Process Controller Performance Monitoring and Assessment

Alan J. Hugo
Control Arts Inc.

Introduction

Estimates of the percentage of industrial process controllers with performance problems are surprisingly high – various studies indicate that anywhere from 66%[1] to 80%[2] of controllers are not performing as well as they should. These controllers can have a significant detrimental effect on plant profitability, both in increased product variance and increased reaction times. In addition, the regulatory control layer is the foundation of all advanced applications; without a well tuned regulatory layer, DMC's, RTO's, ERP's, etc. can (and often do) fail[1].

A primary difficulty of controller performance monitoring is the shear number of loops to be monitored - a typical large processing operation consists of hundreds of control loops, often operating under varying conditions. The majority of the controllers use the PID algorithm, but there may also be advanced multivariate model-based controllers and other application specific controllers. Maintenance of these loops is generally the responsibility of either a control engineer or an instrument technician, but other responsibilities, coupled with the tediousness of consistently monitoring a large number of loops, often results in control problems being overlooked for long periods of time.

However, this task is well suited for automation. The data already resides in the DCS or plant historian, and plant tests are not required, as it is the closed-loop response of the process that is of interest. A complication arises from the fact that the any deviation from setpoint is a function of both the controller performance *and* the plant disturbance spectrum. Any controller performance methodology must separate out the effects of plant disturbances (which are external to the controller) from tuning, equipment problems, and out-of-service issues.

Controller performance assessment has been an area of active research for the last ten years, and several advanced algorithms are available in commercial software packages. These packages enable control engineers to easily and accurately obtain controller quality metrics without performing plant tests, and to monitor all aspects of their control loops. Discussed in this article are the various controller performance techniques, their applicability to the requirements of control engineers, and their advantages and limitations.

Performance Assessment Requirements

Ideally, any controller performance assessment technique should have the following attributes:

1. Independent of disturbance or setpoint spectrums. Both the disturbances and setpoint changes can vary widely in a plant, and the measurement should be insensitive to the time period when the data was taken.
2. Does not require plant tests. This requirement is generally met, as the user is interested in the closed-loop behavior of the process. However, closed-loop data can be information poor, and any performance assessment technique should include tests on the accuracy of the results.
3. Able to be automated. The large numbers of loops in a plant necessitate that at least part of the controller performance assessment be done automatically.
4. Requires minimum specification of process dynamics. The metric should not require a specification of what it is testing – the user may want to know the fidelity of the model in any model based controller, and, in the case of PID controllers, may not have a model.
5. Absolute or non-arbitrary measure. The metric should compare the current quality of control to some universal standard.
6. Sensitive to detuning or process model mismatch or equipment problems only. The metric should give an indication of only those things that the control engineer can affect.

7. Indicative of why the controller is performing poorly. Ideally, the measures should indicate what should be done to improve control, whether the problem is due to poor tuning, valve sticking, or oscillations from an unknown source.
8. Measures the improvement in profit due to the controller. This may be separate from measuring reduction in variance, as a major profit contribution for some controllers is pushing the process to constraints.

Current software packages generally meet requirements 1-6 above. Requirement 7 is only partially met – identifying new tuning parameters or a process model strictly from closed-loop data is the function of a self-tuning regulator (which has found very limited success in industry). The main difficulty in Requirement 8 is defining a base case, which is an activity that is best done off-line. However, performance assessment techniques can indicate whether advanced control can reduce the variance over the current PID controllers.

Controller Performance Assessment Standards

There are many ways to assess the quality of process controllers, but in general they explicitly or implicitly involve a comparison of the current quality of control to some standard. As shown in Figure 1, various standards can be ranked based on the tightness of control. One decision the control engineer must make is to which standard is most applicable. Note that the standards in Figure 1 use some definition of "best" control – generally this is the controller that minimizes the deviation from setpoint variance given the inherent limitations of the controller. While these "best" controllers could never be implemented (they would be very sensitive to model mismatch, and are almost unstable), they do represent an ultimate standard, and do not require any weighting or judgement factors.

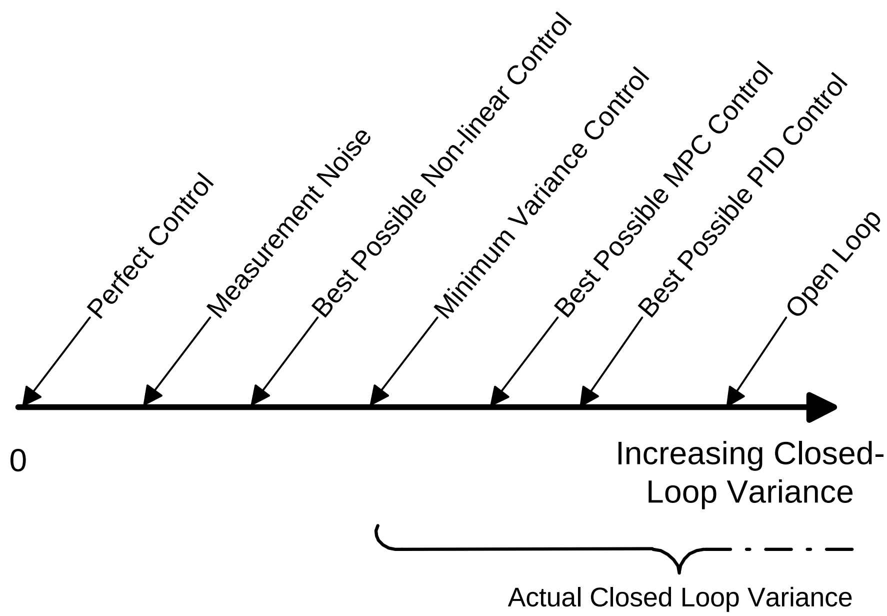

Figure 1: Most controller performance assessment techniques involve comparing the current behavior to some standard. The perceived performance, and the amount of information, is heavily dependent on which standard is used.

Perfect Control

While this may appear to be an unrealistic standard, it is in fact commonly invoked, at least implicitly. Assessing controllers based on output variance implicitly compares the performance to zero variance. Clearly this is too high a standard, but more importantly, output variance is largely a function of the setpoint change and disturbance spectrums (see Figure 2). Often justification of advanced control is done using before/after comparisons of the product variance, but this measure is unreliable unless the setpoint changes and disturbances are the same. At best these spectrums are unknown, at worst they vary widely – many processes can have very non-normal disturbances, with long periods of no disturbance punctuated with extreme disturbances due to step composition changes or equipment failures.

The other difficulty with this metric is that it does not indicate the amount of the variance that is a result of poor control, and what amount is due to other factors such as disturbances and measurement noise. It therefore gives a poor indication whether retuning or controller redesign will improve the product quality.

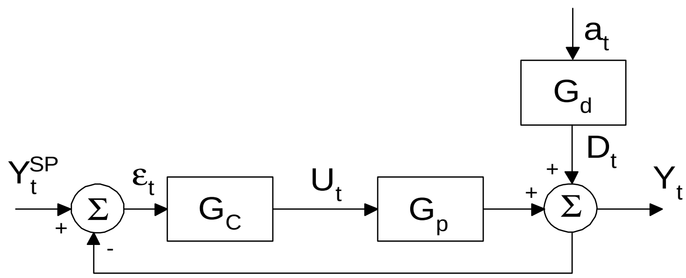

Figure 2. Closed Loop block diagram in standard form. Note that the output measurement Y_t is a direct function of the disturbance D_t.

Best Possible Non-linear Controller

There are 2 fundamental limitations to controller quality. Obviously, no controller can have tighter control than the random measurement error variance. Secondly, the deadtime presents a fundamental limitation to the achievable variance – no controller can affect the process before the deadtime. The best possible non-linear controller therefore represents a lower bound on what is possible using software. However, non-linear controllers are rare in industry due to both their complexity and the difficulty of obtaining a non-linear model. A non-linear performance assessment technique would be similarly complicated, and for these reasons, there does not appear to be any performance index based on non-linear controllers.

Minimum Variance Control

For linear systems, a minimum variance controller results in the smallest possible closed-loop variance. These controllers require a perfect process model (G_p in Figure #2), and a perfect disturbance model (G_d), and will result in complete cancellation of the error (other than measurement noise) one sample time after the process deadtime.

While the controller itself may require specification of process and disturbance transfer functions, Harris and Desborough[3] showed that the closed-loop response of a minimum variance controller may be determined using *only* closed-loop data and a estimation of the process deadtime. They defined a controller performance index h as the ratio of the variance that would be obtained if a minimum variance controller were applied to the system (s^2_{MVC}) to the actual variance of the closed-loop data (s^2_{ACT}):

$$h^{MVC} = 1 - \frac{s^2_{MVC}}{s^2_{ACT}} \qquad (1$$

This expression has many desirable theoretical and practical properties. It can be calculated relatively simply from a closed loop data set with a minimum amount of process knowledge (just the deadtime), and represents a lower bound on what can theoretically be obtained with linear controllers.

Because of its theoretical and practical advantages, this measure, or variants of it, is used in virtually all industrial controller assessment packages[4,5,6,7]. One disadvantage is that the user must specify the process deadtime or a prediction horizon. These will be available if modeling tests have been performed, and can be estimated otherwise. Table 1 gives values for common chemical processes.

Loop Type	Prediction Horizon (minutes)
Pressure	2
Liquid Flow	1
Temperature	5-10
Steam or Gas Flow	5
Level	2

Table 1: Prediction Horizons for common loops (from Thornhill et. al.[8]). The prediction horizon is an extended deadtime estimate that compensates for the limitations of PID control.

In reality, minimum variance controllers are never implemented industrially, and questions arise as to their applicability as a control standard. Minimum variance controllers require perfect models of disturbances (generally a first or second order model), but the vast majority of industrial controllers implicitly contain a step-disturbance model only. If the actual process disturbance is first or second order (which is often the case), the MVC Index can indicate that the controller is behaving poorly, even if there is no *process* model mismatch or move suppression. In this case, the MVC Index is indicating that there is a *disturbance* model mismatch, but there is little the control engineer can or should do to rectify this.

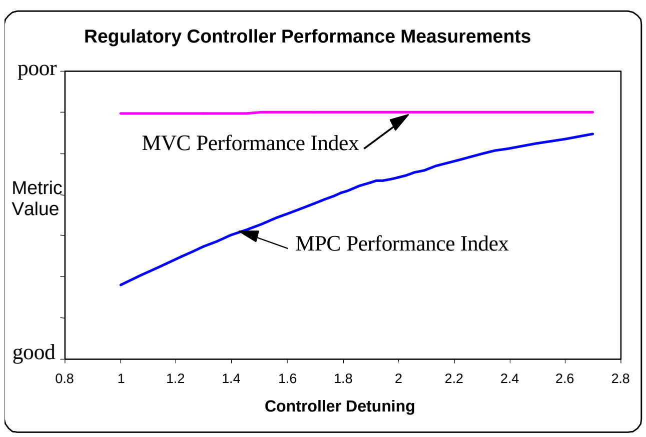

Figure 3: Performance indices for regulatory controllers subject to a 1st order disturbance. The Minimum Variance Controller Index can be very insensitive to controller tuning; the Model Predictive Controller Index is unable to detect the effects of detuning.

The MVC Index can then fail at Requirements 1 and 6 – it will give varying results depending on the actual disturbance that enters the process, and it is not always sensitive to *process* model mismatch and controller detuning (see Figure 3). Indeed, Miller and Desborough [1] found that the minimum variance performance index fails to give consistent results in 40% of the loops they examined.

Furthermore, the minimum variance performance index represents an often unachievably high standard – many industrial practitioners recommend either modifying the standard or increasing the apparent deadtime[*] in order to obtain more reasonable results[9]. An example of this may be seen in Table 1: the recommended prediction horizons are considerably longer than actual process deadtimes. Moreover, these prediction horizons are somewhat arbitrary measures, and thus do not meet Requirement 5.

Best Possible MPC Controller

A more industrially relevant index is one that accounts for the simplified step-disturbance model of Model Predictive Controllers (i.e., DMC). This index compares the current variance to the variance which would occur if an MPC were applied that had no process model error or move suppressions. This index explicitly addresses the fact that disturbance model in MPC's may differ from the true disturbance. Its advantages are that it is much more sensitive to the things the control engineer can change - process model mismatch and controller detuning, and it is much less sensitive to the underlying disturbance spectrums. In addition, the index may be formulated in the same way as the MVC Index:

$$h^{MPC} = 1 - \frac{s^{2}_{MPC}}{s^{2}_{ACT}} \qquad (2$$

As shown in Figure 4, the MPC response generally is not as tight as a Minimum Variance Controller, so the MPC performance index represents a more attainable standard. Thus it is not necessary to "tune" the performance index to attain reasonable results.

With some mild assumptions, the user need only specify the same process parameter as for the Minimum Variance Controller Index – just the process deadtime. In other words, the MPC Index requires the same input, and can be interpreted the same, as the MVC Index, but results in much more useful results. This formulation is available in some performance assessment packages[4].

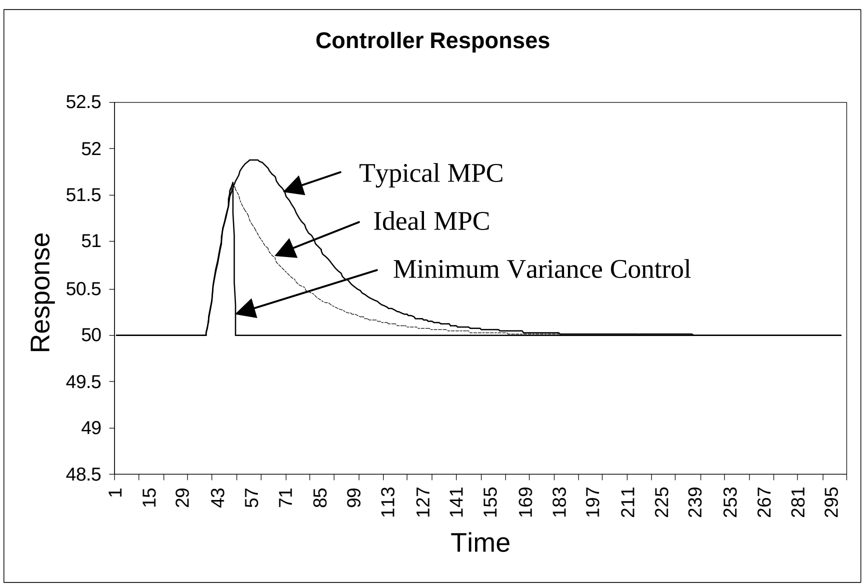

Figure 4: Response for a 1st order disturbance. A Minimum Variance Controller is able to predict the full disturbance, and is therefore able to fully counteract it immediately after the deadtime. In contrast, a Model Predictive Controller always underestimates the disturbance, and there requires a considerable length of time to remove it.

[*] Increasing the apparent deadtime improves the index as it increases the estimated minimum variance, thus making the current controller quality look better in comparison

Best Possible PID Control

While the previous index represents an achievable standard, some 97%[1] of the loops in the process industry use PID algorithms, and a more reasonable standard would be the measurement variance if the best possible PID controller were applied to the process. The main difference between a PID controller and an MPC controller is the lack of deadtime compensation of PID controllers – process with long deadtimes can never be controlled as effectively with a PID controller as they can with an MPC.

Again, it is possible to define a controller index similar to the ones above as:

$$h^{PID} = 1 - \frac{s^2_{PID}}{s^2_{ACT}} \qquad (3$$

Here s^2_{PID} is the variance of the process that would occur if the best possible PID controller were applied. As before, under certain mild assumptions, the process deadtime is the only parameter necessary to calculate this index. Again, this index is available in some controller performance assessment packages[4].

If the PID performance index indicates good tuning, but the user is unsatisfied with the output quality, then the MPC controller performance index may be calculated to determine whether there would be any improvement if a model predictive controller was applied to the process (Figure 8).

Open-Loop

Obviously, the variance of the open-loop process is a very perfunctory standard. Nonetheless, it is somewhat surprising that many control loops do not meet even this criterion; one study[10] has found that up to 80% of controllers lead to an *increase* in variance over open-loop. The open-loop standard is however useful for determining whether *any* control should be applied – the usual benefits of better control have to be balanced against costs for measurement, control valve, and installation and tuning.

Surge Vessel Level Control

Level control of surge vessels, as shown in Figure 5, is fundamentally different from control of other processes, and therefore requires different assessment techniques. As opposed to the normal control objectives of keeping a measurement at setpoint, the purpose of a surge vessel level control is to dampen the changes in controlled flow while keeping the liquid level in the vessel between limits. However, this is contrary to the implicit design of PID controllers – they are, after all, designed to keep a measurement at setpoint. Not surprisingly, it is common to see tight level controllers (the vessel is then in effect acting as a pipe), with poor flow variation damping, and subsequent increased variation in product quality.

One proposed method of evaluating level controller performance is to compare the level fluctuation to the maximum possible level fluctuation, with the idea that the more the level fluctuates the more the disturbance is attenuated. This is flawed reasoning, as the level does not need to fluctuate to constraints to dampen all disturbances. For instance, consider an inlet flow that oscillates in a sinusoidal fashion. If the amplitude is small enough, the outlet flow may remain constant (i.e., perfect damping), with the vessel level oscillating at the same frequency as the inlet flow. For small enough amplitude, the level will not hit upper and lower bounds, and any controller performance index that was based on the maximum fluctuation alone would erroneously indicate sub-optimal control.

A much more reliable index is based on a comparison of the current level fluctuation to one that would be obtained if an optimal level controller were applied on the process. Optimal level controllers are in effect non-linear constrained optimizers, where the objective is to minimize the maximum rate of change of the outlet flow subject to the level remaining within bounds. Again, the performance index may be calculated in the standard fashion:

$$h^{LVL} = 1 - \frac{s^2_{OPT}}{s^2_{ACT}} \quad (8$$

As before, s^2_{OPT} is the variation in controlled flow if an optimal level controller was applied to the system, and s^2_{ACT} is the actual variation in controlled flow.

Again, this index may be calculated using only closed loop data (in this case the level and either the inlet or outlet flows), and the vessel dimensions. This index is also found in some performance assessment packages[4].

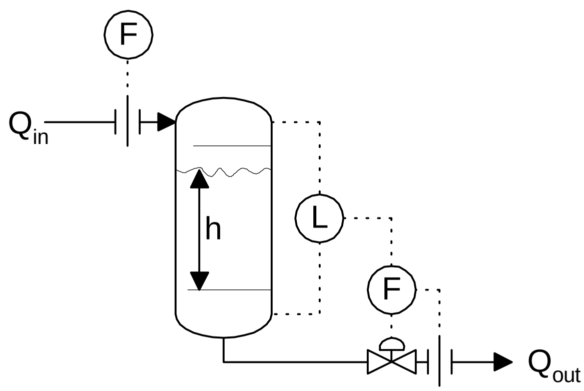

Figure 5: The purpose of a surge vessel is to transfer variation in uncontrolled flow to variation in level. Performance indices for this situation need to address the unique objectives of the controller.

Multivariable Control

There is considerable incentive for extending the univariate controller performance measures to the multivariable case, both for maintaining these controllers and evaluating their economics, but the solutions thus far have the are difficult to implement. One disadvantage is that the user must specify or determine the plant interactor matrix, which is a function of the all plant dynamics (gains and time constants in addition to deadtimes). More importantly, the algorithms do not account for constraints, and controlling against constraints is usually the *raison d' être* of multivariable model predictive controllers[11,12].

To date, commercial packages do not contain algorithms for assessing multivariable controllers. However, it is possible to use the single loop techniques on each of the outputs of a multivariable controller, although the results will be somewhat biased. Following is an examination of each aspect of multivariable model predictive control, and the applicability of single-loop performance assessment to the multivariable case.

<u>Multiple Inputs Affecting Multiple Outputs</u>

Multivariable controllers in general use all the inputs to control all the outputs, but the control engineer is mainly interested in how well *each* output is controlled to its setpoint, and this is exactly what single loop performance indices measure. The results will however be biased somewhat as more than one input can affect each output, with the amount of biasing dependent on the amount of process coupling. Fortunately, most process are not tightly coupled, as it is usually the case that each major output is controlled largely by one input.

Optimization Layer

Most MPC packages contain an optimization layer (usually an LP) which sends setpoints down to a regulatory layer. This is not a problem for performance assessment packages, as it does not matter if the setpoint for the regulatory controller comes from an operator or a computer program.

Performance assessment techniques do not in general determine whether the LP is generating an economically optimal solution. However, as the LP and the regulatory controller generally use the same model steady-state gains, a good performance index indicates that both the regulatory controller and LP are using accurate models.

Input (or manipulated variable) Constraints

For single loop control, care must be exercised when evaluating the performance to check that the input was not saturated, as the controller can not be expected to control well in this circumstance. This same caution applies to multivariable control as well, particularly in checking the dominant input(s). In general, multivariable model predictive controllers should handle input saturation better than single loop controllers should, as they should reduce other variables (i.e., feed) if the main quality handles (i.e., reflux or fuel gas) become saturated.

Output (or controlled variable) Constraints

The upper and lower bounds of output constraints are translated into setpoints by the optimization layer, and therefore present no difficulty for performance assessment techniques. In contrast to the input constraint case, it is preferable if these values are at a constraint. This is because the controller will allow the process to float between constraints if possible, which does not stress the capabilities of the controller.

Input and Output Tuning Parameters

Model Predictive Controllers clearly have different tuning handles than PID controllers, but this is irrelevant to performance assessment algorithms. All these parameters tune (or detune) the controller; performance algorithms tell if the controller is detuned, but they are not designed to provide information on *how* it is detuned.

Economic Benefits of Tighter Control

Tighter control has two direct economic benefits. First, reducing the variance allows the process mean to be moved closer to a constraint while still maintaining product quality specifications. The second benefit is due to reducing the *irreversible losses of process variation*[13]. In both cases, performance indices are invaluable in evaluating the economic incentives of improved control.

To illustrate the first benefit, consider the effect of lowering the standard deviation from s_{ACT} to s_{OPT} (see Figure 6). If the current setpoint is set so that the process meets a quality constraint 99.5% of the time (i.e., a 3-sigma limit), then reducing the process variance would allow changing the setpoint by:

$$3*\left(s_{ACT} - s_{OPT}\right) \qquad (4$$

The $\left(s_{ACT} - s_{OPT}\right)$ term may be calculated from a modified performance assessment equation of the form:

$$h' = 1 - \frac{s_{OPT}}{s_{ACT}} \qquad (5$$

Here s_{OPT} is the standard deviation of the optimal controller as from equations 1,2, or 3, and s_{ACT} is the current actual process standard deviation. Note that this modified performance index uses the ratio of the standard deviations rather than the ratio of the variances. Combining equations 4 and 5, and multiplying by the setpoint marginal profit gives the following expression for the benefit of moving the setpoint the maximal amount:

$$J^{\nabla SP} = 3 * h' * s_{ACT} * P^{\nabla SP} \qquad (6$$

Here the factor $P^{\nabla SP}$ represents the profit of moving the setpoint a unit amount.   All the above terms may be calculated directly from current closed-loop data and plant economics.

The irreversible loss of process variation economic cost may also be calculated using a performance assessment measure. Again, the maximum reduction in variance that can be accomplished by control, $s_{ACT} - s_{OPT}$, is found by rearranging Equation 5, and the economic benefit is simply:

$$J^{IRL} = h' * s_{ACT} * P^{IRL} \qquad (7$$

The factor P^{IRL} represents the unit cost of irreversible process variations, and may be determined from the plant economics and correlations relating the plant yield to the quality.

Note that all the above factors may be calculated using closed-loop data, plant economics, and process simulators (i.e., no special plant tests or measurements are required). In practice, it is also necessary to multiply the optimal standard deviation s_{OPT} some factor to account for the fact that real controllers never achieve the performance of ideal ones. Nonetheless, the performance index is a crucial (and generally overlooked) part of the assessment – if the current control is already good, the performance index will be small, and applying advanced control will not result in significant economic gain.

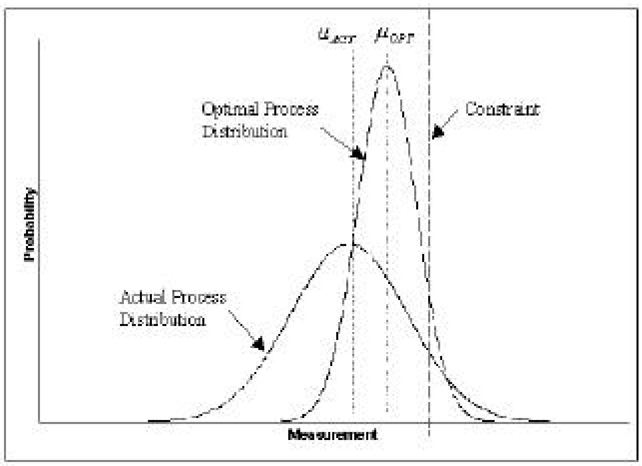

Figure 6: Reducing Variance has 2 economic benefits: the setpoint may be moved closer to an operating constraint, and the irreversible loss from process variance will be reduced.

Diagnostic Tests

While controller assessment techniques can determine which controllers need attention, required also are techniques for determining why a controller is behaving poorly, and what should be done to rectify the situation. Analysis techniques are available for 2 common situations:

1. Detecting valve stiction. Differentiating whether the cause of oscillation is valve stiction or improper tuning can be made by examining the phase shift in the process input/output cross-correlation plot[14]. A phase shift of 90 degrees indicates that the oscillation is likely caused by valve stiction. No phase shift indicates that the oscillation is caused by improper tuning[15].
2. Locating the source of oscillation in a plant. In an interactive plant, it can be difficult to determine if a controller is oscillating because of some flaw within the control loop or because it is responding to external disturbances. One technique for determining the oscillation source relies on the fact that

oscillations are caused by some non-linearity, and the plant tends to linearize responses. It follows therefore that the most non-linear response is the likely source of oscillations[16].

Clearly both these techniques rely on secondary information to draw conclusions about the plant. As such, they rely on a multitude of assumptions, which are not always fully realized, and they are sensitive to poor data sets. There is always a limit on the amount of information that may be obtained from happenstance data, and at some point the user must open the loop and perform plant tests.

Hardware Requirements

Generally, the calculation requirements for performance assessment are not onerous, and may easily be performed on current generation PC's, if not in the DCS system itself. Of more concern is the data-sampling rate. In theory, the data should be sampled as fast as the control interval. In practice, a sampling period of approximately 1/3 of the process open-loop time constant is sufficient.

This fast sampling presents a problem for many DCS systems. Generally, these systems were not designed for fast sampling of multitudes of loops – the fastest sampling is generally 10-15 seconds; adequate for temperature loops, but not for fast loops such as flow control. To overcome this limitation, specialized data collection programs are available[4,5], and in some cases the performance assessment calculations are performed in the fieldbus device itself[17].

Another concern, particularly with data historians, is data compression. In general, performance assessment calculations are very sensitive to data compression (that is, if information is lost when the data is uncompressed). As a general rule, compression should be removed from any points that will be used for performance assessment.

Conclusions

Controller performance assessment should be a part of the control engineer's standard routine assignments, as changing process conditions and equipment degradation result in continual varying of controller effectiveness. Controller performance metrics provide management with indications of controller quality, indicate areas of the plant which require attention, are invaluable for estimating benefits of control, and provide a bound on achievable control. Techniques for assessing model predictive controller performance have been developed and are available in software packages. These techniques do not require plant tests and are very easy to use – some of these packages can be scheduled to automatically calculate the performance of a large number of plant controllers and report the results for quick analysis.

References

1. Miller, R., and Desborough, L., *Increasing Customer Value of Industrial Control Performance Monitoring – Honeywell's Experience*, CPC VI Meeting, 2001.
2. Billkowski, W., *Texas A&M Instrumentation Symposium for the Process Industries*, January 1996.
3. Harris, T., and Desborough, L., *Performance Assessment Measures for Univariate Feedback Control*, Cdn J. of Chem. Eng., **70**, 1992, pp. 1186-1197.
4. Performance Assessment Toolkit, www.ControlArtsInc.com, 2000.
5. Loop Scout, www.IAC.Honeywell.com, 2000.
6. ProcessDoc, www.Matrikon.com, 2000.
7. AspenWatch, www.AspenTech.com, 2000.
8. Thornhill, N.F., Oettiger, M., and Fedenczuk, P., *Refinery-wide control loop performance assessment*, J. of Process Control, **9**, 1999, pp. 109-124
9. Kozub, D.J., *Monitoring and diagnosis of Chemical Processes with automated process control*, CPC V Meeting, 1996.

10. Spencer, M.A., and Elliot, R.M., *Improving instrumentation and control systems performance*, Petroleum Technology Quarterly, Winter 1997/98 pp. 93-97.
11. Harris, T.J., Bourdreau, F., and MacGregor, J.F., *Performance Assessment of Multivariable Feedback Controllers*, Automatica, **11**, pp. 1505-1518.
12. Huang, B., Shah, S.L., and Kwok, E.K., *Good, Bad or Optimal? Performance Assessment of Multivariable Processes*, Automatica, **33**, pp. 1175-1183.
13. Peterson, T., *Irreversible Losses of Process Variation*, www.ControlConsulting.com, 2000.
14. Box, G.E., and Jenkins, G.M., *Time Series Analysis forecasting and control*, Holden Day, 1976.
15. Horch, A., Extensions of a Static Performance Index and Detection of Static Friction in Process Controllers, Ph.D. Thesis, Royal Institute of Technology, 1998.
16. Thornhill, N.F., Shah, S.L., and Huang, B. *Detection of Distributed Oscillations and Root-Cause Diagnosis*, CHEMFAS, 2001
17. Blevins, T., and Ochoa, D., *Statistical analysis rides the bus*, InTech, March 2000.

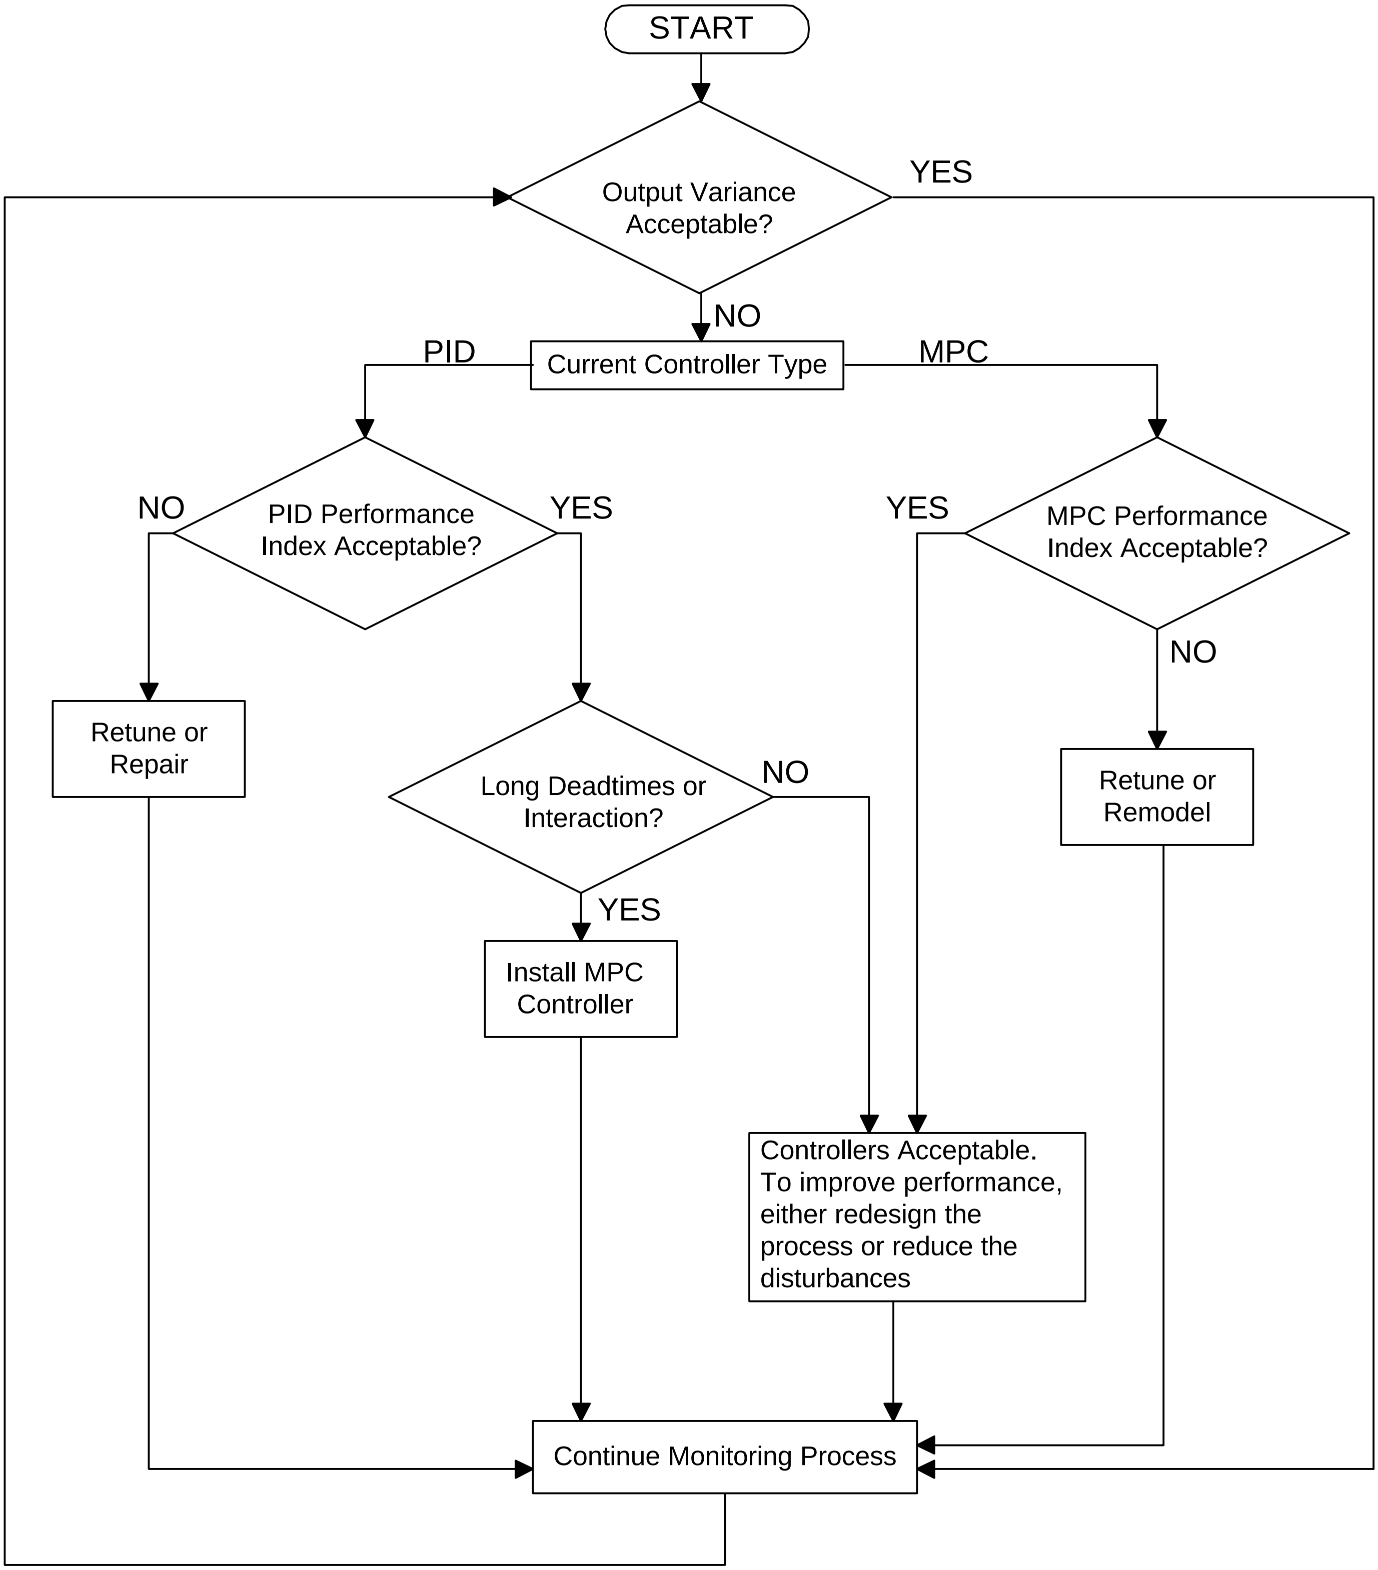

Figure 8: Performance Indices are just part of controller evaluation and maintenance. They do, however, provide key information for analysis and further investigation.