", Real-Time Systems, vol. 35, pp. 239–272, 2007.

[19] M. Boyer, N. Navet, M. Fumey, "Experimental assessment of timing verification techniques for AFDX", Embedded Real-Time Software and Systems (ERTS 2012), Toulouse, France, February 1-3, 2012.

[20] W.A. Broock, J.A. Scheinkman, W.D. Dechert, and B. LeBaron. A test for independence based on the correlation dimension. Econometric Reviews, 15(3):197–235, 1996.

[21] S. Wallis, "Binomial confidence intervals and contingency tests: mathematical fundamentals and the evaluation of alternative methods", Journal of Quantitative Linguistics, vol. 20, issue 3, 2013.

[22] S. Mubeen, J. Maki-Turja, M. Sjodin, "Worst-case response-time analysis for mixed messages with offsets in Controller Area Network," IEEE ETFA 2012, Sept. 17-21, 2012