A New SPC Monitoring Method: The ARMA Chart

Wei JIANG

Department of Industrial Engineering
and Engineering Management
The Hong Kong University of Science and Technology
Clear Water Bay
Hong Kong
(*jiangwei@uxmail.ust.hk*)

Kwok-Leung TSUI

School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332
(*ktsui@isye.gatech.edu*)

William H. WOODALL

Department of Statistics
Virginia Institute of Technology and State University
Blacksburg, VA 24061
(*bwoodall@vt.edu*)

We propose a new control chart, the autoregressive moving average (ARMA) chart, based on monitoring an ARMA statistic of the original observations. It is shown that the special cause chart (SCC) of Alwan and Roberts and the EWMAST chart of Zhang are special cases of the ARMA chart. Simulation studies show that the ARMA chart is competitive to the optimal exponentially weighted moving average chart for iid observations and better than the SCC and EWMAST charts for autocorrelated observations. We develop an informal procedure to determine the appropriate parameter values of the proposed chart based on two signal-to-noise ratios. Two real examples are discussed to demonstrate the advantages of the new chart.

KEY WORDS: Average run length; Control chart; Quality control.

The statistical control chart is an effective tool for achieving process stability. For monitoring autocorrelated observations, various control charts have been developed to detect shifts in the mean of the process. Among those that have been widely discussed are the special cause chart (SCC) (Alwan and Roberts 1988; Wardell, Moskowitz, and Plante 1994) and the EWMAST chart (Schmid 1997; Zhang 1998; Adams and Tseng 1998; Lu and Reynolds 1999a,b). The basic idea of the SCC chart involves filtering techniques to whiten an autocorrelated process (if the process parameters can be accurately estimated) and then monitoring the residuals by traditional control charts. This chart is shown to be effective when detecting large shifts. On the other hand, the EWMAST chart applies the exponentially weighted moving average (EWMA) statistic directly to the autocorrelated process without identifying the process parameters and is shown to be efficient in some parameter regions.

Taking the autocorrelation structure into account, this article proposes a new charting technique based on an autoregressive moving average (ARMA) statistic, the ARMA chart. This new chart provides a more flexible choice of parameters to relate the autocorrelation structure of the statistic to the chart performance and includes the SCC chart and the EWMAST chart as special cases. It is shown that an ARMA chart with appropriate parameter values will outperform both the SCC and EWMAST charts for autocorrelated processes. In Section 1, two real examples are presented to illustrate the problem under consideration. In Section 2, we explain the model of the proposed ARMA chart and illustrate it with an iid example. In Section 3, we discuss the application of the proposed chart to autocorrelated processes. In Section 4, we present the results of simulation studies for both iid and autocorrelated observations. In Section 5, we revisit the previous examples by the proposed ARMA chart and demonstrate the design procedure of the new chart for a given process. Finally, new research directions are discussed and concluding remarks are given in Section 6. The appendix contains a Markov chain approach to evaluate the run-length distribution of the ARMA chart for iid observations.

1. MOTIVATING EXAMPLES

As motivation for this work, the following two real examples demonstrate circumstances under which the proposed method is applicable. The first example is taken from Pandit and Wu's (1983, pp. 491–492) book, which presents the micrometer readings in the quality-control inspection of two diameters on the same machined part manufactured on a single-spindle automatic screw machine. Figure 1 shows the last 360 observations (deviations from target) of Diameter 2. Based on the process knowledge and the data, the system is considered to be in a stable state in the first period of 120 observations in the dataset, and the objective of SPC tools

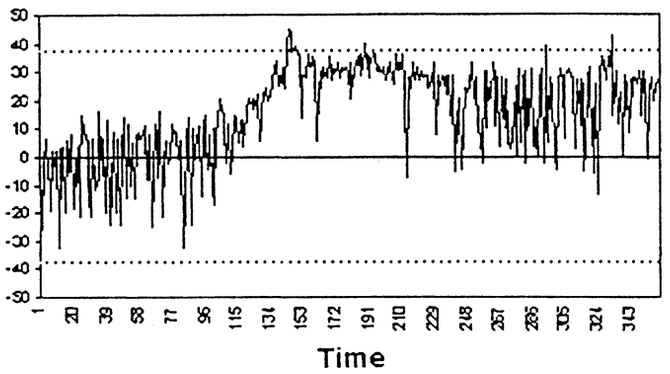

Figure 1. Observation Deviations: The First Example.

is to detect any departure from this stable state.

Figure 1 also presents the 3σ control limits applied to the dataset (traditional Shewhart X chart). The process standard deviation is estimated as $\hat{\sigma}_X = 12.43$ using the average moving range (the moment estimate is 12.49). It is found that the process mean may have increased in the second period of 120 observations since the X chart signals at observations 144, 145, 146, and 148. To investigate this possibility, the first 120 observations were used to fit an ARMA process. Box–Pierce's χ^2 test (Box, Jenkins, and Reinsel 1994) shows that the following ARMA(3, 2) process is appropriate: $X_t + .9202X_{t-1} + .8851X_{t-2} - .0976X_{t-3} = a_t + .9677a_{t-1} + .9133a_{t-2}$, where a_t is the white noise with $\hat{\sigma}_a = 11.95$. Figure 2(a) shows the SCC chart with 3σ limits applied on the residuals of the fitted ARMA(3, 2) model. Like the X chart, the SCC chart signals at observations 144, 145, and 146, indicating a possible increase of the process mean.

For comparison, an EWMAST chart is also applied to the original data. Figure 2(b) shows the EWMAST chart with $\lambda = .2$ recommended by Zhang (1998) [the control limits are adjusted to $\pm 2.572\sigma$ so that the in-control average run length (ARL) is maintained at 370]. The EWMAST chart signals a mean shift at observation 123, which indicates that the EWMAST chart may be more sensitive than the SCC chart in detecting the shift in this example.

A different example, however, also taken from Pandit and Wu's book, shows that the EWMAST chart may not always detect shifts earlier than the SCC chart. Figure 3 (for illustration, only the last 60 observations are displayed) shows the observations (deviation from target) from a mechanical dynamic system consisting of a mass, a dashpot, and a spring (Pandit and Wu 1983, p. 291). The model estimated using the first 100 observations is an ARMA(2, 1) process, $X_t - 1.4385X_{t-1} + .6000X_{t-2} = a_t + .5193a_{t-1}$, where $\hat{\sigma}_X = 9.130$ and $\hat{\sigma}_a = 2.212$. A mean shift of $1\sigma_X$ is manually added to the last 30 observations; neither the process with nor the process without the shift exhibits out-of-control signals when the X chart is applied (Fig. 3). The SCC chart with 3σ control limits and the EWMAST chart with $\lambda = .2$ and adjusted control limits $\pm 2.556\sigma$ for an in-control ARL of 370 are applied to the data and shown in Figure 4. In this case, we find that the EWMAST chart cannot detect the shift, but the SCC chart signals at the first observation after the shift.

Although the average performance of a control chart cannot be compared based on a single set of data, the preceding examples illustrate that the chart performance may differ for different underlying processes. As explained in subsequent sections, relative performance of the EWMAST chart and the SCC chart critically depends on the parameters of the underlying process. In this article, we propose a new control chart, the ARMA chart, which includes the EWMAST and SCC charts as special cases and allows us to optimize the chart performance based on the process parameters.

2. AUTOREGRESSIVE MOVING AVERAGE CHART

We now introduce the new ARMA chart for monitoring the mean of a stochastic process. Independent, identically distributed (iid) processes are considered in this section, and autocorrelated processes will be discussed in the next section.

Suppose that we are monitoring an iid process, $a_1, a_2, \ldots$, with normality, an in-control mean of 0, and variance σ_a^2. We wish to detect shifts in the mean of the process. The successive values to be plotted on an ARMA chart are defined to be the result of a generalized first-order autoregressive moving average [ARMA(1, 1)] process applied to the iid process; that is,

$$\begin{aligned} Z_t &= \theta_0 a_t - \theta a_{t-1} + \phi Z_{t-1} \\ &= \theta_0(a_t - \beta a_{t-1}) + \phi Z_{t-1}, \end{aligned} \tag{1}$$

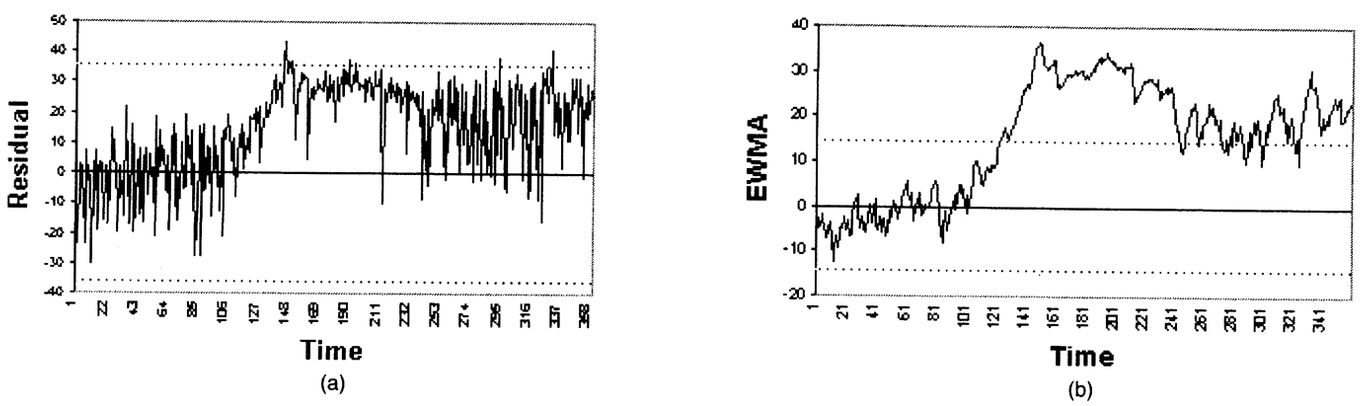

Figure 2. Residual Chart Versus EWMAST Chart (the first example): (a) SCC Chart; (b) EWMAST Chart.

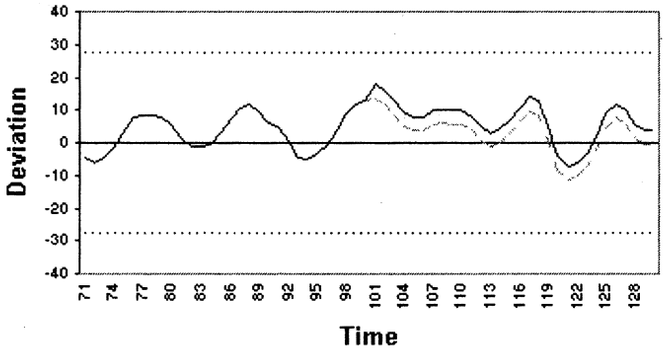

Figure 3. Observation Deviations: The Second Example: ———, Without Shift; ——— , With Shift.

where $\beta = \theta/\theta_0$ and θ_0 is chosen so that the sum of the coefficients is unity when Z_t is expressed in terms of a_t's; that is,

$$\left.\frac{\theta_0 - \theta B}{1 - \phi B}\right|_{B=1} = \frac{\theta_0 - \theta}{1 - \phi} = 1,$$

where B is the backshift operator, $Ba_t = a_{t-1}$. Thus, $\theta_0 = 1 + \theta - \phi$. To guarantee that the monitoring process is reversible and stationary, we have the constraints that $|\beta| < 1$ and $|\phi| < 1$. The ARMA chart signals when $|Z_t| > L\sigma_Z$. Note that the ARMA chart reduces to the EWMA chart when $\theta = 0$ with $\phi = 1 - \lambda$. Thus, the ARMA chart can be considered as an extension of the EWMA chart.

By denoting $W_t = Z_t/\theta_0$, which is a standard ARMA(1, 1) process, it is easy to show that the steady-state variance and the first-lag autocovariance of the monitoring statistic are

$$\sigma_Z^2 = \left[\frac{2(\theta - \phi)(1 + \theta)}{1 + \phi} + 1\right]\sigma_a^2, \tag{2}$$

and

$$\gamma(1) = \frac{(1 + \theta)(1 - \phi)(\phi - \theta)}{1 + \phi}\sigma_a^2. \tag{3}$$

Furthermore, it can be shown that, for a given ϕ, the variance of the monitoring process is minimized at $\phi - \theta = 1 + \theta$; that is, $\theta = -(1 - \phi)/2$. It is interesting to note that this is the condition for a positive first-lag autocovariance $\gamma(1)$. The autocorrelations of the monitoring process are $\rho_1 = \phi - \theta\theta_0/\sigma_Z^2$ and $\rho_k = \phi^{k-1}\rho_1 (k > 1)$.

Suppose that there is a step shift μ in the mean of the original process a_t occurring at t_0. Then the corresponding mean shift on the monitoring process becomes

$$\mu_{t_0} = \theta_0\mu, \quad \mu_{t_0+k} = (\theta_0 - \theta)\mu + \phi\mu_{t_0+k-1}, \quad k \geq 1. \tag{4}$$

The simulated example from Lucas and Saccucci (1990) illustrates the construction of the ARMA chart. Table 1 shows the iid observations, the EWMA chart statistics, and the corresponding ARMA chart statistics. The target value and the standard deviation are chosen to be 0 and 1, respectively. The process is in control for the first 10 observations, and the mean shift occurs at the 11th observation (see Fig. 5).

The parameters of the EWMA chart are chosen to be optimal in detecting one standard deviation of mean shift with an in-control ARL of 500. According to Lucas and Saccucci (1990), the parameters are $\lambda = .15$ and $L = 2.913$, giving actual control limits of $\pm.829$.

For the ARMA chart, the parameters are chosen as $\phi = .85$, which corresponds to $\lambda = .15$ of the EWMA chart, and $\theta = -.03$ (the choice of parameters will be discussed in Sec. 4). The control limits $\pm.725$ are chosen so that the in-control ARL remains at 500. It follows that $\theta_0 = 1 - .03 - .85 = .12$ in this example. Therefore, the ARMA statistic can be calculated as $Z_t = .85Z_{t-1} + .12a_t + .03a_{t-1}$. For example (see Table 1), the ARMA statistic of the 7th run is computed as $Z_7 = .85Z_6 + .12a_7 + .03a_6 = .85 \times (-.311) + .12 \times 1.5 + .03 \times (-1.2) = -.120$.

To see how the ARMA and EWMA charts react to mean shifts, we know a shift of one standard deviation occurs at the 11th run. As shown in Table 1 and Figure 6, the ARMA chart and the EWMA chart both signal at the 16th run.

To illustrate the difference between the two charts, we consider a smaller mean shift of .75 standard deviations. As shown in Table 1 and Figure 6(a) and (b), the EWMA chart signals the shift at the 18th observation, but the ARMA chart detects the shift at the 17th. Although in this example the ARMA chart is more sensitive in detecting a small mean shift, this does not imply that the average performance of the ARMA chart is better. In general, for iid processes, although the ARMA chart can sometimes be more efficient than the EWMA chart, the ARL performance of the two charts is comparable in most cases. For autocorrelated pro-

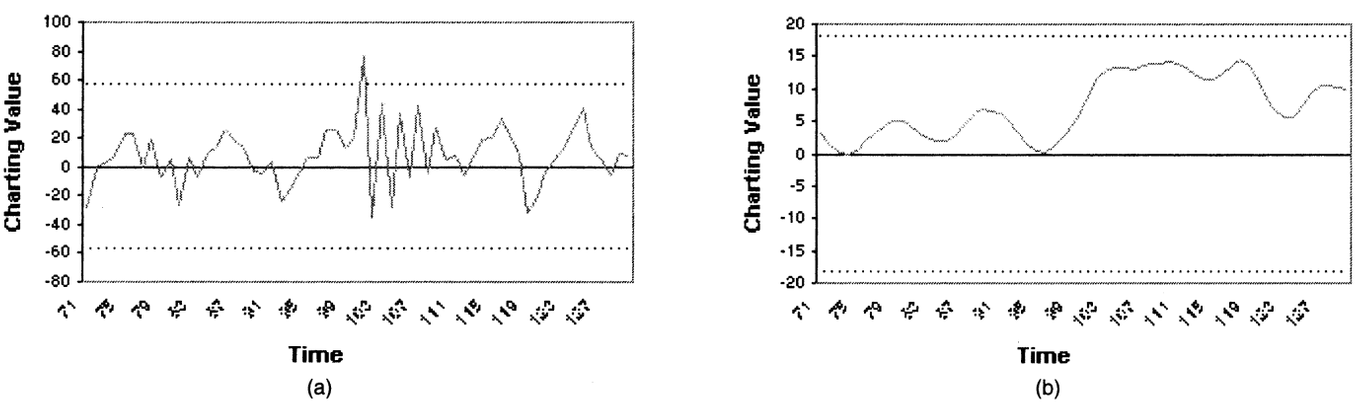

Figure 4. Residual Chart Versus EWMAST Chart (the second example): (a) SCC Chart; (b) EWMAST Chart.

Table 1. The Comparison of EWMA and ARMA Charts

	Obs.	EWMA	ARMA	Obs.	EWMA	ARMA
Obs. no.		$1\sigma_a$ shift			$.75\sigma_a$ shift	
1	1.0	.15	.12	1.0	.15	.12
2	−.5	.053	.072	−.5	.053	.072
3	0.0	.045	.046	0.0	.045	.046
4	−.8	−.082	−.057	−.8	−.082	−.057
5	−.8	−.190	−.168	−.8	−.190	−.168
6	−1.2	−.341	−.311	−1.2	−.341	−.311
7	1.5	−.065	−.120	1.5	−.065	−.120
8	−.6	−.145	−.129	−.6	−.145	−.129
9	1.0	.026	−.008	1.0	.026	−.008
10	−.9	−.113	−.085	−.9	−.113	−.085
11*	1.2	.084	.045	.95	.047	.015
12	.5	.147	.134	.25	.077	.071
13	2.6	.515	.441	2.35	.418	.350
14	.7	.543	.537	.45	.423	.422
15	1.1	.626	.609	.85	.487	.474
16	2.0	.832	.791	1.75	.676	.639
17	1.4	.917	.900	1.15	.747	.733
18	1.9	1.065	1.035	1.65	.883	.856
19	.8	1.025	1.033	.55	.833	.843

* Mean shift starts.

cesses, however, the ARMA chart outperforms the EWMA chart in many situations. Next we discuss how the ARMA chart can be applied to autocorrelated processes.

3. ARMA CHART FOR AUTOCORRELATED PROCESSES

3.1 Variance and Covariance Structure

In this section, we investigate the performance of the ARMA chart applied to a known process model. Now consider the application of the ARMA chart (with parameters ϕ and θ) to a stationary process. [Following the same notation of the EWMAST chart proposed by Zhang (1998), we denote this control chart as the ARMAST chart.] Assume that the underlying process X_t is characterized by the autocorrelation structure $\rho(\tau)$ with $\rho(\tau) = \gamma(\tau)/\gamma(0)$ and $\gamma(\tau) = \text{cov}[X_t, X_{t+\tau}]$. It follows that the ARMA statistic Z_t can be represented by

$$Z_t = \theta_0 X_t + \alpha \sum_{k=1}^{t-1} \phi^{k-1} X_{t-k},$$

where $\alpha = \phi\theta_0 - \theta$.

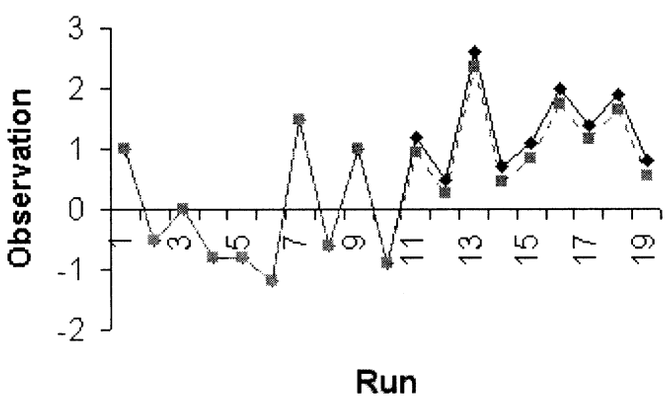

Figure 5. IID Example: —♦—, 1 Sigma Shift; - - ■ - -, .75 Sigma Shift.

Similar to the derivation of Zhang (1998), the covariance of Z_t can be obtained as

$$\begin{aligned}
&\text{cov}[Z_t, Z_{t+\tau}] \\
&= \text{cov}\left[\theta_0 X_t + \alpha \sum_{k=1}^{t-1} \phi^{k-1} X_{t-k}, \right. \\
&\qquad \left. + \theta_0 X_{t+\tau} + \alpha \sum_{k=1}^{t+\tau-1} \phi^{k-1} X_{t+\tau-k}\right] \\
&= \theta_0^2 \gamma(\tau) + \theta_0 \alpha \\
&\quad \times \left[\sum_{k=1}^{t-1} \phi^{k-1}\gamma(\tau+k) + \sum_{k=1}^{t+\tau-1} \phi^{k-1}\gamma(\tau-k)\right] \\
&\quad + \alpha^2 \sum_{i=1}^{t-1} \sum_{j=1}^{t+\tau-1} \phi^{i+j-2}\gamma(\tau-j+i) \\
&= \sigma_X^2 \left\{\theta_0^2\rho(\tau) + \theta_0\alpha\left[\sum_{k=1}^{t-1}\phi^{k-1}\rho(\tau+k) \right.\right. \\
&\qquad \left. + \sum_{k=1}^{t+\tau-1}\phi^{k-1}\rho(\tau-k)\right] \\
&\qquad \left. + \alpha^2 \sum_{i=1}^{t-1}\sum_{j=1}^{t+\tau-1}\phi^{i+j-2}\rho(\tau-j+i)\right\}.
\end{aligned}$$

When $\tau = 0$, this reduces to the variance of the monitoring process Z_t,

$$\sigma_Z^2 = \left\{\theta_0^2 + 2\theta_0\alpha\sum_{k=1}^{t-1}\phi^{k-1}\rho(k) + \alpha^2\sum_{i=1}^{t-1}\sum_{j=1}^{t-1}\phi^{i+j-2}\rho(j-i)\right\}\sigma_X^2.$$

Similarly, the steady-state variance is

$$\begin{aligned}
\sigma_Z^2 &= \left\{\theta_0^2 + 2\theta_0\alpha\sum_{k=1}^{\infty}\phi^{k-1}\rho(k) \right. \\
&\qquad \left. + \alpha^2\sum_{i=1}^{\infty}\sum_{j=1}^{\infty}\phi^{i+j-2}\rho(j-i)\right\}\sigma_X^2 \\
&= \left\{\theta_0^2 + 2\theta_0\alpha\sum_{k=1}^{\infty}\phi^{k-1}\rho(k) \right. \\
&\qquad \left. + \frac{\alpha^2}{1-\phi^2}\rho(0) + 2\alpha^2\phi^{-2}\sum_{k=1}^{\infty}\phi^k\rho(k)\sum_{l=1}^{\infty}\phi^{2l}\right\}\sigma_X^2 \\
&= \left\{\theta_0^2 + \frac{\alpha^2}{1-\phi^2} + 2\left(\theta_0\alpha + \frac{\phi\alpha^2}{1-\phi^2}\right)\sum_{k=1}^{\infty}\phi^{k-1}\rho(k)\right\}\sigma_X^2,
\end{aligned}$$

where $\sum_{k=1}^{\infty}\phi^k\rho(k)$ converges because $|\rho(k)| < 1$ $(k > 0)$. Note that the preceding steady-state variance reduces to Equation (2) when $X_t = a_t$ and $\rho(k) = 0$.

When the original process is an ARMA(1, 1) process with parameters u and v, the steady-state variance can be sim-

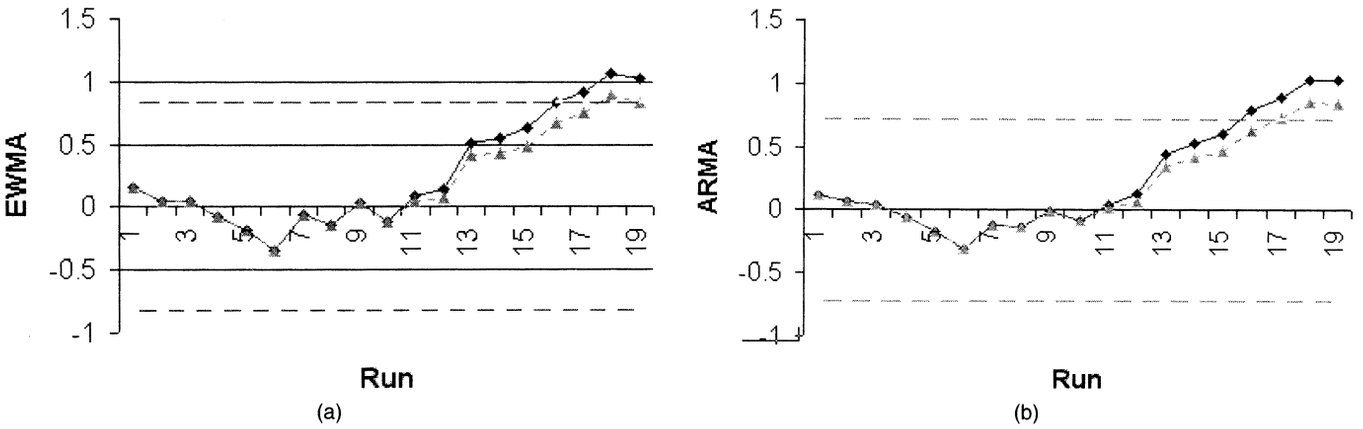

Figure 6. EWMA Chart Versus ARMA Chart: (a) EWMA Chart; (b) ARMA Chart.

plified in terms of the parameters of the process. Suppose that the underlying process is $X_t - uX_{t-1} = a_t - va_{t-1}$ with $|u| < 1$ and $|v| < 1$. It follows that the variance and first-lag correlation coefficient are

$$\sigma_X^2 = \frac{1 - 2uv + v^2}{1 - u^2}\sigma_a^2$$

and

$$\rho(1) = u - v\sigma_a^2/\sigma_X^2,$$

respectively, and $\rho(k) = u\rho(k-1)$ $(k \geq 2)$. By substituting these terms into the steady-state variance, we have

$$\sigma_Z^2 = \left\{\theta_0^2 + \frac{\alpha^2}{1-\phi^2} + 2\left(\theta_0\alpha + \frac{\phi\alpha^2}{1-\phi^2}\right)\frac{\rho(1)}{1-\phi u}\right\}\sigma_X^2.$$

In addition, it can be shown that the application of the ARMA chart to an ARMA(1, 1) process will result in a generalized ARMA(2, 2) process; that is,

$$\begin{aligned} Z_t &= \frac{\theta_0 - \theta B}{1 - \phi B}X_t \\ &= \frac{\theta_0 - \theta B}{1 - \phi B}\,\frac{1 - vB}{1 - uB}a_t \\ &= (\phi + u)Z_{t-1} - \phi u Z_{t-2} + \theta_0 a_t \\ &\quad - (\theta + \theta_0 v)a_{t-1} + \theta v a_{t-2}. \end{aligned} \tag{5}$$

Whenever a step shift μ occurs at $t = t_0$ in the underlying process X_t, the shift pattern of the monitoring process Z_t after the occurrence follows from Equation (4). Solving this difference equation, the shift pattern following t_0 is

$$\mu_{t_0+k} = [1 + (\theta - \phi)\phi^k]\mu, \qquad k \geq 0. \tag{6}$$

It is interesting to note that this shift pattern depends only on the shift size and the charting parameters ϕ and θ.

In general, because the autocorrelation structure of the ARMA chart on an ARMA(1, 1) process depends on the parameters of the charting process (ϕ and θ) as well as those of the original process (u and v), the performance of the ARMA chart depends on all four parameters. It is therefore hard to characterize the performance of the ARMA chart. The performance of a special class of ARMA charts has actually been studied extensively in the literature, however. As shown in Equation (5), if the parameters of the ARMA chart are chosen as $\phi = v$ and $\beta = \theta/\theta_0 = u$, then the monitoring process reduces to $Z_t = \theta_0 a_t$. In addition, the mean shift pattern of the monitoring process is given in Equation (6) with $\phi = v$ and $\theta/\theta_0 = u$. By comparing the monitoring process and the mean shift pattern of the class of ARMA chart with those of the SCC chart of Alwan and Roberts (1988) and Wardell et al. (1994), they are essentially the same except for a scaling constant θ_0. Therefore, the performance of the ARMA chart with $\phi = v$ and $\theta/\theta_0 = u$ is identical to the performance of the SCC chart applied to an ARMA(1, 1) process with parameters u and v.

As pointed out in Section 1, another special class of ARMA charts is the EWMA chart ($\phi = 1 - \lambda$ and $\theta = 0$). The performance of the EWMA chart applied to an ARMA(1, 1) process was first studied by Wardell et al. (1992, 1994) and further investigated and denoted by the EWMAST chart of Zhang (1998). Schmid (1997) and VanBrackle and Reynolds (1997) also investigated the application of the EWMA chart to autocorrelated processes. Wardell et al. (1992) showed that the performance of the SCC chart is better than the EWMAST chart (which was denoted by EWMA in their article) in some cases, but worse in others. Relative performance critically depends on the parameters of the original process (u and v), and neither chart is uniformly better than the other. Because both the SCC and EWMAST charts are special cases of the ARMA chart, it is possible to derive an ARMA chart with appropriate parameter values (ϕ and θ) that outperforms both the SCC and EWMA charts.

Theoretically, the parameters of the ARMA chart (ϕ and θ) can be chosen to optimize the ARL performance of the monitoring process. It is difficult to derive the optimal parameters analytically, however. We propose a heuristic strategy to choose appropriate ARMA chart parameters in practice.

3.2 Choosing ARMA Chart Parameters

As mentioned previously, the shift pattern on the charting process Z_t after the time of the shift t_0 follows Equation (6);

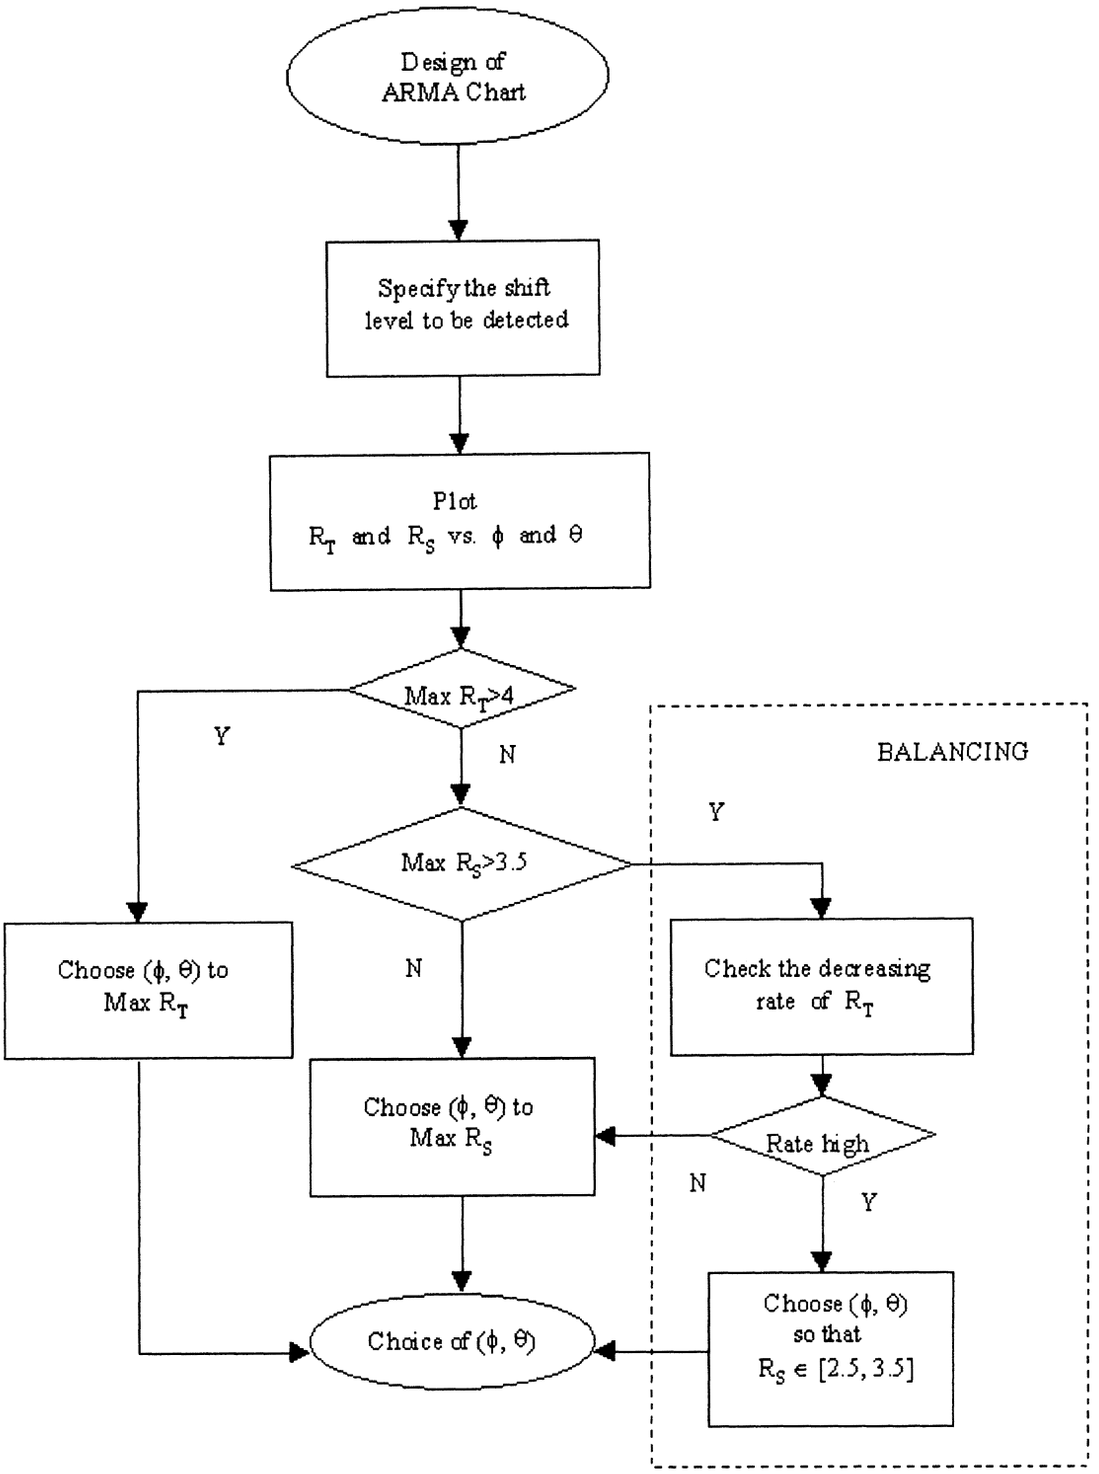

Figure 7. Parameter Design of ARMA Charts.

that is, the transient shift is $\mu_T = \theta_0\mu$ and the steady-state shift is $\mu_S = \mu$. We thus define the transient and steady-state signal-to-noise ratios as

$$R_T = \mu_T/\sigma_Z \tag{7}$$

and

$$R_S = \mu_S/\sigma_Z, \tag{8}$$

respectively. When the shift magnitude is measured in terms of the standard deviation of X_t, these two ratios become

$$R_T = \mu_T \Big/ \sqrt{\theta_0^2 + \frac{\alpha^2}{1-\phi^2} + 2\left(\theta_0\alpha + \frac{\phi\alpha^2}{1-\phi^2}\right)\sum_{k=1}^{\infty}\phi^{k-1}\rho(k)} \tag{9}$$

and

$$R_S = \mu_S \Big/ \sqrt{\theta_0^2 + \frac{\alpha^2}{1-\phi^2} + 2\left(\theta_0\alpha + \frac{\phi\alpha^2}{1-\phi^2}\right)\sum_{k=1}^{\infty}\phi^{k-1}\rho(k)}. \tag{10}$$

When the underlying process is ARMA(1, 1), the equations reduce to

$$R_T = \mu_T \Big/ \sqrt{\theta_0^2 + \frac{\alpha^2}{1-\phi^2} + 2\left(\theta_0\alpha + \frac{\phi\alpha^2}{1-\phi^2}\right)\frac{\rho(1)}{1-\phi u}}$$

and

$$R_S = \mu_S \Big/ \sqrt{\theta_0^2 + \frac{\alpha^2}{1-\phi^2} + 2\left(\theta_0\alpha + \frac{\phi\alpha^2}{1-\phi^2}\right)\frac{\rho(1)}{1-\phi u}}.$$

It is believed that these two ratios are critical to choosing appropriate parameters of an ARMA chart for autocorrelated processes. The transient ratio measures the capability of the chart to detect the shift in the first few runs and is more appropriate for detecting large shifts. If the chart fails to signal the shift in the early runs, the steady-state ratio becomes important for detecting the shift efficiently in later runs.

From our experience, if the transient ratio can be tuned to a high enough value (say 4 to 5) by choosing appropriate ARMA chart parameters, the corresponding chart will be able to detect the shift quickly. On the other hand, if this ratio is smaller than 3, the shift will likely be missed at the transient state and needs to be detected in the later runs. In this case, the steady-state ratio becomes more important for detecting the shift efficiently at the steady state. The steady-state ratio should not be tuned too high, however, because it may result in an extremely small transient ratio and make the transition of the shifts (6) from the transient state to the steady state very slow. To make the chart detect the shift efficiently in the steady state, a balance is needed to make a trade-off between the transient ratio and the steady-state ratio when choosing the charting parameters. Generally, a value of R_S around 3 is appropriate for balancing the values of R_T and R_S. Based on these guidelines, we have developed a heuristic algorithm to choose appropriate ARMA chart parameters, which is given in Figure 7. This algorithm is used to determine the parameters in the simulation study for autocorrelated processes in Section 4.

It is important to note that the implementation of the algorithm relies on two quantities of interest, the shift level μ and the ratio σ_X/σ_Z through (9)–(10). The ratio σ_X/σ_Z can be calculated based on the process parameters and charting parameters (ϕ, θ). In practice, the process model is often unknown and needs to be estimated. Then this ratio can be estimated based on process parameter estimates. Alternatively, we can estimate σ_X and σ_Z directly from the original in-control trial data and the ARMA statistics and then calculate the ratio using sample standard deviations. Therefore, as long as the process is stationary, it is not necessary to identify and estimate the underlying process to implement the algorithm for choosing appropriate charting parameters in practice. We will illustrate the design procedure step by step in Section 5.

4. SIMULATION STUDIES

4.1 IID Observations

We will first study the performance of the ARMA chart under iid processes. Lucas and Saccucci (1990) showed that, for iid processes, the optimal choice of the parameter λ depends on the size of the mean shift to be detected. According to their table 4, for example, the optimal λ value for a mean shift of one standard deviation is .15. This λ value corresponds to $\phi = 1 - \lambda = .85$ for the ARMA chart. Therefore, in our simulation study, we focus on how the value of θ affects the performance of the ARMA chart for the given optimal value of ϕ. Table 2 shows the ARL of ARMA charts for various values of θ. The ARL's were obtained based on 250,000 simulation runs with control limits placed at L standard deviations so that the in-control ARL is approximately 500. The simulation error (shown inside the parentheses) is negligible.

In the simulation study, the EWMA chart (with $\theta = 0$) is found to be very close to optimal in terms of minimizing the out-of-control ARL's. For example, as shown in Table 2, for the detection of a mean shift of one standard deviation, the optimal choice of $\theta = -.03$ gives an ARL = 9.99, which is only slightly better than the EWMA chart with $\theta = 0$. For

Table 2. ARMA Charts Compared with the Corresponding Optimal EWMA Chart (ϕ = .85) for Detecting Mean Shift of $1.0\sigma_a$

			ARMA					
Charts	EWMA λ = .15	θ	−.075	−.05	−.03	.03	.10	.30
L	2.913		2.832	2.843	2.867	2.952	3.023	3.080
μ = .0	502		502	501	501	500	500	498
	(1.00)		(1.00)	(1.00)	(1.00)	(1.00)	(1.00)	(1.00)
μ = .5	35.5		35.0	34.9	35.0	36.4	39.9	62.6
	(.06)		(.06)	(.06)	(.06)	(.06)	(.07)	(.15)
μ = 1.0	10.06		10.12	10.02	9.99	10.19	10.8	15.1
	(.01)		(.01)	(.01)	(.01)	(.01)	(.01)	(.03)
μ = 2.0	3.92		4.19	4.06	3.98	3.88	3.84	4.16
	(.00)		(.00)	(.00)	(.00)	(.00)	(.00)	(.01)
μ = 3.0	2.54		2.91	2.76	2.66	2.44	2.25	2.02
	(.00)		(.00)	(.00)	(.00)	(.00)	(.00)	(.00)
μ = 4.0	1.96		2.37	2.21	2.10	1.83	1.57	1.27
	(.00)		(.00)	(.00)	(.00)	(.00)	(.00)	(.00)

Table 3. ARL Performance of ARMA Charts for Detecting μ = 1.0

ϕ/θ	−.6	−.5	−.3	−.2	−.15	−.1	−.05	−.03	0	.05	.1	.3
.95	46.0	38.4	25.0	19.2	16.8	14.3	12.2	11.6	11.2	11.4	11.9	16.2
.85	29.0	22.2	14.4	12.0	10.3	10.2	10.02	9.99	10.03	10.23	10.8	15.1
.63	22.4	18.0	13.2	12.3	12.2	12.3	12.7	13.0	13.3	14.2	15.4	24.0
.30	23.2	20.3	18.8	19.9	21.0	22.6	24.7	25.6	27.3	30.4	34.1	55.3
.05	26.3	24.7	27.3	31.9	35.7	38.8	43.2	45.2	48.6	54.3	61.2	92.6

detection of shifts other than one standard deviation, the ARMA chart gives different ARL performance from the EWMA chart. When the shift is small, the ARMA chart performs better than the EWMA chart. When the shift is large, the reverse is true. Overall, the ARL is sensitive to the choice of θ. For $\theta = .3$, the ARL increases to 15.1 at a shift of $1\sigma_a$ but decreases from 2.10 to 1.27 at a shift of $4\sigma_a$. Nevertheless, the EWMA chart is nearly optimal among all charts of the ARMA type for detecting a mean shift of $1\sigma_a$. Similar comparison results are observed for other optimal EWMA charts when detecting shifts in the mean of .5, 2, 3, and 4 standard deviations (see Jiang 1999 for detailed results).

Note that the preceding simulation study was performed under the assumption that the optimal choice of ϕ (or $1-\lambda$) for the ARMA chart is the same as the optimal λ value in the EWMA chart. To verify this assumption for $\mu = 1.0$, a grid search was conducted over the joint space of ϕ and θ. The results are summarized in Table 3, which is consistent with our assumption. Because this study is limited, more work is needed to determine the optimal choices of ϕ and θ jointly.

As shown in the appendix, a Markov chain approach can be used to approximate the run-length distribution of the ARMA charts for iid observations and the approximation accuracy is quite adequate. In summary, the performance of the ARMA chart is very similar to that of the EWMA chart when the parameters of both charts are optimized. The optimal ARMA chart can be slightly better than the EWMA chart, but the difference is small. Although the EWMA chart is a special class of the ARMA chart, it seems that the choice of the additional parameter θ does not improve the

Table 4. Comparisons of ARL's: ARMAST, EWMAST, and SCC on ARMA(1, 1) Process

	Process parameters		Charting parameters					
Shift	u	v	ϕ	θ	ARMAST	EWMAST	SCC	WBM
.0	−.95	.0	.0	−.49	370	370	370	370
.5					2.65	4.31	2.67	3.16(2)
1					1.42	2.20	1.42	2.00(2)
2					1.00	1.29	1.00	2.00(2)
3					1.00	1.01	1.00	2.00(2)
.0	−.475	.0	.9	.1	370	370	370	370
.5					13.2	14.7	65.5	17.1(8)
1					4.78	4.97	11.4	6.27(4)
2					2.31	2.32	2.20	2.79(2)
3					1.64	1.63	1.35	2.02(2)
.0	.475	.0	.9	.1	370	370	370	370
.5					65.6	83.3	253	65.6(32)
1					20.3	22.4	118	25.5(16)
2					6.61	6.17	22.6	9.78(8)
3					3.67	3.40	4.20	5.50(4)
.0	.95	.0	.92	.4	370	370	370	370
.5					226	237	331	247(128)
1					102	108	139	136(64)
2					25.8	25.7	1.08	60.7(32)
3					8.65	8.30	1.00	36.7(16)
.0	.475	−.9	.9	.1	380	370	370	—
.5					84.7	105	109	—
1					25.4	29.8	22.8	—
2					7.94	7.68	2.79	—
3					4.29	4.02	1.01	—
.0	.95	.45	−.9	.1	378	370	370	—
.5					224	226	350	—
1					95.4	97.5	275	—
2					23.6	21.9	43.5	—
3					5.14	7.15	1.30	—
.0	.95	−.90	−.9	36.1	370	370	370	—
.5					42.8	240	42.8	—
1					1.00	110	1.00	—
2					1.00	26.4	1.00	—
3					1.00	8.48	1.00	—

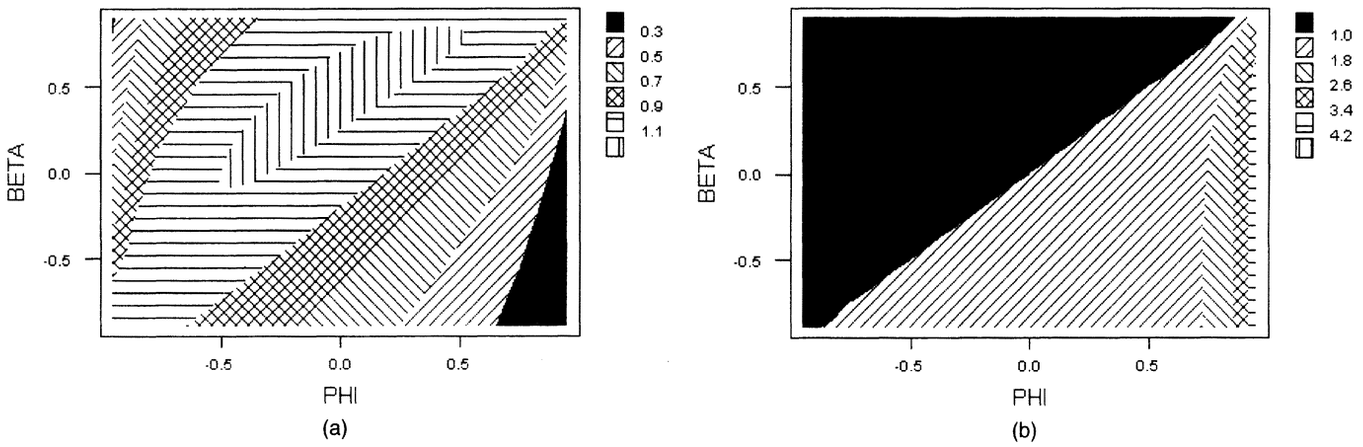

Figure 8. Plots of Signal-to-Noise Ratios for u = .475, v = 0, μ = 1: (a) Value of R_T; (b) Value of R_S.

ARL performance of the EWMA chart significantly when the underlying process is iid. As we shall show, however, when the underlying process is autocorrelated the ARMA chart can be significantly better than the EWMA chart if the parameters are chosen appropriately.

4.2 Autocorrelated Observations

In this section we study how the ARMAST chart performs when applied to an autocorrelated process. Simulation studies by Wardell et al. (1992, 1994) and Zhang (1998) showed that the SCC chart and the EWMAST chart do not uniformly outperform each other. In general, the SCC chart performs better when the absolute value of the first-lag autocorrelation is higher and the EWMAST chart performs better when the absolute value is lower. In this study, we consider several typical autocorrelated processes in which the first four are AR(1) processes and the last three are selected from ARMA(1, 1) processes. When the shift is measured in terms of the standard deviation of the original ARMA process, Table 4 shows that the EWMAST chart is much better than the SCC chart in some cases (e.g., $u = .475, v = 0$) but much worse in some others (e.g., $u = .95, v = -.9$). In the cases, for example, $u = .95, v = -.0$ and $u = .475, v = -.9$, either one is slightly better for smaller shifts and worse for larger shifts than the other. As explained in Section 3.2, because the EWMAST and SCC charts are special cases of the ARMA chart, it is possible to improve the chart performance by choosing the chart parameters (ϕ and θ) appropriately. Table 4 shows the comparison of the ARMAST chart with the EWMAST and SCC charts. The parameters of the ARMAST chart in Table 4 were determined based on the flow chart in Figure 7. The ARMAST chart with appropriately chosen parameters either outperforms or performs competitively with the better of the EMMAST and SCC charts, which is consistent with our expectation. In addition, the performance of the ARMAST chart for the last case ($\phi = .95, \theta = -.9$) is exactly the same as that of the SCC chart as we expected because $\phi = v$ and $\beta = \theta/\theta_0 = u$.

Table 5. Comparison of ARL's of Alternative Charts for an AR(1) Process With u = .9

Charts	L	Shift 0	1	2	3
Shewhart of AR Residuals (SCC)	3.00	370	223	11	1
CUSUM of AR Residuals, K = .5	4.78	370	130	17	1
CUSUM of AR Residuals, K = .125	12.1	370	79	26	12
Run sums	106	380	80	28	17
EWMAST λ = .2	2.4	389	84	20	7.3
ARMAST ϕ = .9, θ = .4	2.49	372	75	18.2	6.4

To further understand the results, we investigate how the performance depends on the signal-to-noise ratios discussed in Section 3.2. For the autocorrelated process with $u = .475$ and $v = 0$, Figure 8 shows the transient and the steady-state ratios for different charting parameters (ϕ and β) when $\mu = 1.0$. Figure 8 shows that the EWMAST chart with $\lambda = .2$ ($\phi = .8, \beta = 0$) has a lower transient ratio (.40) than that (1.14) of the SCC chart ($\phi = 0, \beta = .475$). The steady-state ratio of the EWMAST chart is 2.10, however, which is much larger than that (.60) of the SCC chart. This may explain why the EWMAST chart outperforms the SCC chart. For the ARMAST chart with appropriately chosen parameters ($\phi = .9, \beta = .5$; i.e., $\theta = .1$), both the transient and steady-state ratios (.52 and 2.58, respectively) are higher than those of the EWMAST chart. As a result, the ARMAST chart significantly improves the efficiency of detecting small mean shifts.

Other procedures have been proposed to handle autocorrelated data. For example, the weighted batch mean (WBM) chart proposed by Runger and Willemain (1995) applies the idea of reducing autocorrelation among batch means and monitoring each batch mean using traditional control charts. Runger and Willemain (1995) provided a theoretical study of the optimal weights for AR(1) processes. Table 4 shows the performance of the WBM chart with optimal batch size indicated in parentheses. The WBM chart is not as competitive as the ARMAST chart in all AR(1) processes. For ARMA(1, 1) processes, it is difficult to derive the optimal batch size and weights for the WBM chart. Theoretically, time series modeling of the underlying process is needed to get the appropriate batch size and optimal weights.

In addition, the run-sums technique proposed by Willemain and Runger (1998) is also compared with the

Table 6. Comparison of ARL's With Atienza et al.'s (1998) Approach for AR(1) Processes

ϕ	Charts	L	Shift 0	.5	1	2	3
.00	$\lambda_{LS,max}$	3.41	373	31.4	9.87	3.26	1.83
	$\bar{\lambda}_{LS}$	1.64	377	30.0	14.3	7.14	4.78
	ARMAST	2.76	370	30.7	9.46	3.82	2.59
	$\phi = .85, \theta = -.03$						
.50	$\lambda_{LS,max}$	3.46	382	72.3	22.6	5.86	2.12
	$\bar{\lambda}_{LS}$	1.64	378	54.0	24.8	11.5	7.42
	ARMAST	2.50	388	64.5	20.5	7.66	4.83
	$\phi = .92, \theta = -.02$						
.90	$\lambda_{LS,max}$	3.43	374	184	60.7	2.63	1.00
	$\bar{\lambda}_{LS}$	1.14	381	151	65.5	26.8	14.9
	ARMAST	1.67	370	150	65.6	26.7	15.4
	$\phi = .99, \theta = .04$						

ARMAST chart. Table 5 shows the comparison of the ARMAST chart with the charts given by Willemain and Runger (1998) for an AR(1) process with $u = .9$. The run-sums chart is competitive to the optimal cumulative sum (CUSUM) chart ($K = .125$) of residuals for small shifts (e.g., $\mu = 1$) but is much worse than the SCC chart and the CUSUM chart of residuals with $K = .5$ for large shifts (e.g., $\mu = 3$). On the other hand, the ARMAST chart with $\phi = .9, \theta = .4$ has the best performance for detecting small shifts and is better than the run-sums chart and the CUSUM chart ($K = .125$) of residuals for detecting large shifts. Nevertheless, the ARMA chart is worse than the SCC chart for detecting large shifts because the ARMA chart is designed to detect small shifts.

It is important to note that the ARMA chart can be designed for detecting a particular level of shifts efficiently. Table 6 presents another comparison with the approach of Atienza, Tang, and Ang (1998), which is based on the time series procedure for detecting outliers and level shifts (LS) proposed by Tsay (1988). With each new observation at time t, a least squares estimate of the shift at time d is computed and standardized as $\lambda_{LS,d}$. Then two monitoring statistics $\lambda_{LS,max} = \max_{1 \le d \le t} \lambda_{LS,d}$ and $\bar{\lambda}_{LS} = \sum_{d=1}^{t} \lambda_{LS,d}/t$ are derived for testing if there is a level shift. Atienza et al. recommended using $\lambda_{LS,max}$ because it is sensitive in detecting small shifts. The ARMA charts designed for detecting small shifts are found to be better than or competitive to their approach but worse in detecting large shifts. Actually the ARMA chart performs between that of the two charts based on $\lambda_{LS,max}$ and $\bar{\lambda}_{LS}$ but has much simpler derivations.

5. EXAMPLE REVISITED

So far, our analysis focuses on the design and analysis of the ARMA chart when the model of the underlying process is known or can be accurately estimated. In practice, it is possible to design the ARMA chart without knowing the process model. We now use the second example in Section 1 to illustrate the design procedure in Figure 7 when the process model is unknown. (Note that using real data can only provide clues to practical implementation, however, because the true underlying models are always unknown.)

Given the first 100 stable observations from the process, our objective is to design an ARMA chart to quickly detect a mean shift of one standard deviation of the process. Recall that

$$\theta = \frac{\beta(1-\phi)}{1-\beta} \quad \text{and} \quad \theta_0 = \frac{1-\phi}{1-\beta}.$$

Following simple calculations, the standard deviation of the original process, σ_X, can be estimated by the sample autocovariance of the data, and the standard deviation of the charting process, σ_Z, can be estimated from the sample autocovariance of the ARMA statistics. The ratio of σ_X/σ_Z can then be estimated and the two signal-to-noise ratios defined in Equations (7) and (8) can be obtained as $R_T = \theta_0 \sigma_X/\sigma_Z$ and $R_S = \sigma_X/\sigma_Z$. Figure 9 shows the two signal-to-noise ratios in different regions of the charting parameters. When $\phi = -.8$ and $\beta = .75$ (i.e., $\theta = 5.4$), the transient ratio can reach more than 3.8, while the steady-state ratio remains around .5. These parameters may have advantages to detect shifts larger than $1\sigma_X$ very quickly without loss of much efficiency when detecting smaller shifts. Figure 10(a) shows the plot of this ARMA chart with adjusted control limits ($\pm 2.998\sigma$). It signals the shift right after the shift occurrence as we expect. Its performance is quite similar to the SCC chart [Fig. 4(a)]. For comparison, the EWMAST chart with $\lambda = .2$, which corresponds to the ARMAST chart with $\phi = .8$ and $\beta = .0$, has ratios .30 and

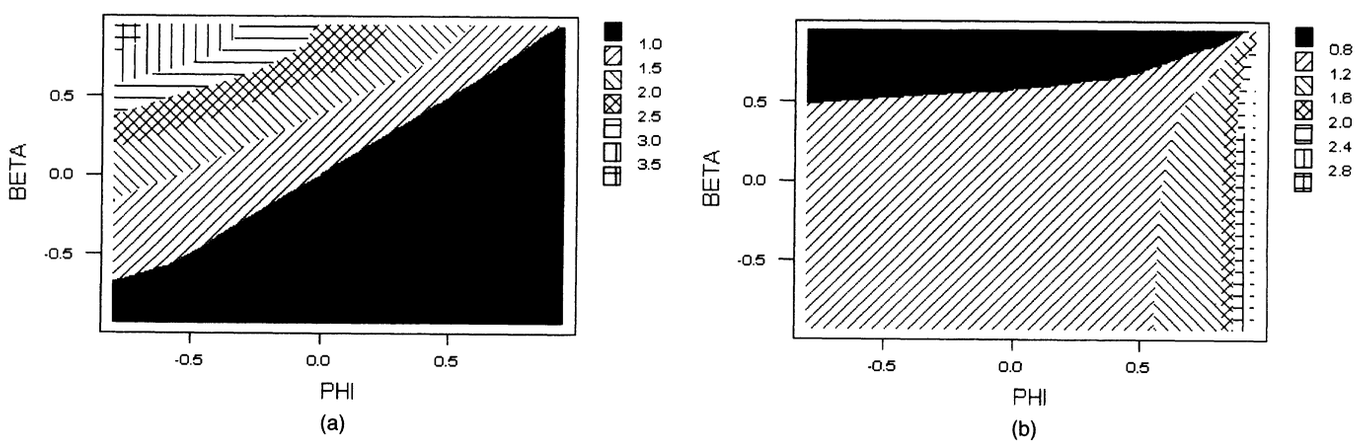

Figure 9. Plots of Signal-to-Noise Ratios for the Second Example: (a) Value of R_T; (b) Value of R_S.

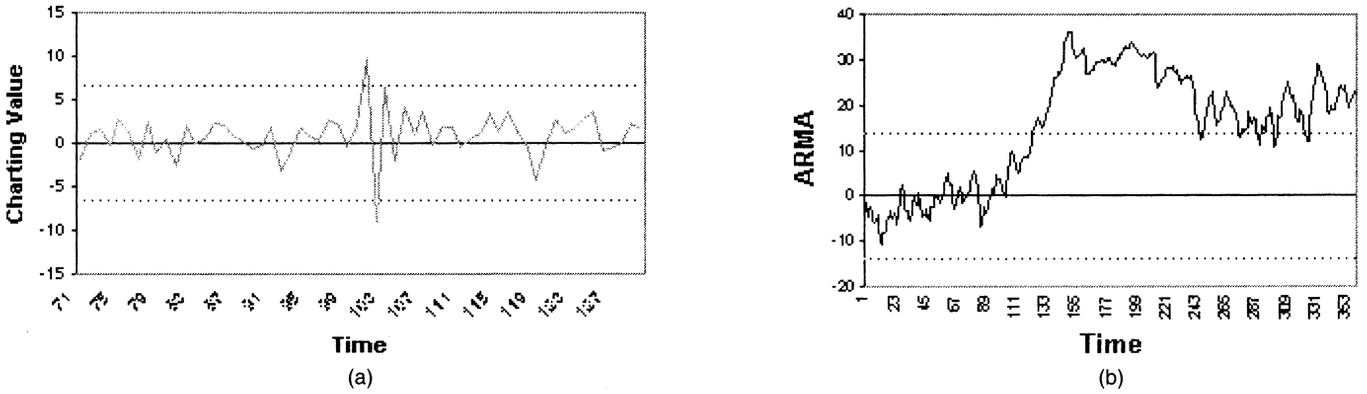

Figure 10. ARMA Charts for the Examples: (a) Example 2; (b) Example 1.

1.51. This may explain why the EWMAST chart has low detection capability and could not signal any shift within the following 30 observations.

Figure 10(b) shows an ARMAST chart (with ± 2.568 control limits) that adds one more term (θ) to the EWMAST chart, $\theta = -.086$, and may detect the possible increase one step before the EWMAST chart with the same ϕ value for the first numerical example.

6. CONCLUSIONS AND FUTURE RESEARCH

By extending the EWMA control chart, a general control-charting technique, the ARMA chart, is proposed in this article. The idea is to monitor the ARMA statistic of the underlying process. For monitoring iid processes, the performance of the proposed ARMA chart is found to be comparable to the optimal EWMA chart. For monitoring autocorrelated processes, the SCC chart and the EWMAST chart are found to be special cases of the ARMA chart. An informal strategy is developed to determine the appropriate parameter values. Based on some simulation studies, the ARMA chart with the appropriately chosen parameters outperforms the EWMAST, SCC, and other charts proposed for autocorrelated data in many cases.

The current study focuses on the first-order ARMA chart. It can be easily extended to higher-order ARMA models. It can be shown that the EWMA chart applied to the residuals from the SCC chart (Lin and Adams 1996; Lu and Reynolds 1999b) can be modeled as a special case of the higher-order ARMA charts. The design of higher-order ARMA charts may involve too many parameters for tuning, however, which would make it too complicated for implementation. Among them, the AR(2) chart, which has only two parameters, is still promising for further investigations.

ACKNOWLEDGMENTS

We thank the editor, an associate editor, and two referees for helpful suggestions that significantly improved the quality of this article. This work was supported by NSF-DMI grants 9908013 and 9908032 and Research Grant Council CERG grant #HKUST6134198E.

APPENDIX: RUN-LENGTH DISTRIBUTION OF THE ARMA CHART IN IID CASE

For EWMA charts, run-length distributions have been studied extensively. Several approximation approaches have been proposed to analyze the ARL performance. Crowder (1987) applied the numerical method to solve an integral equation for the approximation. Lucas and Saccucci (1990) used a Markov chain model to investigate the ARL values and the design strategies. We shall approximate the ARL of the ARMA chart on an iid process using the Markov chain method.

It is easy to see that the distribution of the ARMA(1, 1) statistic in Equation (1) depends on both Z_{t-1} and a_{t-1}; thus Z_t does not have the Markov property. The random vector $\mathbf{W}_t = (Z_t, a_t)'$ is a Markov chain, however, and it can be written as

$$\mathbf{W}_t = \lambda \mathbf{Y}_t + (1-\lambda)\mathbf{W}_{t-1}, \tag{A.1}$$

where $\lambda = 1 - \phi$ and $\mathbf{Y}_t = [(\theta_0 a_t - \theta a_{t-1})/\lambda, (a_t - \phi a_{t-1})/\lambda]'$. To evaluate the ARL for the ARMA chart with control limits $\pm L_Z$, we need to set large control limits $\pm L_a$ for a_t so that the in-control region can be segmented within a two-dimensional rectangle, $[-L_Z, L_Z] \times [-L_a, L_a]$.

Following Lucas and Saccucci (1990) and Runger and Prabhu (1996), the in-control region is divided into $M = t_1 \times t_2 = (2m_1+1) \times (2m_2+1)$ subregions of width $(2\delta_1) \times (2\delta_2)$. The control variable $\mathbf{W}_i = (Z_i, a_i)'$ is said to be in transient state (j) at time (i) if $SZ_j - \delta_1 < Z_i \le SZ_j + \delta_1$ and $Sa_j - \delta_2 < a_i \le Sa_j + \delta_2$, for $1 \le j \le M$, where SZ_j and Sa_j are the partition points of the region. Then the transition probability matrix at each run becomes

$$P = \begin{pmatrix} R & (I-R)\cdot \mathbf{1} \\ 0^T & \mathbf{1} \end{pmatrix},$$

where the submatrix R contains the probabilities of going from one transient state to another, I is the identity matrix, and $\mathbf{1}$ is a column vector of ones. To simplify the calculations, shift magnitudes and control limits are scaled in terms of process standard deviation. When the process mean is μ, the transition probabilities can be calculated as

$$\begin{aligned} p_{jk} &= \Pr\{\text{going to } S_k | \text{in } S_j\} \\ &= \Pr\{SZ_k - \delta_1 < Z_t \le SZ_k + \delta_1, Sa_k - \delta_2 \end{aligned}$$

$$
\begin{aligned}
&< a_t + \mu \le Sa_k + \delta_2 | Z_{t-1} = SZ_j, a_{t-1} = Sa_j\} \\
= \;& \Pr\{SZ_k - \delta_1 < \phi SZ_j + \theta_0 a_t - \theta Sa_j + \theta_0 \mu \\
\le \;& SZ_k + \delta_1, Sa_k - \delta_2 < a_t + \mu \le Sa_k + \delta_2\} \\
= \;& \Pr\{\min[b_U, \max(b_L, c_L)] < a_t \\
\le \;& \max[b_L, \min(b_U, c_U)]\},
\end{aligned}
$$

where $b_L = (SZ_k - \phi SZ_j + \theta Sa_j - \theta_0\mu - \delta_1)/\theta_0, b_U = (SZ_k - \phi SZ_j + \theta Sa_j - \theta_0\mu + \delta_1)/\theta_0$, and $c_L = Sa_j - \mu - \delta_2, c_U = Sa_j - \mu + \delta_2$. With this notation and the same arguments of Lucas and Saccucci (1990), both the in-control and out-of-control ARL's can be evaluated as ARL $= p^T \cdot (I - R)^{-1} \cdot 1$, with the means being 0 and μ, respectively.

Similarly to Runger and Prabhu (1996), the algorithm is implemented in SAS/IML on a UNIX Workstation, where the control limits for a_t are chosen as ± 4.0 because Z_t would likely go outside the control limits whenever a_t exceeds these limits. The results from this implementation agree with the simulation results quite well. Table A.1 shows the analytical results when $\phi = .85$. Compared with Table 2, the discrepancies are within 3% of the simulated ARL's.

The approximation accuracy of the algorithm depends critically on the capability of the memory. When m_1 and m_2 increase, the transition matrix can become quite large and a large memory is needed. To achieve better precision with limited memory space, ARL(m) = ARL(m_1, m_2) is evaluated for $m = m_1 = m_2 = 10, 12, 14, 16$, and 18. Following the method of Brook and Evans (1972), the continuous-state ARL is approximated as the least squares estimate of the intercept of the quadratic equation ARL(m) = ARL$(\infty) + B/m + C/m^2$.

Note that the Markov chain approximation for the MEWMA charts considered by Runger and Prabhu (1996) and that for the univariate ARMA chart here are different. For the MEWMA chart, the control limit is defined for plotting T_t^2, which is a quadratic form of the variables. For the ARMA chart, the control limits are defined as $|Z_t| > L_Z\sigma_Z$ and $|a_t| > L_a\sigma_a$. The consequence is that their grids are inside a circle, whereas our grids are defined in a rectangle. More memory space is needed for implementing our algorithm.

Table A.1. ARMA Charts Compared With the Corresponding Optimal EWMA Chart (ϕ = .85) for Detecting μ = 1.0: Analytical Results

	EWMA		ARMA					
Charts	λ = .15	θ	−.075	−.05	−.03	.03	.10	.30
L	2.913		2.832	2.843	2.868	2.952	3.023	3.080
μ = .0	499		503	498	496	501	503	508
μ = .5	36.2		35.8	35.5	35.9	36.7	40.6	62.0
μ = 1.0	10.3		10.3	10.2	10.1	10.8	11.0	15.6
μ = 2.0	3.97		4.25	4.11	4.01	3.92	3.85	4.16
μ = 3.0	2.56		2.94	2.78	2.69	2.47	2.25	2.00
μ = 4.0	2.01		2.31	2.22	2.11	1.86	1.58	1.25

[Received October 1998. Revised December 1999.]

REFERENCES

Adams, B. M., and Tseng, I.-T. (1998), "Robustness of Forecast-Based Monitoring Schemes," *Journal of Quality Technology*, 30, 328–339.

Alwan, L. C., and Roberts, H. V. (1988), "Time-Series Modeling for Statistical Process Control," *Journal of Business & Economic Statistics*, 6, 87–95.

Atienza, O. O., Tang, L. C., and Ang, B. W. (1998), "A SPC Procedure for Detecting Level Shifts of Autocorrelated Processes," *Journal of Quality Technology*, 30, 340–351.

Box, G. E. P., Jenkins, G. M., and Reinsel, G. C. (1994), *Time Series Analysis Forecasting and Control* (3rd ed.), Englewood Cliffs, NJ: Prentice-Hall.

Brook, D., and Evans, D. A. (1972), "An Approach to the Probability Distribution of CUSUM Run Lengths," *Biometrika*, 59, 539–549.

Crowder, S. V. (1987), "A Simple Method for Studying Run Length Distributions of Exponentially Weighted Moving Average Control Charts," *Technometrics*, 29, 401–407.

Jiang, W. (1999), "Charting Techniques for Integrated APC and SPC Environments," unpublished Ph.D. dissertation, The Hong Kong University of Science & Technology, Dept. of Engineering and Engineering Management.

Lin, S. W., and Adams, B. M. (1996), "Combined Control Charts for Forecast-Based Monitoring Schemes," *Journal of Quality Technology*, 28, 289–301.

Lu, C. W., and Reynolds, M. R., Jr. (1999a), "Control Charts for Monitoring the Mean and Variance of Autocorrelated Processes," *Journal of Quality Technology*, 31, 259–274.

——— (1999b), "EWMA Control Chart for Monitoring the Mean of Autocorrelated Processes," *Journal of Quality Technology*, 31, 166–188.

Lucas, J. M., and Saccucci, M. S. (1990), "Exponentially Weighted Moving Average Control Schemes: Properties and Enhancements," *Technometrics*, 32, 1–12.

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

Runger, G. C., and Prabhu, S. (1996), "A Markov Chain Model for the Multivariate Exponentially Weighted Moving Average Control Chart," *Journal of the American Statistical Association*, 91, 1701–1706.

Runger, G. C., and Willemain, T. R. (1995), "Model-Based and Model-Free Control of Autocorrelated Processes," *Journal of Quality Technology*, 27, 283–292.

Schmid, W. (1997), "On EWMA Charts for Time Series," in *Frontiers in Statistical Quality Control* (Vol. 5), eds. H. J. Lenz, P.-Th. Wilrich, Heidelberg: Physica-Verlag, pp. 115–137.

Tsay, R. S. (1988), "Outliers, Level Shifts, and Variance Changes in Time Series," *Journal of Forecasting*, 7, 1–20.

VanBrackle, L. N., III, and Reynolds, M. R., Jr. (1997), "EWMA and CUSUM Control Charts in the Presence of Correlation," *Communications in Statistics—Simulation and Computation*, 26, 979–1008.

Wardell, D. G., Moskowitz, H., and Plante, R. D. (1992), "Control Charts in the Presence of Data Correlation," *Management Science*, 38, 1084–1105.

——— (1994), "Run-Length Distributions of Special-Cause Control Charts for Correlated Observations," *Technometrics*, 36, 3–17.

Willemain, T. R., and Runger, G. C. (1998), "Statistical Process Control Using Run Sums," *Journal of Statistical Computation and Simulation*, 61, 361–378.

Zhang, N. F. (1998), "A Statistical Control Chart for Stationary Process Data," *Technometrics*, 40, 24–38.