Journal of Process Control 13 (2003) 155–168

JOURNAL OF PROCESS CONTROL

www.elsevier.com/locate/jprocont

A practical approach for large-scale controller performance assessment, diagnosis, and improvement

Michael A. Paulonis, John W. Cox*

Eastman Chemical Company, Kingsport, TN 37662, USA

Received 18 January 2002; received in revised form 14 March 2002; accepted 18 March 2002

Abstract

Eastman Chemical Company has developed a large-scale controller performance assessment system spanning over 14,000 PID controllers in 40 plants at 9 sites worldwide. Controllers can be sorted in order of performance to quickly identify which need attention. Performance history is available to track improvement or degradation in performance for a single controller or an entire plant. Diagnostic aids are available for both novices and experts to substantially reduce troubleshooting time. E-mail reports are automatically generated and sent to subscribers to keep them informed of relevant changes with minimal investment of their time. The user interface is web-based to allow universal access to any employee. Use of the system has dramatically increased controller optimization productivity. © 2002 Elsevier Science Ltd. All rights reserved.

Keywords: PID controller; Performance assessment; Troubleshooting; Diagnosis

1. Introduction

The late 1990s were a time of change at Eastman Chemical Company. Product prices were dropping due to ever-present competition as well as world overcapacity built with strong profits from the last industry upcycle. Raw material prices were increasing due to increasing cost of petroleum feedstocks. Investors were much more interested in chasing the rising spiral of technology stocks rather than sectors like basic materials. The joint arrival of these market forces triggered development of a "perfect storm" for chemicals manufacturing. The business strategies put in place to deal with the situation had a very common theme, "do more with less". Higher quality and production rate was demanded with fewer people, less energy, less raw materials, and importantly, lower capital investment.

For the Advanced Controls Technology group at Eastman, these changes meant a dramatic reduction in control strategy development for new plants. Fortunately, the demand for process control work in existing plants was higher than ever, as production managers were eager to gain the benefits of control improvements that could be delivered with little or no cash out the door. The renewed focus on existing operations revealed numerous opportunities for process control improvement. During these improvement projects, it was very typical to find poorly performing control loops. The most common problem was oscillation as a result of valve hysteresis or stiction. Indeed, the problem was alarmingly common. It quickly became clear that looking for oscillating or poorly performing control loops one-by-one could consume the full resources of the control group and limit investigation to a fraction of all processes within the company. A vision began to form of a controller performance assessment tool that would enable efficient detection and diagnosis of problems in the many thousands of control loops in service worldwide at Eastman.

Studies of controller performance assessment algorithms began to appear in the early 1990s after the work of Harris [1], in which closed-loop time-series controller data were analyzed to benchmark controller performance against minimum variance control. Desborough and Harris [2], Kozub and Garcia [3], and Stanfelj et al. [4] extended this concept with a performance index. Comprehensive reviews of performance index algorithms are provided by Qin [5] as well as Huang and Shah [6].

* Corresponding author. Tel.: +1-423-224-0409; fax: +1-423-224-0453.

E-mail address: jwcox@eastman.com (J.W. Cox).

0959-1524/02/$ - see front matter
PII: S0959-1524(02)00018-5

A performance assessment system making use of a performance index and other analyses of closed-loop process data was reported by Jofriet et al. [7]. However, this system required the deadtime of each loop to be specified by the user, creating a significant burden. An open loop bump test and analysis for each controller was suggested to determine this parameter. Kozub [8] describes an alternative industrial application where a deadtime specification is not necessarily required, but a specification of the desired autocorrelation function of the controller error is required. In some sense, this variation requires even more effort per loop to configure, but it also makes the benchmark performance more relevant to industrial operation. Thornhill et al. [9,10] introduced a significant advance in which default parameters for the performance index algorithm were shown to be useful for various generic categories of control loops. This work lowered the barrier to large-scale implementation of a performance index.

Thornhill and Hagglund [11] extended performance assessment in a different direction with a method to detect and diagnose oscillations in control loops. Together with a performance index, this paper presented an outline of a complete approach to controller assessment.

Miller et al. [12] described a comprehensive system for controller performance assessment developed by Honeywell Hi-Spec Solutions (Thousand Oaks, CA). This system is now offered to the process industries as a service called Loop Scout™. Other industrial applications of performance assessment as well as the challenge of developing a large-scale, automated system are discussed by Harris et al. [13].

Eastman considered using the Loop Scout™ service in 1999, but chose not to for three primary reasons:

- Automated data collection was limited to Honeywell control systems.
- Substantial amounts of process data would have to be sent to Honeywell, requiring complicated approvals.
- The cost to assess loops worldwide was prohibitively high given the emphasis on reducing business expenses.

At that time, we refined our vision for large-scale controller performance assessment at Eastman. Features we felt were necessary included:

- Interfaces to all Eastman's DCS systems, PI systems, and our in-house advanced control platform.
- A friendly user interface, providing accessibility to all company personnel.
- Minimal client and server configuration.
- Periodic assessment of loops and retention of performance history.
- Ranking loops by performance.
- Preliminary problem diagnosis for poorly performing loops.
- Reports which could be generated interactively or predefined reports which could be e-mailed to users on a schedule.
- Ability for users to add comments and documents to the system and link them to loops or loop tuning changes.

A search of the marketplace in 1999 did not reveal any system that we could purchase that included most or all of these features. As a result, we chose to develop one of our own.

2. System architecture

A schematic diagram of our system is shown in Fig. 1. Selected portions of the system are numbered in the figure and are discussed in more detail below.

2.1. Block 1—data interface

Closed-loop controller data is automatically captured directly from Eastman control systems at regular intervals. We have developed data interfaces for most of the control systems in use at Eastman. We considered keeping the data interface simpler by using data from Eastman's company-wide system of PI historians, but the typical data compression was far too great for the assessment analyses. As OPC becomes more prevalent, a universal data interface may become practical. Data capture is handled in a distributed manner, with processing power at each control system being used to collect the data and push it to the assessment system. This approach helps the system scale nicely to many thousands of loops. An interface to PI and to Eastman's advanced control platform is also available for assessments on demand. Controller tuning is captured on a daily basis and pushed to the assessment system.

2.2. Block 2—computation engine

Computational software of our own design does performance assessment on controller data shortly after capture. The full assessment includes:

- Time-series trends (full and zoomed).
- Descriptive statistics for controller setpoint, measurement, output, and error.
- Setpoint crossings (number of times measurement crosses setpoint).
- Univariate and bivariate density of error and output.
- Closed-loop impulse response.

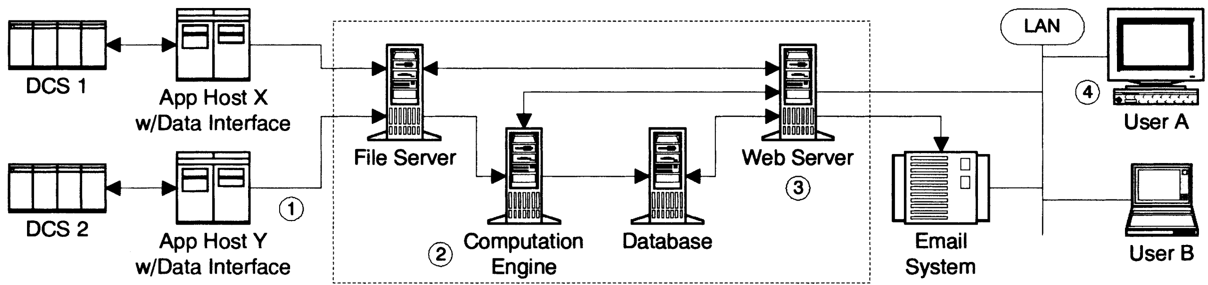

Fig. 1. Schematic diagram of the controller performance assessment system.

- Harris-style extended horizon performance index.
- Power spectrum.
- Oscillation detection and characterization.
- Controller idle index [14].
- Cross-correlation between error and output.
- Diagnosis information in text form.
- A single composite score ranking the loop performance.

Numeric and text information from the assessments is written to a database. The computation engine is also available for assessments on demand. We can quickly and easily modify assessment analyses or add new ones since we are in control of the source code.

2.3. Block 3—web server

A web server is at the heart of the assessment system. A web-based user interface was chosen to eliminate the need for distribution and installation of client software and also because users are very familiar with web pages, reducing the learning curve for use of the system. Very capable software is available for web-database integration and web server-side programming, making a web server a strong application platform.

2.4. Block 4—user interface

The user interacts with the system via a web browser. Hyperlinked web pages allow the user to navigate the system. Web forms are used for user input. Web pages with text, tables, images, and applets are used to present the system output. In addition, users can subscribe to a variety of daily or monthly reports for process areas in which they have interest. These reports are sent to the user via HTML e-mail.

2.5. Scaling the system

The dotted box in Fig. 1 shows separate computers performing individual functions within the assessment system. This configuration is shown for illustration only and is flexible. For a small system, all functions could be performed in a single computer or some number less than four as shown. For a very large system, individual functions such as web serving could take place in multiple computers with load-sharing capability.

3. Features and capabilities

The assessment system has many features designed to help users identify controllers needing attention, diagnose controller problems, and track controller performance over time. These features include e-mail reports, detailed performance assessments, and detailed controller information (includes tuning history and analysis, performance history and problem diagnosis, loop configuration information, etc.). E-mail reports for individual process areas are sent daily or monthly, based on the report type, to interested users who subscribe. To illustrate the capabilities of the assessment system, consider the following descriptions of how these features can be used.

3.1. Daily tuning change report

A daily tuning change report is sent to a subscriber's e-mail inbox once per day and lists controllers having tuning changes during the past day. With this report, users are quickly alerted to controller tuning changes in their area. This is important because tuning changes are often indications of significant problems, and communication of tuning changes to all appropriate personnel by other means has typically been poor, at least at Eastman. Tuning changes are often the first attempt to remedy controller problems because they are easy and cheap. However, in most cases, tuning changes are not the best response to the event that made good controller performance become bad. The daily tuning change report is a reliable and timely way to alert users about potential problems in their process area.

The example in the next section illustrates the use of an actual daily tuning change report. A report excerpt is shown in Fig. 2. For the process area in question, loop LCDC801 had a tuning change in the last day.

3.2. Tuning and performance history

For individual loops, users can pull up a web page that contains tuning history, tuning analysis, performance history, performance analysis, and loop configuration information. The web page can be requested interactively through the performance assessment web site or it can be requested from a hyperlink that appears in HTML e-mail or other documents capable of including links. In the case of the tuning change report in Fig. 2, the link is the tag name. To further investigate the LCDC801 tuning change, the user simply clicks on the tag name while viewing the e-mail report. A web browser pops up and loads a web page full of tuning, performance, and configuration information. An excerpt from this web page is shown in Fig. 3. The following abbreviations are used in the performance history table in Fig. 3: PV—process variable (measurement), OP—controller output, Err—controller error, Std Dev—standard deviation, CLPA—closed-loop performance assessment (a Harris-style performance index), Osc Diag—diagnosis from oscillation analysis, SP Cross—count of times the measurement crosses the setpoint.

The tuning history table shows that the previous tuning had been in place for over 2 years. It seems unlikely that the tuning would have persisted that long if it were bad. The tuning analysis does not indicate that the tuning was changed to unreasonable values. It was simply a small change, probably an experiment.

A glance further down the page at the performance history table reveals that perhaps the tuning is not the problem after all. Performance measures are available weekly, except for the time between 25 May and 19 August, when the process was shut down. There was a dramatic change in the controller performance when the process came back up. The CLPA statistic jumped towards 1, indicating a substantial departure from efficient control. The oscillation detection analysis also began to signal a hardware problem (“hdw” in the “Osc Diag” column), where before it had indicated no problem or some oscillating disturbances.

However, the table also shows that the mean error for the loop has not increased substantially. It is a little higher than previous, but still averaging about 1% of the level span. This level performance is adequate for the reactor. Thus, for this case, the investigation reveals that the loop might have some problems, but they do not appear to be serious at this point. It would be appropriate to follow up on the LCDC801 loop at a convenient time in the future to find out if the valve may need some maintenance.

The following loops had tuning changes in the past day or are new to the database in the past day. You can click on the tagname to get more information. If you made a tuning change, please click the tagname, claim the tuning in the tuning history table, and add a few notes about why you changed it.

Tagname	Description	Gain %/%	Integral rpt/sec	Derivative sec	Filter smooth
LCDC801	DR1 LEVEL CONTROL	0.50	0.00100	0.00	0.00

Fig. 2. Excerpt from a daily tuning change report delivered by HTML e-mail.

3.3. Performance detail report

The performance detail report is another available e-mail report that is extremely useful to plant personnel interested in minimizing process variability and optimizing performance of critical loops. This monthly report contains summary statistics, loop performance rankings, and a section highlighting changes in loop performance for an entire plant or process area. While this report can be used in many different ways, its strength is in identifying the worst performing loops out of the many loops (hundreds to thousands) in the process area. Technical personnel familiar with the process area can pick out loops that need attention (those rated as poor or fair performers which are critical to the process operation) in a matter of minutes. Once a problem loop has been identified, historical information useful for problem diagnosis is easily accessible, again with a hyperlink to the web-based loop information page.

Consider the following example to illustrate the use of an actual performance detail report. A very small excerpt from the main section of the report (which lists loops in ascending order of performance) is shown in Fig. 4. The following abbreviations are used in Fig. 4 in addition to those already listed for Fig. 3: Perf. Class—performance classification, Crit—criticality, Pct Osc Hdw Diag—percentage of assessments where the oscillation analysis results in a diagnosis of hardware problems, Num Assess—number of assessments.

As a user scans the report, attention is drawn to loop TCGH606. This loop is noted because it is a reactor temperature loop; its performance is critical to the process performance. A loop of this importance should be doing better than “fair”. The report currently shows the criticality of this loop to be “Avg” (average). All loops default to average criticality until someone familiar with the process area specifies criticality as Low, Average, High, or Vital. In this case, this loop would be considered to have vital criticality and can be changed by the user for future reports. The criticality can be used for reference in scanning or can be used as search criteria for interactive database queries.

To get a better handle on what “fair” means, a novice user could click on the column heading above “fair” and get a pop-up help window in their web browser. A portion of the help window is shown in Fig. 5. According to the help, the score of 35 is on the low end of fair and is almost considered poor (40+). The user concludes that this loop definitely needs attention.

To investigate further, the user clicks on the TCGH606 tag name to view the detailed loop information web page. The performance history table and performance analysis from the page are shown in Fig. 6. Most of the abbreviations present in Fig. 6 have already been defined for Figs. 3 and 4. There is one new abbreviation used in the performance diagnostics table in Fig. 6, Spec. Hdw Diag (%). This is the percentage of assessments where the spectral analysis results in a diagnosis of hardware problems.

The performance history table shows a fairly consistent preliminary diagnosis of hardware problems (in

DR1 LEVEL CONTROL (LCDC801)

Area: TED_B270

Tuning History

To reorder the table, click on the desired table category.

Date	Gain %/%	Integral rpt/sec	Derivative sec	Filter smooth	By Whom	Notes
11-Oct-2001	0.50	0.00100	0.00	0.00	Unknown-Claim it!	none
09-May-1999	0.80	0.00100	0.00	0.00	Unknown-Claim it!	none
14-Jan-1999	0.80	0.00060	0.00	0.00	Unknown-Claim it!	none

Tuning Analysis

The current tuning parameter values fall within typical bounds for this type of controller. This does not necessarily mean that the tuning is optimal, just that the tuning meets some generic tuning criteria.

Performance History

Click on a table category to see a data plot. Any values in red below were flagged as an indication of potentially poor performance. Note: the most recent 10 sets of performance assessment data have been returned. You can also request the last 10, 25, 50, 75, 100 sets of performance data or view the complete performance history for this loop

Date	Mean PV	Mean OP	Mean Err (%)	OP Std Dev (%)	CLPA	Osc Diag	SP Cross
04-Oct-2001	60.01	33.96	0.98	1.64	**0.80**	hdw	103
27-Sep-2001	60.01	35.39	1.18	1.92	**0.76**	hdw	**67**
20-Sep-2001	60.01	37.23	0.76	1.24	**0.68**	hdw	**90**
26-Aug-2001	56.00	33.42	1.29	1.92	**0.83**	hdw	**63**
19-Aug-2001	55.00	32.71	0.64	1.11	0.43	hdw	281
25-May-2001	55.00	32.40	0.86	1.30	0.17	none	828
18-May-2001	54.98	34.27	0.88	1.44	0.15	dstb	823
11-May-2001	55.99	30.93	0.86	1.29	0.17	dstb	859
03-May-2001	55.99	30.98	0.81	1.42	0.12	none	923
26-Apr-2001	50.98	32.14	0.98	1.43	0.17	dstb	841

Fig. 3. Excerpt from a "complete loop information" report.

Tagname	Description	Perf. Class	Score	Crit	Mean Err	OP Std Dev	Mean OP	Mean CLPA	Pct Osc Hdw Diag	Mean SP Cross	Num Assess
FCJE747	JR3T EG SPRAY DILUTION	fair	36.3	Avg	42.01	2.32	37	0.03	0.0	2108	5
TCTE600B	T-NORTH HOT OIL LOOP	fair	35.9	Avg	0.20	9.75	25	0.88	100.0	144	3
TCQ60418B	QN04 TOP REACTOR BED TEMP	fair	35.8	Avg	5.63	8.72	200	0.68	0.0	26	2
TCGH606	GR1 DOWTHERM TEMP CONTROL	fair	35.6	Avg	0.14	9.19	65	0.84	71.4	117	7

Fig. 4. Excerpt from a performance detail report.

Fair - Loops in this category are not performing up to potential and should be improved. Control is probably being maintained in a broad sense, but deviation from setpoint is likely to be degrading process performance. The measurement may be cycling. These loops should be investigated further. Many times, the problem is not difficult to find and improvement can be obtained without a lot of effort. Occasionally, level loops with a flow smoothing objective that are performing adequately end up in this category.
(Score range 20-40)

Fig. 5. Pop-up help for performance classification.

the "Osc Diag" column). The text analysis of the performance history helps explain the tabular results for a novice user. The contents of the detailed loop information are suggesting a hardware problem.

3.4. Detailed performance assessment

The user can choose to view a detailed performance assessment of closed-loop operating data for TCGH606 by clicking on the link shown at the bottom of Fig. 6. This produces a detailed performance assessment (shown in Fig. 7) using the captured dataset from the latest automated assessment on 9 October. The main part of the detailed assessment includes: (1) a summary text section containing overall performance and problem diagnosis information, (2) time-series plots of the closed-loop setpoint, measurement, and output data, (3) numerical statistics, and (4) graphical analyses. To assist novice users in understanding the different graphs and statistics, pop-up help is available by clicking on portions of the detailed assessment. The pop-up help includes descriptions of the item and, for graphical analyses, shows example patterns and their interpretation. Graphical patterns often provide additional diagnosis information, although the user is currently responsible for matching the pattern to a diagnosis.

For the assessment shown in Fig. 7, several inferences can be made:

- The text summary section at the top pretty much confirms the conclusions formed from the loop information page (Fig. 6). The performance could probably be greatly improved and the problem seems to be hardware in the loop (only one of the tests suggests a problem other than hardware).
- The top right graph shows the joint probability distribution between the controller error and the output. The graph is shown in color for users and color is important for detecting patterns. The grayscale representation in Fig. 7 makes the pattern difficult to see. Users would see warm colors (yellow, orange, red) in a ring around the center and cool colors (blue, green) in the center and outside. This ring pattern is a strong indicator of hardware problems in self-regulating loops such as the temperature loop in question.
- The y–y plot below the large time-series plots is a zoomed section of the time-series. The computer chooses an "interesting" section of the data. The pattern in the zoomed plot suggests a hysteresis problem in the hardware. The controller output needs to move quite a bit (after changing direction) before the measurement (PV) starts to move in the same direction. Note that in the grayscale representation of the plot, the PV is the darkest line, the controller output (OP) is the lightest line, and the setpoint (SP) is the remaining line.
- The diagnosis of a disturbance problem by the spectral analysis is a result of the analysis not finding harmonic peaks in the spectrum. Typically, hardware problems will result in process nonlinearities that produce spectral harmonics. In some sense, there is a disturbance in this loop in the form of cyclic setpoint changes, but this is not the dominant feature.

The prevalence of hardware problem diagnoses from the performance history and the detailed assessment are convincing evidence that TCGH606 does indeed have a hardware problem. The fact that the spectral analysis

Date	Mean PV	Mean OP	Mean Err (%)	OP Std Dev (%)	CLPA	Osc Diag	Spec	SP Cross
09-Oct-2001	221.51	57.25	0.14	8.01	**0.77**	hdw	dstb	138
04-Oct-2001	224.44	65.25	0.12	4.53	**0.82**	hdw	tune	224
16-Sep-2001	221.01	57.80	0.14	4.72	**0.82**	hdw	hdw	123
07-Sep-2001	220.69	55.57	0.15	4.93	**0.83**	hdw	hdw	145
31-Aug-2001	221.52	57.55	0.13	4.58	**0.82**	hdw	hdw	165
23-Aug-2001	223.50	62.78	0.16	5.07	**0.76**	hdw	hdw	140
16-Aug-2001	231.82	78.39	0.13	3.98	**0.84**	hdw	hdw	143
08-Aug-2001	222.09	74.75	1.11	31.87	**0.86**	dstb	dstb	**86**
15-Jul-2001	233.30	4.02	1.54	306.89	**0.98**	dstb	dstb	**16**
01-Jul-2001	229.95	87.46	0.16	4.53	**0.84**	hdw	dstb	124

Plot performance data

Performance Analysis from 01-Jul-2001 to 09-Oct-2001

Keep in mind that the analysis below is based on an average of multiple performance assessments taken over a given time window. Current performance may be significantly different if process or other conditions have changed. Also, the analyses are based on general criteria and are not always accurate. Use this information as one of many tools to help you find and fix significant loop problems.

Basic Performance Info for TCGH606		
Parameter	**Value**	**Analysis**
No. Assess.	10	The relative number of assessments can impact the performance analysis results (e.g., performance results from just a few assessments may not accurately represent overall loop performance).
Lo/Avg./Hi PV	220.69 224.98 233.30	PV average and range give some idea of region of operation included in performance assessments.
Lo/Avg./Hi OP	4.02 60.08 87.46	OP average and range give some idea of region of operation included in performance assessments.

Performance Metrics: How is TCGH606 performing?		
Parameter	**Value**	**Analysis**
CLPA	0.83	High average CLPA value suggests that the measurement deviation from setpoint has patterns that the controller should be removing. Controller performance could be greatly improved.
Mean Error (%)	0.38	Low average error between the setpoint and measurement suggests performance is good.
Overall Score	39	This overall score suggests FAIR performance. Some aspect of performance could be greatly improved. Optimizing this loop is likely to lead to some measurable improvement in related process performance.

Performance Diagnostics: If performance is poor, what may be causing the performance problem?		
Parameter	**Value**	**Analysis**
Avg. SP Cross	131	Somewhat low average SP crossings suggests a problem such as process disturbances, sluggish tuning, or a poorly performing valve.
Osc. Hdw Diag. (%)	80.0	High value suggests potential valve/hardware problems.
Spec. Hdw Diag. (%)	50.0	High value suggests potential valve/hardware problems.
Avg. Idle Index	-0.03	An idle index in this range produces no diagnosis.

Detailed Performance Assessment

Run a detailed assessment using automated assessment data from 09-Oct-2001

Fig. 6. Excerpt from a complete loop information report for loop TCGH606.

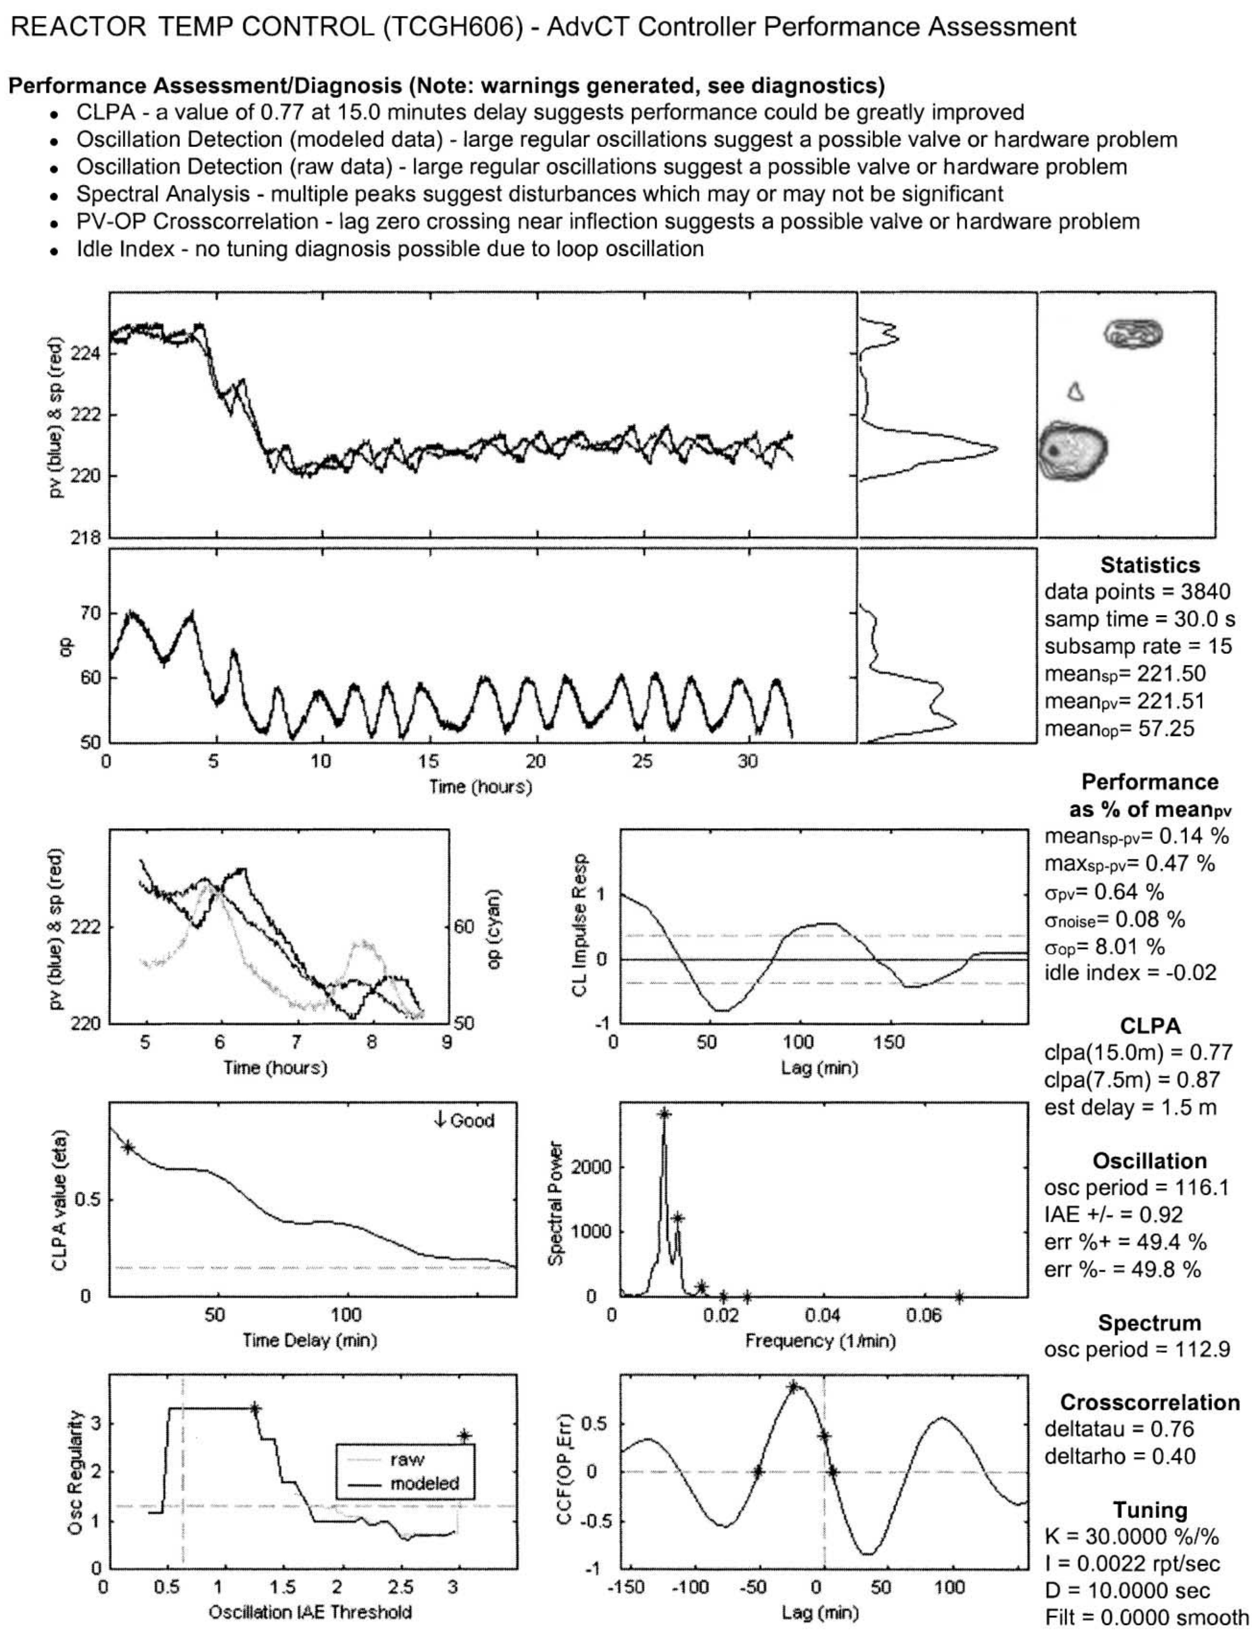

Fig. 7. Excerpt from a detailed performance assessment for loop TCGH606.

diagnosis in the detailed assessment does not agree with the majority of the other results does not cast significant doubt on the potential for a hardware problem. Rather, it is an example of how the user should review all the available information and make conclusions based on a preponderance of the evidence.

Given the importance of the loop, the logical next step is to attempt to verify the suspected hardware problem. It is worth noting that an experienced user of the system may stop at this point and conclude that there is a hardware problem and request that the valve be fixed. Such a conclusion could be aided by knowing, for example, that similar valves in temperature service in the plant have shown hysteresis problems.

3.5. Verifying the preliminary diagnosis

Assuming that the user is not completely convinced of the problem (or if a mistake in diagnosis would be very costly), they might request a preliminary field check of the control valve hardware. A quick check of the TCGH606 valve by an instrument mechanic reveals that the transducer has adequate supply pressure and reasonable output pressure. The valve is a typical sliding stem globe valve that "looks OK". The valve has no positioner. These preliminary findings are typical when trying to verify hardware problems; even loops that have significant problems often do not appear to have problems when checked in a cursory way. Occasionally, a bypassed positioner, leaking diaphragm, or other obvious problem is found, but this is relatively rare. At this point, a more rigorous check of the valve is necessary to confirm the problem. Options at this stage of the investigation include:

- Check and refine the preliminary diagnosis by recording the stem position and transducer pressure output during normal operation using appropriate instrumentation.
- Check and refine the preliminary diagnosis by performing a full field test of the valve and transducer when the process is down or the valve is bypassed.
- Perform step tests of the valve during operation in manual mode to generate more (circumstantial) evidence of a valve problem.
- Accept the preliminary diagnosis as correct and final given the overwhelming evidence in the assessment results.

Given process considerations, the user selects the full field test option, which is performed during the next process shutdown. The key result from the field test is shown in Fig. 8. The transducer shows some slight miscalibration, but the valve shows about 8% hysteresis, poor calibration, and some nonlinearity. This hysteresis measurement is right in line with the assessment results and is clearly the source of the poor performance. The valve needs to be pulled and repaired during the shutdown. In this case, a positioner should also be added for this critical loop to achieve maximum performance.

This example illustrates how easy it is to find an important problem using the performance detail report, as well as the troubleshooting time that can be saved by utilizing the historical performance and detailed controller information. The computer-generated preliminary diagnoses are not always correct. However, the chance that the user will draw good conclusions is improved by having a number of diagnoses and a number of assessments over time.

3.6. Performance summary report

The features described above are geared toward detecting and diagnosing problems with individual control loops. The performance summary report is intended to serve a complementary purpose—help users measure and track the collective performance of loops in their process area. With this monthly report, users can view trends in overall ratings of loop performance by type (flow, pressure, etc.), as well as benchmark performance of loops in their process area with similar loops in other Eastman process areas. The strengths of this report are in monitoring performance on an area-wide scale, mea-

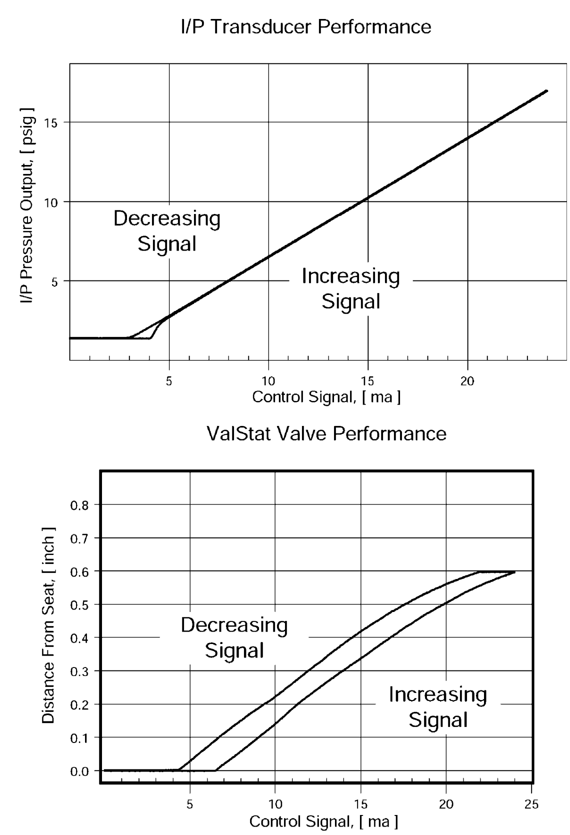

Fig. 8. Results from a field test on a valve with suspected problems.

suring the success of loop maintenance/improvement efforts, and identifying loop types needing the most improvement.

Consider the following example to illustrate the use of an actual performance summary report. Upon receiving the report, the user quickly peruses the entire report, finding a variety of graphs and statistics. She finds that over the past 3 months, about 3000 assessments have been carried out on almost 350 loops in her area. This seems interesting, but her attention is quickly drawn to the graph shown in Fig. 9, the overall performance of controllers by type over the 3-month period. A significant percentage of controllers have poor or fair rankings. The user is not sure of what the number in parentheses after the loop type means, but she reads the explanation above the graph in the report and finds that it is sort of a "grade" for the controller type as a whole, where higher numbers are better grades. It appears that many loops of each type are getting failing grades.

Another graph in the report, shown in Fig. 10, provides some encouragement. The grades for each of the loop types are plotted by month for the last year. The current performance may look bad, but the graph shows that there has been steady improvement.

The final section of the report contains a series of graphs like the one shown in Fig. 11, indicating how controllers of a given type from the user's process area rank against those across Eastman Chemical Company facilities. Of particular interest is the graph for pressure controllers, given that the historical performance for pressure controllers shown in Fig. 10 did not show the same improvement over time that the other types did. The user is somewhat encouraged upon seeing that her pressure controllers with a collective score of 63 are not too far away from the best-performing group of pressure controllers within the company, which have a score of 70.

This example shows the utility of the performance summary report for users concerned about area-wide controller performance. In a short time, users can gauge how the loops in their area are doing, as well as identify major improvement opportunities. Built-in to the report are measures that can be directly used to initiate and monitor loop maintenance and improvement efforts.

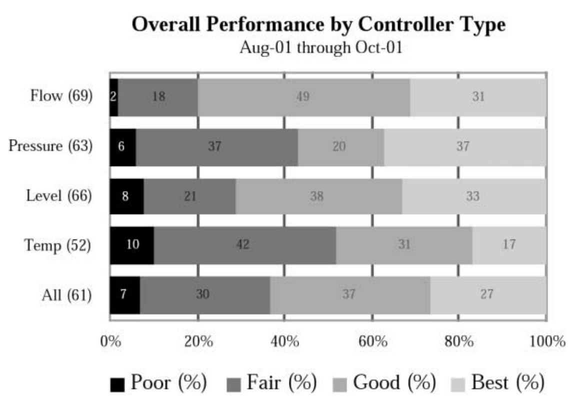

Fig. 9. Excerpt from the performance summary report showing overall performance of 350 loops of various types over a three-month period.

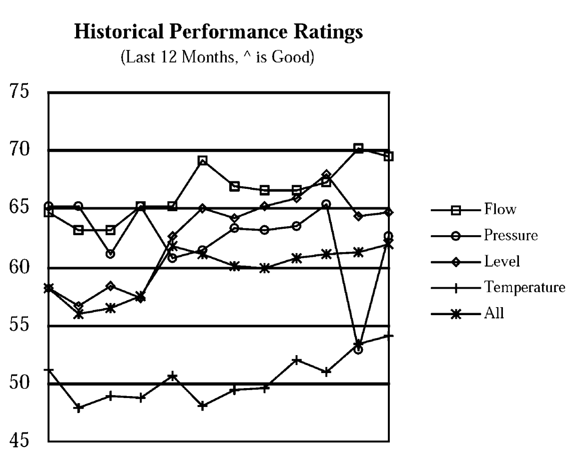

Fig. 10. Excerpt from the performance summary report showing a single-number performance statistic over time for various categories (types) of controllers.

4. Eastman controller demographics

As of January 2002, the Eastman performance assessment system contains daily tuning history on over 14,000 controllers, and weekly performance history is available on almost 9000 controllers. These controllers originate from several brands of control systems and are operating in 40 distinct plant/process areas located at 9 different Eastman sites throughout the world. The bulk of this large-scale system has been in place since 1999.

4.1. Controller types

Fig. 12 shows the percentage of the 14,000 controllers by controller category. Not surprisingly, flow con-

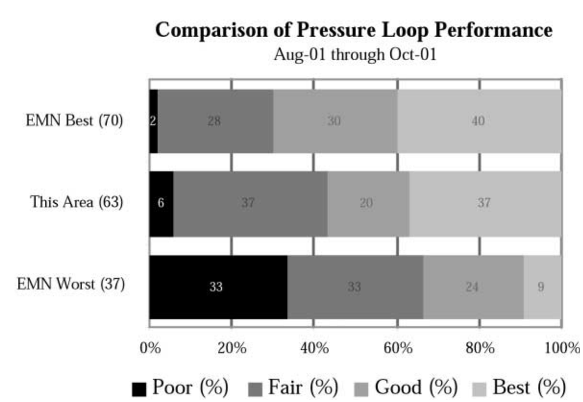

Fig. 11. Excerpt from the performance summary report showing the overall performance of pressure controllers as compared to the best and worst ranked areas within the company.

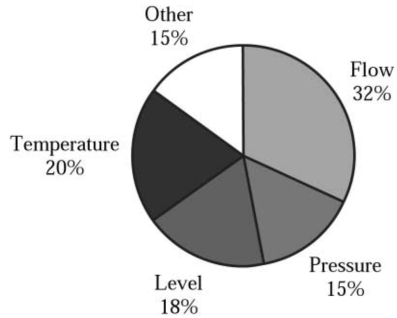

Fig. 12. Percentage of 14,000 Eastman controllers by type.

trollers make up the largest percentage of controllers. The ''other'' category in Fig. 12 is significant and includes analyzer controllers, power controllers, weight controllers, speed controllers, and other miscellaneous types of controllers.

4.2. Controller performance ratings

In the Eastman performance assessment system, a controller's overall performance is rated with a continuous numeric score and the scores are divided into one of four classifications: poor, fair, good, or best. The numeric score is derived from a weighted sum of the most important performance statistics and diagnostic results. For the 9000 controllers having weekly performance history, Table 1 shows the percentage of each controller type rated in each of the 4 performance classifications, along with a description of typical performance for each classification. The percentages of controllers in the poor and fair categories show the potential for controller improvement and optimization. Our experience has shown that many of these under performing controllers, especially those with a poor rating, have hardware problems (valve/positioner/transducer). Eastman has traditionally placed a low priority on loop hardware maintenance, but results from the assessment system are changing that philosophy.

Note that the performance classification percentages in Table 1 are for controllers operating in automatic. Related to this, we have found that approximately 30% of all the PID controllers at Eastman are in manual at any given time. This percentage is similar to published statistics for the process industries, which imply that a serious problem exists. At Eastman, the vast majority of controllers in manual mode are not in that mode as a result of performance problems; some are due to processes not running, some are due to a process operating mode that does not require certain controls, some are ''dummy'' controllers that just hold data, and others are abandoned in place due to process changes over time. This situation may not be representative for all companies, but conclusions should be drawn carefully regarding the significance of the percentage of controllers in manual. In our case, the statistic suggests more about better utilizing our supply of controllers than it does about controller performance.

Table 1
Percentage of 9000 Eastman controllers by performance classification

Class	Controller type					Class description
	Flow	Pressure	Level	Temperature	All	
Best	39	24	27	13	28	These loops are performing well and do not need attention. They are typically tracking the setpoint well, with very few or no significant deviations.
Good	35	29	30	28	31	These loops are performing adequately, but probably have some component of performance that could be improved. Benefit to cost ratio for making improvements to these controllers is likely to be small unless the tolerance for deviation from setpoint for the loop is unusually low.
Fair	23	30	34	43	31	These loops are not performing up to potential and should be improved. Control is probably being maintained in a broad sense, but deviation from setpoint is likely to be degrading process performance. These loops should be investigated further.
Poor	3	17	9	16	10	These loops typically have a serious performance problem. The loop may be cycling strongly, may have large and frequent deviations from setpoint, or may not be tracking the setpoint at all. Many of these loops may have low criticality. Improvements to critical loops in this category will often lead to significant process performance improvements.

5. Summary of system benefits and issues

Development and use of the performance assessment system has generated many benefits and revealed several key issues. These are outlined below.

5.1. Benefits summary

Before development of the performance assessment system at Eastman, poorly performing controllers were found one-by-one in areas where a process control engineer was working. Control loop maintenance (much less optimization) was reduced to reactively troubleshooting controllers that caused enough problems to make themselves obvious. The performance assessment system has ushered in a new era for us. With the performance assessment system:

- The controllers in a process area can be ranked in order of performance. Plant personnel can easily (almost effortlessly) obtain these rankings.
- A wealth of controller tuning, performance, and configuration information is instantly accessible for troubleshooting controller problems.
- The overall performance of controllers in an entire plant or process area can be easily monitored over time.
- Users can be automatically alerted about changes (e.g., controller tuning changes) that are indicative of potential problems.
- Comparison and benchmarking of analogous controllers at different Eastman plant sites is possible because of the universal user access and the large-scale assessment system.

As a result of this functionality, significant productivity gains and other benefits have been realized. The most significant benefits are:

- Huge amounts of process data are transformed into concise information valuable to a variety of plant personnel (area managers, staff engineers, etc.).
- Controllers needing attention can be identified quickly from the hundreds of controllers in process areas, making ongoing controller improvement programs feasible. More problem controllers are fixed and optimized, reducing process variability and operator intervention requirements.
- The process for troubleshooting problem controllers is streamlined, cutting troubleshooting time typically by two-thirds.
- Monitoring and benchmarking of area-wide controller performance fosters large-scale controller improvement efforts, resulting in positive economic impact on processes.
- Use of the system has strengthened communication between control engineers, staff engineers, area managers, and maintenance forces. Input from these personnel is essential to the various components of controller improvement.

The economic benefits resulting from performance assessment are difficult to quantify on a loop-by-loop basis. More often, each problem loop is contributing in a complicated way to poor overall process performance. After finding and fixing problem loops throughout a plant, 6 months to a year worth of data shows reduced off-class production, reduced product property variability, and occasionally lower operating costs or improved production rate. Controller alarms and operator interventions are also typically reduced, which enhances safety and frees some operator time for other value-adding tasks.

As an example, one Eastman process area has been using performance assessment for almost two years. The area has approximately 400 active PID controllers. Assessments identified many loops with hardware problems and approximately 40 loops have had repair or replacement of the valve, positioner, transducer, air supply, or DCS output board. Over the last year, off-class production due to process-control-related causes has been reduced by 53% (540 klbs/year). The standard deviation of the primary specification for material produced in this area has been reduced by 38%. The area has advanced from the 40th percentile to the 75th percentile of all Eastman process areas worldwide in overall controller performance. The areas above it typically are new plants with new equipment and instrumentation. Other problem loops have been identified in this process area, but have not yet been fixed pending a shutdown or availability of funds. We expect that further significant performance gains will accrue when these remaining loops are repaired.

5.2. Issues summary

A primary goal in developing the assessment system was to enable efficient application to a large number of controllers (20,000–30,000 envisioned). Any manual effort required to incorporate loops quickly becomes burdensome and costly with such a large number of loops. We have been able to design the system so it operates effectively with only the specification of a loop type. Initial configuration of the loop type is handled by the system administrators and can often be determined directly from the ISA standard tag name of the loop. The list of loops can be extracted automatically from the DCS or imported from a DCS report in text format.

Users, or process personnel, do not have to do any initial configuration. However, we have provided the ability for users to later customize individual loop configuration, as this can provide additional value. Some examples of this include the ability to specify controller criticality and to adjust weightings on how the performance statistics affect individual loop performance classifications. It is usually not difficult to get users to specify criticality or weighting since they do it naturally if presented a list of problem controllers sorted with only the default criticality and weight. Their process experience tells them which variables can (or should) float and which must be tightly controlled. If they specify their preferences once, they save themselves time in the long run as the list will be sorted more appropriately the next time they look at it.

Due in part to it being a highly automated large-scale application, the assessment system is not intended to be a complete substitute for human oversight, especially with regard to problem diagnosis. We have chosen to offer diagnosis information in a probabilistic framework. We have a number of diagnosis tests, and the results of each are presented. In some cases, the results are completely consistent, but in many cases there is some disagreement. Furthermore, we perform assessments weekly and retain the results so we have a matrix of diagnoses over time and over tests. Each of the diagnoses can be viewed as a "vote" intended to convince the user of the existence of a specific problem. We have found this approach to result in good user conclusions, especially for controllers with hardware problems. Our experience also indicates that reasonable diagnosis inaccuracy can be tolerated if there is a good relationship between the users and the performance assessment experts. In cases where the automated diagnosis does not provide a clear indication of the problem, at least the existence of a problem has been identified and field tests ranging from manual output steps to valve stem position measurement can be carried out to determine the problem.

It is worth noting that when computing an overall performance score for a loop, we also use the "matrix" approach of various statistics over a number of assessments. Since we capture the assessment data on a schedule, we will occasionally get data that include atypical operation such as startups, shutdowns, grade transitions, large upsets, etc. Assessments on this type of data will produce results that are also atypical when compared with the majority of assessments. Therefore, we use outlier detection and removal techniques to get a more robust composite score that is more representative of typical operation.

The performance assessment system is just one part of a process needed to optimize controller performance. Little gain will be realized by introducing the system without instituting an Observe-Orient-Decide-Act type of work process as well. Performance assessment brings huge efficiency increases to the "observe" and "orient" tasks, making decisions easier. However, action is critical. Where action involves repair or replacement of control loop hardware, good feedback about the benefit of this investment is needed to keep the decisions being made in favor of action. At Eastman, we are beginning to establish a culture of loop hardware maintenance as the benefits of this action are reflected in controller performance improvements that now can be measured and communicated.

6. Future enhancements

Work is ongoing to enhance the Eastman performance assessment system in the areas of problem diagnosis, closed-loop identification, and detection of plant-wide distributed disturbances.

While it seems unrealistic to completely automate controller problem diagnosis, we plan to improve problem diagnosis capability by employing nonlinear time-series signatures, automated pattern recognition, and an improved diagnostic rule base.

The number one request of users is specific guidance or recommendations for tuning parameters. We plan to improve our closed-loop model identification such that the model confidence is suitable for making tuning recommendations. Improvements in this area will also add accuracy to problem diagnoses related to poor tuning.

There is also significant interest in improving disturbance diagnosis by identifying sets of loops that appear to share a common disturbance and even identifying a loop that may be the root cause of the distributed disturbance. This capability is somewhat outside the core of performance assessment, but it integrates nicely with the spectral analysis that we already do.

One area of current research interest that is not on our short-term enhancement list is assessment of multivariable controller performance. While this would be of value, we believe that our limited resources would be better utilized in making improvements that apply to the much more common single-loop controller. We expect to be adding multivariable functionality in the longer term as multivariable controllers become more numerous and as the multivariable assessment algorithm development matures.

Acknowledgements

The authors gratefully acknowledge the support and technical input of the Eastman Advanced Controls Technology group, in particular the invaluable help of

Jim Downs, Ernie Vogel, and Joey Watson. Additionally, we acknowledge Nina Thornhill whose research and practical applications in the controller performance assessment field have been of great benefit to us.

References

[1] T.J. Harris, Assessment of control loop performance, Can. J. Chem. Eng. 67 (1989) 856–861.

[2] T.J. Desborough, L.D. Harris, Performance assessment measures for univariate feedback control, Can. J. Chem. Eng. 70 (1992) 1186–1197.

[3] Kozub, D.J. Garcia, C.E. Monitoring and diagnosis of automated controllers in the chemical process industries, Proc. AIChE Annual Meeting, St. Louis, MO, USA, 1993.

[4] N. Stanfelj, T.E. Marlin, J.F. MacGregor, Monitoring and diagnosing process control performance: the single-loop case, Ind. Eng. Chem. Res. 32 (1993) 301–314.

[5] S.J. Qin, Control performance monitoring—a review and assessment, Computers and Chem. Eng. 23 (1998) 173–186.

[6] B. Huang, S.L. Shah, Performance Assessment of Control Loops, Springer-Verlag, London, 1999.

[7] P. Jofriet, C. Seppala, M. Harvey, B. Surgenor, T. Harris, An expert system for control loop performance, Pulp and Paper Canada 97 (6) (1996) T207–T211.

[8] Kozub, D.J. Controller performance monitoring and diagnosis: experiences and challenges, in: Proc. of the Fifth International Conference on Chemical Process Control, Tahoe City, CA, USA, 1996, pp. 83–96.

[9] N.F. Thornhill, M. Oettinger, P. Fedenczuk, Performance assessment and diagnosis of refinery control loops, AIChE Symposium Series 94 (320) (1998) 373–379.

[10] N.F. Thornhill, M. Oettinger, P. Fedenczuk, Refinery-wide control loop performance assessment, J. Proc. Control 9 (1999) 109–124.

[11] N.F. Thornhill, T. Hagglund, Detection and diagnosis of oscillation in control loops, Control Eng. Practice 5 (10) (1997) 1343–1354.

[12] Miller, R.M., Timmons, C.F., Desborough, L.D. CITGO's experience with controller performance assessment, Proc. NPRA 1998 Computer Conference, San Antonio, TX, USA, 1998.

[13] T.J. Harris, C.T. Seppala, L.D. Desborough, A review of performance monitoring and assessment techniques for univariate and multivariate control systems, J. Proc. Control 9 (1999) 1–17.

[14] T. Hagglund, Automatic detection of sluggish control loops, Control Eng. Practice 7 (12) (1999) 1505–1511.