© 1992 American Statistical Association and
the American Society for Quality Control

TECHNOMETRICS, FEBRUARY 1992, VOL. 34, NO. 1

A Multivariate Exponentially Weighted Moving Average Control Chart

Cynthia A. Lowry

Finance and Decision Science Department
Texas Christian University
Fort Worth, TX 76129

William H. Woodall

Department of Management Science and Statistics
University of Alabama
Tuscaloosa, AL 35487-0226

Charles W. Champ and Steven E. Rigdon

Department of Mathematics and Statistics
Southern Illinois University at Edwardsville
Edwardsville, IL 62026-1653

A multivariate extension of the exponentially weighted moving average (EWMA) control chart is presented, and guidelines given for designing this easy-to-implement multivariate procedure. A comparison shows that the average run length (ARL) performance of this chart is similar to that of multivariate cumulative sum (CUSUM) control charts in detecting a shift in the mean vector of a multivariate normal distribution. As with the Hotelling's χ^2 and multivariate CUSUM charts, the ARL performance of the multivariate EWMA chart depends on the underlying mean vector and covariance matrix only through the value of the noncentrality parameter. Worst-case scenarios show that Hotelling's χ^2 charts should always be used in conjunction with multivariate CUSUM and EWMA charts to avoid potential inertia problems. Examples are given to illustrate the use of the proposed procedure.

KEY WORDS: Average run length; Hotelling's T^2 chart; Multivariate CUSUM; Statistical process control.

1. INTRODUCTION

With modern data-acquisition equipment and on-line computers used during production, it is now common to monitor several correlated quality characteristics simultaneously. Various types of multivariate quality-control charts have been proposed to take advantage of the relationships among the variables. Alt (1984) and Jackson (1985) reviewed much of the literature on this topic.

Suppose that the $p \times 1$ random vectors $\mathbf{X}_1$, $\mathbf{X}_2$, $\mathbf{X}_3$, . . . , each representing the p quality characteristics to be monitored, are observed over time. These vectors may represent individual observations or sample mean vectors. To study the performance of the various multivariate control charts, it will be assumed that $\mathbf{X}_i$, $i = 1, 2, \ldots$, are independent multivariate normal random vectors with mean vectors $\boldsymbol{\mu}_i$, $i = 1, 2, \ldots$, respectively. If the independence assumption is not met, then, as in the univariate case, average run length (ARL) properties of the charts can be greatly affected and out-of-control signals can become meaningless. For autocorrelated multivariate observations, an extension of Alwan and Roberts's (1988) approach is required using the theory of multiple time series. For simplicity, it is assumed that each of the random vectors has the known covariance matrix Σ. In practice it will be necessary to collect data over a substantial amount of time when the process is in control to estimate Σ. This data could also be used to check the assumptions of independence and multivariate normality. Further it is assumed without loss of generality that the in-control process mean vector is $\boldsymbol{\mu}_0 = (0, 0, \ldots, 0)' = \mathbf{0}$. Multivariate control charts for the mean vector are designed to detect shifts over time from this in-control vector.

Hotelling's (1947) multivariate control-chart procedure signals that a statistically significant shift in the mean has occurred; that is, it gives an out-of-control signal, as soon as

$$\chi_i^2 = \mathbf{X}_i' \, \Sigma^{-1} \, \mathbf{X}_i > h_1, \qquad (1.1)$$

where $h_1 > 0$ is a specified control limit. Because this procedure is based on only the most recent observation, it is insensitive to small and moderate shifts in the mean vector. Several alternative procedures

that use additional information from the recent history of the process have been proposed.

Woodall and Ncube (1985) recommended monitoring each of the p quality characteristics individually with cumulative sum (CUSUM) charts, and using the covariance matrix Σ to study the ARL performance of the minimum of the p univariate CUSUM run lengths. In practice, simultaneous univariate charts are frequently applied to correlated observations but usually without analyzing the resulting ARL performance.

Healy (1987) showed that a univariate CUSUM chart based on a linear combination of the variables can be used in the multivariate situation, if one is interested in detecting a shift in the mean vector in only one specified direction. Hawkins (1991) extended this approach to the case in which several directions of interest are specified and showed that this method can be more effective than that of Woodall and Ncube (1985).

Crosier (1988) proposed two multivariate CUSUM charts. The one with the better ARL properties is based on the statistics

$$C_i = \{(\mathbf{S}_{i-1} + \mathbf{X}_i)' \, \Sigma^{-1} \, (\mathbf{S}_{i-1} + \mathbf{X}_i)\}^{1/2}$$

and

$$\mathbf{S}_i = \mathbf{0}, \quad \text{if } C_i \le k_1,$$

$$= (\mathbf{S}_{i-1} + \mathbf{X}_i)(1 - k_1/C_i), \quad \text{if } C_i > k_1,$$

$i = 1, 2, \ldots$, where $\mathbf{S}_0 = \mathbf{0}$ and $k_1 > 0$. This multivariate CUSUM scheme, denoted by MCUSUM, signals when

$$Y_i = \{\mathbf{S}_i' \, \Sigma^{-1} \, \mathbf{S}_i\}^{1/2} > h_2, \tag{1.2}$$

where $h_2 > 0$.

Pignatiello and Runger (1990) also proposed two multivariate CUSUM charts. The chart with the better ARL performance, MC1, is based on the following vectors of cumulative sums:

$$\mathbf{D}_i = \sum_{j=i-\ell_i+1}^{i} \mathbf{X}_j$$

and

$$\mathrm{MC}_i = \max\{0, (\mathbf{D}_i' \, \Sigma^{-1} \, \mathbf{D}_i)^{1/2} - k_2\ell_i\}, \tag{1.3}$$

where $k_2 > 0$ and

$$\ell_i = \ell_{i-1} + 1, \quad \text{if } \mathrm{MC}_{i-1} > 0$$

$$= 1, \quad \text{otherwise},$$

$i = 1, 2, 3, \ldots$. An out-of-control signal is given as soon as $\mathrm{MC}_i > h_3$, where $h_3 > 0$.

In their other approaches Crosier (1988) and Pignatiello and Runger (1990) found that ordinary one-sided univariate CUSUM's based on successive values of χ_i and χ_i^2, respectively, do not have good ARL properties. Observations in varying directions from the target vector can cause the univariate CUSUM statistic to increase in value and result in unnecessarily frequent out-of-control signals for these charts when the process is actually in control.

In this article, a new multivariate procedure is proposed. This procedure, which we refer to as the MEWMA chart, is a multivariate extension of the univariate exponentially weighted moving average (EWMA) chart. In our opinion, this MEWMA chart, defined in (2.2) and (2.3), is a more straightforward generalization of the corresponding univariate procedure than the multivariate CUSUM statistics in (1.2) and (1.3).

MacGregor and Harris (1990) identified two uses of the univariate EWMA statistic. First, it can be used, under the assumption of independence of the observations, as a tool for detecting when special causes of variation enter into a system. This is the usual control-charting situation. Second, the EWMA statistic is the optimal predictor of the next observation from a first-order integrated moving average process, a simple but frequently applicable model for processes that exhibit continual drift. In this second form, the EWMA statistic can be used as part of a control algorithm to adjust the process, thereby reducing the mean squared error around the target. In this article, we discuss only the generalization of the first use of the EWMA statistic—that is, as a tool for monitoring the stability of a process. Forecasting applications of multivariate EWMA statistics were presented by Enns, Machak, Spivey, and Wrobleski (1982), Harvey (1986), and Jones (1966).

Design of MEWMA charts is discussed in Section 2, with examples given in Section 3. The ARL performance of the MEWMA chart is compared to that of some of the other multivariate control charts in Section 4. We briefly consider worst-case scenarios in which the multivariate CUSUM procedures and the MEWMA chart can be slow in reacting to process shifts in Section 5. We consider several approaches to the interpretation of out-of-control signals from multivariate control charts such as the MEWMA chart in Section 6. Our conclusions are presented in Section 7.

In Appendix A, we derive the covariance matrix of the vector of EWMA's used with the MEWMA procedure. This matrix is used in Section 2 to obtain the MEWMA control-chart statistic. It is shown in Appendix B that, if $\boldsymbol{\mu}_i = \boldsymbol{\mu}$, $i = 1, 2, \ldots$, the ARL performance of the MEWMA chart depends on the mean vector $\boldsymbol{\mu}$ and covariance matrix Σ only through

the value of the noncentrality parameter λ, where

$$\lambda = (\boldsymbol{\mu}' \sum^{-1} \boldsymbol{\mu})^{1/2}. \quad (1.4)$$

It is much easier to make ARL comparisons among several multivariate control charts if all of the charts have this property.

2. THE MEWMA PROCEDURE

Suppose that we observe $X_1, X_2, \ldots$ in the univariate case—that is when $p = 1$. The univariate EWMA chart is based on the values

$$Z_i = rX_i + (1 - r)Z_{i-1}, \quad (2.1)$$

$i = 1, 2, \ldots$, where $Z_0 = \mu_0 = 0$ and $0 < r \leq 1$. Roberts (1959) showed that, if $X_1, X_2, \ldots$ are iid $N(0, \sigma^2)$ random variables, then the mean of Z_i is 0 and the variance is $\sigma^2_{Z_i} = \{r[1 - (1 - r)^{2i}]/(2 - r)\}\sigma^2$, $i = 1, 2, \ldots$. Thus, when the in-control value of the mean is 0, the control limits of the EWMA chart are often set at $\pm L\sigma_{Z_i}$, where L and r are the parameters of the chart. Lucas and Saccucci (1990) discussed the choice of r and L for the univariate EWMA chart in detail, although their control limits were based on the asymptotic form of σ_{Z_i}; that is, $\sigma_{Z_i} \cong [r/(2 - r)]^{1/2}\sigma$.

In the multivariate case, a natural extension is to define vectors of EWMA's,

$$\mathbf{Z}_i = R\mathbf{X}_i + (I - R)\mathbf{Z}_{i-1}, \quad (2.2)$$

$i = 1, 2, \ldots$, where $\mathbf{Z}_0 = \mathbf{0}$ and $R = \text{diag}(r_1, r_2, \ldots, r_p)$, $0 < r_j \leq 1, j = 1, 2, \ldots, p$. The MEWMA chart gives an out-of-control signal as soon as

$$T_i^2 = \mathbf{Z}_i' \Sigma_{\mathbf{Z}_i}^{-1} \mathbf{Z}_i > h_4, \quad (2.3)$$

where h_4 (>0) is chosen to achieve a specified in-control ARL and $\Sigma_{\mathbf{Z}_i}$ is the covariance matrix of $\mathbf{Z}_i$, which is derived in Appendix A. If there is no a priori reason to weight past observations differently for the p quality characteristics being monitored, then $r_1 = r_2 = \cdots = r_p = r$. Note that the multivariate CUSUM procedures weight past observations in the same way for each quality characteristic. In this case, as shown in Appendix B, the ARL performance of the MEWMA chart depends only on the noncentrality parameter, λ, in (1.4). Of course, the user may use unequal weighting constants, but then the ARL depends on the direction of the shift, not just on the value of the noncentrality parameter. If the variables being monitored are not of equal importance and the desired ARL performance is such that the ARL should not be a function of λ, then the method of Hawkins (1991) is recommended. Another possibility, proposed by Tsui and Woodall (1991), is to use a different matrix in Equation (2.3) to calculate the quadratic form of the MEWMA chart.

If $r_1 = r_2 = \cdots = r_p = r$, then the MEWMA vectors can be written as

$$\mathbf{Z}_i = r\mathbf{X}_i + (1 - r)\mathbf{Z}_{i-1}, \quad (2.4)$$

$i = 1, 2, \ldots$, and from Appendix A we have

$$\Sigma_{\mathbf{Z}_i} = \{r[1 - (1 - r)^{2i}]/(2 - r)\}\,\Sigma. \quad (2.5)$$

Analogous to the situation in the univariate case, the MEWMA chart is equivalent to Hotelling's χ^2 chart if $r = 1$.

As MacGregor and Harris (1990) pointed out for the univariate case, using the exact variance of the EWMA statistic leads to a natural fast initial response for the EWMA chart. Thus initial out-of-control conditions are detected more quickly. This is also true for the MEWMA chart. Because, however, it may be more likely that the process will stay in control for a while and then shift out of control, we will assume for chart design and in our ARL comparisons that the asymptotic (as $i \to \infty$) covariance matrix, that is,

$$\Sigma_{\mathbf{Z}_i} = \{r/(2 - r)\}\,\Sigma, \quad (2.6)$$

is used to calculate the MEWMA statistic in (2.3) unless otherwise indicated.

Table 1 contains the ARL profiles of several MEWMA charts for varying values of r. As illustrated in this table, smaller values of r are more effective in detecting small shifts in the mean vector. This is analogous to the univariate case. Each MEWMA ARL value in Table 1, as well as those in later tables, was obtained using at least 6,000 simulations. The standard errors of the estimates of the ARL's are close to 1% of the estimated ARL. The appropriate values of h_4 were also obtained using simulation. The simulation program used in our article is available on request from C. Lowry.

Table 1. ARL Values for MEWMA Charts (p = 2)

	r			
λ	.8	.6	.4	.2
.0	200.	200.	199.	201.
.5	95.5	73.6	51.9	35.1
1.0	28.1	19.3	13.2	10.1
1.5	10.3	7.24	5.74	5.50
2.0	4.75	3.86	3.54	3.80
2.5	2.75	2.53	2.55	2.91
3.0	1.91	1.88	2.04	2.42
	h_4			
	10.58	10.53	10.29	9.65

Table 2. Optimal MEWMA Control Schemes (in-control ARL ≅ 200)

		p		
λ		2	3	4
.5	r	.06	.06	.06
	h_4	7.70	9.80	11.66
	ARL_{min}	26.5	30.0	32.3
	range of r	.04–.06	.04–.06	.04–.08
1.0	r	.16	.16	.14
	h_4	9.35	11.52	13.34
	ARL_{min}	9.95	11.0	12.0
	range of r	.10–.20	.10–.20	.10–.16
1.5	r	.24	.22	.20
	h_4	9.90	11.96	13.84
	ARL_{min}	5.47	6.02	6.45
	range of r	.18–.30	.20–.30	.18–.30
2.0	r	.34	.30	.28
	h_4	10.17	12.31	14.25
	ARL_{min}	3.53	3.87	4.20
	range of r	.24–.34	.24–.34	.22–.34

Table 2 shows the optimal values of r in detecting shifts of various sizes for in-control ARL = 200. The simpler Hotelling χ^2 chart is effective for detecting large shifts, so values of $\lambda > 2$ are not considered. The value of r is optimal in the sense that the ARL at the shift of interest is minimized for the given in-control ARL. An in-control ARL of 200 is considered for $p = 2, 3$, and 4. These results are analogous to those obtained in the univariate case by Lucas and Saccucci (1990) in that smaller values of r are more effective in detecting the smaller shifts. The ranges of r values given in Table 2 are those values for which the minimum out-of-control ARL's at the shift of interest were virtually identical. In the cases considered in Table 2, the optimal value of r does not vary much as the number of variables changes.

To aid in designing MEWMA charts, Table 3 shows the ARL performance of several MEWMA charts with varying in-control ARL's. The cases $p = 2, 3$, and 4 are considered, with $r = .1$ and in-control ARL's of 50, 100, 500, and 1,000. Charts with $r = .10$ and an in-control ARL of 200 can be found in Tables 5–10 (Sec. 4). We use $r = .10$ because this value is effective in detecting small shifts in the mean vector. In addition, it is shown in Section 4 that the ARL performance of the MEWMA charts with $r = .10$ compares favorably with that of the multivariate CUSUM charts. From the ARL results in Table 3 one can verify that the logarithm of the in-control ARL is very close to a linear function of h_4. This fact can be used to approximate appropriate control limits for other in-control ARL's.

3. TWO EXAMPLES

In this section, we present two examples. The first is a numerical example of the application of an MEWMA chart. The second example is on the design of an appropriate MEWMA chart.

Example 1: A numerical example of an application of the MEWMA chart is given in Table 4 using simulated data from a similar example given by Crosier (1988). The bivariate normal distribution is considered with unit variances and a correlation coefficient of .5. The process mean is on target at (0, 0) for the first five observations and then shifts to (1, 2) for the last five observations. In Table 4, the values of (X_1, X_2) are the observations, the values (Z_1, Z_2) correspond to the MEWMA vector in (2.4) with $r = .10$, and the values of T_i^2 were obtained using Equation (2.3) with either the covariance matrix in (2.5) or (2.6). The use of Equation (2.5) provides the natural head-start feature for the MEWMA chart. The values of h_4 were obtained using simulation to provide in-control ARL's of 200. The MEWMA chart based on (2.5) signals out-of-control after the ninth observation, whereas the MEWMA chart based on (2.6) signals after the tenth observation. A plot of the MEWMA chart based on (2.6) is shown in Figure 1. In Crosier's (1988) example, his MCUSUM procedure, which also had an in-control ARL of 200, signaled after the tenth observation. Note that when the MEWMA chart signals, the MEWMA vector elements, Z_1 and Z_2, give some indication of the direction of the shift. Crosier (1988) made this same point regarding his MCUSUM vector. Further information on the interpretation of out-of-control signals is given in Section 6.

Table 3. ARL Values for MEWMA Charts

	$p = 2$ and $r = .10$				$p = 3$ and $r = .10$				$p = 4$ and $r = .10$			
λ	$h_4 = 5.35$	$h_4 = 7.00$	$h_4 = 10.75$	$h_4 = 12.34$	$h_4 = 7.08$	$h_4 = 9.00$	$h_4 = 13.10$	$h_4 = 14.78$	$h_4 = 8.70$	$h_4 = 10.80$	$h_4 = 15.16$	$h_4 = 16.94$
0.0	50.7	100.	501.	999.	50.5	100.	502.	1007.	50.7	101.	497.	995.
0.5	15.8	21.4	39.5	51.1	17.5	23.7	45.6	61.2	18.8	26.1	52.3	68.0
1.0	7.14	8.63	12.1	13.7	7.91	9.61	13.5	15.3	8.62	10.4	14.5	16.5
1.5	4.61	5.35	7.03	7.69	5.12	5.95	7.66	8.40	5.54	6.42	8.20	8.99
2.0	3.43	3.91	4.97	5.39	3.81	4.32	5.43	5.83	4.16	4.70	5.79	6.23
2.5	2.78	3.16	3.90	4.23	3.07	3.48	4.26	4.54	3.31	3.74	4.51	4.81
3.0	2.36	2.64	3.27	3.49	2.59	2.91	3.52	3.75	2.81	3.12	3.74	3.97

Table 4. Numerical Example of the MEWMA Scheme

	Observations		MEWMA vector		MEWMA statistic	
i	X_1	X_2	Z_1	Z_2	T_i^2 (2.5)	T_i^2 (2.6)
1	−1.19	.59	−.12	.06	3.29	.62
2	.12	.90	−.10	.14	3.17	1.09
3	−1.69	.40	−.25	.17	7.37	3.45
4	.30	.46	−.20	.20	5.26	3.00
5	.89	−.75	−.09	.10	1.09	.71
6	.82	.98	.00	.19	1.28	.92
7	−.30	2.28	−.03	.40	5.66	4.37
8	.63	1.75	.04	.53	8.32	6.78
9	1.56	1.58	.19	.64	9.65*	8.20*
10	1.46	3.05	.32	.88	17.21	15.12
Control limit					$h_4 = 8.79$	$h_4 = 8.66$

*Out-of-control signal.

Example 2: The approach we recommend for the design of MEWMA charts is a straightforward generalization of the approach presented by Crowder (1989) and Lucas and Saccucci (1990) for the design of univariate EWMA charts. The value of r used depends on the size of the shift in the mean vector to be detected quickly, and h_4 is chosen to achieve the desired in-control ARL. As an example, consider the case in which $p = 2$, the desired in-control ARL is 200, and shifts of size $\lambda = 1.0$ are considered to be important. From Table 2, $r = .10$ is a reasonable choice for the smoothing constant r. From Table 5 (Sec. 4) the appropriate value for the control limit is $h_4 = 8.66$ if (2.6) is used to calculate the MEWMA statistic or $h_4 = 8.79$ if (2.5) is used. As discussed in Section 5, it is helpful to use a Hotelling χ^2 limit in conjunction with the MEWMA chart.

4. AVERAGE RUN LENGTH COMPARISONS

In this section, the following control procedures are compared on the basis of their ARL performance:

1. Hotelling's χ^2 chart
2. Crosier's (1988) CUSUM (MCUSUM) based on (1.2)
3. Pignatiello and Runger's (1990) multivariate CUSUM (MC1) based on (1.3)
4. The MEWMA chart with the covariance matrices in (2.5) used to calculate the statistic in (2.3)
5. The MEWMA chart with the covariance matrix in (2.6) used to calculate the statistic in (2.3)

For the cases in which the mean vector is off target, it is assumed that the shift has occurred prior to the application of the chart. Steady-state ARL comparisons would be helpful in comparing performance under delayed shifts in the mean vector. The ARL comparison is simplified because it can be demonstrated that the performance of each of these charts

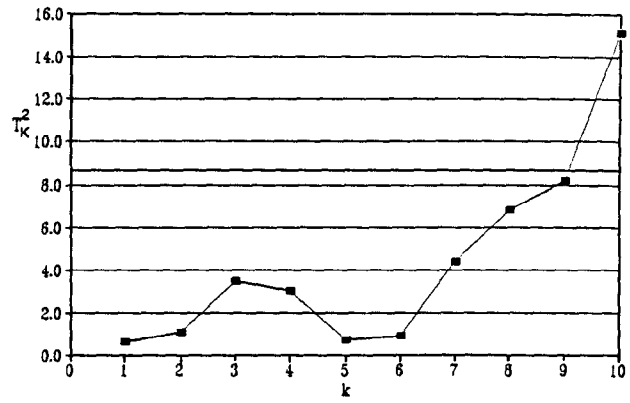


Figure 1. Plot of MEWMA's Based on (2.6) for Example 1.

depends on the values of $\boldsymbol{\mu}$ and Σ only through the value of the noncentrality parameter λ, which is defined in (1.4). This is true for Hotelling's χ^2 chart because, for a constant mean vector, the values of χ_i^2 are independent and have a noncentral chi-squared distribution with p df and noncentrality parameter λ. Crosier (1988) and Pignatiello and Runger (1990) presented proofs for their procedures, and a proof is outlined in Appendix B for the MEWMA chart. If the charts did not have this property, then their relative performance might vary depending on Σ and, even for a given covariance matrix, Σ, one chart might be more effective than another in detecting shifts in some directions and less effective in other directions.

Tables 5–10 show ARL comparisons for $p = 2$, 3, 4, 5, 10, and 20, respectively. Each chart is designed so that the on-target ARL is approximately 200. The ARL values for the multivariate CUSUM procedures were given by Crosier (1988) and Pignatiello and Runger (1990), respectively. Except for Hotelling's χ^2 chart, computer simulation was used in each case.

It is clear from this comparison that the Hotelling's χ^2 chart is the least effective chart in detecting shifts from the in-control mean vector unless the shift is very large. The MEWMA chart using the covariance matrix in (2.5)—that is, the exact covariance matrix—is more effective in detecting an initial out-of-control condition for reasons discussed earlier. The

Table 5. ARL Comparisons for $p = 2$

λ	χ^2; $h_1 = 10.6$	MCUSUM; $k_1 = .50$, $h_2 = 5.50$	MC1; $k_2 = .50$, $h_3 = 4.75$	MEWMA (2.5); $r = .10$, $h_4 = 8.79$	MEWMA (2.6); $r = .10$, $h_4 = 8.66$
.0	200.	200.	203.	200.	200.
.5	116.	28.8	31.3	25.3	28.1
1.0	42.0	9.35	9.28	7.76	10.2
1.5	15.8	6.94	6.23	4.07	6.12
2.0	6.9	4.20	3.69	2.59	4.41
2.5	3.5	3.26	2.91	1.89	3.51
3.0	2.2	2.78	2.40	1.50	2.92

Table 6. ARL Comparisons for $p = 3$

λ	χ^2; $h_1 = 12.85$	MCUSUM; $k_1 = .50$, $h_2 = 6.88$	MC1; $k_2 = .50$, $h_3 = 5.48$	MEWMA (2.5); $r = .10$, $h_4 = 10.97$	MEWMA (2.6); $r = .10$, $h_4 = 10.79$
.0	201.	200.	200.	202.	200.
.5	130.	32.7	33.5	28.5	31.8
1.0	52.6	11.2	10.1	8.64	11.30
1.5	20.5	6.69	5.66	4.47	6.69
2.0	8.8	4.70	4.00	2.84	4.86
2.5	4.4	3.83	3.17	2.08	3.83
3.0	2.6	3.17	2.63	1.62	3.20

Table 8. ARL Comparisons for $p = 5$

λ	χ^2; $h_1 = 16.75$	MCUSUM; $k_1 = .50$, $h_2 = 9.46$	MEWMA (2.5); $r = .10$, $h_4 = 14.74$	MEWMA (2.6); $r = .10$, $h_4 = 14.56$
.0	200.	200.	201.	201.
.5	145.	37.2	34.3	37.7
1.0	68.1	14.0	10.2	12.9
1.5	28.5	8.25	5.18	7.63
2.0	12.4	6.20	3.28	5.49
2.5	5.99	4.94	2.34	4.32
3.0	3.31	4.04	1.80	3.59

ARL performance of the MEWMA chart based on the asymptotic covariance matrix in (2.6) is very similar to that of the two multivariate CUSUM's. We recommend using the covariance matrix in Equation (2.5) to calculate the MEWMA statistic. If the process is initially out of control, then the ARL performance will be very good. If the process is initially in control but later shifts, then the ARL values for the MEWMA chart using (2.6) suggest that the MEWMA chart will react as quickly as the multivariate CUSUM charts.

Again, it is important to emphasize that there are two implicit assumptions in ARL comparisons based on the noncentrality parameter such as those in this section. First, it is assumed that any shift from the in-control mean vector, regardless of the size of the shift, is to be detected as quickly as possible. Second, it is assumed that a shift to $\boldsymbol{\mu} = \boldsymbol{\mu}_1$ is to be detected as quickly as a shift to $\boldsymbol{\mu} = \boldsymbol{\mu}_2$ if $\boldsymbol{\mu}_1' \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_1 = \boldsymbol{\mu}_2' \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_2$; that is, the ARL is a function of the noncentrality parameter λ. If these are not reasonable assumptions in a practical application, then a more general approach to the definition of "in control" and "out of control," such as that discussed by Woodall and Ncube (1985) and Mohebbi and Hayre (1989), is required. If directions of shifts of interest can be specified, then we recommend the approach of Hawkins (1991).

5. INERTIA PROBLEMS

Yashchin (1987) pointed out that univariate EWMA schemes can have "inertia" in reacting to shifts in the mean relative to other charts such as the two-sided CUSUM chart. For example, consider the EWMA chart with $\sigma = 1$, $r = .05$, and control limits $\pm .4189$ recommended by Lucas and Saccucci (1990) in the situation for which the in-control ARL is 500 and a shift in the mean of .5 is to be detected quickly. Then, if the process is actually on target but $Z_i = -.3$, we must have $X_{i+1} \geq 14.08$ to obtain a value of Z_{i+1} over the upper control limit. Thus, if the EWMA statistic is on one side of the center line when a shift in the other direction occurs, the result may be a delayed detection of the shift. This problem can be partially alleviated by using a Shewhart rule in combination with the EWMA chart, an enhancement discussed by Lucas and Saccucci (1990).

The same type of inertia problem can occur with the multivariate CUSUM charts and with the MEWMA chart. The multivariate CUSUM procedure of Pignatiello and Runger (1990) can, at least theoretically, build up an arbitrarily large amount of inertia. Thus a Shewhart-type rule should be implemented by using a Hotelling χ^2 control limit with the charts; that is, a signal occurs if the multivariate chart signals or if $\mathbf{X}_i' \boldsymbol{\Sigma}^{-1} \mathbf{X}_i > h_5$, where $h_5 > 0$. Crosier (1988) mentioned such a rule to be used in conjunction with his multivariate CUSUM. The use of the Hotelling's χ^2 limit does not, however, completely solve the inertia problem because, if inertia has built up, then a sequence of relatively large shifts from the target vector that does not trigger the χ^2 chart may still not result in an out-of-control signal from the multivariate CUSUM or MEWMA component of the combined procedure.

If the Hotelling χ^2 limit is used with the MEWMA chart, then there is a trade-off between protection

Table 7. ARL Comparisons for $p = 4$

λ	χ^2; $h_1 = 14.86$	MCUSUM; $k_1 = .50$, $h_2 = 8.15$	MEWMA (2.5); $r = .10$, $h_4 = 12.93$	MEWMA (2.6); $r = .10$, $h_4 = 12.73$
.0	200.	200.	201.	201.
.5	138.	34.2	32.1	34.7
1.0	60.9	12.2	9.46	12.1
1.5	24.6	7.42	4.84	7.23
2.0	10.6	5.42	3.07	5.18
2.5	5.19	4.34	2.22	4.10
3.0	2.93	3.65	1.72	3.41

Table 9. ARL Comparisons for $p = 10$

λ	χ^2; $h_1 = 25.91$	MCUSUM; $k_1 = .50$, $h_2 = 14.9$	MC1; $k_2 = .50$, $h_3 = 5.48$	MEWMA (2.5); $r = .10$, $h_4 = 22.91$	MEWMA (2.6); $r = .10$, $h_4 = 22.67$
.0	200.	200.	202.	201.	200.
.5	162.	43.2	43.9	44.3	48.1
1.0	92.8	18.6	12.6	12.8	15.9
1.5	44.7	11.8	7.66	6.28	9.16
2.0	20.6	8.79	5.66	3.96	6.55
2.5	9.9	7.03	4.55	2.79	5.15
3.0	5.2	5.86	3.84	2.14	4.28

Table 10. ARL Comparisons for p = 20

λ	χ^2; $h_1 = 40.0$	MCUSUM; $k_1 = .50$, $h_2 = 24.70$	MEWMA (2.5); $r = .10$, $h_4 = 37.32$	MEWMA (2.6); $r = .10$, $h_4 = 37.01$
.0	200.	200.	200.	201.
.5	174.	55.9	58.1	64.1
1.0	117.	27.2	16.4	20.0
1.5	66.2	17.9	7.96	11.3
2.0	34.3	13.3	4.96	8.00
2.5	17.4	10.8	3.43	6.26
3.0	9.1	9.1	2.59	5.19

against inertia and the quick detection of small shifts in the mean vector. This is because the MEWMA control limit h_4 must be increased slightly once a χ^2 limit is used to maintain the desired in-control ARL. The smaller the χ^2 limit h_5, the greater the protection against inertia and the more h_4 must be increased. At this time ARL tables are not available for the design of MEWMA charts with the Hotelling χ^2 limit. We have found, however, that one can maintain the desired in-control ARL, say ARL_0, by increasing the value of h_4 by 5% and using a value of h_5 such that the probability that the χ^2 limit is exceeded is $1/(5 \cdot ARL_0)$. This is only a rough rule of thumb. Of course, as an alternative one could use simulation to select values for h_4 and h_5 and then to evaluate the ARL performance of the combined procedure for various shifts in the mean vector to see if it is satisfactory.

6. INTERPRETATION OF OUT-OF-CONTROL SIGNALS

The interpretation of out-of-control signals from multivariate control charts can be quite difficult. For example, if a multivariate chart signals due to the joint effect of two or more correlated variables, the univariate charts on the separate variables may not even signal. This topic was discussed by Alt (1984), Chua and Montgomery (1991), Doganaksoy, Faltin, and Tucker (1991), Jackson (1985), Murphy (1987), and Pignatiello and Runger (1990), among others. Jackson (1985) and Pignatiello and Runger (1990) recommended monitoring the individual variables and/or the principal components with univariate charts. Alt (1984) and Doganaksoy et al. (1991) also recommended examining the individual variables but used Bonferroni limits for the univariate charts. The Bonferroni limits used by Doganaksoy et al. (1991) are less conservative than those recommended by Alt (1984) and thus more likely to indicate that a particular variable is out of control.

We recommend monitoring the principal components if these are interpretable. Otherwise, we recommend monitoring the original variables using joint univariate charts. All of the methods that have been proposed thus far classify individual variables as being in control or out of control based on only the most recent observation—that is, the one that produced the out-of-control signal. The univariate EWMA statistics are calculated as part of the MEWMA method, however, so we recommend using these values to decide whether or not a particular variable has shifted away from its target. Because this diagnosis is not performed until after an out-of-control signal from the MEWMA chart, the user is free to choose a decision criterion without affecting the false-alarm rate of the MEWMA chart. The MEWMA vector in (2.4) gives an indication of the direction of the shift but does not provide an accurate estimate of the current process mean vector.

In Example 1 in Section 3, the MEWMA vector is $(.32, .88)'$ when the MEWMA chart based on (2.6) signals. The corresponding z scores based on the variance of the univariate EWMA statistics are 1.39 and 3.84, respectively. Thus there is somewhat weak evidence that a shift in the first variable has occurred but strong evidence of a shift for the second variable. Recall that in this example there was a shift in the mean of both variables to the mean vector $(1, 2)'$. We also recommend, like Doganaksoy et al. (1991), ranking the variables in terms of the evidence that a shift in the mean has occurred.

As a graphical method, the polyplot approach of Blazek, Novic, and Scott (1987) is particularly helpful, although we recommend that the T^2 statistic used in the polyplot be replaced by a more effective MEWMA statistic.

7. CONCLUSION

Our ARL comparisons indicate that the MEWMA procedure can perform better than the multivariate CUSUM procedures of Crosier (1988) and Pignatiello and Runger (1990) when the process is initially out-of-control and it performs roughly the same if the shift in the mean vector is delayed. Inertia problems can delay reaction to a shift when using multivariate CUSUM or MEWMA charts, so Hotelling's χ^2 limits should always be used in conjunction with these charts to help to prevent such delays.

We have given some guidelines on the construction of MEWMA charts and on the interpretation of out-of-control signals. Additional work in both of these areas is now in progress.

ACKNOWLEDGMENTS

The research for this article was partially supported by a grant from the ALCOA Foundation. We are grateful to the editor, the referees, and the associate editors for making many helpful comments and suggestions on earlier versions of this article. We also thank Brian Jackson, a graduate student at Texas Christian University, for his assistance with the computer simulations.

APPENDIX A: DERIVATION OF THE COVARIANCE MATRIX FOR Z_i

By repeated substitution in Equation (2.2), it can be shown that

$$\mathbf{Z}_i = \sum_{j=1}^{i} R(I - R)^{i-j}\mathbf{X}_j.$$

Thus

$$\Sigma_{\mathbf{Z}_i} = \sum_{j=1}^{i} \text{var}[R(I - R)^{i-j}\mathbf{X}_j]$$

$$= \sum_{j=1}^{i} [R(I - R)^{i-j}\,\Sigma\,(I - R)^{i-j}R].$$

Because R and $(I - R)$ are diagonal matrices, the (k, ℓ)th element of $\Sigma_{\mathbf{Z}_i}$ is

$$r_k r_\ell[1 - (1 - r_k)^i(1 - r_\ell)^i]/[r_k + r_\ell - r_k r_\ell]\sigma_{k,\ell}, \quad \text{(A.1)}$$

where $\sigma_{k,\ell}$ is the (k, ℓ)th element of Σ. If $r_1 = r_2 = \cdots = r_p = r$, then the expression in (A.1) simplifies to $\{r[1 - (1 - r)^{2i}]/(2 - r)\}\sigma_{k,\ell}$ so that

$$\Sigma_{\mathbf{Z}_i} = \{r[1 - (1 - r)^{2i}]/(2 - r)\}\,\Sigma. \quad \text{(A.2)}$$

The covariance matrix is derived here under the assumption that the control rule in (2.3) is ignored, but it offers some guidance on the type of control rule to be used.

APPENDIX B: PROOF REGARDING ARL PERFORMANCE

This appendix outlines a proof that the ARL performance of the MEWMA procedure depends on $\boldsymbol{\mu}$ and Σ only through the value of the noncentrality parameter $\lambda = (\boldsymbol{\mu}'\Sigma^{-1}\boldsymbol{\mu})^{1/2}$. The results of Crosier (1988) can be applied in a straightforward way to establish these results provided we show that the values of the MEWMA statistic in (2.3) are invariant to any full-rank transformation of the data—that is, to $\mathbf{X}_i^* = M\mathbf{X}_i$, $i = 1, 2, \ldots$, where M is a $p \times p$ full-rank matrix.

We have $\mathbf{Z}_i = \Sigma_{j=1}^{i} r(1 - r)^{i-j}\mathbf{X}_j$. Thus if we define $\mathbf{Z}_i^* = \Sigma_{j=1}^{i} r(1 - r)^{i-j}M\mathbf{X}_j$, then $\mathbf{Z}_i^* = M\mathbf{Z}_i$. It follows easily that $T_i^{*2} = \mathbf{Z}_i^{*\prime}\,\Sigma_{\mathbf{Z}_i^*}^{-1}\,\mathbf{Z}_i^* = \mathbf{Z}_i'\,\Sigma_{\mathbf{Z}_i}^{-1}\,\mathbf{Z}_i = T_i^2$, $i = 1, 2, \ldots$.

The results of Crosier (1988) can now be applied. Those interested in more details regarding this proof are referred to Lowry (1989).

[*Received April 1989. Revised September 1991.*]

REFERENCES

Alt, F. B. (1984), "Multivariate Quality Control," in *The Encyclopedia of Statistical Sciences*, eds. S. Kotz, N. L. Johnson, and C. R. Read, New York: John Wiley, pp. 110–122.

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

Blazek, L. W., Novic, B., and Scott, D. M. (1987), "Displaying Multivariate Data Using Polyplots," *Journal of Quality Technology*, 19, 69–74.

Chua, M., and Montgomery, D. C. (1991), "A Multivariate Quality Control Scheme," *Quality and Reliability Engineering International*.

Crosier, R. B. (1988), "Multivariate Generalizations of Cumulative Sum Quality Control Schemes," *Technometrics*, 30, 291–303.

Crowder, S. V. (1989), "Design of Exponentially Weighted Moving Average Schemes," *Journal of Quality Technology*, 21, 135–162.

Doganaksoy, N., Faltin, F. W., and Tucker, W. T. (1991), "Identification of Out-of-Control Quality Characteristics in a Multivariate Manufacturing Environment," *Communications in Statistics—Theory and Methods*, 20, 2775–2790.

Enns, P. G., Machak, J. A., Spivey, W. A., and Wrobleski, W. J. (1982), "Forecasting Applications of an Adaptive Multiple Exponential Smoothing Model," *Management Science*, 28, 1035–1044.

Harvey, A. C. (1986), "Analysis and Generalisation of a Multivariate Exponential Smoothing Model," *Management Science*, 32, 374–380.

Hawkins, D. M. (1991), "Multivariate Quality Control Based on Regression-Adjusted Variables," *Technometrics*, 33, 61–75.

Healy, J. D. (1987), "A Note on Multivariate CUSUM Procedures," *Technometrics*, 29, 409–412.

Hotelling, H. (1947), "Multivariate Quality Control—Illustrated by the Air Testing of Sample Bombsights," in *Techniques of Statistical Analysis*, eds. C. Eisenhart, M. W. Hastay, and W. A. Wallis, New York: McGraw-Hill, pp. 111–184.

Jackson, J. E. (1985), "Multivariate Quality Control," *Communications in Statistics—Theory and Methods*, 14, 2657–2688.

Jones, R. H. (1966), "Exponential Smoothing for Multivariate Time Series," *Journal of the Royal Statistical Society*, Ser. B, 28, 241–251.

Lowry, C. A. (1989), "A Multivariate Exponentially Weighted Moving Average Control Chart," unpublished Ph.D. dissertation, University of Southwestern Louisiana, Dept. of Statistics.

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

MacGregor, J. F., and Harris, T. J. (1990), Discussion of "Exponentially Weighted Moving Average Control Schemes: Properties and Enhancements," by J. M. Lucas and M. S. Saccucci, *Technometrics*, 32, 23–26.

Mohebbi, C., and Hayre, L. (1989), "Multivariate Control Charts: A Loss Function Approach," *Sequential Analysis*, 8, 253–268.

Murphy, B. J. (1987), "Selecting Out-of-Control Variables with the T^2 Multivariate Quality Control Procedures," *The Statistician*, 36, 571–583.

Pignatiello, J. J., Jr., and Runger, G. C. (1990), "Comparisons of Multivariate CUSUM Charts," *Journal of Quality Technology*, 22, 173–186.

Roberts, S. W. (1959), "Control Chart Tests Based on Geometrics Moving Averages," *Technometrics*, 1, 239–250.

Tsui, K., and Woodall, W. H. (1991), "Multivariate Control Charts Based on Loss Functions," unpublished manuscript.

Woodall, W. H., and Ncube, M. M. (1985), "Multivariate CUSUM Quality Control Procedures," *Technometrics*, 27, 285–292.

Yashchin, E. (1987), "Some Aspects of the Theory of Statistical Control Schemes," *IBM Journal of Research and Development*, 31, 199–205.