UNIVERSITÉ DU LUXEMBOURG

Part of this presentation is based on a paper to appear at ERTSS'2014: "Timing verification of automotive communication architectures using quantile estimation" co-authored with Shehnaz LOUVART (Renault), Jose VILLANUEVA (Renault), Sergio CAMPOY-MARTINEZ (Renault) and Jörn MIGGE (RealTime-at-Work).

Quantile-based performance evaluation on CAN

Nicolas NAVET

14th International CAN Conference
Paris, November 12-13, 2013.

November, 12 2013

1 Outline

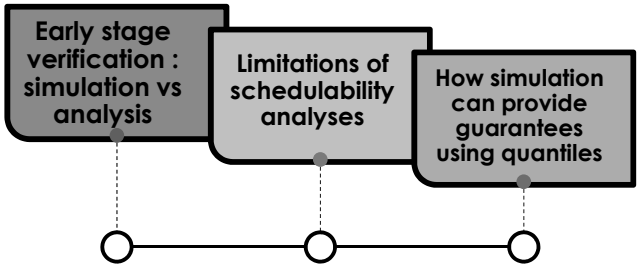

Schedulability analysis *"mathematic model of the worst-case possible situation"* VS **Simulation** *"program that reproduces the behavior of a system"*

$$K_i^k(t) \stackrel{\text{def}}{=} \underbrace{\left\lfloor \frac{J_i^k + \varphi_i^k(\phi^i)}{T_i^k} \right\rfloor}_{\text{max number of instances that can accumulate at critical instants}} + \underbrace{\left\lfloor \frac{t - \varphi_i^k(\phi^i)}{T_i^k} \right\rfloor + 1}_{\text{max number of instances arriving after critical instants}}$$

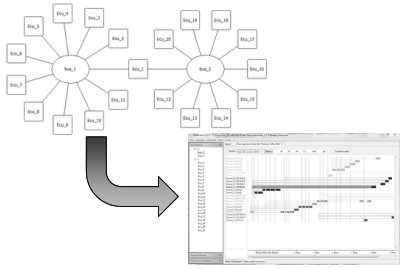

☺ Upper bounds on the perf. metrics → Safe (TBD)

☺ Analysis is known to be correct → Safe (TBD)

☹ Pessimistic → over-dimensioning

☹ Gap between models and real systems!

☹ Do not provide much information since a single trajectory is studied

☺ Models close to real systems

☺ Fine grained information

☹ Upper bounds are out of reach! → Unsafe (TBD)

☹ Model correctness is unsure

Beware of verification models !

"Schedulability analysis ensures safety!"
Our view: it might not be so…

1. Analytic models are pessimistic (except in the "ideal" case)
2. Analytic models are unrealistic (except in the "ideal" case)
3. Analytic models and their implementation can be flawed

"Simulation cannot provide firm guarantees"
Our view: it might not be so…

4. It is possible to verify correctness of simulation models
5. User- chosen guarantees can be enforced with proper methodology, e.g. with quantiles

Typical CAN-based automotive system

49% load – 500kbit/s
72 frames – 10ms to 2sec

38% load – 500kbit/s
44 frames – 10ms to 2sec

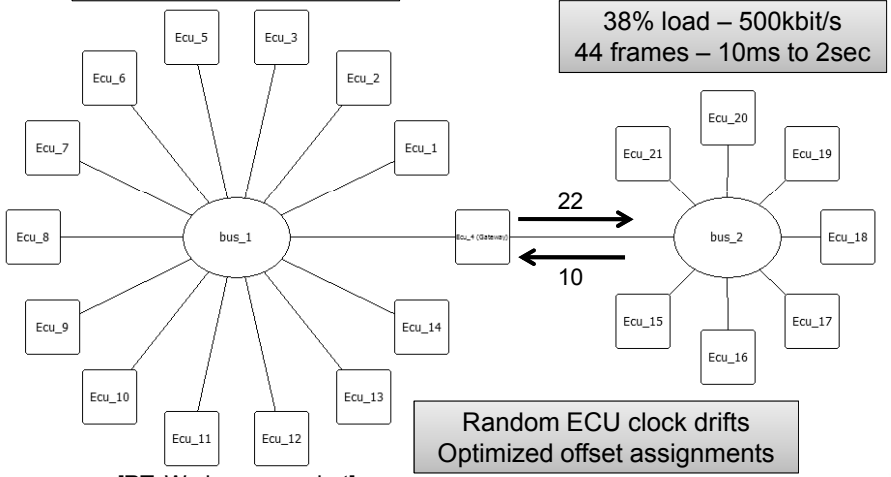

Random ECU clock drifts
Optimized offset assignments

[RTaW-sim screenshot]

Performance metric: frame response time ≈ communication latency
"Time from transmission request until frame received by consuming nodes"

Ideal CAN = *"whenever message arbitration starts on the bus, the highest priority message queued on each node is entered into arbitration"*

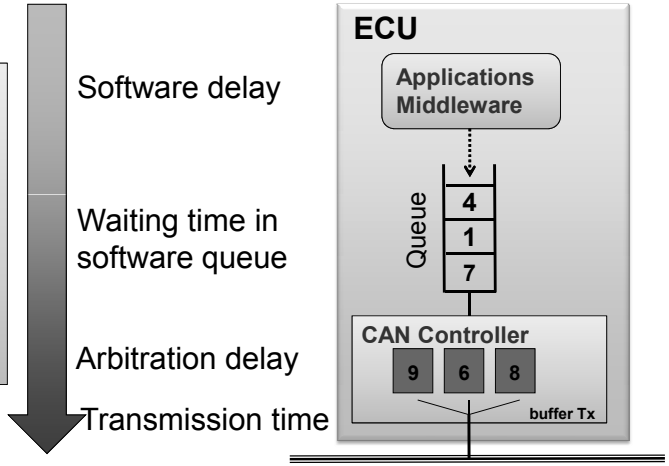

1

Analytic models are pessimistic (except in the "ideal" case)

Frame response time distribution

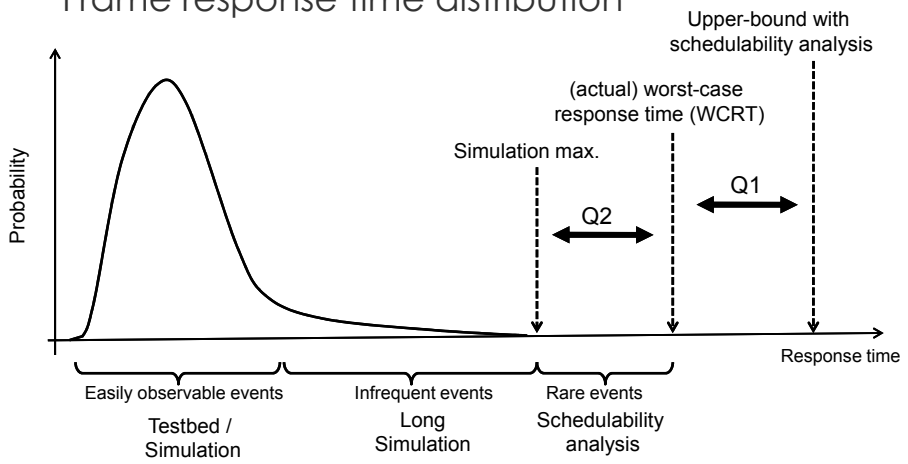

Q1: pessimism of schedulability analysis ?!

Q2: distance between simulation max. and WCRT ?!

(Typical) Frame response time distribution

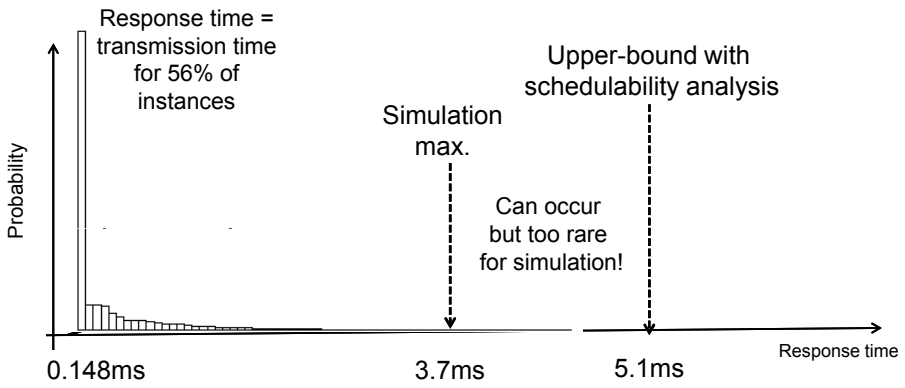

Medium priority frame on a 50% loaded 500kbits/bus with offsets

Q1 : Pessimism of CAN schedulability analysis ?
Q2: distance with simulation ?

Case 1: ideal communication stacks + no gateway →
the computed upper-bound can occur (and be re-simulated)

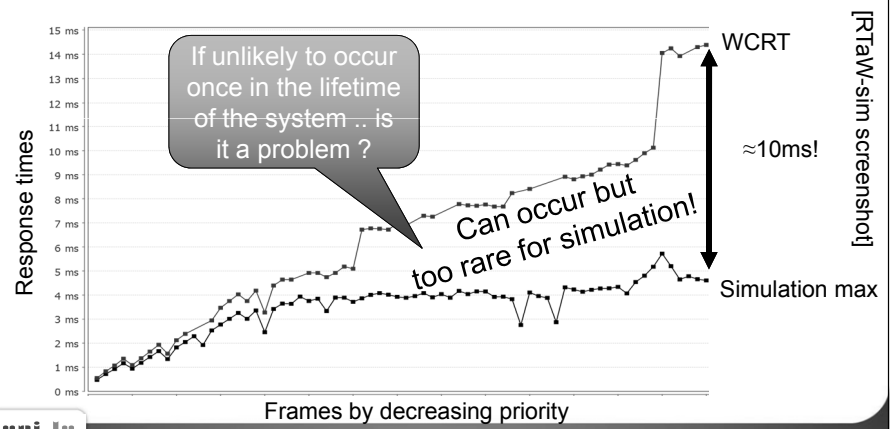

Q1 : Pessimism of CAN schedulability analysis ?
Q2: distance with simulation ?

Case 1: ideal communication stacks + no gateway →
the computed upper-bound can occur (and be re-simulated)

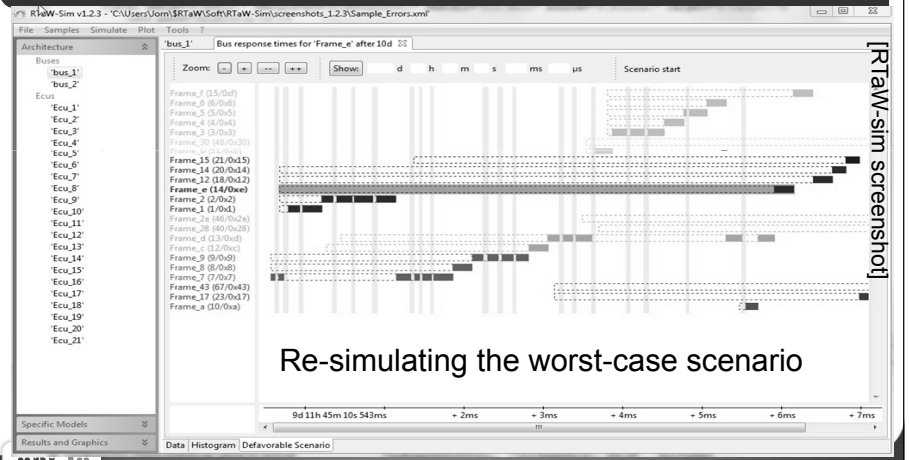

Q1 : Pessimism of CAN schedulability analysis ?
Q2: distance with simulation ?

Case 2: perfect communication stacks + **gateway** →
the computed upper-bounds do not occur for forwarded frames
in the general case

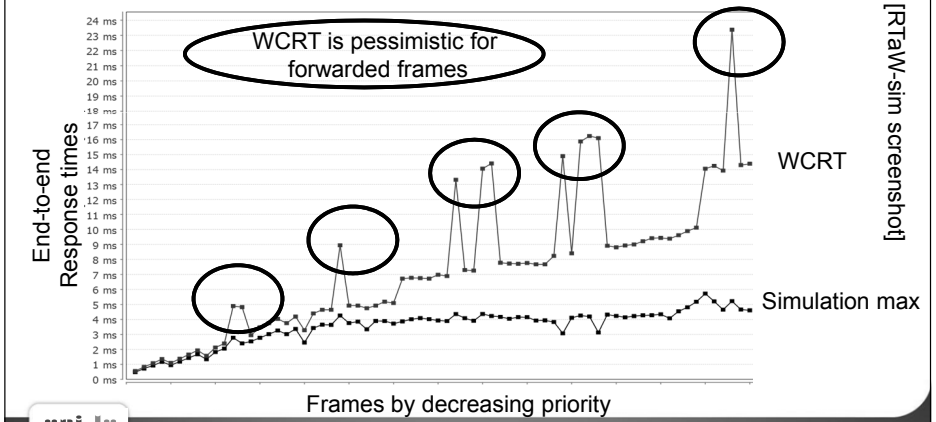

Q1 : Pessimism of CAN schedulability analysis ?
Q2: distance with simulation ?

Case 3: non-ideal communication stacks
the computed upper-bounds do not occur
in the general case – analysis are in general very pessimistic !

How Pessimistic ?!

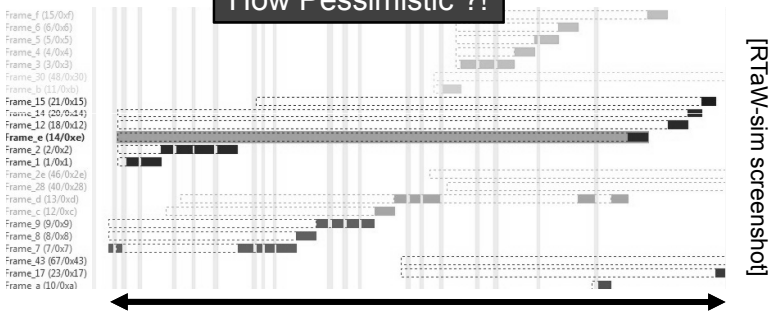

[RTaW-sim screenshot]

Up to the longest possible busy-period on the bus
≈ Worst-case response time of lowest priority frame in the ideal case

2

Analytic models are not realistic
(it the system has not been conceived with schedulability in mind)

Departure from ideal CAN: HW and SW

1. **Non-HPF waiting queues [5,6]**
2. **Frame queuing not done in priority order by communication task**
3. **Non abortable transmission requests [9]**
4. **Not enough transmission buffers [8,10]**
5. **Delays in refilling the buffers [11]**
6. **Delay data production / transmission request**

...

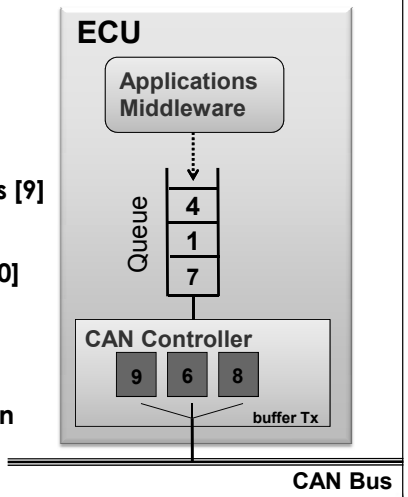

Departure from ideal CAN: frame transmission patterns

7. **code upload or segmented messages**
8. **Autosar-like mixed transmission models**
9. **Diagnostic requests**
10. **Transmission errors (probabilistic model ?! [1])**
11. **Aperiodic traffic (probabilistic model ?! [2])**
12. **Gatewayed traffic**

...

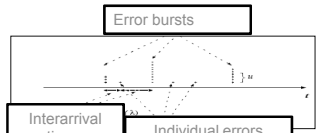

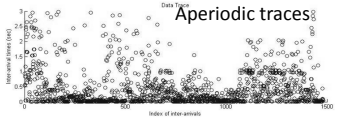

Higher load level calls for more realistic models

If the analytic model does not capture accurately all the characteristics of the system, then the results will be wrong … in an unpredictable manner

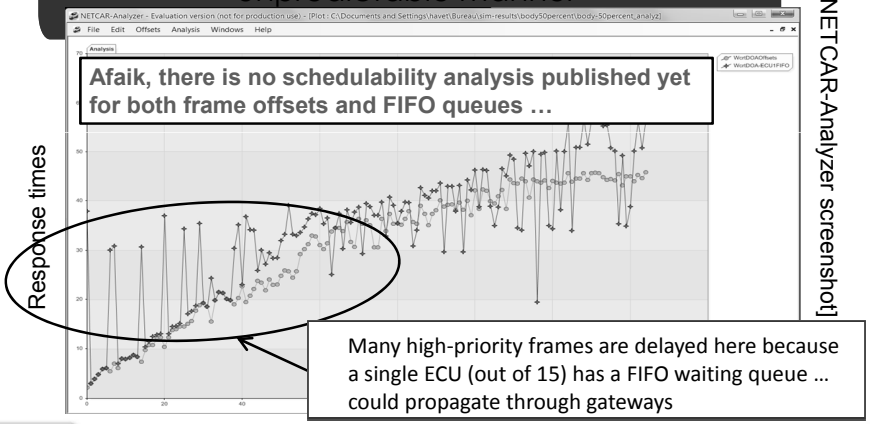

NETCAR-Analyzer screenshot]

About the suitability of schedulability analysis for non-ideal architectures..

- ✓ Good news: many works try to bridge the gap between analytic models and real systems [Ref.1 to 12]
- ✓ Bad news #1: not everything is covered, no integrated framework (first step in [6])
- ✓ Bad news #2: many existing analyses are conservative (= inaccurate), thus hardly usable for highly-loaded systems.
- ✓ Bad news #3: comprehensive and exact analysis would be overly complex (e.g. as in [9]) and intractable!

Personal view : both accurate and comprehensive analyses are out of reach … if you need analysis, you have to conceive the systems accordingly

3

And, schedulability analysis can be flawed …

What's different from other software (e.g. a simulator) ?

- ✓ Analysis are complex and error prone. remember "CAN analysis refuted, revisited, etc" [14] ?!
- ✓ Implementations are error prone: analyses complexity, floating-point arithmetic !, how to check correctness ?, not many end-users, cost-pressure, etc …
- ✓ Solutions ?
  - peer-review of the WCRT analyses is needed
  - coarse-grained / conservative but simple as far as possible: e.g., [5,6] vs [9]
  - no black-box software − documentation of implemented analyses and underlying hypotheses
  - rational arithmetic (w. float for Design Space Exploration)
  - cross-validation between tools / techniques on benchmarks

4

Simulation models validity can be questioned as well, after all …

Validating a network simulator ?

- ✓ Cross-validation by re-simulating worst-case situation from schedulability analysis (when possible)
- ✓ Cross-validation by comparison with real communication traces : e.g., comparing inter-arrival times distribution
- ✓ Checking a set of correctness properties on simulation traces

And model parameters must be realistic: transmission patterns, transmission errors, clock drifts, communication stacks, etc → analysis of communication traces is helpful here

5

Simulation can provide guarantees with proper methodology

Using quantiles means accepting a ***controlled*** risk

Quantile Q_n: $P[\text{ response time} > Q_n] < 10^{-n}$

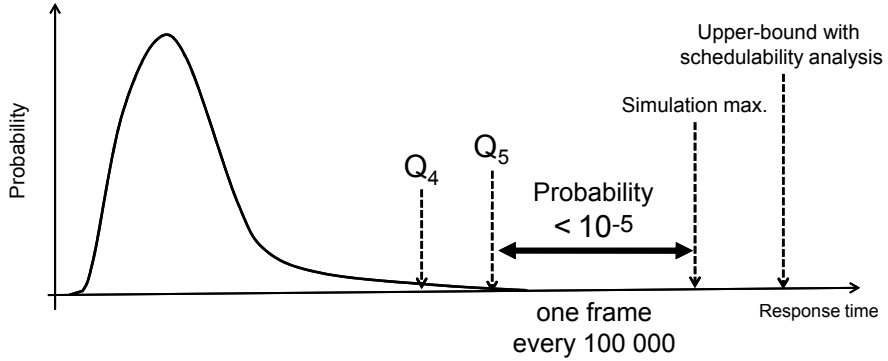

✓ Convergence unlike max → reproducibility & controllability
✓ No extrapolation here, won't help to say anything about what is too rare to be in simulation traces

1) How often performance objectives can be violated ?

Quantile	One frame every …	Mean time to failure Frame period = 10ms	Mean time to failure Frame period = 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
…	…	…	

Warning : successive failures in some cases might be temporally correlated, this must be ruled out …

2) Determine the minimum simulation length

- ✓ not obvious because non-Gaussian and possibly non i.i.d.
- ✓ time needed for quantile convergence
- ✓ reasonable # of values: a few tens …

Tool support can help here: e.g. numbers in gray should not be trusted

Reasonable values for Q5 and Q6 (with periods <500ms) are obtained in a few hours of simulation (with a high-speed simulation engine) – e.g. 2 hours for a typical automotive setup

[RTaW-sim screenshot]

Max, Q6, Q5, Q3 on our example…

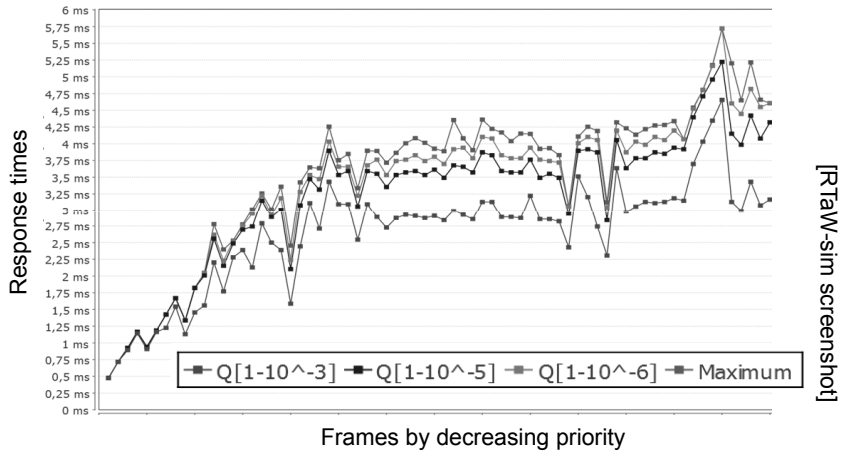

[RTaW-sim screenshot]

Concluding remarks

1 There is gap between analytic models and real (non-ideal) systems

- ✓ pessimistic at best, unsafe if assumptions not met
- ✓ no dramatic improvements in sight
- ✓"analyzability" should be a design constraint if needed

2 Simulation is a practical alternative even for critical systems .. some precautions needed

- ✓ Determine quantile wrt criticality, and simulation length wrt to quantile
- ✓ Simulator and models validation
- ✓ High-performance simulation engine needed for higher quantiles

References

References

Most available from
http://nicolas.navet.eu

[1] N. Navet, Y-Q. Song, F. Simonot, “Worst-Case Deadline Failure Probability in Real-Time Applications Distributed over CAN (Controller Area Network)“, Journal of Systems Architecture, Elsevier Science, vol. 46, n°7, 2000.

[2] 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.

[3] 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.

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

[5] 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.

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

[7] 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.

[8] M. D. Natale, “Evaluating message transmission times in Controller Area Networks without buffer preemption”, in 8th Brazilian Workshop on Real-Time Systems, 2006.

[9] 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.

[10] U. Keskin, R. Bril, and J. Lukkien, “Evaluating message transmission times in Controller Area Network (CAN) without buffer preemption revisited”, to appear in Proc. of the 9th IEEE International Workshop on Factory Communication System (WFCS 2012), May 21-24, 2012, Lemgo/Detmold, Germany.

[11] D. Khan, R. Bril, N. Navet, “Integrating Hardware Limitations in CAN Schedulability Analysis”, WiP at the 8th IEEE International Workshop on Factory Communication Systems (WFCS 2010), Nancy, France, May 2010.

[12] R. Davis, N. Navet, “Traffic Shaping to Reduce Jitter in Controller Area Network (CAN)“, ACM SIGBED Review, Volume 9, Issue 4, pp37-40, November 2012.

[13] N. Navet, H. Perrault, “CAN in Automotive Applications: a look forward”, 13th International CAN Conference, Hambach Castle, March 5-6, 2012.

[14] R.I. Davis, A. Burns, R.J. Bril, J.J. Lukkien. “Controller Area Network (CAN) Schedulability Analysis: Refuted, Revisited and Revised”, Real-Time Systems, Volume 35, Number 3, pp. 239-272, April 2007.