

Multivariate Statistical Process Control Using LASSO

Changliang Zou¹ and Peihua Qiu²

¹*Department of Statistics, Nankai University, China*

²*School of Statistics, University of Minnesota, USA*

Abstract

This paper develops a new multivariate statistical process control (SPC) methodology based on adapting the LASSO variable selection method to the SPC problem. The LASSO method has the *sparsity* property that it can select exactly the set of nonzero regression coefficients in multivariate regression modeling, which is especially useful in cases when the number of nonzero coefficients is small. In multivariate SPC applications, process mean vectors often shift in a small number of components. Our major goal is to detect such a shift as soon as it occurs and identify the shifted mean components. Using this connection between the two problems, a LASSO-based multivariate test statistic is proposed, which is then integrated into the multivariate EWMA charting scheme for Phase II multivariate process monitoring. It is shown that this approach balances protection against various shift levels and shift directions, and hence provides an effective tool for multivariate SPC applications.

Keywords: Hotelling's T^2 statistic; Model selection; Multivariate statistical process control; Penalized likelihood; Regression-adjusted variables.

1 Introduction

In *multivariate statistical process control* (MSPC), one monitors several quality characteristics of a process. The fundamental tasks of MSPC are to determine whether a multivariate process mean $\boldsymbol{\mu}$ has changed; to identify when a detected shift in $\boldsymbol{\mu}$ occurred; and to isolate the shifted components of $\boldsymbol{\mu}$. Methods for accomplishing these tasks are usually derived under the assumption that the observed measurement vectors $\mathbf{X}_i = (X_{1i}, \dots, X_{pi})'$ are distributed $N_p(\boldsymbol{\mu}_0, \boldsymbol{\Sigma})$ for $i = 1, 2, \dots, \tau$, and $N_p(\boldsymbol{\mu}_1, \boldsymbol{\Sigma})$ for $i = \tau+1, \dots, n$, with $\boldsymbol{\mu}_0$ and $\boldsymbol{\Sigma}$ known. Throughout we take $\boldsymbol{\mu}_0 = \mathbf{0}$ without loss of generality. Then a portmanteau test for detecting a mean shift occurring at τ is based on testing $H_0 : \boldsymbol{\mu} = \mathbf{0}$ versus $H_1 : \boldsymbol{\mu} \neq \mathbf{0}$, using the likelihood ratio test (LRT) statistic $n\bar{\mathbf{X}}'\boldsymbol{\Sigma}^{-1}\bar{\mathbf{X}}$ where $\bar{\mathbf{X}} = \sum_{i=1}^n \mathbf{X}_i/n$. This procedure assumes that the covariance matrix does not change. Replacing $\boldsymbol{\Sigma}$ with the sample covariance matrix \mathbf{S} results in the Hotelling's T^2 statistic. Based on such test statistics, several MSPC control charts have been proposed, in the framework of cumulative sum (CUSUM) or exponentially weighted moving average (EWMA). Most charting statistics take quadratic forms of the related test statistics; for instance, Healy (1987), Croisier (1988), Pignatiello and Runger (1990), Hawkins (1991, 1993), Lowry et al. (1992), Runger and Prabhu (1996), Zamba and Hawkins (2006), and Hawkins and Maboudou-Tchao (2007). After a control chart signals a mean shift, a separate diagnostic procedure is often used for identifying the shift time and the component(s) of $\boldsymbol{\mu}$ that shifted. Commonly used diagnostic procedures include the ones based on decomposition of T^2 (e.g., Mason et al. 1995, 1997; Li et al. 2008) and various step-down procedures (e.g., Hawkins 1993; Mason and Young 2002; Sullivan et al. 2007).

MSPC control charts with quadratic charting statistics are powerful when one is interested in detecting shifts that occur in majority components of $\boldsymbol{\mu}$. In practice,

however, shifts often occur in only a few of the mean components. In such cases, more powerful control charts are possible. To be specific, if we know that the potential shift is in one specific direction \mathbf{d} , then the corresponding hypothesis testing problem becomes $H_0 : \boldsymbol{\mu} = \mathbf{0}$ versus $H'_1 : \boldsymbol{\mu} = \delta \mathbf{d}$, where δ is an unknown constant. The corresponding LRT statistic is $n(\mathbf{d}'\boldsymbol{\Sigma}^{-1}\bar{\mathbf{X}})^2/\mathbf{d}'\boldsymbol{\Sigma}^{-1}\mathbf{d} \stackrel{H_0}{\sim} \chi_1^2$. In the case when only one measurement component shifts but the component index is unknown, the testing problem becomes

$$H_0 : \boldsymbol{\mu} = \mathbf{0} \quad \text{versus} \quad H''_1 : \boldsymbol{\mu} = \delta \mathbf{d}_1 \quad \text{or} \quad \boldsymbol{\mu} = \delta \mathbf{d}_2 \quad \text{or} \quad \dots \quad \text{or} \quad \boldsymbol{\mu} = \delta \mathbf{d}_p,$$

where \mathbf{d}_j is the p -dimensional vector with 1 at the j -th element and 0 elsewhere. In this case, likelihood ratio arguments would lead to the following p test statistics:

$$V_j = \sqrt{n}(\mathbf{d}'_j\boldsymbol{\Sigma}^{-1}\bar{\mathbf{X}})/(\mathbf{d}'_j\boldsymbol{\Sigma}^{-1}\mathbf{d}_j)^{\frac{1}{2}}, \quad \text{for } j = 1, \dots, p. \quad (1)$$

It can be checked that V_j s are just the regression-adjusted variables defined in Hawkins (1991). Hawkins (1991; 1993) suggested a MSPC control chart using $\max_{j=1, \dots, p} |V_j|$ as a charting statistic, and showed that this chart was more effective than the one based on T^2 when the potential shift occurs in only one measurement component. Furthermore, after a shift is signaled, $\arg \max_j |V_j|$ can be used as a diagnostic tool to specify the shifted component. On the other hand, as noted by Hawkins (1991), the chart based on $\mathbf{V} = (V_1, \dots, V_p)$ can perform poorly, as when the shift occurs in several highly correlated components and when the shift is proportional to the two leading principal components of the covariance matrix. These findings are consistent with the properties of LRT with misspecified alternatives. Moreover, when the shift occurs in more than one component, diagnosis based on \mathbf{V} could be misleading, because it cannot indicate multiple shift components correctly.

Based on the above discussion, it is natural to consider more directions as *a priori* alternatives and to construct corresponding test statistics. For instance, we

can assume that two or more measurement components may shift; the resulting control charts would alleviate certain problems mentioned above. However, such control charts are infeasible when the dimension p is large, since the total number of possible shift directions would increase exponentially with p . Furthermore, their shift-detection power would decline substantially when p increases, due to the well-known problem of multiple tests. Regarding post-detection diagnosis, conventional approaches (e.g., the decomposition of T^2 procedure and the step-down test) are theoretically sound, but are inefficient when p is large. For instance, the decomposition of T^2 procedure considers $p!$ different decompositions of T^2 , which is computationally expensive in such cases. Certain parameters in these approaches (e.g., the threshold values) affect their diagnostic ability significantly, but are generally difficult to determine. See Sullivan et al. (2007) for related discussion.

In this paper, we propose a MSPC chart that has three features: (1) It is capable of detecting shifts in one or more components efficiently; (2) its computation is reasonably simple; and (3) it provides an effective post-signal diagnostic tool for identifying shifted measurement components. By adapting certain recent results on variable selection in multiple regression to the SPC problem, we propose a novel MSPC approach that has all these three properties. This approach integrates a *least absolute shrinkage and selection operator* (LASSO) test statistic (cf., Tibshirani 1996) into a multivariate EWMA charting scheme for on-line process monitoring, which is easy to implement using the *least angle regression* (LARS) algorithm (cf., Efron et al. 2004). Because of certain good properties of LASSO, it can determine the shift direction “automatically,” and can offer balanced protection against various possible shifts. Moreover, due to the sparsity property of LASSO estimators that some of their components would be exactly zero after certain tuning parameters are chosen properly, the proposed control chart can provide an effective diagnostic tool as well.

We describe in detail the proposed MSPC control chart in Section 2, and investigate its numerical performance in Section 3. In Section 4, we demonstrate the proposed method using a real-data example from manufacturing. Several remarks conclude the article in Section 5. Technical details are included in an appendix, available online as supplementary materials.

2 Methodology

Our proposed MSPC methodology is described in four parts. Subsection 2.1 provides a brief introduction to the LASSO in the context of variable selection in multiple regression. We demonstrate its connection to the MSPC problem and propose a LASSO-based test statistic in Subsection 2.2. In Subsection 2.3, a Phase II control chart is constructed by integrating the LASSO-based test statistic into the EWMA charting scheme. Finally, its diagnostic property is discussed in Subsection 2.4.

2.1 LASSO for regression variable selection

Consider the following multiple linear regression model:

$$y_i = \mathbf{Z}_i\boldsymbol{\beta} + \varepsilon_i, \quad \text{for } i = 1, 2, \dots, n,$$

where y_i , \mathbf{Z}_i , and $\boldsymbol{\beta}$ are the response variable, vector of predictors, and vector of regression coefficients, respectively, and ε_i are i.i.d. random errors with distribution $N(0, \sigma^2)$. In practice, some predictors do not provide much useful information about the response, given the remaining predictors; thus, they can be deleted from the model. To this end, stepwise or all subset selection procedures along with a model selection criterion (e.g., AIC or BIC) are usually used. Such model selection procedures are practically useful; but they have several limitations, including lack of

stability (cf., Breiman 1996), lack of theoretical properties (cf., Fan and Li 2001), and extensive computation (cf., Tibshirani 1996). To overcome these limitations, some authors suggest using the following penalized least squares method (or equivalently, the penalized likelihood method in the normal error distribution case):

$$\sum_{i=1}^n (y_i - \mathbf{z}_i \boldsymbol{\beta})^2 + n \sum_{j=1}^p g_{\gamma_j} (|\boldsymbol{\beta}^{(j)}|),$$

where $\boldsymbol{\beta}^{(j)}$ denotes the j -th component of $\boldsymbol{\beta}$, γ_j are the penalty parameters (or called the regularization parameters, see Bickel and Li 2006), and g_{γ_j} are the penalty functions. When $g_{\gamma_j}(|\boldsymbol{\beta}^{(j)}|) = \gamma |\boldsymbol{\beta}^{(j)}|$ where γ is a constant parameter, the corresponding penalized least squares method is called LASSO (cf., Tibshirani 1996). Besides LASSO, another major penalized least squares method is the so-called *smoothly clipped absolute deviation* method (SCAD, Fan and Li 2001). See Fan and Li (2006), Wang and Leng (2007), and Zou and Li (2008) for related discussion. Among other good properties, Fan and Li (2001) have shown that, with the proper choice of the penalty functions and regularization parameters, the penalized likelihood estimators would perform asymptotically as well as if the correct submodel was known, which is referred to as the oracle property in the literature. When the penalty functions are chosen to be continuous, the coefficient estimates that correspond to insignificant predictors would shrink towards 0 as γ increases, and these coefficient estimates could be exactly 0 if γ is sufficiently large. In addition, the penalized likelihood approach, especially those with LASSO-type penalty functions, enjoys efficient computation using Efron et al.'s (2004) LARS algorithm. See Zou and Li (2008) for detailed discussion about computational issues of LASSO and SCAD.

2.2 LASSO-based testing

In the variable selection problem discussed in the previous subsection, an *a priori* assumption is that some components of the coefficient vector $\boldsymbol{\beta}$ are zero. In high-dimensional cases, it is often reasonable to assume that only a few coefficients are non-zero, which is the so-called sparsity characteristic. Analogously, in the process monitoring problem, we assume in MSPC that only a few components in the shift vector $\boldsymbol{\delta} = \boldsymbol{\mu}_1$ are expected to be non-zero when a shift occurs. Therefore, the penalized likelihood method should have a potential in solving the MSPC problem, which is investigated below.

As described in Section 1, we are interested in testing $H_0 : \boldsymbol{\mu} = \mathbf{0}$ versus $H_1 : \boldsymbol{\mu} \neq \mathbf{0}$ in the MSPC testing problem. After a constant term is ignored, the corresponding penalized likelihood function can be written as

$$PL(\boldsymbol{\mu}) = n(\bar{\mathbf{X}} - \boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1} (\bar{\mathbf{X}} - \boldsymbol{\mu}) + n \sum_{j=1}^p g_{\gamma_j} (|\boldsymbol{\mu}^{(j)}|).$$

If the adaptive LASSO (ALASSO) penalty function (cf., Zou 2006) is used, then it becomes

$$PL(\boldsymbol{\mu}) = n(\bar{\mathbf{X}} - \boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1} (\bar{\mathbf{X}} - \boldsymbol{\mu}) + n\gamma \sum_{j=1}^p \frac{1}{|\bar{\mathbf{X}}^{(j)}|^a} |\boldsymbol{\mu}^{(j)}|, \quad (2)$$

where $a > 0$ is a pre-specified constant and $\boldsymbol{\mu}^{(j)}$ is the j -th component of $\boldsymbol{\mu}$. As recommended by Zou (2006) and Wang and Leng (2007), we fix $a = 1$ in this paper. The ALASSO penalty used in (2) is slightly different from the traditional one by Tibshirani (1996). The latter uses the same amount of shrinkage for each regression coefficient. Because of that, its coefficient estimator cannot be as efficient as the oracle estimator (Fan and Li 2001), and its model selection results could be inconsistent in certain cases. As a comparison, the ALASSO applies different amount of shrinkage to different regression coefficients. Consequently, its LASSO estimator is asymptotically

unbiased and has certain oracle properties. See Fan and Li (2001), Zou (2006), and Zhao and Yu (2006) for related discussion. The ALASSO estimator of $\boldsymbol{\mu}$ is defined by

$$\hat{\boldsymbol{\mu}}_\gamma = \arg \min PL(\boldsymbol{\mu}).$$

It can be checked that $\hat{\boldsymbol{\mu}}_\gamma$ is the same as the one-step sparse estimator in Zou and Li (2008) (see also Bühlmann and Meier 2008), and it is also the same as the least squares approximation estimator in Wang and Leng (2007). Because of its sparsity assumption (i.e., some of its components are exactly zero after γ is properly chosen), it is an ideal estimator of the shift direction in the MSPC problem. Then, a LASSO-based test statistic can be defined as

$$\tilde{T}_\gamma = \frac{n(\hat{\boldsymbol{\mu}}'_\gamma \boldsymbol{\Sigma}^{-1} \bar{\mathbf{X}})^2}{\hat{\boldsymbol{\mu}}'_\gamma \boldsymbol{\Sigma}^{-1} \hat{\boldsymbol{\mu}}_\gamma}.$$

Theoretically, $\hat{\boldsymbol{\mu}}_\gamma \neq \mathbf{0}$ almost surely. For completeness, \tilde{T}_γ can be defined to be any negative number when $\hat{\boldsymbol{\mu}}_\gamma = \mathbf{0}$. Obviously, \tilde{T}_γ can be regarded as a data-driven version of the test statistics in (1), with the shift direction estimator $\hat{\boldsymbol{\mu}}_\gamma$ substituting for the pre-specified shift directions \mathbf{d}_j .

Before \tilde{T}_γ can be used for MSPC, the regularization parameter γ should be chosen properly, since it plays an important role in balancing robustness and sensitivity of \tilde{T}_γ to various shifts. By Theorem 1 in Wang and Leng (2007) and Theorem 5 in Zou and Li (2008), we can obtain results in the following proposition without much difficulty, which provide us some insights about selection of γ . Without loss of generality, hereafter we assume that, under H_1 , the first p_0 components of $\boldsymbol{\mu}$ are nonzero and the remaining $p - p_0$ components are all zero, where $0 < p_0 \leq p$.

Proposition 1 *Assume that there are two regularization parameters γ_{1n} and γ_{2n} satisfying the conditions that $1/(n\gamma_{jn}) = o(1)$, for $j = 1, 2$, and $\gamma_{1n}/\gamma_{2n} = o(1)$. Then,*

(i) under H_0 , $P(\tilde{T}_{\gamma_{2n}} \geq 0)/P(\tilde{T}_{\gamma_{1n}} \geq 0) = o(1)$, and

(ii) under H_1 , if $\boldsymbol{\mu}$ satisfies $\min\{\boldsymbol{\mu}^{(j)}, j \leq p_0\} = O(n^s)$ and $\max\{\boldsymbol{\mu}^{(j)}, j \leq p_0\} = O(n^t)$, where $-\frac{1}{2} < s \leq t < 0$, then $P(\tilde{T}_{\gamma_{jn}} \geq 0) - 1 = o(1)$ if and only if $n^{-2t}\gamma_{jn} = o(1)$, for $j = 1$ or 2 . Furthermore, when $n^{-s+\frac{1}{2}}\gamma_{2n} = o(1)$, $\tilde{T}_{\gamma_{1n}}$ and $\tilde{T}_{\gamma_{2n}}$ converge to infinity with the same rate.

By this proposition, from an asymptotic viewpoint, selection of γ should depend on the shift size. For detecting a large shift (i.e., t is large in part (ii) of the proposition), γ should be chosen large, and small for detecting a small shift. One can choose γ is to use CV, GCV, AIC, BIC, or other model selection criteria. However, our simulation results show that these criteria tailored for estimation often do not produce a powerful test. This finding is not surprising because similar conclusions have been made in the nonparametric regression testing problem. There, the power and size of a typical test would depend on the bandwidth used in regression function estimation (cf., Bickel and Li 2006), and it is recognized that optimal bandwidth for nonparametric curve estimation may not be optimal for testing (cf., Hart 1997).

To overcome this difficulty, we follow the approach proposed by Horowitz and Spokoiny (2001) in the nonparametric testing setup, by combining several values of γ to make the resulting test nearly optimal. Let $\Gamma_q = \{\gamma_j, j = 1, \dots, q\}$ be a set of penalty parameters, where q is a pre-specified constant. Then, the modified penalized test statistic is defined as

$$\tilde{T} = \max_{j=1, \dots, q} \frac{\tilde{T}_{\gamma_j} - \mathbb{E}(\tilde{T}_{\gamma_j})}{\sqrt{\text{Var}(\tilde{T}_{\gamma_j})}},$$

where $\mathbb{E}(\tilde{T}_{\gamma_j})$ and $\text{Var}(\tilde{T}_{\gamma_j})$ are respectively the mean and variance of \tilde{T}_{γ_j} under H_0 . Ideally, q should be chosen to be the number of non-zero components of $\boldsymbol{\mu}$, and γ_j to be the penalty parameter that is good for detecting $\boldsymbol{\mu}$ with j non-zero components.

In practice, however, $\boldsymbol{\mu}$ is unknown. To overcome this difficulty, we propose a simple and efficient method to determine Γ_q , as described below.

Let us first rewrite the ALASSO-type penalized likelihood (2) as

$$PL(\boldsymbol{\alpha}) = n(\bar{\mathbf{X}} - \boldsymbol{\Lambda}\boldsymbol{\alpha})'\boldsymbol{\Sigma}^{-1}(\bar{\mathbf{X}} - \boldsymbol{\Lambda}\boldsymbol{\alpha}) + n\gamma \sum_{i=1}^p |\boldsymbol{\alpha}^{(i)}|, \quad (3)$$

where $\boldsymbol{\alpha}^{(i)} = \boldsymbol{\mu}^{(i)}/|\bar{\mathbf{X}}^{(i)}|$ and $\boldsymbol{\Lambda} = \text{diag}(|\bar{\mathbf{X}}^{(1)}|, \dots, |\bar{\mathbf{X}}^{(p)}|)$. This is exactly a LASSO-type penalized likelihood function. According to Zou et al. (2007), for given $\bar{\mathbf{X}}$ in (3), there is a finite sequence

$$\tilde{\gamma}_0 > \tilde{\gamma}_1 > \dots > \tilde{\gamma}_K = 0 \quad (4)$$

such that (1) for all $\gamma > \tilde{\gamma}_0$, $\hat{\boldsymbol{\alpha}}_\gamma = 0$, and (2) in the interval $(\tilde{\gamma}_{m+1}, \tilde{\gamma}_m)$, the active set $\mathfrak{B}(\gamma) = \{j : \text{sgn}[\boldsymbol{\alpha}_\gamma^{(j)}] \neq 0\}$ and the sign vector $\mathbf{S}(\gamma) = \{\text{sgn}[\boldsymbol{\alpha}_\gamma^{(1)}], \dots, \text{sgn}[\boldsymbol{\alpha}_\gamma^{(p)}]\}$ are unchanged with γ . These $\tilde{\gamma}_m$'s are called transition points because the active set changes at each $\tilde{\gamma}_m$. By Efron et al. (2004), the random integer K can be larger than p . Thus, we suggest using $\tilde{\gamma}_{m_j^{\text{last}}}$, for $j = 1, \dots, q$, to construct Γ_q , where m_j^{last} is the index of last $\tilde{\gamma}$ in the sequence $\{\tilde{\gamma}_0, \tilde{\gamma}_1, \dots, \tilde{\gamma}_K\}$ defined in (4) that the corresponding active set contains exactly j elements. By Lemma 7 in Zou et al. (2007), $\tilde{T}_{\tilde{\gamma}_{m_j^{\text{last}}}}$ are well defined, because the ‘‘one at a time’’ condition (cf., Efron et al. 2004) holds almost everywhere, where ‘‘one at a time’’ means that two active sets of two consecutive $\tilde{\gamma}$ s differ on at most a single index. The following asymptotic result provides some insight about the use of the transition points for constructing Γ_q .

Proposition 2 *Under H_1 , assume that $\min\{\boldsymbol{\mu}^{(j)}, j \leq p_0\} = O(n^s)$, where $-\frac{1}{2} < s < 0$. Then, the test using $\tilde{T}_{\tilde{\gamma}_{m_{p_0}^{\text{last}}}}$ as a testing statistic would be asymptotically more powerful than the test using T^2 .*

In practice, p_0 is often unknown. So, we suggest combining all $\tilde{\gamma}_{m_j^{\text{last}}}$, for $j =$

$1, \dots, q$, in our test. The resulting test statistic is

$$\tilde{T}_L = \max_{j=1, \dots, q} \frac{\tilde{T}_{\tilde{\gamma}_j^{\text{last}}} - \text{E}(\tilde{T}_{\tilde{\gamma}_j^{\text{last}}})}{\sqrt{\text{Var}(\tilde{T}_{\tilde{\gamma}_j^{\text{last}})}}. \quad (5)$$

Note that $\tilde{T}_{\tilde{\gamma}_j^{\text{last}}}$ is exactly the T^2 test statistic. Therefore, test (5) should share certain properties of the T^2 test and certain properties of the likelihood ratio tests constructed when the potential shift has some specific structures (cf., expression (1)).

Next, we make several remarks about the transition points used in (5). First, $\tilde{\gamma}_m$ can be easily determined by the LARS algorithm (cf., Efron et al. 2004). The whole testing procedure (5) involves $O(np + p^3)$ computations, which are about the same as the computational cost of a least squares regression with p covariates. This property is especially desirable for on-line process monitoring. Second, each $\tilde{\gamma}_m$ corresponds to a different active set $\mathfrak{B}_{\tilde{\gamma}_m}$. Thus, by combining different $\tilde{\gamma}_m$'s, test (5) actually combines some cases with different numbers of shifted components, and these cases are determined by the observed data. Third, based on Lemma 1 in Appendix, for any $\gamma \in (\tilde{\gamma}_{m_j^{\text{last}}}, \tilde{\gamma}_{m_j^{\text{last}}-1})$, we have $l(\hat{\boldsymbol{\mu}}_\gamma) > l(\hat{\boldsymbol{\mu}}_{\tilde{\gamma}_{m_j^{\text{last}}}})$, where $l(\boldsymbol{\mu}) = n(\bar{\mathbf{X}} - \boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1} (\bar{\mathbf{X}} - \boldsymbol{\mu})$ and $\hat{\boldsymbol{\mu}}_\gamma = \boldsymbol{\Lambda} \hat{\boldsymbol{\alpha}}_\gamma$. That is, in the interval $[\tilde{\gamma}_{m_j^{\text{last}}}, \tilde{\gamma}_{m_j^{\text{last}}-1})$, $\hat{\boldsymbol{\mu}}_{\tilde{\gamma}_{m_j^{\text{last}}}}$ is the ‘‘optimal’’ estimator in terms of the likelihood value in the case when j components are shifted, which provides a strong evidence that the transition points should be used in our test.

2.3 Phase II MSPC

In this part, we propose a framework to construct Phase II MSPC control charts using the LASSO-based test statistic (5). We consider EWMA-type control charts only in this paper; but, other types of control charts (e.g., CUSUM) can be developed in a similar way.

Let \mathbf{X}_j be the j -th observed measurement vector collected over time. Then, a

multivariate EWMA (MEWMA) sequence of statistics (cf., Lowry et al. 1992) is often defined as

$$\mathbf{U}_j = \lambda \mathbf{X}_j + (1 - \lambda) \mathbf{U}_{j-1}, \quad \text{for } j = 1, 2, \dots, \quad (6)$$

where $\mathbf{U}_0 = \mathbf{0}$ and λ is a weighting parameter in $(0, 1]$. For each \mathbf{U}_j , we suggest computing q LASSO estimators $\hat{\boldsymbol{\mu}}_{j, \tilde{\gamma}_{m_k^{\text{last}}}}$, for $k = 1, 2, \dots, q$, from the penalized likelihood function

$$(\mathbf{U}_j - \boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1} (\mathbf{U}_j - \boldsymbol{\mu}) + \gamma \sum_{k=1}^p \frac{1}{|\mathbf{U}_j^{(k)}|} |\boldsymbol{\mu}^{(k)}|. \quad (7)$$

Our proposed control chart signals a shift if

$$Q_j = \max_{k=1, \dots, q} \frac{W_{j, \tilde{\gamma}_{m_k^{\text{last}}}} - \text{E}(W_{j, \tilde{\gamma}_{m_k^{\text{last}}}})}{\sqrt{\text{Var}(W_{j, \tilde{\gamma}_{m_k^{\text{last}}}})}} > L, \quad (8)$$

where

$$W_{j, \gamma} = \frac{2 - \lambda}{\lambda [1 - (1 - \lambda)^{2j}]} \frac{(\mathbf{U}_j' \boldsymbol{\Sigma}^{-1} \hat{\boldsymbol{\mu}}_\gamma)^2}{\hat{\boldsymbol{\mu}}_\gamma' \boldsymbol{\Sigma}^{-1} \hat{\boldsymbol{\mu}}_\gamma},$$

and $L > 0$ is a control limit chosen to achieve a given in-control (IC) average run length (ARL). In practice, $(2 - \lambda) / \{\lambda [1 - (1 - \lambda)^{2j}]\}$ can be replaced by its asymptotic form $(2 - \lambda) / \lambda$, as conventionally done in the EWMA literature. Hereafter, the above control chart is called LASSO-based multivariate EWMA (LEWMA) chart.

In LEWMA (6)–(8), selection of λ mainly depends on a target shift (cf. Proposition 1). It should be chosen large if the target shift is large and small otherwise, as in conventional EWMA charts (e.g., Lucas and Saccucci 1990; Prabhu and Runger 1997). Regarding q , recall that an implicit *a priori* assumption underlying the LEWMA chart is that a potential shift occurs only in a small number of measurement components. If prior information indicating potential shifts in at most r components, with $1 \leq r \leq p$, then our numerical studies show that using $q = r + 1$ or $q = r + 2$ provides satisfactory performance in practice (cf., Section 3). When prior information is unavailable,

numerical results (Section 3) show that the LEWMA chart with $q = p$ performs reasonably well in all cases considered there.

In (8), quantities $E(W_{j, \tilde{\gamma}_{m_k^{\text{last}}}})$ and $\text{Var}(W_{j, \tilde{\gamma}_{m_k^{\text{last}}}})$ are usually unknown. The following result suggests replacing them by their approximations.

Proposition 3 *When the process is IC, $E(W_{j, \tilde{\gamma}_{m_k^{\text{last}}}})$ and $\text{Var}(W_{j, \tilde{\gamma}_{m_k^{\text{last}}}})$ do not depend on λ and j .*

By Proposition 3, $E(W_{j, \tilde{\gamma}_{m_k^{\text{last}}}})$ and $\text{Var}(W_{j, \tilde{\gamma}_{m_k^{\text{last}}}})$ can be approximated by the empirical expectation and variance of $W_{j, \tilde{\gamma}_{m_k^{\text{last}}}}$ computed from simulated IC measurement vectors, since \mathbf{U}_j would be the same as \mathbf{X}_j when $\lambda = 1$. The control limit L in (8) can also be determined by simulation as follows. First, choose an initial value for L . Then, compute the IC ARL of the LEWMA chart based on a large number of replicated simulation runs in which the IC observations are generated from the IC distribution of the process. If the computed IC ARL value is smaller than the nominal IC ARL value, then increase the value of L . Otherwise, choose a smaller L value. The above process is repeated until the nominal IC ARL is reached and is within a desired precision. In this process, some numerical searching algorithms, such as the bisection search, can be applied (cf., e.g., Qiu 2008). For given λ , Σ and a desired IC ARL, computation involved in finding L is not difficult, partly due to the fact that the LARS algorithm used in LASSO computation is efficient. For instance, when IC ARL=200 and $p = 15$, it requires about 15 minutes to complete the bisection searching procedure based on 25,000 simulations, using a Pentium-M 1.6MHz CPU. Since this is a one-time computation before the Phase II online process monitoring, it is convenient to accomplish. Fortran code for implementing the proposed procedure is available from the authors upon request.

2.4 Post-signal diagnostic

Assume that the LEWMA chart (6)–(8) triggers a mean shift signal at the k -th observation. Estimates of the shift locations and identification of the specific measurement components that have shifted will help engineers to eliminate the root causes of the shift in a timely fashion. With respect to the shift location, the generalized maximum likelihood estimation approach for change-point detection can be used (cf., e.g., Zamba and Hawkins 2006). After an estimate of the shift location is obtained, which is denoted as $\hat{\tau}$, we have $(k - \hat{\tau})$ out-of-control (OC) observations in which a few measurement components have shifts in their means. Among these $(k - \hat{\tau})$ observations, a few might actually be IC observations, because $\hat{\tau}$ is only an estimator of the true shift location τ . After a mean shift is detected and the shift location is estimated, the LASSO methodology can be used for specifying the shifted components, by choosing one of the LASSO estimators (cf., (4)) with a model selection criterion (e.g., C_p , GCV, AIC, or BIC). To be specific, our proposed diagnostic algorithm tries to find $\hat{\boldsymbol{\mu}}_{\gamma^*}$

$$\gamma^* = \arg \min_{\gamma} (k - \hat{\tau})(\bar{\mathbf{X}}_{\hat{\tau},k} - \hat{\boldsymbol{\mu}}_{\gamma})' \boldsymbol{\Sigma}^{-1} (\bar{\mathbf{X}}_{\hat{\tau},k} - \hat{\boldsymbol{\mu}}_{\gamma}) + \eta \cdot \widehat{\text{df}}(\hat{\boldsymbol{\mu}}_{\gamma}), \quad (9)$$

$\bar{\mathbf{X}}_{\hat{\tau},k} = \sum_{i=\hat{\tau}+1}^k \mathbf{X}_i / (k - \hat{\tau})$, $\hat{\boldsymbol{\mu}}_{\gamma}$ is the ALASSO estimator from $\bar{\mathbf{X}}_{\hat{\tau},k}$, η is a parameter, and $\widehat{\text{df}}(\hat{\boldsymbol{\mu}}_{\gamma})$ is the number of nonzero coefficients in $\hat{\boldsymbol{\mu}}_{\gamma}$. If AIC is used, then $\eta = 2$. It equals $\ln(k - \hat{\tau})$ if BIC is used. By Theorem 3 in Zou et al. (2007), $\hat{\boldsymbol{\mu}}_{\gamma^*}$ is one of $\hat{\boldsymbol{\mu}}_{\hat{\tau}_1}, \dots, \hat{\boldsymbol{\mu}}_{\hat{\tau}_K}$.

In the literature, it is well demonstrated that BIC tends to identify the true sparse model well if the true model is included in the candidate set (cf., e.g., Yang 2005; Zou et al. 2007; Wang et al. 2007). However, unlike the fixed sample case, in the current SPC problem, we usually do not have a large sample (relative to the shift magnitude) to implement the diagnostic procedure, because the whole point of SPC

is to detect the shift (and hence stop the process) as quickly as possible. Based on extensive simulations, we find that the conventional BIC criterion does not perform well in certain cases (not reported here). Instead, the diagnostic procedure (9) would perform well if it uses the risk inflation criterion (RIC) proposed by George and Foster (1994) based on the minimax principle, in which $\eta = 2 \ln(p)$. Since some components of $\hat{\boldsymbol{\mu}}_{\gamma^*}$ would be exactly zero, we can simply take its nonzero components as the shifted measurement components, without making any extra tests that are necessary in most existing diagnostic methods, such as the decomposition of T^2 method and the step-down method. As a side note, in cases when estimation of the shift location τ is not our major concern, after the LEWMA chart gives an OC signal at the k -th observation, we can use \mathbf{U}_k directly to implement the diagnostic procedure. Namely, we can find $\hat{\boldsymbol{\mu}}_{\gamma^*}$ so that

$$\gamma^* = \arg \min_{\gamma} \frac{2 - \lambda}{\lambda[1 - (1 - \lambda)^{2k}]} (\mathbf{U}_k - \hat{\boldsymbol{\mu}}_{\gamma})' \boldsymbol{\Sigma}^{-1} (\mathbf{U}_k - \hat{\boldsymbol{\mu}}_{\gamma}) + \eta \cdot \widehat{\text{df}}(\hat{\boldsymbol{\mu}}_{\gamma}), \quad (10)$$

where as before, $\eta = 2 \ln(p)$. Compared to (9), procedure (10) is more convenient to use and is faster as well because the LASSO estimators from \mathbf{U}_k have been computed in the monitoring process. Our numerical studies show that (9) and (10) provide similar diagnostic results.

3 A Simulation Study

We present some simulation results in this section regarding the numerical performance of the proposed LEWMA control chart and the corresponding diagnostic procedure. In the LEWMA chart, the parameter q needs to be chosen, although $q = p$ is always an option. We first investigate the potential effect of q on the performance of the LEWMA chart in the following setting. Let $p = 15$ and the IC measurement distribution be $N_p(\mathbf{0}, \boldsymbol{\Sigma})$, where $\boldsymbol{\Sigma} = (\sigma_{ij})$ is chosen to be $\sigma_{ii} = 1$ and $\sigma_{ij} = 0.75^{|i-j|}$,

for $i, j = 1, 2, \dots, p$. In the LEWMA chart, λ is chosen to be 0.2, and its IC ARL is fixed at 500. Two representative shifts are considered. The first one has three shifted components $\delta_2 = 0.25$, $\delta_3 = 0.25$ and $\delta_8 = 0.50$, where δ_j denotes the shift size of the j -th component, and the second one has all eight odd components shifted of the same size 0.25. The two plots in Figure 1 show the OC ARLs of the LEWMA chart when it is used for detecting the two shifts and when q changes. The plots show that the LEWMA chart performs the best when $q = 5$ and $q = 10$, respectively, and its performance is close to optimal when $q = p$. Therefore, $q = p$ is indeed a reasonable choice when no prior information is available about the number of shifted components.




Figure 1: OC ARLs of the LEWMA chart when q changes, $p = 15$ and two representative shifts are detected. Plot (a): the shift has three shifted components $\delta_2 = 0.25$, $\delta_3 = 0.25$ and $\delta_8 = 0.50$. Plot (b): all eight odd components are shifted of 0.25.

Next, we compare the LEWMA chart with the conventional MEWMA chart (i.e., the multivariate EWMA chart using a quadratic charting statistic) and REWMA chart (i.e., the multivariate EWMA chart using \mathbf{V}). We first consider a low-dimensional case when $p = 5$ that was discussed in Hawkins (1991). In all three charts, $\lambda = 0.2$ and the IC ARL is fixed at 500. The true shift location τ is fixed at 25, and any

simulations in which shift signals occur before τ are discarded. In the LEWMA chart, to investigate its sensitivity to selection of q , we consider two different q values: $q = 3$ and $q = p = 5$. Twenty different shifts are considered, which include the following cases: shifts in a single component, equal shifts in a pair of components, and shifts in the principle component directions of \mathbf{X} . As was done in Hawkins (1991), all shifts are scaled so that the largest displacement in \mathbf{V} equals one. The simulation results over 20,000 replications are presented in Table 1. Control limits L of the three control charts are included in the last row of this table. In order to assess the overall performance of the charts, besides OC ARLs, we also compute their relative mean index (RMI) values. The RMI index of a control chart is suggested by Han and Tsung (2006), and is defined as

$$\text{RMI} = \frac{1}{N} \sum_{l=1}^N \frac{\text{ARL}_{\delta_l} - \text{MARL}_{\delta_l}}{\text{MARL}_{\delta_l}}, \quad (11)$$

where N is the total number of shifts considered, ARL_{δ_l} is the OC ARL of the given control chart when detecting shift δ_l , and MARL_{δ_l} is the smallest OC ARL among all OC ARL values of the charts considered when detecting shift δ_l . So, $\frac{\text{ARL}_{\delta_l} - \text{MARL}_{\delta_l}}{\text{MARL}_{\delta_l}}$ can be considered as a relative efficiency measure of the given control chart, compared to the best chart, when detecting shift δ_l , and RMI is the average of all such relative efficiency values. By this index, a control chart with a smaller RMI value is considered better in its overall performance. From Table 1, we observe the following results. First, LEWMA offers a good balance between sensitivity and robustness to potential shifts. This can be seen from cases when REWMA outperforms MEWMA in quite large margins or vice versa. In such cases, LEWMA may not be best, but it is always close to the best. Second, in several cases in which LEWMA performs the best, the margins are generally quite small. Third, in terms of the RMI index, LEWMA performs the best overall. Fourth, the OC ARL values of the LEWMA chart when $q = 3$ are close to those when $q = 5$, although the LEWMA chart with $q = 3$

performs slightly better when the number of shifted components is small and slightly worse when the number of shifted components is large. Fifth, the last shift considered in the table is large in magnitude and MEWMA performs the best as expected. The LEWMA chart with $q = p$ also performs well in this case. In Table 1, λ is fixed at 0.2 for all three charts. To diminish the possible effect of λ on the performance of the three charts, we also compute the optimal OC ARLs with respect to λ (i.e., the smallest OC ARLs when λ changes) for all three charts in the cases considered in Table 1. The results are presented in Table S1 given in the online supplementary materials. They are slightly better than those in Table 1, as expected; but, the conclusions made from Table 1 are also true in Table S1.

Table 1: OC ARL values of control charts MEWMA, REWMA, and LEWMA in the case when $p = 5$, $\lambda = 0.2$, and IC ARL=500. Numbers in parentheses are standard errors. They are shown as “.00” when they are smaller than 0.01.

shifts					MEWMA	REWMA	LEWMA	
X_1	X_2	X_3	X_4	X_5			$q = 3$	$q = 5$
0.91	0.00	0.00	0.00	0.00	17.3 (.09)	14.4 (.07)	14.6 (.07)	14.9 (.07)
0.00	0.36	0.00	0.00	0.00	17.0 (.08)	13.3 (.06)	13.9 (.06)	14.3 (.07)
0.00	0.00	0.48	0.00	0.00	17.3 (.09)	13.8 (.07)	14.6 (.07)	15.0 (.07)
0.00	0.00	0.00	0.34	0.00	17.2 (.09)	13.4 (.06)	14.2 (.07)	14.6 (.07)
0.00	0.00	0.00	0.00	0.46	17.9 (.09)	14.1 (.07)	14.9 (.07)	15.2 (.07)
0.36	0.36	0.00	0.00	0.00	15.0 (.07)	13.5 (.06)	13.4 (.06)	13.7 (.06)
0.54	0.00	0.54	0.00	0.00	12.8 (.06)	12.9 (.06)	12.4 (.05)	12.4 (.05)
0.32	0.00	0.00	0.32	0.00	15.2 (.07)	13.1 (.06)	13.4 (.06)	13.6 (.06)
0.49	0.00	0.00	0.00	0.49	13.2 (.06)	12.6 (.06)	12.4 (.06)	12.5 (.06)
0.00	0.54	0.54	0.00	0.00	8.79 (.03)	12.3 (.06)	8.75 (.03)	8.88 (.03)
0.00	1.60	0.00	1.60	0.00	3.48 (.00)	9.03 (.03)	3.59 (.00)	3.57 (.00)
0.00	0.28	0.00	0.00	0.28	13.0 (.06)	10.5 (.05)	11.1 (.05)	11.3 (.05)
0.00	0.00	0.28	0.28	0.00	13.0 (.06)	10.5 (.05)	11.3 (.05)	11.4 (.05)
0.00	0.00	1.26	0.00	1.26	4.28 (.01)	8.70 (.03)	4.41 (.01)	4.34 (.01)
0.00	0.00	0.00	0.56	0.56	8.60 (.03)	12.2 (.06)	8.60 (.03)	8.70 (.03)
0.01	-0.15	0.07	0.17	-0.09	15.0 (.07)	11.5 (.05)	12.3 (.06)	12.5 (.06)
0.07	-0.13	-0.40	0.19	0.35	10.1 (.04)	12.4 (.06)	10.0 (.04)	10.1 (.04)
0.40	0.63	-0.57	0.47	-0.68	4.74 (.01)	9.28 (.04)	4.97 (.01)	4.94 (.01)
-1.11	0.26	-0.17	0.34	-0.04	11.8 (.05)	14.6 (.07)	12.5 (.05)	12.2 (.05)
2.51	7.11	7.05	7.11	7.08	1.19 (.00)	7.52 (.02)	2.88 (.00)	1.30 (.00)
RMI					0.123	0.514	0.103	0.046
Control Limit					18.13	3.425	5.181	5.262

Next, we consider the case when $p = 15$ and all IC parameters are chosen to be the same as those in the example of Figure 1. In this example, besides the OC ARL, the “worst case” OC ARL is also evaluated. The “worst case” scenario is that, immediately before a shift occurs, the control charting statistic is at the value that would maximize the OC ARL (cf., e.g., Yashchin 1993; Woodall and Mahmoud 2005). The “worst case” OC ARL reflects the worst case performance of a given control chart. The simulation results based on 10,000 replications are presented in Table 2, where 27 shifts with various numbers of shifted components are considered. In this example, we simply choose $q = p$ in the LEWMA chart. The results show that the LEWMA chart is either the best or near the best in all cases, in terms of both OC ARL and “worst case” OC ARL. Even in the last two cases with large shifts in all fifteen components, the LEWMA chart still performs reasonably well. RMI values in the table confirm the claimed advantage of the LEWMA chart that it does offer a balanced protection against various shifts.

In all previous examples, shifts are pre-specified. For a more objective comparison of different control charts, next we consider an example in which r indices are randomly generated from $\{1, \dots, p\}$ without replacement as the indices of the shifted components, and r shift magnitudes are randomly generated from $N(\delta, \sigma^2)$. A sequence of multivariate normal vectors are generated, with the shift vector specified above. We generate control charts to this sequence and obtain their corresponding run-lengths. This simulation is repeated N times, and we compute the RMI value for each control chart using (11), except that ARLs in that formula are replaced by those run-lengths. The resulting performance measure is referred to as relative run length index (RRLI). By properly choosing δ , σ^2 and a large N , this index would be a good measure of the relative run length performance of the considered control charts. Figure 2 shows the RRLI values (in log-scale) of the LEWMA, REWMA and

Table 2: OC ARL and “worst case” OC ARL values of control charts MEWMA, REWMA, and LEWMA in the case when $p = 15$, $\lambda = 0.2$, and IC ARL=500. All numbers in parentheses are standard errors. For a given shift, components not specified are all 0.

shifts	OC ARL			“worst case” OC ARL		
	MEWMA	REWMA	LEWMA	MEWMA	REWMA	LEWMA
$\delta_1 = 0.50$	62.5 (.58)	39.8 (.35)	40.8 (.35)	66.3 (.58)	43.1 (.35)	44.7 (.36)
$\delta_1 = 1.00$	11.2 (.06)	7.84 (.04)	8.11 (.04)	14.9 (.06)	10.6 (.04)	10.9 (.04)
$\delta_3 = 0.50$	34.1 (.29)	21.5 (.16)	22.5 (.17)	38.6 (.29)	24.9 (.16)	26.4 (.20)
$\delta_3 = 1.00$	7.26 (.03)	5.41 (.02)	5.62 (.02)	10.6 (.03)	7.72 (.02)	8.02 (.02)
$\delta_1 = 0.50; \delta_2 = 0.25$	106 (1.01)	138 (1.30)	109 (1.00)	110 (1.02)	143 (1.34)	114 (1.03)
$\delta_1 = 0.50; \delta_2 = 0.50$	57.3 (.52)	91.0 (.88)	57.7 (.51)	61.5 (.53)	93.4 (.89)	62.3 (.50)
$\delta_1 = 0.50; \delta_2 = 0.75$	21.2 (.16)	19.7 (.14)	17.8 (.12)	25.5 (.16)	23.2 (.14)	21.3 (.12)
$\delta_1 = 0.50; \delta_3 = 0.25$	39.3 (.34)	29.9 (.24)	30.2 (.23)	43.8 (.34)	33.1 (.25)	33.9 (.25)
$\delta_1 = 0.50; \delta_3 = 0.50$	18.0 (.13)	14.9 (.10)	14.8 (.09)	22.2 (.13)	18.3 (.10)	18.3 (.10)
$\delta_1 = 0.50; \delta_3 = 0.75$	9.78 (.05)	7.92 (.04)	8.02 (.04)	13.4 (.05)	10.8 (.04)	10.9 (.04)
$\delta_3 = 0.50; \delta_8 = 0.25$	25.5 (.19)	20.3 (.15)	20.0 (.14)	29.9 (.20)	23.7 (.15)	23.4 (.14)
$\delta_3 = 0.50; \delta_8 = 0.50$	14.4 (.09)	13.8 (.09)	12.7 (.07)	18.5 (.09)	17.1 (.09)	16.1 (.08)
$\delta_3 = 0.50; \delta_8 = 0.75$	8.78 (.04)	8.04 (.04)	7.70 (.03)	12.3 (.04)	10.9 (.04)	10.5 (.04)
$\delta_1 = 0.50; \delta_2 = 0.25; \delta_3 = 0.25$	103 (.99)	127 (1.23)	104 (1.01)	106 (.99)	130 (.24)	107 (1.02)
$\delta_1 = 0.25; \delta_2 = 0.25; \delta_3 = 0.50$	55.9 (.52)	44.0 (.39)	44.1 (.37)	60.5 (.53)	47.3 (.39)	47.5 (.39)
$\delta_2 = 0.50; \delta_3 = 0.25; \delta_8 = 0.25$	33.5 (.28)	34.5 (.29)	30.2 (.24)	37.7 (.28)	38.2 (.30)	34.1 (.24)
$\delta_2 = 0.25; \delta_3 = 0.25; \delta_8 = 0.50$	25.4 (.20)	21.1 (.16)	20.6 (.14)	29.8 (.20)	24.6 (.16)	24.0 (.15)
$\delta_7 = 0.50; \delta_8 = 0.25; \delta_9 = 0.50$	24.3 (.18)	25.4 (.20)	22.4 (.16)	28.5 (.19)	28.5 (.20)	26.3 (.17)
$\delta_7 = 0.25; \delta_8 = 0.75; \delta_9 = 0.50$	23.1 (.18)	60.4 (.54)	26.5 (.20)	27.2 (.18)	61.1 (.55)	28.9 (.23)
$\delta_6 = 0.50; \delta_8 = 0.25; \delta_{10} = 0.50$	20.6 (.15)	18.1 (.13)	17.1 (.11)	21.6 (.16)	20.2 (.16)	19.0 (.15)
$\delta_6 = 0.25; \delta_8 = 0.75; \delta_{10} = 0.50$	6.96 (.03)	6.95 (.03)	6.49 (.03)	10.3 (.04)	9.86 (.03)	9.55 (.03)
Even components shift 0.25	15.9 (.11)	23.8 (.18)	17.2 (.11)	20.9 (.11)	27.6 (.18)	21.9 (.12)
Even components shift 0.50	4.60 (.02)	6.50 (.03)	4.90 (.02)	7.50 (.02)	9.22 (.04)	7.78 (.02)
Odd components shift 0.25	17.1 (.11)	24.4 (.20)	17.9 (.12)	20.2 (.10)	26.8 (.17)	21.3 (.11)
Odd components shift 0.50	4.75 (.02)	6.64 (.03)	5.03 (.02)	7.33 (.02)	9.63 (.04)	7.67 (.04)
Even (odd) components shift 0.50 (0.25)	13.7 (.09)	24.3 (.19)	16.8 (.11)	17.7 (.09)	27.4 (.18)	21.0 (.11)
Even (odd) components shift 0.25 (0.50)	12.2 (.07)	22.5 (.17)	15.1 (.10)	16.4 (.09)	25.8 (.16)	18.8 (.10)
RMI	0.164	0.251	0.040	0.165	0.180	0.030
Control limit	34.75	3.749	4.950	34.75	3.749	4.950

MEWMA charts when r changes, $p = 15$ (plot (a)), and $p = 30$ (plot (b)). In this example, we choose $\delta = 0.25$, $\sigma = 0.5$, $N=50,000$ in plot (a), $N = 100,000$ in plot (b), and $q = p$ in the LEWMA chart. From the plots, we can see that (i) for small r (e.g., $r \leq 2$), the LEWMA and REWMA charts have similar performance and both significantly outperform the MEWMA chart, (ii) for moderate r (e.g., $2 < r \leq p/2$), the LEWMA chart performs the best, and (iii) for large r (e.g., $r > p/2$), the MEWMA chart performs slightly better than the LEWMA and both of them are much better than the REWMA chart. These findings confirm our statement made earlier that the

LEWMA chart is effective when the sparsity condition holds and it offers a protection against various OC conditions.

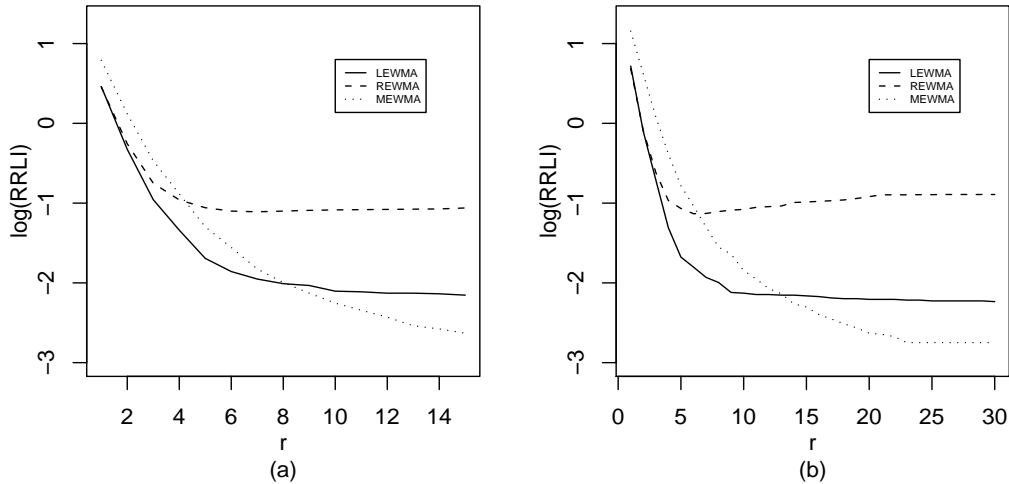


Figure 2: RRLI comparison (in log-scale) of the LEWMA, REWMA and MEWMA charts when $p = 15$ (plot (a)) and $p = 30$ (plot (b)). In the plots, r denotes the number of shifted components.

We now compare the proposed diagnostic procedure (9) with the existing step-down procedure. The step-down procedure depends heavily on the pre-specified type I error probability, which is set to be 0.10 here for illustration. The following three schemes are considered: the MEWMA chart with the step-down procedure (denoted as “step-down”), the MEWMA chart with procedure (9) (denoted as “MEWMA-(9)”), and the LEWMA chart with procedure (9) (denoted as “LEWMA-(9)”). Simulation results in the settings of Tables 1 and 2 are presented in Table S2 (given in the online supplemental materials) and Table 3, respectively. In the two tables, the columns labelled “C” present relative frequencies that the diagnostic procedures identify shifted measurement components correctly, and the columns labelled “I” denote relative frequencies that at least one shifted component is missed and at least one identified shifted component is false. So, for a given diagnostic procedure, it performs better in a given case if its value in column “C” is larger and its value in column “I” is smaller. The results show that the proposed LASSO-based approach has comparable

diagnostic ability to that of the step-down procedure. In many situations, especially when the dimension p is large (cf., Table 3), the LASSO-based approach outperforms the step-down procedure. For instance, in the first row of Table 3, the relative frequency of the MEWMA-(9) scheme is $\hat{\pi} = 0.54$ ($SE = \sqrt{\hat{\pi}(1 - \hat{\pi})/10000} = 0.005$) for identifying shifted measurement components correctly. As a comparison, the corresponding relative frequency and its standard error of the step-down scheme are 0.07 and 0.003, respectively. So, the MEWMA-(9) scheme is significantly better in this case. After taking into account its computational advantage, we think that the LASSO-based approach provides a reasonable diagnosis tool for MSPC.

Table 3: Diagnostic results of the proposed procedure (9) and the step-down procedure in the case when $p = 15$, $\lambda = 0.2$, and IC ARL=500. In the table, columns labeled ‘‘C’’ present relative frequencies of identifying shifted measurement components correctly, and the columns labelled ‘‘I’’ denote relative frequencies that at least one shifted components is missed and at least one identified shifted component is false. For a given entry with relative frequency $\hat{\pi}$, its standard error can be computed by the formula $\sqrt{\hat{\pi}(1 - \hat{\pi})/10000}$.

shifts	step-down		MEWMA-(9)		LEWMA-(9)	
	C	I	C	I	C	I
$\delta_1 = 0.50$	0.07	0.08	0.54	0.11	0.55	0.13
$\delta_1 = 1.00$	0.07	0.04	0.57	0.08	0.61	0.06
$\delta_3 = 0.50$	0.05	0.09	0.53	0.12	0.55	0.12
$\delta_3 = 1.00$	0.07	0.06	0.58	0.08	0.61	0.07
$\delta_1 = 0.50; \delta_2 = 0.25$	0.07	0.22	0.34	0.24	0.32	0.28
$\delta_1 = 0.50; \delta_2 = 0.50$	0.07	0.13	0.40	0.18	0.37	0.22
$\delta_1 = 0.50; \delta_2 = 0.75$	0.07	0.23	0.19	0.26	0.22	0.27
$\delta_1 = 0.50; \delta_3 = 0.25$	0.06	0.28	0.22	0.37	0.27	0.34
$\delta_1 = 0.50; \delta_3 = 0.50$	0.06	0.36	0.24	0.34	0.29	0.32
$\delta_1 = 0.50; \delta_3 = 0.75$	0.06	0.36	0.16	0.35	0.20	0.32
$\delta_3 = 0.50; \delta_8 = 0.25$	0.05	0.44	0.15	0.35	0.20	0.34
$\delta_3 = 0.50; \delta_8 = 0.50$	0.06	0.34	0.27	0.29	0.30	0.30
$\delta_3 = 0.50; \delta_8 = 0.75$	0.07	0.35	0.23	0.29	0.27	0.29
$\delta_1 = 0.50; \delta_2 = 0.25; \delta_3 = 0.25$	0.08	0.32	0.19	0.30	0.18	0.34
$\delta_1 = 0.25; \delta_2 = 0.25; \delta_3 = 0.50$	0.06	0.48	0.04	0.36	0.06	0.36
$\delta_2 = 0.50; \delta_3 = 0.25; \delta_8 = 0.25$	0.06	0.55	0.11	0.42	0.12	0.43
$\delta_2 = 0.25; \delta_3 = 0.25; \delta_8 = 0.50$	0.05	0.59	0.05	0.43	0.07	0.42
$\delta_7 = 0.50; \delta_8 = 0.25; \delta_9 = 0.50$	0.05	0.51	0.05	0.42	0.06	0.44
$\delta_7 = 0.25; \delta_8 = 0.75; \delta_9 = 0.50$	0.09	0.30	0.28	0.25	0.24	0.30
$\delta_6 = 0.50; \delta_8 = 0.25; \delta_{10} = 0.50$	0.04	0.62	0.08	0.47	0.13	0.43
$\delta_6 = 0.25; \delta_8 = 0.75; \delta_{10} = 0.50$	0.07	0.55	0.12	0.43	0.09	0.47

In the previous examples, the IC and OC distributions are assumed multivariate normal and the IC parameters are all known. In the online supplementary materials, a numerical example is presented to study the performance of the three control charts LEWMA, MEWMA and REWMA when these assumptions are violated, which shows that the three control charts are all affected by violations of the assumptions and their sensitivity to the violations is comparable.

4 An Application

We illustrate the proposed method using a dataset from an aluminum smelter that produces metallic aluminum from dissolved alumina through a chemical reaction process. The dataset contains five variables—the contents of SiO_2 , Fe_2O_3 , MgO , CaO , and Al_2O_3 (labeled as X_1 , X_2 , X_3 , X_4 , and X_5) in the cryolite/alumina mixture that goes through the chemical reaction process. These measures are important to the quality of metallic aluminum produced. More detailed information about this data can be found in Qiu and Hawkins (2001) and Zamba and Hawkins (2006).

The dataset has 189 vectors. Figure 3 shows the time series of the raw data. Similar to Qiu and Hawkins (2001), we first pre-whiten the dataset by removing substantial autocorrelation in the original measurements. Then, since our method relies on the joint normality of the 5 variables, we transform X_1 and X_5 to the log scale, as was done in Zamba and Hawkins (2006). For the transformed data, Mardia’s multivariate normality test (Mardia 1970) gives the p-value of about 0.27, which suggests that the multivariate normality assumption is valid in this case. Some summary statistics of the transformed data are presented in Table S3 included in the online supplementary materials, which shows that the correlation matrix $\hat{\Sigma}$ contains several large entries. Therefore, multivariate control charts should be more appropriate to

use in this example than univariate control charts.




Figure 3: The raw aluminum smelter data.

We then apply the three control charts LEMMA, REWMA and MEWMA to this dataset. In the LEMMA chart, we set $\lambda = 0.2$, $q = 5$, and IC ARL=500. Its control limit is computed to be $L = 5.143$, using estimated IC parameters from the first 75 vectors of the data. Note that a calibration sample of this size is smaller than one would like to fully determine the IC measurement distribution. Based on the simulation results shown in Figure S1 in the online supplementary materials, the actual IC ARLs of the three procedures could be quite different from their nominal IC ARLs in such cases. So, this example is just for illustrating the use of the proposed method in a real-world setting. In practice, a larger calibration sample (say 500 vectors; cf., Figure S1) is desired. Figure 4 shows the resulting LEWMA chart (solid curve connecting the dots), along with its control limit (solid horizontal line). The corresponding REWMA chart (dashed curve connecting little circles), MEWMA chart (dotted curve connecting little diamonds), and their control limits (horizontal dashed and dotted lines, respectively) are also presented in the figure. All three charts signal at the 83rd observation and remain above their control limits in the remainder of the sequence, but the LEMMA chart gives a more convincing evidence of the shift because its charting statistic values after the 83rd observation are much higher than its control limit, compared to the other two charts. We also applied the anti-rank CUSUM chart proposed by Qiu and Hawkins (2001) to this dataset; with a reference value of 0.5, that chart does not give a signal until the 91st observation.

We think that the anti-rank chart is relatively inefficient here because it does not use the normality assumption that approximately is satisfied for the transformed data. Then, we use the proposed diagnostic procedure (9) to identify the shifted variables, after the change-point estimate of $\hat{\tau} = 81$ is obtained. Table 4 tabulates the resulting five LASSO estimates $\hat{\boldsymbol{\mu}}_{\tilde{\gamma}_j}$, for $j = 1, \dots, 5$, and the corresponding values of $D_j = 2(\bar{\mathbf{X}}_{81,83} - \hat{\boldsymbol{\mu}}_{\tilde{\gamma}_j})' \boldsymbol{\Sigma}^{-1}(\bar{\mathbf{X}}_{81,83} - \hat{\boldsymbol{\mu}}_{\tilde{\gamma}_j}) + 2 \ln(5) \cdot j$ (denoted as D-value in Table 4). These values indicate that the shift may have occurred in the last two variables CaO and Al₂O₃.

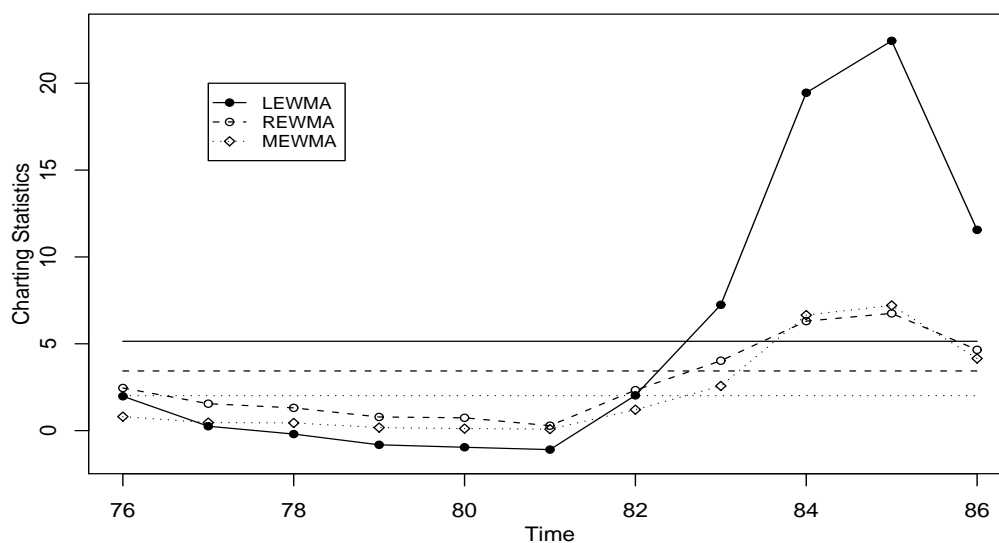


Figure 4: LEWMA, REWMA and MEWMA control charts for monitoring the aluminum smelting process. The solid, dashed and dotted horizontal lines indicate their control limits, respectively.

Table 4: Diagnostic results of procedure (9) about the aluminum smelter data.

j	$\hat{\boldsymbol{\mu}}_{\tilde{\gamma}_j}$				D-value	
1	0.000	0.000	0.000	0.000	-2.383	17.332
2	0.000	0.000	0.000	0.977	-2.970	13.914
3	0.248	0.000	0.000	1.134	-3.112	15.388
4	0.799	0.000	-0.423	1.747	-3.671	14.552
5	1.263	0.778	-1.184	1.984	-3.689	16.094

5 Concluding Remarks

In this paper, we have proposed a new framework for MSPC using LASSO estimators. A LASSO-based multivariate test statistic is proposed, and it is combined with a multivariate EWMA procedure for Phase II process monitoring. Because of certain good properties of the LASSO estimators, the proposed control chart has a balanced protection against various shifts. We also proposed a LASSO-based diagnostic procedure that is convenient to use. Numerical studies show that the proposed monitoring and diagnosis approaches would be effective in industrial applications.

The current version of the proposed LEWMA chart is designed for detecting mean shifts only, and it is based on the assumptions that the multivariate observations are independent of each other, they follow normal distributions, and their IC distribution is known. We believe that, after certain modifications to the LASSO penalized likelihood function (2), the proposed method should be able to handle cases in which monitoring both the mean vector and the covariance matrix is of interest (cf., e.g., Huwang et al. 2007). In such cases, we would face a much higher dimensional problem (i.e., $[p(p+3)/2]$ -dimensional case); we expect that the potential improvement by LEWMA, compared to its peers, would be even larger. In addition, the multivariate normality assumption could be readily relaxed to the elliptical distribution assumption (cf., Liu and Singh 1993). It requires much future research to (i) further extend the proposed approach to nonparametric multivariate SPC cases (cf., Liu 1995; Qiu and and Hawkins 2001, 2003) and (ii) accommodate possible correlation among observations observed at different time points. Future research includes a self-starting version of the LEWMA chart (cf., Hawkins and Maboudou-Tchao 2007) and a study of its properties in cases when the IC parameters in the measurement distribution are unknown (cf., Jones et al. 2001).

Acknowledgments The authors thank the editor, the associate editor, and two referees for many constructive comments and suggestions which greatly improved the quality of the paper. This research is supported in part by the grant DMS-0721204 from NSF of USA and the grant 10771107 from NNSF of China.

References:

- Bickel, P. J., and Li, B. (2006), “Regularization in Statistics,” *Test*, 15, 271–344.
- Breiman, L. (1996), “Heuristics of Instability and Stabilization in Model Selection,” *The Annals of Statistics*, 24, 2350–2383.
- Bühlmann, P., and Meier, L. (2008), “Discussion of ‘One-step Sparse Estimates in Nonconcave Penalized Likelihood Models’,” *The Annals of Statistics*, 36, 1534–1541.
- Croisier, R. B. (1988), “Multivariate Generalizations of Cumulative Sum Quality-Control Schemes,” *Technometrics*, 30, 243–251.
- Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004), “Least Angle Regression,” *The Annals of Statistics*, 32, 407–489.
- Fan, J., and Li, R. (2001), “Variable Selection via Nonconcave Penalized Likelihood and Its Oracle Properties,” *Journal of the American Statistical Association*, 96, 1348–1360.
- Fan, J., and Li, R. (2006), “Statistical Challenges with High Dimensionality: Feature Selection in Knowledge Discovery,” in *Proceedings of the Madrid International Congress of Mathematicians 2006*.
- George, E. I., and Foster, D. P. (1994), “The Risk Inflation Criterion for Multiple Regression,” *The Annals of Statistics*, 22, 1947–1975.
- Han D., and Tsung, F. (2006), “A Reference-Free Cuscore Chart for Dynamic Mean Change Detection and a Unified Framework for Charting Performance Comparison,” *Journal of the American Statistical Association*, 101, 368–386.
- Hart, J. D. (1997). *Nonparametric Smoothing and Lack-of-Fit Tests*, Springer, New York.
- Hawkins, D. M., (1991), “Multivariate Quality Control Based on Regression-Adjusted Variables,” *Technometrics*, 33, 61–75.
- Hawkins, D. M., (1993), “Regression Adjustment for Variables in Multivariate Quality Control,” *Journal of Quality Technology*, 25, 170–182.
- Hawkins, D. M., and Maboudou-Tchao, E. M. (2007), “Self-Starting Multivariate Exponentially Weighted Moving Average Control Charting,” *Technometrics*, 49, 199–209.
- Healy, J. D. (1987), “A Note on Multivariate CUSUM Procedure,” *Technometrics*, 29, 409–412.
- Horowitz, J. L., and Spokoiny, V. G. (2001), “An Adaptive, Rate-Optimal Test of a Parametric Mean-Regression Model Against a Nonparametric Alternative,” *Econometrica*, 69, 599–631.
- Huwang, L., Yeh, A. B., Wu, C. (2007), “Monitoring Multivariate Process Variability for Individual Observations,” *Journal of Quality Technology*, 39, 258–278.
- Jones, L. A., Champ, C. W., and Rigdon, S. E. (2001), “The Performance of Exponentially Weighted Moving Average Charts with Estimated Parameters,” *Technometrics*, 43, 156–167.
- Knight, K., and Fu, W. (2000), “Asymptotics for Lasso-Type Estimators,” *The Annals of Statistics*, 28, 1356–1378.
- Li, J., Jin, J., and Shi, J. (2008), “Causation-Based T^2 Decomposition for Multivariate Process Monitoring and Diagnosis,” *Journal of Quality Technology*, 40, 46–58.
- Liu, R. (1995), “Control Charts for Multivariate Processes,” *Journal of the American Statistical Association*, 90, 1380–1388.
- Liu, R., and Singh, K. (1993), “A Quality Index Based on Data Depth Tests,” *Journal of the American Statistical Association*, 88, 252–260.

- Lowry, C. A., Woodall, W. H., Champ, C. W., and Rigdon, S. E. (1992), "Multivariate Exponentially Weighted Moving Average Control Chart," *Technometrics*, 34, 46–53.
- Lucas, J. M., and Saccucci, M. S. (1990), "Exponentially Weighted Moving Average Control Scheme Properties and Enhancements," *Technometrics*, 32, 1–29.
- Mardia, K.V. (1970), "Measures of Multivariate Skewness and Kurtosis With Applications," *Biometrika*, 57, 519–530.
- Mason, R. L., Tracy, N. D., and Young, J. C. (1995), "Decomposition of T^2 for Multivariate Control Chart Interpretation," *Journal of Quality Technology*, 27, 99–108.
- Mason, R. L., Tracy, N. D., and Young, J. C. (1997), "A Practical Approach for Interpreting Multivariate T^2 Control Chart Signals," *Journal of Quality Technology*, 29, 396–406.
- Mason, R. L., and Young, J. C. (2002), *Multivariate Statistical Process Control With Industrial Application*, Philadelphia: SIAM.
- Osborne, M., Presnell, B., and Turlach, B. (2000), "A New Approach to Variable Selection in Least Squares Problems," *IMA Journal of Numerical Analysis*, 20, 389–403.
- Pignatiello, J. J., and Runger, G. C. (1990), "Comparison of Multivariate Cusum Charts," *Journal of Quality Technology*, 22, 173–186.
- Prabhu, S. S., and Runger, G. C. (1997), "Designing a Multivariate EWMA Control Chart," *Journal of Quality Technology*, 29, 8–15.
- Qiu, P. (2008), "Distribution-Free Multivariate Process Control Based on Log-Linear Modeling," *IIE Transactions*, 40, 664–677.
- Qiu, P., and Hawkins, D. M. (2001), "A Rank-Based Multivariate CUSUM Procedure," *Technometrics*, 43, 120–132.
- Qiu, P., and Hawkins, D. M. (2003), "A nonparametric multivariate CUSUM procedure for detecting shifts in all directions," *JRSS (Series D) - The Statistician*, 52, 151–164.
- Runger, G. C., and Prabhu, S. S. (1996), "A Markov Chain Model for the Multivariate Exponentially Weighted Moving Averages Control Chart," *Journal of the American Statistical Association*, 91, 1701–1706.
- Sullivan, J. H., Stoumbos, Z. G., Mason, R. L., and Young, J. C. (2007), "Step-Down Analysis for Changes in the Covariance Matrix and Other Parameters," *Journal of Quality Technology*, 39, 66–84.
- Tibshirani, R. J. (1996), "Regression Shrinkage and Selection via the LASSO," *Journal of the Royal Statistical Society: Series B*, 58, 267–288.
- Wang, H., and Leng, C. (2007), "Unified Lasso Estimation by Least Square Approximation," *Journal of the American Statistical Association*, 102, 1039–1048.
- Wang, H., Li, R., and Tsai, C. L. (2007), "On the Consistency of SCAD Tuning Parameter Selector," *Biometrika*, 94, 553–568.
- Woodall, W. H., and Mahmoud, M. A. (2005), "The Inertial Properties of Quality Control Charts," *Technometrics*, 47, 425–436.
- Yang, Y. (2005), "Can the Strengths of AIC and BIC Be Shared?—A Conflict between Model Identification and Regression Estimation," *Biometrika*, 92, 937–950.
- Yashchin, E. (1993), "Statistical Control Schemes: Methods, Applications and Generalizations," *International Statistical Review*, 61, 41–66.
- Zamba, K. D., and Hawkins, D. M. (2006), "A Multivariate Change-Point for Statistical Process Control," *Technometrics*, 48, 539–549.
- Zhao, P., and Yu, B. (2006), "On Model Selection Consistency of Lasso," *Journal of Machine Learning Research*, 7, 2541–2563.
- Zou, H. (2006), "The Adaptive Lasso and Its Oracle Properties," *Journal of the American Statistical Association*, 101, 1418–1429.
- Zou, H., Hastie, T., and Tibshirani, R. (2007), "On the 'Degrees of Freedom' of Lasso," *The Annals of Statistics*, 35, 2173–2192.
- Zou, H., and Li, R. (2008), "One-step Sparse Estimates in Nonconcave Penalized Likelihood Models," *The Annals of Statistics* (with discussion), 36, 1509–1533.