Estimation methods for statistical process control

Schoonhoven, M.

Publication date
2011

Chapter 1

Introduction

This chapter gives a short introduction to Shewhart control charts, an overview of new developments and an outline of the thesis.

1.1 Shewhart control charts

Processes are subject to variation. Whether or not a given process is functioning normally can be evaluated with control charts. Such charts show whether the variation is entirely due to common causes or whether some of the variation is due to special causes. Variation due to common causes is inevitable: it is generated by the design and standard operations of the process. When the process variation is due to common causes only, the process is said to be in statistical control. In this case, the process fluctuates within a predictable bandwidth. Special causes of process variation may consist of such factors as extraordinary events, unexpected incidents, or a new supplier for incoming material. For optimal process performance, such special causes should be detected as soon as possible and prevented from occurring again. Control charts are used to signal the occurrence of a special cause. The power of the control chart lies partly in its simplicity: it consists of a graph of a process characteristic plotted through time. The control limits in the graph provide easy checks on the stability of the process (i.e. no special causes present). The concept of control charts originates with Shewhart (1931) and has been extensively discussed and extended in numerous textbooks (see e.g. Duncan (1986), Does et al. (1999) and Montgomery (2009)).

In the standard situation, 20-30 samples of about five units are taken initially to construct a control chart. When a process characteristic is a numerical variable, it is standard practice to control both the mean value of the characteristic and its spread. The control limits of the statistic of interest are calculated as the average of the sample mean or standard deviation plus or minus a multiplier times the standard deviation of the statistic. The spread parameter of the process is controlled first, followed by the location parameter. An example of such a combined standard deviation and location chart is given in Figure 1.1.

The general set-up of a Shewhart control chart for the dispersion parameter is as follows. Let Y_{ij}, $i = 1, 2, 3, ...$ and $j = 1, 2, ..., n$, denote samples of size n taken in sequence of the process variable to be monitored. We assume the Y_{ij}'s to be independent and $N(\mu, (\lambda\sigma)^2)$ distributed, where λ is a constant. When $\lambda = 1$, the standard deviation of the process is in control; otherwise the standard deviation has changed. Let $\hat{\sigma}_i$ be an estimate of $\lambda\sigma$ based on the i-th sample Y_{ij}, $j = 1, 2, ..., n$. Usually, $\lambda\sigma$ is estimated by the sample standard deviation S. When the in-control σ is known, the process standard deviation can be monitored by plotting $\hat{\sigma}_i$ on a standard deviation control chart with respective upper and lower control limits

$$UCL = U_n\sigma,\ LCL = L_n\sigma, \tag{1.1}$$

where U_n and L_n are factors such that for a chosen type I error probability α we have

$$P(L_n\sigma \leq \hat{\sigma}_i \leq U_n\sigma) = 1 - \alpha.$$

When $\hat{\sigma}_i$ falls within the control limits, the spread is deemed to be in control.

For the location control chart, the Y_{ij}'s, $i = 1, 2, 3, ...$ and $j = 1, 2, .., n$, again denote samples of the process variable to be monitored. In this case, we assume the Y_{ij}'s to be independent and $N(\mu + \delta\sigma, \sigma^2)$ distributed, where δ is a constant. When $\delta = 0$, the mean of the process is in control; otherwise the process mean has changed. Let $\bar{Y_i} = \frac{1}{n}\sum_{j=1}^{n} Y_{ij}$ be an estimate of $\mu + \delta\sigma$ based on the i-th sample Y_{ij}, $j = 1, 2, ..., n$. When the in-control μ and σ are known, the process mean can be monitored by plotting $\bar{Y}_i$ on a location control chart with respective upper and lower control limits

$$UCL = \mu + C_n\sigma/\sqrt{n},\ LCL = \mu - C_n\sigma/\sqrt{n}, \tag{1.2}$$

where C_n is the factor such that for a chosen type I error probability α we have

$$P(LCL \leq \bar{Y}_i \leq UCL) = 1 - \alpha.$$

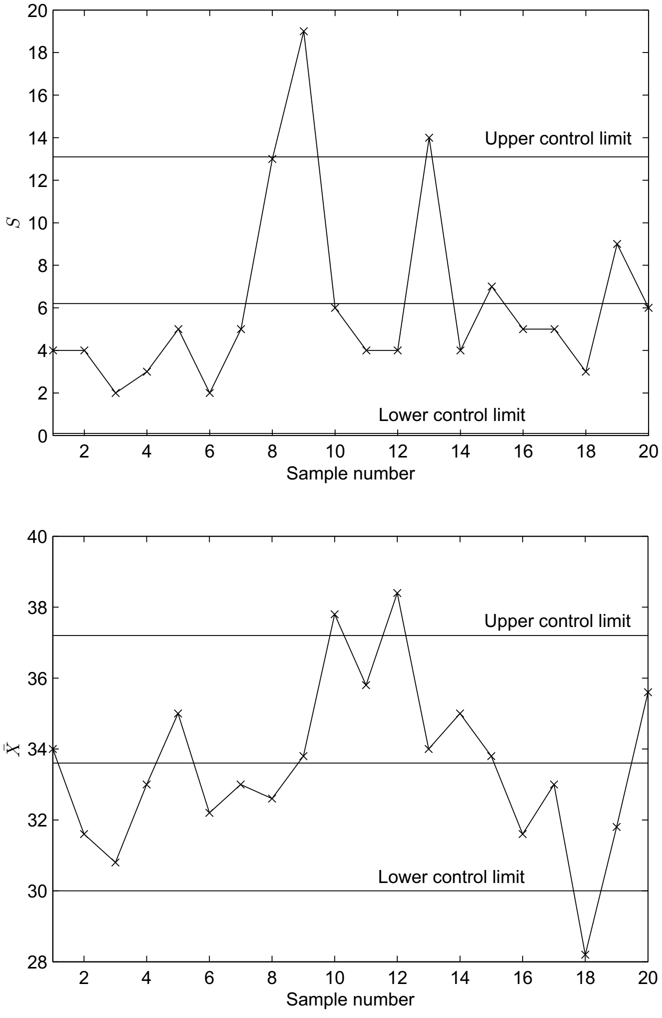

Figure 1.1: Standard deviation and location control chart

When $\bar{Y}_i$ falls within the control limits, the location of the process is deemed to be in control.

The performance of the spread control chart is evaluated in the same way as that of the location control chart. We define E_i as the event that $\hat{\sigma}_i$ ($\bar{Y}_i$) falls beyond the control limits, $P(E_i)$ as the probability that $\hat{\sigma}_i$ ($\bar{Y}_i$) falls beyond the limits and RL as the run length, i.e. the number of samples drawn until the first $\hat{\sigma}_i$ ($\bar{Y}_i$) falls beyond the limits. When σ (μ, σ) is known, the events E_i are independent, and therefore RL is geometrically distributed with parameter $p = P(E_i) = \alpha$. It follows that the average run length (ARL) is given by $1/p$ and that the standard deviation of the run length (SDRL) is given by $\sqrt{1-p}/p$.

In practice, the in-control process parameters are usually unknown. Therefore, they must be estimated from k samples of size n taken when the process is assumed to be in control. This stage in the control charting process is called Phase I (cf. Woodall and Montgomery (1999) and Vining (2009)). The monitoring stage is denoted by Phase II. The samples used to estimate the process parameters are denoted by X_{ij}, $i = 1, 2, ..., k$ and $j = 1, 2, ..., n$. Define $\hat{\sigma}$ and $\hat{\mu}$ as the unbiased estimates of σ and μ respectively, based on the X_{ij}. The control limits are estimated by

$$\widehat{UCL} = U_n\hat{\sigma},\ \widehat{LCL} = L_n\hat{\sigma} \tag{1.3}$$

for the standard deviation control chart and

$$\widehat{UCL} = \hat{\mu} + C_n\hat{\sigma}/\sqrt{n},\ \widehat{LCL} = \hat{\mu} - C_n\hat{\sigma}/\sqrt{n} \tag{1.4}$$

for the location control chart.

Note that U_n, L_n and C_n in (1.3) and (1.4) are not necessarily the same as in (1.1) and (1.2) and might be different even when the probability of signaling is the same. Below, we describe how we evaluate the standard deviation control chart with estimated parameters. The location chart with estimated parameters is evaluated in the same way. Let F_i denote the event that $\hat{\sigma}_i$ is above $\widehat{UCL}$ or below $\widehat{LCL}$. We define $P(F_i|\hat{\sigma})$ as the probability that sample i generates a signal given $\hat{\sigma}$, i.e.

$$P(F_i|\hat{\sigma}) = P(\hat{\sigma}_i < \widehat{LCL} \text{ or } \hat{\sigma}_i > \widehat{UCL}|\hat{\sigma}). \tag{1.5}$$

Given $\hat{\sigma}$, the distribution of the run length is geometric with parameter $P(F_i|\hat{\sigma})$. Consequently, the conditional ARL is given by

$$E(RL|\hat{\sigma}) = \frac{1}{P(F_i|\hat{\sigma})}. \tag{1.6}$$

In contrast with the conditional RL distribution, the unconditional RL distribution takes into account the random variability introduced into the charting procedure through parameter estimation. It can be obtained by averaging the conditional RL distribution over all possible values of the parameter estimates. The unconditional p is

$$p = E(P(F_i|\hat{\sigma})), \tag{1.7}$$

the unconditional average run length is

$$ARL = E(\frac{1}{P(F_i|\hat{\sigma})}) \tag{1.8}$$

and the unconditional standard deviation of the run length is determined by

$$\begin{aligned} SDRL &= \sqrt{Var(RL)} \\ &= \sqrt{E(Var(RL|\hat{\sigma})) + Var(E(RL|\hat{\sigma}))} \\ &= \sqrt{2E(\frac{1}{p(F_i|\hat{\sigma})})^2 - (E\frac{1}{p(F_i|\hat{\sigma})})^2 - E\frac{1}{p(F_i|\hat{\sigma})}}. \end{aligned} \tag{1.9}$$

Quesenberry (1993) showed that for the $\overline{X}$ and X control charts the unconditional ARL is higher than in the (μ,σ)-known case. Furthermore, a higher in-control ARL is not necessarily better because the RL distribution will reflect an increased number of short RL's as well as an increased number of long RL's. He concluded that, if limits are to behave like known limits, the number of samples (k) in Phase I should be at least 400/(n-1) for $\overline{X}$ control charts and 300 for X control charts. Chen (1998) studied the unconditional RL distribution of the standard deviation control chart under normality. He showed that if the shift in the standard deviation in Phase II is large, the impact of parameter estimation is small. In order to achieve a performance comparable with known limits, he recommended taking at least 30 samples of size 5 and updating the limits when more samples become available. For permanent limits, at least 75 samples of size 5 should be used. Thus, the situation is somewhat better than for the $\overline{X}$ control chart with both process mean and standard deviation estimated.

1.2 Contributions and thesis outline

Jensen et al. (2006) conducted a literature survey of the effects of parameter estimation on control chart properties and identified the following issue for future research: *"The effect of using robust or other alternative estimators has not been studied thoroughly. Most evaluations of performance have considered standard estimators based on the sample mean and the standard deviation and have used the same estimators for both Phase I and Phase II. However, in Phase I applications it seems more appropriate to use an estimator that will be robust to outliers, step changes and other data anomalies. Examples of robust estimation methods in Phase I control charts include Rocke (1989), Rocke (1992), Tatum (1997), Vargas (2003) and Davis and Adams (2005). The effect of using these robust estimators on Phase II performance is not clear, but it is likely to be inferior to the use of standard estimates because robust estimators are generally not as efficient"* (Jensen et al. 2006, p. 360). This recommendation is the main subject of the thesis. In particular, we will study alternative estimators in Phase I and we will study the impact of these estimators on the performance of the Phase II control chart.

Chen (1998) studied the standard deviation control chart when σ is estimated by the pooled sample standard deviation ($\tilde{S}$), the mean sample standard deviation ($\bar{S}$) or the mean sample range ($\bar{R}$) under normality. He showed that the performance of the charts based on $\tilde{S}$ and $\bar{S}$ is almost identical, while the performance of the chart based on $\bar{R}$ is slightly worse. Rocke (1989) proposed robust control charts based on the 25% trimmed mean of the sample ranges, the median of the sample ranges and the mean of the sample interquartile ranges in contaminated Phase I situations. Moreover, he studied the use of a two-stage procedure whereby the initial chart is constructed first and then subgroups that seem to be out of control are excluded. Rocke (1992) gave the practical details for the construction of these charts. Wu et al. (2002) considered three alternative statistics for the sample standard deviation, namely the median of the absolute deviation from the median (MDM), the average absolute deviation from the median (ADM) and the median of the average absolute deviation (MAD), and investigated their effect on $\overline{X}$ control chart performance. They concluded that, if there are no or only a few contaminations in the Phase I data, ADM performs best. Otherwise, MDM is the best estimator. Riaz and Saghir (2007 and 2009) showed that the statistics for the sample standard deviation based on Gini's mean difference and the ADM are robust against non-normality. However, they only considered the situation where a large number of samples is available in Phase I and did not consider contaminations in Phase I. Tatum (1997) clearly distinguished two types of disturbances: diffuse and localized. Diffuse disturbances are outliers that are spread over multiple samples whereas localized disturbances affect all observations in a single sample. He proposed a method, constructed around a variant of the biweight A estimator, that is resistant to both diffuse and localized disturbances. A result of the inclusion of the biweight A estimator is, however, that the method is relatively complicated in its use. Apart from several range-based methods, Tatum did not compare his method with other methods for Phase I estimation. Finally, Boyles (1997) studied the dynamic linear model estimator for individuals charts (see also Braun and Park (2008)).

In Chapter 2 we compare an extensive number of Phase I estimators that have been presented in the literature and a number of variants of these statistics. We study their effect on the Phase II performance of the standard deviation control chart. The estimators considered are $\tilde{S}$, $\bar{S}$, the 25% trimmed mean of the sample standard deviations, the mean of the sample standard deviations after trimming the observations in each sample, $\bar{R}$, the sample interquartile range, Gini's mean difference, the MDM, the ADM, the MAD, and the robust estimator of Tatum (1997). Moreover, we propose a robust estimation method based on the mean absolute deviation from the median supplemented with a simple screening method. The performance of the estimators is evaluated by assessing the mean squared error (MSE) of the estimators under normality and in the presence of various types of contaminations. Finally, we assess the Phase II performance of the control charts by means of a simulation study.

Most of the standard deviation estimators presented in Chapter 2 are robust against *either* diffuse disturbances, i.e. outliers spread over the samples, *or* localized disturbances, which affect an entire sample. In Chapter 3 we therefore propose an algorithm that is robust against *both* types of disturbances. The method is compared with the pooled standard deviation (because this estimator is most efficient under normality), the robust estimator of Tatum (1997) and several adaptive trimmers. The performance of the estimators is evaluated by assessing the MSE of the estimators in several situations. Furthermore, we derive factors for the Phase II limits of the standard deviation control chart and assess the performance of the Phase II control charts by means of a simulation study.

As noted earlier, the dispersion parameter of the process is controlled first, followed by the location parameter.

So far the literature has proposed several alternative robust location estimators. Rocke (1989) proposed the 25% trimmed mean of the sample means, the median of the sample means and the mean of the sample medians. Rocke (1992) followed with the practical details for the construction of the corresponding charts. Alloway and Raghavachari (1991) constructed a control chart based on the Hodges-Lehmann estimator. Tukey (1997) and Wang et al. (2007) developed the trimean estimator, which is defined as the weighted average of the median and the two other quartiles. Finally, Jones-Farmer et al. (2009) proposed a rank-based Phase I location control chart. Based on this control chart, they define the in-control state of a process and identify an in-control reference dataset to estimate the location parameter.

In Chapter 4 we consider several robust location estimators as well as several estimation methods based on a Phase I analysis, whereby a control chart is used to study a historical dataset retrospectively and thus identify disturbances. In addition, we propose a new type of Phase I analysis. The methods are evaluated in terms of their MSE and their effect on $\overline{X}$ Phase II control chart performance. We consider situations where the Phase I data are uncontaminated and normally distributed, as well as various types of contaminated Phase I situations.

The results of Chapter 4 indicate that the $\overline{X}$ Phase II control chart (with σ known) based on the new estimation method performs well under normality and outperforms the other charts when contaminations are present in Phase I. However, the results indicate that the effect of estimating the process location on the performance of the $\overline{X}$ Phase II control chart is more limited than the effect of the standard deviation estimator. Chapter 5 therefore looks at the effect of alternative standard deviation estimators under various Phase I scenarios.

In Chapter 5 we develop an estimation method to derive the standard deviation for the $\overline{X}$ control chart when both μ and σ are unknown. Apart from the new method, several alternative estimation methods are included in the comparison. The methods are evaluated in terms of their MSE and their effect on $\overline{X}$ Phase II control chart performance. We again consider the situation where the Phase I data are uncontaminated and normally distributed, as well as various types of contaminated Phase I situations.

The material presented in Chapters 2-5 has led to four papers in various stages of publication. The analysis in Chapter 2 has been published in the *Journal of Quality Technology* (Schoonhoven et al. (2011b)). A follow-up paper based on Chapter 3 has been accepted for publication in *Technometrics*, with minor revisions (Schoonhoven and Does (2011a)). The work in Chapter 4 has been published in the *Journal of Quality Technology* (Schoonhoven et al. (2011a)) and a follow-up article based on the material in Chapter 5 has been submitted to the *Journal of Quality Technology* (Schoonhoven and Does (2011b)).