Time-Exponentially Weighted Moving Histograms (TEWMH) for Self-Adaptive Systems

Michael Menth, Jens Milbrandt, and Jan Junker
Department of Distributed Systems, Institute of Computer Science, University of Würzburg, Germany
Email: {menth,milbrandt,junker}@informatik.uni-wuerzburg.de

***Abstract*— The distribution of a stationary point process can be sampled by an ordinary histogram. If the distribution of the process varies over time, a static histogram still yields results that are averaged over time since the beginning of the data collection. In this paper, we propose the time-exponentially weighted moving histogram (TEWMH) to derive an estimate for the time-dependent distribution of an instationary point process. The importance of the samples decays exponentially over time such that young samples contribute more to the empirical distribution than old ones. The strength of the decay can be controlled by a simple parameter which determines the memory of the histogram. We present a simple implementation of the TEWMH such that this mechanism can be well applied in practice. The empirical distribution serves for the derivation of other time-dependent statistical measures such as time-dependent percentiles of the observed random variable. These provide useful feedback in self-controlled adaptive systems. We illustrate the application of the TEWMH for experience-based admission control (EBAC) and show its benefits.**

Keywords: Measurement Techniques, Quality and Performance in Autonomic Systems

I. Introduction

Histograms collect data of an observed random variable X to obtain an estimate of the distribution of X. They allow the derivation of statistical measures like the mean, the standard deviation, or various percentiles. These measurement data may be used for the calibration of adaptive systems. Technical systems may be even self-adaptive if they adapt autonomously to new conditions, i.e., they observe their environment and adapt to changes. The observation of changes is problematic with conventional histograms if they collect data since the start of the system: the more data are already collected, the more data with a different characteristic are needed to effect a significant change in the distribution. Thus, it is desirable to give more importance to young samples and less to old ones in order to propagate the properties of the latest samples faster into the distribution obtained by the histogram.

In this paper, we propose the time-exponentially weighted moving histogram (TEWMH). The importance of the collected samples decays exponentially over time and the strength of the devaluation of old samples is controlled by a simple memory parameter. Its implementation is simple such that it can be well used in technical systems. We further illustrate the application of the TEWMH for experience-based admission control (EBAC) which is an example for a self-adaptive system. Admission control admits or rejects flows for high quality transport over a link or through a network. The decision is based on the bandwidth requested by the flows and on the bandwidth of the overall admitted flows. If the flows use only half of their requested and reserved rate, AC can admit about the double amount of flows, i.e., the available bandwidth is safely overbooked. To that end, EBAC collects samples of the utilization of the overall flow reservations. It derives the inverse of the 99%-quantile and takes the inverse as the overbooking factor. The quantile leads to a more conservative overbooking factor than the mean. This design decision has been made to avoid congestion due to statistical fluctuations. If a conventional static histogram is used for the implementation, EBAC cannot adapt quickly to new traffic properties while the TEWMH allows for a fast reaction regardless of how many samples have already been collected by the histogram.

The paper is structured as follows. Section II adapts the time-exponentially weighted moving average (TEWMA) for rates from [1] to averages. It is similar to the TEWMH, however, it is less complex. Section III explains possible implementation of static and dynamic histograms and it presents the TEWMH. Section IV introduces the fundamentals of EBAC, illustrates the application of the TEWMH, and points out its benefit. Finally, we summarize this work and draw our conclusions in Section V.

II. Time-Exponentially Weighted Moving Average (TEWMA)

Similar but simpler than the derivation of the current distribution of a random variable X is the estimation of its current mean. In [1], four different methods have been presented to calculate time-dependent rates. We adapt them for the computation of time-dependent means and discuss their pros and cons.

A. Stochastic Point Processes

In the following, we observe a stochastic point process X_i, $i \in \mathbb{N}_0$. Arrivals of size X_i happen at times $t_i \in \mathbb{R}_0$, $i \in \mathbb{N}_0$. It describes, e.g., the packet sizes and the temporal structure of a packet arrival process. If the process is stationary, the distribution for all X_i are identical. In practice, however, we might observe that the X_i depend on the daytime. For instance,

This work was funded by Siemens AG, Munich. The authors alone are responsible for the content of the paper.

the traffic volume in a network varies over time and other aggregate properties like the traffic composition are also time-dependent.

B. Static Mean

The static mean is calculated by $E[X] = \frac{1}{n}\sum_{0\leq i<n} X_i$. The drawback of this equation is that all samples from the past are considered to calculate the mean. Therefore, if the distribution of the X_i change, the obtained mean value hardly changes if the number n of already collected samples is large. Thus, the calculation of the conventional definition of the mean does not allow to track temporary changes of the mean value $E[X]$.

C. Time-Dependent Average Based on Discrete Intervals (Avg-DI)

To calculate a time-dependent average based on discrete intervals (Avg-DI), the time axis is slotted into equidistant intervals of length Δt. We define the set of indices with all arrivals in the j-th interval by $\mathcal{I}_j = \{i : j\cdot\Delta t \leq t_i < (j+1)\cdot\Delta t\}$. The empirical mean for interval j is then calculated by $E[X]_j = \frac{1}{|\mathcal{I}_j|}\sum_{i\in\mathcal{I}_j} X_i$. The value $E[X]_j$ is taken as an estimate for the time-dependent average during the $(j+1)$-th interval. We define the memory M of time-dependent means by the time a sample contributes to the resulting average. In the case of Avg-DI, a sample contributes exactly $M = \Delta t$. Avg-DI suffers mainly from two drawbacks. If the interval length Δt is short, the statistical base for the computation of $E[X]_j$ is extremely weak. If Δt is long, the average values are late or even obsolete since only the average values for the last past interval are available.

D. Exponentially Weighted Moving Average (EWMA)

The timeliness of the mean values can be improved by the exponentially weighted moving average (EWMA) without disregarding the past samples. It starts with $E[X]_0 = X_0$ and calculates the succeeding values recursively by $E[X]_i = w\cdot E[X]_{i-1} + (1-w)\cdot X_i$. The average of the last arrival instant is valid until the next one. In contrast to Avg-DI, samples always contribute to the EWMA but their strength is devaluated by w whenever a new sample arrives. Thus, the memory depends on the temporal structure of the process. The average memory of the EWMA is $M = \frac{1}{\lambda}\cdot(1-w)\cdot\sum_{i=0}^{\infty} w^i = \frac{1}{\lambda}$ whereby λ is the interarrival rate of the process. Thus, the memory depends on the arrival rate only. The weight parameter w just controls the averaging strength of the EWMA, i.e., large w lead to slowly changing results and small w lead to quickly changing results. The EWMA comes with a semantic problem. It disregards the interarrival time between the samples, i.e., the impact of a new sample on the new mean is the same whether the last calculation of the EWMA has been long time ago or only recently.

The EWMA was introduced by [2] and this mechanism has been studied quite intensively especially in the field of economics for chart analysis [3]–[8]. The EWMA is also used in many technical documents of the IETF [9], [10], the most prominent one is probably the obsolete estimation of the round trip time for TCP in [11].

E. EWMA Based on Discrete Intervals (EWMA-DI)

EWMA-DI is a combination of Avg-DI and EWMA: it applies the EWMA to the time-dependent averages obtained from Avg-DI. The corresponding memory is $M = \Delta t\cdot(1-w)\cdot\sum_{i=0}^{\infty} w^i = \Delta t$ which is now independent of the arrival rate λ. However, this method suffers from the same disadvantages as Avg-DI and its averaging strength depends on the choice of w. In addition, samples from an interval with only a few arrivals have more impact on the mean than others.

F. EWMA Based on EWMA-Based Sums (EWMA-ES)

The static mean $E[X] = \frac{S[X](t)}{N(t)}$ is calculated as the sum of all already arrived sample sizes $S[X](t)$ divided by the sum of all arrivals $N(t)$ (counter). We enhance this concept by devaluating $S[X](t)$ and $N(t)$ in regular intervals Δt by a parameter $0 < w < 1$. This resembles a combination of the static mean and EWMA-DI. EWMA-ES combines the positive properties of EWMA-DI and the static mean as it is time-dependent and reacts immediately when new samples are collected. The average memory of the EWMA-ES is $M = \Delta t\cdot\sum_{i=0}^{\infty} w^i - \frac{\Delta t}{2} = \Delta t\cdot(\frac{1}{1-w} - \frac{1}{2})$. The design and the memory analysis of the EWMA-ES look cumbersome, but the EWMA-ES serves as comprehensible predecessor for the TEWMA since the limit for $\Delta t \rightarrow 0$ of EWMA-ES with a constant memory leads to the TEWMA.

G. Time-Exponentially Weighted Moving Average (TEWMA)

The time-exponentially weighted moving average (TEWMA) is an elegant enhancement of the EWMA-ES. It is defined by $E[X](t) = \frac{S[X](t)}{N(t)}$ with $S[X](t_0) = X_0$ and $N(t_0) = 1$. The sums $S[X](t)$ and $N(t)$ are updated by

$$S[X](t_i) = S[X](t_{i-1})\cdot e^{-\gamma\cdot(t_i - t_{i-1})} + X_i \quad (1)$$

$$N(t_i) = N(t_{i-1})\cdot e^{-\gamma\cdot(t_i - t_{i-1})} + 1 \quad (2)$$

whenever a new arrival occurs. The memory of the TEWMA is exactly $M = \int_0^{\infty} e^{-\gamma\cdot t}dt = \frac{1}{\gamma}$ and an exact half-life period $T_H = \frac{ln(2)}{\gamma}$ of the sample can be derived by the equation $\frac{1}{2} = e^{-\gamma\cdot T_H}$. Like EWMA-ES, TEWMA improves the timeliness of the calculated mean without disregarding the temporal structure of the process and in addition, it is independent of the sampling time interval Δt.

H. Comparison of the Averaging Methods

We illustrate the impact of the temporal structure of the process on the time-dependent average for some of the above presented concepts. We consider four different but very similar realizations of a single stochastic process that produce two samples within the first three potential arrival instants $t \in \{1,2,3\}$. Thus, the arrival rate for that interval is $\lambda = \frac{2}{3}$ which determines the average memory of the EWMA. To be comparable from a memory point of view, we set $\Delta t = \frac{1}{\lambda}$ for the interval-based mean and $\gamma = \lambda$ for the TEWMA.

Table I shows that the static mean is unaware of the order and temporal structure of the stochastic process. Avg-DI at time instant $t = 2$ is based on the average of the samples at

$t=0$ and $t=1$ and, therefore, it yields delayed results. The results of the EWMA depend significantly on the value of w and are insensitive to the temporal structure of the process. The TEWMA produces timely results and respects both the trend and the temporal structure of the process, i.e., it matters whether a sample arrives at $t=0$ or at $t=1$.

TABLE I

COMPARISON OF AVERAGING METHODS WITH REGARD TO THE TEMPORAL STRUCTURE OF THE MEASURED PROCESS.

$X(0)$	$X(1)$	$X(2)$	Static mean	Avg-DI	EWMA $w=\frac{2}{3}$	EWMA $w=\frac{1}{3}$	TEWMA $\gamma=\lambda$
4	–	1	2.5	4	3	2	1.63
–	4	1	2.5	4	3	2	2.01
–	1	4	2.5	1	2	3	2.99
1	–	4	2.5	1	2	3	3.37

III. Time-Exponentially Weighted Moving Histogram (TEWMH)

In this section, we discuss the implementation of time-dependent histograms according to the concepts of time-dependent averages. Some of these concepts are trivial, others are new. In particular, the new TEWMH is a non-trivial adaptation of the TEWMA concept to histograms.

A. Static Histograms

A histogram discretizes a certain value range $[v_{low}, v_{high}]$ into n_{bins} equidistant subintervals numbered from 0 to $n_{bins}-1$, the so-called bins. Each of them is associated with a counter c_j, $0 \leq j < n_{bins}$. A random variable X_i is collected in the histogram by incrementing the counter for bin $\lfloor (X_i - v_{low}) \backslash (\frac{v_{high}-v_{low}}{n_{bins}}) \rfloor$ by 1 and by incrementing the total number of hits n_{hits} by 1, too. If the random variable X_i lies outside the considered interval, counter 0 or $n_{bins}-1$ is incremented. All collected values contribute equally to the relative frequency $h_j = \frac{c_j}{n_{hits}}$ of their corresponding bins. Therefore, the static histogram is not able to represent the latest trends of the observed process in the empirical distribution.

B. Time-Dependent Histogram Based on Discrete Intervals (Hist-DI)

Time-dependent histograms based on discrete intervals (Hist-DI) are analogous to time-dependent averages based on discrete intervals (Avg-DI). The data are collected during an interval of length Δt and during that time, the relative frequencies from the previous interval are taken as results. This approach comes with the same disadvantages as Avg-DI: for small intervals the data lack statistical significance and for large intervals the results are late.

C. Exponentially Weighted Moving Histograms (EWMH)

The exponentially weighted moving histogram (EWMH) starts with an increment of 1 for the counter c_j corresponding to the first sample to assure that the sum of the counters is 1 which is an invariant of this method. It continues with increments of $(1-w)$ combined with a devaluation of each counter by w whenever a random variable is observed. The counters c_j contain directly the relative frequencies, therefore, a very large w is recommended. The EWMH inherits the disadvantages of the EWMA, e.g., the memory of this approach depends on the arrival rate λ of the process. Moreover, the EWMH requires very large computation overhead since all counters must be devaluated whenever a new observation is made.

D. EWMH Based on Discrete Intervals (EWMH-DI)

Similarly to EWMA-DI, EWMH-DI is a combination of Hist-DI and EWMA: it applies the EWMA process to the counter values c_j obtained from EWMA-DI. EWMH-DI inherits its memory from EWMA-DI and suffers from the same shortcomings.

E. EWMH Based on EWMA-Based Counters (EWMH-EC)

The EWMH-EC works basically like the EWMA-ES. It increments the corresponding counters by 1 whenever a random variable is observed and devaluates all counters c_j and n_{hits} of the static histogram in regular intervals of length Δt. The EWMH-EC is simpler to implement than Hist-DI and EWMH-DI because it requires only a single data structure, and it needs less computation power than the pure EWMH since counters are devaluated only after Δt time. The EWMH-EC inherits its memory from the EWMA-ES. The limit for $\Delta t \rightarrow 0$ of the EWMH-EC with a constant memory leads to the TEWMH presented in the next section.

F. Time-Exponentially Weighted Moving Histograms (TEWMH)

The TEWMH is basically derived from the TEWMA. Whenever a new random variable X_i is observed, the corresponding counter c_j and n_{hits} are set to

$$c_j(t_i) = c_j(t_{i-1}) \cdot e^{-\gamma \cdot (t_i - t_{i-1})} + 1 \quad (3)$$

$$n_{hits}(t_i) = n_{hits}(t_{i-1}) \cdot e^{-\gamma \cdot (t_i - t_{i-1})} + 1 \quad (4)$$

$$c_{n \neq j}(t_i) = c_{n \neq j}(t_{i-1}) \cdot e^{-\gamma \cdot (t_i - t_{i-1})}, \quad (5)$$

i.e., the other counters $c_{n \neq j}$ are just devaluated by $e^{-\gamma \cdot (t_i - t_{i-1})}$. This makes the approach as computationally expensive as the pure EWMH method since all counters are updated upon a new observation. We get rid of scaling down the other counters by leaving them untouched and scaling up the increments in Equation (3) and 4 instead. Thus, only two counters need to be incremented in case of a new observation:

$$c_j(t_i) = c_j(t_{i-1}) + e^{\gamma \cdot (t_i - t_R)} \quad (6)$$

$$n_{hits}(t_i) = n_{hits}(t_{i-1}) + e^{\gamma \cdot (t_i - t_R)} \quad (7)$$

taking the last reset instant t_R into account. When the number of hits $n_{hits}(t_i)$ becomes too large, the counters are reset by $c_j(t_i) = \frac{c_j(t_i)}{n_{hits}(t_i)}$ and $n_{hits}(t_i) = 1$, and the new reset time $t_R = t_i$ is stored.

G. Derivation of Percentiles

The percentile or quantile of a distribution regarding a random variable X is defined by

$$X_p = \min\left(x : P(X \leq x) \geq p\right). \quad (8)$$

An estimation of the time-dependent quantile $X_p(t)$ can be derived from the TEWMH or EWMH-EC by

$$X_p(t) = \min\left(v_{low} \cdot j : \sum_{0 \leq i < j} c_i(t) \geq p \cdot n_{hits}(t)\right). \quad (9)$$

Thus, the relative frequency of the smallest bins is summed up such that their sum is equal to or larger than p. The lower bound of the next largest bin yields the desired percentile value.

IV. Application of TEWMH to Experience-Based Admission Control (EBAC)

In this section, we show the application of the TEWMH in the context of experience-based admission control (EBAC). We briefly review the concept of EBAC that we first presented in [12]. We calibrated its basic parameters in [13] and investigated its behavior in the presence of traffic changes in [14]. In this paper, we illustrate the impact of the histogram type on the reaction time of EBAC and show the benefit of TEWMH.

A. Basic EBAC Mechanism

The idea of EBAC is briefly described as follows. An admission control (AC) entity limits the access to a link l with capacity $c(l)$ and records all admitted flows $f \in \mathcal{F}(t)$ at any time t together with their requested peak rates $\{r(f) : f \in \mathcal{F}(t)\}$. When a new flow f_{new} arrives, it requests a reservation for its peak rate $r(f_{new})$. If

$$r(f_{new}) + \sum_{f \in \mathcal{F}(t)} r(f) \quad \leq \quad c(l) \cdot \varphi(t) \cdot \rho_{max} \quad (10)$$

holds, admission is granted and f_{new} joins $\mathcal{F}(t)$. If flows terminate, they are removed from $\mathcal{F}(t)$. The experience-based overbooking factor $\varphi(t)$ is calculated by statistical analysis and indicates how much more bandwidth than $c(l)$ can be safely allocated for reservations. The maximum link utilization threshold ρ_{max} limits the traffic admission such that the expected packet delay W exceeds a maximum delay threshold W_{max} only with probability p_W.

The reserved bandwidth of all flows is $R(t) = \sum_{f \in \mathcal{F}(t)} r(f)$ while $C(t)$ denotes the unknown mean rate of the traffic aggregate $\mathcal{F}(t)$. The intention of EBAC is to derive a suitable overbooking factor $\varphi(t)$ that takes advantage of the peak-to-mean-rate ratio (PMRR) $K(t) = \frac{R(t)}{C(t)}$ of the traffic aggregate. EBAC makes traffic measurements $M(t)$ at an appropriate time scale on the link and samples the reservation utilization $U(t) = M(t)/R(t)$ by a histogram [13] to derive the p-percentile $U_p(t)$ of the empirical distribution of U. Its reciprocal yields the time-dependent overbooking factor $\varphi(t) = 1/U_p(t)$ as the contents of the histogram depends on the current time t.

B. Impact of the Histogram Type on the EBAC Reaction Time

In [13], EBAC worked well when we used static histograms for the measurement of the utilization of the reservations since the flows had constant traffic properties and the statistical properties of the collected utilization values did not change over time. However, in the presence of traffic changes, the percentile of the utilization must quickly reflect the new conditions, thus, the applied histogram must be time-dependent. To keep things simple, we study the behavior of EBAC on a single link with a capacity of 10 Mbit/s. As AC only reacts to avoid congestion, we operate the link under saturated conditions, i.e., we trigger so many flow requests that many of them are rejected.

In [14] we investigated EBAC for very fast traffic changes, i.e., the PMRR $K(t)$ of the traffic aggregate suddenly increased or decreased. A PMRR increase means that less traffic is sent than before and a PMRR decrease denotes more traffic from the admitted flows. The adaptation of the overbooking factor $\varphi(t)$ works very fast when the PMRR $K(t)$ suddenly decreases, but is is relatively slow when it increases. Very fast decreases of PMRR may result from coordinated QoS attacks of all admitted traffic sources, which is an extreme and also not very realistic scenario.

In this paper, we investigate the impact of a slow increase of the PMRR $K(t)$ on the reaction time of the EBAC. We have almost the same experiment settings like in [14] where also the detailed traffic model is described. We simulate the EBAC with the following time-dependent parameters. Until simulation time $t = 230$ s, new flows have a PMRR of $k(f) = 2$ and afterwards new flows have a PMRR of $k(f) = 4$. We perform this experiment 50 times and present averaged results in Figures 1(a) - 1(d). The figures show the link bandwidth $c(l)$, the reserved rate $R(t)$, the measured rate $M(t)$, the resulting overbooking factor $\varphi(t)$ and the effective PMRR $K(t)$ of the admitted flows. Figure 1(a) shows that the overbooking factor $\varphi(t)$ approximates the PMRR $K(t)$ quite well as long as $K(t)$ is constant. As soon as the PMRR of the entire aggregate $K(t)$ increases since old flows with $k(f) = 2$ terminate and new flows with $k(f) = 4$ are admitted, the reservation utilization $U(t)$ decreases, but a significant amount of samples is required that the quantile $U_p(t)$ decreases, too. This delays the increase of the overbooking factor $\varphi(t)$. The reserved capacity grows with the overbooking factor because we assumed a sufficiently high request rate such that free capacity can be used whenever available by EBAC. Since $\varphi(t)$ is underestimated compared to $K(t)$, the measured traffic rate $M(t)$ drops under 7.5 Mbit/s for a while. The results in Figure 1(a) belong to EBAC implementations with EWMH-EC with a memory of $M = 14$ s and a devaluation interval of $\Delta t = 10$ s. This produces the same curves like the TEWMH with the same memory. Figures 1(b) - 1(d) show the results for time-dependent histograms of type EWMH-EC with the same memory of $M = 14$ s but longer devaluation intervals of $\Delta t = 110, 210, 310$ s. The resource utilization for large values of Δt is lower than for small ones which is disadvantageous in a situation where traffic is

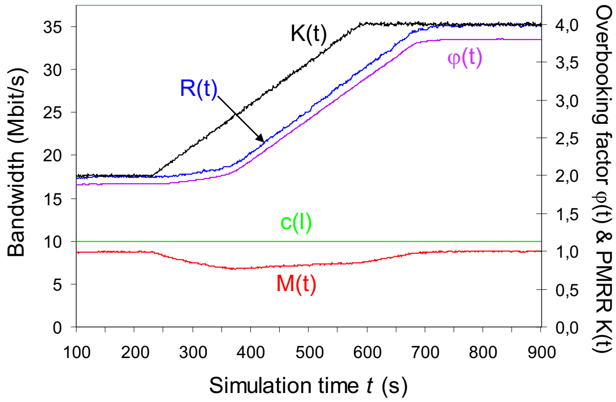

(a) $\Delta t = 10$ s.

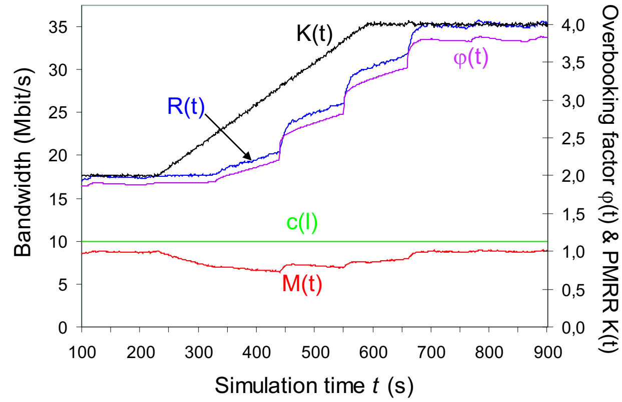

(b) $\Delta t = 110$ s.

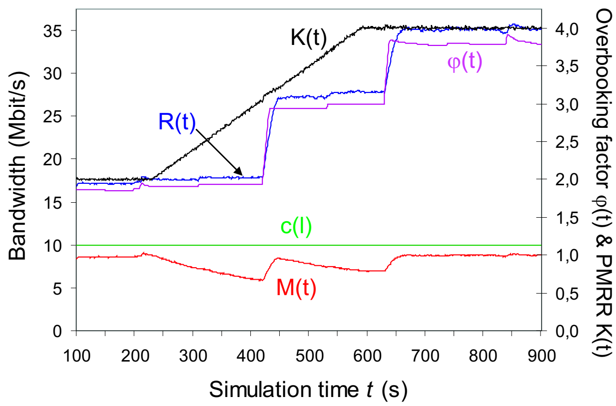

(c) $\Delta t = 210$ s.

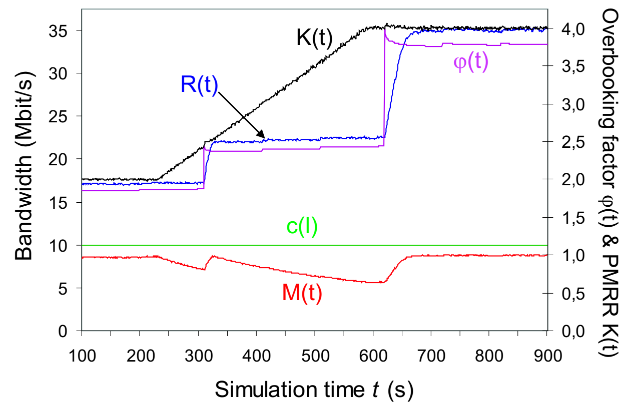

(d) $\Delta t = 310$ s.

Fig. 1. Impact of length Δt of the devaluation interval for EWMH-EC histograms with a memory of $M = 14$ s on the reaction time of EBAC if the peak-to-mean rate ratio of new flows suddenly changes from $k(f) = 2$ to $k(f) = 4$ at time $t = 230$ s.

blocked. The curves for $\varphi(t)$ and the accompanying $R(t)$ tend to have steps whose corners approach $K(t)$. The size of the steps clearly correlates with the devaluation interval because the percentile $U_p(t)$ increases when the counters with high utilization values in the histogram are devaluated. At these time instants, there is no safety margin between $\varphi(t)$ and $K(t)$ anymore, which might lead to QoS violations. Histograms with large values of Δt have a very small devaluation factor w for a short memory of $M = 14$ s such that the overall sum $n_{hits}(t)$ in the histogram becomes extremely small. As a consequence, sufficiently many new samples are required to produce a good estimate for the quantile $U_p(t)$. We performed the same experiments with the TEWMH, too, with different thresholds to normalize the counters. They all lead to the same smooth results as in Figure 1(a) but without updating all counters within relatively short intervals of $\Delta t = 10$ s.

After all, the presented curves show that the EWMH-EC is well feasible, but it must be carefully parameterized, i.e., its devaluation interval Δt must not be chosen too long compared to its memory M. Very short intervals increase the computation overhead for the devaluation of the counters. The TEWMH is simpler since it does not require any other parameters besides the length of its memory γ. Its quantile reacts rather quickly to changes of the traffic properties compared to the one from the EWMH-EC with large Δt and improves thereby the timeliness of the histogram without sacrificing the statistical significance of its values.

V. Conclusion

The conventional calculation of average values and the conventional evaluation of histograms to derive empirical distributions respect all samples equally. As a consequence, their results do not react quickly to changes of the statistical properties of the observed random variable. In earlier work, we proposed several methods to derive time-dependent rates of a point process $X(t)$. In this paper, we adapted these concepts to calculate time-dependent averages. We extended them to time-dependent histograms to obtain time-dependent empirical distributions that allow, e.g., the derivation of time-dependent percentiles of X. The most elegant method is the time-exponentially weighted moving histogram (TEWMH). In contrast to other methods, it provides timely results since the impact of collected samples is immediately seen in the empirical distribution. It has a relatively low computational overhead and it is simple to use as it relies only on the single parameter γ that controls the length of its memory. To illustrate the usefulness of the TEWMH, we used it to make experience-based admission control (EBAC) self-adaptive to changes of the traffic properties. Our experiments showed that the TEWMH leads to stable results while other time-dependent histogram types can lead to artifacts in the system that are caused by additional implementation parameters of the these histograms.

Due to its simplicity regarding both its implementation and parametrization, the TEWMH is a good concept for measurements in self-adaptive systems as they require timely feedback for quick reactions to changes in the observed environment.

Acknowledgements

The authors would like to thank Prof. Phouc Tran-Gia for the stimulating environment which was a prerequisite for that work.

References

[1] R. Martin and M. Menth, "Improving the Timeliness of Rate Measurements," in 12^{th} *GI/ITG Conference on Measuring, Modelling and Evaluation of Computer and Communication Systems (MMB) together with* 3^{rd} *Polish-German Teletraffic Symposium (PGTS)*, Dresden, Germany, Sept. 2004, pp. 145 – 154.

[2] S. W. Roberts, "Control Chart Tests Based on Geometric Moving Averages," *Technometrics (Journal of the American Statistical Association)*, vol. 1, no. 3, pp. 239 – 250, 1959.

[3] S. J. Hunter, "The Exponentially Weighted Moving Average," *Journal of Quality Technology*, vol. 18, no. 4, pp. 203 – 210, 1986.

[4] S. V. Crowder, "A Simple Method for Studying Run-Length Distributions of Exponentially Weighted Moving Average Charts," *Technometrics (Journal of the American Statistical Association)*, pp. 401 – 407, 1987.

[5] J. M. Lucas and M. S. Saccucci, "Exponentially Weighted Moving Average Control Schemes: Properties and Enhancements," *Technometrics (Journal of the American Statistical Association)*, vol. 32, no. 1, pp. 1 – 29, 1990.

[6] S. H. Steiner, "Exponentially Weighted Moving Average Control Charts with Time Varying Control Limits and Fast Initial Response," *Journal of Quality Technology*, vol. 31, pp. 75 – 86, 1999.

[7] J. E. Wieringa, "Statistical Process Control for Serially Correlated Data," PhD thesis, University of Groningen, Department of Econometrics, Feb. 1999.

[8] A. Chen and E. A. Elsayed, "Design and Performance Analysis of the Exponentially Weighted Moving Average Mean Estimate for Processes Subject to Random Step Changes," *Technometrics (Journal of the American Statistical Association)*, vol. 44, no. 4, 2002.

[9] D. O. Awduche, L. Berger, D.-H. Gan, T. Li, V. Srinivasan, and G. Swallow, "RFC3209: RSVP-TE: Extensions to RSVP for LSP Tunnels," Dec. 2001.

[10] Y. B. et al., "RFC3290: An Informal Management Model for Diffserv Routers," http://www.ietf.org/rfc/rfc3290.txt, May 2002.

[11] Information Sciences Institute, "RFC793: Transmission Control Protocol," September 1981.

[12] J. Milbrandt, M. Menth, and S. Oechsner, "EBAC - A Simple Admission Control Mechanism," in 12^{th} *IEEE International Conference on Network Protocols (ICNP)*, Berlin, Germany, Oct. 2004.

[13] M. Menth, J. Milbrandt, and S. Oechsner, "Experience-Based Admission Control (EBAC)," in 9^{th} *IEEE Symposium on Computers and Communications (ISCC)*, Alexandria, Egypt, June 2004, pp. 903 – 910.

[14] J. Milbrandt, M. Menth, and J. Junker, "Performance of Experience-Based Admission Control in the Presence of Traffic Changes," in 5^{th} *IFIP-TC6 Networking Conference (Networking)*, Coimbra, Portugal, May 2006.