Quality Engineering, 17:481–489, 2005

ISSN: 0898-2112 print/1532-4222 online
DOI: 10.1081/QEN-200068575

Quality Quandaries*: The Effect of Autocorrelation on Statistical Process Control Procedures

Søren Bisgaard
Eugene M. Isenberg School of Management, University of Massachusetts Amherst and Institute for Business and Industrial Statistics, University of Amsterdam, The Netherlands

Murat Kulahci
Department of Industrial Engineering, Arizona State University

INTRODUCTION

The context in which statistical process control methods are employed has changed dramatically over the past few decades. Of course, the traditional applications remain important. But increasingly networked sensors—thermo couples, pressure gages, flow meters and even video cameras—are connected to powerful computers running sophisticated software to monitor complex processes. In this modern environment, individual observations arrive one by one at whatever sampling rate we choose, every second, minute, hour, or day. In this context, the classical Shewhart chart is not always so useful. Charts designed for sequential use, such as the individuals moving range chart, the cumulative sum (Cusum) chart, or the exponentially weighted moving average (EWMA) chart are often more appropriate.

In the application of each of these standard charts, it is assumed that the individual observations are statistically independent. However, the more frequently we sample a process, the more likely it is the observations are serially correlated. Now, some of the assumptions we make such as normality of the data are relatively unimportant, but independence is not. Serial correlation can have profound effects on the performance of a control chart. Indeed, positive autocorrelation, the most common form of serial correlation in industrial processes, may completely alter the performance of a control chart and must be dealt with appropriately. In this column, we provide an example of a serially correlated process and demonstrate what can be done to modify the standard charts to make them work.

*Edited by Søren Bisgaard

Address correspondence to Søren Bisgaard, Eugene M. Isenberg School of Management, University of Massachusetts Amherst.

EXAMPLE: TEMPERATURE CONTROL OF A CERAMIC FURNACE

In a recent consulting engagement with a large ceramics manufacturer, an engineer despairingly declared that the statistical process control methods he was taught as part of a Six Sigma program were useless. The methods, he explained, gave completely ridiculous answers. When applying a regular individual's control chart to a steady process exhibiting only a small and, from an engineering point of view, completely tolerable amount of variation, the process was pronounced out of control all the time.

We will consider 80 consecutive hourly temperature readings $z_t, t = 1, \ldots, 80$ from a thermo couple placed inside the large ceramics furnace. See Appendix for a listing of the data. This 80 hour time segment, shown as time series plot in Figure 1, was selected because it represented a stable period that the process engineers wanted to use as a baseline for good process performance.

The three sigma control limits for an individuals control chart are computed as

$$\text{Upper Control Limit (UCL)} = \bar{z} + 3\frac{\overline{MR}}{d_2} \quad (1)$$
$$\text{Centerline} = \bar{z}$$
$$\text{Lower Control Limit (LCL)} = \bar{z} - 3\frac{\overline{MR}}{d_2}$$

where $\bar{z} = n^{-1}\sum z_t$ and $\overline{MR}$ is the average of moving ranges, MR, typically computed from a moving window of m observations, see e.g. Montgomery (2001).

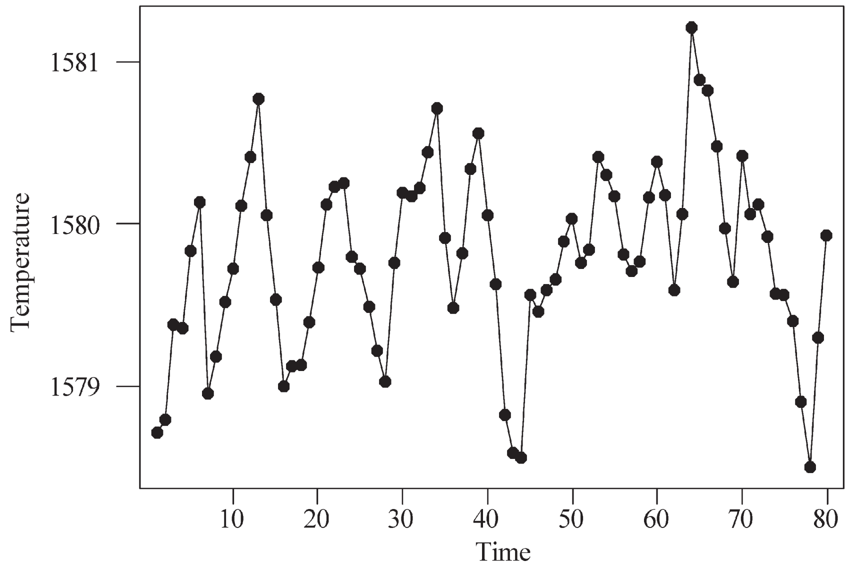

Figure 1. A times series plot of 80 consecutive hourly temperature observations from a ceramic furnace.

If $m = 2$, the default value in many software packages including MINITAB, the moving ranges are the absolute values of the differences between successive observations $MR_t = |z_t - z_{t-1}|$ and $d_2 = 1.128$. The reason for using a moving range with $m = 2$ to estimate the process variability is that it provides a good estimate of the short-term process variability, even if the process mean should be slowly drifting over time.

An individuals control chart based on a moving range of $m = 2$ for the 80 consecutive hourly temperature readings is shown in Figure 2. We notice that despite its stable appearance, the process is seriously out of control. As explained above, this was the reason the process control engineer concluded that statistical process control was useless for his process. It is also easy to see that his objections were sensible and should receive serious consideration. Indeed, his process did not change much more than plus and minus one degree over this three day period. Further, considering the large size of this furnace, any attempt to counteract such small changes whenever the control chart provided an out of control signal would likely cause more variability, not less. Deming called this tampering.

So, what is the problem? The problem, of course, is not with the individuals control chart per se, but with its application to this process. We will discuss this in more detail in the next section, but the observations in this case seriously violate the important assumption of statistical independence. Indeed, the keen observer will notice that Figure 1 shows a slow moving wavy pattern rather than being completely random. Given the physics of the circumstances, this is to be expected. With a sampling rate of one observation per hour of the temperature of something extremely large such as a furnace, the temperature at time t and that of one hour later at time $t + 1$, are obviously going to be related. Things don't change that fast!

Another chart that is popular for the control of industrial processes where the individual observations arrive one by one is the Exponentially Weighted Moving Average (EWMA) chart, see e.g. Box and Luceno (1997). The EWMA, $\tilde{y}_t$, is computed sequentially

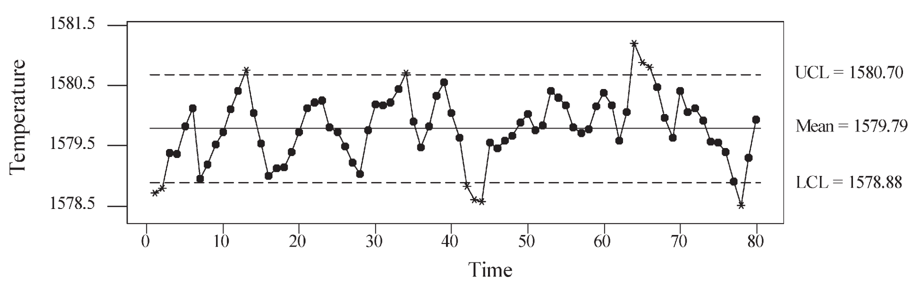

Figure 2. An individuals control chart of the 80 consecutive hourly temperature observations from a ceramic furnace.

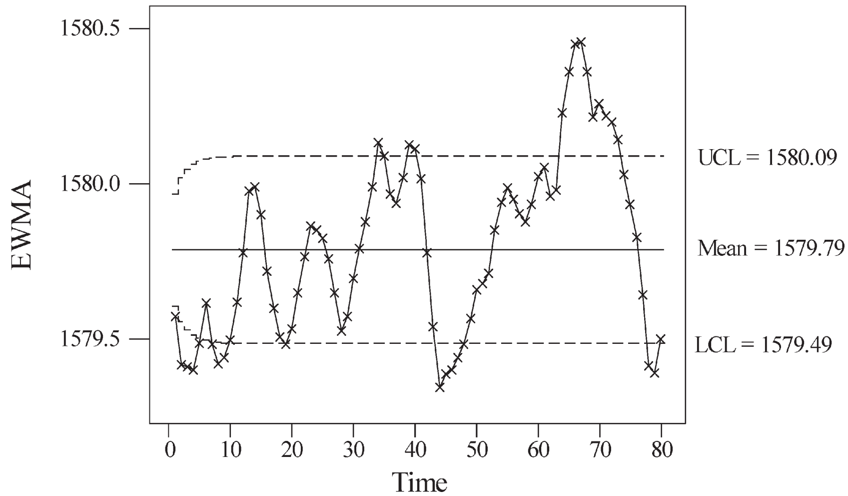

Figure 3. An EWMA of hourly temperature observations from the ceramic furnace.

as a linear interpolation between the present observation z_t and $\tilde{y}_{t-1}$, the previous EWMA

$$\tilde{y}_t = \lambda z_t + (1 - \lambda)\tilde{y}_{t-1}, \tag{2}$$

where λ is a constant $0 \leq \lambda \leq 1$. Hunter (1986) has shown that for independent and normally distributed data, the control limits for the EWMA $\tilde{y}_t$ are given by

$$\begin{aligned} \mathrm{UCL}_t &= \bar{z} + 3\hat{\sigma}\sqrt{\frac{\lambda}{2-\lambda}\left[1-(1-\lambda)^{2t}\right]} \\ \mathrm{CL} &= \bar{z} \\ \mathrm{LCL}_t &= \bar{z} + 3\hat{\sigma}\sqrt{\frac{\lambda}{2-\lambda}\left[1-(1-\lambda)^{2t}\right]}, \end{aligned} \tag{3}$$

where the estimate of the process variability, $\hat{\sigma}$, typically is estimated using the same method as for the individuals control chart.

Figure 3 shows an EWMA for the furnace data with the typical default value $\lambda = 0.2$. We see that the process also with this chart appears to be out of control. Again, we attribute this to the observations not being independent, rather than the process not be stationary and in control.

SERIAL CORRELATION

The basis for a process being in statistical control is that its joint probability distribution is *stationary*. It is not, as some mistakenly may think, that the observations are independent. For practical purposes, we interpret stationarity to imply that the first two moments, the mean, the variance, and the cross correlation between observations from different points in time, are constant. This is sometimes referred to as weak stationarity; see Box et al. (1994) for a precise definition. Although there are some rigorous statistical tests that can be performed, often, a visual inspection of the time series plot will provide important information as to whether the process is stationary or not. In our case, since the time series plot does not raise serious concerns regarding the stationarity of the data, we will proceed with our analysis. Once we fit the time series model, we will come back to the issue of stationarity and provide a test based on the parameter estimates of the proposed model.

For the present process we can readily see that the observations are correlated over time. Figure 4 shows a panel of four plots: (*i*) the temperatures z_t versus the lag one temperatures, z_{t-1}, (*ii*) the temperatures, z_t

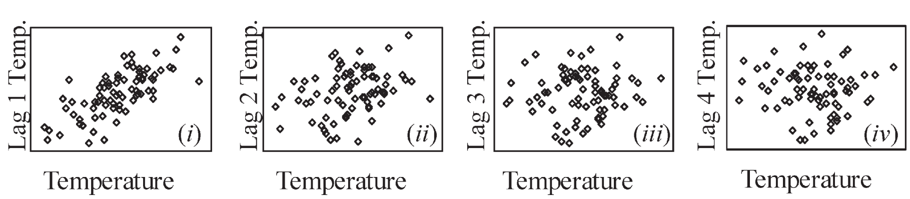

Figure 4. The correlation between observations (*i*) one time unit apart, (*ii*) two time units apart, (*iii*) three time units apart, and (*iv*) four time units apart.

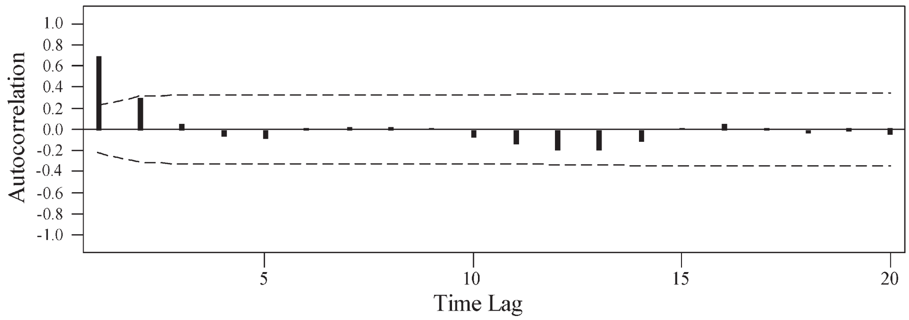

Figure 5. The autocorrelation function of the furnace temperature.

versus the lag 2 temperatures, z_{t-2}, (*iii*) the temperatures, z_t, versus the lag 3 temperatures, z_{t-3}, and (*iv*) the temperatures, z_t versus and the lag 4 temperatures, z_{t-4}. From these four plots, we see that observations one time unit apart are highly positively correlated, observations two time units apart are also correlated, but less so, and after three and four time lags, they are more or less uncorrelated. This kind of correlation is called autocorrelation. An alternative representation of the autocorrelation is provided in Figure 5. In this plot, only the correlation coefficients for different lags are presented. Specifically, the correlation shown as scatter plots in Figure 4 (*i*) between z_t and z_{t-1} is about 0.7, and the correlation between z_t and z_{t-2} shown in Figure 4 (*ii*) is about 0.3. These lag 1 and lag 2 autocorrelations are shown in Figure 5 as the two first bars of the length 0.7 and 0.3 respectively. Further, the two dotted horizontal lines in Figure 5 indicate confidence intervals around zero. Thus the lag 1 autocorrelation is clearly significantly different from zero, and the lag 2 autocorrelation is borderline significant.

A TIMESERIES MODEL OF THE FURNACE DATA

Stationary autocorrelated process data can often be modeled via an Autoregressive Moving Average (ARMA) time series model, see Box et al. (1994). The identification of the particular model within this general class of models is determined by looking at the Autocorrelation Function (ACF) and the Partial Autocorrelation Function (PACF).

To appreciate what the PACF is, consider a pth order autoregressive process AR(p). This model is like a regular regression equation except we regress the current observations z_t on the past p values, $z_{t-1}, \ldots, z_{t-p}$. That is, if $\tilde{z}_t = z_t - \mu$ is the current observation's deviation from the process mean, μ, then

$$\tilde{z}_t = \phi_1 \tilde{z}_{t-1} + \phi_2 \tilde{z}_{t-2} + \cdots + \phi_p \tilde{z}_{t-p} + a_t \qquad (4)$$

where a_t is assumed to be independent white noise errors, $a_t \sim N(0, \sigma_a^2)$. The kth order partial autocorrelation measures the additional correlation between $\tilde{z}_t$ and $\tilde{z}_{t-k}$, after adjustments have been made for the intermediate observations $\tilde{z}_{t-1}, \tilde{z}_{t-2}, \ldots, \tilde{z}_{t-k+1}$. In other words, the lag k partial autocorrelation can be thought of as the last regression coefficient ϕ_{kk} if we progressively for $k = 1, 2, \ldots$ fit regression equations

$$\tilde{z}_t = \phi_{k1} \tilde{z}_{t-1} + \cdots + \phi_{kk} \tilde{z}_{t-k} + a_t \qquad (5)$$

to the data. Thus, by the time we fit too many terms, the partial regression coefficients ϕ_{kk} will approximately be zero. For example, if we fit an AR(2) model to data that

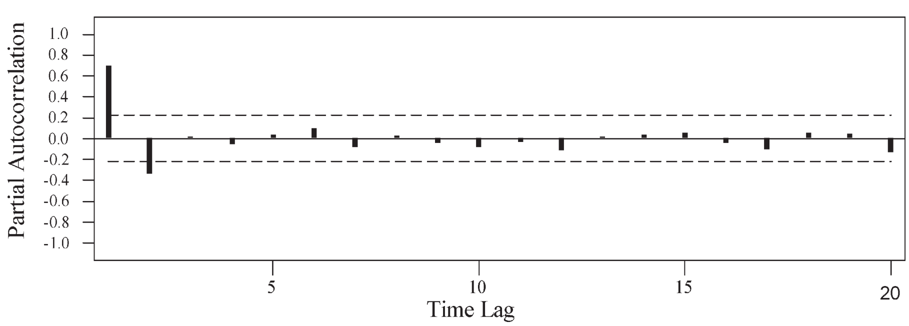

Figure 6. The partial autocorrelation function for the furnace data.

truly follows an AR(2) model, then the regression coefficients $\hat{\phi}_{kk}$ for lag $k = 3, 4, \ldots$ will be zero.

In the present case the ACF in Figure 5 looks like a damped sine function. The PACF in Figure 6 shows that the two first partial autocorrelation coefficients are larger than the two standard error limits, and hence deemed significant. After that, the PACF cuts off. To identify the particular type of ARMA model we use Table 3.3 of Box et al. (1994). From that table we see that a pattern of an exponentially decaying or sine wave decaying ACF and a PACF that cuts off after lag 2 suggest an AR (2) model. Thus, we tentatively fit the following AR (2) model

$$\tilde{z}_t = \phi_1 \tilde{z}_{t-1} + \phi_2 \tilde{z}_{t-2} + a_t \cdot \qquad (6).$$

This model can also be written directly in terms of the data z_t, which is how MINITAB parameterizes stationary AR models. Thus, if we use the substitution $\tilde{z}_t = z_t - \mu$ we get

$$\begin{aligned} z_t - \mu &= \phi_1(z_{t-1} - \mu) + \phi_2(z_{t-2} - \mu) + a_t \\ z_t &= \mu - \phi_1\mu - \phi_2\mu + \phi_1 z_{t-1} + \phi_2 z_{t-2} + a_t \qquad (7) \\ z_t &= \text{constant} + \phi_1 z_{t-1} + \phi_2 z_{t-2} + a_t, \end{aligned}$$

where the constant $= \mu - \phi_1\mu - \phi_2\mu$ or $\mu = \text{constant}/(1 - \phi_1 - \phi_2)$.

The fitting of an AR(2) time series model requires non-linear iterative estimation, but is easily facilitated by many software packages including MINITAB 13.3 used here. Table 1 provides the final estimates of the parameters.

For this AR(2) model to be stationary, it is required that the coefficients satisfy the relations,

$$\begin{aligned} \phi_2 + \phi_1 &< 1 \\ \phi_2 - \phi_1 &< 1 \qquad (8) \\ -1 < \phi_2 &< 1. \end{aligned}$$

Thus with $\hat{\phi}_2 + \hat{\phi}_1 = 0.61$, $\hat{\phi}_2 - \hat{\phi}_1 = -1.35$ and $\hat{\phi}_2 = -0.3722$, we conclude that the process is indeed stationary and therefore, as far as we are concerned, in statistical control. (For more information, see Box et al., 1994)

Table 1
Estimated coefficients for an AR(2) process

Coefficient	Estimate	Standard error	t-value	p-value
$\hat{\phi}_1$	0.9824	0.1062	9.25	0.000
$\hat{\phi}_2$	−0.3722	0.1066	–3.49	0.001
Constant	615.836	0.042		
$\hat{\mu}$	1579.79	0.11		
$\hat{\sigma}_a^2$	0.1403			

To check the model, we show in Figure 7 a normal plot of the residuals, and in Figure 8 a plot of the residuals in time order. Both plots indicate that the model fits the data well. The ACF and the PACF of the residuals provide a further check. Ideally, if the model fits well, all serial correlation would have been removed from the data and the residuals behave like white noise. Figures 9 and 10 show the ACF and the PACF for the residuals after fitting the AR (2) model to the furnace data. Both the ACF and the PACF are essentially zero for all lags.

Now that we have identified the furnace temperature to be a stationary AR(2) process and confirmed that the model fits the data well, we can compute the process's stationary variance. The variance inflation factor (see Box et al., (1994)) for an AR(2) process with $\phi_1 = 0.9824$ and $\phi_2 = -0.3722$ is

$$\begin{aligned} \sigma_z^2 &= \left(\frac{1-\phi_2}{1+\phi_2}\right) \frac{\sigma_a^2}{\{(1-\phi_2)^2 - \phi_1^2\}} \\ &= \left(\frac{1+0.3722}{1-0.3722}\right) \frac{\sigma_a^2}{\{(1+0.3722)^2 - 0.9824^2\}} \qquad (9) \\ &= 2.38146\sigma_a^2. \end{aligned}$$

Thus, the variance of the furnace temperature process is approximately 2.4 times larger than the residual variability, σ_a^2. This partly explains why the individuals control chart in Figure 2 underestimated the true process variability, and hence gave so many false alarms. Further, from Table 1 we have $\hat{\sigma}_a^2 = 0.1403$. Thus, $\hat{\sigma}_z = \sqrt{2.38146 \times 0.1403} \cong 0.5780$. This compares well with the estimate of the overall sample standard deviation, $s = \sqrt{\sum (z_t - \bar{z})^2/(80-1)} \cong 0.5685$.

AN ALTERNATIVE INDIVIDUALS CONTROL CHART FOR THE FURNACE PROCESS

A number of considerations go into setting up a control chart. One is to reduce the number of false alarms; we do not want the chart to signal that the process is out of control when it is not. Another consideration is to make the chart sensitive enough to get a quick, valid alarm to a real change in the process. Unfortunately, these two considerations are to some degree at cross-purposes. When we try to reduce the false alarm rate, we desensitize the control chart and hence, delay its ability to signal fast when a real change has happened. Another more practical consideration is that the chart should be easy for the user to interpret. In the present context, it is desirable that the control

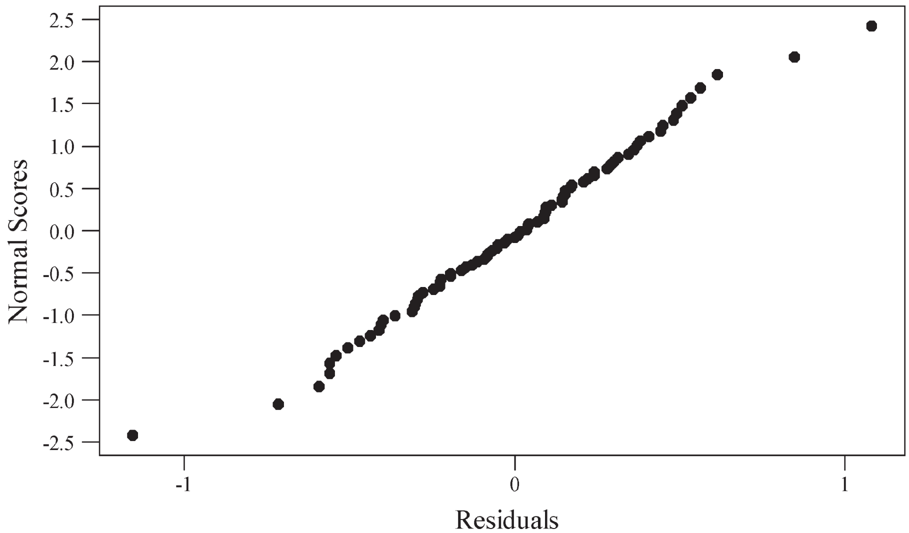

Figure 7. Normal plot of residuals after fitting an AR(2) model to the furnace data.

chart shows the actual temperature, because that carries with it important information for the engineer.

Given these considerations, we suggest using an individuals control chart with appropriately inflated control limits. The inflated control limits are computed as

$$\begin{aligned} \mathrm{UCL} &= \bar{z} + 3\hat{\sigma}_z = 1579.79 + 3 \times 0.5780 = 1581.52 \\ \mathrm{Centerline} &= \bar{z} = 1579.79 \qquad (10) \\ \mathrm{LCL} &= \bar{z} - 3\hat{\sigma}_z = 1579.79 - 3 \times 0.5780 = 1578.06. \end{aligned}$$

Figure 11 shows the modified individuals control chart for the furnace process. Note that the process now appears to be in statistical control. A major reason for recommending a modified individuals control chart is that it provides comfort to the control engineer because it shows the actual temperatures backed with control limits that have physical meaning. Admittedly, this chart is perhaps not as sensitive and as quick to detect a real change as other types of control charts. However, in the current context, the overriding concern was to avoid that the control engineers dismissed the use of control charts all together.

In the above discussion, we have outlined how to provide a theoretical foundation for inflating the

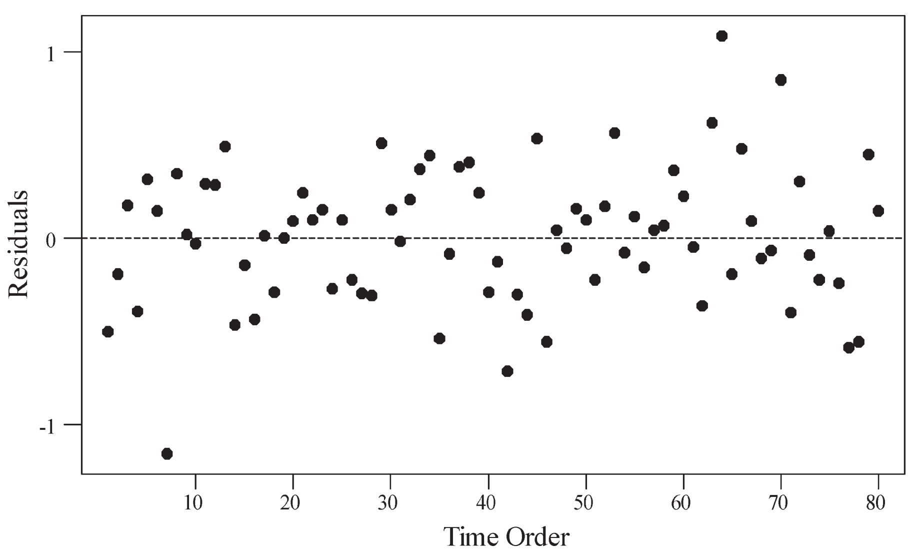

Figure 8. Time series plot of the residuals.

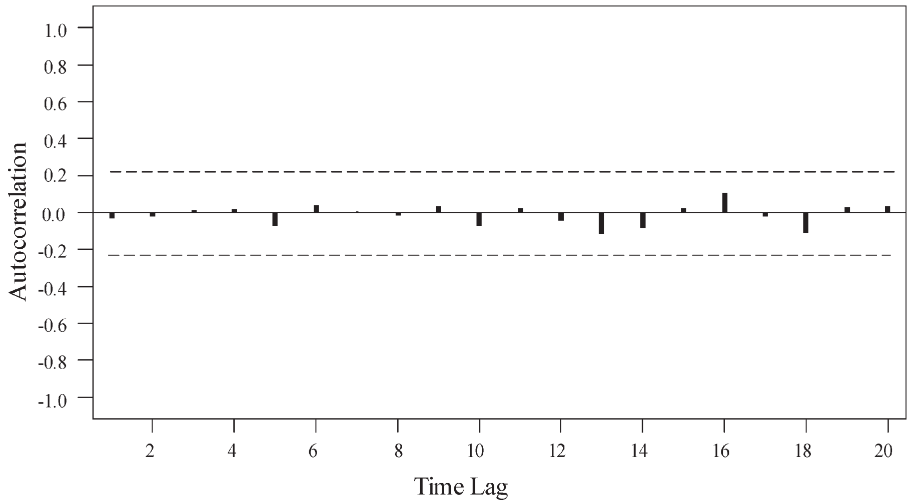

Figure 9. The ACF of the residuals after fitting an AR(2) model to the furnace data.

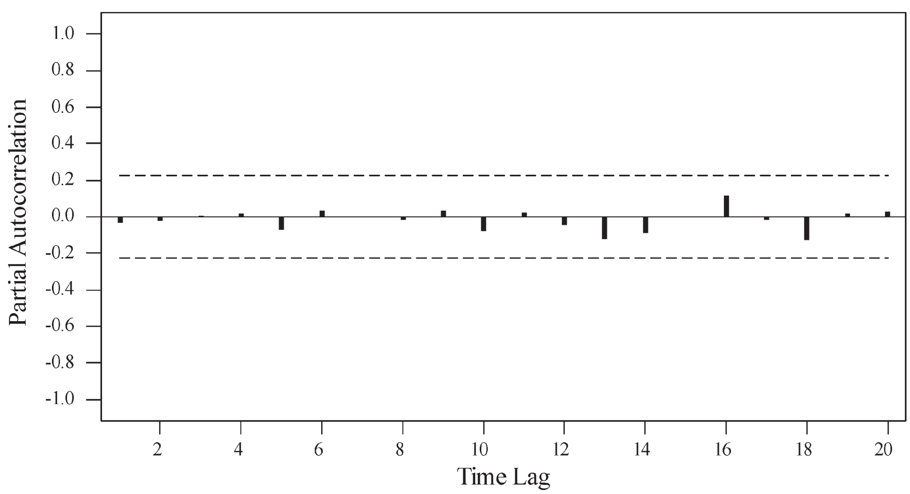

Figure 10. The PACF of the residuals after fitting an AR (2) model to the furnace data.

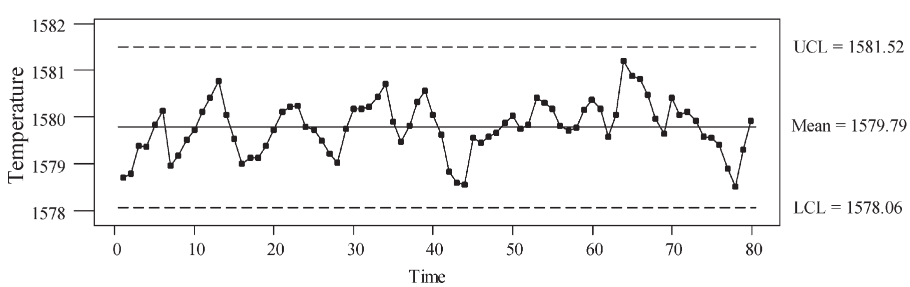

Figure 11. An individuals control chart for the furnace process with inflated limits to accommodate for the serial correlation.

control limits on an individuals control chart via ARMA time series modeling of a serially correlated but stationary process. We have done so because the control engineer was unhappy with our more practical first suggestion, which simply was to multiply the estimated sigma based on the moving range ($\approx$0.30) by a factor of two. As it turns out, this suggestion was very nearly what we ended up using based on the time series

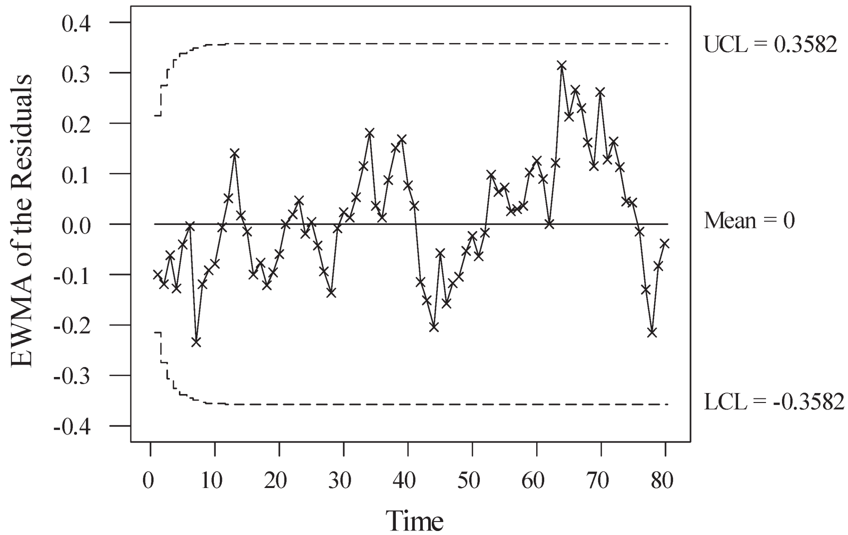

Figure 12. EWMA of the residuals after fitting an AR(2) model to the furnace data.

approach. However, rather than simply guessing the inflation factor, a better alternative is to use the overall sample standard deviation for a sufficiently long, good, and stable period. In the current case, using the 80 observations as such a period, we would have concluded that $\hat{\sigma}_z \approx 0.57$, and again we would have been very close to what we found using the ARMA time series approach.

ANOTHER ALTERNATIVE USING RESIDUALS

In using the inflated limits for the individuals control chart, we emphasized the importance of reducing the false alarm rate, and making the chart easy to interpret. However, this approach desensitizes the chart and will likely increase the average run length (ARL) to signal an alarm in case of a real change. An alternative approach, first suggested by Berthouex, Hunter, and Pallesen (1978), see also Montgomery (2001), is to use a standard control chart intended for independent observations applied to the residuals after fitting a time series model. In the present case, the residuals, $\hat{a}_t$, are computed by rewriting the time series model as $\tilde{z}_t - \hat{\phi}_1\tilde{z}_{t-1} - \hat{\phi}_2\tilde{z}_{t-2} = \hat{a}_t$, or in terms of the actual observations as,

$$\begin{aligned}\hat{a}_t &= (z_t - \hat{\mu}) - \hat{\phi}_1(z_{t-1} - \hat{\mu}) - \hat{\phi}_2(z_{t-2} - \hat{\mu}) \\ &= z_t - \hat{\phi}_1 z_{t-1} - \hat{\phi}_2 z_{t-2} - \hat{\mu}(1 - \hat{\phi}_1 - \hat{\phi}_2).\end{aligned} \tag{11}$$

For the current process, we could use an individuals control chart, a cumulative sum (CUSUM) chart or an EWMA chart. The residuals are not on a meaningful scale. Hence the practical interpretation argument for using the individuals control chart no longer applies. We therefore suggest using an EWMA chart, as show in Figure 12. This chart, computed with the default $\lambda = 0.2$, and forcing the mean to be zero, now shows that the process is in statistical control.

CONCLUSION

In modern applications of statistical process control, autocorrelation is increasingly becoming a fact of life and must not be ignored. In this article, we have demonstrated with an industrial example how to detect autocorrelation, illustrated its consequences for a few standard control charts, and showed a couple of ways to alleviate the problem. As demonstrated, modern software packages such as MINITAB, make it relatively easy to perform the computations needed when dealing with auto-correlated processes and using ARMA time series models. By providing the data, we hope this detailed example will serve as a hands-on introduction to the use of time series approaches to statistical process control.

ACKNOWLEDGMENTS

Søren Bisgaard was supported by the Isenberg Program for Technology Management, the Isenberg School of Management, University of Massachusetts Amherst, and by the European Community (EC) through the Thematic Network–Pro–ENBIS–EC contract number G6RT-CT-200 + -05059. MINITAB was used for all the computations.

REFERENCES

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

Box, G. E. P., Luceño, A. (1997). *Statistical Control by Monitoring and Feedback Adjustment*. New York: Wiley.

Berthouex, P. M., Hunter, W. G., Pallesen, L. (1978). Monitoring a Sewage Treatment Plant: Some Quality Control Aspects. *Journal of Quality Technology* 10:139–149.

Hunter, J. S. (1986). The Exponentially Weighted Moving Average. *Journal of Quality Technology* 18(4):203–210.

Montgomery, D. C. (2001). *Introduction to Statistical Quality Control*, 4th ed. New York: Wiley.

APPENDIX

Table A.1
Furnace temperature data: 80 consecutive hourly observations reading horizontally line-by-line

1578.71	1578.79	1579.38	1579.36	1579.83	1580.13	1578.95
1579.18	1579.52	1579.72	1580.11	1580.41	1580.77	1580.05
1579.53	1579.00	1579.12	1579.13	1579.39	1579.73	1580.12
1580.23	1580.25	1579.80	1579.72	1579.49	1579.22	1579.03
1579.76	1580.19	1580.17	1580.22	1580.44	1580.71	1579.91
1579.48	1579.82	1580.34	1580.56	1580.05	1579.63	1578.82
1578.59	1578.56	1579.56	1579.46	1579.59	1579.66	1579.89
1580.03	1579.76	1579.84	1580.41	1580.30	1580.17	1579.81
1579.71	1579.77	1580.16	1580.38	1580.18	1579.59	1580.06
1581.21	1580.89	1580.82	1580.48	1579.97	1579.64	1580.42
1580.06	1580.12	1579.92	1579.57	1579.56	1579.40	1578.90
1578.50	1579.30	1579.93				