QUALITY AND RELIABILITY ENGINEERING INTERNATIONAL
Qual. Reliab. Engng. Int. 2003; **19**:337–353 (DOI: 10.1002/qre.586)

Special Issue

A Comparison of Shewhart Individuals Control Charts Based on Normal, Non-parametric, and Extreme-value Theory[‡]

M. B. (Thijs) Vermaat[1], Roxana A. Ion[2], Ronald J. M. M. Does[1,3,*,†] and Chris A. J. Klaassen[1,3]

[1] *Institute for Business and Industrial Statistics of the University of Amsterdam, IBIS UvA BV, Plantage Muidergracht 24, 1018 TV Amsterdam, The Netherlands*
[2] *Eurandom, PO Box 513, 5600 MB Eindhoven, The Netherlands*
[3] *Korteweg-de Vries Institute for Mathematics of the University of Amsterdam, Plantage Muidergracht 24, 1018 TV Amsterdam, The Netherlands*

Several control charts for individual observations are compared. Traditional ones are the well-known Shewhart individuals control charts based on moving ranges. Alternative ones are non-parametric control charts based on empirical quantiles, on kernel estimators, and on extreme-value theory. Their in-control and out-of-control performance are studied by simulation combined with computation. It turns out that the alternative control charts are not only quite robust against deviations from normality but also perform reasonably well under normality of the observations. The performance of the Empirical Quantile control chart is excellent for all distributions considered, if the Phase I sample is sufficiently large. Copyright © 2003 John Wiley & Sons, Ltd.

KEY WORDS: Shewhart control charts; statistical process control; non-parametrics; bootstrap; kernel estimators; extreme-value theory

INTRODUCTION

Since Shewhart originated the concept of the control chart in the early 1920s, it has become a powerful tool in statistical process control[1]. Shewhart-type control charts consist of a graph with time on the horizontal axis and a control characteristic (individual measurements or statistics such as mean or range) on the vertical axis. Control limits drawn provide easy checks on the stability of the process, that is, they signal the presence of special causes. The charts are usually constructed using 20–30 initial samples of about five items each, which in general are supposed to arise from purely random sampling. A treatment of statistical aspects of control charts in this typical textbook situation was given by Does and Schriever[2]. Up-to-date books on statistical process control also incorporating more advanced methods based on practice, are Wheeler[3], Quesenberry[4], Does *et al.*[5], and Montgomery[6].

In practice, situations frequently arise that require a charting procedure for individual measurements. Charting of individual observations has received extensive attention in the literature. In Roes *et al.*[7] and

*Correspondence to: Professor Dr R. J. M. M. Does, IBIS UvA, Plantage Muidergracht 24, 1018 TV Amsterdam, The Netherlands.
†E-mail: rjmmdoes@science.uva.nl
‡This paper is based on a presentation given at the second ENBIS Conference, Rimini, September 2002.

Received 30 January 2003
Revised 9 May 2003

Reynolds and Stoumbos[8,9] statistical aspects of charts of individual observations have been studied. Usually, the underlying distribution function is assumed to be normal, for the evaluation of the statistical performance. Nevertheless, there has always been some concern about this normality assumption. Especially when individual measurements are used, the normality assumption is risky. In Borror *et al.*[10] an EWMA control chart is designed that is robust to non-normality. In Stoumbos and Reynolds[11] the effects of both non-normality and autocorrelation on the performance of various individual control charts are studied. In a recent paper, Woodall and Montgomery[12] stated that

> There would appear to be an increasing role for non-parametric methods, particularly as data availability increases. Abundant data would cause the loss of power associated with non-parametric methods to become less of an issue

(see also Stoumbos *et al.*[13]). If large data sets are available, an attractive approach to parametric statistical inference would be to first use these large data sets to study distributional form.

However, in the present paper, we consider the non-parametric situation for individual measurements in which the underlying distribution function, denoted by F, is assumed to be unimodal, but otherwise unknown. This means that we include distributions that have an increasing–decreasing density, such as the normal, Student t, uniform, exponential, Laplace, and logistic distribution. It has been shown by Ion[14] that for densities with more than one mode, any Shewhart control chart for individual measurements is inappropriate, or at least suboptimal.

In our comparison we will use the standard, essentially parametric Shewhart control chart with control limits based on the average of the moving ranges of the individual measurements. The alternative control charts we will compare are charts based on empirical quantiles, which are related to the bootstrap method, charts based on kernel estimators, and charts based on extreme-value theory. The availability of modern computing power in statistical process control enables one to apply these computationally intensive techniques from mathematical statistics.

The ascertainment of the control limits is based on the observations obtained in the so-called Phase I, in which the data are collected from the production process and parameters are estimated (cf. Woodall and Montgomery[12]). In the present article, we consider the monitoring phase which is usually called Phase II. In most evaluations and comparisons of performance of control charts in Phase II, it is assumed, as noted by Woodall and Montgomery[12], that the in-control parameters are known, which is not the case in practice. For this reason, the statistical performance of the classical and newly proposed control charts will be studied by simulating the average and standard deviation of the in-control and out-of-control run length in Phase II of the control charts with the control limits determined by in-control observations from Phase I.

The use of all control charts was demonstrated by a real-life example from a printers assembling company in the master thesis of Vermaat (2003). In Quesenberry[15] it has been shown by simulation that in order to estimate control limits for individual measurements sufficiently accurately, one needs rather large sample sizes like 300 observations and more. Our simulation results support this for all control charts considered.

It turns out that the control chart based on the average of the moving ranges is suboptimal compared with the newly proposed control charts, except for independently, normally distributed random variables. However, even under normality, the alternative charts have quite good performance, especially when a sufficient amount of data are available. The control chart based on empirical quantiles is excellent for all distributions considered.

This paper is organized as follows. In the next section, the four control charts are defined. The results are given of an extensive Monte Carlo study based on 10 000 simulations for six different distributions and for several Phase I sample sizes. This number of simulations turned out to be large enough for our simulation method. Finally, conclusions are given.

DESCRIPTION OF THE CONTROL CHARTS

We will consider the usual control charts with a lower control limit (LCL) and an upper control limit (UCL). This means that if the measurement value X is lower than LCL or higher than UCL, then the process is called out of control.

Classical individuals control chart based on the average of the moving ranges

If the distribution function F is assumed to be normal, then the traditional Shewhart individuals control chart has limits defined by

$$\mathrm{UCL} = \mu + \Phi^{-1}\left(1 - \frac{\alpha}{2}\right)\sigma$$

and

$$\mathrm{LCL} = \mu + \Phi^{-1}\left(\frac{\alpha}{2}\right)\sigma$$

where Φ^{-1} is the standard normal quantile function, and where μ is the mean and σ is the standard deviation of the normal distribution function F. Level α is the false alarm rate. Typically, μ and σ are unknown. However, we shall assume that they can be estimated via a Phase I sample $X_1, X_2, \ldots, X_k$ of independently and identically distributed random variables. Classical estimators of μ and σ are the sample mean $\bar{X}_k = \sum_{i=1}^{k} X_i/k$ and the sample standard deviation $S_k = \sqrt{\sum_{i=1}^{k}(X_i - \bar{X}_k)^2/(k-1)}$. The sample standard deviation is asymptotically efficient for independently and identically distributed normal random variables, but it has the disadvantage that it is sensitive to trends and oscillations. Consequently, when such phenomena might occur, we have to use estimators of the standard deviation that are less sensitive to these deviations, cf. Kamat[16]. The average of the moving ranges can be scaled by $2/\sqrt{\pi}$ to obtain such an estimator, and is defined by

$$\overline{\mathrm{MR}}_k = \frac{1}{k-1}\sum_{i=2}^{k}|X_i - X_{i-1}|$$

In Duncan[17] the individuals control chart is mentioned with control limits based on the average of the moving ranges (AMR) defined by

$$\mathrm{UCL}_{\mathrm{AMR}} = \bar{X}_k + \Phi^{-1}\left(1 - \frac{\alpha}{2}\right)\frac{\sqrt{\pi}}{2}\overline{\mathrm{MR}}_k$$

and

$$\mathrm{LCL}_{\mathrm{AMR}} = \bar{X}_k + \Phi^{-1}\left(\frac{\alpha}{2}\right)\frac{\sqrt{\pi}}{2}\overline{\mathrm{MR}}_k$$

In practice, this AMR control chart is the standard chart for individual observations. The reasons for this are, first, the constant by which the average of the moving ranges has to be multiplied in order to obtain an unbiased estimator of σ is quite similar for probability distributions having density curves of very different shapes, cf. Burr[18]. Second, the AMR control chart tends to perform reasonably well for moderate Phase I sample sizes more or less independently of the probability distribution the observations stem from. This is thoroughly studied in chapter five of Wheeler[3]. In this paper we compare the performance of the AMR control chart to three competitors.

A more exact version of the individuals control chart is developed in Roes *et al.*[7]. However, due to limitations in the size of this paper we will restrict our attention to the control chart described above.

Empirical quantile control chart (bootstrap)

A natural estimator of the q-quantile of the distribution function F is the empirical quantile $\hat{F}_k^{-1}(q)$, which is defined as

$$\hat{F}_k^{-1}(q) = \inf\{x \mid \hat{F}_k(x) \geq q\}, \quad 0 < q < 1$$

where $\hat{F}_k$ is the empirical distribution function that puts mass $1/k$ at each X_i, $1 \leq i \leq k$, i.e.

$$\hat{F}_k(x) = \frac{1}{k}\sum_{i=1}^{k} I_{\{X_i \leq x\}}, \quad -\infty < x < \infty$$

with I the indicator function, i.e. $I_{\{x \leq y\}} = 1$ if $x \leq y$ holds and 0 otherwise. Now, an obvious estimator of the upper control limit is based on the empirical quantile (EQ)

$$\text{UCL}_{\text{EQ}} = \hat{F}_k^{-1}\left(1 - \frac{\alpha}{2}\right) = X_{(\lceil (1-\alpha/2)k \rceil)}$$

with $X_{(1)} \leq X_{(2)} \leq \cdots \leq X_{(k)}$ denoting the order statistics of the initial sample $X_1, X_2, \ldots, X_k$ and $\lceil y \rceil$ the smallest integer not smaller than y. The lower control limit estimated by the empirical quantile is defined by

$$\text{LCL}_{\text{EQ}} = X_{(\lfloor (\alpha/2)k+1 \rfloor)}$$

where $\lfloor y \rfloor$ denotes the largest integer not larger than y. The exceedance probability for this estimation method is not always smaller than α. To be more specific:

$$\text{P}(X > \text{UCL}_{\text{EQ}} \text{ or } X < \text{LCL}_{\text{EQ}}) = 2\frac{\lfloor (\alpha/2)k \rfloor + 1}{k+1} \in \left(\alpha - \frac{\alpha}{k+1}, \alpha + \frac{2-\alpha}{k+1}\right]$$

To overcome this problem another estimation method is proposed in Ion[14], which guarantees that the exceedance probability equals at most the significance level α. However, the resulting alternative does not have a better performance than the EQ control chart for shifts in the mean. Therefore, we will not consider this alternative.

It can be shown that the EQ control chart is a special case of the bootstrap control chart defined by Willemain and Runger[19]. The philosophy of the bootstrap approach to statistical problems is to replace the unknown distribution function F of a random variable X by an empirical distribution function. In control charts for individual observations one would like to use the $\text{UCL}_{\text{EQ}} = F^{-1}(1 - \alpha/2)$. Consequently, the bootstrap approach applied to control charts for individual observations, with a Phase I sample $X_1, X_2, \ldots, X_k$, yields $\hat{F}_k^{-1}(1 - \alpha/2)$ as an estimate of the upper control limit. For the more general situation of control charts for averages the bootstrap method is studied in Liu and Tang[20]. Note that in Jones and Woodall[21] several control charts based on the bootstrap are compared.

Kernel control charts

In Rosenblatt[22] and Parzen[23] kernel estimators of the density function f are defined as

$$\hat{f}_w(x) = \frac{1}{k}\sum_{i=1}^{k}\frac{1}{h}w\left(\frac{x - X_i}{h}\right), \quad -\infty < x < \infty$$

where the kernel w is non-negative such that $\int_{-\infty}^{\infty} w(x)\,\text{d}x = 1$, and where the bandwidth h is positive and small. Consequently, with $W(x) = \int_{-\infty}^{x} w(y)\,\text{d}y$ the distribution function corresponding to the density w,

$$\hat{F}_w(x) = \frac{1}{k}\sum_{i=1}^{k} W\left(\frac{x - X_i}{h}\right), \quad -\infty < x < \infty$$

is a kernel estimator of the distribution function F. As in Reiss[24] a smooth alternative to the conventional sample quantile function may be defined by

$$\hat{F}_w^{-1}(q) = \inf\left\{x \,\middle|\, \frac{1}{k}\sum_{i=1}^{k} W\left(\frac{x - X_i}{h}\right) \geq q\right\}, \quad 0 < q < 1$$

The choice of the bandwidth h is more important than the choice of the kernel w. A relatively large value of h gives too much smoothness, and a relatively small value of h gives big fluctuations. Several choices of the kernel are possible, e.g. the Gaussian kernel and the so-called Epanechnikov kernel, i.e.

$$w(x) = \begin{cases} \frac{3}{4\sqrt{5}}\left(1 - \frac{x^2}{5}\right), & \text{if } |x| < \sqrt{5} \\ 0, & \text{if } |x| \geq \sqrt{5} \end{cases}$$

In this paper we restrict attention to the Epanechnikov kernel because it gives better results than the Gaussian kernel (cf. Ion[14]).

In view of Azzalini[25], the optimal choice of the bandwidth h is $h = Ck^{-1/3}$, where C is a constant which depends on σ, the standard deviation of F. Azzalini did some numerical work using the Epanechnikov kernel, and he concluded that good values for the constant C are between σ and 2σ for a large number of distributions. Based on an extensive simulation study we found that the best choice in our situation is $C = 2\sigma$. Since σ is unknown, we estimate it by the sample standard deviation $S_k = \sqrt{\sum_{i=1}^{k} (X_i - \bar{X}_k)^2/(k-1)}$. Note, that we could also use other estimators for σ like the one based on the average of the moving ranges. So for the Epanechnikov kernel (EK) the control limits are:

$$\mathrm{UCL}_{\mathrm{EK}} = \inf\left\{x \,\middle|\, \frac{1}{k}\sum_{i=1}^{k} W\left(\frac{x - X_i}{2k^{-1/3}S_k}\right) \geq 1 - \alpha/2\right\}$$

$$\mathrm{LCL}_{\mathrm{EK}} = \sup\left\{x \,\middle|\, \frac{1}{k}\sum_{i=1}^{k} W\left(\frac{x - X_i}{2k^{-1/3}S_k}\right) \leq \alpha/2\right\}$$

Extreme-value theory control chart

Extreme-value theory yields another method to estimate the control limits of a Shewhart control chart. To the best of our knowledge, Shewhart control charts based on this theory have not been considered before.

Define

$$M_k^{(r)} = \frac{1}{m}\sum_{j=1}^{m}(\log X_{(k-j+1)} - \log X_{(k-m)})^r$$

and

$$\bar{M}_k^{(r)} = \frac{1}{m}\sum_{j=1}^{m}(\log X_{(j)} - \log X_{(m+1)})^r$$

where the integer r takes the values $r = 1, 2$, and m is the number of upper respectively lower order statistics used in the estimation of the control limits.

Extreme-value theory deals with the tail behaviour of distributions. These tails can be modelled by an extreme-value distribution, which is determined by an extreme-value index γ (cf. Dekkers *et al.*[26]). If we do not make any assumptions on γ, we may use the moment estimator of it, defined by

$$\hat{\gamma}_k = M_k^{(1)} + 1 - \frac{1}{2}\left\{1 - \frac{(M_k^{(1)})^2}{M_k^{(2)}}\right\}^{-1}$$

Furthermore, the q-quantile of the distribution function F is estimated in Dekkers *et al.*[26] as

$$\hat{F}_k^{-1}(1 - q; \hat{\gamma}_k) = X_{(k-m)} + \frac{(m/(kq))^{\hat{\gamma}_k} - 1}{\hat{\gamma}_k}(1 - (\hat{\gamma}_k \wedge 0))X_{(k-m)}M_k^{(1)}$$

with $0 < q < 1$. Recall that $x \wedge y$ and $x \vee y$ denote the minimum and maximum, respectively, of x and y. Consequently, the upper control limit of the extreme-value theory control chart (EV) may be defined by

$$\mathrm{UCL}_{\mathrm{EV}} = \hat{F}_k^{-1}\left(1 - \frac{\alpha}{2}; \hat{\gamma}_k\right)$$

In a similar way, the lower control limit may be taken to be

$$\mathrm{LCL}_{\mathrm{EV}} = X_{(m+1)} + \frac{(m/(k\alpha/2))^{\bar{\gamma}_k} - 1}{\bar{\gamma}_k}(1 - (\bar{\gamma}_k \wedge 0))X_{(m+1)}\bar{M}_k^{(1)}$$

with $\bar{\gamma}_k$ defined as

$$\bar{\gamma}_k = \bar{M}_k^{(1)} + 1 - \frac{1}{2}\left\{1 - \frac{(\bar{M}_k^{(1)})^2}{\bar{M}_k^{(2)}}\right\}^{-1}$$

Of course, the sequence $m = m(k)$ has to be chosen appropriately in order to obtain good performance of $\mathrm{UCL}_{\mathrm{EV}}$ and $\mathrm{LCL}_{\mathrm{EV}}$. In the numerical evaluation of the extreme-value theory control chart we have considered several values of m. We found that $5 \vee (k/500)$ is a reasonable choice.

SIMULATIONS

In order to study the performance of the control charts introduced in the section 'Description of the control charts' we have conducted a simulation experiment, for some choices of the distribution function F and of the Phase I sample size k. To describe our simulation procedure let us assume as above that we have a sample $X_1, X_2, \ldots, X_k$ of size k from a distribution function F. This training sample is used to estimate the LCL and UCL by $\widehat{\mathrm{LCL}}$ and $\widehat{\mathrm{UCL}}$, respectively. Note, that the process generating this training sample is in statistical control. We do not study robustness properties, regarding moderately out-of-control situations during Phase I. We denote by $\hat{p}_k(X_1, \ldots, X_k)$ the conditional probability given the training sample, that a new independent random variable X from the same distribution F exceeds this upper control limit or is below the lower control limit, i.e.

$$\hat{p}_k(X_1, \ldots, X_k) = \mathrm{P}(X > \widehat{\mathrm{UCL}} \text{ or } X < \widehat{\mathrm{LCL}} \mid X_1, \ldots, X_k)$$

Given the training sample and hence the control limits, the run length RL of the resulting control chart is a random variable with a geometric distribution with parameter $\hat{p}_k(X_1, \ldots, X_k)$ and consequently with average run length

$$\mathrm{E}(\mathrm{RL} \mid X_1, \ldots, X_k) = 1/\hat{p}_k(X_1, \ldots, X_k)$$

Note that this is a conditional average run length and hence a random variable. Taking the expectation over the training sample $X_1, \ldots, X_k$ we get the unconditional average run length

$$\mathrm{ARL} = \mathrm{E}(\mathrm{RL}) = \mathrm{E}(\mathrm{E}(\mathrm{RL} \mid X_1, \ldots, X_k)) = \mathrm{E}\frac{1}{\hat{p}_k(X_1, \ldots, X_k)}$$

Since this expectation cannot be computed directly, we have simulated it by generating at each instance 10 000 training samples, by computing for each training sample $(x_1, \ldots, x_k)$ the average run length $1/\hat{p}_k(x_1, \ldots, x_k)$, and by averaging these average run lengths over the 10 000 training samples. In the same way it can be seen that the unconditional standard deviation of the average run length equals

$$\begin{aligned}\mathrm{SDRL} &= \sqrt{\mathrm{Var}(\mathrm{RL})}\\ &= \sqrt{\mathrm{E}(\mathrm{Var}(\mathrm{RL} \mid X_1, \ldots, X_k)) + \mathrm{Var}(\mathrm{E}(\mathrm{RL} \mid X_1, \ldots, X_k))}\\ &= \sqrt{2\mathrm{E}\left(\frac{1}{\hat{p}_k(X_1, \ldots, X_k)}\right)^2 - \left(E\frac{1}{\hat{p}_k(X_1, \ldots, X_k)}\right)^2 - \mathrm{E}\frac{1}{\hat{p}_k(X_1, \ldots, X_k)}}\end{aligned}$$

Moreover, we have simulated this SDRL by also computing for each training sample $(1/\hat{p}_k(x_1, \ldots, x_k))^2$, by averaging these squares over the 10 000 training samples, and by substituting the appropriate averages into the above formula for the SDRL.

To study the performance of the control chart in an out-of-control situation we have studied shifts in the mean. We have considered shifts $\delta\sigma$ with δ ranging from 0.25 through 5 where σ is the standard deviation of the studied distribution. Of course, we also have considered $\delta = 0$ in order to study the in-control properties of the control charts. Given the 10 000 training samples $(x_1, \ldots, x_k)$, we have calculated for each shift the average run length ARL and the standard deviation of the run length SDRL as described above.

This procedure of simulating the performance of control charts differs apparently from the one used by e.g. Quesenberry[15], where once the $\widehat{\mathrm{LCL}}$ and $\widehat{\mathrm{UCL}}$ have been obtained from the training sample, one realization of the run length is again obtained by simulation. However, this second simulation step is not necessary, as we have seen above, since the conditional average run length can be computed exactly. Of course, there is no point in simulation when exact computation is possible.

In our simulations the false alarm rate α is chosen to be equal to 0.0027, since this value yields the traditional 3σ limits in the classical Shewhart control chart. Consequently, if the variation of $\hat{p}_k(X_1, \ldots, X_k)$ would have been small the mean (ARL) and the standard deviation (SDRL) of the run length should both have been close to 370 under $\delta = 0$.

The simulations have been done for six different choices of distribution function F, namely the normal, Student's t with four degrees of freedom, uniform, exponential, Laplace, and logistic distribution and for sample sizes k equal to 250, 500, 1000, 2500, 5000, and 10 000. The shift in the mean of size $\delta\sigma$ is done for 17 different values of δ, namely 0 (0.25) 3.5, 4, and 5.

The results of the simulations are presented for each distribution by means of two complementary figures indicating the ARL and SDRL. Although we have simulated k at six different sample sizes, we present here the results only for $k = 1000$ and under normal F also for $k = 250$. A complete survey of the simulation results may be found in the master thesis of Vermaat (2003). The simulation results are valid for all values of σ, because all methods are scale invariant.

Under normality the control chart based on the average of the moving ranges has to perform closely to optimal. This is confirmed by Figure 1, which also shows that the other control charts perform rather badly. Apparently, the sample size $k = 250$ is too small for them. The behaviour of all four control charts under the non-normal distributions is rather bad overall. The same conclusions hold for $k = 500$. Note, that for these 'small' values of k the EQ control chart has control limits equal to the smallest and largest observation. From $k = 1000$ onwards the differences between the control charts become clearer. In Figure 2 we see that all control charts behave quite similarly under normality.

In Figure 3 the results for a t_4 distribution are given. Indeed, the AMR control chart generates a lot of false alarms. The behaviour of the other three control charts is quite reasonable for the in-control situation. However, few false alarms tend to create few real alarms under shifts.

In Figure 4 the results for uniform distributions are given. Only the EQ control chart has finite ARL and SDRL. The other control charts are not even applicable when the underlying distribution is uniform.

Figure 5 presents the results for exponential distributions. We see, that the EQ control chart has a maximum in the ARL around $\delta = 0.25$. This maximum is due to the fact that the EQ control chart estimates an LCL within the support of the exponential distribution. If the process shifts towards the UCL when δ increases, the LCL is of course harder to violate and the UCL easier. Since the density of the exponential is larger near the LCL than near the UCL, this causes the probability of an alarm to decrease and hence the ARL to increase for small δ.

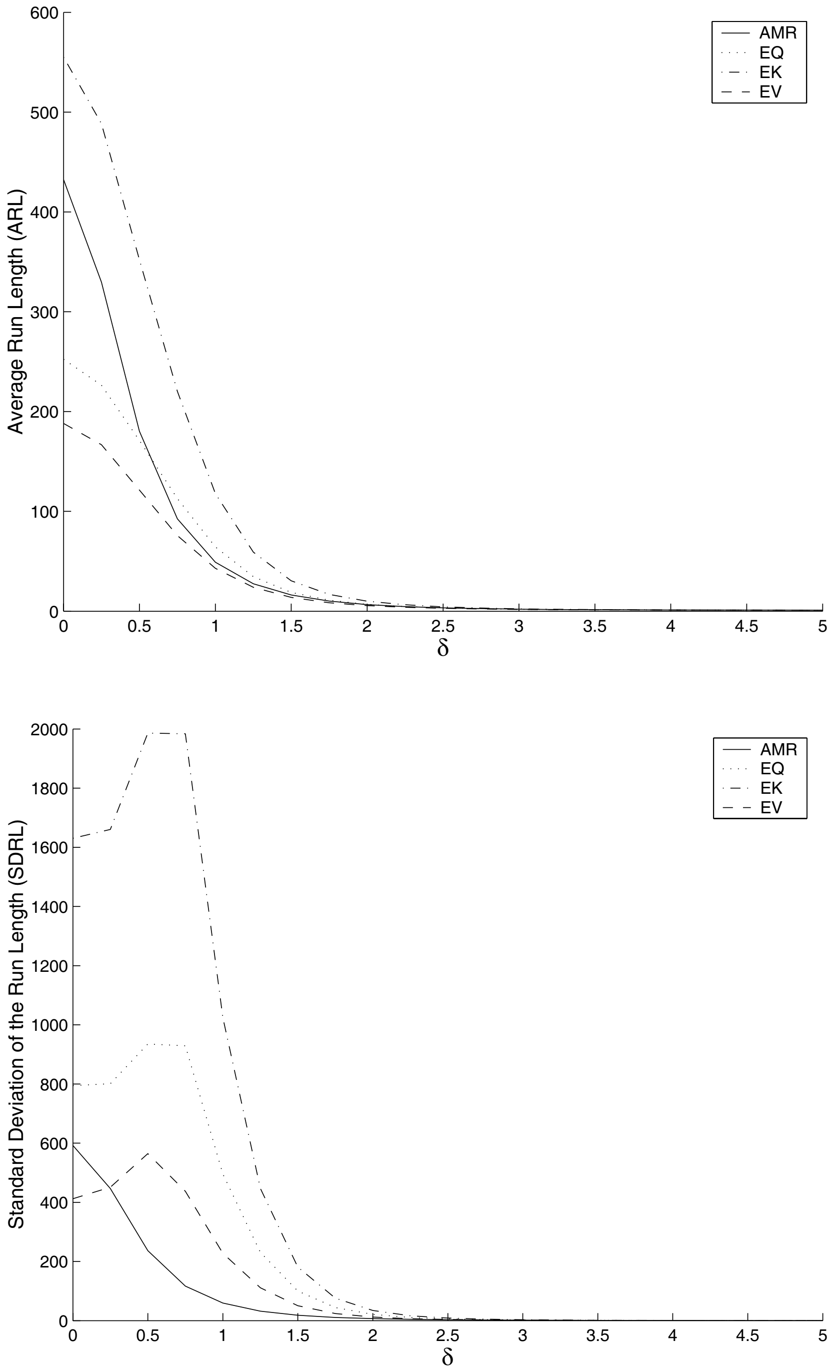

Figure 1. The ARL and the SDRL of the four control charts under normality for $k = 250$ and standardized shifts δ

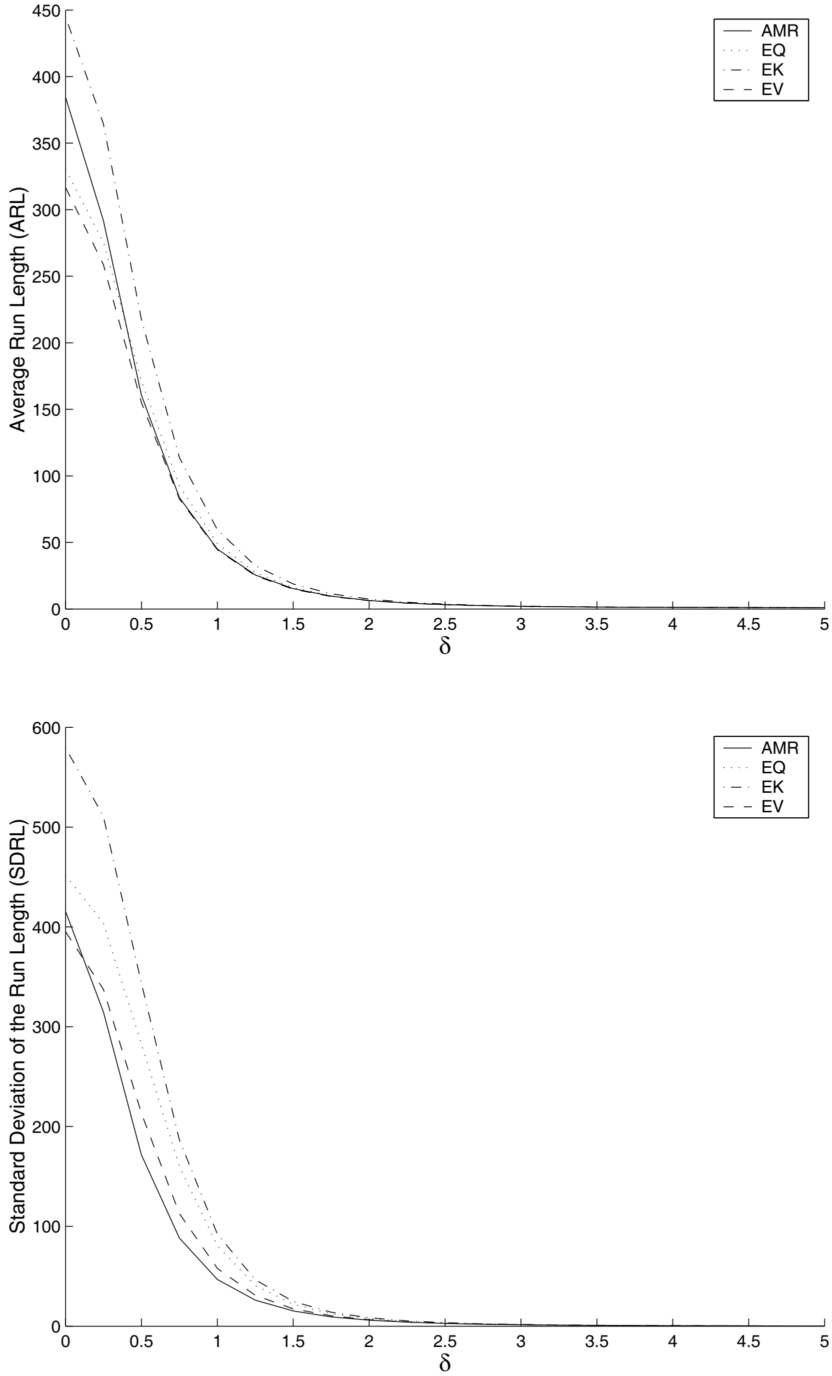

Figure 2. The ARL and the SDRL of the four control charts under normality for $k = 1000$ and standardized shifts δ

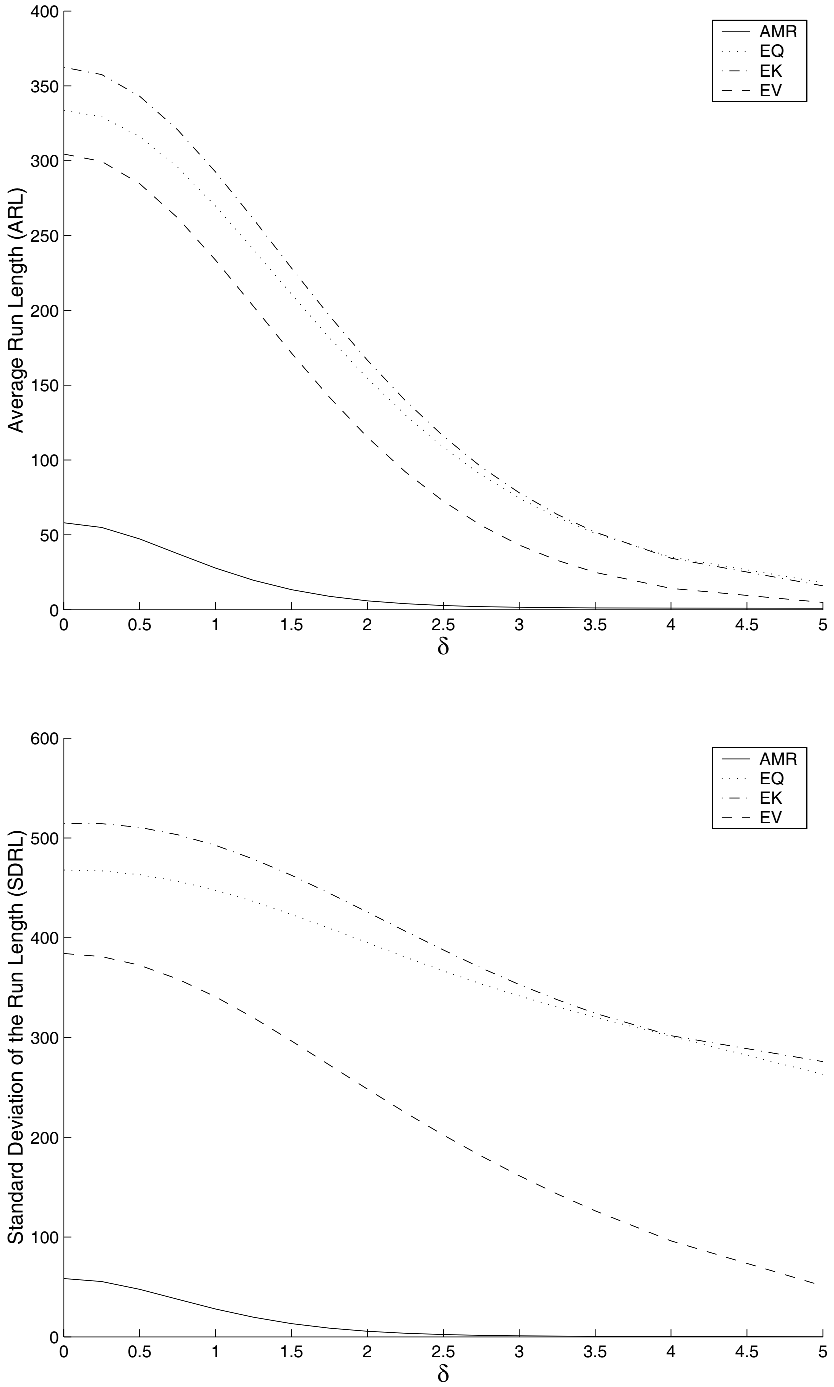

Figure 3. The ARL and the SDRL of the four control charts under a Student t_4 distribution for $k = 1000$ and standardized shifts δ

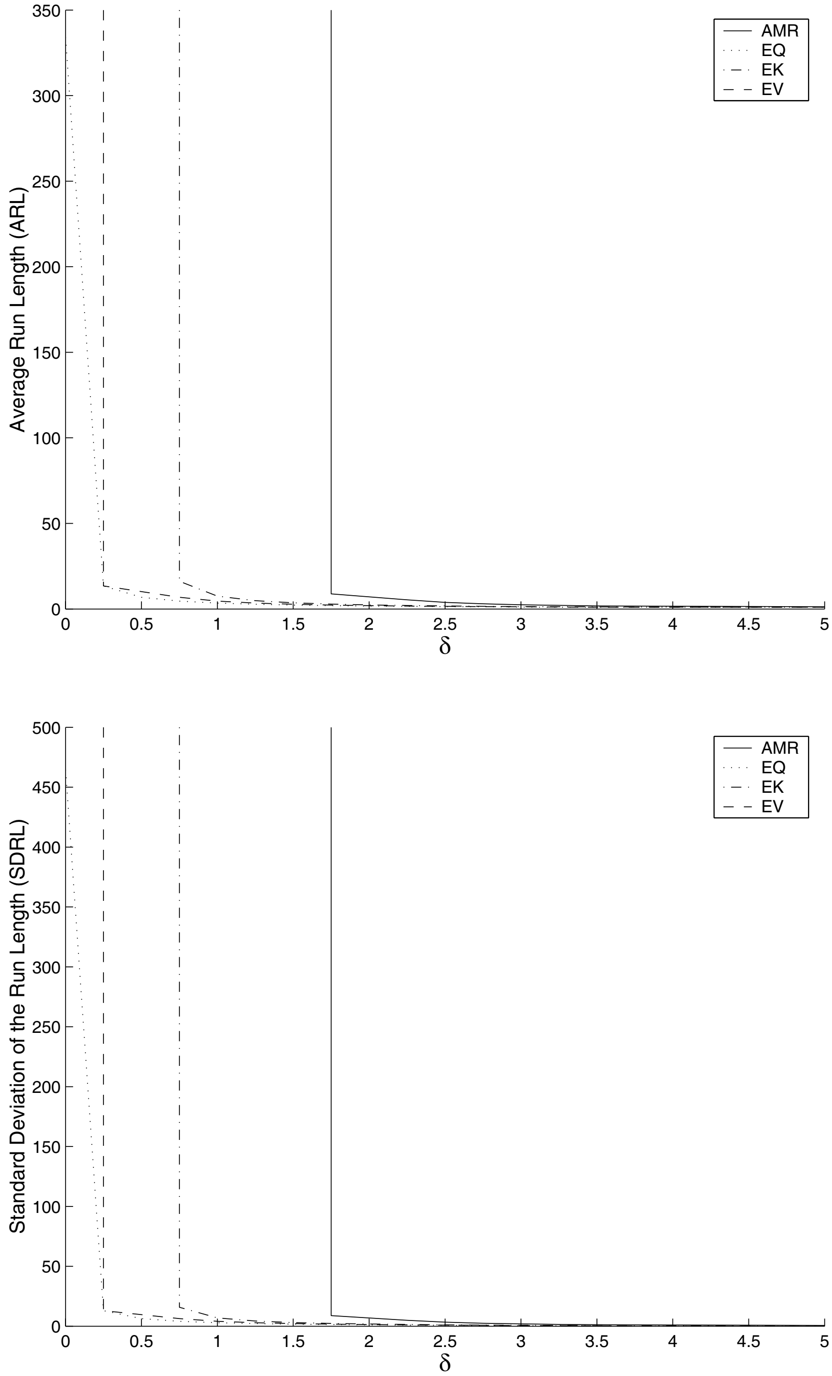

Figure 4. The ARL and the SDRL of the four control charts under uniformity for $k = 1000$ and standardized shifts δ

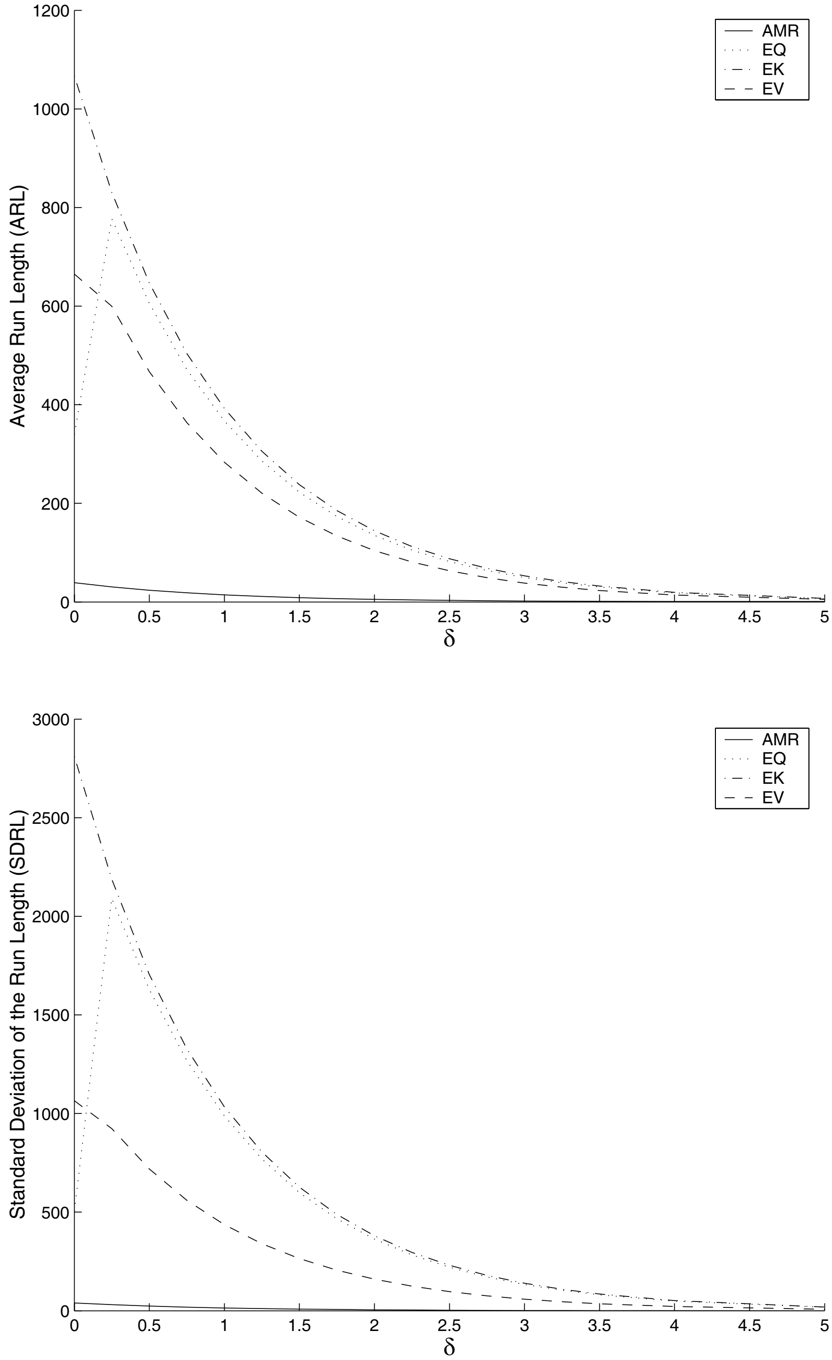

Figure 5. The ARL and the SDRL of the four control charts under an exponential distribution for $k = 1000$ and standardized shifts δ

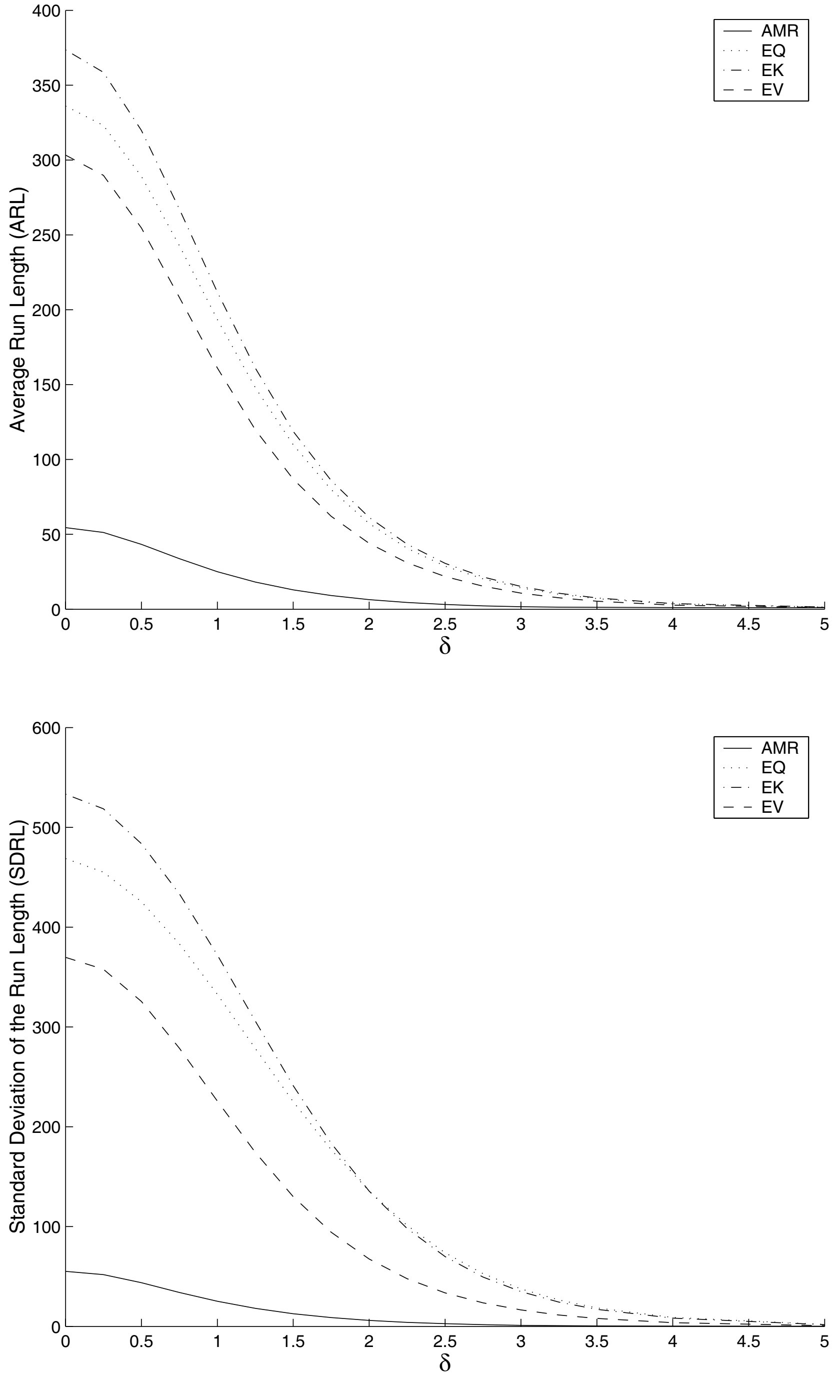

Figure 6. The ARL and the SDRL of the four control charts under a Laplace distribution for $k = 1000$ and standardized shifts δ

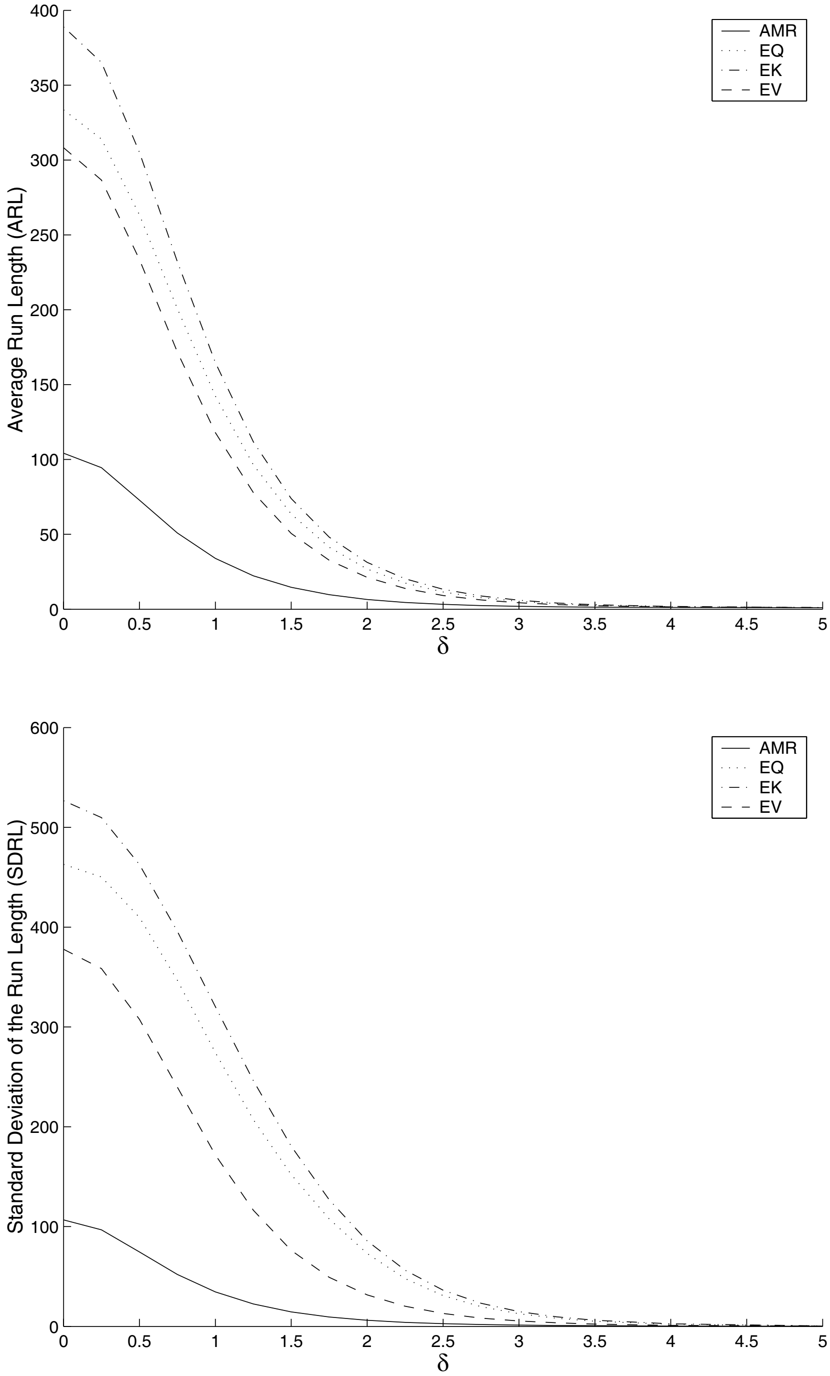

Figure 7. The ARL and the SDRL of the four control charts under a logistic distribution for $k = 1000$ and standardized shifts δ

Because the estimates of the LCL in the EK and the EV control charts typically fall outside the support of the exponential distribution, these control charts do not show this phenomenon. The EK control chart and the EQ control chart show almost identical performance for $\delta > 0.25$. Note, that the behaviour of the AMR control chart for $\delta = 0$ is really bad.

In Figures 6 and 7 the results are given for Laplace and logistic distributions, respectively. We see again that the AMR control chart performs very poorly for $\delta = 0$. The other control charts perform quite similarly. Note, that although the shape of a logistic distribution is comparable with a normal distribution, the tails are completely different.

For all studied distributions it holds that if the sample size increases the results of the EQ, the EK, and the EV control charts get closer to each other. The EQ control chart is the only one which behaves reasonably for a uniform distribution. The AMR control chart generates a lot of false alarms for non-normal distributions.

CONCLUSIONS

If the distribution function F is normal, then the AMR control chart based on the average of the moving ranges behaves quite well, but under non-normal distributions its performance is extremely bad in the in-control situation. Hence it is reasonable to consider other control charts based on e.g. non-parametric and extreme-value theory.

The EQ control chart has the advantage that it is easy to compute and distribution-free in the in-control situation. Simulations show that for a broad range of distributions we get almost the same results. When we look to the control charts based on the EK and EV theory, respectively, we may conclude that their behaviour is quite similar to that of the EQ control chart, except for the uniform distribution. The EQ control chart turns out to be the best of all proposed control charts. It can be improved only partially for some specific situations, e.g. when normality holds.

It should be noted that the non-parametric and the EV control charts need more than 500 and preferably at least 1000 observations from Phase I in order to attain reasonable performance. Such a limitation is not present for the AMR control chart, which may work well in some situations even for a training set of size much smaller than 250, cf. Wheeler[3].

Acknowledgements

This paper is supported by funding under the European Commission's Fifth Framework 'Growth' Programme via the Thematic Network 'Pro-ENBIS' contract reference: G6RT-CT-2001-05059. The authors are solely responsible for the content and it does not represent the opinion of the Community; the Community is not responsible for any use that might be made of data therein. The second author also is supported financially by the Netherlands Organization for Scientific Research, NWO. The authors thank the referee and the guest editor Shirley Coleman for their useful comments and suggestions.

REFERENCES

1. Shewhart WA. *Economic Control of Quality of Manufactured Product*. Van Nostrand: Princeton, NJ, 1931.
2. Does RJMM, Schriever BF. Variables control chart limits and tests for special causes. *Statistica Neerlandica* 1992; **46**:229–245.
3. Wheeler DJ. *Advanced Topics in Statistical Process Control*. SPC Press: Knoxville, TN, 1995.
4. Quesenberry CP. *SPC Methods for Quality Improvement*. Wiley: New York, 1997.
5. Does RJMM, Roes KCB, Trip A. *Statistical Process Control in Industry*. Kluwer Academic: Dordrecht, 1999.
6. Montgomery DC. *Introduction to Statistical Quality Control*. Wiley: New York, 2001.
7. Roes KCB, Does RJMM, Schurink Y. Shewhart-type control charts for individual observations. *Journal of Quality Technology* 1993; **25**:188–198.

8. Reynolds MR Jr, Stoumbos ZG. Monitoring the process mean and variance using individual observations and variable sampling intervals. *Journal of Quality Technology* 2001; **33**:181–205.
9. Reynolds MR Jr, Stoumbos ZG. Individuals control schemes for monitoring the mean and variance of processes subject to drifts. *Stochastic Analysis and Applications* 2001; **19**:863–892.
10. Borror CM, Montgomery DC, Runger GC. Robustness of the EWMA control chart to non-normality. *Journal of Quality Technology* 1999; **31**:309–316.
11. Stoumbos ZG, Reynolds MR Jr. Robustness to non-normality and autocorrelation of individuals control charts. *Journal of Statistical Computation and Simulation* 2000; **66**:145–187.
12. Woodall WH, Montgomery DC. Research issues and ideas in statistical process control. *Journal of Quality Technology* 1999; **31**:376–386.
13. Stoumbos ZG, Reynolds MR Jr, Ryan TP, Woodall WH. The state of statistical process control as we proceed into the 21st century. *Journal of the American Statistical Association* 2000; **95**:992–998.
14. Ion RA. Non-parametric statistical process control. *PhD Thesis*, University of Amsterdam, Amsterdam, The Netherlands, 2001.
15. Quesenberry CP. The effect of sample size on estimated limits for $\bar{X}$ and X control charts. *Journal of Quality Technology* 1993; **25**:237–247.
16. Kamat AR. On the mean successive difference and its ratio to the root mean square. *Biometrika* 1953; **40**:116–127.
17. Duncan AJ. *Quality Control and Industrial Statistics*. R. D. Irwin: Homewood, IL, 1965; 40.
18. Burr IW. The effect of non-normality on constants for X-bar and R-charts. *Industrial Quality Control* 1967; **23**:563–569.
19. Willemain TR, Runger GC. Designing control charts using an empirical reference distribution. *Journal of Quality Technology* 1996; **28**:31–38.
20. Liu RY, Tang J. Control charts for dependent and independent measurements based on bootstrap methods. *Journal of the American Statistical Association* 1996; **91**:1694–1700.
21. Jones LA, Woodall WH. The performance of bootstrap control charts. *Journal of Quality Technology* 1998; **30**:362–375.
22. Rosenblatt M. Remarks on some non-parametric estimates of a density function. *Annals of Mathematical Statistics* 1956; **27**:832–837.
23. Parzen E. On estimation of a probability density function and mode. *Annals of Mathematical Statistics* 1962; **33**:1065–1076.
24. Reiss RD. *Approximate Distributions of Order Statistics*. Springer: New York, 1989; 253.
25. Azzalini A. A note on the estimation of a distribution function and quantiles by a kernel method. *Biometrika* 1981; **68**:326–328.
26. Dekkers ALM, Einmahl JHJ, De Haan L. A moment estimator for the index of an extreme-value distribution. *Annals of Statistics* 1989; **17**:1833–1855.

Authors' biographies

M. B. (Thijs) Vermaat obtained a Masters Degree in Econometrics and Operations Research at the University of Groningen in 2002 and a Master Degree in Statistics at the same University in 2003. Currently, he is a PhD student at the University of Amsterdam and a Consultant in Industrial Statistics at the Institute for Business and Industrial Statistics. His research interests are control charts, extreme-value theory, Bernstein approximations, and Six Sigma.

Roxana A. Ion obtained a Masters Degree in Mathematics with specialization in Applied Mathematics at the University of Bucharest in 1996. After graduation she started working as a PhD student in Statistics at the University of Amsterdam. In December 2001 she obtained her PhD in Statistics at the University of Amsterdam. The title of her thesis is 'Non-parametric statistical process control'. Afterwards she was postdoc at the European research center EURANDOM, where she did research and consultancy within joint projects with industry. Currently, she is assistant professor at the Department of Technology Management of the Technical University of Eindhoven. Her research interests are non-parametric statistics, statistical process control, extreme-value theory, and reliability theory.

Ronald J. M. M. Does obtained his MSc degree (cum laude) in Mathematics at the University of Leiden in 1976. In 1982 he defended his PhD entitled 'Higher order asymptotics for simple linear rank statistics' at the same university. From 1981 to 1989, he worked at the University of Maastricht, where he became Head of the Department of Medical Informatics and Statistics. In that period his main research interests were medical statistics and psychometrics. In 1989 he joined Philips Electronics as a Senior Consultant in Industrial Statistics. Since 1991 he has been Professor of Industrial Statistics at the University of Amsterdam. In 1994 he founded the Institute for Business and Industrial Statistics, which operates as an independent consultancy firm within the University of Amsterdam. The projects at this institute involve SPC, Taguchi and Shainin methods, and Six Sigma. His current research activities lie in the design of control charts for non-standard situations and the improvement of statistical methods in Six Sigma.

Chris A. J. Klaassen obtained his MSc degree in Mathematics at the University of Nijmegen in 1974. In 1980 he defended his PhD entitled 'Statistical performance of location estimators' at the University of Leiden. Having held positions in Statistics at the Mathematical Center in Amsterdam and the Mathematics Department of the University of Leiden, he was appointed as full Professor in Mathematical Statistics at the University of Amsterdam in 1990. Occasionally he acts as advisor for the Institute for Business and Industrial Statistics, which operates as an independent consultancy firm within the University of Amsterdam. His main interest is semiparametric statistics, but he has also published on e.g. Edgeworth expansions, finite sample inequalities in estimation, GARCH processes, the bootstrap, and acceptance sampling.