10

Loop and Performance Assessment

10.1 Introduction

The design, tuning, and implementation of control strategies and controllers is only the first phase in the solution of a control problem. The second phase includes operation, supervision, and maintenance. This phase has traditionally been handled manually, but the interest for automatic supervisory functions has increased significantly in recent years because of the reduction of personnel in the process industry.

This chapter treats methods for commissioning, supervision, and diagnosis of control loops. The adaptation methods presented in Chapter 9 were divided into two categories, tuning on demand and continuous adaptation. Procedures for supervision and diagnosis can be classified in the same way. We call them loop assessment and performance assessment. Loop assessment procedures are used to investigate properties of the control loop, e.g., signal levels, noise levels, nonlinearities, and equipment conditions. Performance assessment procedures are used to supervise the control loops during operation and ensure that they meet the specifications. Failure to meet the specifications may be caused by equipment problems, nonlinearities, or other variations in process dynamics or the surroundings.

The chapter begins with a presentation of problems occurring in valves. These problems are identified as one of the major reasons for bad control loop performance. Sections 10.3 and 10.4 treat loop assessment and performance assessment, respectively. Tuning and diagnosis have many aspects in common. These aspects are discussed in Section 10.5.

10.2 Valves

Control valves are subject to wear. After some time in operation, this wear results in friction and hysteresis that deteriorates the control performance. Furthermore, valves are often both nonlinear and over-sized. Therefore, valves

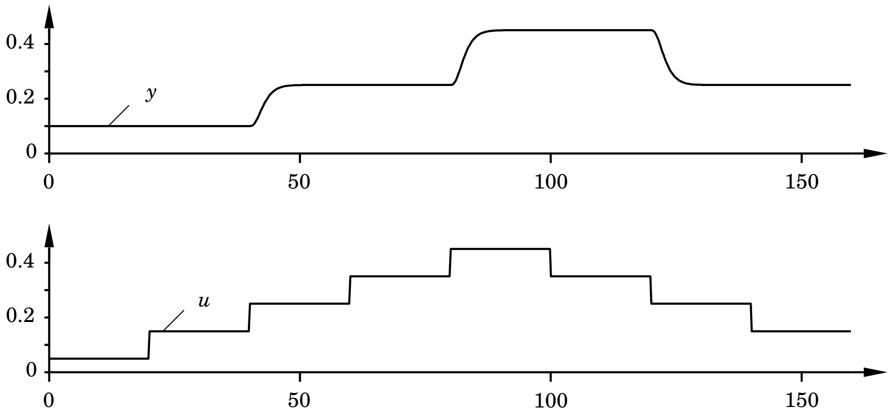

Figure 10.1 Procedure to check the amount of valve friction. The upper diagram shows process output y and the lower diagram shows control signal u.

have been identified as the major source of problems at the loop level in process control. We therefore devote a section to these problems.

Friction in the Valve

High friction in the valve is a common cause of problems. There is, of course, always static friction (stiction) in the valve, but if valve maintenance is insufficient, friction may be so large that the control performance degrades. The amount of friction can easily be measured by making small changes in the control signal and observing how the process outputs react. The procedure is shown in Figure 10.1. In the figure, the process output only responds to the control signal when the changes in the control signal are large enough to overcome the static friction.

Friction in the valve results in stick-slip motion. This phenomenon is illustrated in Figure 10.2. Suppose that the valve is stuck at a certain position due to friction. If there is a control error, the integral action of the controller will cause the controller output to increase until the pressure in the actuator is high enough to overcome the static friction. At this moment, the valve moves (slips) to a new position where it is stuck again. This valve position is normally such that the process output is moved to the other side of the set point, which means that the procedure is repeated. The process output will therefore oscillate around the set point. The pattern in Figure 10.2, where the measurement signal is close to a square wave and the control signal is close to a triangular wave, is typical for stick-slip motion.

Many operators detune the controller when they see oscillations like the one in Figure 10.2, since they believe that the oscillations are caused by a bad controller tuning. Unfortunately, most adaptive controllers do the same. What should be done when a control loop starts to oscillate is to first determine the cause of the oscillation. A good way to do this is presented in Figure 10.3.

The first problem to determine is whether the oscillations are generated inside or outside the control loop. This can be done by disconnecting the feed-

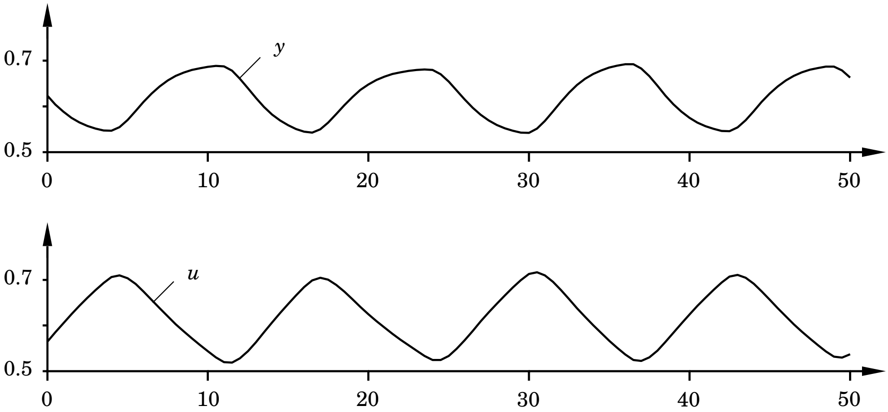

Figure 10.2 Stick-slip motion caused by valve friction and integral action. The upper diagram shows process output y, and the lower diagram shows control signal u.

back, e.g., by switching the controller to manual mode. If the oscillation is still present, the disturbances must be generated outside the loop, otherwise they were generated inside the loop. There might be a situation when the control loop oscillates because of valve friction even when the controller is in manual mode, namely, if the friction occurs in the pilot valve of the positioner instead of the valve itself.

If the disturbances are generated inside the loop, the cause can be either friction in the valve or a badly tuned controller. Whether friction is present or not can be determined by making small changes in the control signal and checking if the measurement signal follows, as shown in Figure 10.1. If friction is causing the oscillations, the solution to the problem is valve maintenance.

If the disturbances are generated outside the control loop, one should try, of course, to find the source of the disturbances and try to eliminate it. This is not always possible, even if the source is found. One can then try to feed the disturbances forward to the controller and in this way reduce their effect on the actual control loop. See Section 5.6.

Hysteresis in the Valve

Because of wear, there is often hysteresis (backlash) in the valve or actuator. The amount of hysteresis can be measured as shown in Figure 10.4. The experiment starts with two step changes in the control signal in the same direction. The hysteresis gap will close if the first step is sufficiently large. This means that the second step is performed without hysteresis. The third step is then made in the opposite direction. The control signal then has to pass the whole gap before the valve moves. If the last two steps are of the same size, the hysteresis is $\Delta y/K_p$, where Δy is the difference between the process outputs after the second and the third step (see Figure 10.4), and K_p is the static process gain (also easily obtained from Figure 10.4).

The hysteresis can also be determined from a continuous sweep over parts of the operating range. Figure 10.5 shows the process outputs from a process

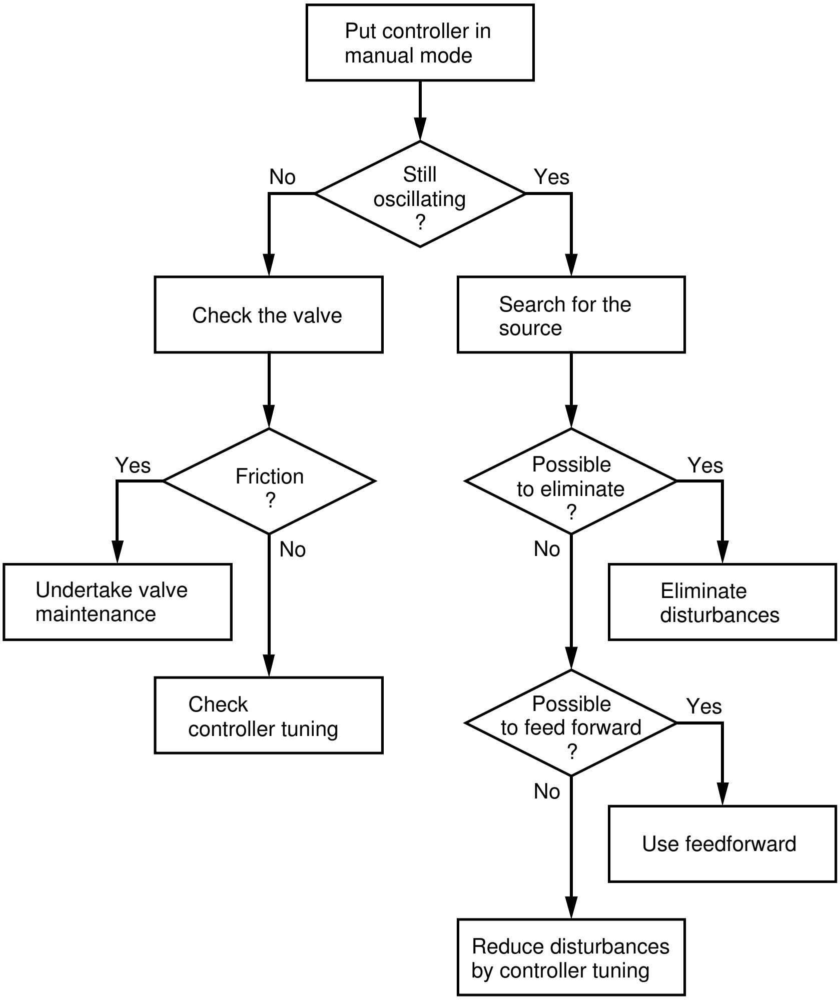

Figure 10.3 Diagnosis procedure to discover the cause of oscillations and recommended actions to eliminate them.

with friction and a process with hysteresis, respectively, when the process input is ramped from zero to one and then back to zero again. The corresponding phase plots are presented in Figure 10.6. One can easily measure the amount of hysteresis from the phase plot. Sweeps of this type are conveniently done during commissioning.

Figure 10.7 shows closed-loop control of a process with 10 percent hysteresis in the valve. The process is

$$P(s) = \frac{1}{(1+0.05s)^2} e^{-0.3s},$$

and the controller is a PI controller with parameters $K = 0.35$ and $Ti = 0.15$. The control signal has to travel through the gap in order to move the valve.

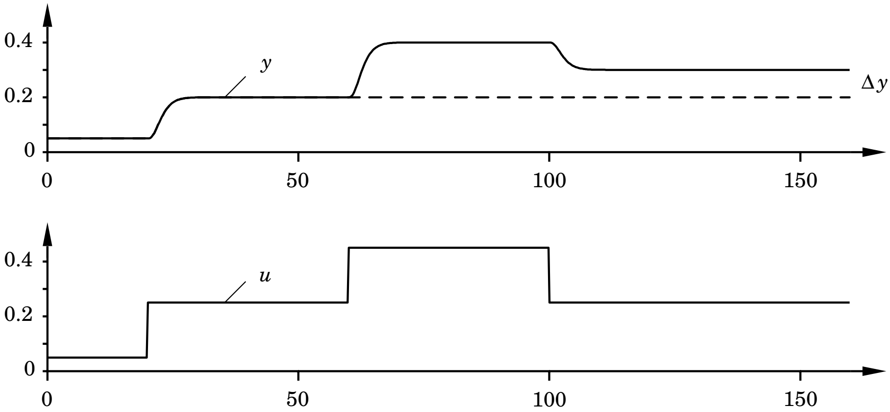

Figure 10.4 Procedure to check valve hysteresis. The upper diagram shows process output y, and the lower diagram shows control signal u.

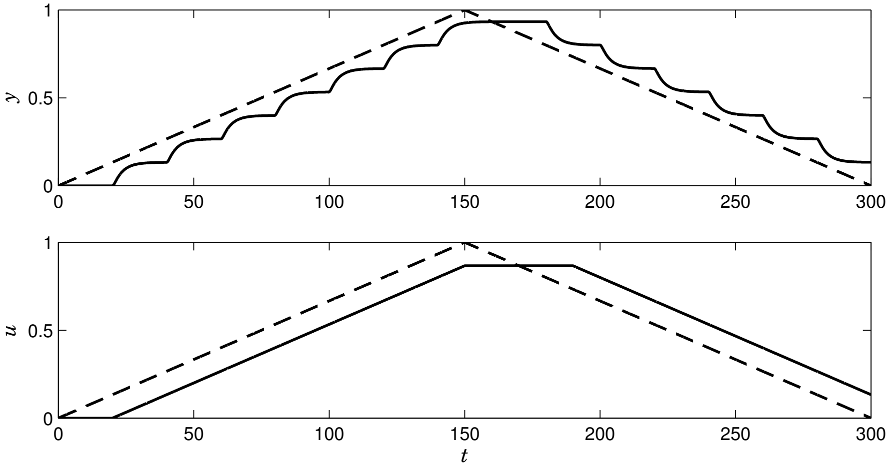

Figure 10.5 Process outputs (solid lines) and control signals (dashed lines) for process with friction (upper graph) and hysteresis (lower graph).

Therefore, we get the typical linear drifts in the control signal as shown in Figure 10.7.

If a relay auto-tuner is applied to a process with hysteresis, the estimated process gain will be smaller than the true value. This gives a too large controller gain. An auto-tuner based on a step-response experiment will work properly if the gap is closed before the step-response experiment is performed. (Compare with the second step in Figure 10.4).

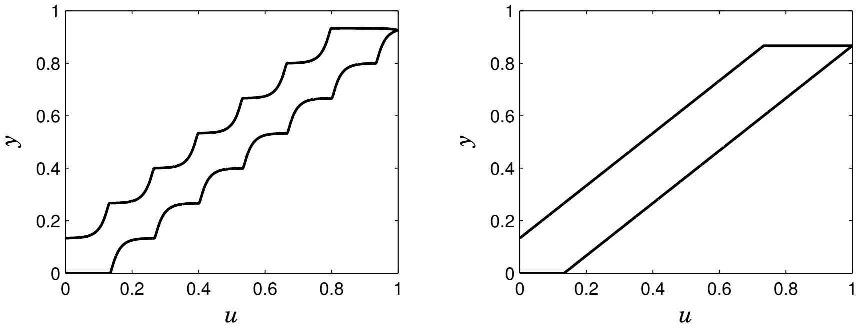

Figure 10.6 Phase plots of the signals in Figure 10.5 for the process with friction (left) and hysteresis (right).

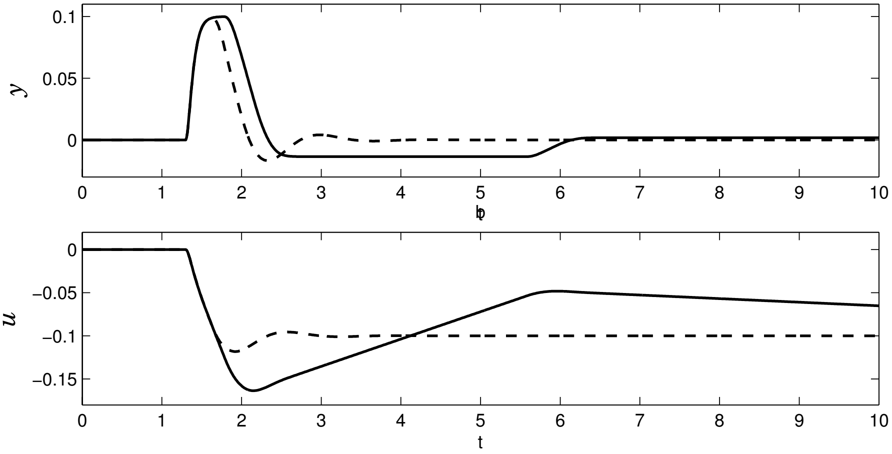

Figure 10.7 Closed-loop control with valve hysteresis. The upper diagram shows process output y, and the lower diagram shows control signal u. The dotted lines show control without hysteresis. The solid lines show control with a hysteresis of 10 percent (0.1).

10.3 Loop Assessment

This section suggests tests that are useful to perform on the control loop. These tests should be performed regularly, and especially in connection with controller tuning. The tests for friction and hysteresis, presented in Section 10.2, are two important loop assessment procedures. The experiments suggested in Section 2.7 to obtain the process dynamics are also loop assessment procedures for tuning the controllers. The checks and tests added in this section are basic, but often forgotten or neglected.

Signal Ranges

The signal range of the measurement signal is related to the resolution of the sensor. A large signal range means that the resolution becomes low. To obtain

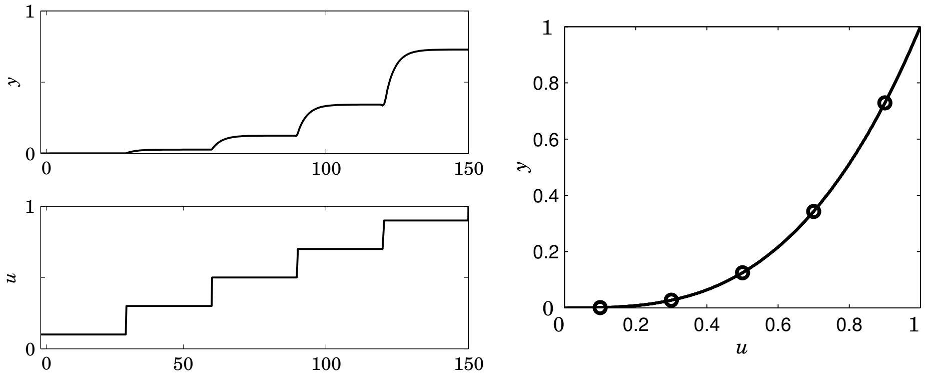

Figure 10.8 The left diagrams show a procedure to determine the static process characteristic. Control signal u is changed stepwise, and the corresponding changes in process output y are determined. The right diagram shows the static process characteristic, i.e., process output y as function of control signal u.

a high resolution, it is therefore important to restrict the signal range to those values that are relevant for the control.

If the final control element is a valve, the output range is determined by the size of the valve. Valves are normally over-sized. The main reasons are insecurity among engineers combined with a fear of installing a valve that is too small to deliver the maximum possible flows.

A large valve has not the same accuracy as a smaller one. The friction and backlash problems discussed in the previous section are more severe if the valve is over-sized.

If the signal ranges are properly chosen and if the process is linear, the ideal static process gain is $P(0) = 1$. If the static gain is one, the measurement signal reaches its maximum value when the control signal is at its maximum value. Because of over-sized valves, the static process gain is often larger than one in process control applications.

Static Input-Output Relations

From a control point of view, it is desirable to have a linear static input-output relation. This relation is, however, often nonlinear, mainly because of a nonlinear valve characteristic. Nonlinearities may also occur in sensors or in the process itself.

If the process is nonlinear, the control may be improved using gain scheduling or other forms of linearization. As pointed out in Section 9.3, it is important to understand the cause of the nonlinearity in order to determine a suitable gain-scheduling reference.

The static characteristic of the process can be obtained by determining the static relation between the control signal and the measured signal. This can be done by performing step changes in the control signal and measuring the corresponding changes in process output; see Figure 10.8.

The characteristic shown in Figure 10.8 is obviously nonlinear. It has a higher gain at larger control signals. If the stationary values of the measured signal are plotted against the control signal, we obtain the static process characteristic. See Figure 10.8. A plot like this reveals whether gain scheduling is suitable or not.

Disturbances

Another important issue to consider before tuning the controller is the disturbances acting on the control loop. We have pointed out that it is important to know if the major disturbances are set-point changes (the servo problem) or load disturbances (the regulator problem).

It is also important to investigate the level of the measurement noise and its frequency content. Compare with Section 2.6. If the noise level is high, it may be necessary to filter the measurement signal before it enters the control algorithm. This is an easy way to get rid of high-frequency noise. If there are disturbances with a large frequency content near the ultimate frequency, it is not possible to use low-pass filtering to remove them. Feedforward is one possibility, if the disturbances can be measured at their source. Notch filters can be used if the noise is concentrated in a narrow frequency range. See Section 2.6 where noise modeling and measurements were discussed.

10.4 Performance Assessment

The loop assessment, followed by appropriate actions like valve maintenance, selection of signal ranges, linearization of nonlinearities, and controller tuning, should leave the control loop in good shape.

After some time in operation, the performance may, however, deteriorate because of variations in the process and the operation. Therefore, it is important to supervise the control loops and detect these degradations. This supervision has traditionally been made by humans, but the reduction of personnel in the process industry combined with increasing quality demands have been a driving force behind developing procedures for automatic performance monitoring and assessment. This section provides some procedures for automatic supervision of control loop performance.

The Static Input-Output Relation

If a detector for stationarity is available, it is simple to keep a statistic for the fraction of time that the system is stationary. The static input-output relation can then be obtained simply by logging the process input and output during stationary conditions. To obtain good data the signals should be filtered with respect to the time scale of the closed loop. Graphs like the ones shown in Figure 10.9 are then obtained. From these curves it can be determined whether the major variations in the output are due to set-point changes or load disturbances, i.e., whether we are dealing with a servo problem or a regulation problem. We have a servo problem if the experimental data gives a well-defined curve and a regulation problem if there is no definite relation between inputs and outputs. A simple statistic of the fraction of the total time when there are

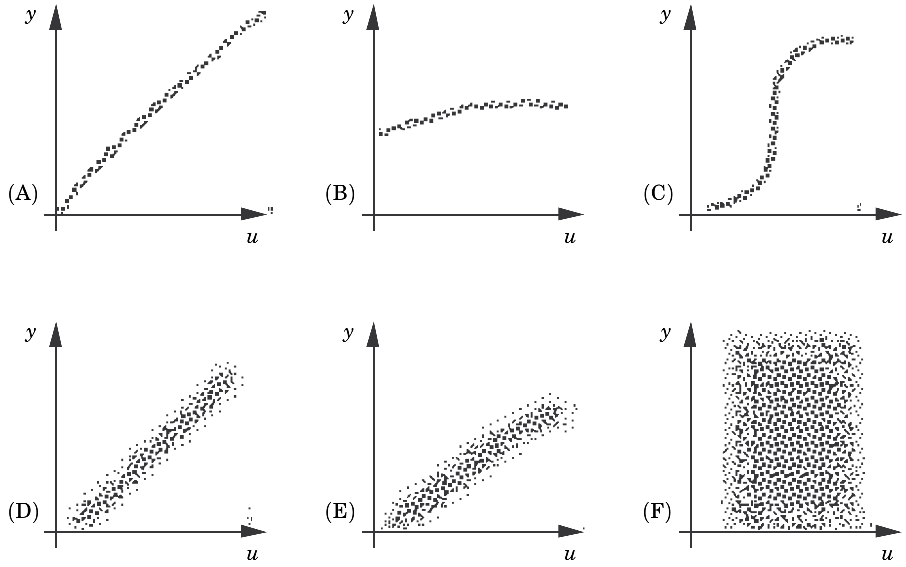

Figure 10.9 Examples of static input-output data logged during normal operation. The results shown in A, B, and C indicate a pure servo problem. The results in F indicate a pure regulation problem. Case D and E are mixed cases. Case B indicates poor resolution of the sensor, and case E indicates poor actuator sizing.

set-point changes or transients due to set-point changes is also a useful indicator. Of course, there are also systems that are mixtures of servo and regulation problems.

For a servo problem the variations in the static gain of a system can also be determined. This gives a valuable indication as to whether gain scheduling is required. The static gain curve can also be used for diagnostic purposes. Changes in the curve indicate changes in the process. By comparing the slope of the static gain curve with the incremental process gain measured during tuning or adaptation, we can also get indications of whether there is some hysteresis in the loop or not. It also indicates if actuators are properly sized.

Model-Based Diagnosis

Most automatic supervisory procedures are, in principle, based on the idea shown in Figure 10.10. If a model of the process is available, the control signal can be fed to the input of the process model. By comparing the output of the model with the true process output, one can detect when the process dynamics change. If the model is good, the difference between the model output and the process output (e) is small. If the process dynamics change, e will no longer be small, since the two responses to the control signal are different.

Harris Index

One of the most widely applied supervisory functions is based on the Harris index. The idea is to calculate the variance of the process output, either on line or off line, and then compare it with the minimum variance obtainable. The

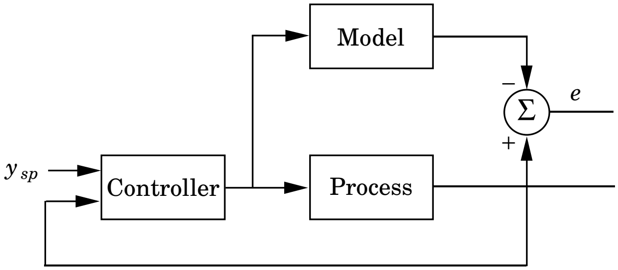

Figure 10.10 Model-based fault detection.

problem was discussed in Section 2.6. The Harris index is defined as

$$I_H = 1 - \frac{\sigma_{MV}^2}{\sigma_y^2},$$

where σ_{MV}^2 is the minimum variance of the process output, and σ_y^2 is the actual process output variance. The Harris index, I_H, takes values between zero and one. If the index is close to zero, the actual variance is close to the minimum variance, which means that the control loop behaves satisfactorily. If the actual variance is large, the Harris index is close to one.

The method requires that the minimum variance σ_{MV}^2 is known. A nice feature of the method is that the minimum variance can be determined from the deadtime only, which means that the modeling can be made relatively simple. A drawback is that the minimum variance normally cannot be achieved with a controller as simple as the PID controller, which means that it is difficult to determine reasonable values of the Harris index. Furthermore, even if it is possible to obtain minimum variance control, this control is often undesirable since it may be very aggressive.

For these reasons, many variations of the Harris index have been presented where the minimum variance σ_{MV}^2 is replaced with the variance obtained using other design objectives and where the limitations to the PID control structure are taken into account. The main drawback of these approaches is that they require a more accurate process model.

The performance monitoring tools based on the Harris index approach provides information about the loop performance compared to some ideal performance. There is no intention to detect any causes of possible bad performance. There are other performance monitoring tools that do not look at the overall performance, but instead try to detect certain types of problems. Some of these are discussed in the following subsections.

Oscillating Control Loops

The most serious problem at the loop level is that many control loops oscillate. There are several possible causes of these oscillations; see Section 10.2. One reason might be that an oscillating load is disturbing the loop. Low-frequency load disturbances are eliminated efficiently by the controller, since a controller with integral action gives a high loop gain at low frequencies. Since the process normally has a low-pass character, high-frequency load disturbances are

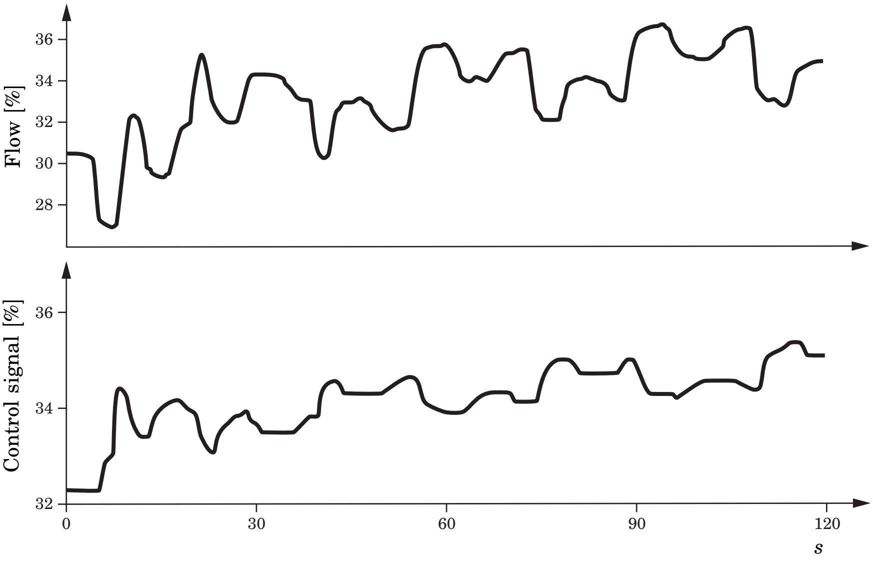

Figure 10.11 Stick-slip motion in a flow control loop.

filtered by the process. Therefore, high-frequency components in the measurement signal are normally not introduced in the process but in the sensor or on the connections between the sensor and the controller. Since they do not contain any valuable information about the status of the process, they should be filtered out by the controller. It is also important not to transfer these signals to the controller output, since they may cause wear on the actuating equipment.

Disturbances with much energy near the ultimate frequency ω_u are too fast to be treated efficiently by the controller, and they are too slow to be filtered out. These disturbances might even be amplified because of the feedback.

A badly tuned controller may be another reason for oscillations, in particular in nonlinear plants where a change in operating point might result in a too high loop gain. However, controllers in process control plants are often tuned conservatively, and bad controller tuning is not the most likely cause of oscillations.

The most common reason for oscillations in control loops is, however, friction in the valve, resulting in “stick-slip” motion as discussed in Section 10.2.

Detection Oscillations in control loops can be detected in several ways. One way is to make a spectral analysis of the measured signal and look for peaks in the spectrum. A difficulty is that the oscillations often are far from pure sine waves, which means that no distinct peaks appear in the spectrum.

Figure 10.11 shows a recording from a flow control loop in a paper mill with high valve stiction. The figure shows the result of a step change in the set point. The controller used was a PI controller with gain $K = 0.30$ and integral time $T_i = 34$ s. Notice that the oscillations are far from a pure sine wave. A retuning of the controller gave controller parameters $K = 0.19$ and $T_i = 2$ s. Notice that the integral time was decreased from 34 s to 2 s! A step

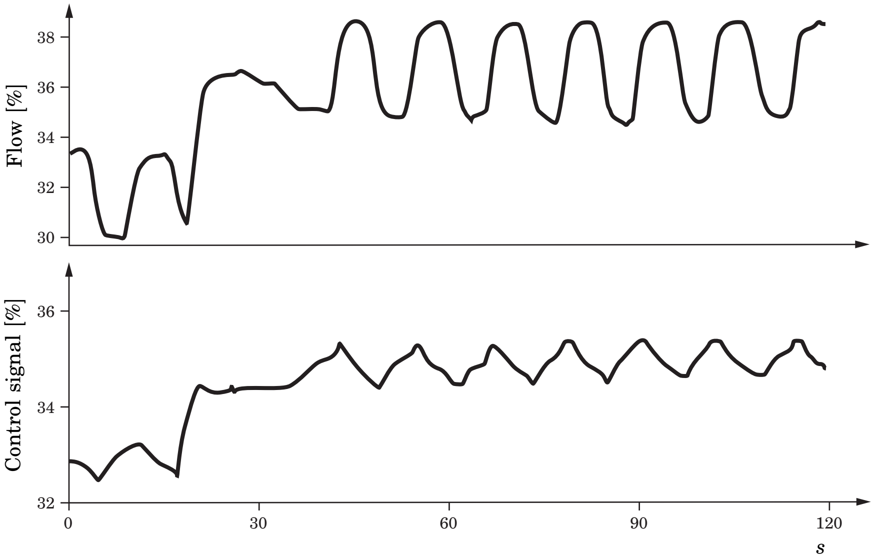

Figure 10.12 Stick-slip motion in a flow control loop – retuned controller.

response experiment using the new controller settings is shown in Figure 10.12. The settling time is significantly shorter than in Figure 10.11. It is also more obvious that the oscillations really are caused by friction, since the typical pattern of the measurement signal is close to a square wave and the control signal is close to a triangular wave.

Another approach to detect oscillations is to investigate the characteristics of the control error. The idea behind this detection procedure is to study the magnitude of the integrated absolute error (IAE) between successive zero crossings of the control error, i.e.,

$$IAE = \int_{t_{i-1}}^{t_i} |e(t)|dt, \tag{10.1}$$

where t_{i-1} and t_i are two consecutive instances of zero crossings. It is assumed that the controller has integral action, so that the average error is zero.

During periods of good control, the magnitude of the control error is small and the times between the zero crossings are relatively short. This means that the IAE values calculated from (10.1) are small when control is good.

When a load disturbance occurs, the magnitude of $e(t)$ increases, and there is a relatively long period without zero crossings. This means that the corresponding IAE value becomes large.

When the control loop starts to oscillate, there will be a high frequency of large IAE values. This observation is used to detect oscillations in the control loop.

EXAMPLE 10.1—PULP CONCENTRATION CONTROL

The following example is taken from a pulp concentration control section in a paper mill, where pulp is diluted with water to a desired concentration.

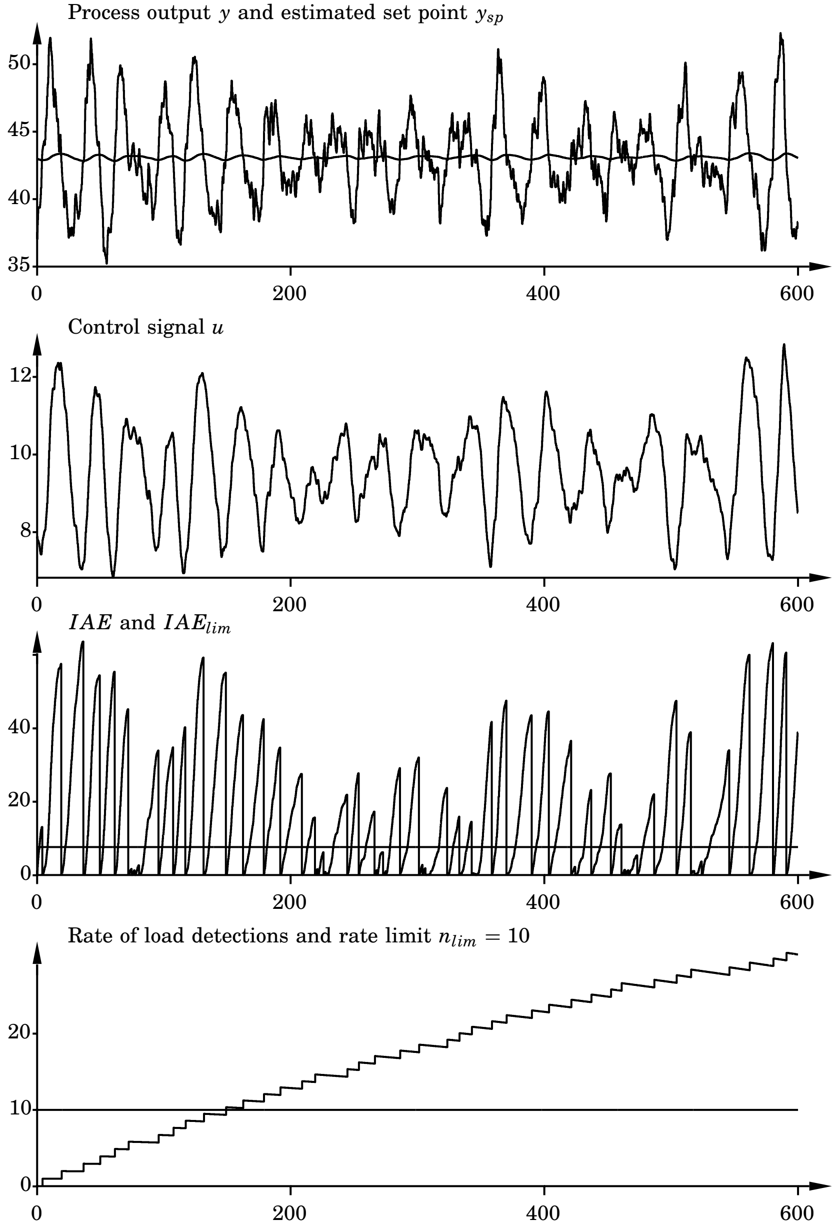

Figure 10.13 The oscillation detection procedure applied on a pulp concentration control loop.

The water valve had too high friction, and an oscillation detection procedure was connected to the controller. The controller was a PI controller with gain $K = 0.33$ and integral time $T_i = 24$ s.

Figure 10.13 shows 10 minutes of data from the concentration control loop. The first graph shows the process output, the pulp concentration in percent. Because of high friction in the water valve, the process is oscillating with an amplitude of a few percent. The first graph also shows an estimate of the set point, since this variable was not recorded. The estimate is simply obtained by a low-pass filtering of the process output.

The second graph shows the control signal in percent. It is obvious that the controller tries to eliminate the oscillation but without success.

The third graph shows the IAE calculated between successive zero crossings of the control error. The graph also shows IAE_{lim}, which is the limit of what is considered large values of IAE. In this implementation, the value of IAE_{lim} is determined automatically from the controller parameters in each loop. The IAE values are significantly larger than IAE_{lim}, indicating that the loop is oscillating.

The fourth graph finally shows the rate of load detections and the rate limit $n_{lim} = 10$. The rate exceeds the rate limit after about three minutes, and the detection procedure gives an alarm.

This example shows how the oscillation detection procedure manages to detect oscillations in control loops. The actual oscillations are easily noticed in Figure 10.13. However, process operators seldom have access to these kinds of graphs, but are often left with a bar graph with a low resolution. The present oscillation had been present for a long time without being discovered by the process operators. □

Diagnosis Since a control loop may oscillate for various reasons, it is important not only to detect the oscillation, but also to find the reason for oscillations. This can be done manually as described in Section 10.2.

Attempts have also been made to develop procedures for automatic diagnosis. Here, the difference in the spectrum can be used. When a control loop oscillates because of too high loop gain, the control error is often close to a sine wave, resulting in one single peak in the spectrum. The same holds in most cases when the loop is oscillating because of external disturbances. However, when the control loop is oscillating because of valve stiction, several peaks in the spectrum can be found.

Sluggish Control Loops

Oscillations in control loops are common, but the opposite situation is also common, namely, that the control loops are sluggish because of conservative tuning. This causes unnecessarily large and long deviations from the set point at load disturbances.

The main reason for the controllers being conservatively tuned is lack of time. The engineers tune the controllers until they are considered “good enough.” They do not have the time to optimize the control. Many controllers are tuned once they are installed, and then never again. To retain stability when operating conditions change, the controllers are tuned for the “worst case.” A better solution would, of course, be to use gain scheduling and perhaps adaptation. When a controller is retuned, it is mostly because the process conditions cause oscillatory control. In other words, when the controllers are retuned, they are detuned. When the process conditions change to sluggish control, the controller is normally not retuned again.

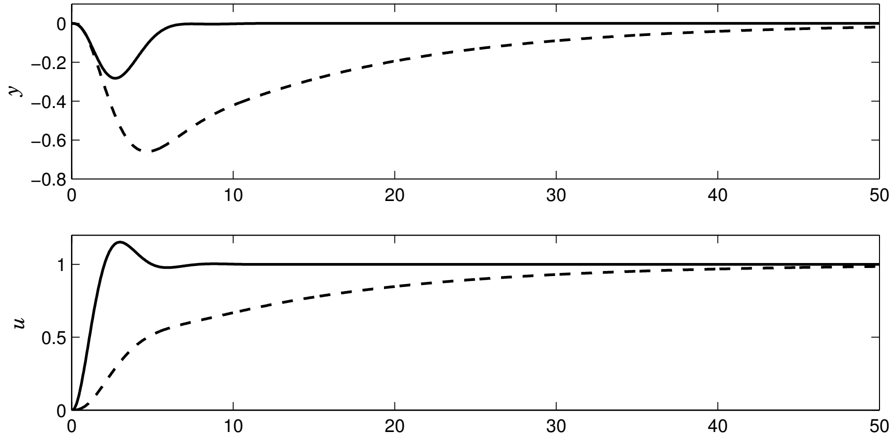

Figure 10.14 A good (solid lines) and a sluggish (dashed lines) response to a step change in load at the process input.

Detection Figure 10.14 shows two responses to load disturbances in the form of step changes at the process input. One response is good, with a quick recovery without any overshoot. The second response is very sluggish. One feature that characterizes the second response is that there is a long period where both process output y and control signal u drift slowly in the same direction. This feature is used for detection.

Both responses have an initial phase where the two signals go in opposite directions, i.e., $\Delta u \Delta y < 0$, where Δu and Δy are the increments of the two signals. What characterizes the sluggish response is that after this initial phase there is a very long time period where the correlation between the two signal increments is positive. This observation forms the base for the Idle index, which expresses the relation between the times of positive and negative correlation between the signal increments.

To form the Idle index, the time periods when the correlations between the signal increments are positive and negative, respectively, are first calculated. The following procedures are updated every sampling instant

$$t_{\text{pos}} = \begin{cases} t_{\text{pos}} + h & \text{if } \Delta u \Delta y > 0 \\ t_{\text{pos}} & \text{if } \Delta u \Delta y \leq 0 \end{cases}$$

$$t_{\text{neg}} = \begin{cases} t_{\text{neg}} + h & \text{if } \Delta u \Delta y < 0 \\ t_{\text{neg}} & \text{if } \Delta u \Delta y \geq 0, \end{cases}$$

where h is the sampling period. The Idle index I_I is then defined by

$$I_I = \frac{t_{\text{pos}} - t_{\text{neg}}}{t_{\text{pos}} + t_{\text{neg}}}. \tag{10.2}$$

Note that I_I is bounded to the interval $[-1, 1]$. A positive value of I_I close to 1 means that the control is sluggish. The Idle index for the sluggish response in

Figure 10.14 is $I_I = 0.82$. A negative value of I_I close to -1 may be obtained in a well-tuned control loop. The Idle index for the good response in Figure 10.14 is $I_I = -0.63$. However, negative Idle indices close to -1 are also obtained in oscillatory control loops. Therefore, it is desirable to combine the Idle index calculation with an oscillation detection procedure like the one described above.

Calculation of the Idle index can be made both off line and on line using a recursive version. Since the method is based on the characteristics of signal increments, it is sensitive to noise. Therefore, it is important to filter the signals properly before they are differentiated.

EXAMPLE 10.2—CONTROL OF A HEAT EXCHANGER
This example is taken from an industrial heat exchanger. The control objective is to control the water temperature on the secondary side by controlling the water steam flow on the primary side.

The upper graphs in Figure 10.15 show load responses obtained with a conservatively tuned PI controller. The controller parameters were $K = 0.01$ and $T_i = 30s$. The signals are relatively noisy because of the low resolution, 1 percent, of the controller output. The control is sluggish. This is also well reflected by the Idle index, which was calculated to $I_I = 0.8$.

The controller structure was changed to a PID controller and tuned properly, resulting in the controller parameters $K = 0.025$, $T_i = 8s$, and $T_d = 2s$. The improved control behavior is illustrated in the lower graphs in Figure 10.15. The recovery after load disturbances is significantly faster, still without any noticeable overshoot. The integral gain k_i is increased by almost a factor of ten. The improvements are also demonstrated by the Idle, which index that was reduced to $I_I = 0.3$. □

10.5 Integrated Tuning and Diagnosis

The diagnosis procedures are related to the adaptive techniques in several ways. We have pointed out the importance of checking valves before applying an automatic tuning procedure. If not done, the automatic tuning procedure will not provide the appropriate controller parameters. For this reason, it would be desirable to have these checks incorporated in the automatic tuning procedures. Such devices are not yet available, and the appropriate checks, therefore, must be made by the operator.

The on-line detection methods are related to the continuous adaptive controller. The adaptive controller monitors the control loop performance and changes the controller parameters, if the process dynamics change. The performance assessment procedures also monitor the control-loop performance. They give an alarm instead of changing the controller parameters if the process dynamics change. As an example, in Figure 10.3 we have seen that it is important to determine *why* the performance has changed before actions are taken. Most adaptive controllers applied to a process with stiction will detune the controller, since they interpret the oscillations as caused by a badly tuned controller. Consequently, it is desirable to supply the adaptive controllers with

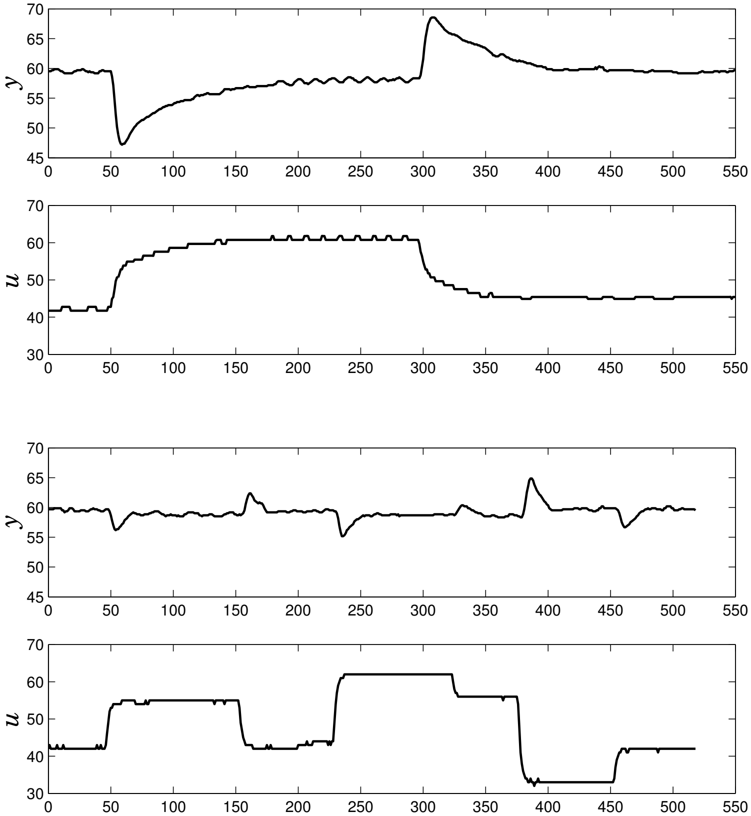

Figure 10.15 Control of a heat exchanger. The graphs show responses to load disturbances for a sluggish control loop with Idle index $I_I = 0.8$ (upper), and a properly tuned loop with Idle index $I_I = 0.3$ (lower).

on-line detection methods, so that reasons for bad control-loop performance, other than poor controller tuning, are detected. The lack of these kinds of detection procedures in adaptive controllers are perhaps the major reason for the relatively few applications of continuous adaptive control available today.

10.6 Summary

It is important to make an assessment of the control loop before tuning the controller. This assessment includes checks of equipment such as sensors and valves, signal ranges, nonlinearities, and disturbances.

When the loop assessment and the controller tuning is performed, the control loop should behave well. Due to changes in the process and its operation, the control loop may degrade after some time in operation. It is therefore important to supervise the control loops. This is traditionally done by humans, but methods for automatic supervision are becoming more and more used in process control.

In this section, some examples of loop and performance monitoring tools have been presented. The section has only provided a short overview of the area. It has focused on methods for the single loops only. In recent years, many attempts have been made to derive methods for the performance monitoring of process sections including several control loops. However, these procedures are seldom general, but often developed for specific plants.

10.7 Notes and References

Early work on fault-detection was done by [Himmelblau, 1978]. Problems associated with the control valves were brought to a broader audience in the early nineties; see [Ender, 1993; Bialkowski, 1994]. At that time there was also an awareness that it was beneficial to assess the performance of the control loops; see [Shinskey, 1990; Shinskey, 1991a; Åström, 1991]. The Harris index [DeWries and Wu, 1978], [Harris, 1989] is based on comparison with performance obtained by minimum variance control [Åström, 1970]. The concept has been extended and applied in various process control applications; see e.g. [Desborough and Harris, 1992; Stanfelj *et al.*, 1993; Harris *et al.*, 1996; Kozub and Garcia, 1993; Kozub and Garcia, 1996; Harris *et al.*, 1996; Owen *et al.*, 1996; Lynch and Dumont, 1996; Harris *et al.*, 1999; Thornhill *et al.*, 1999]. The oscillation detection procedure is described in [Hägglund, 1995] and [Thornhill and Hägglund, 1997], and the Idle index is presented in [Hägglund, 1999]. Good surveys of the area are presented in [Qin, 1998; Huang and Shah, 1999; Horch, 2000]. A method for reducing the effect of friction in valves was developed by [Hägglund, 2002].