

Aperiodic traffic in response time analyses with adjustable safety level

Dawood A. KHAN
INRIA / INPL
Vandoeuvre-lès-Nancy, France
dawood.khan@loria.fr

Bernard BAVOUX
PSA Peugeot-Citroën
Vélizy-Villacoublay, France
bernard.bavoux@mps.com

Nicolas NAVET
INRIA / RealTime-at-Work
Vandoeuvre-lès-Nancy, France
nnavet@loria.fr

Jörn MIGGE*
RealTime-at-Work
Rome, Italy
jorn.migge@realtimeatwork.com

July 15, 2009

Abstract

In distributed real-time systems it is crucial to ensure the temporal validity of the data exchanged among the nodes. Classically, the frame Worst Case Response Time (WCRT) analyses, and the software tools which implement them, do not take into account the aperiodic traffic. One of the main reasons for this is that the aperiodic traffic is generally very difficult to characterize (i.e., the arrival patterns of the aperiodic frames). The consequence of this is that one tends to underestimate the WCRT, which may have an impact on the overall safety of the system.

In this paper, we propose a probabilistic approach to model the aperiodic traffic and integrate it into response time analysis. The approach allows the system designer to choose the safety level of the analysis based on the system's dependability requirements. Compared to existing deterministic approaches the approach leads to more realistic WCRT evaluation and thus to a better dimensioning of the hardware platform.

1 Introduction

Context of the study. In the field of real-time systems, methods to assess the real-time performances of periodic activities (tasks, messages) have been extensively studied. Response times, worst-case or average, and jitters can be evaluated by simulation or analysis for a wide range of scheduling policies provided that the activation patterns of the tasks and messages are well identified. The problem is more intricate for aperiodic activities since, in many practical cases, it is difficult to have a precise knowledge of their activation pattern and because deterministic WCRT analyses have not been conceived to handle aperiodic activities. For example, the arrival pattern of aperiodic frames in the body network of a vehicle is hard to predict, as it is dependent on the user interactions. However aperiodic frames of higher priority exchanged among the Electronic Control Units (ECUs) in the body network of a vehicle can delay periodic traffic. Indeed, most often the Controller Area

*Jörn MIGGE was with PSA Peugeot-Citroën at the time the study was undertaken.

Network (CAN) priority bus is used and the aperiodic frames do not necessarily get the lowest priority levels¹ assigned to them.

Problem definition. In this paper, we address the problem of evaluating response times when both periodic and aperiodic activities are taken into account. Activities are termed frames in rest of the paper, because the approach will be developed and illustrated on the CAN bus, but our approach equally holds for tasks. The increase in the WCRT of the periodic frames which may be caused by the higher priority aperiodic frames could be critical for hard real-time systems as it could lead to the violation of the deadlines. Besides, large response times of aperiodic frames may jeopardize the execution of a function or may even raise safety concerns in some cases (e.g. headlights flashes in a vehicle). In addition, low responsiveness is negatively perceived by the user. It is worth mentioning that activities that are periodic by essence are sometimes implemented in an aperiodic manner in order to save resources.

Whatever the exact approach, one of the main step is to derive a model of the arrival patterns for aperiodic activities, what will be called in the following the aperiodic Work Arrival Function (WAF). Then, this aperiodic WAF has to be integrated into the response time analysis. There are however difficulties:

- obtaining aperiodic data (i.e., by measurements or simulation),
- modelling aperiodic data,
- integrating the model into schedulability analysis.

What we are discussing in this paper is not how to obtain data but how to model it and integrate it into schedulability analysis.

Ways to handle aperiodic traffic and their limitations. There are two classical approaches to handle the aperiodic traffic:

- A worst-case deterministic approach : aperiodic frames are considered as periodic frames with their periods equal to the minimum inter-arrival times, this is the well known sporadic model [8]. However, in many cases, the minimum inter-arrival time is so small that the resulting workload is unrealistic, and often greater than 100% [10].
- An average-case probabilistic approach : the aperiodic traffic is modelled according to a probabilistic inter-arrivals process, the next step is then to estimate the “probable” number of arrivals in a given interval of time. This approach is clearly not suited to real-time systems because it largely underestimates the arrivals of aperiodic traffic which can occur in small time intervals².

A basic probabilistic framework was set for inclusion of aperiodic frames in a controlled manner using a threshold value in [4]. This paper builds upon this framework and discusses precisely the mechanism of deriving the aperiodic WAF, as well as it removes some assumptions placed in [4]. In particular, we show that in our specific context it is not necessary that the different streams of aperiodic frames are modelled individually.

¹Because of the incremental design process, in-house usages or constraints of the cooperation process between carmakers and suppliers, priorities on the CAN bus do not necessarily reflect the criticality of the frames (i.e., importance from a functional point of view, deadline constraint).

²According to the principle of large deviations: the smaller the interval, the larger (in proportion) the deviation to the mean [7].

Overview of our approach. We do not assume any prior knowledge of the aperiodic frame activation pattern, however we assume that it is possible to monitor the system, or a simulation model of it, and gather data about the arrival times of aperiodic frames. Then, from the measurements, we build a probabilistic model of the aperiodic inter-arrival times under the form of an empirical frequency histogram or a distribution obeying a closed-form equation whenever possible. The next step is to derive deterministic WAFs from the probability distribution of the aperiodic frame inter-arrival times. A general mechanism is provided enabling to derive the deterministic WAF from the underlying probabilistic distributions of the aperiodic traffic even given in form of empirical histograms, which is worthy in practice since aperiodic arrivals do not necessarily obey a closed-form equation. Another advantage is that the technique is independent of the scheduling and can be used whatever the policy (preemptive, non-preemptive, fixed priority, dynamic-priority, etc) and whatever the task model. All in all, we believe that our proposal offers a better solution for taking into account aperiodic traffic in systems with dependability constraints, compared to worst-case and average case probabilistic approaches.

Organization of the paper. The rest of this paper is organized as follows: the modeling of the aperiodic traffic is described in Section 2. The aperiodic work arrival function is discussed in Section 3. Section 4 describes the schedulability analysis and a small case study. Finally, the conclusion is given in Section 5.

2 Modeling aperiodic traffic

The data used in this study comes from measurements³ taken on-board of a PSA vehicle but because of confidentiality reasons we have obscured the characteristics which could reflect about the design at PSA Peugeot Citroën. In order to model the interarrival times of the aperiodic traffic, we first analyse some important structural properties of the data (e.g., linear and non-linear correlation) then find out the probability distribution that best fits our data.

2.1 Exploratory data analysis

The presence of linear and non-linear dependencies in the data would impact its modelling because it would imply a departure from the i.i.d. property (independent and identically distribution). To test these two kind of dependencies, as classically done in exploratory data analysis, we make use of some visual confirmatory tests, the “run sequence plot” and “lag plot” here, as well as the auto-correlation and BDS test (Brock, Dechert, Scheinkman, see[2]).

2.1.1 Run sequence plot

The run sequence plot displays an observed univariate data in a time sequence. It helps to detect outliers and shifts in the process. Figure 1(upper) is a run sequence plot of our data trace where the data points are indexed by their order of occurrence. The plot indicates that data does not have any long term shifts in heights over time.

³What was measured are the times at which the frames started to be transmitted and not the times at which the transmission requests were issued. Especially when the network is loaded, the two can be significantly different because of frames transmissions being delayed by higher priority frames. This could be taken into account by studying the busy periods on the bus and constructing a worst-case activation process. This is however outside the scope of the paper and will be addressed in future work. We would like to thank Reinder J. Bril for pointing this out.

2.1.2 Lag plot

A lag plot helps to gain some insight into whether a data set or time series is random or not. Random data should not exhibit any visually identifiable structure in the lag plot. Figure 1(lower) is a lag plot of our data trace (here the lag is chosen equal to 1: $x = X_{k+1}$ and $y = X_k$, where X_k is the k th observation). Since the lag plot appears to be structureless, the randomness assumption cannot be rejected.




Figure 1: Visual analysis of captured data trace. The upper graphic is a run sequence plot where the x-axis is the index of the data points and the y-axis is the time till the next aperiodic arrival expressed in seconds. In the lower graphics, a lag plot, both axes indicates the time till the next aperiodic arrival in seconds.

2.1.3 Autocorrelation analysis

The autocorrelation analysis detects the existence of serial correlations in a data trace. Precisely the correlation of order k indicates the linear relationship that may exist between data values separated by k positions. The first 100 correlation coefficients of the data trace are shown in figure 2 associated with the thresholds beyond which the values are statistically significant (1% significance level here). The graphic visualization of the correlation coefficients makes it possible to evaluate the importance and the duration of the temporal dependencies. Here, serial correlations in the aperiodic traffic are relatively limited:

- limited in frequency: on the entire aperiodic traffic, there are only 19 significant auto-correlations coefficients until a lag of 100,

- limited in intensity: the few significant auto-correlations are below 0.2 which is insufficient to be used at ends of predictions.




Figure 2: Auto-correlation of captured data trace.

These autocorrelations can probably be explained by the fact that the activation of certain functions of the vehicle requires the transmission of several consecutive frames, but, the instants of activations of the functions have small correlations. Also, the spike that can be observed around the lag 50 is likely due to a periodic frame that has not been properly filtered out in the data trace.

2.1.4 BDS analysis

Auto-correlation has the limitation that it can only test the linear dependency in the data. In order to test for non-linear dependencies a more general statistical test than the auto-correlation must be used. One such test is the BDS test [2] which employs the concept of spatial correlation from chaos theory to test the hypothesis that the values of a sequence, in this paper inter-arrival times, are independent and identically distributed (i.i.d.). Deviation from the i.i.d. case will be caused by the non-stationarity of the process (e.g., existence of trends), or the fact that there are linear or non-linear dependencies in the data.

We carried out the BDS test for various combinations of its parameters m and δ (for example for $m = 2$ and $\delta = 3$) as recommended by the authors of the test. For certain combinations we could not reject the hypothesis that the data points are i.i.d. at the 1% confidence level. The results of auto-correlation analysis and BDS test enable us to conclude that it is possible in our specific context

to model the aperiodic inter-arrival traffic by a random variable obeying a memory-less probabilistic distribution without diverging from reality.

2.2 Distribution fitting

We now need to find the probability distribution and its parameters which models the experimental data the most accurately. After having drawn aside certain possibilities for obvious reasons (for example, the normal law because its density function of density is not monotonously decreasing), we tested distributions identified by adjusting their parameters according to the principle of the maximum of likelihood (MLE). Specifically, we have successively considered the exponential law, the log-normal law and the Weibull law. The exponential law was plausible a priori taking into account the decrease of the density which one can observe in the data trace, the two other laws have been chosen for their well-known flexibility.

2.2.1 Probability plots for visual selection

The distribution of the observed data is plotted against a theoretical distribution in such a way that the points should form approximately a straight line. Departures from this straight line indicate departures from the specified distribution. If the probability plot is approximately linear, the underlying distribution is close to the theoretical distribution. What can be observed in figure 3 is that the Weibull law is the distribution that best fits the data. This visual conclusion is confirmed by statistical acceptance tests discussed in the next paragraph.




Figure 3: Probability plots for 3 candidate distributions, from top to bottom, the exponential law, the log-normal law and the Weibull Law.

2.2.2 Acceptance test

In previous section evaluation of the quality of results was done visually. In this section we use the statistical tests to verify the assumption that data trace follow a particular distribution. Specifically, we are using the χ^2 and Kolmogorov-Smirnov "goodness-of-fit" tests [6, 3]. The best results were obtained using the Weibull law, followed at some distance by the log-normal law. The conclusion of

the two tests is that one cannot reject the assumption that the data follows a Weibull distribution at a significance level of 1%. For a broad data sample collected on a real system, and not artificially generated data, it is a conclusive result.

Figure 4 presents the real data trace and an "artificial" trace generated by a Weibull law with MLE-fitted parameters. It is observed that some "patterns" present in the real trace disappear and that the simulated trace is more homogeneous in time, but overall adequacy of the modeling seems good. From the analysis, carried out in this section, we can conclude that in our specific context the Weibull distribution provides a satisfactory model for the aperiodic traffic interarrival times, followed by log-normal and exponential distributions at some distance.

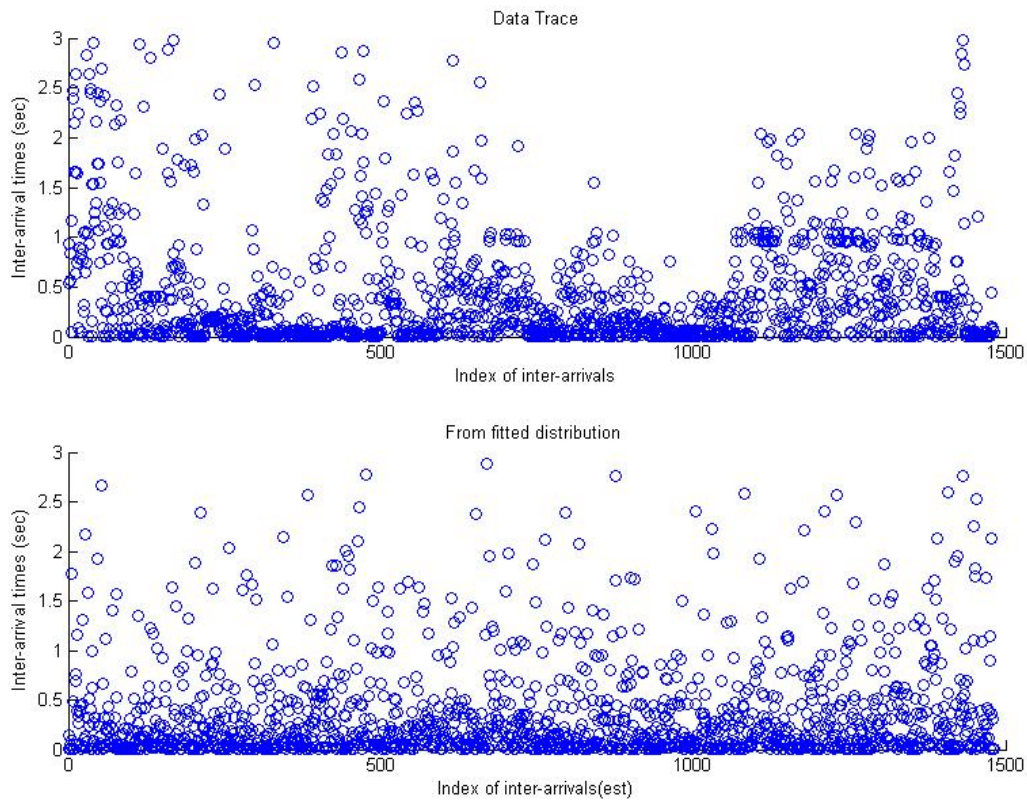


Figure 4: Comparison between the captured data trace and a random trace generated by a Weibull model with MLE-fitted parameters.

2.2.3 Parameter estimation without data trace

Because of cost and design time constraints, it is not always possible to derive the interarrival model from a real data trace, or traces of simulation. This is often the case in automobile projects. In such a situation, as an approximation, a solution is to set the parameters of the distribution based on already known parameters corresponding to another electronic architectures. In the following, we show how to adapt a Weibull⁴ model to a new intensity of the aperiodic traffic.

The expected value of a random variable obeying a Weibull law is:

⁴The case of single parameter distribution such as the exponential law is trivial, a similar approach can be used for the log-normal law.

$$E(X) = \lambda \Gamma(1 + \frac{1}{k}) \quad (1)$$

where λ is the scale parameter, k is the shape parameter of the Weibull law and the Gamma function is an extension of the factorial function to the real and complex numbers. There exist many, more or less precise, approximations to calculate the gamma function. One good approximation is given by the following formula:

$$\Gamma(z) \approx \sqrt{\frac{2\pi}{z}} \left(\frac{1}{e} \left(z + \frac{1}{12z - \frac{1}{10z}} \right) \right)^z \quad (2)$$

To adjust the expected value of the Weibull law for a new vehicle project, one simply has to change the scale parameter λ to the targeted intensity of the aperiodic traffic. In first approximation, we assume here that the shape of the distribution should not change very importantly from project to project and so set the parameter k . This assumption should be verified in the light of the analysis of additional data traces but this is left as a future work. The network load of the aperiodic traffic, denoted ρ , obeys the relation:

$$\rho = \left(\frac{1}{E(X)} \right) \cdot \bar{A} \quad (3)$$

where \bar{A} is the average transmission time of an aperiodic frame. From equations 1, 2 and 3, one obtains:

$$\lambda = \left(\frac{1}{\Gamma(1 + \frac{1}{k}) \cdot \rho} \right) \cdot \bar{A} \quad (4)$$

By replacing the values of network load, ρ , and average transmission time, \bar{A} , by the values which correspond to the automotive network that one wants to model, one obtains the new value of λ .

3 The aperiodic work arrival function $S(t)$

$S(t)$ is the aperiodic work arrival function which gives us the number of aperiodic frames in a time interval t and that will be used in the response time analysis. $S(t)$ is an increasing "staircase" function such that the "jumps" in the function correspond to the arrival of an aperiodic frame. To construct this function, we propose to discretize the time and calculate the value taken by $S(t)$ for each value of t between 1 and n where n , expressed in milliseconds, is the largest value that we may reasonably require during the computation of a response time. For example, one can set $n = 1000\text{ms}$ if the largest period of activity on the bus (*i.e.*, the largest busy period) does not exceed a second.

3.1 Safety threshold α for $S(t)$

We denote by $X(t)$ the stochastic process which counts the number of aperiodic frames in time interval t . For example, in the data trace which we studied in the preceding sections, inter-arrivals would be controlled by a Weibull law. The idea is to find the "smallest" $\hat{S}(t)$ such that the probability of $X(t)$ being larger than or equal to $\hat{S}(t)$ is lower than a threshold value α fixed by the designer. Formally, we are looking for:

$$\hat{S}(t) = \min\{S(t) \mid Pr[X(t) \geq S(t)] \leq \alpha\} \quad (5)$$

For example, if one sets $\alpha = 0.01$ it means that in no more than 1% of its trajectories the stochastic process $X(t)$ induces more aperiodic traffic than $\hat{S}(t)$. If $X(t)$ models the real aperiodic traffic accurately, the number of aperiodic frames integrated in the calculation of the response time of a periodic frame will have more than 99 percent chances to be higher than what each instance of the frame will undergo. Of course, the choice of α depends on the dependability objectives of (SIL, System Integrity Level, for example) but $\alpha = 10^{-4}$ seems a reasonable value in the context of a body network that will be considered in the experiments hereafter.

3.2 Computation of $S(t)$

We need a way to evaluate $Pr[X(t) \geq S(t)]$ at each time instant t . Two cases arise. The first case corresponds to the situation where one can calculate explicitly:

$$Pr[X(t) \geq S(t)] = \int_t^\infty f(u)du = 1 - F(t) \quad (6)$$

as it is possible for the exponential law. However, this is not analytically feasible when there is no closed-form solution to the integral. We should then either resort to numerical integration or to monte-carlo simulation. In the following, we choose to develop the later possibility because it is simple and completely independent of the underlying distribution.

3.2.1 Graphical illustration

Figure 5 illustrates the computation of $S(t)$ for a specific value of t , here $t = 5$:

$$\hat{S}(5) = \min\{S(5) \mid Pr[X(5) \geq S(5)] \leq \alpha\} \quad (7)$$

The probability $Pr[X(5) \geq S(5)]$ can be found using:

$$Pr[X(5) \geq S(5)] = \int_5^\infty f(u)du = 1 - F(5) \quad (8)$$

where $F(t)$ is the CDF and $f(u)$ is the probability density function of the distribution followed by the aperiodic frame interarrival times.




Figure 5: Graphical representation of algorithm for computation of $S(5)$. It consists in finding the smallest value of k using the CDF of the interarrival distribution according to equations 5 and 6.

3.2.2 The case of discrete distributions

In case the distribution followed by the aperiodic frame interarrival is discrete, the equation 6 can be easily evaluated by summation of probabilities. Algorithm 1 describes the process of finding $S(t)$ for discrete distributions where α is the targeted safety level, Δ is the discrete time step, θ is the set of parameters of the distribution modeling the aperiodic frame arrivals (for example, chosen according to the Maximum Likelihood Estimate principle, MLE) and T is the maximum time one may reasonably require for the computation of a response time (e.g., 1 second for an automotive network).

```

Input:  $T, \alpha, \Delta, \theta$ 
Output:  $S(t)$ : The work arrival function
foreach each  $IDX \in 0:\Delta:T$  do
     $k = 0$ ;
    while  $1 - CDF(\theta, IDX, k) > \alpha$  do
         $k = k + 1$ ;
    end
     $S(IDX) = k$ ;
end

```

Algorithm 1: Algorithm for computing $S(t)$ in the discrete case.

3.2.3 Monte-Carlo simulation

We do not always have a discrete distribution modelling the data nor a continuous distribution such that equation 5 can be evaluated analytically. We need an alternate method to evaluate equation 6 in such cases. This can be done with numerical integration techniques or using Monte Carlo simulation method. The latter approach is described in algorithm 2 where α is the safety level, Δ is the discrete time step, θ is the set of parameters of the aperiodic frame arrival distribution, T is the time horizon,

N is the number of random samples⁵ to be drawn for the Monte-Carlo simulation. Basically, $S(t)$ is computed for each time unit by drawing N values from the probabilistic distribution modelling the aperiodic frame arrival process and checking if the accumulated probability value smaller than the probability value for which we are evaluating $S(t)$.

```

Input:   $T, \alpha, \Delta, \theta, N$ 
Output:  $S(t)$ : The work arrival function
 $index = 0$ ;
 $Data = random(\theta, N)$ ;
foreach each  $IDX \in 0:\Delta:T$  do
    foreach  $i=1:N$  do
         $AccTime = 0$ ;
         $k = 0$ ;
        while  $AccTime < IDX$  do
             $AccTime = AccTime + Data[index]$ ;
             $index = index + 1$ ;
             $k = k + 1$ ;
        end
    end
     $S(IDX) = k$ ;
end

```

Algorithm 2: Deriving $S(t)$ by Monte Carlo Simulation.

As an illustration of the approach, we derived $S(t)$ in the cases where the aperiodic interarrival distribution obeys 1) an exponential law 2) a Weibull law 3) a log-normal law. The number of random draws of the Monte-Carlo simulations (parameter N in algorithm 2) is set to 5 million for each distribution. For all three distributions, the parameters are fitted using MLE against the data traces and the three distributions lead to the same average intensity. What can be observed is that the distribution, and not only the average intensity of the aperiodic traffic, plays a major role in the shape and height of the aperiodic WAF.




Figure 6: Aperiodic work arrival function up to time $t = 125ms$ for various interarrival distributions.

⁵Central Limit Theorem tells us that the convergence rate is of order $N^{1/2}$ where N is the number of random draws, which means that adding one significant digit requires increasing N by a factor 100. The value of N should be set depending on the threshold α and accuracy objectives.

4 Experimentation

4.1 Basic schedulability analysis

Classically, schedulability analyses for real-time communication networks assume periodic or sporadic streams of frames [9, 5]. In this paper, for the sake of simplicity, we make use of a sufficient but not necessary schedulability test⁶ presented in [5] as the framework to include aperiodic WAF into the schedulability analysis. However, the approach would remain similar with the sufficient and necessary test proposed in the aforementioned paper.

In the following, we re-use the concepts and notations from [5]. The worst-case response time of frame m is made up of several elements:

1. An upper bound on the queuing jitter J_m ,
2. The longest transmission time C_m ,
3. The waiting delay w_m at the sending end, that is the longest time that the frame can wait before it starts being successfully transmitted (i.e., before it wins the arbitration on the CAN bus). This delay is given by equation 10,

The waiting delay w_m includes the interference due to the aperiodic frames of higher priority than m , which is given by the function $N_m^{\alpha, \mathcal{M}}(t)$ defined as follow:

$$N_m^{\alpha, \mathcal{M}}(t) = S_{\mathcal{M}}^{\alpha}(t) \cdot \max_{j \in HpAf(m)} C_j \quad (9)$$

where \mathcal{M} is the aperiodic interarrival model, α the chosen safety threshold, $S_{\mathcal{M}}^{\alpha}(t)$ the corresponding aperiodic WAF and $HpAf(m)$ is the set of aperiodic frames having higher priority than frame m . It has to be pointed out that the definition of $N_m^{\alpha, \mathcal{M}}(t)$ makes the conservative assumptions that all aperiodic frames are of higher priority than m . How this can be improved is left for future work (see the discussion in the conclusion).

As classically done, the waiting delay w_m can be determined with the following recurrence relation:

$$\begin{aligned} w_m^{n+1} &= N_m^{\alpha, \mathcal{M}}(w_m^n) + \max(B_m, C_m) \\ &+ \sum_{\forall k \in hp(m)} \lceil \frac{w_m^n + J_k + \tau_{bit}}{T_k} \rceil C_k \end{aligned} \quad (10)$$

where $hp(m)$ is the set of frames with priority higher than m , and $\max(B_m, C_m)$ corresponds to the longest possible time for which an invocation of frame m can be blocked either by lower priority messages or due to the previous invocation of the same frame. The recurrence relation goes on until $J_m + w_m^{n+1} + C_m > D_m$ or $w_m^{n+1} = w_m^n$. In the former case, the frame is not schedulable while in the latter case the worst-case response time of the frame is given by:

$$R_m = J_m + w_m + C_m \quad (11)$$

⁶This test is applicable when deadlines do not exceed their periods.

4.2 Case study




Figure 7: Aperiodic WAFs for inclusion in WCRT computations up to time 250ms. The y-axis starts from value 1.

In this section, we illustrate the analysis on an typical 125Kbit/s automotive body network. To overcome the confidentiality issue that prevent us from giving the details of real sets of messages, we used *Netcarbench*[1], a GPL-licensed software that generates sets of messages according to parameters defined by the user (network load, number of ECUs, distribution of the periods of the frames, etc.). The *Netcarbench* input file used in this experiment is available upon request. The generated periodic message sets under study consists of 120 CAN frames with deadlines equal to periods and data payload ranging from 1 to 8 bytes. The total periodic load is equal to 41.4% while the aperiodic network load is set to 3% with aperiodic frames containing at most 7 bytes of data. The aperiodic WAFs corresponding to this experimental setup for two interarrival distributions and a safety threshold α equal to 10^{-4} are shown in figure 7.

Figure 8 shows the worst-case response times of the highest priority periodic frames for the two aperiodic WAFs considered here and a safety threshold α equal to 10^{-4} . The WCRT of the frames are computed with the software *NETCAR-Analyzer* from *RealTime-at-Work* whose purpose is to analyse the performances of CAN-based communication systems and optimize their design and configuration (e.g., choices for the message priorities and offsets, waiting queue policy and length, etc). Even in this context where the periodic load is moderate (41.4%) and the aperiodic traffic is limited (3%), one observes that aperiodic traffic rather significantly impacts the worst-case response times of the periodic frames. For instance, the WCRT for the frame with id 96 raises from 96.1ms without aperiodic traffic to 116.3ms with 3% aperiodic traffic obeying the Weibull law (+21%). Besides, we see from figure 8 that the exact model of the aperiodic traffic plays some role and thus cannot be overlooked.

5 Conclusion

In this paper, we develop a new approach for integrating the aperiodic traffic in response time analyses. The main interest of the proposal is that the overestimation of the aperiodic traffic is kept to the




Figure 8: Worst-case response times (WCRT) after inclusion of 3% of aperiodic traffic with different aperiodic frame interarrival models. The frames (x-axis) are sorted by decreasing priorities. The steep increase of the WCRT at time 50 and 100ms can be explained because, in that particular message set, there are several messages with periods equal to 50 and 100ms (resp. 4 and 13), and, at these points, new instances of these frames are causing additional delays.

minimum that still enables the system to meet some chosen dependability requirements.

However, the resulting response time estimation can be pessimistic especially for high priority frames when there is a large volume of aperiodic traffic, as we have assumed that all aperiodic frames are of higher priority. It is possible to be less pessimistic by modelling each aperiodic stream individually (or consider several distinct groups of aperiodic streams) and integrate only the higher priority aperiodic WAFs into the schedulability analysis. However, we believe that this more fine-grained approach will not be always practical since it requires significant modelling efforts and large quantity of data traces. We are working on a more global scheme that would still respect the safety threshold while being as accurate as possible (*i.e.*, discard as much as possible of the lower priority aperiodic traffic).

Some of the results shown here hold under the assumption that the aperiodic interarrivals are independent and identically distributed (i.i.d.). In practice, this assumption can be easily tested using statistical tests such as the BDS test statistics but it is clear that it may not hold for all kinds of systems and workloads. It would be interesting to study, for instance by simulation, how departure from the i.i.d. property impacts the accuracy of the results. How non-i.i.d. streams can be handled is another follow-up to this study.

Acknowledgements The authors would like to thank the anonymous reviewers and Reinder J. Brill (Eindhoven University of Technology) for their useful comments.

References

- [1] C. Braun, L. Havet, and N. Navet. NETCARBENCH: A benchmark for techniques and tools used in the design of automotive communication systems. In *7th IFAC International Conference on Fieldbuses and Networks in Industrial and Embedded Systems*, pages 321–328, 2007.
- [2] 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.
- [3] B. Brumback and M. Srinath. A Chi-Square test for fault-detection in Kalman filters. *Automatic Control, IEEE Transactions on*, 32(6):552–554, Jun 1987.
- [4] A. Burns, G. Bernat, and I. Broster. A probabilistic framework for schedulability analysis. In *Proceedings of the Third International Conference on Embedded Software (EMSOFT 2003)*, pages 1–15, 2003.
- [5] R.I. Davis, A. Burn, R.J. Bril, and J.J. Lukkien. Controller Area Network (CAN) schedulability analysis: Refuted, revisited and revised. *Real-Time Systems*, 35:239–272, 2007.
- [6] J. Millard and L. Kurz. The Kolmogorov-Smirnov tests in signal detection (corresp.). *IEEE Transactions on Information Theory*, 13(2):341–342, Apr 1967.
- [7] N. Navet, L. Cucu, and R. Schott. Probabilistic estimation of response times through large deviations. In *WiP of 28th IEEE Real-Time Systems Symposium (RTSS'2007 WiP)*. IEEE, 2007.
- [8] M. Spuri and G. Buttazzo. Scheduling aperiodic tasks in dynamic priority systems. *Real-Time Systems*, 10:179–210, 1996.
- [9] K. Tindell, A. Burns, and A.J. Wellings. Calculating Controller Area Network (CAN) message response times. *Control Engineering Practice*, 3(8):1163 – 1169, 1995.
- [10] Y. Zhang, D.K. Kreckler, C. Gill, C. Lu, and G.H. Thaker. Practical schedulability analysis for generalized sporadic tasks in distributed real-time systems. *Real-Time Systems*, 0:223–232, July 2008.