Designing $\overline{X}$ Charts for Known Autocorrelations and Unknown Marginal Distribution

Huifen Chen and Yuyen Cheng
Department of Industrial and Systems Engineering, Chung-Yuan University
200 Chung-Pei Rd., Chungli, 32023, TAIWAN

Abstract

In the design of the $\overline{X}$ control chart, both the sample size m of $\overline{X}$ and the control-limit factor k (the number of standard deviations from the center line) must be determined. We address this problem under the assumption that the quality characteristic follows an autocorrelated process with known covariance structure but unknown marginal distribution shape. We propose two methods for determining m and k, chosen to minimize the out-of-control ARL (average run length) while maintaining the in-control ARL at a specified value. Method 1 calculates the ARL values as if the sample means were independent normal random variables; Method 2 calculates the ARL values as if the sample means were an AR(1) process. Method 2 outperforms Method 1 when the correlation and mean shift are both high. We also modify Methods 1 and 2 with a minimum sample size of 30; the modification moves the in-control ARL closer to the specified value. Our numerical results show that the modified Method 2 performs better than two previous design procedures, especially when the correlation is high.

Keywords: average run length, covariance stationary, optimization, SPC

1 Introduction

Control charts send a signal when observed process data $\{X_t, t = 1, 2, \ldots\}$ appear to be, in some sense, out of control. When *out of control* is based on a shift in the process mean, it is natural to send a signal when the data stray far from the target mean μ. When the process data $\{X_t\}$ are assumed to be stationary, except for the instant when the mean shift occurs, an $\overline{X}$ chart can be used, where $\overline{X}_j = m^{-1} \sum_{t=(j-1)m+1}^{jm} X_t$ is the jth sample average of m contiguous discrete-time observations, $j = 1, 2, \ldots$ (or, analogously for continuous-time data, the time average $\overline{X}_j = m^{-1} \int_{(j-1)m}^{jm} X_t dt$, where X_t is the quality measurement at time t). The signal is sent when the first $\overline{X}_j$ does not lie between the lower and upper control limits, typically $\mu - d_l$ and $\mu + d_u$, where d_l and d_u are positive constants. The design of $\overline{X}$ charts involves determining appropriate values of the sample size m and these two constants. The chart is *symmetric* when $d_l = d_u$.

We consider the design of symmetric $\overline{X}$ control charts when the in-control process data $\{X_t\}$ are assumed to be stationary with known mean μ, known standard deviation σ, known autocorrelations ρ_h for $h = 1, 2, \ldots$, but with unknown marginal-distribution shape. In addition, the process mean is assumed to shift to $\mu + \delta\sigma$ with no change in values of the other moments when the process goes out of control. These assumptions imply that for any sample size m the variance of the sample mean is

$$\sigma_{\overline{X}}^2 = \frac{\sigma^2}{m}\Big[1 + 2\sum_{h=1}^{m-1}(1 - \frac{h}{m})\rho_h\Big] \tag{1}$$

both when the process is in control and after it is out of control. We define the optimal $\overline{X}$ chart to minimize ARL_δ subject to a specified value L for ARL_0. Here ARL_0 denotes the in-control average run length (ARL), i.e., the expected number of observations X_t until the signal is sent when the process is in control with mean μ; and ARL_δ denotes the out-of-control ARL, i.e., the ARL when the process data have mean $\mu + \delta\sigma$, where δ is any non-zero constant. We seek design algorithms that, given values μ, σ, $\{\rho_h\}$, δ, and L, compute the optimal integer sample size m^* and optimal positive real-number distance d^* so that the signal is sent when the first sample mean arising from a sample of size m^* lies outside $\mu \pm d^*$. To keep our discussion independent of the scaling of X_t, we use k rather than d, where $d = k\sigma_{\overline{X}}$.

Three previous papers also design control charts for the mean μ under the assumption that the process data are autocorrelated with a known in-control mean, known variance, and known autocorrelations, but with unknown marginal distribution. Runger and Willemain (1995) design $\overline{X}$ charts by choosing the sample size m to obtain lag-1 sample-mean autocorrelation $\mathrm{Corr}(\overline{X}_j, \overline{X}_{j+1}) \approx 0.1$ and then choosing k to obtain $\mathrm{ARL}_0 = L$ under the assumption that the sample means are independent and identically distributed (iid) normal random variables; we refer to this $\overline{X}$ chart as R&W. Kim et al. (2006) design MFC (model-free CuSum) charts with sample size $m = 1$ and control limits that obtain $\mathrm{ARL}_0 = L$ using a functional central limit theorem and the known asymptotic variance constant. Kim et al. (2007) design DFTC (distribution-free tabular CuSum) charts that use sample means arising from samples of size m to obtain lag-1 sample-mean autocorrelation $\mathrm{Corr}(\overline{X}_j, \overline{X}_{j+1}) \leq 0.5$ and control limits that obtain $\mathrm{ARL}_0 = L$.

Procedures that assume known autocorrelations and unknown marginal distribution can be useful in at least three ways. First, as mentioned in the conclusions of Kim et al. (2007), the procedures "can be used as the foundation for the ultimate development of an SPC (statistical process control) procedure for correlated processes that can be directly applied in practice." Second, the procedures *can* be applied directly; for example, because autocorrelations can arise from the logical flow of the system (as in queueing data) while the marginal distribution depends upon the stochastic components of the system (such as service times), a change in the system (such as a new customer class) can have known autocorrelations but unknown marginal distribution. Third, the procedures extend the important special case in which only iid data are considered.

Applications that use autocorrelated data are described in Pandit and Wu (1983), Hahn (1989), Koo and Case (1990), Tucker et al. (1993), English and Case (1994), Wardell et al. (1994), Runger and Willemain (1995), Faltin et al. (1997), and Boyles (2000).

The structure of this paper is as follows. Section 2 describes the optimization model for determining the values of m and k. Two methods for computing the optimal values of m and k are proposed and are further modified with a sample size at least 30 so that the in-control ARL is closer to the specified value. Numerical results show that the modified Method 2 performs better than the modified Method 1 when the autocorrelation is high and the mean shift is moderate. Section 3 empirically compares the performance of the modified Method 2 to that of the R&W and DFTC charts; the MFC chart, dominated by the DFTC chart, is not considered. Section 4 gives our summary and conclusions.

2 A New Design Procedure for the $\overline{X}$ Chart

The $\overline{X}$ chart is a useful SPC tool for monitoring the process mean. It works as follows. First, the products are divided into consecutive samples of m quality measurements $X_1, \ldots, X_m$. For each sample, the sample mean $\overline{X}$ is computed. If it falls outside the control limits $\mu \pm k\sigma_{\overline{X}}$, an out-of-control signal is recorded in the control chart. (Recall that $\sigma_{\overline{X}}$ can be computed using Equation (1) because σ and $\{\rho_h\}$ are known.) The quality-control engineers then determine whether there is an assignable cause. If an assignable cause is identified, appropriate action is taken to tune the production process and restore the in-control state.

The control-chart design parameters m and k are chosen to minimize the out-of-control ARL while maintaining the in-control ARL at a specified value L. We describe the optimization model for determining the values of m and k in Section 2.1. Sections 2.2 and 2.3 propose two methods for computing the optimal values of m and k. Section 2.4 compares Methods 1 and 2 empirically and presents modifications of both methods that bring the in-control ARL closer to L.

2.1 The Optimization Model for Setting m and k

A common performance measure in control-chart design is the ARL. A good control scheme results in a long ARL when the process is in control and a short ARL when the process is out of control. We would like the values of m and k to depend on both the in-control and out-of-control ARLs. Therefore, we choose the values of the $\overline{X}$-chart design parameters m and k to minimize the out-of-control ARL while keeping the in-control ARL at a specified value L. That is, the values of m and k are chosen to satisfy the following optimization criterion:

$$\begin{aligned} \min \quad & \text{ARL}_\delta \\ \text{s.t.} \quad & \text{ARL}_0 = L \ , \\ & m \in \{1, 2, \ldots, L\}, \quad k > 0. \end{aligned} \tag{2}$$

Both ARL_0 and ARL_δ are measured in terms of the number of observations rather than the number of $\overline{X}$ points. For example, suppose that the quality measurements are independently and normally distributed, $m = 5$, and $k = 3$. Then the average number of $\overline{X}$ charting points until a false alarm occurs is $[1+\Phi(-3)-\Phi(3)]^{-1}= 370$, where $\Phi(\cdot)$ is the cumulative distribution function (cdf) of the standard normal distribution (Montgomery, 2005). Hence, the number of observations until a false alarm occurs is $\text{ARL}_0 = (5)(370) = 1850$. The fixed value of the in-control ARL enables us to seek the values of the design parameters that result in the lowest possible out-of-control ARL. This is reasonable because when δ is large, the shift is easier to detect, and hence the sample size m need not be large.

Ideally the values of m and k would be chosen to maximize ARL_0 and minimize ARL_δ simultaneously. Unfortunately, for a fixed value of k, as the sample size m increases, both ARL_0 and ARL_δ increase. Similarly, for any fixed m, as k increases, both ARL_0 and ARL_δ increase. Therefore, there are no paired values of m and k that would simultaneously optimize both ARL_0 and ARL_δ.

Figure 1 illustrates the tradeoff between ARL_0 and ARL_δ for independently and normally distributed quality measurements $\{X_t\}$ and the shift $\delta = 2$. Contour plots are shown for $\text{ARL}_0 = 10^2$, 10^3, and 10^4 (dot curves) and $\text{ARL}_\delta = 2, 3, \ldots, 10$ (solid curves), calculated using Equation (3) in Section 2.2 below. The triangles in the ARL_0 curves and the circles in the ARL_δ curves denote all possible (m, k) combinations for the contour plots. Clearly, as we increase the sample size m while keeping k constant (e.g., by going from Point A to Point B), both ARL_0 and ARL_δ increase. Similarly, as we increase k while keeping m constant (e.g., by going from Point B to Point C), both ARL_0 and ARL_δ increase.

Since no control scheme can optimize ARL_0 and ARL_δ simultaneously, we seek the control scheme that minimizes the value of ARL_δ for a fixed ARL_0. The resulting values of

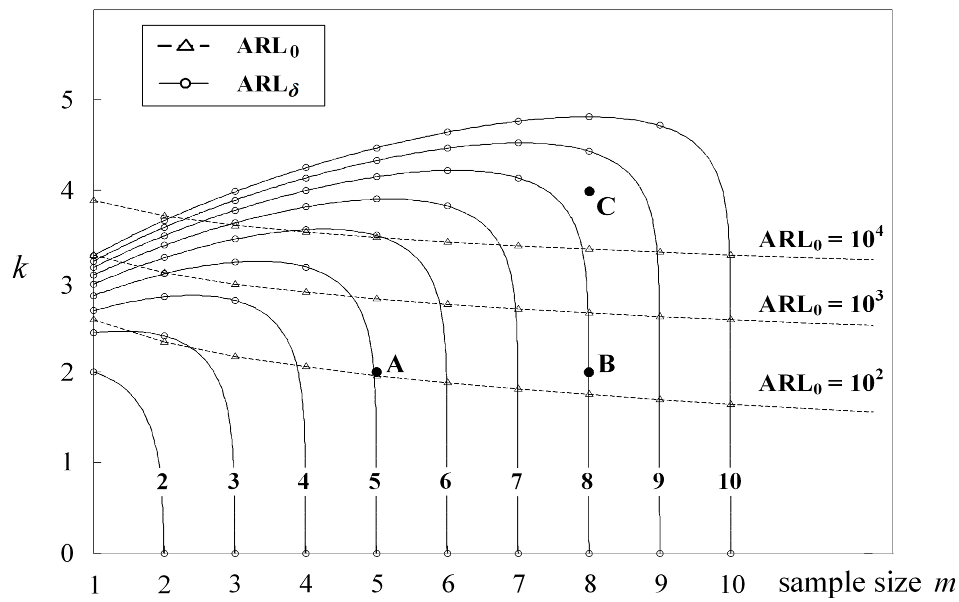

Figure 1: Contour plots on the $(m,\, k)$ plane for $\mathrm{ARL}_0 = 10^2$, 10^3, and 10^4 and $\mathrm{ARL}_\delta = 2$, 3,. . ., 10, where the quality characteristic is iid normal and the shift $\delta = 2$

m and k then satisfy Equation (2). Figure 2 shows these optimal values for $\mathrm{ARL}_0 = 10^3$. The optimal point is $(m,\, k) = (3,\, 2.97)$, corresponding to the minimum $\mathrm{ARL}_\delta = 4.35$.

The ARLs in the optimization model depend on the marginal distribution, as do the optimal sample size m^* and the optimal number k^* of standard deviations from the center line. Because the marginal distribution is assumed to be unknown, Methods 1 and 2 below use the central limit theorem to obtain approximate normality of the sample means.

2.2 Method 1: Independent Normal Sample Means

Let $\{\overline{X}_1,\, \overline{X}_2, \ldots\}$ denote successive nonoverlapping sample means, each arising from a sample of size m, to be plotted in the $\overline{X}$ control chart. Method 1 assumes that the sample means are iid normal, corresponding to a large sample size m with mean $\mathrm{E}(\overline{X})$ and variance $\sigma^2_{\overline{X}}$, as shown in Equation (1). In this case, the run length in units of m follows a geometric distribution and the out-of-control ARL is

$$\mathrm{ARL}_\delta = \frac{m}{\mathrm{P}\{\overline{X} \notin \mu \pm k\sigma_{\overline{X}} \mid \mathrm{E}(\overline{X}) = \mu + \delta\sigma\}} = \frac{m}{1 + \Phi(-k - \delta\sqrt{c}) - \Phi(k - \delta\sqrt{c})}\ , \qquad (3)$$

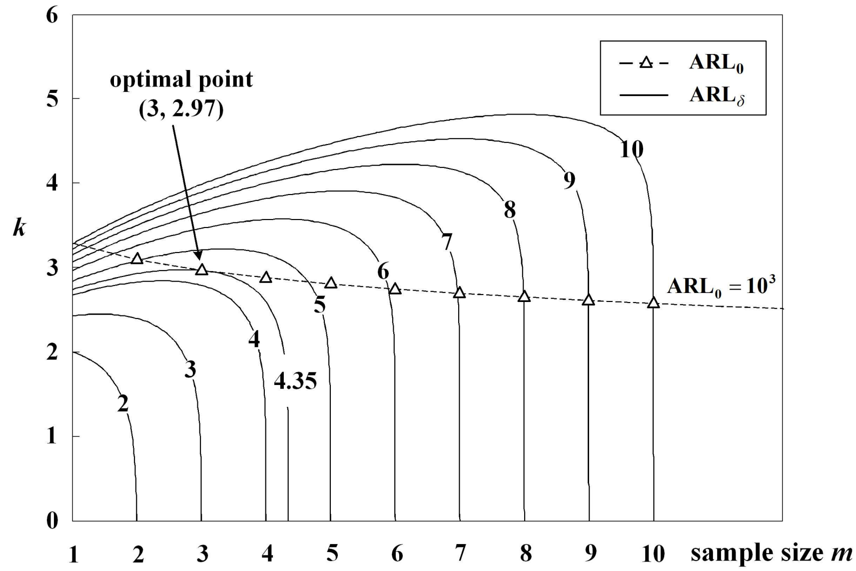

Figure 2: ARL$_\delta$ contour plots and the corresponding optimal point on the (m, k) plane for the constraint ARL$_0 = 10^3$, where the quality measurements are iid normal and the shift $\delta = 2$

where $c = \sigma^2/\sigma^2_{\overline{X}} = m/[1 + 2\sum_{h=1}^{m-1} (1 - h/m)\rho_h]$. (See Runger and Willemain, 1995.) When the process is in control, the ARL is $\text{ARL}_0 = m/[1 + \Phi(-k) - \Phi(k)] = m/[2\Phi(-k)]$. (Recall that both ARL$_0$ and ARL$_\delta$ are measured in terms of the number of observations.)

Using the above formulas for ARL$_0$ and ARL$_\delta$, we find the values of m and k that satisfy the optimization model in Equation (2). The value of k that meets the constraint $\text{ARL}_0 = L$ is

$$k = -\Phi^{-1}[m/(2L)] \,, \tag{4}$$

where Φ^{-1} is the inverse function of Φ. Substituting $k = -\Phi^{-1}[m/(2L)]$ into Equation (3), we can find the value of m that minimizes ARL$_\delta$ and compute the corresponding k value.

In summary, given $\{\rho_h\}$, δ, and L, Method 1 performs a one-dimensional search on m to determine $m^* = \text{argmin}_m\{\text{ARL}_\delta(m)\}$, where $\text{ARL}_\delta(m)$ is computed using Equation (3) with sample size m and factor $k = -\Phi^{-1}[m/(2L)]$. Using Equation (4), the method then determines the optimal value $k^* = -\Phi^{-1}[m^*/(2L)]$. Because there are local minima, an explicit search over $\{1, 2, \ldots, L\}$ is necessary to determine m^*, unless care is taken to develop a more-efficient search. Computation time is negligible.

2.3 Method 2: AR(1) Sample Means

The advantage of Method 1 is its simplicity. However, there are drawbacks. One is that Method 1 ignores correlations between successive sample means, which may be high when m is small and the autocorrelations $\{\rho_h\}$ are high. This phenomenon may cause the values of m and k obtained from Method 1 to be far from the true optimal values. Method 2 seeks to remedy this problem by modeling the sample means $\{\overline{X}_1, \overline{X}_2, \ldots\}$ as an AR(1) process with its parameters matching the mean (μ or $\mu + \delta\sigma$), variance, and lag-1 autocorrelation of the sample means. We choose the AR(1) model because the corresponding ARL_0 and ARL_δ can be computed numerically by a Markov chain approach. The values of m and k satisfying Equation (2) can then be computed numerically as well. By expanding on this approach, one can devise more complicated models (e.g., AR(p) with $p > 1$) that may better match the autocovariance structure of the sample means. However, the corresponding ARL_0 and ARL_δ would be harder to compute and may need to be estimated via simulation experiments.

The AR(1) data process $\{Z_t\}$ is a time series process with $(Z_t - \mu_z) = \phi_z(Z_{t-1} - \mu_z) + \epsilon_t$, where $|\phi_z| < 1$, $\mu_z = \mathrm{E}(Z_t)$ for $t = 1, 2, \ldots$, and the random error ϵ_t is independently distributed as $\mathrm{N}(0, \sigma_\epsilon^2)$. The marginal distribution of the AR(1) process is $\mathrm{N}(0, \sigma_\epsilon^2/(1-\phi_z^2))$ and the lag-h autocorrelation is $\phi_z^{|h|}$. The AR(1) model has three parameters: the AR(1) marginal mean μ_z, the lag-1 autocorrelation ϕ_z, and the variance σ_ϵ^2 of the random error. To fit an AR(1) model to the sample means, we impose three requirements on the AR(1) parameters:

$$
\begin{aligned}
\mu_z &= \mathrm{E}(\overline{X}) = \begin{cases} \mu & \text{if } \mathrm{ARL}_0 \text{ in Equation (2) is computed} \\ \mu + \delta\sigma & \text{if } \mathrm{ARL}_\delta \text{ in Equation (2) is computed} \end{cases}, \\
\phi_z &= \mathrm{Corr}(\overline{X}_1, \overline{X}_2) = \frac{\sum_{h=1}^{m} h\,\rho_h + \sum_{h=1}^{m-1} h\,\rho_{2m-h}}{m + 2\sum_{h=1}^{m-1}(m-h)\rho_h}, \\
\sigma_\epsilon^2 &= (1-\phi_z^2)\sigma_{\overline{X}}^2,
\end{aligned}
\tag{5}
$$

where the lag-1 autocorrelation ϕ_z can be computed analytically, because the autocorrelations $\{\rho_h\}$ are known.

Since the sample means $\{\overline{X}_1, \overline{X}_2, \ldots\}$ are assumed to follow an AR(1) process, the average run length can be computed numerically. Lucas and Saccucci (1990) propose a Markov-chain approximation for computing the average run length of EWMA (exponentially weighted moving average) control charts. Since the AR(1) process behaves like an EWMA for independent normal data, we can use the Markov-chain approximation to compute ARL_0 and ARL_δ subject to Equation (2).

In summary, given μ, σ, $\{\rho_h\}$, δ, and L, Method 2 performs a two-dimensional search on (m, k) to determine (m^*, k^*) that minimizes $\text{ARL}_\delta(m, k)$ subject to $\text{ARL}_0(m, k) = L$ over the positive integers m and positive real numbers k, where $\text{ARL}_0(m, k)$ and $\text{ARL}_\delta(m, k)$ are the ARL_0 and ARL_δ values corresponding to the sample size m and factor k. For any pair (m, k), Method 2 computes $\sigma_{\overline{X}}(m)$ using Equation (1), $\phi_z(m)$ and $\sigma_\epsilon^2(m)$ using Equation (5), and $\text{ARL}_0(m, k)$ and $\text{ARL}_\delta(m, k)$ using Lucas and Saccucci (1990) with means μ and $\mu + \delta\sigma$, respectively. (For clarity, we denote $\sigma_{\overline{X}}$, ϕ_z, and σ_ϵ as functions $\sigma_{\overline{X}}(m)$, $\phi_z(m)$, and $\sigma_\epsilon(m)$ of the sample size m.) The two-dimensional search can, for example, enumerate m over the set $\{1, 2, \ldots, L\}$, for each m finding $k^*(m)$ such that $\text{ARL}_0(m, k^*(m)) = L$, as determined by a one-dimensional root-finding search on k. Though the computation time is longer than for Method 1, the exact time required and Markov-chain approximation error depend upon how the state space is truncated.

2.4 Comparisons of Methods 1 and 2

We begin this section by testing the performance of Methods 1 and 2, first using an AR(1) process and then using an ARTA (AutoRegressive To Anything) process (Cario and Nelson, 1996). We show that for both processes, Method 2 performs better than Method 1, though Method 1 has the advantage of computational simplicity. Moreover, Method 2 outperforms Method 1 when the marginal distribution of the quality characteristic X is normal and both the autocorrelation and shift are large. When the marginal distribution is nonnormal and the shift δ is large, the ARL_0 values computed by Methods 1 and 2 differ from the specified value L.

At the end of the section, we modify Methods 1 and 2 by requiring the sample size to be at least 30. In both modified methods, the ARL_0 values are closer to L than in the corresponding unmodified method. Numerical results show that the modified Method 2 performs slightly better than the modified Method 1.

The first numerical comparison experiment employs AR(1) data. Suppose that when the process is in control, the quality characteristic measurements $\{X_t\}$ follow an AR(1) process with mean μ, variance σ^2, and lag-1 autocorrelation ρ_1. Kang and Schmeiser (1987) have shown that the sample means $\{\overline{X}_1, \overline{X}_2, \ldots\}$ then follow an ARMA(1, 1) process. In this special case, we can compute the ARL numerically using a two-dimensional Markov-chain approximation approach presented in Jiang et al. (2000).

The following parameter values are employed in the AR(1) comparison test: the lag-1 autocorrelation $\rho_1 \in \{0, 0.25, 0.5, 0.9, 0.95, 0.99\}$, the shift $\delta \in \{0.25, 0.5, 0.75, 1, 1.5, 2, 2.5, 3, 4\}$, and the desired average run length $L = 10000$, for a total of 54 ($= 6 \cdot 9$) experimental points. (Note that the values of ARLs, and hence those of the Method-1, Method-2, and true optimal solutions m^* and k^*, are independent of the location parameter μ and scale parameter σ.) Table 1 shows the results. Columns 1 and 2 list the values of ρ_1 and δ; columns 3 to 6, the Method-1 m^* and k^* values and those of the corresponding ARL_0 and ARL_δ; columns 7 to 10, the Method-2 m^* and k^* values and those of the corresponding ARL_0 and ARL_δ; columns 11 to 13, the actual m^* and k^* values and the value of the corresponding ARL_δ. When the actual optimal values of m and k are used, the corresponding ARL_0 value is exactly equal to 10000.

The results in Table 1 show that the performance of Method 2 is equal to or better than that of Method 1 in all cases, and the Method-2 m^* and k^* values result in an ARL_δ that is very close to the true optimal value. (Note that because of rounding effects, in some cases the Method-1 and/or Method-2 m^* and k^* values differ from the true m^* and k^* values, but the corresponding ARL_δ values are the same.) This outcome is foreseeable: Method 2 matches the lag-1 autocorrelation of the sample means and hence yields a lower ARL_δ and an ARL_0 closer to 10000. The relative performance of Method 1 declines as ρ_1 increases. When $\rho_1 = 0$, the sample means are iid normal, and both methods yield the actual optimal values of m and k. When ρ_1 is small or moderate, Method 1 performs almost as well as Method 2 for all values of δ. As ρ_1 increases, the autocorrelation between adjacent sample means increases, and hence Method 2 outperforms Method 1.

Likewise, the performance of Method 1 deteriorates as δ increases. When δ is small,

Table 1: The Method-1 and Method-2 design outputs and the true optimal solutions for AR(1) processes with $\rho_1 = 0$, 0.25, 0.5, 0.9, 0.95, 0.99, and $L = 10000$

		Method 1				Method 2				True optimal solutions		
ρ_1	δ	m	k	ARL_0	ARL_δ	m	k	ARL_0	ARL_δ	m	k	ARL_δ
0	0.25	133	2.476	10^4	202					133	2.476	202
	0.5	45	2.841	10^4	65					45	2.841	65
	0.75	23	3.048	10^4	32					23	3.048	32
	1	14	3.195	10^4	20					14	3.195	20
	1.5	7	3.390	10^4	9.7	same as for Method 1				7	3.390	9.7
	2	4	3.540	10^4	5.9					4	3.540	5.9
	2.5	3	3.615	10^4	3.9					3	3.615	3.9
	3	2	3.719	10^4	2.9					2	3.719	2.9
	4	1	3.891	10^4	1.8					1	3.891	1.8
0.25	0.25	194	2.338	10^4	302	194	2.338	10^4	302	194	2.338	302
	0.5	67	2.711	10^4	98	67	2.711	10^4	98	67	2.711	98
	0.75	35	2.920	10^4	50	35	2.920	10^4	50	35	2.920	50
	1	21	3.076	10^4	30	22	3.062	10^4	30	21	3.076	30
	1.5	10	3.291	10^4	15	11	3.264	10^4	15	11	3.264	15
	2	6	3.432	10001	8.6	6	3.432	10^4	8.6	6	3.432	8.6
	2.5	4	3.540	10001	5.6	4	3.540	9999	5.6	4	3.540	5.6
	3	2	3.719	10002	4.3	3	3.615	9999	3.9	3	3.615	3.9
	4	1	3.891	10013	2.5	1	3.89	10^4	2.5	2	3.719	2.3
0.5	0.25	295	2.177	10^4	475	296	2.175	10^4	475	296	2.175	475
	0.5	105	2.559	10^4	159	106	2.556	10^4	159	106	2.556	159
	0.75	55	2.776	10^4	81	56	2.770	10^4	81	55	2.776	81
	1	34	2.929	10001	49	35	2.920	10^4	49	34	2.929	49
	1.5	16	3.156	10002	24	17	3.138	10^4	24	17	3.138	24
	2	9	3.320	10004	14	10	3.290	9999	14	10	3.290	14
	2.5	5	3.481	10010	9.1	6	3.431	9998	8.9	6	3.431	8.9
	3	2	3.719	10062	7.0	4	3.540	9975	6.1	4	3.540	6.1
	4	1	3.891	10185	3.5	1	3.886	10^4	3.5	1	3.886	3.5
0.9	0.25	940	1.675	10^4	1754	945	1.672	10^4	1754	944	1.673	1754
	0.5	393	2.061	10001	664	396	2.058	10^4	664	399	2.055	664
	0.75	217	2.296	10002	355	222	2.287	10^4	355	224	2.283	355
	1	137	2.465	10004	222	143	2.449	10^4	222	142	2.452	222
	1.5	64	2.727	10016	110	72	2.687	9999	110	70	2.696	110
	2	1	3.891	16833	130	40	2.877	9997	64	40	2.877	64
	2.5	1	3.891	16833	59	22	3.058	9987	41	25	3.021	41
	3	1	3.891	16833	32	1	3.753	10^4	28	9	3.311	28
	4	1	3.891	16833	14	1	3.753	10^4	13	1	3.753	13
0.95	0.25	1351	1.494	10000	2730	1362	1.490	10^4	2730	1367	1.488	2729
	0.5	616	1.869	10002	1108	625	1.863	10^4	1108	623	1.864	1108
	0.75	352	2.106	10004	610	364	2.092	10^4	610	362	2.095	610
	1	224	2.284	10009	388	236	2.263	10^4	387	236	2.263	387
	1.5	1	3.891	25279	606	120	2.511	9999	194	119	2.514	194
	2	1	3.891	25279	231	66	2.713	9994	113	68	2.704	113
	2.5	1	3.891	25279	108	1	3.634	10^4	78	40	2.871	73
	3	1	3.891	25279	60	1	3.634	10^4	47	1	3.634	47
	4	1	3.891	25279	28	1	3.634	10^4	23	1	3.634	23
0.99	0.25	2285	1.204	10003	5959	2357	1.186	10^4	5957	2346	1.188	5956
	0.5	1	3.891	77270	28740	1443	1.459	10^4	3081	1458	1.454	3081
	0.75	1	3.891	77270	14216	932	1.678	10^4	1856	936	1.676	1856
	1	1	3.891	77270	7369	638	1.851	9999	1235	633	1.855	1235
	1.5	1	3.891	77270	2374	317	2.140	9992	648	330	2.125	648
	2	1	3.891	77270	960	1	3.253	10^4	397	108	2.511	390
	2.5	1	3.891	77270	472	1	3.253	10^4	240	17	2.946	239
	3	1	3.891	77270	273	1	3.253	10^4	156	1	3.253	156
	4	1	3.891	77270	129	1	3.253	10^4	87	1	3.253	87

the Method-1 sample size m^* is large, and hence the iid normal presumption is nearly valid. However, as δ increases, the Method-1 m^* decreases, causing the ARL_0 to deviate increasingly from the specified value of 10000.

Method 2 remedies the shortcomings of Method 1 by modeling the sample means $\{\overline{X}_j\}$ as an AR(1) process so that the correlations among the sample means can be taken into account. The computed values of m and k are either identical to or very close to the true optimal values, except for some minor discrepancies when ρ_1 and δ are both large. When ρ_1 is large and m is small, the modeled AR(1) deviates from the true ARMA(1, 1) process of the sample means, causing the Method-2 approximation error to increase. However, even when δ is large enough to result in a small m^*, the magnitude of the approximation error is not significant.

The second numerical comparison experiment employs ARTA data. The ARTA process $\{X_t\}$ of order p, denoted ARTA(p), is a stationary time series transformed from a standardized Gaussian AR(p) process $\{Z_t\}$: $X_t = F^{-1}(\Phi(Z_t))$, where $F(\cdot)$ is the ARTA(p) marginal cdf. (See Cario and Nelson, 1996, for details.) Cario and Nelson (1998) provide ARTAFACTS and ARTAGEN software for fitting and generating from an ARTA(p) process; we use the ARTAGEN software to generate ARTA data in our simulation experiments. We choose the ARTA(p) implementation because it allows the specification of the autocorrelations $\rho_1, \ldots, \rho_p$ of the first p lags to be arbitrary, provided that certain restrictions are observed. First, the autocorrelations may not include certain values in the interval $[-1, 0)$ for the time series to be stable and for the autocorrelation matrix to be positive definite (Ghosh and Henderson, 2003). Second, when specifying autocorrelation values, we must differentiate between the absolute minimum -1 and the effective minimum, which depends on the marginal cdf F. The effective minimum is equal to the absolute minimum value -1 only when the cdf F is symmetric (Chen, 2001). For other marginal cdfs, the effective minimum value is higher. For example, if the marginal distribution is exponential, the effective minimum correlation is -0.645.

In our numerical experiment, we assume that when the process is in control, the quality measurements $\{X_t\}$ follow an ARTA(1) process with lag-1 autocorrelation ρ_1 and the Student-t marginal distribution with 10 degrees of freedom (skewness = 0 and kurtosis =

Table 2: The Method-1 and Method-2 design outputs and the true optimal solutions for ARTA(1) processes with t_{10} marginal distribution, $\rho_1 = 0$, 0.25, 0.5, 0.7, 0.9, and $L = 1000$ (The boxed $\widehat{\mathrm{ARL}}_0$ values satisfy $|\widehat{\mathrm{ARL}}_0 - L| > 100$.)

		Method 1				Method 2				True optimal solution		
ρ_1	δ	m	k	$\widehat{\mathrm{ARL}}_0$	$\widehat{\mathrm{ARL}}_\delta$	m	k	$\widehat{\mathrm{ARL}}_0$	$\widehat{\mathrm{ARL}}_\delta$	m	k	$\widehat{\mathrm{ARL}}_\delta$
0	0.25	66	1.838	995	115					67	1.832	114
	0.5	26	2.226	981	41					28	2.203	42
	0.75	14	2.457	944	22					14	2.477	22
	1	9	2.612	$\boxed{889}$	14	same as for Method 1				10	2.614	14
	1.5	5	2.807	$\boxed{755}$	7					5	2.911	7
	2	3	2.968	$\boxed{593}$	4					4	3.022	5
	4	1	3.291	$\boxed{234}$	1					2	3.448	2
0.25	0.25	88	1.706	997	161	88	1.706	997	160	87	1.713	160
	0.5	36	2.097	982	61	37	2.086	987	60	36	2.103	61
	0.75	20	2.326	956	32	21	2.308	960	32	22	2.305	33
	1	13	2.484	911	20	13	2.484	911	20	13	2.520	21
	1.5	7	2.697	$\boxed{795}$	10	7	2.697	$\boxed{795}$	10	7	2.788	11
	2	4	2.878	$\boxed{632}$	6	4	2.878	$\boxed{632}$	6	5	2.948	7
	4	1	3.291	$\boxed{237}$	1	1	3.289	$\boxed{237}$	1	2	3.500	2
0.5	0.25	120	1.555	994	235	120	1.555	994	235	125	1.535	235
	0.5	53	1.935	991	93	54	1.927	992	92	53	1.938	93
	0.75	30	2.170	969	51	30	2.170	969	50	30	2.183	51
	1	19	2.346	926	32	20	2.326	930	32	21	2.334	32
	1.5	9	2.612	$\boxed{789}$	16	10	2.575	$\boxed{814}$	16	11	2.616	17
	2	4	2.878	$\boxed{568}$	9	6	2.746	$\boxed{681}$	9	7	2.830	10
	4	1	3.291	$\boxed{255}$	1	1	3.277	$\boxed{249}$	1	3	3.296	3
0.7	0.25	159	1.408	1000	338	159	1.408	1000	338	160	1.403	337
	0.5	77	1.768	995	143	77	1.768	995	143	80	1.751	143
	0.75	45	2.005	981	80	46	1.995	981	80	47	1.994	80
	1	29	2.183	947	51	31	2.157	948	51	29	2.205	52
	1.5	13	2.484	$\boxed{811}$	26	16	2.407	$\boxed{851}$	26	17	2.447	27
	2	1	3.291	$\boxed{302}$	21	8	2.647	$\boxed{695}$	14	10	2.696	16
	4	1	3.291	$\boxed{302}$	2	1	3.239	$\boxed{276}$	2	4	3.177	4
0.9	0.25	227	1.208	1002	582	233	1.192	995	582	243	1.167	584
	0.5	135	1.495	1006	299	142	1.468	994	298	142	1.469	299
	0.75	83	1.734	987	179	91	1.689	987	178	99	1.654	181
	1	52	1.943	966	120	62	1.864	969	118	63	1.870	120
	1.5	1	3.291	$\boxed{558}$	103	31	2.149	$\boxed{894}$	60	33	2.172	63
	2	1	3.291	$\boxed{558}$	48	14	2.429	$\boxed{751}$	31	20	2.404	37
	4	1	3.291	$\boxed{558}$	2	1	3.063	$\boxed{382}$	2	9	2.767	9

4), denoted t_{10}. Since we can not compute the ARL analytically for ARTA data, we estimate it via simulation experiments. The other parameter values used for the comparison experiments are as follows: the ARTA(1) lag-1 autocorrelation $\rho_1 \in \{0, 0.25, 0.5, 0.7, 0.9\}$, the shift $\delta \in \{0.25, 0.5, 0.75, 1, 1.5, 2, 4\}$, and the desired ARL_0 value $L = 1000$, for a total of 35 ($= 5 \cdot 7$) experimental points.

Table 2 shows results for the ARTA(1) process with the t_{10} marginal distribution. The columns in Table 2 are identical to those in Table 1, except that the $\widehat{\text{ARL}}_0$ and $\widehat{\text{ARL}}_\delta$ figures in columns 5, 6, 9, 10, and 13 are estimates instead of exact values. To obtain the estimates, we generated 80,000 observations of the run length based on every (m, k) solution computed by Methods 1 and 2 and rounded the resulting values to the nearest integer. The standard errors of each $\widehat{\text{ARL}}_0$ and $\widehat{\text{ARL}}_\delta$ are around 0.05% of the reported value. See Table S1 in the supplement (Chen and Cheng, 2008). The true optimal solutions in columns 11 and 12 are determined through an exhaustive search procedure, with 160,000 to 640,000 observations of the run length for each iterate (m, k), depending on the run-length variation. Common random-number streams are used.

The results in Table 2 show that Methods 1 and 2 both work well for small δ. However, as δ increases (especially when $\delta > 1$), ARL_0 deviates further and further from the specified value $L = 1000$. In the table, we have marked with boxes all $\widehat{\text{ARL}}_0$ values for which the discrepancy exceeds 100. When δ is large, Methods 1 and 2 both yield small sample sizes. Since these sample sizes are too small to apply the central limit theorem, the sample mean is not normally distributed. As a result, even Method 2 cannot overcome the problems posed by large values of δ.

To get around this problem, we now modify Methods 1 and 2 slightly by requiring that the sample size m be at least 30. Under this constraint, the central limit theorem would be more applicable. We adjust the associated value of k accordingly to satisfy the constraint $\text{ARL}_0 = L$ under the iid normal assumption for Method 1 and the AR(1) assumption for Method 2. Note that our lower bound on the sample size m was chosen arbitrarily. Although "$m \geq 30$" seems to work well for many distributions (e.g., a t distribution with at least 2 degrees of freedom) and is suggested by some statistics textbooks (e.g., Ross, 2004, p. 212), the extent of normality approximation achieved by using this lower bound depends on the distribution shape of the quality measurement.

Table 3 lists the results for these modified Methods 1 and 2. All m values that were less than 30 in Table 2 are raised to 30 in Table 3. The modifications occur when the lag-1 autocorrelation ρ_1 is small and the shift δ is large. The resulting $\widehat{\mathrm{ARL}}_0$ values are very close to $L = 1000$. However, modification of the sample size also increases ARL_δ. Furthermore, comparing the modified Methods 1 and 2, we see that both methods have similar results for $\rho \leq 0.7$ because the sample means are nearly independent. (With sample size 30, the correlation between adjacent sample means equals 0.009, 0.023, 0.051, and 0.208 for $\rho =$ 0.25, 0.5, 0.7, and 0.9, respectively.) When $\rho = 0.9$, the modified-Method-1 ARL_0 value is closer to L than the modified-Method-2 ARL_0 value for cases with the same m values. This is because the modified Method 2 considers the positive autocorrelations but the modified Method 1 does not, and hence, with the same sample size, the corresponding modified-Method-2 k value is smaller than that for the modified Method 1. However, the difference in the $\widehat{\mathrm{ARL}}_0$ values is negligible. Overall, even with the restriction $m \geq 30$, the modified Method 2 still performs better than the modified Method 1.

3 Numerical Comparisons with Previous Methods

In this section, we empirically compare the modified Method 2 to the R&W and DFTC charts. We chose the modified Method 2 instead of the modified Method 1 because it performs better and because it works well for nonnormal marginal distributions. Our numerical results show that the modified Method 2 performs better than the R&W and DFTC charts, especially when the correlation is high.

As in Section 2.4, we use AR(1) and ARTA(1) data with the t_{10} marginal distribution to compare the modified Method 2 to the R&W and DFTC charts. The values of the lag-1 autocorrelation ρ_1, shift δ, and specified value L are the same as in Tables 1 and 3.

Table 4 shows the AR(1) results for $\rho_1 = 0$, 0.25, and 0.5; Table 5, for $\rho_1 = 0.9$, 0.95, and 0.99. In both tables, columns 1 and 2 list the values of ρ_1 and δ; columns 3 and 4, the R&W and DFTC ARL_δ; and columns 5 to 7, the modified-Method-2 values of m, ARL_0, and ARL_δ. For each combination of ρ_1 and δ, the lowest and best ARL_δ is marked with a

Table 3: The modified-Method-1 and modified-Method-2 design outputs and the true optimal solutions for ARTA(1) processes with t_{10} marginal distribution, $\rho_1 = 0$, 0.25, 0.5, 0.7, 0.9, and $L = 1000$

		modified Method 1				modified Method 2				true optimal solution		
ρ_1	δ	m	k	$\widehat{\text{ARL}}_0$	$\widehat{\text{ARL}}_\delta$	m	k	$\widehat{\text{ARL}}_0$	$\widehat{\text{ARL}}_\delta$	m	k	$\widehat{\text{ARL}}_\delta$
0	0.25	66	1.838	995	115	same as for the modified Method 1				67	1.832	114
	0.5	30	2.170	982	42					28	2.203	42
	0.75	30	2.170	982	31					14	2.477	22
	1	30	2.170	982	30					10	2.614	14
	1.5	30	2.170	982	30					5	2.911	7
	2	30	2.170	982	30					4	3.022	5
	4	30	2.170	982	30					2	3.448	2
0.25	0.25	88	1.706	997	161	88	1.706	997	160	87	1.713	160
	0.5	36	2.097	982	61	37	2.086	987	60	36	2.103	61
	0.75	30	2.170	979	35	30	2.170	979	35	22	2.305	33
	1	30	2.170	979	31	30	2.170	979	31	13	2.520	21
	1.5	30	2.170	979	30	30	2.170	979	30	7	2.788	11
	2	30	2.170	979	30	30	2.170	979	30	5	2.948	7
	4	30	2.170	979	30	30	2.170	979	30	2	3.500	2
0.5	0.25	120	1.555	994	235	120	1.555	994	235	125	1.535	235
	0.5	53	1.935	991	93	54	1.927	992	92	53	1.938	93
	0.75	30	2.170	969	51	30	2.170	969	51	30	2.183	51
	1	30	2.170	969	35	30	2.170	969	35	21	2.334	32
	1.5	30	2.170	969	30	30	2.170	969	30	11	2.616	17
	2	30	2.170	969	30	30	2.170	969	30	7	2.830	10
	4	30	2.170	969	30	30	2.170	969	30	3	3.296	3
0.7	0.25	159	1.408	1000	338	159	1.408	1000	338	160	1.403	337
	0.5	77	1.768	995	143	77	1.768	995	143	80	1.751	143
	0.75	45	2.005	981	80	46	1.995	981	80	47	1.994	80
	1	30	2.170	952	51	31	2.157	948	51	29	2.205	52
	1.5	30	2.170	952	33	30	2.170	951	32	17	2.447	27
	2	30	2.170	952	30	30	2.170	951	30	10	2.696	16
	4	30	2.170	952	30	30	2.170	951	30	4	3.177	4
0.9	0.25	227	1.208	1002	582	233	1.192	995	582	243	1.167	584
	0.5	135	1.495	1006	299	142	1.468	994	298	142	1.469	299
	0.75	83	1.734	987	179	91	1.689	987	178	99	1.654	181
	1	52	1.943	966	120	62	1.864	969	118	63	1.870	120
	1.5	30	2.170	915	60	31	2.149	894	60	33	2.172	63
	2	30	2.170	915	38	30	2.162	892	38	20	2.404	37
	4	30	2.170	915	30	30	2.162	892	30	9	2.767	9

box. For the R&W and DFTC charts, their values of m and ARL_0 do not depend on δ and hence are listed at the top for each ρ_1 value. The DFTC results are as reported in Tables 3 and 4 of Kim et al. (2007), in which the DFTC parameter K is set to 0.1σ. The R&W sample sizes, calculated so that adjacent sample means having correlation near 0.1, are as reported in Table 3 of Runger and Willemain (1995). The ARL_0 and ARL_δ values for the R&W and modified-Method-2 charts are computed using the Markov-chain approach.

Together, Tables 4 and 5 demonstrate that the modified Method 2 often outperforms the other two methods. The difference is most noticeable in the following two cases: (i) ρ_1 and δ both small or moderate and (ii) ρ_1 large. When $\rho_1 \leq 0.5$ (Table 4), none of the three methods clearly dominates. The modified Method 2 works best for $\delta = 0.5$, 0.75, 1.0, and the combination $\rho_1 = 0.5$, $\delta = 1.5$. On the other hand, the DFTC chart, a tabular CuSum chart, is better at detecting small shifts in the process mean, consistently yielding smaller values of ARL_δ when $\delta = 0.25$. It also does better for certain parameter combinations, such as $\rho_1 = 0$, $\delta = 1.5$ and $\rho_1 = 0$, $\delta = 2$. On the other hand, the DFTC chart loses its advantage as δ increases. As δ approaches 4 and the modified-Method-2 sample size is forcibly increased to 30, the R&W chart works best. When $\rho_1 \geq 0.9$ (Table 5), the modified Method 2 performs better than the other two methods for almost all cases. Comparing the R&W and DFTC charts, the DFTC chart performs better than the R&W chart for small values of δ and worse for large values of δ. One disadvantage of the DFTC is that it yields ARL_0 values that are slightly higher than the specified value 10000.

We next consider nonnormal quality measurements. Table 6 compares the performance of four charts—the R&W, tuned R&W, DFTC, and modified Method 2—for ARTA(1) data with a t_{10} marginal distribution. We incorporate the tuned R&W chart to compensate for two shortcomings of the unmodified R&W chart in handling data of this type. First, as shown in Table 5, with a normally distributed quality characteristic, the sample size is often too large when ρ_1 and δ are both large. Second, with a nonnormally distributed quality characteristic, the sample size is often too small to yield approximately iid normal sample means, and as a result ARL_0 is far from the specified value. In the tuned R&W method, we adjust the sample sizes to bring the associated $\widehat{\mathrm{ARL}}_0$ values to within three standard errors of 1000.

Table 4: Comparisons of three charts: R&W, DFTC and modified Method 2 for AR(1) processes with $\rho_1 = 0$, 0.25, 0.5 and $L = 10000$ (The lowest ARL_δ is boxed.)

ρ_1	δ	R&W	DFTC	Modified Method 2: m	Modified Method 2: ARL_0	Modified Method 2: ARL_δ
		$\text{ARL}_0 = 10^4$ (m=1)	$\widehat{\text{ARL}}_0 = 9585$ (m=1)			
0	0.25	6522	$\boxed{178}$	133	10^4	202
	0.5	2822	72	45	10^4	$\boxed{65}$
	0.75	1184	45	30	10^4	$\boxed{34}$
	1	520	33	30	10^4	$\boxed{30}$
	1.5	119	$\boxed{21}$	30	10^4	30
	2	34	$\boxed{16}$	30	10^4	30
	2.5	$\boxed{12}$	13	30	10^4	30
	3	$\boxed{5.4}$	11	30	10^4	30
	4	$\boxed{1.8}$	8	30	10^4	30
		$\text{ARL}_0 = 10001$ (m=4)	$\widehat{\text{ARL}}_0 = 10846$ (m=1)			
0.25	0.25	4298	$\boxed{270}$	194	10^4	302
	0.5	1175	111	67	10^4	$\boxed{98}$
	0.75	367	69	35	10^4	$\boxed{50}$
	1	134	50	30	10^4	$\boxed{33}$
	1.5	$\boxed{28}$	32	30	10^4	30
	2	$\boxed{10}$	24	30	10^4	30
	2.5	$\boxed{5.6}$	19	30	10^4	30
	3	$\boxed{4.3}$	16	30	10^4	30
	4	$\boxed{4}$	12	30	10^4	30
		$\text{ARL}_0 = 10004$ (m=8)	$\widehat{\text{ARL}}_0 = 11356$ (m=1)			
0.5	0.25	4211	$\boxed{434}$	296	10^4	475
	0.5	1150	180	106	10^4	$\boxed{159}$
	0.75	367	112	56	10^4	$\boxed{81}$
	1	140	82	35	10^4	$\boxed{49}$
	1.5	34	53	30	10^4	$\boxed{31}$
	2	$\boxed{14}$	39	30	10^4	30
	2.5	$\boxed{9}$	31	30	10^4	30
	3	$\boxed{8}$	26	30	10^4	30
	4	$\boxed{8}$	19	30	10^4	30

Table 5: Comparisons of three charts: R&W, DFTC and modified Method 2 for AR(1) processes with $\rho_1 = 0.9, 0.95, 0.99$ and $L = 10000$ (The lowest ARL_δ is boxed.)

				Modified Method 2		
ρ_1	δ	R&W	DFTC	m	ARL_0	ARL_δ
		$\text{ARL}_0 = 10019$ (m=58)	$\widehat{\text{ARL}}_0 = 11668$ (m=7)			
0.9	0.25	4930	$\boxed{1728}$	945	10^4	1754
	0.5	1657	755	396	10^4	$\boxed{664}$
	0.75	650	481	222	10^4	$\boxed{355}$
	1	307	352	143	10^4	$\boxed{222}$
	1.5	112	227	72	9999	$\boxed{110}$
	2	69	167	40	9997	$\boxed{64}$
	2.5	59	133	30	9981	$\boxed{41}$
	3	58	111	30	9981	$\boxed{33}$
	4	58	83	30	9981	$\boxed{30}$
		$\text{ARL}_0 = 10029$ (m=118)	$\widehat{\text{ARL}}_0 = 12032$ (m=15)			
0.95	0.25	5445	2754	1362	10^4	$\boxed{2730}$
	0.5	2058	1250	625	10^4	$\boxed{1108}$
	0.75	893	792	364	10^4	$\boxed{610}$
	1	461	577	236	10^4	$\boxed{387}$
	1.5	194	377	120	9999	$\boxed{194}$
	2	132	278	66	9994	$\boxed{113}$
	2.5	120	223	30	9926	$\boxed{74}$
	3	118	185	30	9926	$\boxed{51}$
	4	118	139	30	9926	$\boxed{35}$
		$\text{ARL}_0 = 10058$ (m=596)	$\widehat{\text{ARL}}_0 = 12735$ (m=74)			
0.99	0.25	6890	6735	2357	10^4	$\boxed{5957}$
	0.5	3518	3383	1443	10^4	$\boxed{3081}$
	0.75	1949	2240	932	10^4	$\boxed{1856}$
	1	1240	1641	638	9999	$\boxed{1235}$
	1.5	728	1065	317	9992	$\boxed{648}$
	2	614	794	30	9011	$\boxed{381}$
	2.5	597	636	30	9011	$\boxed{234}$
	3	596	530	30	9011	$\boxed{162}$
	4	596	402	30	9011	$\boxed{98}$

Columns 1 and 2 of Table 6 list the values of ρ_1 and δ; columns 3 to 5, the ARL_δ estimates for the original R&W, tuned R&W, and DFTC charts; and columns 6 to 8, the modified-Method-2 values of m, $\widehat{\mathrm{ARL}}_0$, and $\widehat{\mathrm{ARL}}_\delta$, which are identical to those in Table 3. The standard-error estimates of $\widehat{\mathrm{ARL}}_0$ and $\widehat{\mathrm{ARL}}_\delta$ are listed in the supplement by Chen and Cheng (2008). The results for each ρ_1 value are displayed in a separate section of Table 6, with the R&W, tuned R&W, and DFTC values of m and $\widehat{\mathrm{ARL}}_0$ listed across the top. Each row in a section corresponds to a different ρ_1, $\delta > 0$ combination, with a box marking the lowest associated value of $\widehat{\mathrm{ARL}}_\delta$. For the t_{10} marginal distribution, we cannot apply the central limit theorem to the unmodified R&W chart because the sample sizes 1, 4, 8, and 17 are too small. As a result, the $\widehat{\mathrm{ARL}}_0$ values 234, 636, 759, and 872 stray far from the specified value $L = 1000$. For the five ρ_1 values $\rho_1 = 0$, 0.25, 0.5, 0.7, and 0.9, however, the tuned R&W sample sizes are 60, 60, 70, 80, and 140, bringing the associated $\widehat{\mathrm{ARL}}_0$ values 997, 994, 994, 997, and 995 closer to L.

Table 6 shows that when the lag-1 autocorrelation ρ_1 is high (e.g., 0.7 and 0.9), the modified Method 2 performs better than the other methods for most δ values. When ρ_1 and δ are small or moderate, the modified Method 2 performs best. When ρ_1 is small or moderate and δ is large, the modified-Method-2 sample sizes are increased to 30 (Column 6 of Table 6) to enable application of the central limit theorem. Therefore, the resulting ARL_δ increases. In this case, the DFTC chart is superior to the modified Method 2. Overall, the original and tuned R&W charts are not appealing since the original R&W chart may not meet the specified ARL_0 value and the tuned R&W chart has the highest ARL_δ for many combinations of ρ_1 and δ. We also empirically compared the modified Method 2 to the R&W and DFTC charts using ARTA(1) processes with the exponential marginal distribution, where the values of ρ_1 and δ are as in Table 6. (See the supplement by Chen and Cheng, 2008.) Our empirical results yielded similar conclusions as those from Table 6 and hence are not presented here.

Table 6: Comparisons of four charts: R&W, R&W with tuned sample sizes, DFTC, and modified Method 2 for ARTA(1) processes with t_{10} marginal distribution, $\rho_1 = 0$, 0.25, 0.5, 0.7, 0.9, and $L = 1000$ (The lowest $\widehat{\text{ARL}}_\delta$ is boxed.)

					Modified Method 2		
ρ_1	δ	R&W	Tuned R&W	DFTC	m	$\widehat{\text{ARL}}_0$	$\widehat{\text{ARL}}_\delta$
		$\widehat{\text{ARL}}_0 = 234$ (m=1)	$\widehat{\text{ARL}}_0 = 997$ (m=60)	$\widehat{\text{ARL}}_0 = 999$ (m=1)			
	0.25	—	115	$\boxed{105}$	66	995	115
	0.5	—	61	44	30	982	$\boxed{42}$
0	0.75	—	60	$\boxed{28}$	30	982	31
	1	—	60	$\boxed{21}$	30	982	30
	1.5	—	60	$\boxed{13}$	30	982	30
	2	—	60	$\boxed{10}$	30	982	30
	4	—	60	$\boxed{5}$	30	982	30
		$\widehat{\text{ARL}}_0 = 636$ (m=4)	$\widehat{\text{ARL}}_0 = 994$ (m=60)	$\widehat{\text{ARL}}_0 = 1100$ (m=1)			
	0.25	—	170	$\boxed{154}$	88	997	160
	0.5	—	69	66	37	987	$\boxed{60}$
0.25	0.75	—	60	42	30	979	$\boxed{35}$
	1	—	60	$\boxed{30}$	30	979	31
	1.5	—	60	$\boxed{20}$	30	979	30
	2	—	60	$\boxed{15}$	30	979	30
	4	—	60	$\boxed{7}$	30	979	30
		$\widehat{\text{ARL}}_0 = 759$ (m=8)	$\widehat{\text{ARL}}_0 = 994$ (m=70)	$\widehat{\text{ARL}}_0 = 1196$ (m=1)			
	0.25	—	254	$\boxed{234}$	120	994	235
	0.5	—	96	102	54	992	$\boxed{92}$
0.5	0.75	—	72	64	30	969	$\boxed{51}$
	1	—	70	47	30	969	$\boxed{35}$
	1.5	—	70	$\boxed{30}$	30	969	$\boxed{30}$
	2	—	70	$\boxed{22}$	30	969	30
	4	—	70	$\boxed{11}$	30	969	30
		$\widehat{\text{ARL}}_0 = 872$ (m=17)	$\widehat{\text{ARL}}_0 = 997$ (m=80)	$\widehat{\text{ARL}}_0 = 1184$ (m=3)			
	0.25	—	374	347	159	1000	$\boxed{338}$
	0.5	—	144	160	77	995	$\boxed{143}$
0.7	0.75	—	92	103	46	981	$\boxed{80}$
	1	—	82	75	31	948	$\boxed{51}$
	1.5	—	80	49	30	951	$\boxed{32}$
	2	—	80	36	30	951	$\boxed{30}$
	4	—	80	$\boxed{18}$	30	951	30
		$\widehat{\text{ARL}}_0 = 970$ (m=58)	$\widehat{\text{ARL}}_0 = 995$ (m=140)	$\widehat{\text{ARL}}_0 = 1292$ (m=9)			
	0.25	679	604	696	233	995	$\boxed{582}$
	0.5	348	$\boxed{298}$	361	142	994	$\boxed{298}$
0.9	0.75	190	190	236	91	987	$\boxed{178}$
	1	119	154	173	62	969	$\boxed{118}$
	1.5	69	141	113	31	894	$\boxed{60}$
	2	59	140	84	30	892	$\boxed{38}$
	4	58	140	43	30	892	$\boxed{30}$

4 Summary, Conclusions, and Future Research

This paper presents two methods for designing $\overline{X}$ charts under the assumption that the quality characteristic follows an autocorrelated process with an unknown marginal distribution shape and a known covariance structure. The design task is to determine the sample size m and the number k of standard deviations away from the center line such that the out-of-control ARL is minimized while the in-control ARL is maintained at a specified value. Method 1 models the sample means as iid normal random variables and Method 2 models them as an AR(1) process. To increase the applicability of the central limit theorem to nonnormal marginal distributions, we have modified Methods 1 and 2 so that the sample size is at least 30.

In our numerical results, we first determined that the modified Method 2 performs better than the modified Method 1. We then compared the modified Method 2 to R&W and DFTC charts. The R&W chart performed best under three conditions: the marginal distribution was close to normal, the autocorrelation was small to moderate, and the shift was large. The DFTC chart performed best under two conditions: the correlation was small to moderate and the shift was small or large. In all other cases, the modified Method 2 performed best.

One direction of future research is to design a procedure for the $\overline{X}$ chart when the data properties are unknown but can be estimated from a set of Phase-I in-control data. In our work, we have considered data with an unknown marginal distribution but a known covariance structure. However, in some cases both the marginal distribution and covariance structure are unknown. In this case, the procedure would use the Phase-I data to estimate the data properties, e.g., $\sigma_{\overline{X}}$, or the entire marginal distribution. The issue is to determine the amount of performance degradation (compared to the current work) as a function of the amount of Phase-I data.

Acknowledgements

This research was conducted with the support of the National Science Council in Taiwan under grant NSC 93-2213-E-033-023. We thank Bruce Schmeiser and David Goldsman for their helpful discussions. We also thank two anonymous referees for their insightful suggestions.

References

Boyles, R.A. (2000). Phase I analysis for autocorrelated processes. *Journal of Quality Technology* **32**, 395–409.

Cario, M.C., Nelson, B.L. (1996). Autoregressive to anything: Time-series input processes for simulation. *Operations Research Letters* **19**, 51–58.

Cario, M.C., Nelson, B.L. (1998). Numerical methods for fitting and simulating autoregressive-to-anything processes. *INFORMS Journal on Computing* **10**, 72–81.

Chen, H. (2001). Initialization for NORTA: Generation of random vectors with specified marginals and correlations. *INFORMS Journal on Computing* **13**, 312–331.

Chen, H., Cheng, Y. (2008). Supplement to "Designing $\overline{X}$ Charts for Known Autocorrelations and Unknown Marginal Distribution". *Technical Report*, Department of Industrial and Systems Engineering, Chung-Yuan University, Taiwan. (Available online via http://140.135.139.141/huifen/index1.htm .)

English, J.R., Case, K.E. (1994). Statistical process control in continuous flow processes—ramp disturbance. *IIE Transactions* **26**, p. 22–28.

Faltin, F.W., Mastrangelo, C.M., Runger, G.C., Ryan, T.P. (1997). Considerations in the monitoring of autocorrelated and independent data. *Journal of Quality Technology* **29**, 131–133.

Ghosh, S., Henderson, S.G. (2003). Behavior of NORTA method for correlated random vector generation as the dimension increases. *ACM TOMACS*, **13**, 276–294.

Hahn, G.J. (1989). Statistics-aided manufacturing: A look into the future. *American Statistician* **43**, 74–79.

Jiang, W., Tsui, K.-L., Woodall, W.H. (2000). A new SPC monitoring method: The ARMA chart. *Technometrics* **42**, 399–410.

Kang, K., Schmeiser, B.W. (1987). Properties of batch means from stationary ARMA time series. *Operations Research Letters* **6**, 19–24.

Kim, S.-H., Alexopoulos, C., Goldsman, D., Tsui, K.-L. (2006). A new model-free CuSum procedure for autocorrelated processes. *Technical Report*, School of Industrial & Systems

Engineering, Georgia Institute of Technology, GA, USA.

Kim, S.-H., Alexopoulos, C., Tsui, K.-L., Wilson, J.R. (2007). A distribution-free tabular CuSum chart for autocorrelated data. *IIE Transactions* **39**, 317–330.

Koo, T.-Y., Case, K.E. (1990). Economic design of X-bar control charts for use in monitoring continuous flow processes. *International Journal of Production Research* **28**, 2001–2011.

Lucas, J.M., Saccucci, M.S. (1990). Exponentially weighted moving average control schemes: Properties and enhancements. *Technometrics* **32**, 1–12.

Montgomery, D.C. (2005). *Introduction to Statistical Quality Control*, 5th edn., New York: Wiley.

Pandit, S.M., Wu, S.M. (1983). *Time Series and System Analysis with Applications*, New York: Wiley.

Ross, S.M. (2004). *Introduction to Probability and Statistics for Engineers and Scientists*, 3rd edn., Elsevier Academic Press.

Runger, G.C., Willemain, T.R. (1995). Model-based and model-free control of autocorrelated processes. *Journal of Quality Technology* **27**, 283–292.

Tucker, W.T., Faltin, F.W., Vander Wiel, S.A. (1993). Algorithmic statistical process control: An elaboration. *Technometrics* **35**, 363–375.

Wardell, D.G., Moskowitz, H., Plante, R.D. (1994). Run-length distributions of special-cause control charts for correlated processes. *Technometrics* **36**, 3–17.