

A PICTURE WORTH A THOUSAND CONTROL LOOPS: AN INNOVATIVE WAY OF VISUALIZING CONTROLLER PERFORMANCE DATA

Warren Mitchell*, David Shook* and Sirish L. Shah

Department of Chemical and Materials Engineering, U. of Alberta Edmonton, Canada T6G 2G6

**Matrikon Inc., Suite 1800, 10405 Jasper Ave., Edmonton, Canada T6N 4A3*

Keywords: Control loop performance assessment; performance monitoring; Fault detection and diagnosis; Tree Maps; Data Visualization

Abstract: Process control performance is a cornerstone of operational excellence in the refining, petrochemicals, pulp and paper and the mineral processing industry. Control performance assessment and monitoring applications have become mainstream in these industries and are changing the maintenance methodology surrounding control assets from predictive to condition based. The large numbers of these assets on most sites compared to the number of maintenance and control personnel has made monitoring and diagnosing control problems challenging. Identifying specific control issues on a plant-wide basis and their root-causes is analogous to looking for a needle in a haystack. Tree Mapping, a data visualization technology developed in 1990 is proposed as one useful solution to this problem and is demonstrated to work very effectively within a condition based maintenance framework. This paper explores the challenges associated with control performance assessment and monitoring and shows how tree mapping technology eases the visualization challenge associated with monitoring hundreds or thousands of control assets. Both regulatory and advanced process control examples are demonstrated.

INTRODUCTION

It has now been well recognized that continuous performance assessment is essential for maintaining the advanced process control (APC) assets in the process industry. The commissioning of elaborate control system platforms e.g. DCS, advanced control applications e.g. Model-based Predictive Control (MPC), and Information management systems e.g. Historians and Databases etc., have become commonplace in the process industry. Incidentally, these investments have led to the accumulation of tremendous amounts of process data with several new data-mining tools and control-relevant techniques for extracting information relevant to controller performance monitoring.

The implementation of an advanced control strategy is not an end in itself. Continuous improvement in process performance must be ensured by constantly monitoring and assessing the performance of the basic control loops. Recent studies show that significant improvements in performance, up to 30% reduction in

variance, can be realized by re-tuning most of the basic control loops in the process industry (Bialkowski, 1993). This same study also notes that only 30% of the control loops perform the control task satisfactorily, *i.e.* reduce the process variability. This means that there is a significant opportunity for improving the operation of a plant.

A common benchmark for control loop performance assessment is minimum variance control. This benchmarking yields a performance index that is defined as the ratio of the minimum output variance (or minimum mean square error) and the actual output variance (or actual output mean square error). The performance index lies between zero and one. An index of one suggests that the controller is at the minimum variance condition. In this case, further reduction in the output variance is not possible by re-tuning the controller. However, the output variance can be reduced by process re-engineering (e.g. reducing process delay, adding more sensors, implementing feed-forward control etc). A performance index close to zero implies that there is a high potential for reducing the output variance by re-tuning the existing controller. For those loops that have been identified to have poor performance, detailed time-domain and frequency-domain analyses are performed. These analyses provide an insight on the sources of poor performance and lead to tuning guidelines for performance improvement, e.g. whether the problem is due to poor tuning, disturbance upset or loop oscillations.

The most attractive feature of performance assessment using minimum variance control as a benchmark is that the performance index can be calculated by using only routine operating data with *a priori* knowledge of time-delay. Harris (1989) showed that a lower bound of process variance (or the minimum variance) under feedback control could be estimated from routine operating data. This lower bound (or minimum variance) can then be used as a reference point to assess current control loop performance. This technique has since attracted significant interest and has been further developed by many researchers (e.g. Desborough and Harris, 1992; Kozub and Garcia, 1993; Lynch and Dumont, 1993; Huang et al. 1995; Thonhill et al., 1995). An expert system for plant-wide feedback control performance assessment has also been implemented at QUNO newsprint mill (Harris et al., 1996). In addition, Huang et al. (1996, 1997) and Harris et al. (1996) have extended this technique to multivariate processes.

PERFORMANCE ASSESSMENT OF UNIVARIATE CONTROL LOOPS USING USER SPECIFIED BENCHMARKS

The increasing level of global competitiveness has pushed chemical plants into high performance operating regions that require advanced process control technology. Consequently, industry has increasing need to upgrade the conventional PID controllers to advanced control systems. The most natural questions to ask for such an upgrading are follows. Has the advanced controller improved the performance as expected? If yes, where is the improvement and can it be justified? Has the advanced controller been tuned to its full capacity? Can this improvement also be achieved by simply re-tuning the existing traditional (e.g. PID) controllers? In other words, what is the cost versus benefit of implementing an advanced controller? Unlike performance assessment using minimum variance control as benchmark, the solution to this problem does not require *a priori* knowledge of time-delays. Two possible relative benchmarks may be chosen: one is the historical data benchmark or reference data set benchmark, and the other is a user-specified benchmark.

The purpose of reference data set benchmarking is to compare performance of the existing controller with the previous controller during the 'normal' operation of the process. This reference data set may represent the process when the controller performance is considered satisfactory with respect to meeting the performance objectives. The reference data set should be representative of the normal conditions that the process is expected to operate at, i.e. the disturbances and set-point changes entering into the process should not be unusually different. This analysis provides the user with a relative performance index (RPI) which compares the existing control loop performance with a reference control loop benchmark chosen by the user. The RPI is bounded by $0 \leq RPI \leq \infty$, with '<1' indicating deteriorated performance, '1' indicating no change of performance, and '>1' indicating improved performance. Figure 1 shows a result of reference data set benchmarking. The impulse response of the reference data smoothly decays to zero, indicating good performance of the controller. After one increases the proportional gain of the controller, the impulse response shows oscillatory behavior, with an RPI=0.4, indicating deteriorated performance due to the oscillation

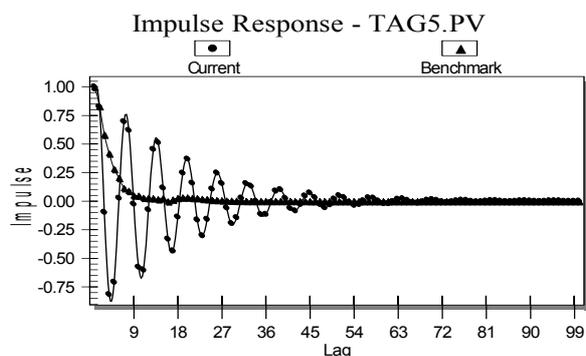


Figure 1: Reference benchmarking based on impulse responses

In some cases one may wish to specify certain desired closed loop dynamics and carry out performance analysis with respect to such desired dynamics. One such desired dynamic benchmark is the closed loop settling time. As an illustrative example, Figure 2 shows a system where a settling time of 10 sampling units is desired for a process with a delay of 5 sampling units. The impulse responses show that the existing loop is close to the desired performance, and the value of RPI=0.9918 confirms this. Thus no further tuning of the loop is necessary.

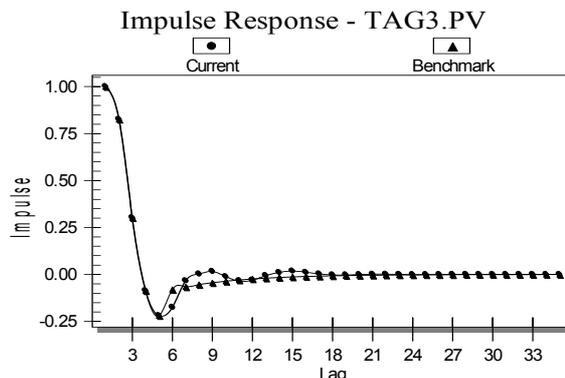


Figure 2: User specified (settling time) benchmark based on impulse responses

CURRENT MAINTENANCE APPROACHES

In modern refining and petrochemical complexes, there is clearly a strong positive correlation between the performance of the process control assets and the financial performance of the business. Virtually every advanced process control project has been justified using predictions of improved business performance resulting from improved control. The process of implementing such projects routinely includes improving the performance of the regulatory controls and in some cases instrumentation.

For the purpose of this discussion, consider the plant control hierarchy (assets) as multiple layers of equipment and technology including: primary sensors and analyzers, end devices (valves, drives, dampers), regulatory controls and advanced control applications such as model predictive control and their associated inferential models (Figure 3). This contains more layers than the standard Purdue Reference Model (PRM) functional hierarchical computer control structure Williams (1989), because it is necessary to show the dependencies among the different components of Levels 2 and 3 of the PRM.

Each device in the plant process control hierarchy impacts the bottom line performance of the unit to which it is associated as its performance degrades. Many authors have documented the root causes of poor control, its effects and ultimate costs to the associated business (Mitchell *et al.*, 2003). These papers discuss the technical, human and business challenges associated with sustaining the performance of process control assets in modern refining and chemical complexes.



Figure 3: Plant Process Control Hierarchy

The current paper assumes that the challenges associated with sustaining the performance of these assets are well understood by the reader and rather than reviewing this material we attempt to bring focus to following question: “Among all of the control assets in the plant, where should maintenance resources be expended in order to have the most significant impact on plant performance?” In this paper we present a new technology that addresses the data visualization problem associated with looking at performance measures from large Model Predictive Control (MPC) applications in addition to hundreds or even thousands of regulatory control loops

Due to the sheer numbers of assets, plant staff remain challenged to maintain large numbers of regulatory control assets and a growing number of MPC applications in the expected ‘optimal’ states. In most refining organizations today, a combination of failure-based and scheduled preventative based maintenance techniques is used in maintaining the plant control performance.

A failure-based maintenance approach is one where the control asset is left un-maintained until one or more failure modes have been observed. As an example, for the regulatory loop, the failure may include valve or sensor problems, improper tuning, disturbance problems, and so on. In a worst case scenario, if the loop is causing significant difficulty for the operator it may be taken out of automatic control and placed in manual until corrective maintenance can be applied. Under this model, maintenance technicians and engineers respond to high priority operator complaints in a “fire-fighting” fashion.

A preventative based maintenance approach involves regularly visiting each asset, assessing its performance and applying necessary maintenance to ensure that it continues to perform optimally. In practice, most facilities have mature preventative maintenance programs. Critical loops and applications tend to be fairly well maintained as they receive the majority of the control engineer or maintenance technician’s time, while others are maintained on a less frequent schedule

if at all. As described, in most plants, the large number of loops per control engineer dictate the ‘fire fighting’ or ‘don’t fix till broke’ maintenance approach.

Recognizing the obvious weaknesses of these maintenance strategies, many refining and chemical companies have now adopted a condition-based approach to maintaining their plant controls. A condition-based maintenance (CBM) program employs a control performance assessment and monitoring application which detects control performance problems at all levels of the plant control hierarchy from control valve to MPC application. Based on the health or condition of the asset, appropriate engineering, maintenance, and operations personnel are alerted to problems and are able to more effectively address control related problems. Significant benefits have been reported by many companies. As examples, Eastman Chemical reported a 53% reduction in off-class production due to process control related issues (Paulonis and Cox, 2003) while Marathon Ashland Petroleum reported a 500 bbl/day increase in throughput on their crude topping unit (Mitchell, 2003). Other financial benefits have been reported by many refining and chemical customers world wide. Other, less tangible benefits reported have been:

- performance of all plant controllers and applications can be viewed on-line and those needing attention are flagged automatically
- visibility of the problem by all those associated with maintaining and supporting the plant controls is improved
- reduction in maintenance costs associated with a given level of plant performance as the work process of maintaining the plant controls and applications is streamlined
- reduced number of alarms and controller interventions by operation staff

CHALLENGES AND REVELATIONS

Over the past several years control performance assessment and monitoring technology has improved. In our opinion, the challenges are no longer related to whether the technology itself is effective, but rather related to the human factors surrounding the use of these applications. How an application integrates with existing work practices and maintenance procedures is the critical success factor. If the application is truly monitoring all of the plant control assets, a tremendous amount of information is generated from huge volumes of plant data. Key to the success of these applications is in the ability of the application to allow the user to quickly identify where to direct his or her attention. Keeping in mind that in a single plant, one may be monitoring upwards of 1000 regulatory controllers, a dozen plant analyzers, online inferential models and 10 or more MPC applications, the important question becomes: “Among the poor performing control assets, where should individuals focus their time?” Earlier versions of control performance applications generated

reports on individual controllers or applications, but required users to access performance measures and reports using a tree hierarchy that mimicked the plant hierarchy (plant-unit-controller) to organize the assets and required the user to drill in on individual controller performance reports to successfully troubleshoot problems. Such an interface is shown in Figure 4.

Though effective, users complained that finding problems with this method was inefficient. They needed a user interface that quickly directed them to important problems. The issue here is not unique to this application, but rather is common to many automatic monitoring problems. Perhaps the best description of the problem is by Tufte (1990):

“...at every screen are two powerful information-processing capabilities, human and computer. Yet all communication between the two must pass through the low-resolution, narrow-band video display terminal, which chokes off fast, precise, and complex communication”

the user can be presented with an unmanageable set of independent problems to address. Prioritizing among assets becomes impossible.

To overcome these limitations, we have found it necessary to deliver the following functionality:

1. A consistent view of different classes of asset, that permits the user to prioritize across assets (*i.e.* regulatory loops, MPC applications, analyzers, online estimators);
2. A high-information-density display, that allows users to make visual comparisons among assets. *“Comparisons must be enforced within the scope of the eyespan, a fundamental point occasionally forgotten in practice”*. Tufte (1990);
3. An interface that enforces the appropriate workflow, presenting problems to users requiring user action and follow-up, and supporting that action. Figure 5 shows a simplified condition-based maintenance workflow. It clear from Figure 5 that the system must present the user with the information required to prioritize, confirm diagnosis and schedule maintenance.

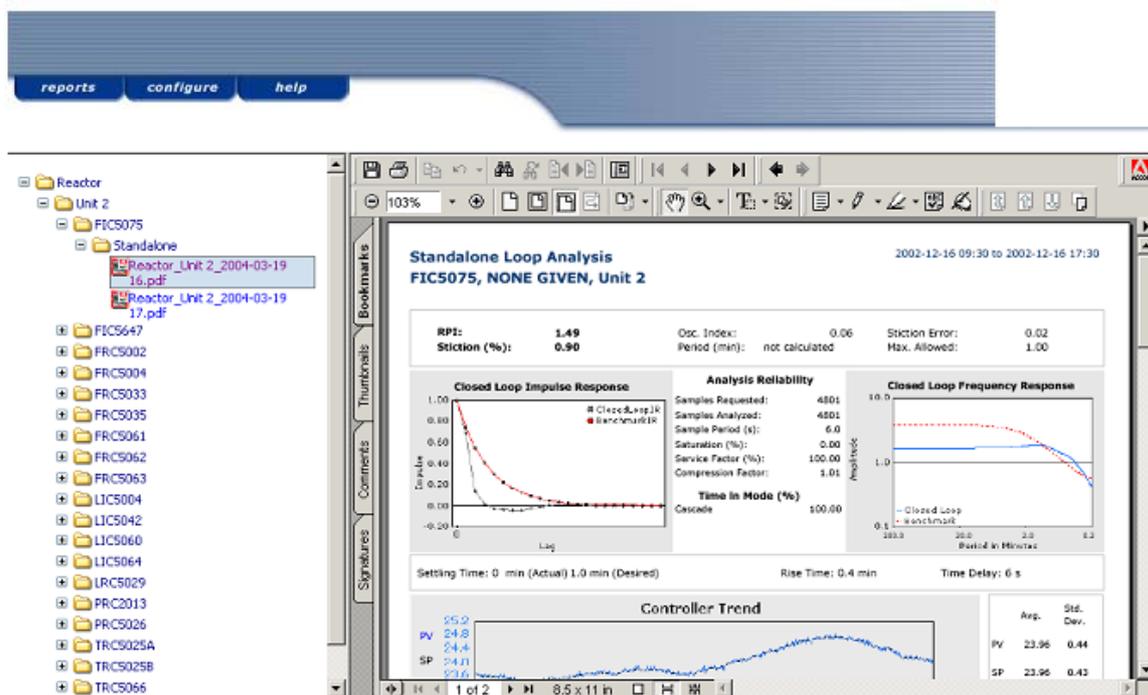


Figure 4: Tree-View of Control Assets

In short, people can only act on information once it has been squeezed through the computer interface. This problem is compounded with complex monitoring solutions, where there are several performance criteria. With the integration of several different classes of assets in a single system, such as with control assets,

TREEMAPPING TECHNOLOGY AS A MEANS OF VISUALIZING CONTROL PERFORMANCE PROBLEMS

A tree map is a nested classification. It can be represented as a diagram that grows from a single trunk and ends with many leaves. Tree maps were intended to

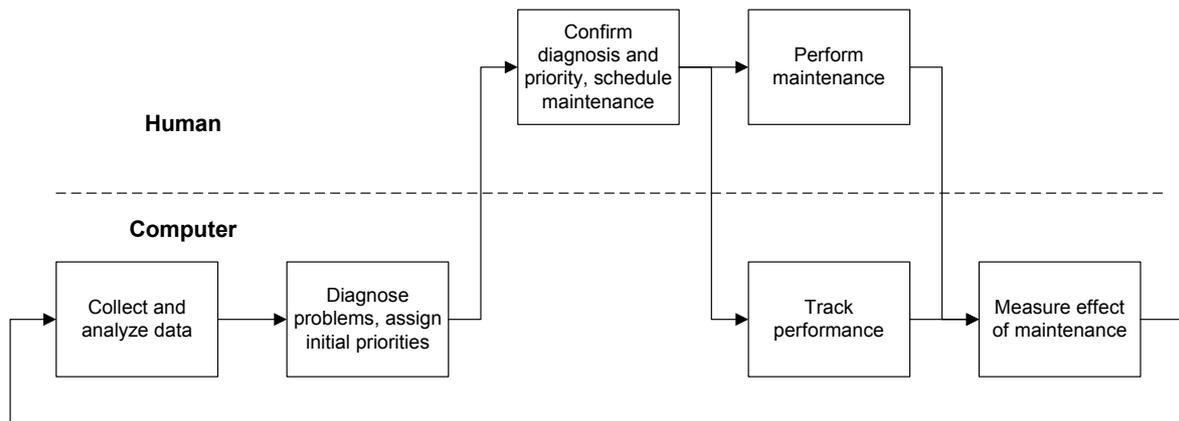


Figure 5: Condition-Based Maintenance Work Cycle

be visualization aids for representing complex organization charts, biological species trees and so on. However, designers of such ‘maps’ found that a large wall/space was necessary to display the whole picture of complex organizations or family trees. But even at that level, only the structural relationship was displayed; additional information such as the size or importance of each node is generally ignored. The control hierarchy in a typical plant can be grouped as shown in Figure 6. Alternately the controllers can be grouped according to the higher level MPC controllers that they serve.

The objective of an effective controller performance visualization chart is to present the information on the performance of the entire plant at a macro level and yet have the ability to display local (micro) details at the unit level or even the loop level. From a visual point of view it is important to provide an indication of the number of control loops in each unit rather than simply displaying them as nodes of the same size. This problem can be solved by using Shneiderman’s (1992) method of representing tree maps as a 2-dimensional rectangular plot within which each node or unit is also represented by a smaller rectangle whose area depends on the number of control loops being monitored in that unit. The idea was conceived when Shneiderman(1992) created an application to track the utilization of a hard drive shared by 14 users. Shneiderman’s work revolved around creating interfaces that present large quantities of data simultaneously in a way that allows users to quickly detect and recognize relevant information and patterns. His work facilitates effective visual comparison by presenting a vast amount of information in a single display. Simple controls allow the user to change the display criteria and filter the set of data viewed.

It is possible to view all of the monitored assets on a large installation in a single screen, thus circumventing the need for “drill-down” and making it possible to prioritize across an entire site.

A controller performance monitoring tool can take advantage of tree mapping technology by using it to provide an information-abundant visual interface that allows users to monitor and assess all of a facility’s control assets in a single view. Rather than depict control assets as a text list, treemapping technology uses shape, size, color and grouping of geometric shapes to effectively impart key performance information related to individual control assets.

Performance monitoring tree maps can present data on all the control assets in a given plant with user-configurable views of key performance statistics such as: service factor, oscillation strength, relative performance index, number of user interventions, number of alarms, valve stiction, process nonlinearity indexes and so on. Such a treemap makes it possible to easily spot controllers that are performing poorly among thousands, then zoom in on a group or hover the mouse over a single asset for more detailed information.

APPLICATION OF TREE MAPPING FOR THE CONDITION-BASED MAINTENANCE OF REGULATORY CONTROL ASSETS

As an illustration, Figure 7 below shows a treemap of the performance of 817 regulatory controllers on a single screen. On this visualization screen map, each controller is represented by a single rectangle, and controllers are grouped by process unit as denoted by the blue/grey border. In this particular view, the size of the rectangle for the controller is determined by the number of operator interventions, and colored by the controllers performance (RPI – Relative Performance Index) while it is active. Controllers that were never active are shown in grey. The user can change the selection criteria using the drop down boxes at the top, and can apply filters shown at the right.

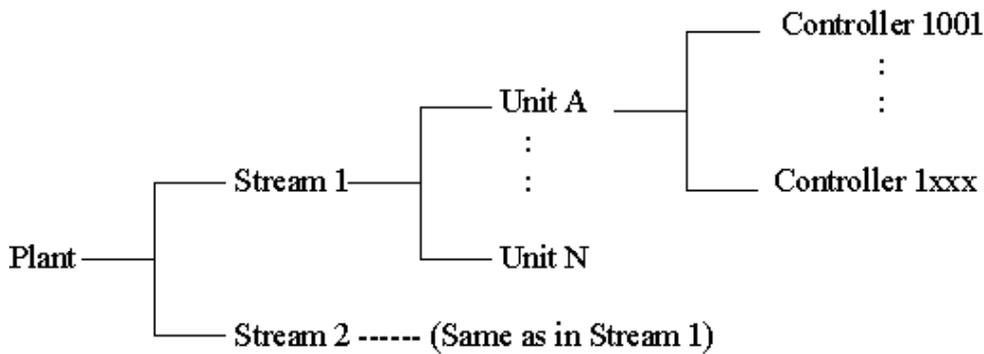


Figure6: Trunk and leaf node tree map of control layout in a typical plant

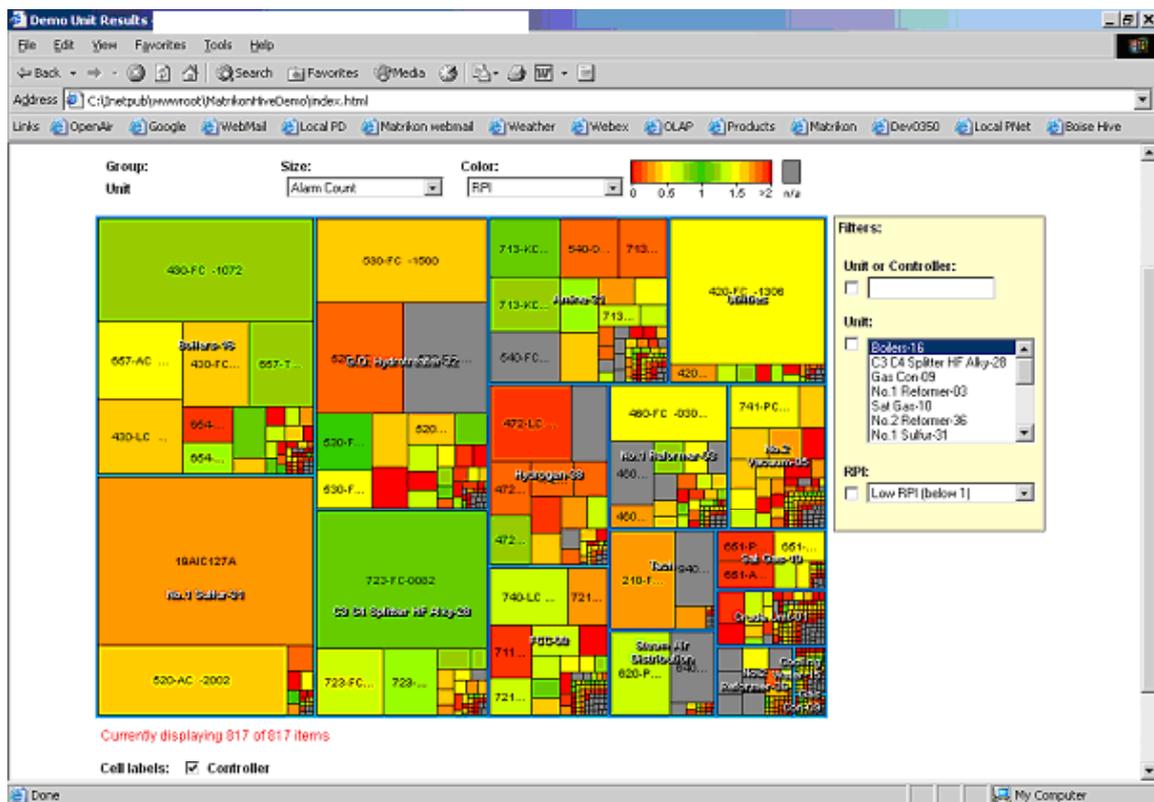


Figure7: Tree Map of Regulatory Control Performance

The advantages of this type of presentation are obvious as the human eye is automatically drawn to the largest and most colorful objects. As the computer screen space (real-estate) and color are allocated based on the filters selected, users can very rapidly identify problems among all the controllers on the site without spending valuable time searching for problems. The controllers generating the greatest number of alarms in this example show up as the largest rectangles on the screen, while the colour spectra represent the RPI (relative performance index) statistic calculated by the online performance monitoring tool (in this particular

case, Matrikon's ProcessDoctor). Rectangles that are red for example indicate controllers that are either settling much faster or slower than their prescribed benchmark. Conceptually, large red blocks denote controllers that require attention; while small green blocks represent controllers that are performing well and can be ignored by the user. Both the size and color filters are user adjustable and can be set for several of the key performance indicators such as– service factor, operator interventions, oscillation strength, Harris Index, %valve stiction, and so on. This provides maintenance, operations and engineering with several different views of the same dataset.

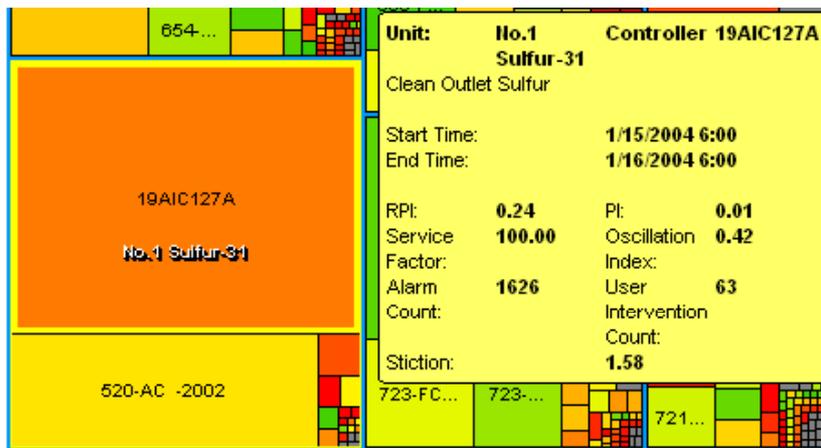


Figure 8: Key Performance Indicator (

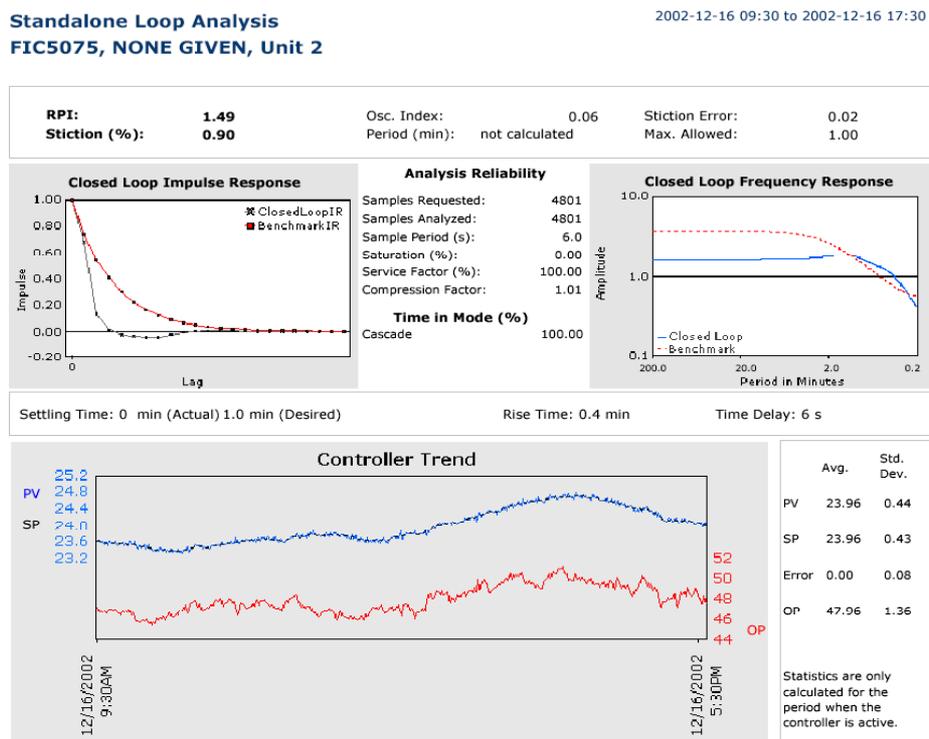


Figure 9: Detailed Diagnostic Report for FIC-5075

Several of the key performance measures calculated for a particular controller can be viewed simply by passing the mouse over the individual rectangles quickly providing users the information they require – as shown in Figure 8.

Additional actions are initiated from the menu that pops up when the user clicks on the controller. In this implementation, users can view a detailed diagnostic report where additional controller performance details are provided as shown in Figure 9. This display helps the user confirm or refute the diagnosis as well as the initial priority set for the specific controller.

The tree mapping technology employed in a state of the art performance monitoring tool should allow users to visually determine the priority of fixing an asset relative to others. One should be able to initiate additional maintenance work processes from this interface, such as creating a maintenance work order or annotating the result if desired. User may then schedule maintenance using a screen similar to the one shown in Figure 10.

After scheduling the asset for maintenance, the system continues to track performance, but does not alert the user of ongoing performance problems until the scheduled date for maintenance has passed. In this way an efficient performance monitoring tool can facilitate

the condition-based maintenance workflow shown earlier in Figure 5.

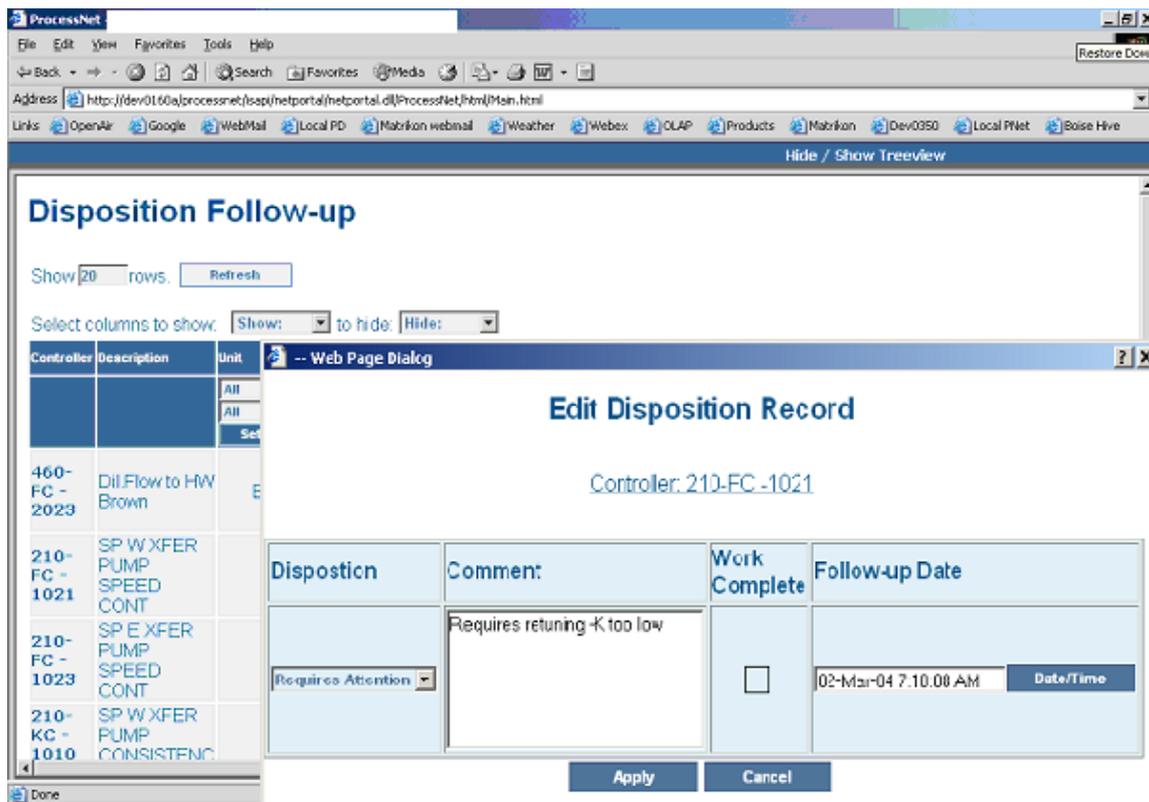


Figure 10: Maintenance Scheduling View

APPLICATION OF TREE MAPPING FOR THE CONDITION-BASED MAINTENANCE OF ADVANCED PROCESS CONTROL (APC) APPLICATIONS

The power of tree mapping is that it permits the user to view a large quantity of information from many assets simultaneously. It is simple to present all of the APC applications in a plant, or even an enterprise, within a single screen, as can be seen in Figure 11.

Figure 11 shows 9 MPC applications across 3 units (Crude, LER, Cat Reformer). Again, the size and color filters allow users to quickly identify issues with constraints, service factor, and model performance, problems commonly experienced with CMPC applications. These summary views are excellent for a quick inspection of the status for a group of applications like the ones above. Tree maps complement the traditional tabular views of data seen in more detailed reports as shown in Figure 12. Tabular views of the data are preferred for a small number of monitored items, where there are several measures per item, as in Figure 12. Tufte (1990) states that “tables usually outperform graphics in small data sets of 20 numbers or less” and that one strength of tables is their ability to display exact numerical values. For that reason, tables are used during drilldown showing detailed information more precisely, than can be accomplished in a tree map.

The tree map is even more effective for examining the performance of the controllers in detail. It is possible to view all of the CV’s (Controlled Variables), MV’s (Manipulated Variables) or DV’s (Disturbance Variables) for all of the APC applications in a plant in a single display, as can be seen in Figure 13. In this example, the CVs are shown grouped by APC application, the size of the boxes are determined by the percent time the CV is outside its limits, and the colour is determined by the model quality for the CV. It is simple to see from the figure that two APC applications have significant limit violation issues: 01ATM_MPC and 14RX_MPC. It is also clear that model performance underlies much of the problem. Using a hierarchical tree interface as was shown in Figures 4 and 6, a minimum of four web page views (plant, unit, MPC application, CV’s) would be required to draw a similar conclusion. The tree map exposes this information immediately. Unlike navigation trees, which become more cumbersome as items are added, the power of the treemaps increases with the number of monitored assets. Figure 13 shows 9 MPC applications with a total of 89 CVs.

From this view, users can drill into a single CV to confirm the problem revealing a set of trends showing the behaviour of the CV in detail over the analysis period. In a similar manner, associated tree maps and

detailed reports analogous to Figures 13 and 9 can provide detailed MV and DV performance.

CONCLUSIONS

In most process plants, unit operators are responsible for monitoring several hundred or even thousands of variables. With such a wide span-of-control, operator responsibility forces them to work in an "alarm-driven" mode and acting only on a post-breakdown maintenance mode to fix a problem. On the other hand performance monitoring tools such as the one presented here have the capability of not only giving a 'macro' or the big-picture view of the state of the entire plant but also efficiently providing access to the local or micro details to inform engineers and operators of any impending difficulties. The issue of control performance is one that is well known in the refining, petrochemical, pulp and paper and mineral processing industry. Modern technologies for monitoring control assets are now available that are transforming the control maintenance model from either preventative or failure based to condition based.

Experience to date with control performance monitoring technologies in refining has been generally positive, with recent industrial feedback bringing focus to a significant area for improvement – visualization of important results.

Information density is a significant problem for any CBM application where large numbers of assets are monitored. These applications assist in the control maintenance work process by transforming volumes of data into valuable performance and diagnostic metrics for individual assets, but as a world-scale plants may have in excess of 2000 regulatory controllers and 10 or more MPC applications, finding the highest priority problems can be challenging. Users report that they do not typically have the time to search for information, yet common user interfaces for such applications require users to search through lists of assets arranged in a tree-like hierarchy. This is time-consuming and can be frustrating for the user, often reducing the effectiveness of the condition based monitoring tool. Tree maps flatten the tree hierarchy displaying key information from each respective leaf on the tree. This allow users to see far more information in context than has been previously possible permitting hundreds or even thousands of monitored assets to be viewed simultaneously. Users can quickly spot, diagnose and prioritize performance problems, expediting the entire control maintenance work process. Applications of tree mapping for both regulatory as well as APC monitoring and diagnosis have been illustrated in a refinery setting.

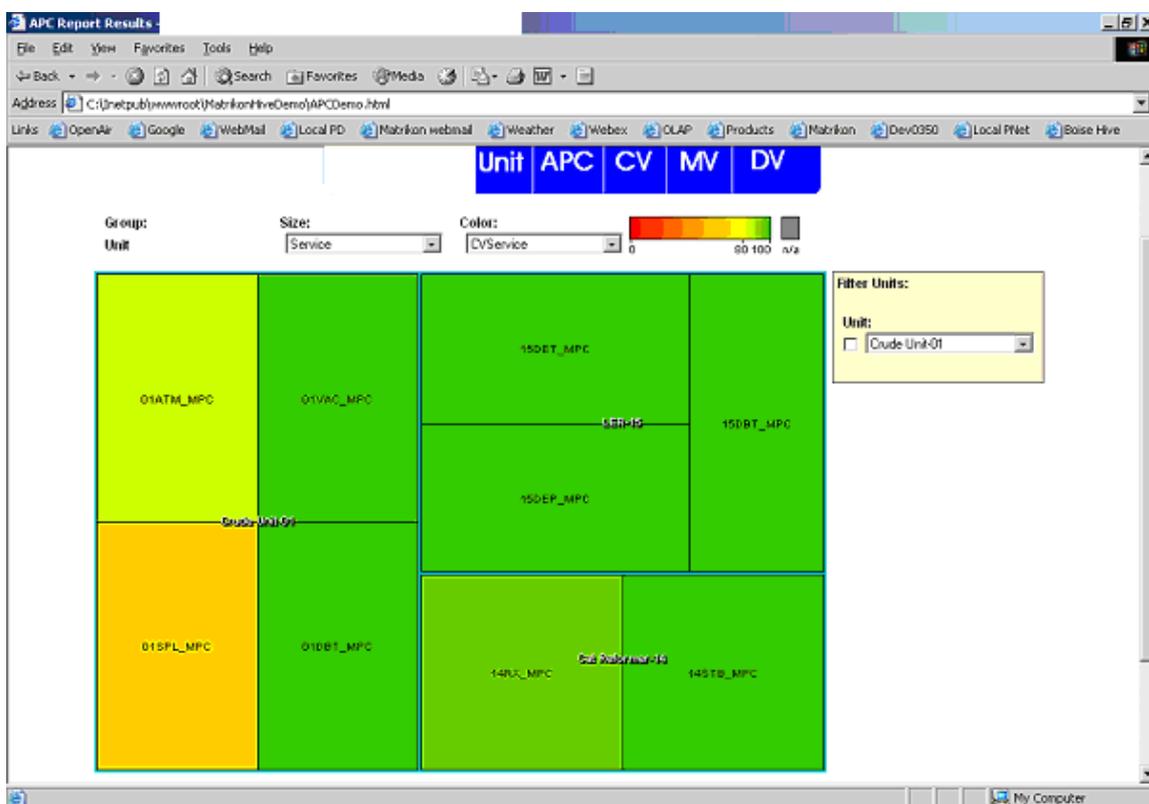


Figure 11: MPC Summary Treemap

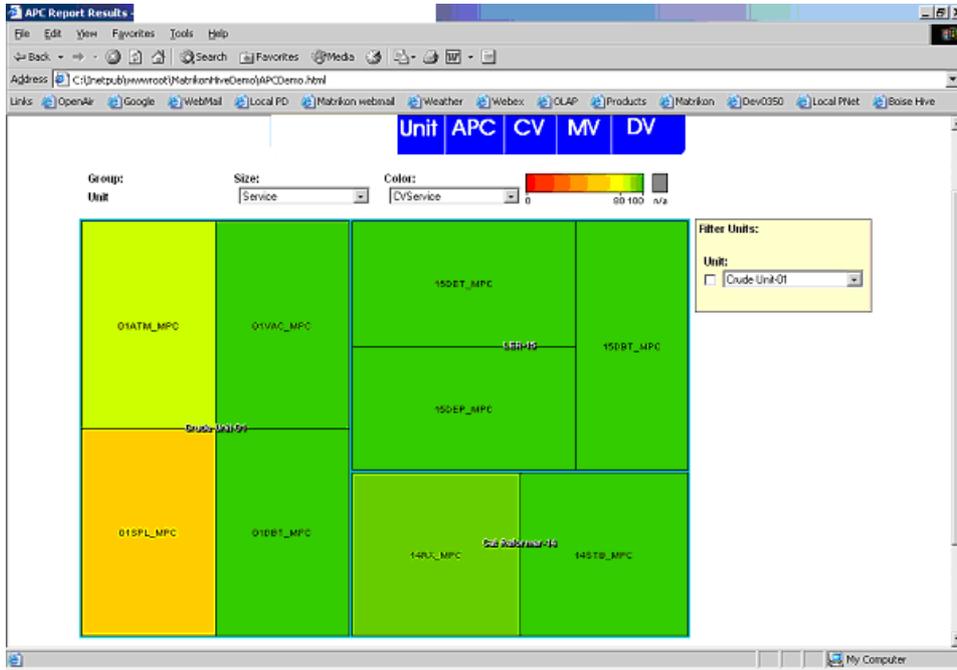


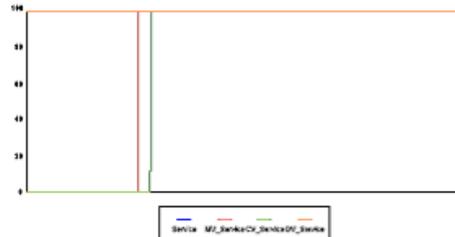
Figure 11: MPC Summary Treemap

APC Performance Monitoring Report

Splitter

3/5/2004

Controller Utilization



Controller Information

Department:	Crude Unit -01
Sample Interval:	1 Minute
Weight:	1

Controller Daily Summary

Service%	MV	CV	DV	Opt	Opt Eff	%MV Limit	%CV Limit	MV RPI	CV RPI	Model Index
100	74.4	71.4	100	100	92	50.6	4.14	0.88	0.72	0.59

Controlled Variables

Service	Service Factor	Effective %Time SF	%Time Opt	%Limit Violation	%Limit Tracking	RPI	Model Index	OSI	Period	Alarms	Operator Interaction
01CC459	100	100	100	0	0	0.95	0.98	0.1	3.8	0	0
01FC492	100	100	100	0	0	0.59	0.82	0.2	1.4	1	0
01FC481	100	100	100	23	19	0.46	0.76	0.14	25	6	0
01FC459.OP	0	0	0	0	0		-0.31	0.7	18	0	0
01DP400	0	0	0	0	0		0.18	0.81	27	0	0
01TI459	100	100	100	0	0	0.76	0.93	0.54	17	0	0
01CC473	100	100	100	6	2	0.87	0.82	0.1	2.1	4	0

Manipulated Variables

Service	Service	Effective Service	%Time Opt	%Time Limit Tracking	At Movesize Limit %	RPI	OSI	Period	Alarms	Operator Action	Prior-ity
01TC437	0	0	0	0	0	0.92	0.13	1.2	0	0	1
01PC420	100	100	100	32	4	0.91	0.06	4.2	3	0	1
01TC478	100	0	100	100	0	0.89	0.14	3.6	0	0	3
01FC459	72	45	100	21	0	0.79	0.07	2.7	0	9	2
01TC487	100	0	100	100	0	0.89	0.19	3.4	0	0	2

Disturbance Variables

Service	Service	OSI	Period	Alarms	Operator action	Priority
01FC301	100	0.11	4.1	0	0	2
01TI302	0	0.81	27	0	0	1

Figure 12: Tabular APC Detail Report

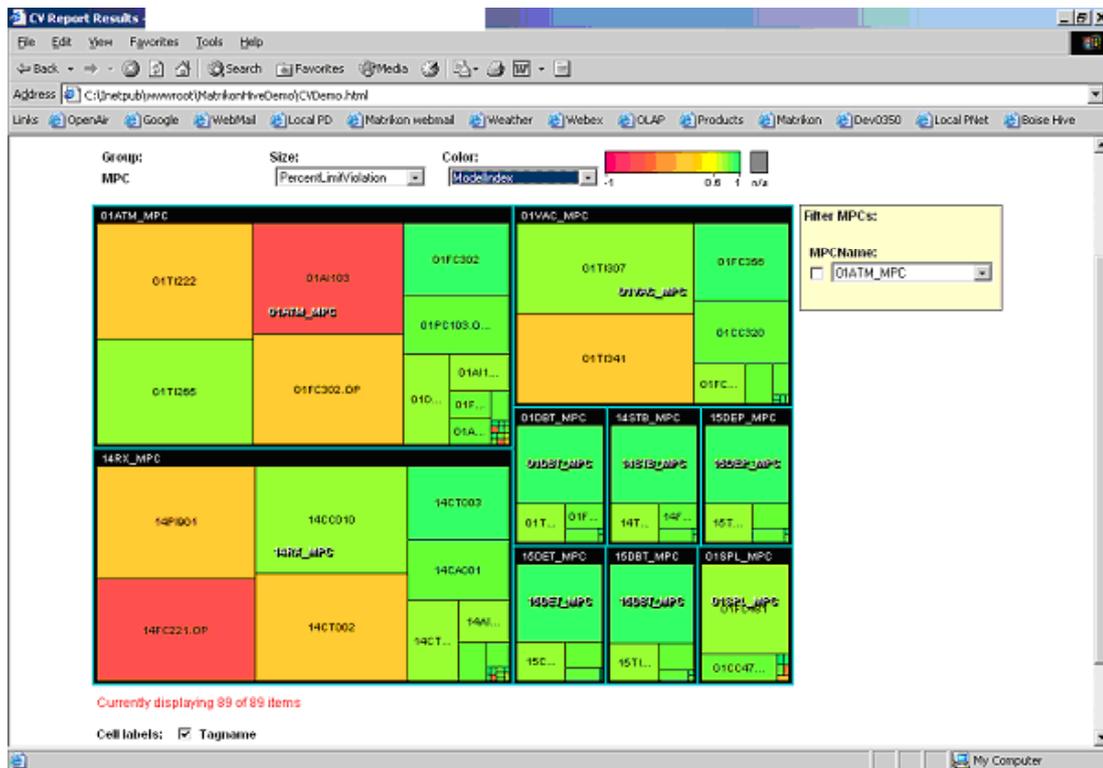


Figure 13: Controlled Variable Treemap

REFERENCES

- Bialkowski, W.L., Elliot, R. Competency in Process Control – Industry Guidelines. Pulp and Paper Canada (1996) 155-160
- Bialkowski, W.L., Dreams vs. reality: A view from both sides of the gap. Pulp and Paper Canada, *94* (11), pp. 19-27., 1993.
- Brown, R. S., The Development of a Rigorous Statistical Test to Audit the Advanced Control Benefits on an FCCU, NPRA Computer Conference (1999)
- Den Bakker, K., Seppala, C., Snoeren, R., Maintaining the Optimum Advanced Process Control. *PTQ*. 107 (2002) 107-115
- Desborough, L., and Harris, T., Performance Assessment Measure for Univariate Feedback Control, *Can J. of Chemical Engineering*, **70**, pp.1186-1197, 1992
- Gao, J., Patwardhan, R., Akamatsu, K., Hashimoto, Y., Emoto, G., Shah, S.L. and Huang, B., Performance Evaluation of two Industrial MPC Controllers, *Control Engineering Practice*, (2003)
- Harris, T., “Assessment of Closed Loop Performance”, *Can. J. Chem. Eng.*, **67**, pp.856—861 (1989).
- Harris, T.J., Boudreau, F., and MacGregor, J.F., Performance Assessment of Multivariable Feedback Controllers, *Automatica* **32**, pp.1505-1518, (1996).
- Huang, B., S.L. Shah, and E.Z. Kwok, On-line Control Performance Monitoring of MIMO Processes, in Proc. Of American Control Conference, Seattle, Washington, pp. 1250-1254, 1995
- Huang, B., S.L. Shah, and E.Z. Kwok, “Good, Bad or Optimal? Performance Assessment of Multivariable Process”, *Automatica*, **33**(6), pp. 1175-1183, 1997
- Huang, B., and Shah, S.L., *Performance Assessment of Control Loops: Theory and Applications*, Springer-Verlag, (1999), October 1999, ISBN: 1-85233-639-0.
- Kendra, S.J. and A. Cinar, “Controller Performance Assessment by Frequency Domain Techniques”, *Journal of Process Control*, *7*(3), pp. 181-194, 1997.
- Kozub, D.J. and C.E. Garcia, “Monitoring and Diagnosis of Automated Controllers in the Chemical Process Industries”, presented at 1993 AIChE Annual Meeting, Nov. 9, 1993.
- Miller, R., And Desborough, L., Increasing Customer Value of Industrial Control Performance, Presented at CPC-6, Tuscon, AZ, Jan. 2001, AIChE Symposium series, 326, 98, 2002
- Mitchell, W.R., Shook, D., Industrial Application of Multivariable Controller Analysis and Monitoring Techniques. ERTC Technology conference, Milan, Italy (2003)
- Mitchell, W.R., Hamel, C., A Condition Based Approach to Control Performance Analysis and Monitoring. NPRA Computer Conference, Sept (2003)
- Paulonis, M.A., Cox, J.W., A practical approach for large-scale controller performance assessment, diagnosis and improvement. *Journal of Process Control* *13* (2003) 155-168
- Shah, S., Patwardhan, R., Huang, B., *Multivariate Controller Performance Analysis: Methods, Applications and Challenges*. Presented at CPC-6, Tuscon, AZ, Jan. 2001, pp 187-219), AIChE Symposium series, 326, 98, 2002
- Shook, D., Khalili, R., Grenier, M., Lariviere, L., Achieving business results through IT: Insuring controller performance benefits with web-enabled monitoring at Noranda, *CIM Bulletin* (2003) 57-59
- Shneiderman, B., *Tree visualization with tree-maps: A 2-d space-filling approach*. *ACM Transactions on Graphics* *11*, 1 (1992), 92-99
- Thornhill N F, Oettinger M, Fedenczuk P, “Refinery-wide control loop performance assessment, *J Process Control* *9* (2): 109-124 April 1999
- Tufte, E. R., *Envisioning Information*, Graphics Press, (1990)
- Tufte, E. R. *The Visual Display of Quantitative Information*, Graphics Press, (1983)
- VanDoren, V.J., Software Review: analysis software answers performance questions, *Control Engineering*, April, pp. 130, 1998
- Vishnubhotla, A., *Frequency and Time Domain Techniques in Control Loop Performance Assessment*, M.Sc Thesis, University of Alberta, Edmonton, Alberta, 1997.
- Willaims, T.J., ed. *A Reference Model For Computer Integrated Manufacturing (CIM)*, ISA, (1989)