

FINAL TECHNICAL REPORT

**“Demonstration of Integrated Optimization Software
at the Baldwin Energy Complex”**

Reporting Period Start Date: February 18, 2004

Reporting Period End Date: November 17, 2007

Principal Author(s): Rob James
John McDermott
Sanjay Patnaik
Steve Piché

Date Report Issued: February 14, 2008

Date Report Revised: September 5, 2008

COOPERATIVE AGREEMENT: DE-FC26-04NT41768

NeuCo, Inc.
Prudential Tower
800 Boylston Street, 30th floor
Boston, MA 02199

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Abstract

This report is the Final Technical Progress Report submitted by NeuCo, Inc., under Award Identification Number, DE-FC26-04NT41768. This award is part of the Clean Coal Power Initiative ("CCPI"), the ten-year, \$2B initiative to demonstrate new clean coal technologies in the field.

This report is one of the required reports listed in Attachment B Federal Assistance Reporting Checklist, part of the Cooperative Agreement. The report covers the whole award period (February 18, 2004 – November 17, 2007) and NeuCo's efforts within design, development, and deployment of on-line optimization systems during that period.

Table of Contents

1	Introduction	3
2	Executive Summary	3
3	Experimental.....	3
3.1	Overview	3
3.2	CombustionOpt and SCR-Opt.....	3
3.2.1	How They Work	3
3.2.2	The User Interface	3
3.3	SootOpt	3
3.3.1	How It Works	3
3.3.2	The User Interface	3
3.4	PerformanceOpt.....	3
3.4.1	How It Works	3
3.4.2	The User Interface	3
3.5	MaintenanceOpt.....	3
3.5.1	How It Works	3
3.5.2	The User Interface	3
4	BEC Data Analysis.....	3
4.1	Overview	3
4.2	BEC Background	3
4.2.1	Historical Events	3
4.2.2	Fuel Heating Value Trend	3
4.2.3	Unit Load Trends	3
4.3	Closed-Loop Optimization Approach	3
4.3.1	Combustion and SCR Optimization on the Cyclones	3
4.3.2	Combustion and Sootblowing Optimization Approach on Unit 3.....	3
4.3.3	Closed-loop Utilization	3
4.4	Closed-Loop Optimization Analysis	3
4.4.1	Comparison of Populations Methodology	3
4.4.2	Modeled Response to Disturbances.....	3
4.5	Advisory Optimization.....	3
4.5.1	Advisory Optimization Examples	3
4.5.2	Summary of Data Analysis	3
5	Managing Tradeoffs and Estimating Optimization Impact	3
5.1	Visualizing Tradeoffs.....	3
6	Conclusion	3
6.1	Estimate of Benefits	3
6.2	Economic Implications	3
6.2.1	BEC Economic Benefits	3
6.2.2	Economic Benefits as Applied to US Fossil Generation.....	3
7	References (N/A)	3
8	List of Acronyms and Abbreviations.....	3
9	Appendices	3

List of Graphical Materials

Figure 1 Overview of the Optimizers at BEC	3
Figure 2 Control Adjustments Before (left of vertical line) and After (right) deploying CombustionOpt on Unit 3. The Overlay (right) y-axis indicates the deviation of the various fuel and air bias and trim settings from "neutral" positions.	3
Figure 3 NOx Pre and Post optimization on Unit .The Overlay (right) y-axis indicates the deviation of the various fuel and air bias and trim settings from their "neutral" positions. The Primary (left) y-axis is NOx in lb/Mbtu.	3
Figure 4 CombustionOpt and SCR-Opt in coordinated action. The Primary (left) y-axis show Ammonia flow in lbs/hr. The squiggly green and blue trends represent the A and B-side reactor ammonia flow measures. The Overlay (right) y-axes represents the deviation of the fuel and air biases from their neutral setpoint.	3
Figure 5 Model Accuracy with Online Learning.	3
Figure 6 Direct Search in Action. Primary (left) y-axis shows NOx in lb/Mbtu (blue squiggly trend). The Overlay (right) y-axis shows the values of a set of air damper position biases as they are adjusted by CombustionOpt in Direct Search Mode.	3
Figure 7 CombustionOpt/SCR-Opt Home page	3
Figure 8 CombustionOpt/SCR-Opt Analysis Screens.....	3
Figure 9 Causality Analysis.....	3
Figure 10 Block diagram of the SootOpt system.....	3
Figure 11 An alternative implementation of SootOpt. In this example, SootOpt computes the optimal cleanliness factors of a criteria based sootblowing system.	3
Figure 12 Expert system for selecting zone to be cleaned.	3
Figure 13 Sample propose rules.....	3
Figure 14 The Blower Selection Component of SootOpt	3
Figure 15 Flow chart of the scenario generator in SootOpt.....	3
Figure 16 The scenario evaluator is used to determine the sootblower activation that minimizes a user-specified cost function.	3
Figure 17 SootOpt Home page.....	3
Figure 18 SootOpt Analysis pages	3
Figure 19 PerformanceOpt - Components in Problem Identification	3
Figure 20 Typical set of modular flowsheets representing a PerformanceOpt model.....	3
Figure 21 PerformanceOpt Home page	3
Figure 22 MaintenanceOpt workflow for problem detection, diagnosis and resolution.....	3
Figure 23 Problems to Diagnose view in MaintenanceOpt.....	3
Figure 24 Potential causes and heuristics for Feedwater heater DCA performance symptom	3
Figure 25 Root cause identification and specification of remedial action in MaintenanceOpt	3
Figure 26 View of diagnosed problems with corrective maintenance actions that require a derate.	3
Figure 27 MaintenanceOpt Home page	3
Figure 28 BEC Coal Analysis Btu/lb (Jan 2004- Dec 2007). The downward trend in fuel energy density could be expected to affect cyclone stability and unit Heat Rate adversely.	3
Figure 29 Unit 1 MW Trend Primary (left) y-axis is in units of Megawatts. Trend is shown to provide general production context for the time span analyzed.....	3
Figure 30 Unit 2 MW Trend. Primary (left) y-axis is in units of Megawatts. Trend is shown to provide general production context for the time span analyzed.....	3
Figure 31 Unit 3 MW Trend. Primary (left) y-axis is in units of Megawatts. Trend is shown to provide general production context for the time span analyze	3
Figure 32 Unit 1 CombustionOpt Master Enable. Primary (left) y-axis is in units of Megawatts (blue trend). CombustionOpt Master Enable Signal is in green with low = Disabled, high = Enabled	3
Figure 33 Unit 2 CombustionOpt Master Enable Left y-axis is in units of Megawatts (blue trend). CombustionOpt Master Enable Signal is in green with low position = Disabled, high position = Enabled	3
Figure 34 Unit 3 CombustionOpt Master Enable Left y-axis is in units of Megawatts (blue trend). CombustionOpt Master Enable Signal is in green with low position = Disabled, high position = Enabled	3

Figure 35 Unit 3 SootOpt Master Enable Left y-axis is in units of Megawatts (blue trend). SootOpt Master Enable Signal is in red with low position = Disabled, high position = Enabled..... 3

Figure 36 Unit 1 % Total MV's Enabled. This value equals the number of fuel and air distribution biases the operator has enabled, divided by the number of biases he has available to enable. For any individual biases to be enabled the Master Enable must be true or "ON"..... 3

Figure 37 Unit 2 % Total MV's Enabled This value equals the number of fuel and air distribution biases the operator has enabled, divided by the number of biases he has available to enable. For any individual biases to be enabled the Master Enable must be true or "ON"..... 3

Figure 38 Unit 3 % Total MV's Enable 3

Figure 39 Layout of Categorized Scatter Plots..... 3

Figure 40 Unit 1 Measured NOx vs. MW..... 3

Figure 41 Unit 2 Measured NOx vs. MW..... 3

Figure 42 Unit 3 Measured NOx vs. MW..... 3

Figure 43 Unit 1 Measured NH3 vs. MW..... 3

Figure 44 Unit 2 Measured NH3 vs. MW..... 3

Figure 45 Unit 1 Measured Heat Rate vs. MW..... 3

Figure 46 Unit 2 Measured Heat Rate vs. MW..... 3

Figure 47 Unit 3 Measured Heat Rate vs. MW..... 3

Figure 48 Unit 1 Measured NOx vs. NH3..... 3

Figure 49 Unit 2 Measured NOx vs. NH3..... 3

Figure 50 Unit 1 Measured NOx vs. Heat Rate 3

Figure 51 Unit 2 Measured NOx vs. Heat Rate 3

Figure 52 Unit 3 Measured NOx vs. Heat Rate 3

Figure 53 Unit 1 Measured NH3 vs. Heat Rate 3

Figure 54 Unit 2 Measured NH3 vs. Heat Rate 3

Figure 55 Unit 1 Modeled NOx as f(Disturbances) vs. MW..... 3

Figure 56 Unit 2 Modeled NOx as f(Disturbances) vs. MW..... 3

Figure 57 Unit 3 Modeled NOx as f(Disturbances) vs. MW..... 3

Figure 58 Unit 1 Modeled Heat Rate as f(Disturbances) vs. MW 3

Figure 59 Unit 2 Modeled Heat Rate as f(Disturbances) vs. MW 3

Figure 60 Unit 3 Modeled Heat Rate as f(Disturbances) vs. MW 3

Figure 61 Unit 1 Modeled NH3 as f(Disturbances) vs. MW..... 3

Figure 62 Unit 2 Modeled NH3 as f(Disturbances) vs. MW..... 3

Figure 63 Unit 1 Cyclone Stoichiometry vs. MW 3

Figure 64 Unit 2 Cyclone Stoichiometry vs. MW 3

Figure 65 Unit 1 Average Cyclone Main Flame Scanner Quality (%) (2005 - present). Orange Line show for visual reference, the red lines indicates were data was not available 3

Figure 66 Unit 2 Average Cyclone Main Flame Scanner Quality (%) (2005 - present) Orange line shown for visual reference, the red lines indicates were data was not available 3

Figure 67 Unit 3 Sootblowing Steam Flow..... 3

Figure 68 Issues to Screen View with FD Fan Amp Increase and Context Data 3

Figure 69 FD Fan Amp Increase Diagnostics View and Context Data 3

Figure 70 Initial Actions Specified for Damaged Over-Fire Air Duct..... 3

Figure 71 JPEG Image of Damaged Over-Fire Air Duct 3

Figure 72 Problem History for Damaged Over-Fire Air Duct..... 3

Figure 73 JPEG Image of Missing Feedwater Heater Nuts..... 3

Figure 74 SootOpt Analysis Screen Showing Furnace-Economizer DPs 3

Figure 75 PerformanceOpt alerts BEC to low condenser cleanliness and the associated impacts . 3

Figure 76 CombustionOpt Analysis Screen Showing Cyclone Flame Quality and Stoichiometries . 3

Figure 77 CombustionOpt Analysis Screen Showing Modeled Flame Quality 3

Figure 78 $NO_x=f(DV's); (S!,C!), (S!,C),(S,C)$ 3

Figure 79 $R=f(DV's); (S!,C!), (S!,C),(S,C)$ 3

Figure 80 CombustionOpt's Home Page 3

Figure 81 What's Going on Now and Why section on CombustionOpt's Home Page..... 3

Figure 82 Optimization Advice section on CombustionOpt's Home Page..... 3

Figure 83 The Change Status Information Collector to Change Priority..... 3

Figure 84 Optimization Benchmarks section on CombustionOpt's Home Page..... 3

Figure 85 SootOpt's Home Page..... 3

Figure 86 What's Going on Now and Why section on SootOpt's Home Page..... 3

Figure 87 Optimization Advice section on SootOpt's Home Page	3
Figure 88 The Change Status Information Collector to Change Priority.....	3
Figure 89 Optimization Benchmarks section on SootOpt's Home Page.....	3
Figure 90 PerformanceOpt's Home Page	3
Figure 91 What's Going on Now and Why section on PerformanceOpt's Home Page.....	3
Figure 92 Controllable Loss parameters monitored on PerformanceOpt Home Page.....	3
Figure 93 Optimization Alerts section on PerformanceOpt Home Page.....	3
Figure 94 The Change Status Information Collector to Change Priority.....	3
Figure 95 Optimization Benchmarks section on PerformanceOpt's Home Page.....	3
Figure 96 MaintenanceOpt's Home Page	3
Figure 97 What's Going on Now and Why section on MaintenanceOpt's Home Page.....	3
Figure 98 Optimization Advice section on MaintenanceOpt's Home Page.....	3
Figure 99 Optimization Benchmarks section on MaintenanceOpt's Home Page.....	3

1 Introduction

The objective of the first CCPI Solicitation (DE-PS26-02NT41428) is to improve emissions, efficiency, maintainability and asset life of coal-based generation and bolster the long-term viability of the United States' abundant coal resources. The first round awards entail a \$1.3 billion cost-shared partnership between the industry and government to demonstrate advanced coal-based power generation technologies that could help meet the President's Clear Skies and Climate Change initiatives.

NeuCo is one of eight companies selected as winners in this initial round. DOE awarded NeuCo a 4-year technology development initiative to design, develop, and demonstrate integrated on-line optimization systems at Dynegy Midwest Generation's Baldwin Energy Complex ("BEC"), which is the host site for the project. The total project budget is approximately \$19 million.

NeuCo is providing 55% of the total project cost; while DOE is providing the remaining 45%. The DOE requires repayment of its investment. This repayment will result from commercial sales of the products NeuCo develops under the project. Dynegy Midwest Generation is contributing the host site, human resources, and engineering support to ensure the project's success.

2 Executive Summary

This project encompassed the design, development, and demonstration of integrated online optimization systems at Dynegy Midwest Generation's Baldwin Energy Complex (BEC) located in Baldwin, Illinois.

The overall project objective was to improve coal-based generation's emission profile, efficiency, maintenance requirements and plant asset life in order to enhance the long-term viability of the United States' abundant coal resources.

Five separate but integrated optimization products were developed, addressing combustion, sootblowing, SCR operations, overall unit thermal performance, and plant-wide availability optimization.

Optimization results are inherently unit-specific and cannot be known for a particular generating unit in advance. However, NeuCo believed that the following were reasonable targets for the completed, integrated set of products:

- Furnace NOx reduction improvement by 5%,
- Heat rate improvement by 1.5%,
- Increase of annual Available MWh by 1.5%,
- Commensurate reductions in greenhouse gases, mercury, and particulates; and
- Commensurate increases in profitability from lower costs, improved reliability, and greater commercial availability.

The goal during Phase I was to establish each system and demonstrate their integration in unified plant optimization. Efforts during Phase I focused on

1. developing, deploying, integrating, and testing prototypes for each of the five products;
2. identifying and addressing issues required for the products to integrate with plant operations; and
3. systematically collecting and assimilating feedback to improve subsequent product releases.

As described in the Phase II continuation application NeuCo successfully achieved the goal for Phase I.

The goal of Phase II was to improve upon the products installed and tested in Phase I and to quantify the benefits of the integrated system.

As this report documents, NeuCo has also successfully achieved the goal for Phase II. The overall results of the project, compared with the project goals, are discussed below.

- **NOx Reduction:** The 5% target for NOx reduction was exceeded with average CEMS and SCR Inlet (furnace) NOx reduction of between 12% and 14%.
- **Heat Rate Improvement:** The optimization systems delivered an average heat rate improvement of between 0.67% and 0.7%. This falls short of the 1.5% heat

rate improvement target largely because Cyclone Stability (availability) and CEMS and SCR Inlet NO_x were prioritized over heat rate in the event they needed to be traded-off with one another. A different prioritization of objectives could have driven a different balance, thereby meeting the target of 1.5% improvement. There were also several factors that could have been masking greater heat rate improvements such as the decrease in fuel density over the course of the project and the impact of actions taken as a result of advice provided by the optimizers that are difficult to quantify.

- **Increased Annual Available MWh:** Although difficult to measure precisely, the target of increasing available MWh's by 1.5% was met by providing prioritized alerts and knowledge-based diagnostics for a wide array of plant equipment and process anomalies; helping the plant to move from high sulfur, high Btu Illinois coal to PRB and run that fuel at low stoichiometries without derates; and improved management of cyclone flame quality as well as improved vigilance with respect to cyclone conditions which avoided some degree of temporary de-rate due to cyclone slag build up.
- **Commensurate Reductions in Greenhouse Gases, Mercury, and Particulates:** Reductions in all three of these indices can be associated directly with the optimization leverage observed in the heat rate and NO_x reductions.
- **Commensurate Increases in Profitability from Lower Costs, Improved Reliability, and Greater Commercial Availability:** Commensurate improvements in costs, reliability and availability resulted from the previously-described benefits. Also playing a role were the sustained operation of the cyclones while using more available, less expensive but off-design fuel; more effective catalytic reduction of NO_x; and the reduced time required to discover, prioritize and diagnose plant equipment issues.

The total annual dollar value of the benefits associated with the products installed, refined, and commercialized at BEC are estimated to range from \$1.8 to \$3.2 million dollars per-unit, and \$7.2 to \$8.1 million dollars per year plant-wide depending on whether CO₂ benefits are included.

The benefits available to the industry based on the results achieved at BEC are estimated at between \$2.3 and \$2.6 billion dollars per year in annual savings across the full combination of unit types, fuel sources, and post-combustion controls characterizing the current US fossil generation fleet. These aggregate benefits are distributed across the categories of fuel efficiency, NO_x reduction, reagent costs, CO₂ emissions, and commercial availability.

The suite of four integrated optimizers commercialized as part of this project are expected to yield well under a one-year payback for average-sized units across all unit types and fuel categories comprising the US fossil power industry. This represents a highly cost-effective way of addressing some of the industry's most pressing challenges – one that complements and enhances a wide variety of other methods such as SCR, alternative fuels, low-NO_x systems, and modern control and instrumentation systems.

In summary, the NeuCo project at BEC was successfully completed and has demonstrated that advanced optimization technologies can play an important role in improving the environmental footprint of coal-based power generation while achieving other important operating objectives.

3 Experimental

3.1 Overview

The Baldwin Energy Complex consists of three coal-fired units. The three units include two cyclone-fired (2x585 MWe) boilers and one tangentially-fired (595 MWe) with low NO_x burners, along with selective catalytic reduction (SCR) units on the cyclone boilers.

During the course of the project, the following five products were developed:

- CombustionOpt
- SCR-Opt
- SootOpt
- PerformanceOpt
- MaintenanceOpt

Figure 1 shows a schematic overview over the Optimizers that were developed, and their impact on the different processes within a power plant.

Different combinations of these products were installed on the three units at the Baldwin Energy Complex:

- Unit 1 (cyclone-fired)
 - CombustionOpt
 - SCR-Opt
 - PerformanceOpt
 - MaintenanceOpt
- Unit 2 (cyclone-fired)
 - CombustionOpt
 - SCR-Opt
 - SootOpt
 - PerformanceOpt
 - MaintenanceOpt
- Unit 3 (tangentially-fired)
 - CombustionOpt
 - SootOpt
 - MaintenanceOpt

This Chapter provides a detailed description, for each product, of the functionality and user interface that evolved during the course of the project.

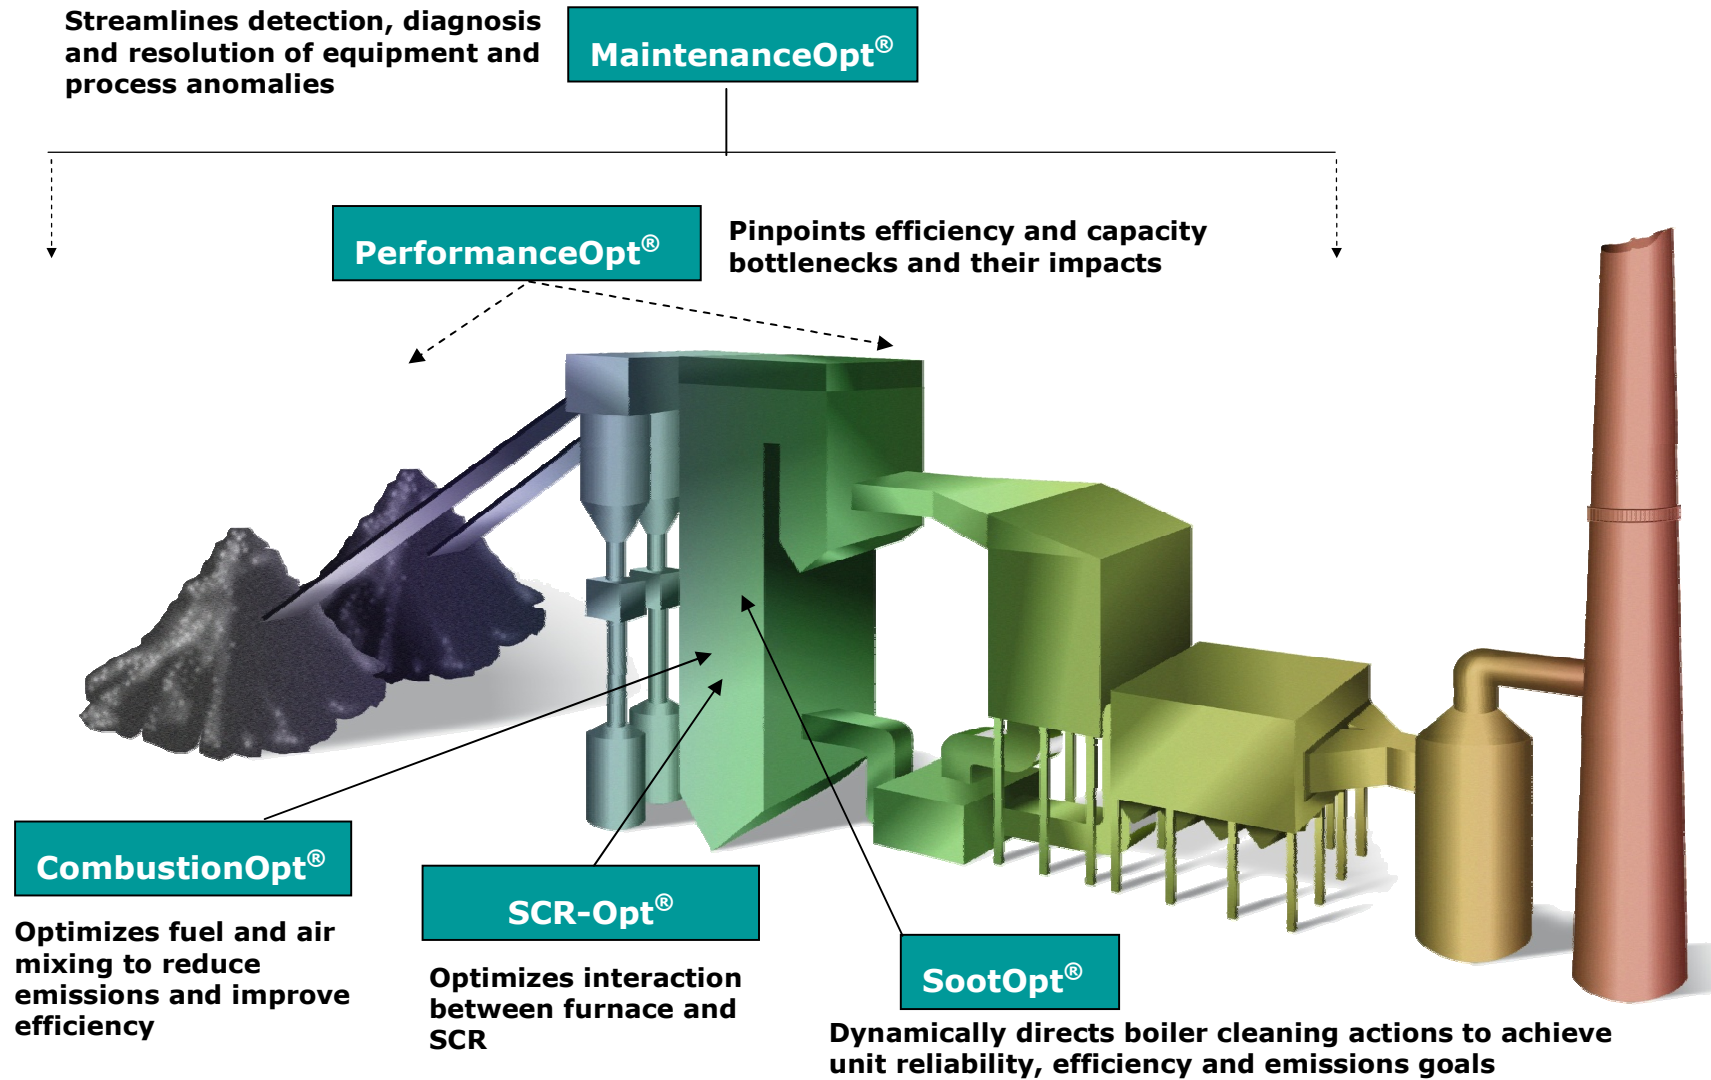


Figure 1 Overview of the Optimizers at BEC

3.2 CombustionOpt and SCR-Opt

3.2.1 How They Work

Because CombustionOpt and SCR-Opt are so tightly integrated, we will describe them together in this section. Both CombustionOpt and SCR-Opt are implemented using neural network technology. A neural network is a function approximator that understands how to change input variables to better meet any set of performance objectives, such as NH₃ NO_x and heat rate, based on real-time and near-history data. Neural networks do not presume any first principle relationships, like those used in multivariate predictive modeling. Instead they capture the real relationships of the input variables and the objective based on current plant data.

3.2.1.1 CombustionOpt in Action

Modern boilers use computerized control systems that automatically deliver the amount of coal (or other fuel) and combustion air needed to make desired amount of power (via the steam turbine). Operators essentially set the amount of Megawatts they would like and the control system drives the coal pulverizers and the fans and air dampers to approximately the right setpoint to meet the demand. These “demand setpoints” are given by a set of curves developed by the boiler designer.

Although straightforward in principle this control process involves literally dozens of different dampers and actuators, and even though there are generally correct “demand setpoints” for each MW setting, a great deal of discretion exists as to how exactly to distribute the quantities of fuel and air. These discretionary settings (also called trim and bias settings) have an impact on the efficiency of combustion and the rate at which byproducts (like NO_x) are produced. The challenge is to know how they exert this effect, which dampers moving in which directions will have what effect on efficiency and NO_x. This knowledge needs to be found not just for one Megawatt setting but for the whole variety of production contexts, and in real-time as the demand for Megawatts varies according to the needs of the transmission system and the final consumers. Add to this complexity the fact that the response of the boiler to these settings changes over time, and the fact that fuel quality often varies greatly and this problem quickly becomes difficult. To sum it up, figuring out how to tune a boiler is an inherently multivariable, non-linear, non-stationary problem.

Only in recent years have such difficult problems had a feasible solution. Machine learning lets human engineers ask the data that the boiler control process generates what effects the adjustment of bias and trim settings have. Regression modeling is a form of machine learning that extracts mathematical relationships from data, without the necessity of any prior understanding of what the relationships might be. Neural networks are powerful learning machines (or regressors) and are used heavily in problems like this across many industries.

Once models containing an understanding of the complex cause and effect relationships hidden in large amounts of data are developed (through neural network model training – or regression), those models can be coupled with an optimizer to find out what bias and trim settings should be used to achieve a set of objectives (which may also be complex) under real conditions in real-time. Because neural networks are robust and easy to use for big regression problems the machine learning-optimization loop can be set in permanent motion, so it is constantly trying to improve its own performance. This is one of ProcessLink’s core technology sets and features prominently in CombustionOpt.

CombustionOpt calculates in real-time the bias and trim settings that improve the mixing of the fuel and air in the furnace leading to reduced furnace NO_x production.

Prior to the CombustionOpt system being turned on, operators generally only make occasional adjustments based on anecdotal observations of how individual moves affect unit performance. In contrast, CombustionOpt makes many changes as the conditions in the boiler change in response to the underlying models' understanding of exactly what changes are needed to meet current performance objectives.

In Figure 2, Fuel-Air (Primary Air for combustion and transport of coal) begin moving when CombustionOpt is turned on. Prior to it being turned on the biases were mostly stationary at values set by operators throughout the course of the shift.

The fine-tuning by CombustionOpt, shown in Figure 2, can have a dramatic impact on the performance of a unit. Figure 3 clearly shows the NO_x measurement (blue line) responding to the change in combustion.

The bias and trim settings shown in Figure 2 and Figure 3 adjust the amount of Secondary Air that is delivered to the flame just above each burner. Primary Air generally serves to transport air (to carry the pulverized coal from the pulverizer to the burner). Secondary air makes up the remainder of the air required for combustion to be sustained. Changes to the relative proportion of Primary to Secondary Air has a significant effect on the properties of the flame, particularly its temperature and oxygen distribution, both of which impact the formation of NO_x. These Secondary Air biases are only a portion of the total list of bias and trim settings CombustionOpt is manipulating, but suffice to visually mark the beginning of optimization activity and to illuminate the complexity of movement needed to achieve optimization. Other biases proportion the amount of the total coal flow needed to make the Megawatts requested, between the upper and lower elevations of the furnace/boiler. This "vertical staging" of combustion is also known to affect NO_x production significantly. Biases that control the overall air/fuel ratio for the boiler, the temperature of the coal-Primary Air mixture leaving the pulverizer and the amount of "burnout" or "overfire" air delivered to the final stages of the furnace combustion process are also typical optimization levers. On a typical unit CombustionOpt manipulates between 25 and 50 of these types of biases, making small step changes, once every few minutes.

Note: Figure 2 and Figure 3 show only a single (albeit highly repeatable) experiment to illustrate the mechanics of combustion optimization. The analysis to follow in this report will also look at the aggregated effect of sustained optimization over time.

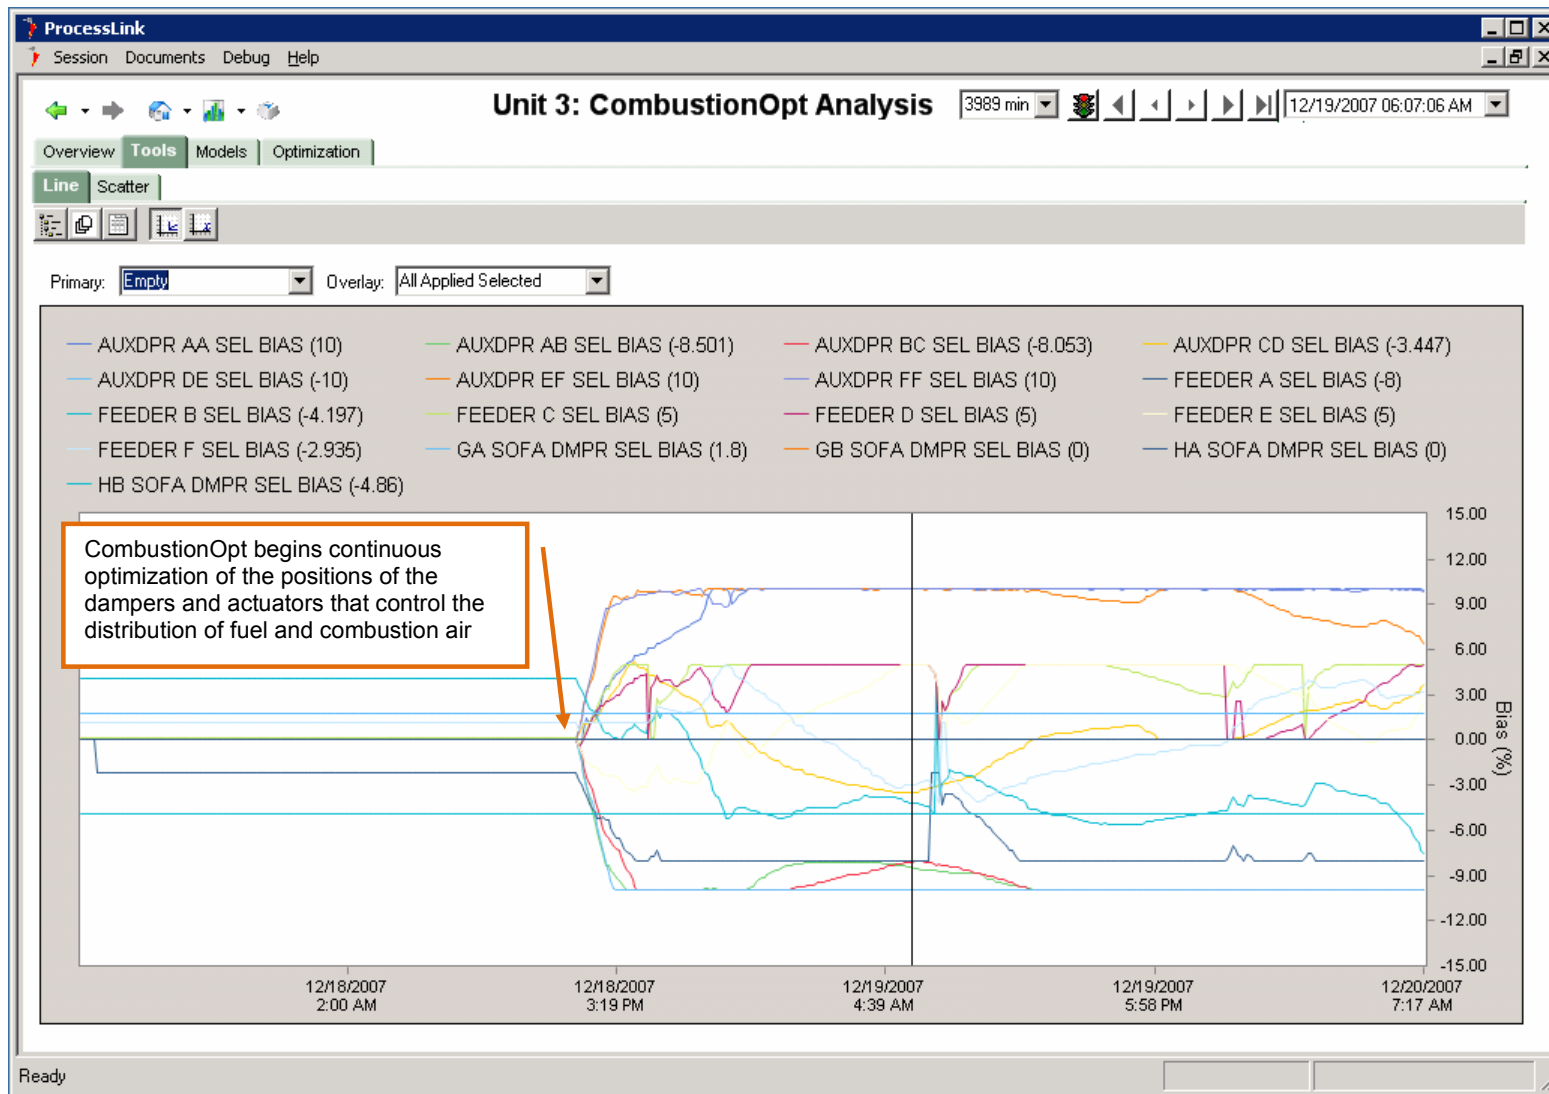


Figure 2 Control Adjustments Before (left of vertical line) and After (right) deploying CombustionOpt on Unit 3. The Overlay (right) y-axis indicates the deviation of the various fuel and air bias and trim settings from "neutral" positions.

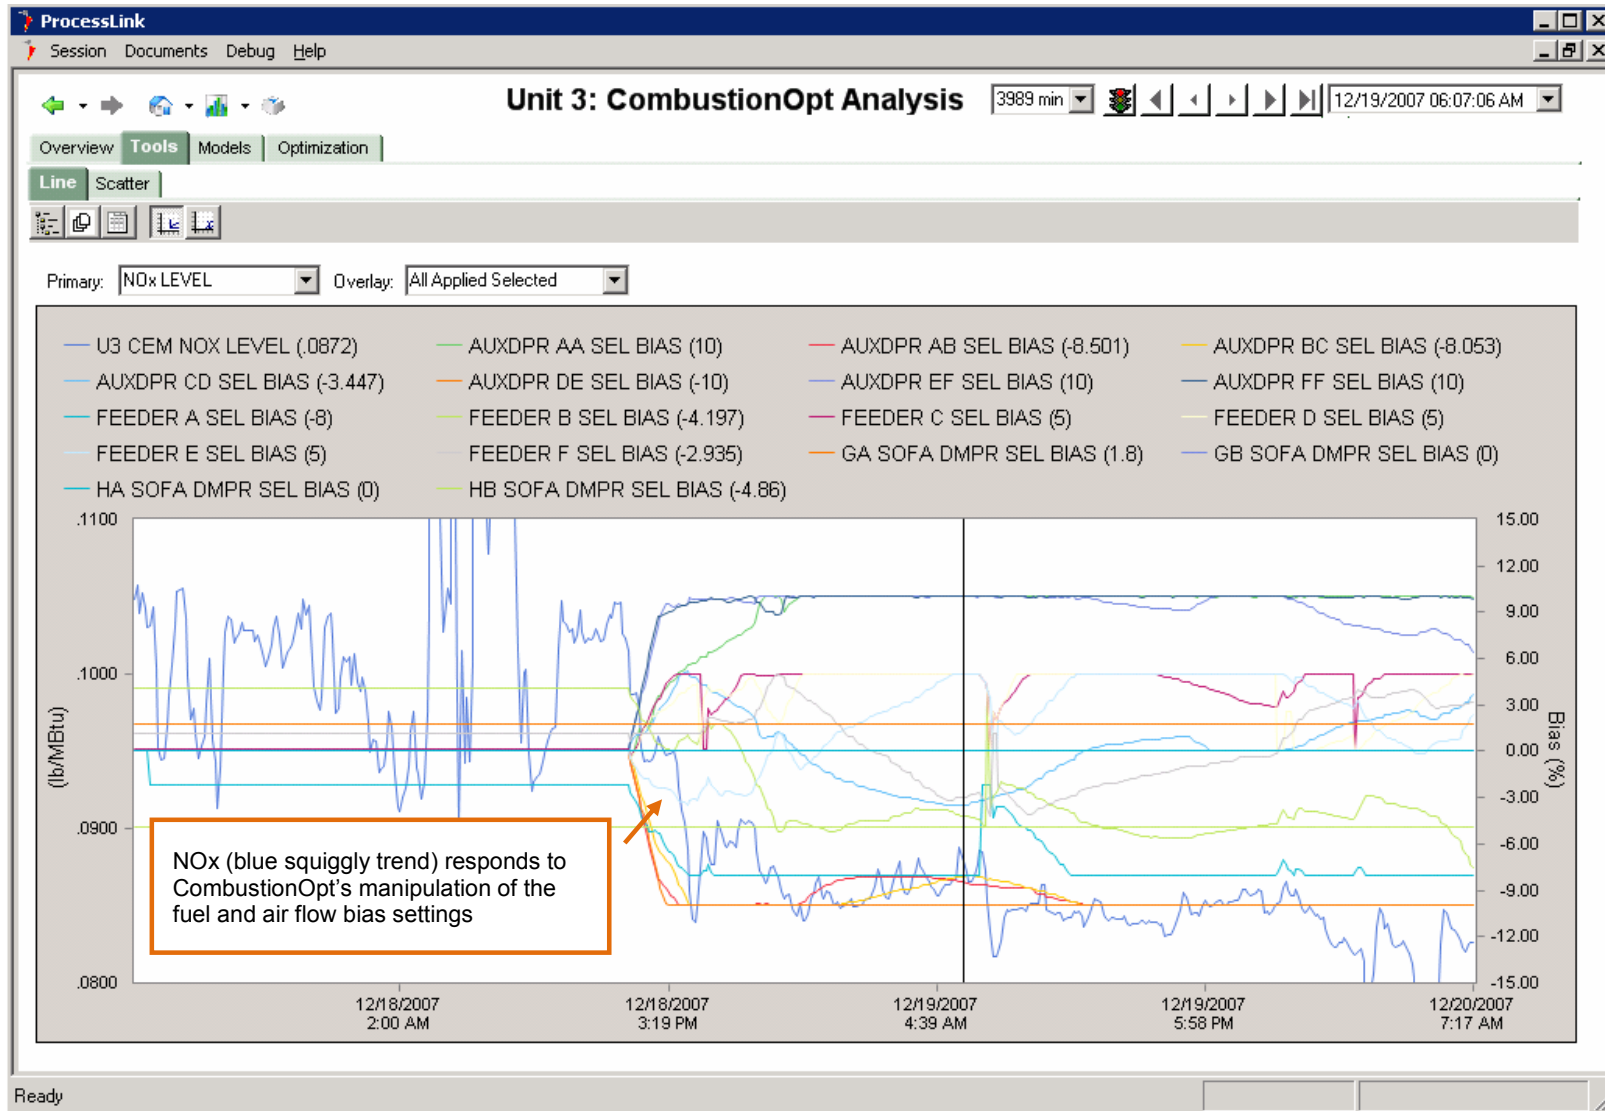


Figure 3 NOx Pre and Post optimization on Unit .The Overlay (right) y-axis indicates the deviation of the various fuel and air bias and trim settings from their "neutral" positions. The Primary (left) y-axis is NOx in lb/Mbtu.

3.2.1.2 CombustionOpt and SCR-Opt in Action

Selective Catalytic Reduction (SCR) uses a Nickel-Cadmium catalyst placed in layered beds and inserted into the flow of combustion exhaust gases leaving the boiler to convert NO_x to N₂ and H₂O. This is achieved using ammonia (NH₃) in some form, as a reagent. SCR's consist of large catalyst arrays built (at great expense) onto the back end of boilers, prior to the stack. To control the reaction occurring inside them, the amount of NO_x entering the reactor is estimated by looking at the amount of air flow through the boiler and a rough ammonia demand setpoint is derived based on the molar chemistry of the reaction (1 mole of ammonia per mole of NO_x), and the desire to achieve some % of NO_x removal. The precise amount of ammonia injected is then adjusted using feedback control to achieve the desired removal rate. Essentially for units with SCR's the operators dial in the desired CEM NO_x rate and the SCR control loop does the rest, adjusting the ammonia flow to maintain that setting. Any excess un-reacted NH₃ leaves the SCR as undesirable "ammonia slip"

The amount of NH₃ being injected into the SCR is directly proportional to the amount of NO_x entering the reactor (though reaction efficiency is also a factor). For a given CEM NO_x setting a reduction of 1% in the amount of NH₃ flowing into the SCR means that 1% less NO_x was seen at the inlet (or the reaction was made 1% more efficient).

The fluid dynamics and control characteristics of SCR's vary significantly, as does their effectiveness. They are affected by stratification in the exhaust gas stream entering the reactor as well as the amount of CO in the gas, and the gas temperature. For this reason on units with SCRs, models and objectives for CombustionOpt and SCR-Opt are integrated and optimization is coordinated so that CombustionOpt is focused on the effect its actions have on NH₃ flow. In this way the characteristics of their interaction between the furnace and the SCR is accounted for. Changes made to the mixing of the fuel and air in the furnace need to lead not only to reduced furnace NO_x production, and combustion efficiency (heat rate) but also increased SCR efficiency.

The impact that these changes can have on performance is best illustrated by looking at actual plant data. In Figure 4, the Cyclone Feeder Speed Biases (regulating fuel flow) and Secondary Air Biases (regulating stoichiometry or fuel to air ratio) are being moved by the integrated CombustionOpt/SCR-Opt system. The specific goals given to this system include maintaining Cyclone Main Flame Scanner Quality, reducing Heat Rate and reducing SCR NH₃ flow, indicating a reduction in SCR Inlet NO_x or an improvement in SCR efficiency, (indicated by the green and blue process trends in the figure below).

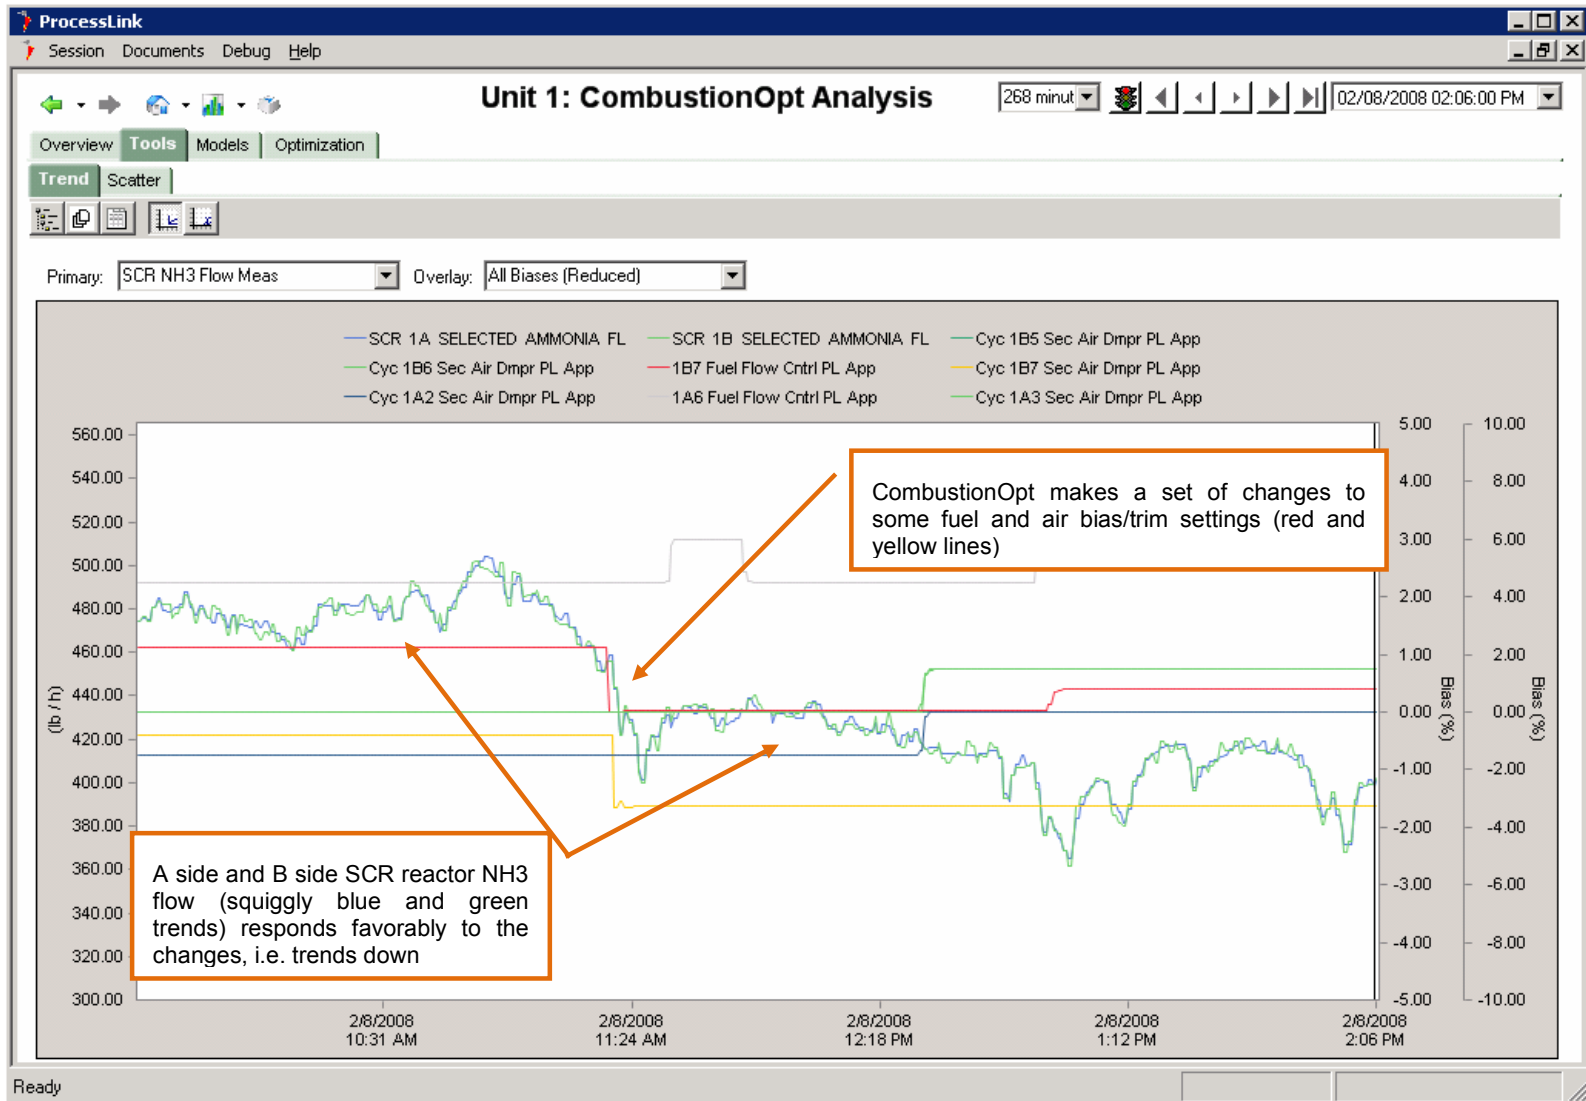


Figure 4 CombustionOpt and SCR-Opt in coordinated action. The Primary (left) y-axis show Ammonia flow in lbs/hr. The squiggly green and blue trends represent the A and B-side reactor ammonia flow measures. The Overlay (right) y-axis represents the deviation of the fuel and air biases from their neutral setpoint.

3.2.1.3 Online Learning

A key feature of NeuCo's neural Optimizers, such as CombustionOpt/SCR-Opt, is the use of Online Learning algorithms, which play a significant role in maintaining the accuracy of the neural models over time and in expediting project start-ups.

Without the use of rigorous, reliable calibration, optimization systems will degrade over time. Degradation that renders initial models obsolete can occur in weeks, requiring an ongoing cycle of recalibration and its associated costs. A regression model-based system must be able to update itself by analyzing new data in order to track plant changes and thereby remain current and accurate. In this way the recommendations prove reliable and the operators remain confident in the technology.

NeuCo's approach to Online Learning uses a combination of design-of-experiments methods with online model training and state-of-the-art model validation techniques. ProcessLink's neural Optimizers automatically adapt to changing plant conditions by continuously learning.

NeuCo's Online Learning uses competing pools of individual models to elect model "committees" of the best performers against the most recent operating data. This ensures that the models used are those that most accurately reflect the generating unit's actual operating conditions.

Figure 5 shows actual and "neural network predicted" CEMS NO_x for two models. The bottom one has re-tuned nightly, the top one was trained to the same accuracy as the first initially, but was then removed from the automatic retuning process. Clearly the response of the process to control changes is not exactly the same as it was when both models were initially trained. The model that has been learning all along has more accurate model predictions, which help ensure that the Optimizer is providing robust and reliable control recommendations for obtaining optimal performance.

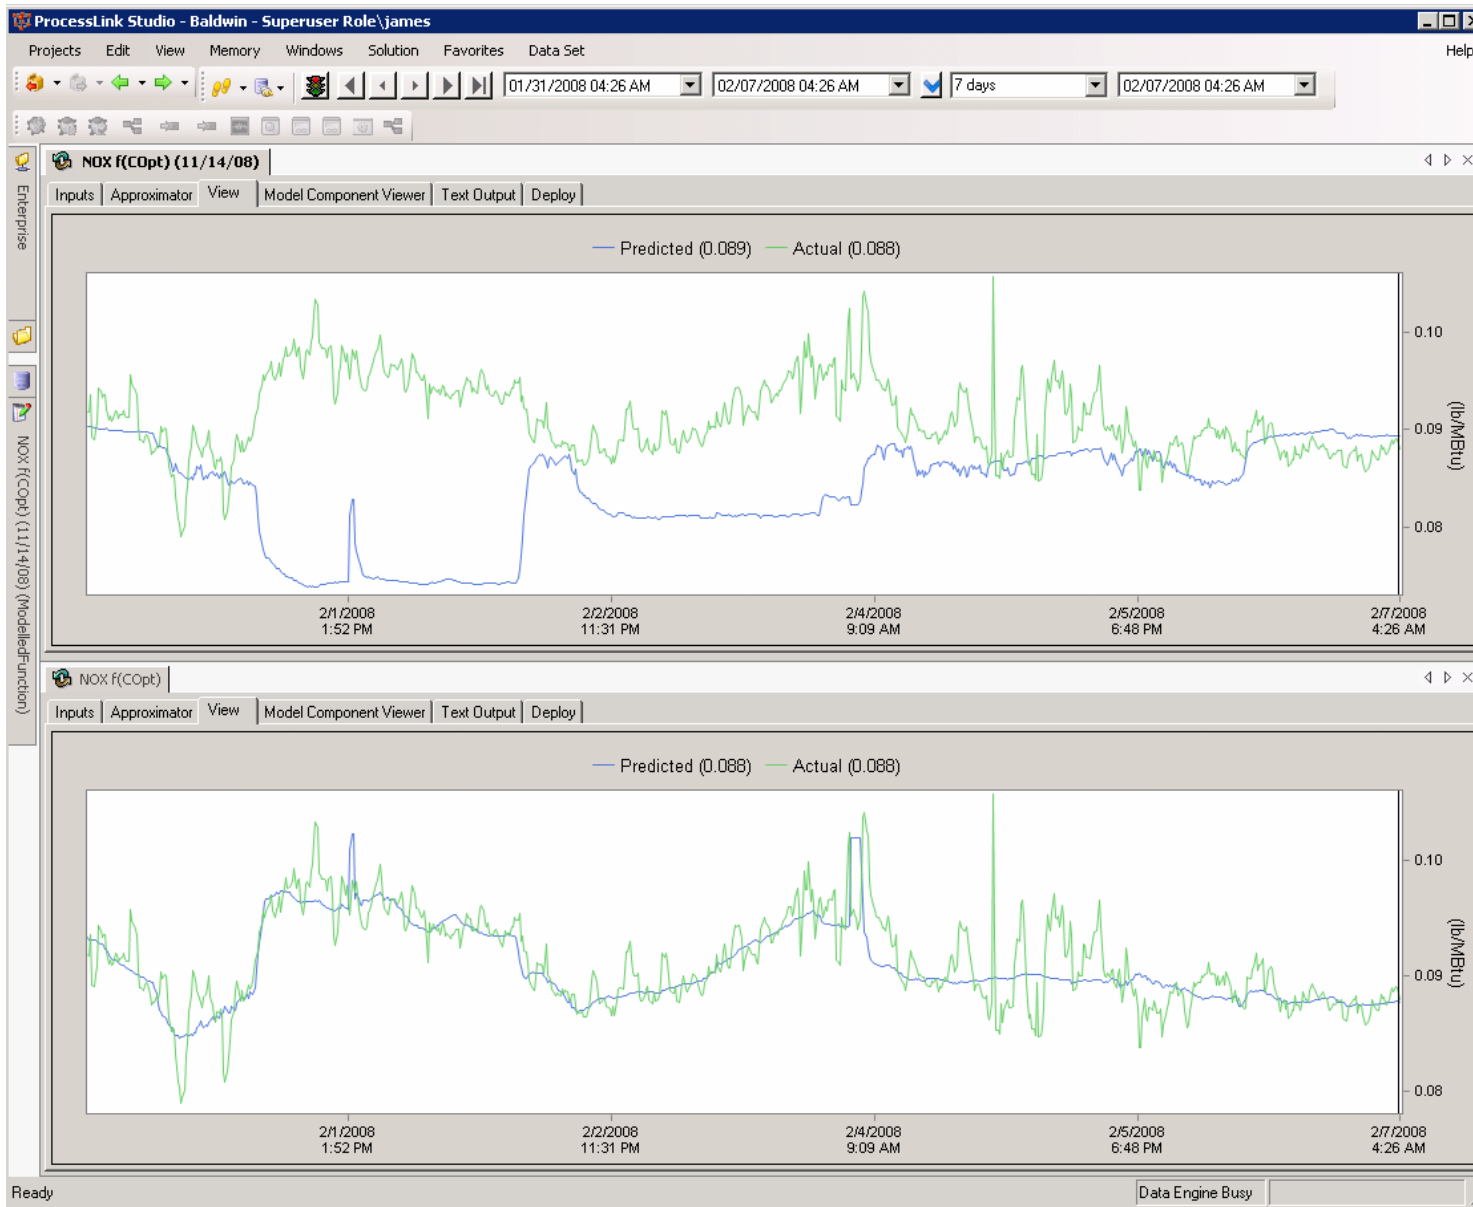


Figure 5 Model Accuracy with Online Learning.

In Figure 5 actual NOx is green, Blue is the NOx prediction given by a neural net model of NOx as a function of the fuel and air distribution biases. The upper plot shows model accuracy before on-line learning (auto-adaptive regression). Lower plot shows fully trained regressor. The time span is 7days. When trained, the model can predict with meaningful accuracy what NOx will do, just by looking at the settings of the fuel and air distribution biases.

The design of experiments methodology provides the exploration component of Online Learning enabling NeuCo's neural models to automatically learn important relationships that may be poorly represented in existing plant data without the assistance of an engineer or an operator. NeuCo's Optimizers continue to use the design of experiment routines throughout its life, when necessary, to explore new input spaces.

Figure 6 shows the direct search algorithm moving selected previously seldom used biases one at time, first in one direction, then either keeping the move, or moving it in the other direction. Where there is no movement it is either waiting for steady-state or cycling through biases not shown here. Once the experiment cycle is complete for all biases, meaning it has tested each bias in turn, it makes a combined move with all the biases at once, in the direction of improvement. This is a fully automated model free algorithm that actively searches for the best response from the actual process using levers that have an unknown effect Useful in its own right (though less efficient than a model-based search), it is also good at generating rich data on which to train neural network models.

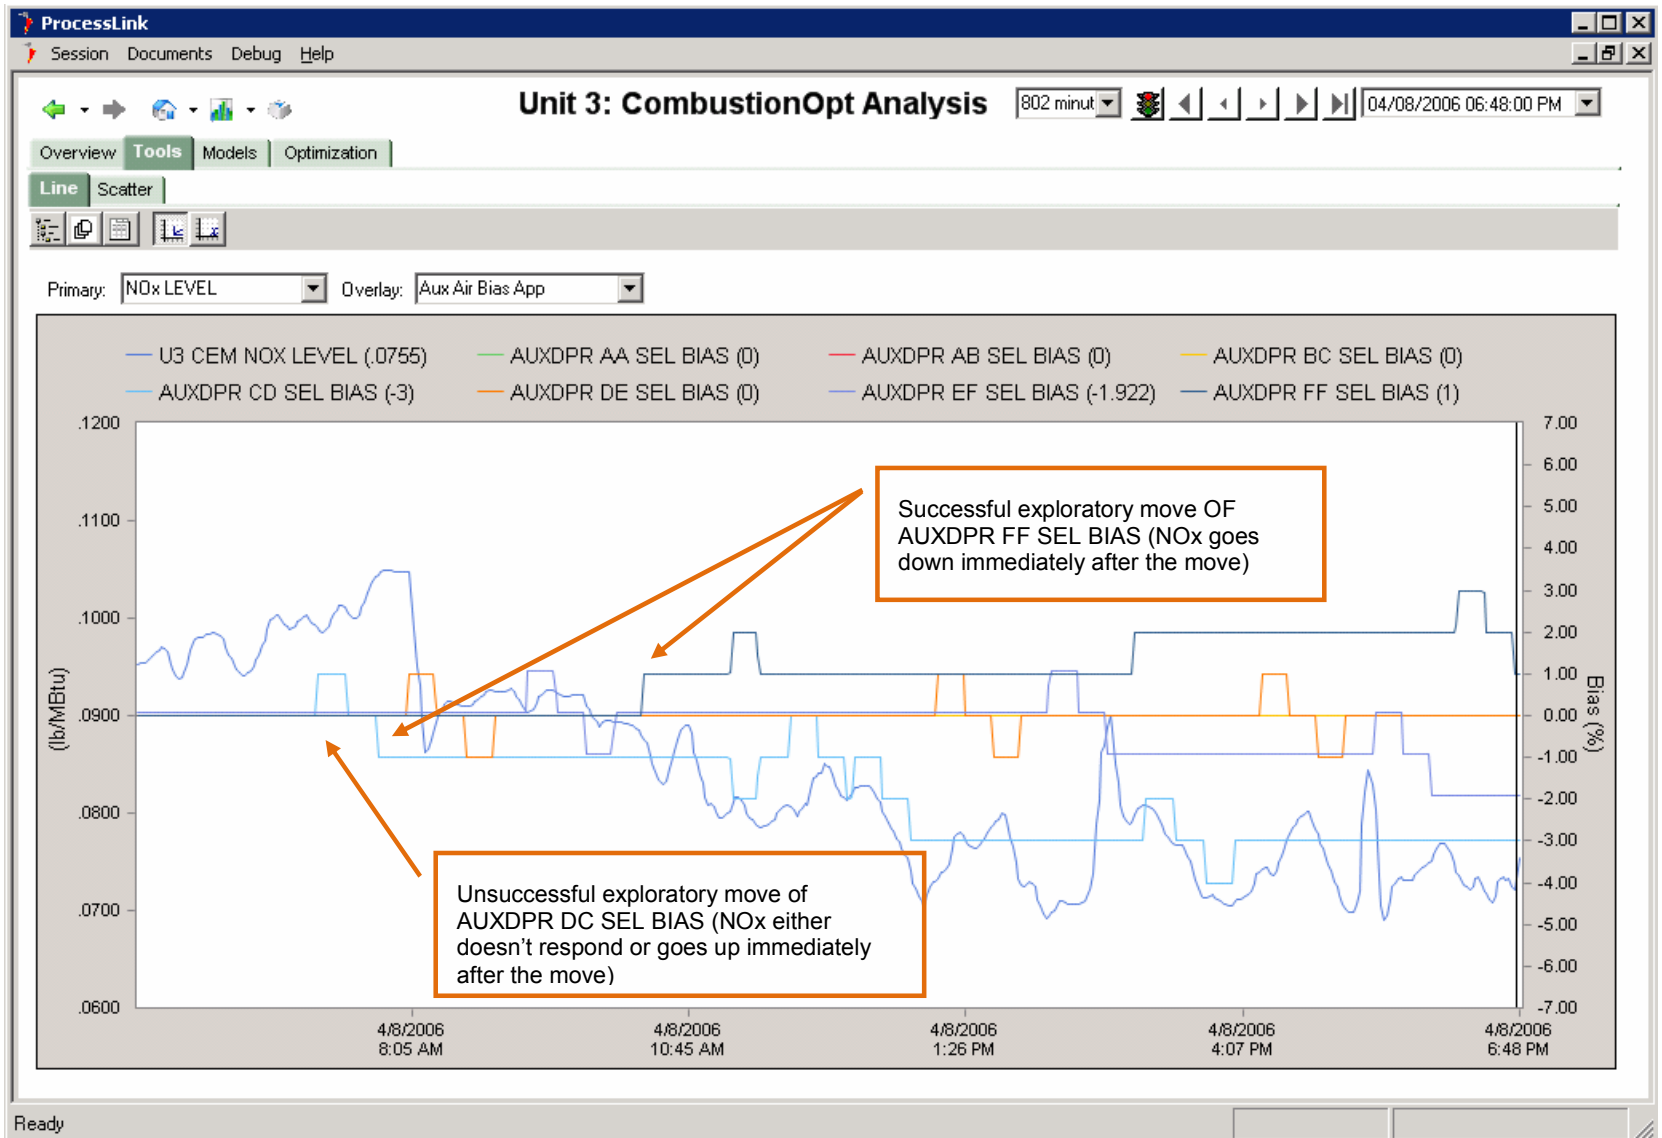


Figure 6 Direct Search in Action. Primary (left) y-axis shows NOx in lb/Mbtu (blue squiggly trend). The Overlay (right) y-axis shows the values of a set of air damper position biases as they are adjusted by CombustionOpt in Direct Search Mode.

3.2.1.4 Advanced Modeling functionality

All of NeuCo's neural Optimizers are able to leverage ProcessLink's advanced modeling functionality to provide the following:

- **Monetized tradeoffs: Where appropriate,** each Optimizer prioritizes its objectives and manages tradeoffs between them based on their monetary value. Where dollar-cost assignment does not completely describe real priority, the features of the utility function are exposed and demystified so that engineering unit priority is assigned with confidence and precision.
- **Condition-based optimization profiles:** Optimization profiles are pre-packaged sets of unit marching orders that include objective definitions, constraints, response times and step sizes. Different operating environments at times require different optimization profiles to obtain the best unit performance. ProcessLink enables the automated modification of optimization profiles based on predetermined cues, and enables Optimizer users to script these condition-based rules.
- **Dynamic and discrete optimization:** each Optimizer can address a variety of operating situations and challenges including fast ramping units and burners/mills out-of-service. ProcessLink's ability to determine the optimal set-point trajectory between current and optimal states and recalculate it with every move, means that the Optimizers can anticipate as far ahead as necessary for dynamic situations. They can also incorporate discrete decision parameters into their optimization profiles for a more comprehensive solution.

3.2.1.5 System flexibility

The flexibility of NeuCo's neural Optimizers enables them to respond to the dynamic nature of today's power plants and to constantly changing objectives and constraints. Each Optimizer leverages ProcessLink's advanced enterprise architecture to offer scalable, extensible and flexible solutions.

The scope of an optimization project can be easily modified or expanded to incorporate new controls and objectives, or to address additional optimization challenges across the plant floor or up the enterprise. Each neural Optimizer can rapidly accommodate itself to changing conditions, inputs, controls and objectives, and supports total flexibility around the definition, relative priority and configuration of objectives and constraints.

The open nature of the ProcessLink architecture ensures that all Optimizers can easily work with existing or future systems, programs and devices. This enables a high degree of integration while lowering costs and minimizing disruption that comes with learning new systems as the power plant's physical make-up and character changes.

3.2.2 The User Interface

CombustionOpt and SCR-Opt share a Home page (see Figure 7 below) that provides three kinds of information that enable users to derive the maximum value from their system:

Optimization Advice

The upper left section of the Home Page shows current advice for how to further optimize the unit, along with tools and information to help prioritize and specify actions. CombustionOpt provides advice about actions that it cannot take directly, but that if taken by the user, can help deliver more benefit.

What's Going on Now and Why

The lower left section of the Home Page provides information that explains why the closed-loop Optimizer made the move it did. It displays the optimization objectives in a way that gives insight into how much attention is being paid to each objective. It also displays what manipulated variables (MVs) were most recently moved by CombustionOpt and SCR-Opt, and how much each MV was moved. The iconized trend in the center is a minimized version of different "trend graphs" that contains key plant state variables and optimization objectives; the values that are trended are specific to each site, but typically include gross MW, net MW, heat rate and NO_x (or other key emissions parameters). These iconized trends allow the user to quickly (by double-clicking) assess how the parameters have been trending over a specific time period.

Optimization Benchmarks

The section on the right of the Home Page shows how the unit has been doing with respect to its optimization objectives as well as how much better it might have done over the past 28 days. The Benchmark bar chart shows three values for each variable, where the variables represent optimization objectives such as NO_x, Heat Rate, and NH₃ Flow. The three values represent the Achievable, Actual, and Baseline benchmarks for that objective. An *achievable benchmark* represents the performance with respect to objectives and constraints that would have been achieved if all known optimization opportunities, both closed-loop and user actuated, had been taken. The *actual benchmark* value represents how the process being optimized has actually been doing, or in other words, the cost actually incurred. The *baseline benchmark* represents the performance with respect to objectives and constraints that would have occurred if there had been no optimization.

Note that in "vs. Achievable", the values for some of the objectives may be negative (indicating that CombustionOpt and/or SCR-Opt needed to incur additional costs with respect to those objectives in order to achieve optimal savings overall.)

The Triband chart in the upper right hand corner of Figure 8 below displays a single value for all objectives, in \$/MWhr, calculated at each moment. These values represent the benefit achieved through optimization (difference between actual and baseline), and the benefit that was not achieved but could have been had all identified optimization actions been taken (difference between achievable and actual.)

CombustionOpt and SCR-Opt also provide many analysis screens. Note that on the screen shown in Figure 8 below there are five tabs (Overview, Tools, Models, Optimization and Benchmarking). In this particular case, the user has navigated to the Optimization tab, and the Objectives/Constraints Sub-tab, where information about the distribution of the optimizer's attention across its list of objectives for each optimization move made in closed-loop is provided.

Other analysis views provide access to basic Trend and Scatter Plotting, as well as a wealth of Model Analysis tools that let the user leverage the knowledge being extracted nightly from plant data by the neural models as they learn online. Figure 9 below shows the Causality Analysis available from the Model Tab. It lets the user select a particular change in a variable of interest displayed in a trend format, the cause of which is then returned by the models as a ranked bar chart of inputs to that change.

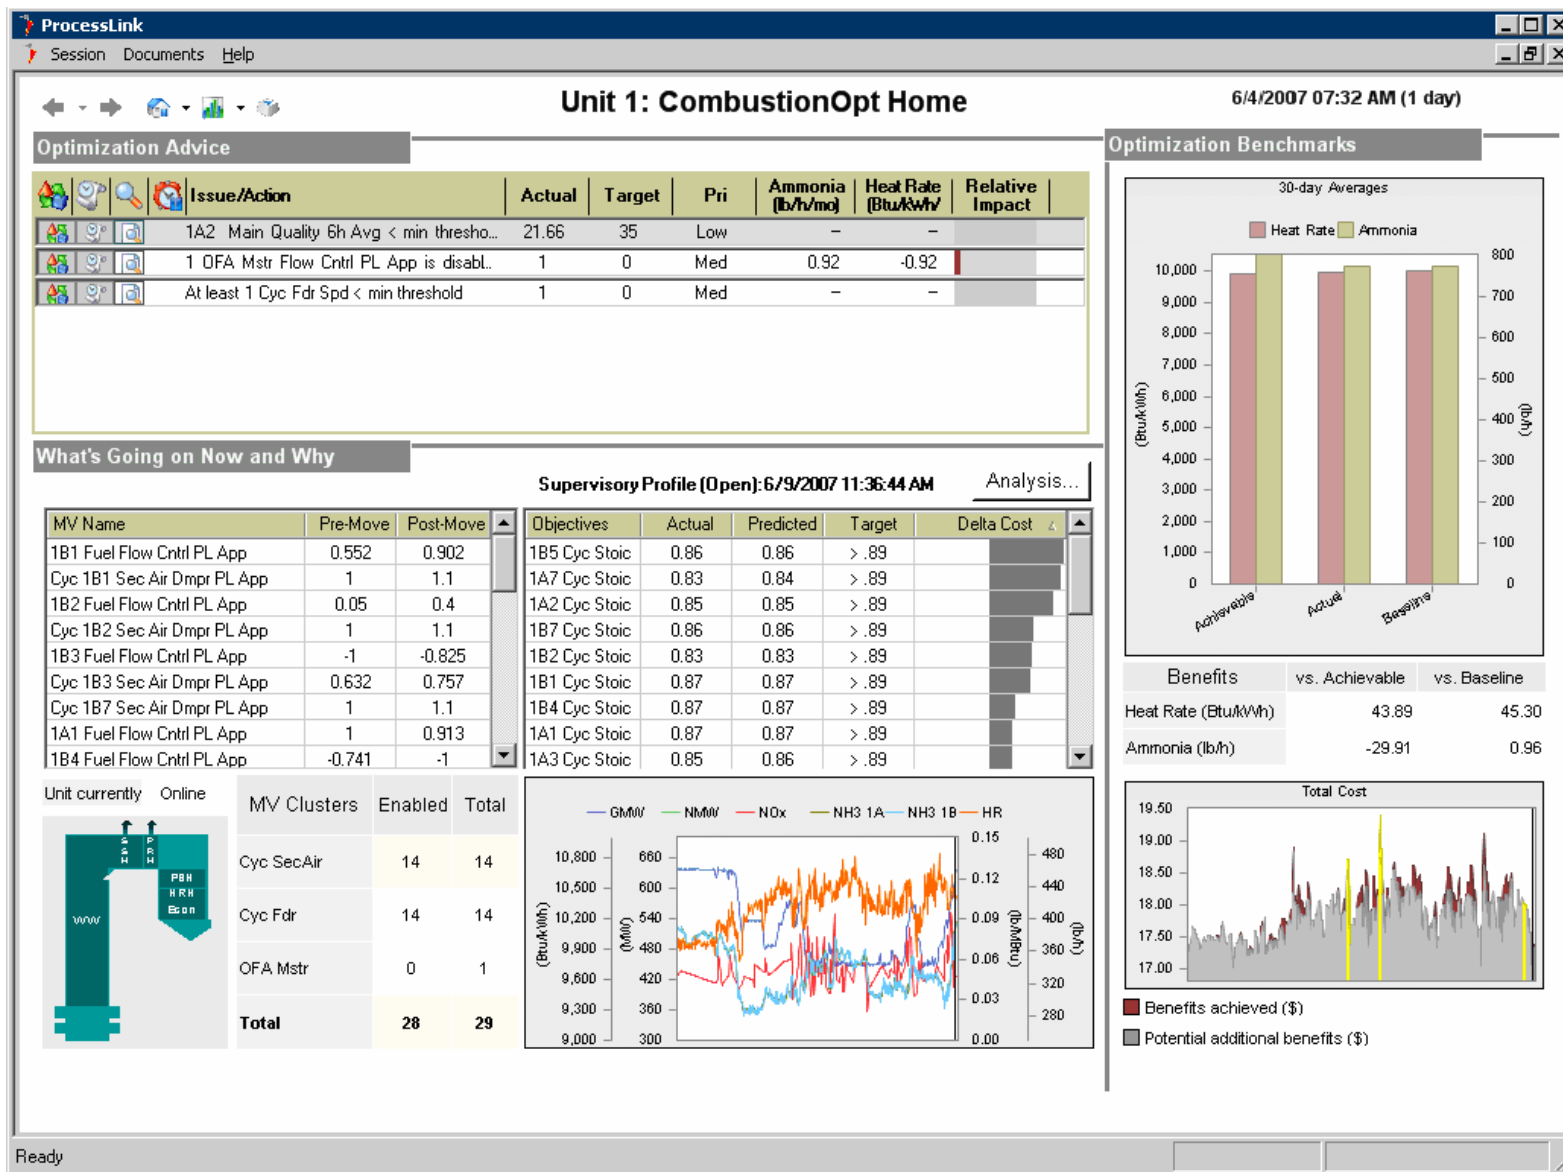


Figure 7 CombustionOpt/SCR-Opt Home page

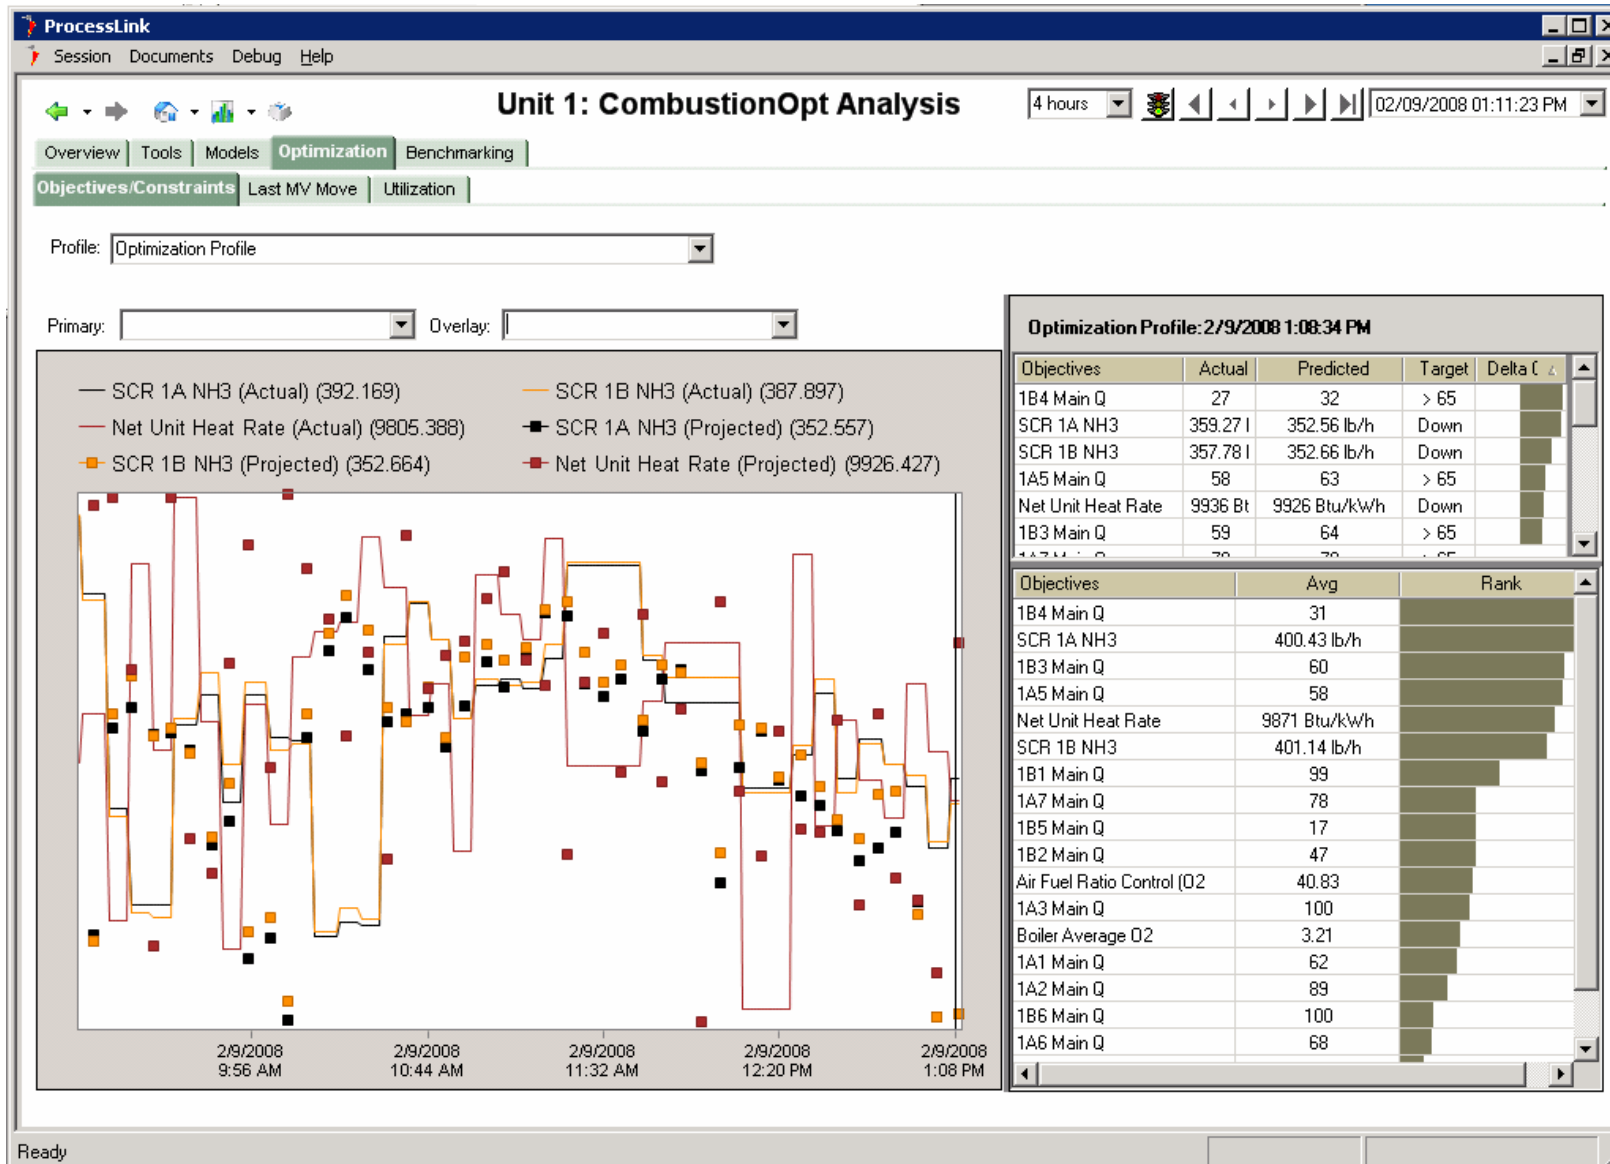


Figure 8 CombustionOpt/SCR-Opt Analysis Screens

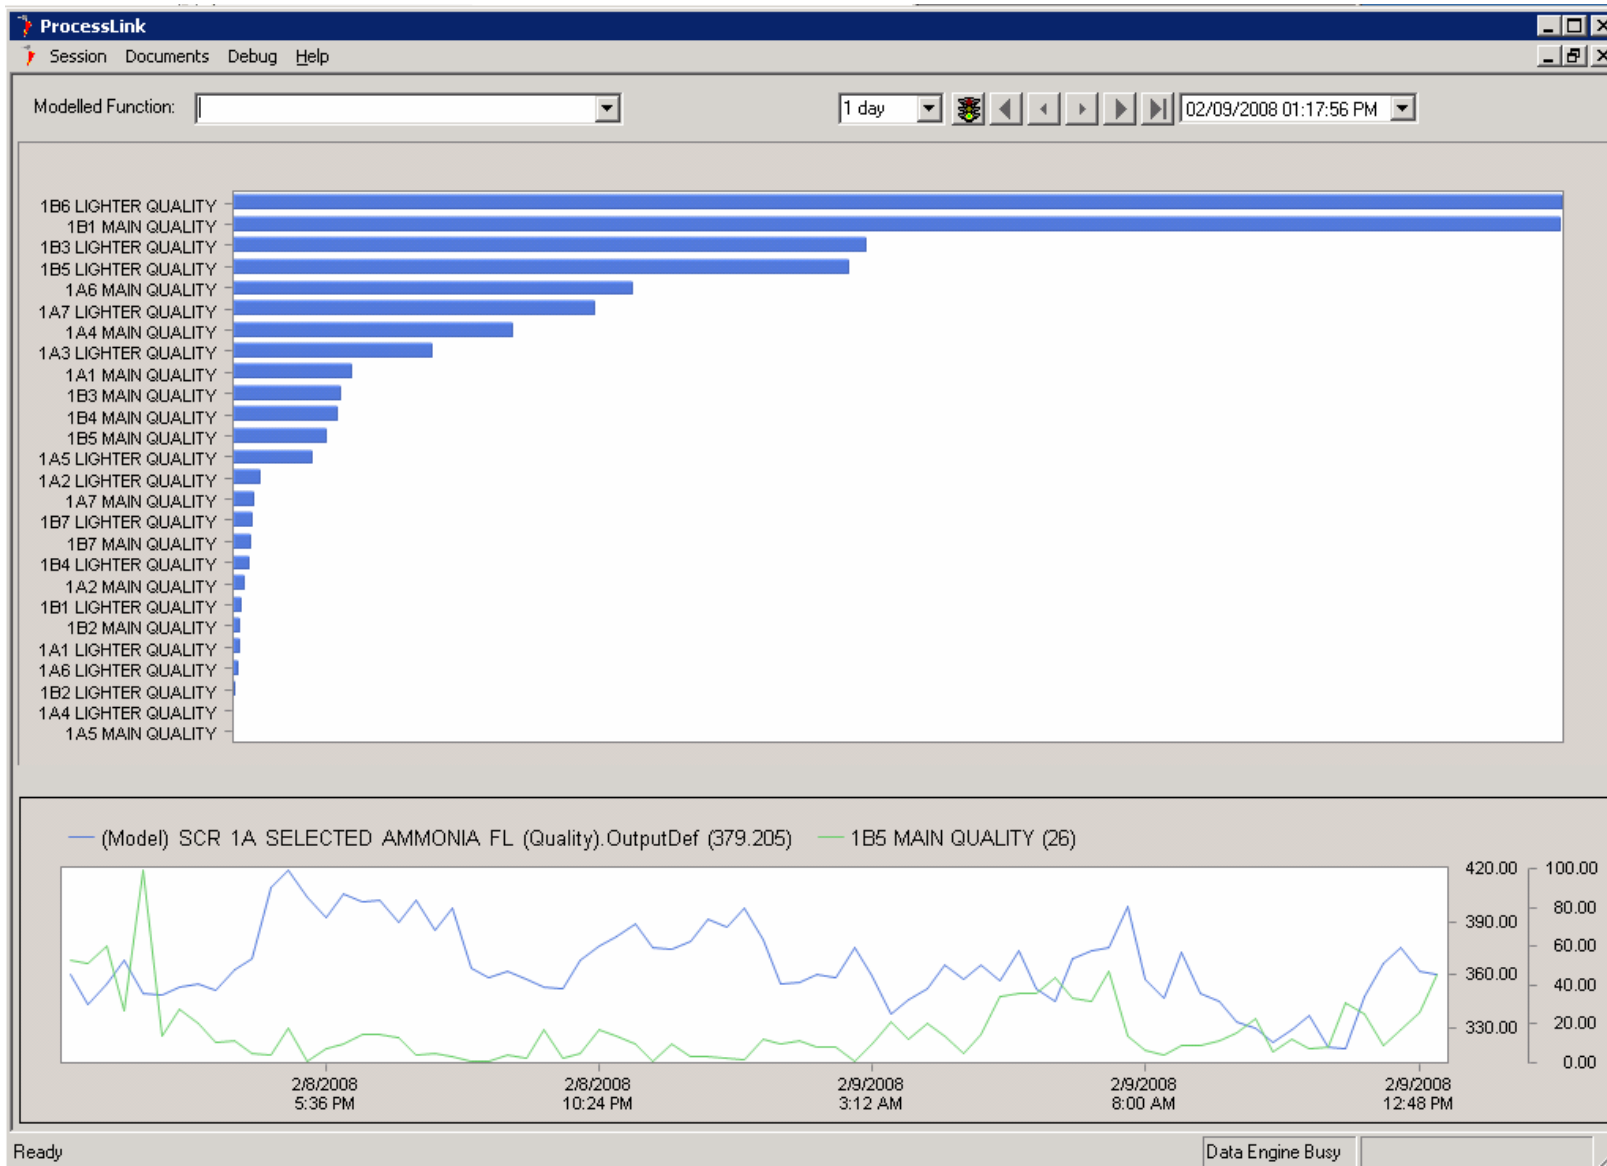


Figure 9 Causality Analysis

3.3 SootOpt

3.3.1 How It Works

SootOpt is a closed-loop soot blowing optimization system that takes globally optimal soot blowing actions that factor in heat rate, emissions, reliability impacts, and operational constraints. SootOpt models the effect of soot blowing activity on heat transfer throughout the furnace and backpass and dynamically determines the boiler cleaning actions to best achieve improved control of steam temps/spray and exit gas temps, reduce heat rate, minimize NOx with the minimize number of cleaning operations.

SootOpt uses a combination of optimization methods in conjunction with available direct measurement and local controls to relate boiler cleaning actions to global performance objectives and to adjust or establish soot blowing requirements. In addition to adaptive modeling techniques, SootOpt leverages customized operational constraints and control considerations, in the form of rules or heuristics, to identify the correct response to different operating conditions, such as when soot blowing is required due to suboptimal steam temperatures or high sprays, or when it should be suspended due to the same, or when soot-cleaning media limitations dictate coordination of activity. Local instrumentation, such as FEGT, heat flux measurements, strain gauges or rigorous first principles-based PerformanceOpt cleanliness factors can be added to increase the resolution of understanding of and response to soot-cleaning actions. Because of the flexibility of the underlying ProcessLink architecture, SootOpt can be deployed to take advantage of whatever equipment is already in place, as long as it can be managed through digital systems. The key strength of SootOpt is its ability to seamlessly combine advanced and legacy controls and instrumentation, adaptive models of global objectives and unit-specific rules to not only consistently execute the appropriate base-line soot-cleaning protocols, but also optimize steam temps/sprays furnace economizer and aph exit gas temperatures, lower NOx and heat rate, and minimize unnecessary operations.

The overall solution results in improved consistency and quality of soot-cleaning decisions, improved insight into soot-cleaning activity and its effects on unit performance, and improved bottom line performance of emissions, heat rate, and reliability indicators.

Figure 10 shows a block diagram overview of the SootOpt system. The operating conditions (including the state variables) are input to a zone selection block that is used to determine which boiler zone needs to be cleaned. As described in the next section, an expert system determines the zone to be cleaned. Once the zone has been determined, an optimization algorithm that utilizes predictive models determines which blower or set of blowers to activate in the associated zone. The algorithm selects the blower that is expected to provide the best boiler performance in the future based upon current operating conditions.

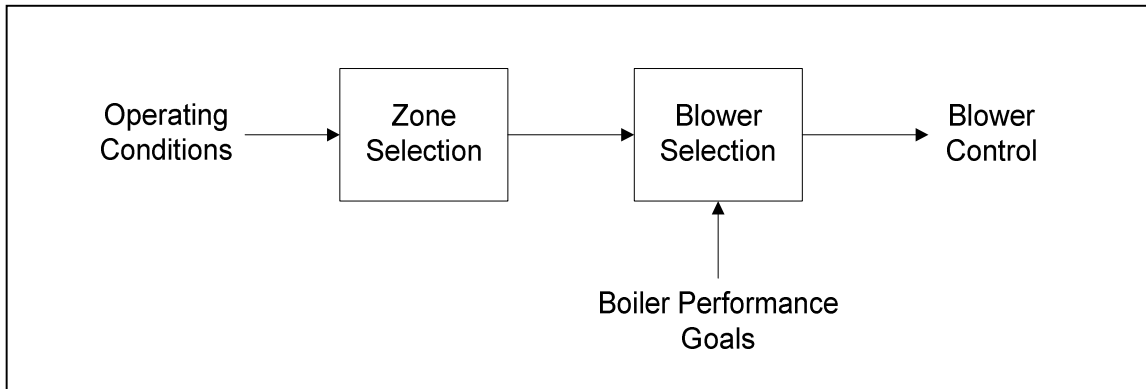


Figure 10 Block diagram of the SootOpt system.

Figure 10 shows the typical implementation of the SootOpt system. However, because SootOpt is configured and implemented in ProcessLink, it is possible to realize different implementations of SootOpt depending upon the requirements of the application.

For example, Figure 11 shows an alternative implementation. In this example, SootOpt computes the cleanliness factors needed for a criteria-based sootblowing system. As shown in Figure 11, a neural network based model determines the effects of varying the cleanliness factors on boiler performance parameters (i.e., heat rate and NOx). Using this model and an optimizer, SootOpt computes the optimal cleanliness factors based upon desired boiler parameters. Other implementations of SootOpt can be realized using ProcessLink depending upon the system requirements.

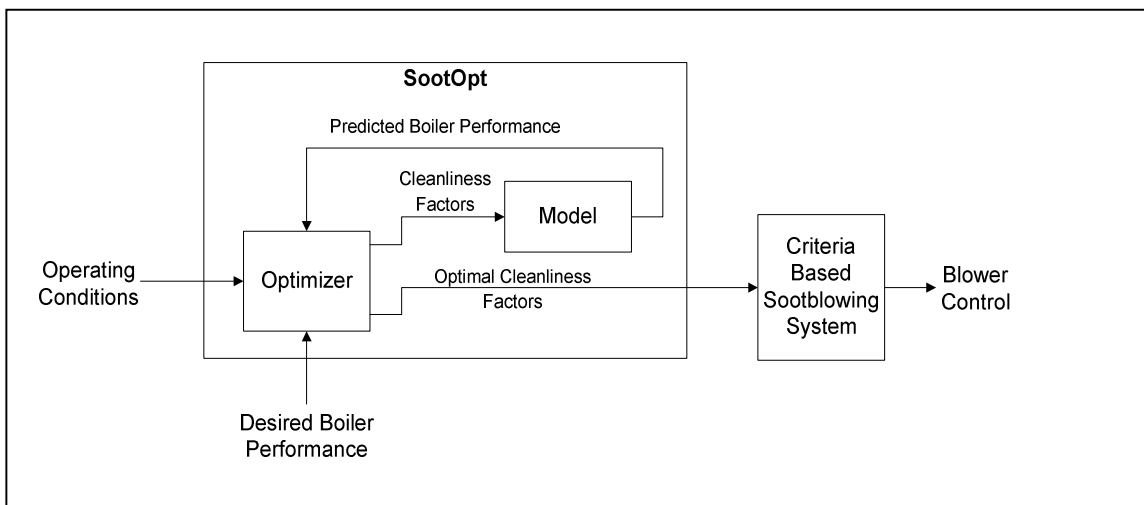


Figure 11 An alternative implementation of SootOpt. In this example, SootOpt computes the optimal cleanliness factors of a criteria based sootblowing system.

Finally, SootOpt can be combined with other NeuCo optimization solutions, such as CombustionOpt, to improve boiler performance. For example, CombustionOpt adjusts the boiler's fuel and air biases to lower NOx and improve heat rate. CombustionOpt computes the resulting fuel and air biases and inputs them to SootOpt, which then takes the effects of these changes into account when determining the optimal sootblowing sequence. Similarly, the sootblowing sequences which SootOpt determines can be input into CombustionOpt so that sootblowing effects are taken into account when adjusting fuel and air biases in the boiler. Because SootOpt and

CombustionOpt both use the ProcessLink platform, integrating these systems is relatively seamless.

3.3.1.1 Zone Selection

To best utilize expert knowledge in the SootOpt system, the boiler and associated sootblowers are divided into zones. For example, a boiler may be divided into furnace, reheat, superheat, economizer, and air preheater zones.

The goal of the zone selection component of Figure 10 is to determine the best boiler zone to clean given current operating conditions. SootOpt accomplishes this objective by using the expert system shown in Figure 12. The expert system is composed of three primary components: the inference engine, the knowledge base of propose rules, and the knowledge base of apply rules. The inference engine is a standard component of ProcessLink and allows the optimizer to achieve prioritized actions based on a knowledge base of rules. The knowledge base (i.e. propose rules and apply rules) may be determined through expert knowledge sources such as engineers, textbooks and journals.

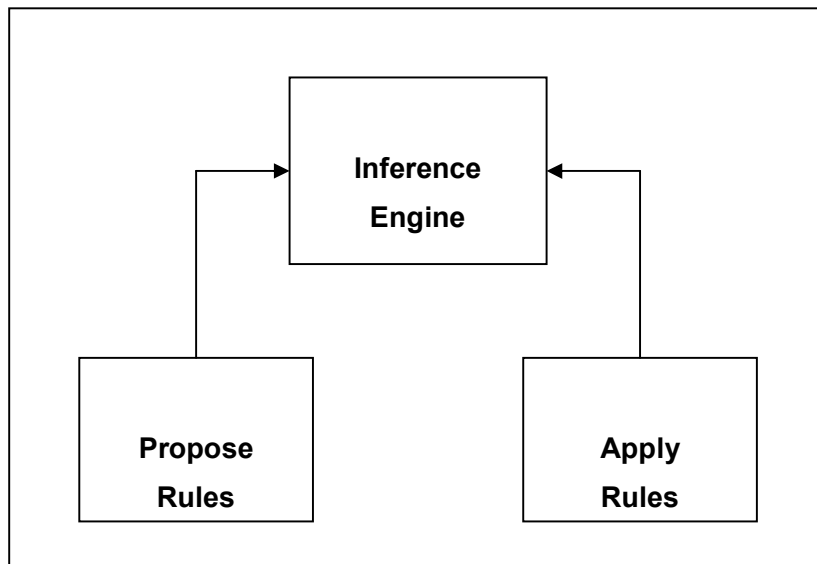


Figure 12 Expert system for selecting zone to be cleaned.

The Propose Rules knowledgebase of Figure 12 is comprised of propose rules. Propose rules are used to generate possible actions that will address any current issues. Figure 12 illustrates a sample set of propose rules. The first rule proposes the action of cleaning the furnace zone to reduce superheat spray because it is currently too high. The inference engine of Figure 12 proposes this rule only if (a) the trigger conditions and (b) the enabling condition are satisfied. For example, with reference to the set of propose rules of Figure 12, one of the trigger conditions that must be satisfied is that the superheat spray is greater than a superheat spray threshold. Before proposing the action, the inference engine uses the enabling condition to determine if sootblowing can currently be initiated in the zone.

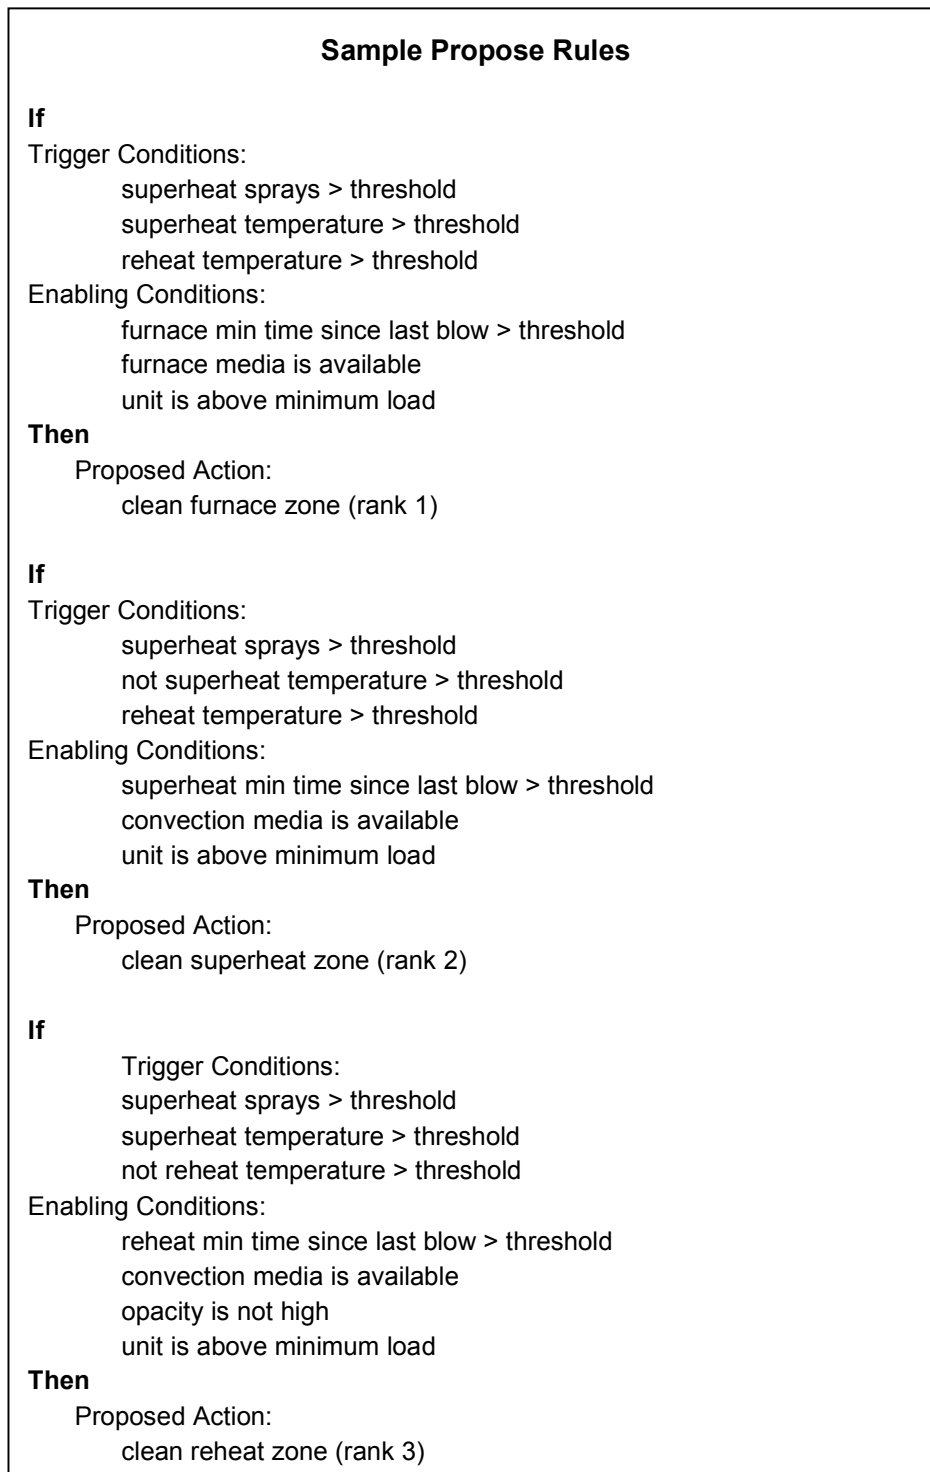


Figure 13 Sample propose rules

The inference engine uses the propose rules to generate a list of potential actions. The apply rules of the rules knowledgebase are subsequently used to determine which of these potential actions should be taken. The simplest form of an apply rule is one that

selects the action with the highest rank. For example, if the first propose rule of the sample shown in Figure 4 is satisfied, then the action of cleaning the furnace zone would be selected because it has the highest rank.

Another advantage of the propose-apply approach described above is that the apply rule can be based upon the dollarized (monetized) effect of a proposed action. For example, the propose rule may be written such that an action is proposed if the predicted savings of the action is above a user specified threshold. Thus, the following propose rule could be constructed:

```
if Dollarized_Effect_of_Cleaning_Furnace > Dollar_Threshold and Cleaning
Furnace Zone is Enabled

then Propose Action of Clean Furnace Zone and Goal Rank is equal to
Dollarized_Effect_of_Cleaning_Furnace.
```

The variable `Dollarized_Effect_of_Cleaning_Furnace` is computed using a model that predicts the effects of furnace cleaning on NO_x emissions and heat rate. The predicted change in NO_x emissions and heat rate is multiplied by the current NO_x credit value and fuel costs to determine the cost savings associated with the cleaning event. In order for this rule to trigger, this savings must be greater than a user specified threshold represented in this rule by `Dollar_Threshold`. Finally, it should be noted that the goal rank is equal to the expected savings variable, `Dollarized_Effect_of_Cleaning`.

Various propose rules that represent cost savings of cleaning different regions of the boiler can be implemented. Notice that these propose rules specify an action that is ranked based on dynamically determined cost savings rather than on a fixed ordering. In order to select the appropriate action, a more sophisticated apply rule needs to be included:

```
if a proposed action with a fixed rank exists (such as those in Figure 13),
then select the action with the lowest rank
else select the proposed action with the highest dollarized rank.
```

Using this approach, the propose rules of Figure 13 are first checked. If no action is proposed by these rules, then the economic savings rules of the type given above are checked. If significant economic savings are found, then the action that results in the largest savings is taken. This example shows the flexibility of the propose-apply approach.

An advantage of the propose-apply approach described above is that the apply rules can be used to effectively combine propose rules. For example, if the same action is proposed by multiple rules, the rank of the proposed action can be re-evaluated by an apply rule and selected if it is higher than the rank of any other action.

Another advantage of the propose-apply approach described above is that the apply rules can be adaptive or based on neural network model(s). For example, the sootblowing optimization system can dynamically adjust the ranks of actions based on boiler performance. Alternatively, neural network models may be used to determine the effects of cleaning a zone on boiler performance. The resulting boiler performance can then be used to adjust the ranks of the possible actions. By separating inferencing into two sets of rules (i.e., propose and apply), the sootblowing optimization system provides great flexibility for appropriately selecting the zone to clean in a boiler.

The expert system used in sootblowing optimization provides several advantages:

1. **Prioritizing Actions:** Engineers can specify an a priori ordering on the various actions that can be taken. Because these priorities may change based upon

current operating conditions, an action's rank can be dynamically changed at run-time by the sootblowing optimization system using apply rules.

2. Rules Design: To simplify knowledge capture, engineers only needed to collect propose and apply rules. Also, it is possible to add rules at any time to a rules database in order to improve performance.
3. Demystification: Using an inference engine, the conditions that resulted in the selection of a zone to be cleaned may be displayed to a user on a computer interface (e.g., a computer monitor). Thus, the expert system approach of the present invention can provide transparency into the operation of the zone selection algorithm.

3.3.1.2 Blower Selection

Given a selected zone for sootblowing, the blower selection component of SootOpt is used to determine which blower to activate within the zone. Figure 14 shows an overview of the blower selection component used in the typical implementation of SootOpt. Blower selection is composed of a scenario generator and a scenario evaluator. The scenario generator creates a complete set of soot blowing scenarios for the zone given current operating conditions. The scenario evaluator then determines which scenario (blower activation) results in the best predicted future boiler performance.

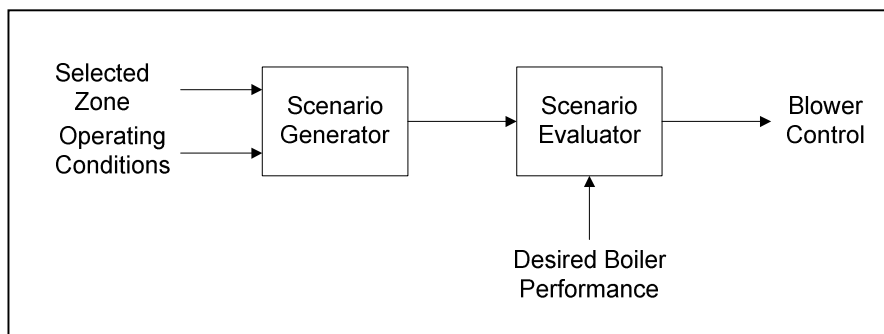


Figure 14 The Blower Selection Component of SootOpt

Figure 15 provides a flow chart of the scenario generator. The scenario generator first determines if any of the blowers have violated a maximum time limit since last blowing in the selected zone. If so, this blower is selected for activation and a single scenario is generated. (If multiple blowers are over the time limit, the blower that is the most over the limit is typically selected for activation.) By monitoring time limits, SootOpt guarantees that any constraints upon the system are observed before attempting to optimize performance.

If no time limits have been violated, the scenario generator identifies all blowers that can be activated using the enabling conditions described in the previous section. Next, SootOpt generates a scenario for activating each identified sootblower. For example, if three sootblowers in the selected zone are enabled, then three separate scenarios would be generated for activating each of these sootblowers. At the end of the scenario generation, a set of activation scenarios are available for evaluation.

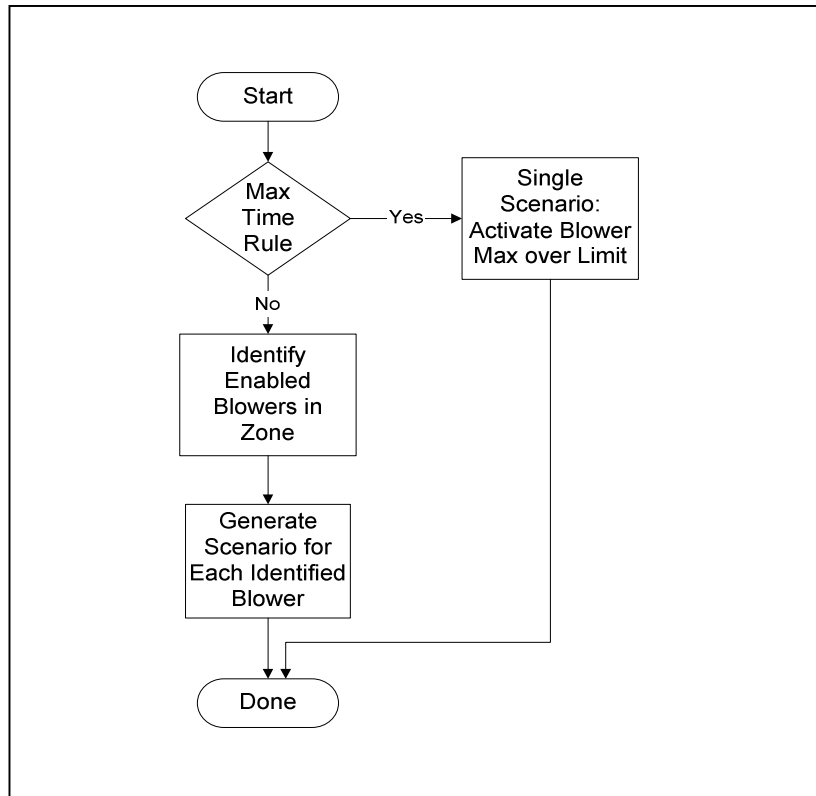


Figure 15 Flow chart of the scenario generator in SootOpt

Figure 16 provides an overview of the scenario evaluator. On the left side of diagram, each of the scenarios identified by the scenario generator is listed. Each scenario contains a list of the history of sootblowing activations such as time since start of last activation of each blower. In addition, the scenario may contain data associated with current operating conditions such as load. In each scenario, one sootblower is selected for activation by the scenario generator; therefore, the history of activation associated with that sootblower is modified to reflect turning the blower on at current time (i.e., time since last activation is modified to be equal to zero).

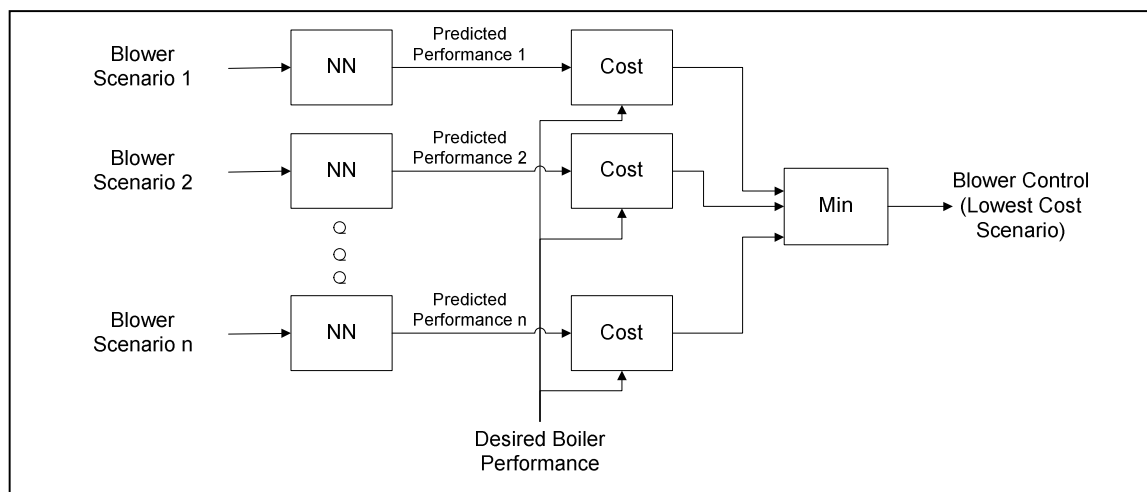


Figure 16 The scenario evaluator is used to determine the sootblower activation that minimizes a user-specified cost function.

Each of the blower scenarios is input to a neural network model that is used to predict future boiler performance. Thus, SootOpt can predict how activating different sootblowers in a zone will affect boiler performance factors such as heat rate and NO_x. It should be noted that an identical neural network model is used to predict the effects of activations on boiler performance. This model is trained upon historical data over a significant period of time. In addition, this model is typically automatically retuned daily so that any changes in boiler performance can be included in the latest blower selection.

As shown in Figure 16, the predicted performance and desired boiler performance parameters are used to compute a cost associated with each of the blower scenarios. The cost function may represent the actual cost associated with boiler performance or it may represent an artificial cost used to achieve a user specified boiler performance. For example, the cost function may be used to compute the cost of the predicted fuel usage and NO_x production. (In this case, heat rate, load, fuel cost and NO_x credit price are needed to compute these costs.) Alternatively, the cost function may be constructed so that heat rate is minimized while NO_x is maintained below a user-defined level. Engineers design the appropriate cost function so that lower cost represents better overall boiler performance.

Once SootOpt computes the cost of each scenario (Figure 16), it identifies the scenario with the lowest cost and selects the blower(s) to be activated. The final step of the SootOpt system is activating the selected sootblower through the communication interface. After activation, SootOpt waits a predetermined amount of time before starting the cycle again. Using this approach, SootOpt achieves optimal sootblowing and selects the lowest cost scenario that observes all system constraints.

3.3.2 The User Interface

Like the other Optimizers, SootOpt has a Home page (see Figure 17 below) that provides three kinds of information that enable users to derive the maximum value from their system:

Optimization Advice

The upper left section shows SootOpt's current advice for how to further optimize the unit, and in those cases where it can determine the impact that its advice has on its various objectives, it displays what that impact will be over the next 30 days. In the example seen here, SootOpt is notifying the user that automated activity has dropped below the user-defined minimum threshold and the furnace and convection areas are in manual control. The priority of the Primary SH and Economizer section activity threshold alerts is higher for this unit than other zones because these are horizontal regions and must be kept reasonably clean to prevent high differential pressures and potential furnace plugging.

What's Going on Now and Why

The lower left section provides information that explains why SootOpt just made the move it did. In this example, on a unit that includes Intelligent Sootblowing controls, the first activity we see here tells us that a rule in the ISB is suspending all cleaning in the furnace because an FEGT measurement is below a certain threshold. We see a list of the heuristics currently active. The story that summarized heuristics tell here, is that we have low FEGT and EEGT temperatures, and low steam temps and cleaning has been essentially suspended.

Optimization Benchmarks

The section on the right shows how the unit has been doing over the past day, week or month with respect to its reliability, efficiency and emissions objectives; it shows what benefits were achieved over that period from the moves SootOpt made (actual minus baseline) and what additional benefits would have been achieved all Recommended Actions in the Optimization Advice table had been taken (achievable minus actual).

SootOpt also provides many analysis screens. Note that on the screen shown in Figure 18 there are five tabs (Overview, Tools, Optimization, Benchmarking). In this particular case, the user has navigated to the Overview Tab, Activity Status Sub-tab, Sequence Operations View, which shows summary statistics regarding the activity of the sootblowing equipment being dispatched by SootOpt.

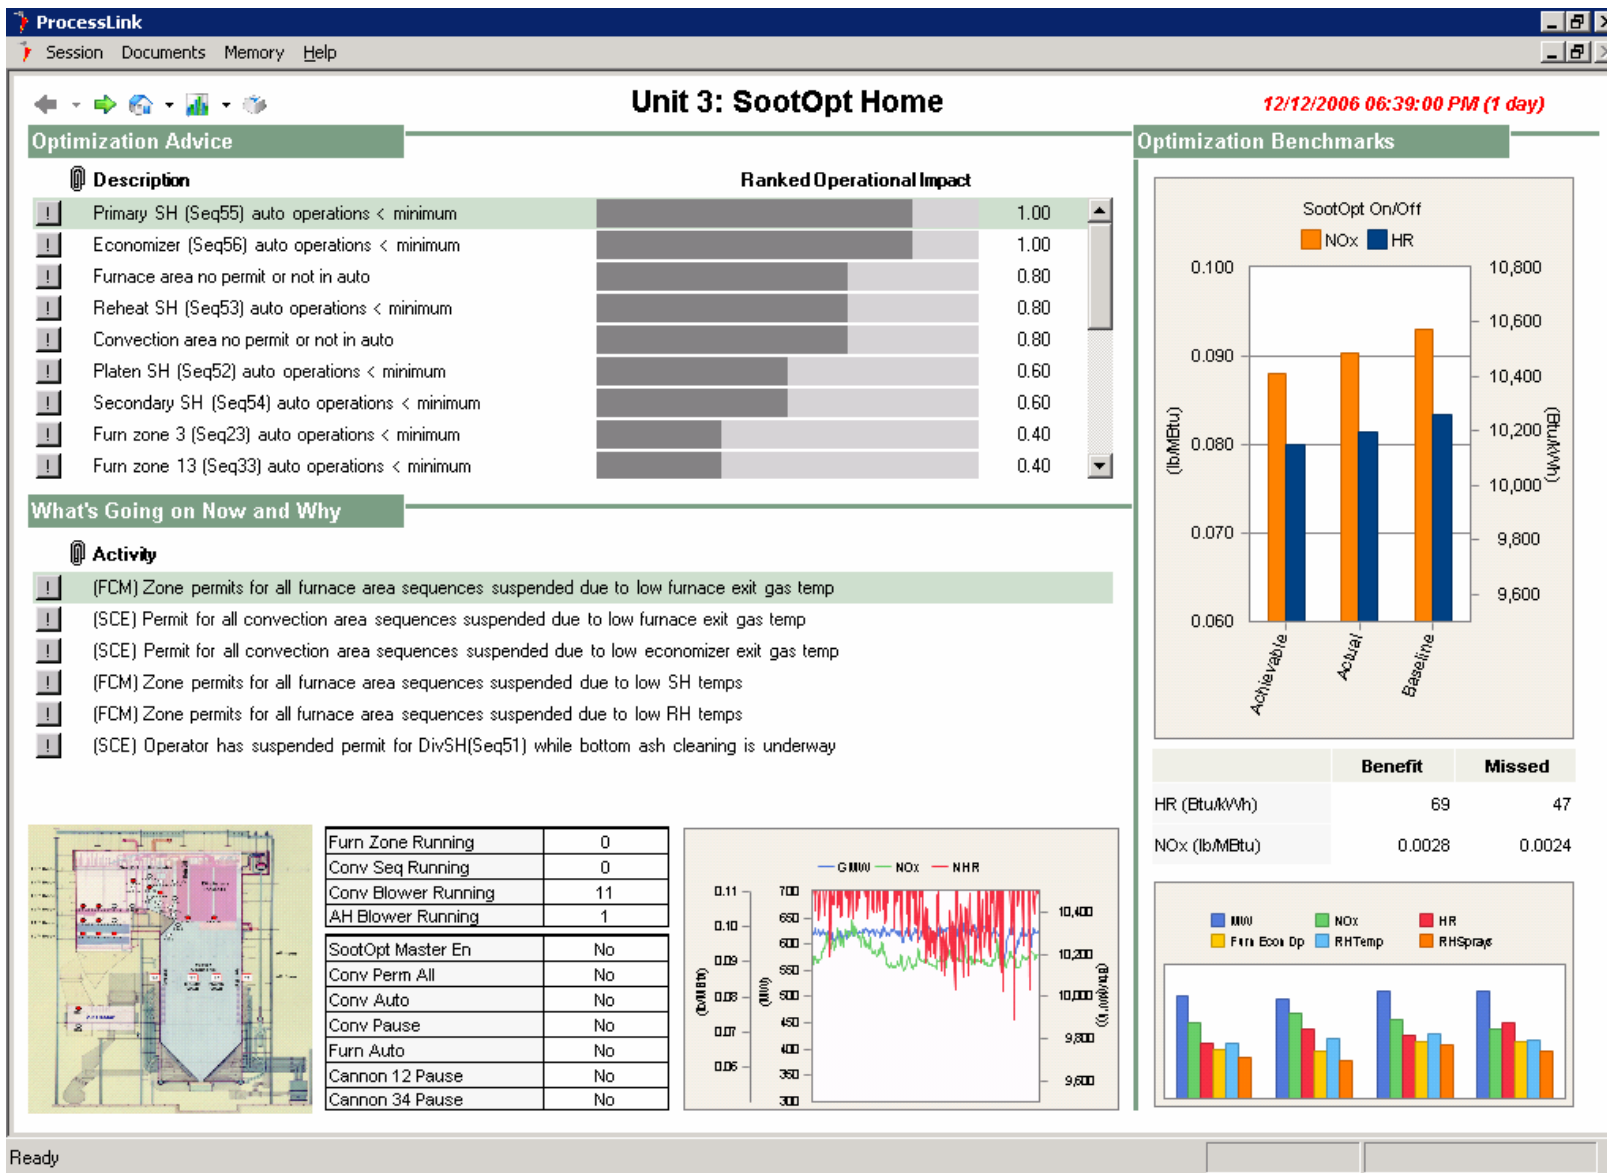


Figure 17 SootOpt Home page

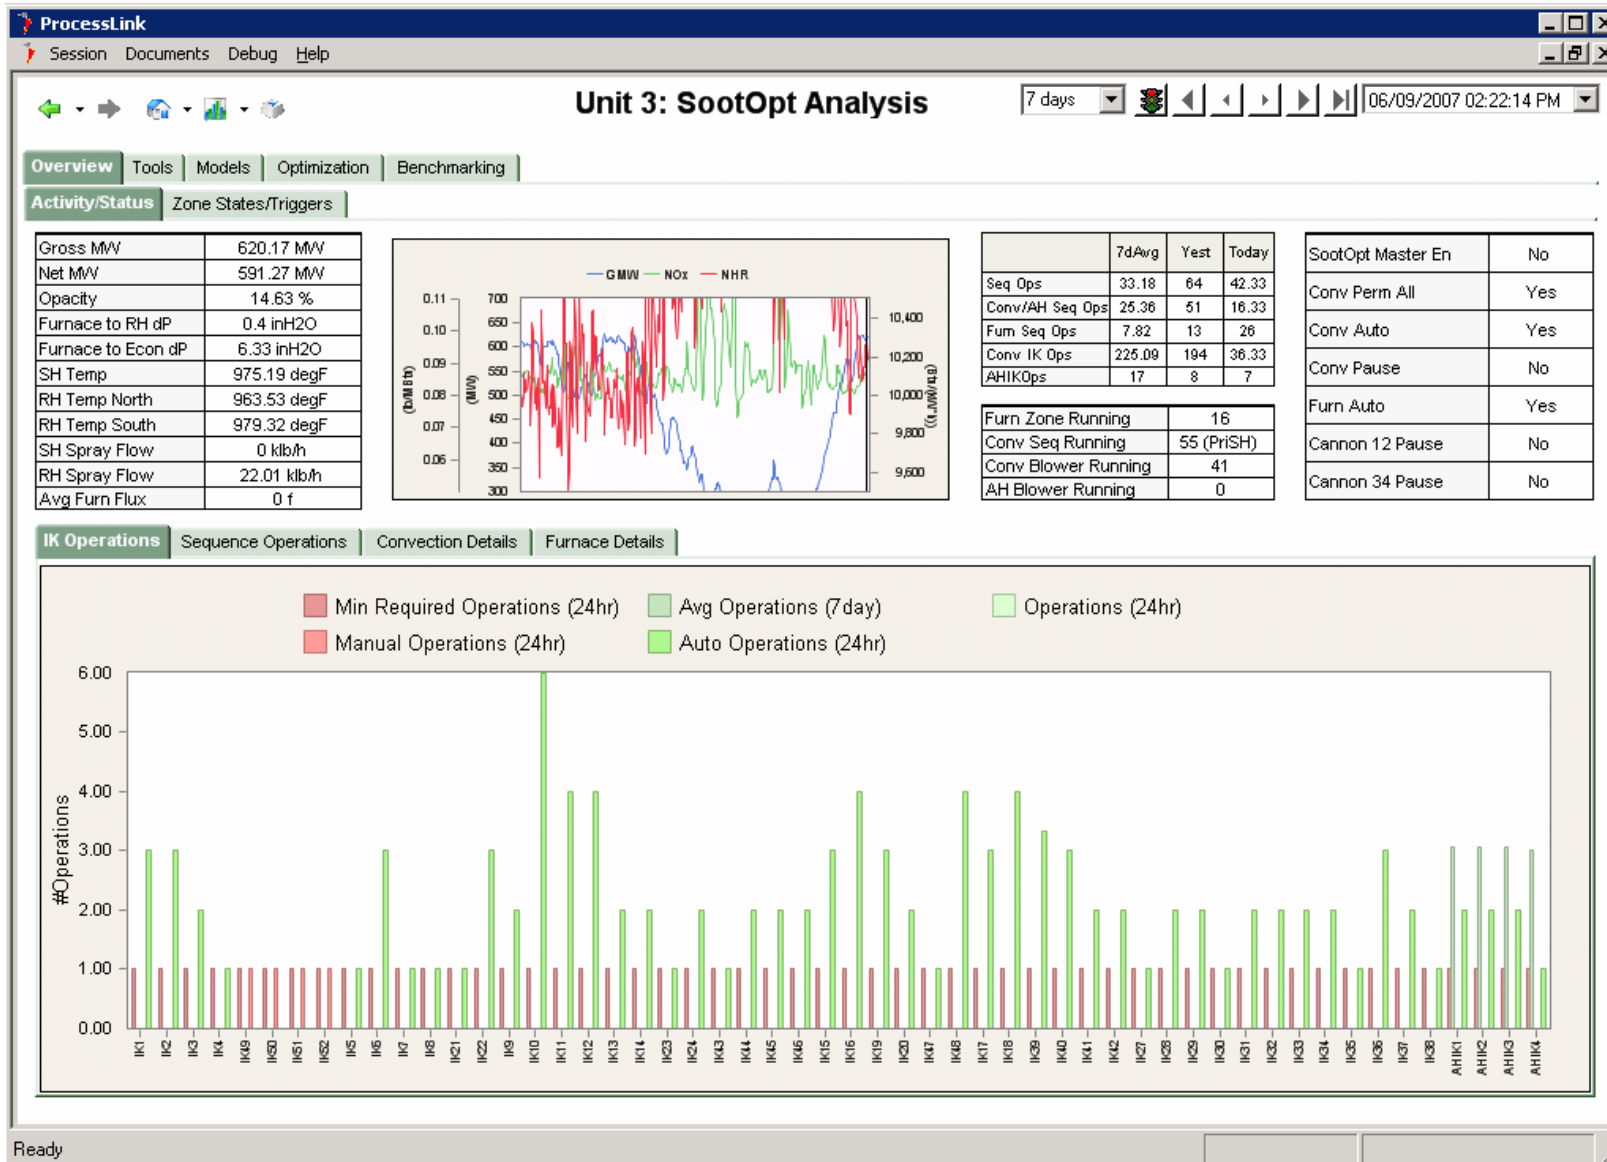


Figure 18 SootOpt Analysis pages

3.4 PerformanceOpt

3.4.1 How It Works

3.4.1.1 Performance Problem Identification

PerformanceOpt is a real-time predictive performance management system that identifies efficiency and capacity losses and enables users to take actions to control those losses and reduce operating costs.

PerformanceOpt identifies problems that are causing performance deficiencies, and determines the efficiency and capacity impacts of each problem. Based on a detailed first-principles model of the unit, PerformanceOpt conducts a full mass and energy balance, calculating on a minute-by-minute basis the results for thousands of variables including process flow rates and conditions, heat transfer rates and sub-system and unit performance metrics. In addition, PerformanceOpt ensures model accuracy and reliability by making use of sophisticated sensor validation mechanisms as well as equipment out-of-service logic.

PerformanceOpt continuously monitors key equipment- and unit-level performance factors and detects, in real-time, when actual performance deviates from what is achievable under current operating conditions. The achievable values are predicted through what-if scenarios that are run with the full-scale model of the unit. For each of the problems that it identifies, PerformanceOpt uses its predictive simulations to determine the potential improvement in efficiency and capacity that would result from resolving that problem. The problem identification workflow is shown in Figure 19 below.

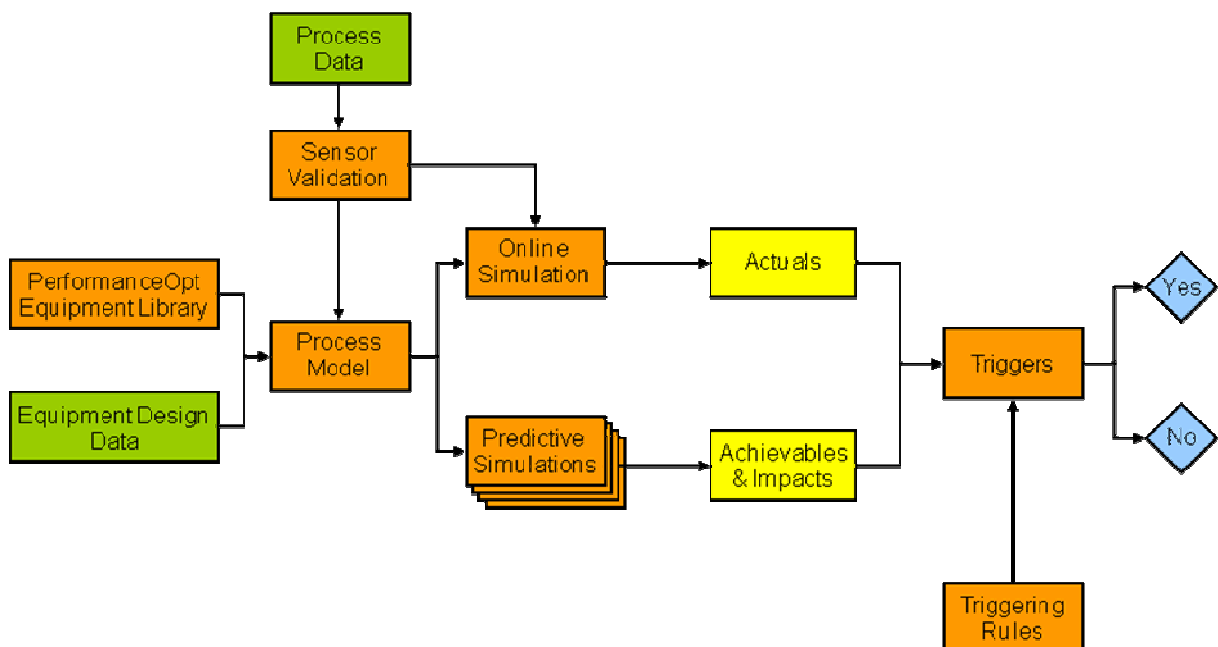


Figure 19 PerformanceOpt - Components in Problem Identification

The first-principles PerformanceOpt model of the integrated combined cycle process is typically comprised of several interconnected flow-sheets that represent the totality of plant equipment, their interconnecting streams, instrumentation, source streams and products. This is the rigorous common model used for both monitoring and predicting performance.

This model runs concurrently in real-time in monitoring mode as well as in predictive mode. In monitoring mode it is configured to calculate the actual unit performance, and in predictive mode it is configured to rigorously calculate the achievable unit performance as well as the efficiency and capacity impacts associated with the deviations between actual and achievable unit performance.

Model scenarios are individual what-if simulations that are configured to correspond to each condition variable of interest (such as Feedwater Heater TTD) in monitoring and improving the unit performance. All of these simulations are run simultaneously and take as input the current plant conditions from the model running in online monitoring mode. They then predict, in real time, the achievable value of each condition variable at current plant conditions and also the potential improvement in heat rate and capacity corresponding to each of the individual performance issues being addressed.

Differences between actual process values (from the monitoring predictions) and achievable process values (from the predictive simulation) above a pre-configured threshold represent the process deviations that allow PerformanceOpt to know that there is a performance gap. Deviations and resulting process impacts form the input for the triggers used by PerformanceOpt to identify if there is a process performance problem that needs to be "alarmed" and brought to the attention of the user.

For problem identification, PerformanceOpt continuously evaluates hundreds of triggers that represent individual equipment performance measures as well as overall process conditions. A PerformanceOpt deployment will typically include numerous instances of a trigger e.g. a FWH TTD trigger for each feedwater heater in the low pressure and high pressure trains. In addition, PerformanceOpt can be configured to include Triggering Rules that serve as pre-requisite or enabling conditions that need to be satisfied for the deviation between achievable and actual values in a performance trigger to be identified as a problem. They help to suppress false positives in problem detection and may include conditions such as unit at full load, process in steady state, etc. as well as combinations of these conditions.

Following problem identification and prioritization based on its calculated impacts, PerformanceOpt facilitates the analysis needed to determine the root cause and identify remedial action by providing the user with intelligently grouped detailed information on measured as well as model-calculated process conditions and equipment performance. The user reviews this information as well as other data that he may gather from the field and diagnoses the problem and takes corrective action. At this point the problem has moved from detection to resolution and PerformanceOpt resumes monitoring the associated performance indicator.

3.4.1.2 Data Validation

Prior to any calculations, all data received from the PerformanceOpt data acquisition system is processed through a sophisticated set of data validation and data substitution algorithms. This helps ensure the integrity of the data being fed into the PerformanceOpt model simulations coming from a source such as the plant data historian which are likely to contain measurements from faulty sensors.

Processing rules for data validation and substitution can be configured in PerformanceOpt at the level of the individual element of process data. These rules can

consist of functions such as clipping, filtering, moving averaging, curve lookup, etc. These functions can be used individually or combined to be used in a sequence on a particular process measurement. Limits for range checks may be defined as static high/low limits or a function of the independent variable. Any PerformanceOpt variable can be used as an independent parameter in the data validation rule for a particular process input.

When a measured input fails the reasonableness check, the input value used in the calculations is replaced with an acceptable replacement value based on the configured rules. The raw input from the data acquisition system is retained so that the user may see the value actually retrieved from the external source. The processed input is assigned a "good", "bad" or "suspect" status and flagged for displaying accordingly in the user interface.

The sensor validation functionality adds reliability to the calculation modules and prevents system shutdowns due to temporary instrument problems. The information on potentially faulty instruments gathered during data validation is available to support instrument maintenance prioritization and scheduling.

3.4.1.3 Equipment Out of Service

In addition to the data validation and substitution mechanisms, PerformanceOpt has the capability to build in logic regarding equipment that may be out of service during plant operation. The PerformanceOpt model calculations know the running status of each piece of equipment. For some equipment, the data source may carry a digital point indicating if the equipment is in service. When this is not available, equipment out-of-service logic is used to set this point based on a data model or rule that looks at the signature of a number of other points from the data source to automatically determine the equipment status.

The ability for the PerformanceOpt models to dynamically detect and compensate for equipment out-of-service helps ensure the accuracy of the simulation results and convergence reliability of the model without causing unnecessary alarms when a piece of equipment is off.

3.4.1.4 Graphical Flowsheeting

PerformanceOpt provides a graphical, easy-to-use flowsheeting environment for constructing, inspecting and maintaining the integrated process model for the complete combined cycle plant configuration. The comprehensive process model is built and configured using drag-and-drop of building blocks from the available library of equipment, instrument and stream models. The PerformanceOpt flowsheeting environment supports the creation of hierarchical, modular, inter-connected sub-flowsheets that facilitate accurate review, better troubleshooting and, most importantly, easier maintainability. Figure 20 shows a typical set of PerformanceOpt flowsheets that make up the integrated process model for the plant.

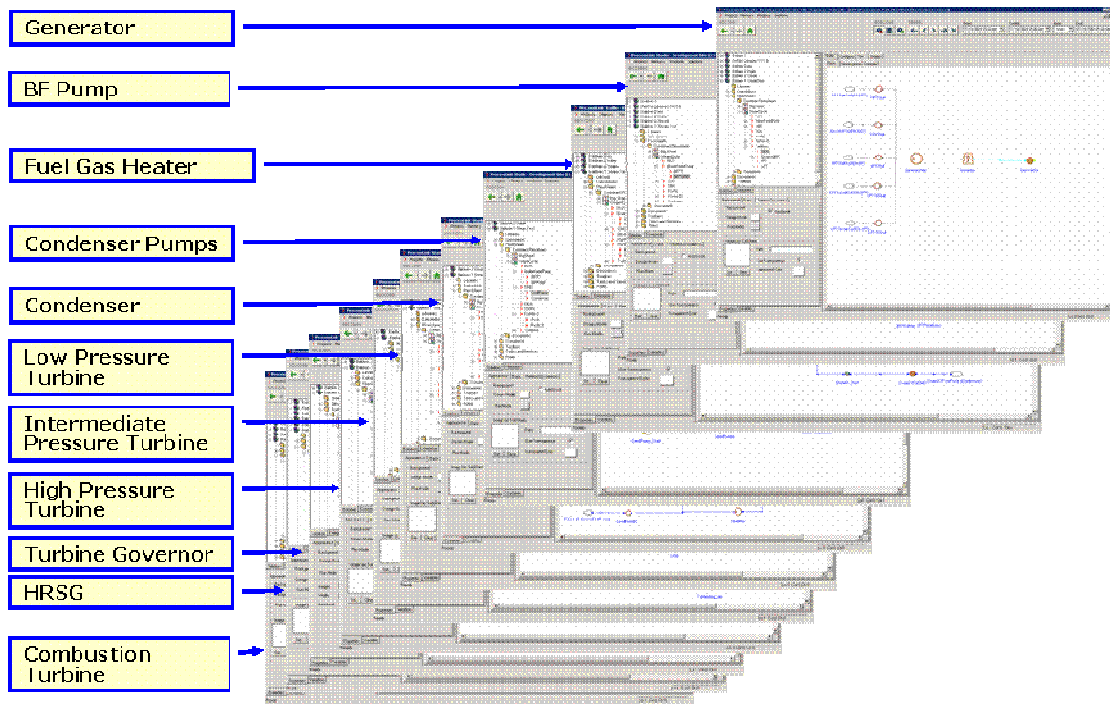


Figure 20 Typical set of modular flowsheets representing a PerformanceOpt model

3.4.1.5 Engineering Library

The PerformanceOpt Engineering Library consists of heat and mass balance models of individual equipment and sub-systems that make up the integrated combined cycle process in a power generation unit. The library also includes various stream types that connect the equipment blocks in a flowsheet representation of the process. The currently supported equipment and stream types are listed below.

Equipment

- Combustor w/ Performance
- Combustion Turbine Compressor and Overall CTG Efficiency
- Feedwater heater w/ performance (when applicable)
- Stream drum (HRSG)
- Convective zone slice (for HRSG)
- Condenser w/ performance
- Steam turbine generator w/ performance
- Mixer/Splitter
- Pump w/ performance
- STG Governing stage
- Pressure drop
- Deaerator w/ performance
- Electrical generator
- Cooling towers w/ performance
- Spray controller
- Steam seal receiver
- Fuel gas heater

Streams, Pressures, Temperatures

- Combustion Turbine Exhaust Temperature
- Reheat temperature
- Gas stream
- Feedwater temperature
- HRSG gas outlet temperature
- HRSG Excess Spray Flow
- Reheat spray flow
- Superheat spray flow
- Water stream
- Shaft stream
- Signal stream
- Wire stream
- Radiation stream (n/a for HRSG)
- Steam Turbine Throttle temperature
- Main Steam temperature
- Throttle pressure
- Air heater cold end average or air inlet temperature (n/a for HRSG)
- Condenser Backpressure
- Condenser make-up flow
- Auxiliary power
- Condenser pressure

In addition, PerformanceOpt also includes the following library of engineering and physical property functions for use in the heat and mass balance calculations:

- ASME 1967 and 1997 Steam Tables
- Psychrometric functions
- HEI 8th edition
- NIST gas property tables

3.4.1.6 Equipment Performance Calculations

In addition to the post-processing performance calculation modules that are included in PerformanceOpt for Boiler Efficiency, Boiler Cleanliness, ASME Turbine Performance, and Heat Rate, there are also equipment-level performance results that are generated during the model simulation and made available to the user.

Boiler Performance

Boiler efficiency is calculated using the heat loss method and is based upon the current draft of ASME PTC 4. Efficiency is also calculated using the input-output method for comparison purposes. Individual losses and credits around the boiler envelope are calculated and displayed along with boiler efficiency. A mass and heat balance for the steam/water process streams and air/gas streams is determined using available measured flow rates, pressures, temperatures, and enthalpies and displayed on a graphic representation of the steam generator.

Steam Turbine Generator

Performance parameters are calculated for the HP, IP, and LP turbine sections. Turbine section efficiencies and turbine heat rate (actual and corrected) are calculated based upon the requirements of ASME PTC 6, Steam Turbines. Corrected turbine performance is based on Group 1 and Group 2 procedures. The level of calculations performed is dependent upon the available instrumentation and are customized to match performance requirements and instrumentation resources.

Deaerator and Other heat Exchangers

PerformanceOpt calculates performance parameters for any deaerator configuration and any other closed or open heat exchanger. Actual heater transfer efficiency and/or effectiveness is calculated and compared with the design values (specified by the manufacturer) which are corrected to the actual operating conditions. Heat exchanger effectiveness, pressure drop, and the capability to simulate plugged tubes are also included. These calculations are based on the requirements of ASME PTC 12.1, Closed Feedwater Heaters, and the HEI Standards for Closed Feedwater Heaters.

Condenser

PerformanceOpt calculates performance parameters for the surface condenser and simulates the impacts associated with degraded conditions. Condenser cleanliness, expected backpressure (at specified cleanliness), and condenser heat load (steam side and/or circulating water side) are calculated based upon the requirements of ASME PTC 12.2, Steam Condensing Apparatus. The capability to calculate condenser performance to compensate for the number of tubes plugged is included. With proper instrumentation, air in-leakage and backpressure effects may also be calculated.

Cooling Tower

Cooling tower cold water temperature, range approach and cooling tower capability are calculated using manufacturers design curves. The calculations are based on the 'Performance Curve Method' specified in the Cooling Tower Institute Test Code, CTI ATC-105.

Pumps

Performance parameters are calculated for the large pump systems. Analysis is performed in two areas for the system: individual pump performance and operating configuration (if applicable). Individual pump performance is determined by calculating actual pump efficiency per ASME PTC 8.2 methods. Expected pump performance is determined from the manufacturer's pump curves (or performance test data if available). Actual performance is compared to expected performance using pump affinity laws to determine pump performance degradation. The method of determining turbine-driven pump performance will depend upon the instrumentation available.

3.4.2 The User Interface

Like the other Optimizers, PerformanceOpt has a Home page (see Figure 21) that provides three kinds of information that enable users to derive the maximum value from their system:

Optimization Advice

The upper left section shows PerformanceOpt's current advice for how to further optimize the unit based on the differences between the actual performance and predicted achievable performance by the PerformanceOpt model. Examples of problems that PerformanceOpt identifies are low feedwater heater cleanliness, high condenser backpressure, etc. In addition, to assist with prioritization and trade-offs, this section shows the impacts, calculated by what-if simulations of the full PerformanceOpt model, on heat rate and capacity that are likely to occur by remedying the identified performance gap. Further, to facilitate quicker analysis of the problem, PerformanceOpt provides problem contextual navigation from this section to equipment views with collections of drill-down information.

What's Going on Now and Why

The lower left section provides access to a variety of information that describes how the unit is currently operating and also explains why PerformanceOpt is giving the advice it is giving. The Controllable Losses table maps the specific equipment problems that PerformanceOpt has identified to the more general controllable loss categories that plant personnel are familiar with. The unit process schematic in this section provides access to more detailed information about how each piece of equipment on the unit is performing.

Optimization Benchmarks

The section on the right shows how the unit has been performing over the recent past period, typically a month, with respect to its heat rate and capacity objectives; it compares actual heat rate and capacity factor over that period to a baseline and an achievable performance standard, determined from what-if simulations of the rigorous PerformanceOpt model.

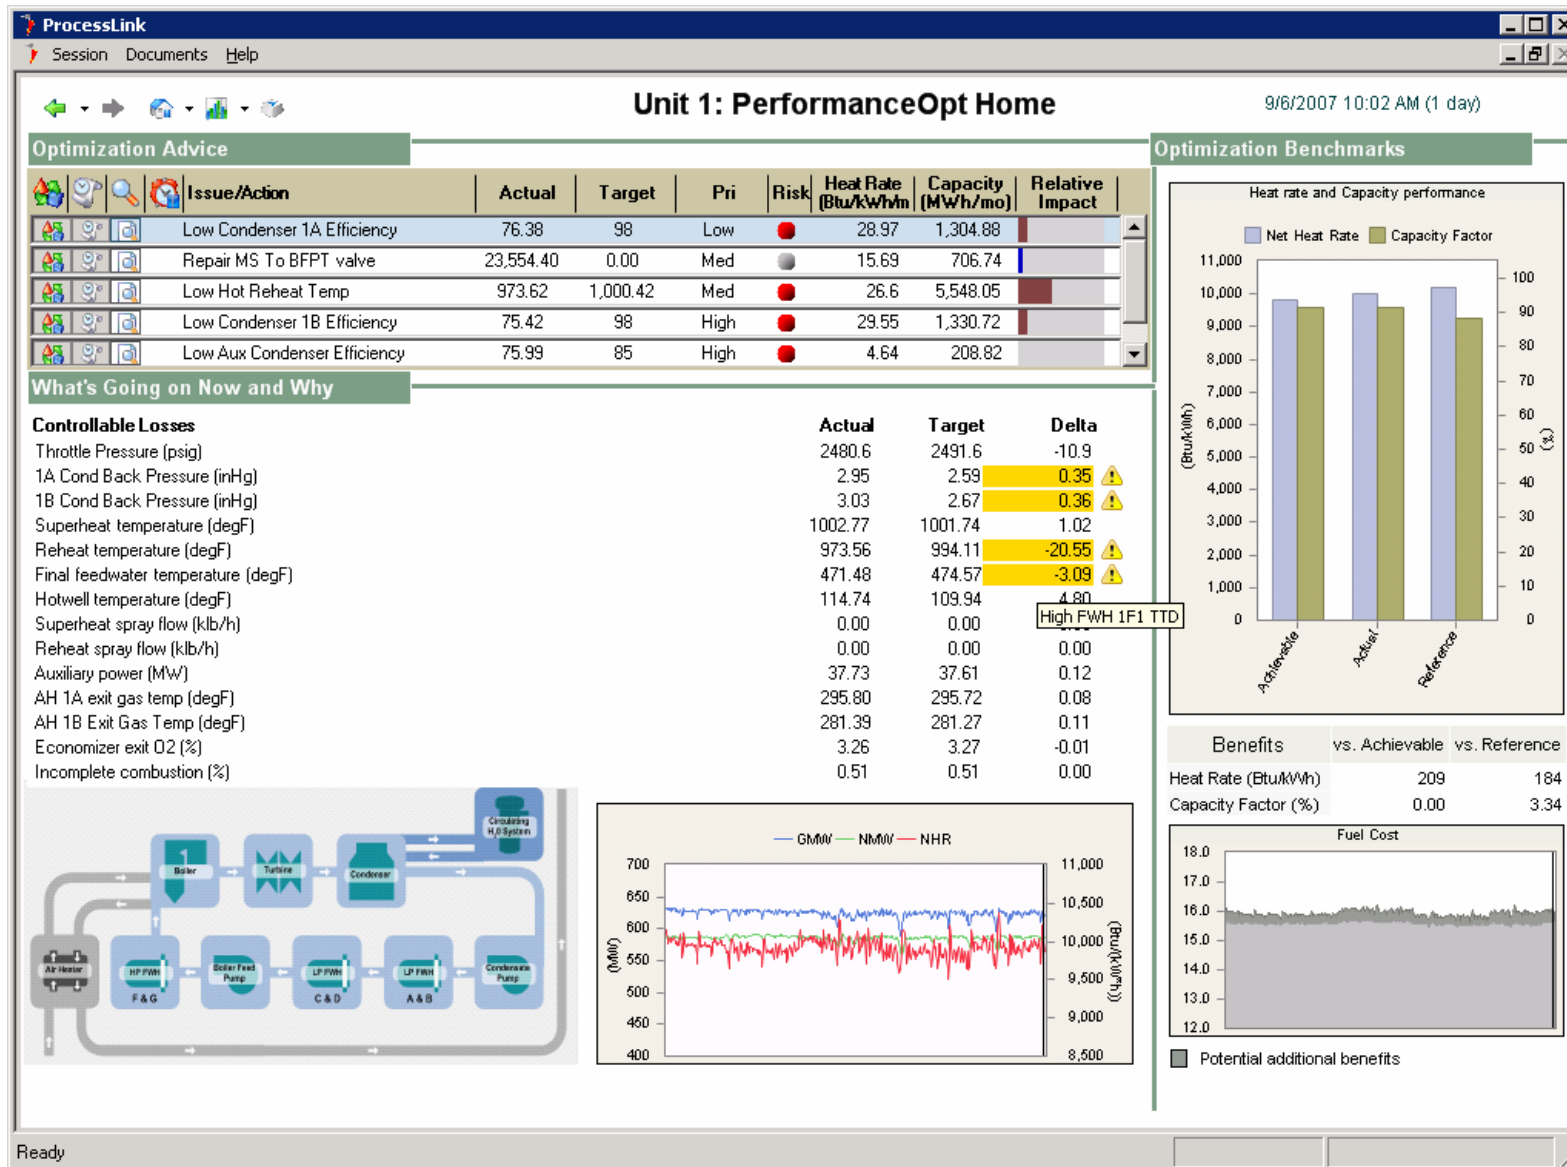


Figure 21 PerformanceOpt Home page

3.5 MaintenanceOpt

3.5.1 How It Works

MaintenanceOpt continuously monitors process and equipment health data looking for anomalies that might indicate the presence of reliability, capacity, or efficiency problems. When anomalies are detected by MaintenanceOpt, the system's heuristics knowledgebase supports the identification of the most likely causes of the anomalies. It also performs an estimate of the reliability, efficiency and capacity impacts to help prioritize the order in which problems should be addressed. The system records the details and events of the workflow steps so that the knowledge is available for future review and re-use by other personnel and also for historical reporting. MaintenanceOpt's real-time decision support interface presents the maintenance problems, their diagnoses, required actions, and impacts and risks.

MaintenanceOpt detects reliability, capacity and efficiency-related anomalies and helps engineers manage the entire lifecycle of a detected problem more efficiently and effectively. After anomalies are detected, it can typically take a lot of time for plant staff to investigate potential problems to see if they are real and to determine how important they are. MaintenanceOpt saves time by displaying all in one place the information required to determine whether the detected anomaly points to a real problem or whether it is a result of bad sensors or the anomaly engine being confused. If the user decides the problem is real, he escalates the problem for diagnosis. The diagnostics knowledgebase identifies possible causes for the problem and draws attention to the data that indicates which cause is most likely. Based on the impacts MaintenanceOpt projects and on other available information, plant engineers assign a priority to the problem and put it on the action list. The workflow supported by MaintenanceOpt is shown in Figure 22.

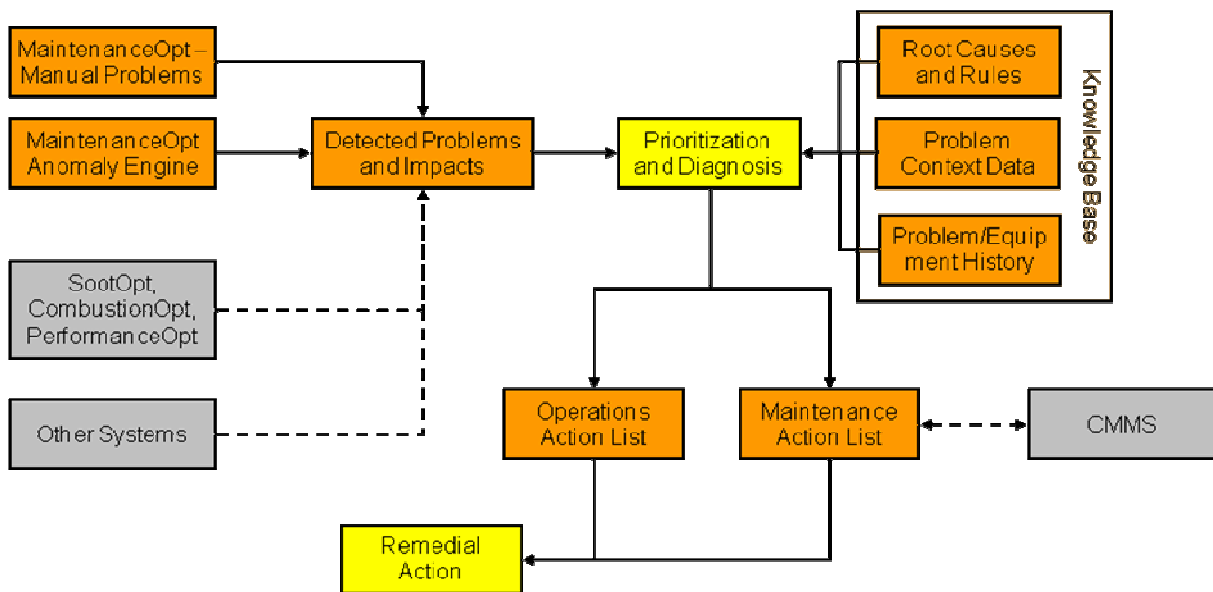


Figure 22 MaintenanceOpt workflow for problem detection, diagnosis and resolution

MaintenanceOpt uses its anomaly detection mechanisms to identify problems that impact reliability, capacity, and efficiency. It can detect both slowly developing problems that have an increasingly negative impact on capacity and efficiency as well

as problems that could have a critical near-term reliability impact. MaintenanceOpt then uses its diagnostics knowledge base containing hundreds of heuristic rules to help users determine the cause of these problems. The anomaly engine and the heuristic rules are incorporated into a robust software infrastructure so that they can be applied plant-wide in real-time.

In addition to displaying the impacts of problems it itself detects, MaintenanceOpt can display the impacts of problems found by other Optimizers such as PerformanceOpt, CombustionOpt and SootOpt using their first principles-, neural network-, and rules-based modeling capabilities. Thus, MaintenanceOpt serves as clearing house for all problems that are impacting plant performance where the appropriate plant personnel can do the relative prioritization, perform diagnosis and identify the appropriate corrective action, and track through scheduling and execution of the maintenance task. Figure 23 shows the Problems to Diagnose view in MaintenanceOpt that typically is the starting point for a user participating in this workflow.

Problems are diagnosed by reviewing and differentiating between the potential causes that are available in the MaintenanceOpt knowledge base for each identified problem symptom. The user does the diagnosis by using the list of likely causes and their associated heuristics, measured and modeled process data that are contextual to the problem, as well as additional data that may need to be obtained through local measurements.

The diagnostic heuristic associated with the list of likely causes guides the user through the root cause analysis by allowing him to rule out causes and/or differentiate between causes. The diagnostic rule is used to not only bring specific elements of the context data to the user's attention but also to highlight additional data required for the analysis. Through this process the user is able to identify the most likely cause and then specify the appropriate remedial action as shown in Figure 24 and Figure 25 below.

Once the user has diagnosed the problem and specified the corrective action, it is put into the Operations Tasks list or the Maintenance Tasks list based on the nature of the remedial action required. Maintenance tasks are further sub-categorized into activities that require no derate, require a derate, or require an outage. As an example, Figure 26 shows a view of diagnosed problems along with their corrective actions that have been identified as requiring a de-rate; the context data used in the diagnosis is still available to user in this view for re-evaluation, if needed.

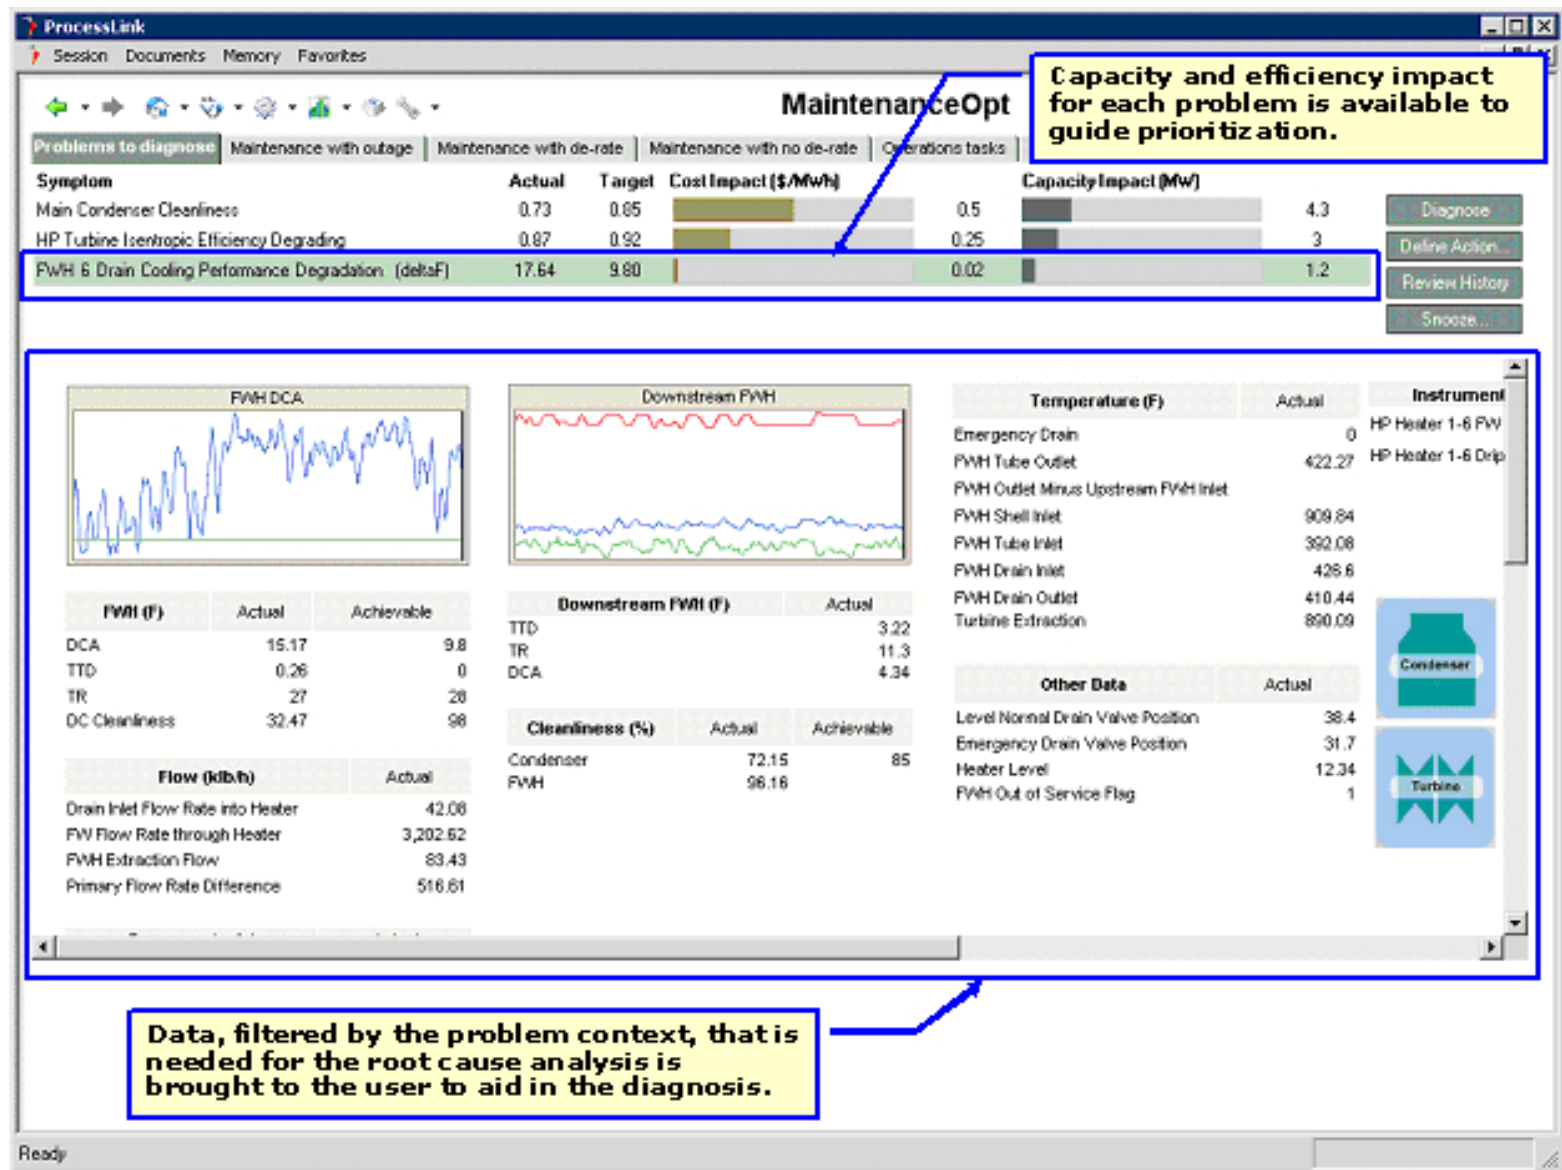


Figure 23 Problems to Diagnose view in MaintenanceOpt

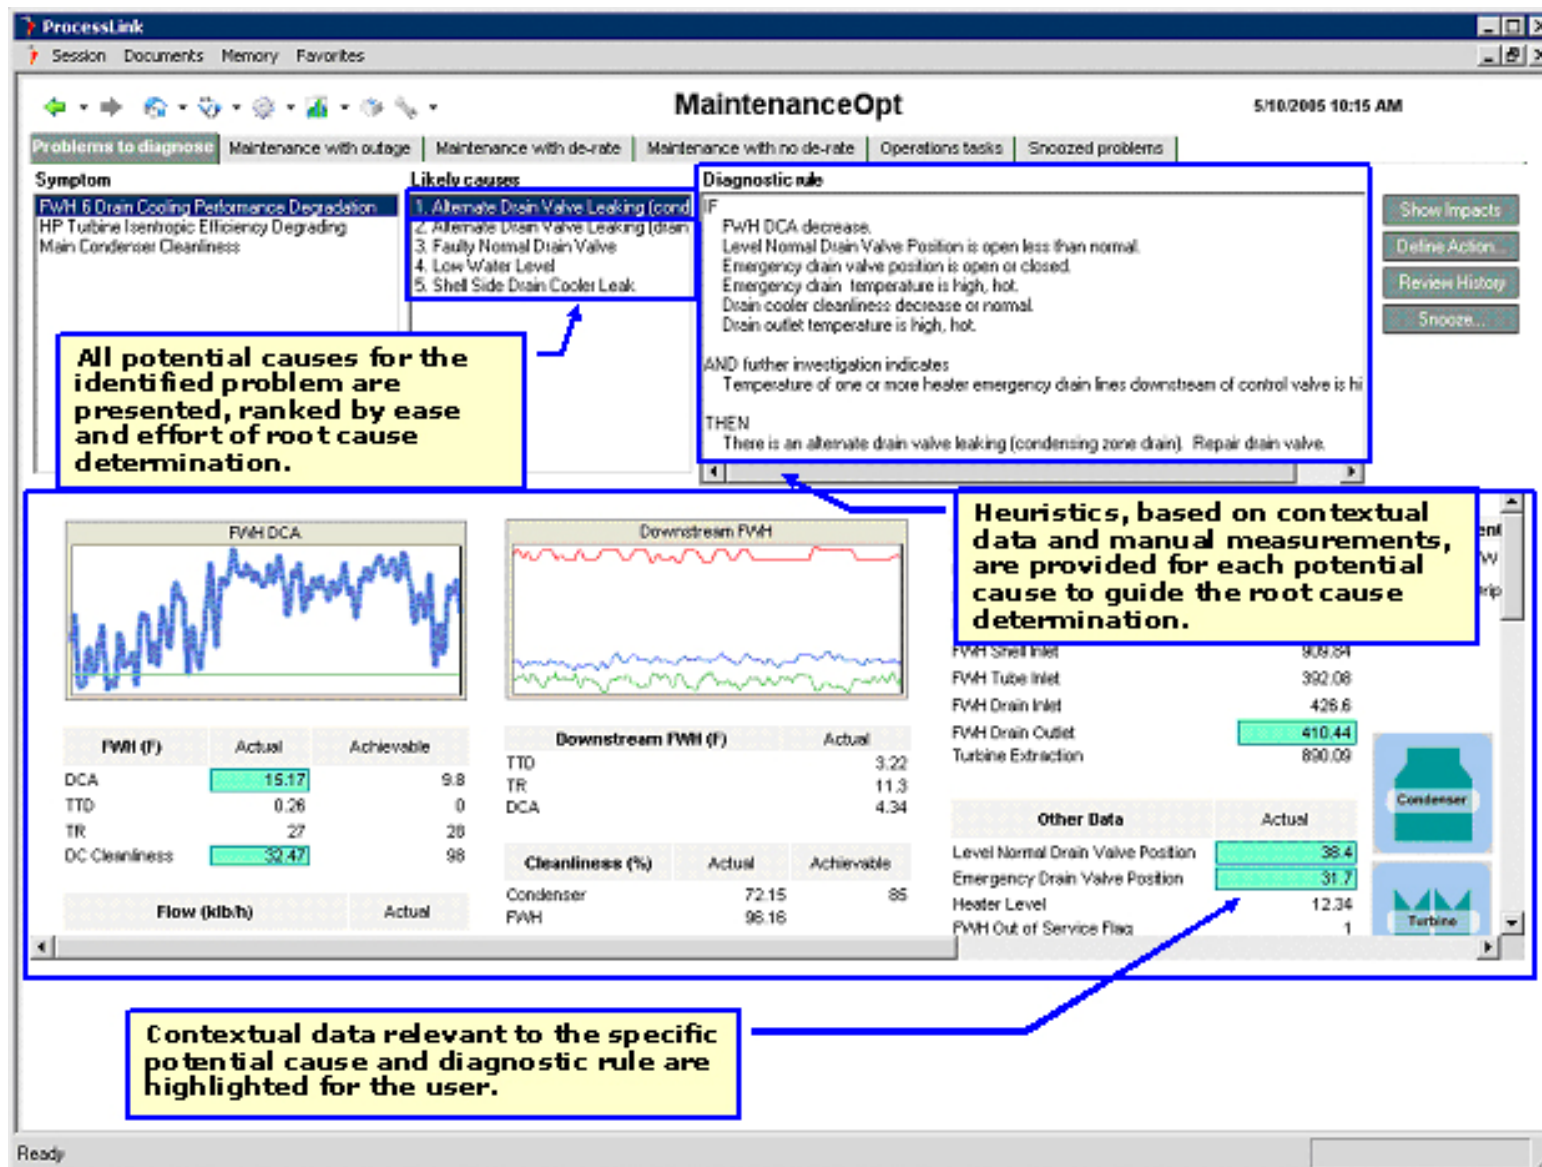


Figure 24 Potential causes and heuristics for Feedwater heater DCA performance symptom

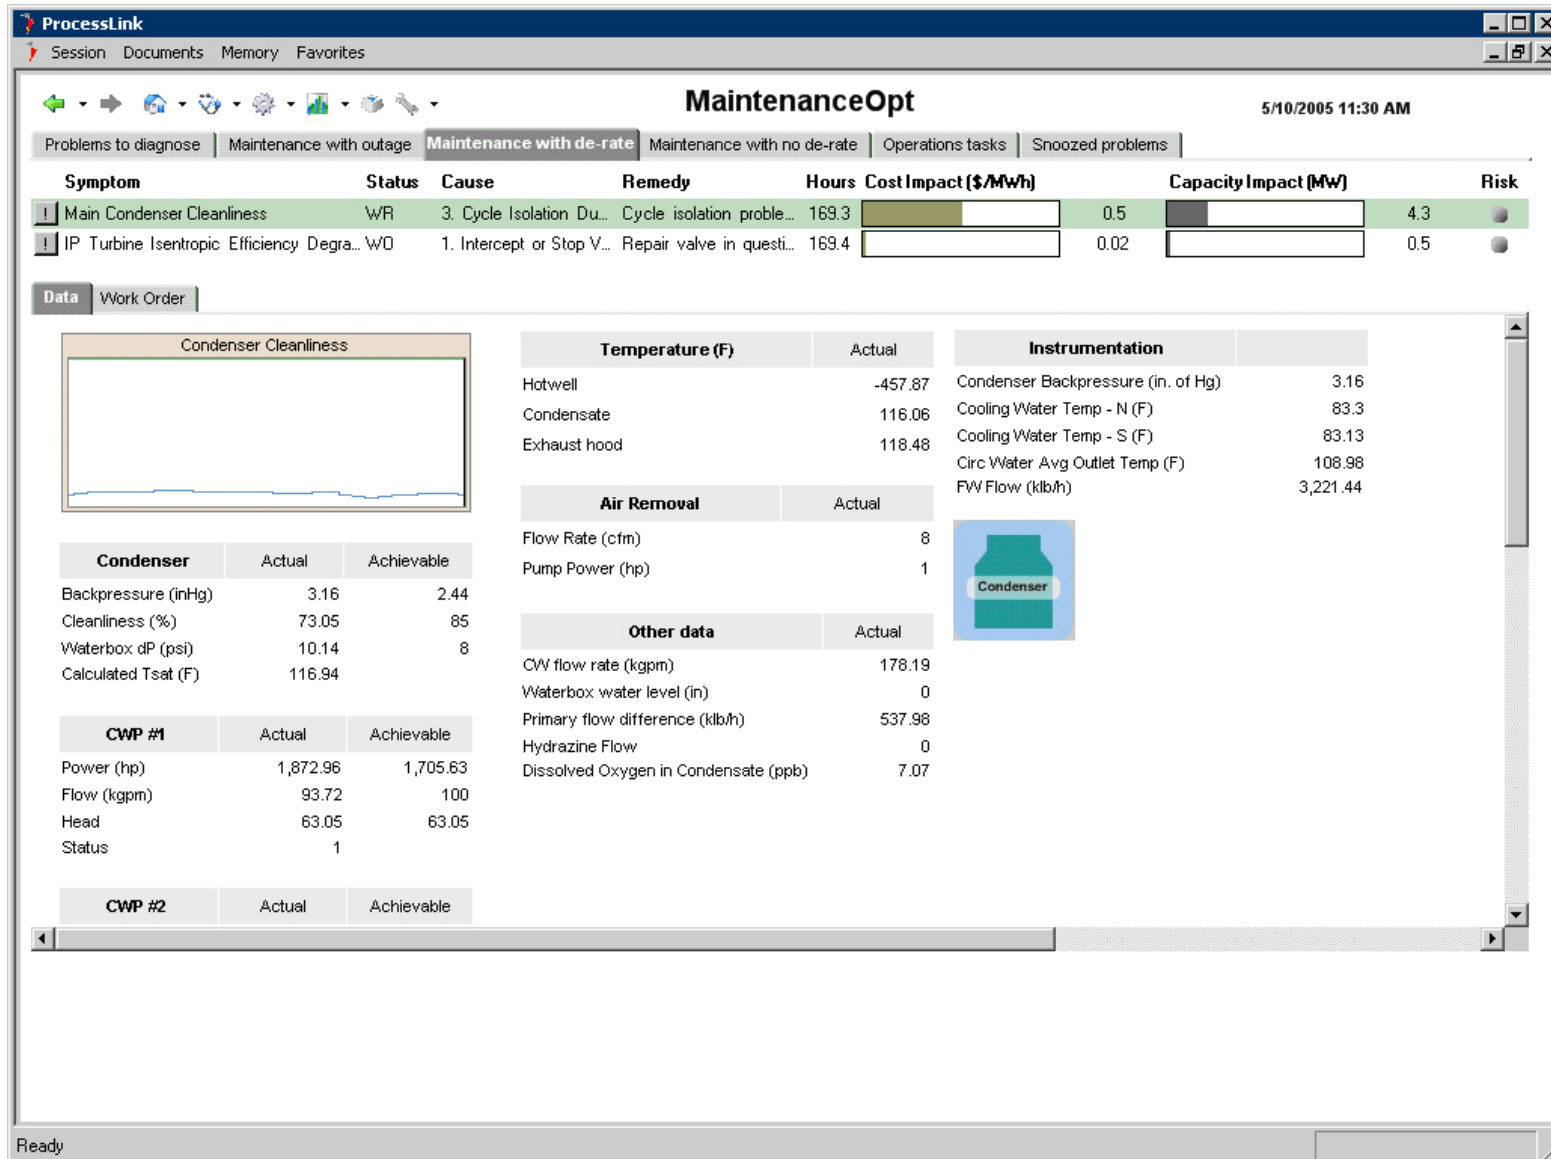


Figure 26 View of diagnosed problems with corrective maintenance actions that require a derate.

3.5.2 The User Interface

Like the other Optimizers, MaintenanceOpt has a Home page (see Figure 27) that provides three kinds of information that enable users to derive the maximum value from their system:

Optimization Advice

This section at the top left provides a summary of all of the issues that are currently being managed in MaintenanceOpt. It provides an overview of the reliability risks and impacts associated with the problems currently in the various states of the problem lifecycle. It also provides drill-down navigation to the MaintenanceOpt Action Lists where the user can access more problem-level detail.

What's Going on Now and Why

This section at the bottom left provides a summary view into the current set of problems being managed in MaintenanceOpt based on affected equipment and priority. In addition, the user is also presented with a consolidated list of instrumentation-related problems.

Optimization Benchmarks

This section on the right shows how the unit has been doing over the past month with respect to optimizing costs and revenue; it compares cost performance, based on fuel, ammonia, NO_x, etc., and revenue performance, based on MWs, over that period to a set of baseline and target performance standards for the unit. In addition this section also benchmarks the efficiency of problem lifecycle management over that period based on the average time that problems remained in the various states (not yet screened, undiagnosed, etc).

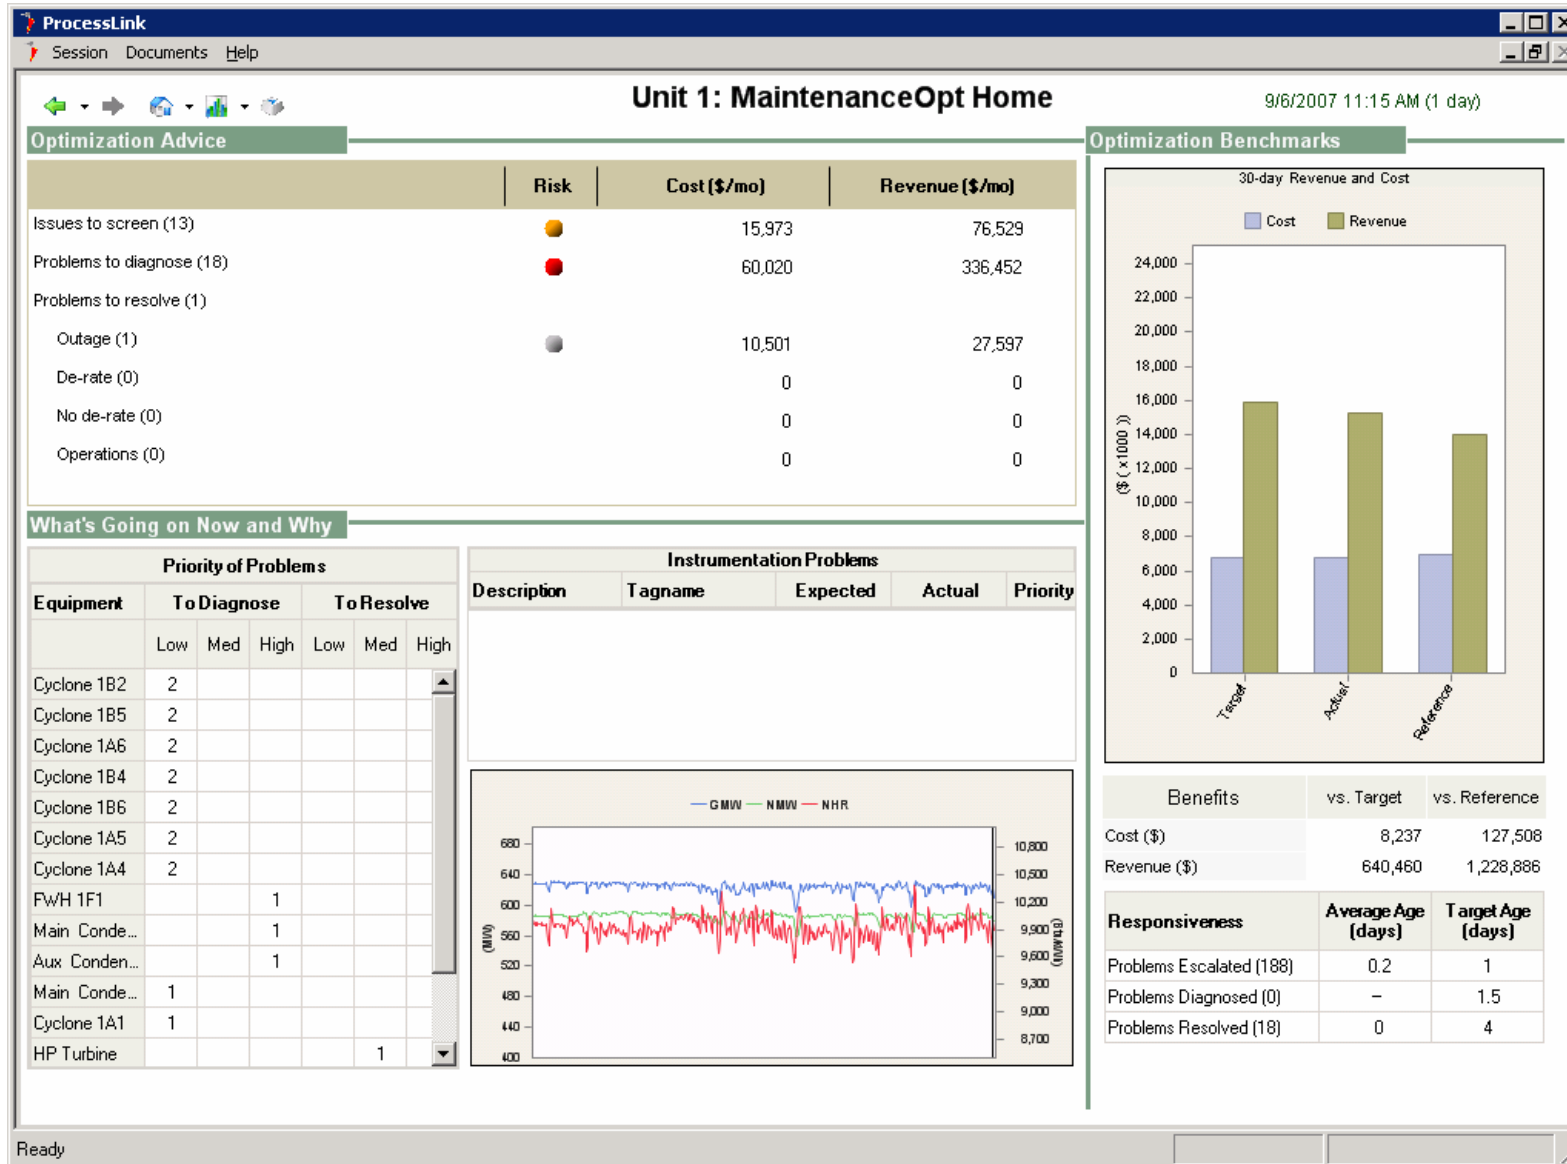


Figure 27 MaintenanceOpt Home page

4 BEC Data Analysis

4.1 Overview

For CombustionOpt, SCR-Opt, and SootOpt, the primary optimization actions are taken in closed-loop, i.e. they directly manipulate or bias equipment. CombustionOpt and SootOpt also have advisory optimization components, and PerformanceOpt and MaintenanceOpt are almost exclusively advisory optimization systems. These advisory systems provide real-time alerts of issues that require a person's attention, along with contextual information, priority assignments, impact estimates and, for alerts where expert knowledge about how to handle them exists, diagnostic and resolution support. They also organize information about problem history by problem type and equipment category, thus helping to capture that knowledge, which is ever more critical with the industry's challenges finding and retaining qualified operations and engineering personnel.

Online closed-loop systems lend themselves well to rigorous data analysis. How to estimate the value of advisory systems is much less straightforward. For this reason, following some background information, this section is divided into two main parts: the first focuses on detailed analyses of the benefits delivered by closed-loop optimization, and the second focuses on examples of valuable alerts and advice provided through advisory optimization.

For clarification purposes, because this project involved so many manipulated and controlled variables (literally hundreds) across all the products installed, and because some products function in closed-loop utilizing high-frequency optimization search algorithms (both model-based and direct) that include the design of experiments, the test plan essentially was to apply optimization across all the manipulated and control variables, then examine the results in key objective dimensions using data collected and the following two analytical methods:

- 1) Comparison of large populations a' posteriori - with one population representing the result of the experiment and one the control (please refer Section 4.4.1 Comparison of Populations Methodology), and
- 2) regression modeling (using neural nets) to correct the experiment and control populations selected for comparison, by removing a known set of strong factors that were not the experimental factor (please refer Section 4.4.2 Modeled Response to Disturbances).

4.2 BEC Background

4.2.1 Historical Events

To add some background and context, following are several important emissions-related plant events that occurred prior to the optimization project:

2002: Switch to 100% low Sulfur PRB coal

2002: EPA Consent Decree established for BEC

2000: Low NOx Burners installed on U3

SCR's on U1&2

- U1 - first three layers installed in the spring of 2003
- U1 - 4th layer added in the spring of 2005
- U1 - has room to install a 5th layer in the future

- U2 - first two layers installed in the spring of 2002
- U2 - 3rd and 4th layers installed in the spring 2004
- U2 - has room to install a 5th layer in the future

Continuously changing fuel purchase/supply.

4.2.2 Fuel Heating Value Trend

Figure 28 shows the record of fuel heating value from the BEC delivery assays for the period between 2004 and 2007. Like many power plants, BEC purchases fuel from multiple distant sources, and often sees variability between shipments from a given source. The trend shows that over the period between 2004 and 2007 the heating value of this coal on average was declining.

It is reasonable to expect that if anything this would put negative pressure on overall plant performance. When taking into account the fact that these units (especially the cyclones) were designed to run high sulfur, high Btu content Illinois coal, with much different mineral content and characteristics than PRB coal, this is especially relevant. In summary, over the period between 2004 and 2007 the site was receiving and burning worse and worse quality coal with respect to expected heat rate performance and cyclone stability. As we look at the plots to follow, and keep in mind the notions of tradeoff relationships and interacting objectives, this downward trend, when **not** matched by downward trend in heat rate and cyclone stability (represented in later figures by Flame Scanner Quality), represents delivered equivalent benefit. This is to say that one benefit delivered, in fact one of the most significant achievements of the project, is that these units were able to burn this low sulfur (but off-design) fuel without de-rating, and while maintaining and even improving their performance.




Figure 28 BEC Coal Analysis Btu/lb (Jan 2004- Dec 2007). The downward trend in fuel energy density could be expected to affect cyclone stability and unit Heat Rate adversely.

4.2.3 Unit Load Trends

Figure 29 - Figure 31 show the unit load trends from roughly 2004 to present.

Observations:

- All three BEC Units are base-loaded, with high capacity factors
- The frequency of outages, planned and not planned, appears lower as time goes on for Units 1 and Unit 2 (some of this is likely due to progress made with the combustion in the cyclones).
- Outage frequency for Unit 3 is lower earlier, then increasing. To some degree this might be expected (high reliability leading to lower reliability through increase wear and tear).

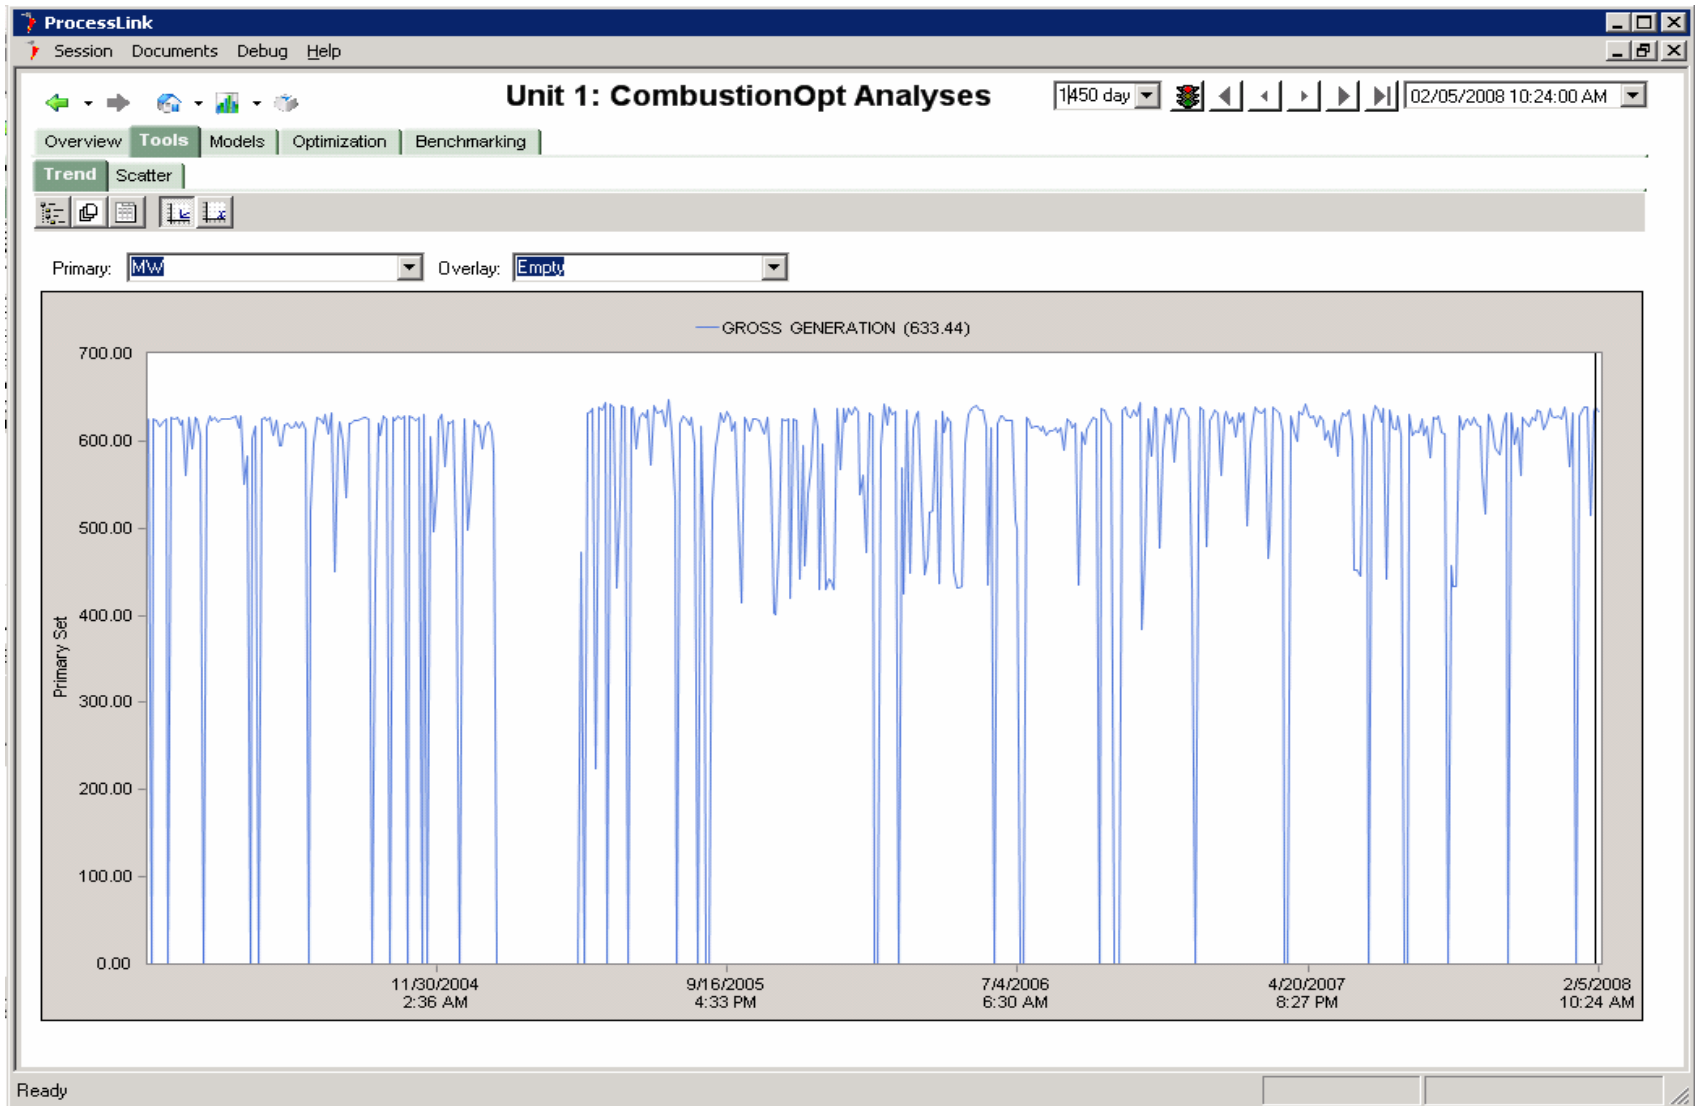


Figure 29 Unit 1 MW Trend Primary (left) y-axis is in units of Megawatts. Trend is shown to provide general production context for the time span analyzed




Figure 30 Unit 2 MW Trend. Primary (left) y-axis is in units of Megawatts. Trend is shown to provide general production context for the time span analyzed




Figure 31 Unit 3 MW Trend. Primary (left) y-axis is in units of Megawatts. Trend is shown to provide general production context for the time span analyze

4.3 Closed-Loop Optimization Approach

4.3.1 Combustion and SCR Optimization on the Cyclones

The Optimization approach for the cyclone units at Baldwin (Units 1 and 2) progressed in the following way:

1. Direct modeling and optimization of SCR Inlet NO_x as function of Cyclone Feeder and Secondary Air biases (SCR Inlet NO_x instrument proved not to be a very good one for estimating furnace NO_x production, NH₃ flow is a more complete estimator).
2. Direct modeling and optimization of SCR NH₃ Flow as function of Cyclone Feeder and Secondary Air biases as well as other biases.
3. Addition of optimization on Cyclone Stoichiometry (as function of individual Cyclone biases, in concert with direct optimization of NH₃ flow and Heat Rate as functions of Cyclone biases).
4. Addition of expert rules to govern profile switching to respond to degrading cyclone slag flow (to prevent it from becoming irreversible).
5. Addition of Flame Scanner Main Quality signals to models of NH₃ and cyclone stoichiometry
6. Direct Neural optimization of Flame Scanner Main Quality as function of individual Cyclone Secondary Air and Feeder Biases.
7. Two tiered scheme optimizing Cyclone Main Flame Scanner Quality as function of individual cyclone Feeder and Secondary Air biases, with flame scanner Quality targets determined through modeling and optimization of NH₃ and Heat Rate as functions of flame quality and MW

This process was driven by an evolving understanding of the following challenges:

- Running cyclones designed for high sulfur coal on PRB, which means pushing the feeder and cyclone mass flows to their limits, under low stoichiometric conditions.
- Optimizing operations in the context of an SCR post combustion NO_x controller.

The specific optimization goal used to address these challenges was resolved to be keeping the flame quality corresponding to unmeasured temperature and the resulting slag consistency high, while reducing the fuel/air ratio.

The data in this report will show how optimization helped to enable BEC to run 100% PRB coal through a boiler designed for bituminous coal, at low stoichiometries. The end results of this are low NO_x, high efficiency SCR duty, reduced NH₃ consumption, positive impacts on heat rate and avoided de-rates.

4.3.2 Combustion and Sootblowing Optimization Approach on Unit 3

For T-fired Unit 3, which is a common type found in the commercial CombustionOpt treated fleet, and for which a good combustion optimization strategy was already well defined, the evolution of the optimization approach consisted primarily of the integration of CombustionOpt and SootOpt. This integration combined the adaptive models used to direct the Combustion fuel and air biases, and sootblowing heat flux and uRatio targets, into one model, with SootOpt using CombustionOpt's MV's as state inputs and vice versa.

4.3.3 Closed-loop Utilization

During the initial phase of this project over 50 manipulate-able variables were identified and installed in the DCS control systems for both Units 1 and 2, and over 25 for Unit 3. CombustionOpt was first installed on Unit 1 with Unit 2 being installed in parallel with some delay. Most teething problems with understanding the cyclones occurred in 2005 culminating in a change of strategy in December to focus more directly on cyclone stability and to add an expert system component.

Note: For an explanation of what is meant by "Closed Loop" and how CombustionOpt, SCR-Opt and SootOpt function please see section 3.2

Unit 3 CombustionOpt was completed somewhat later mostly due to the need for an outage opportunity suitable for making DCS logic changes to support closed-loop optimization, and the fact the unit had a 370+ day record run.

Unit 3 SootOpt was installed in the spring of 2006. In order for SootOpt to be integrated with the existing ISB by Diamond/ASI, NeuCo needed to work with that company to develop an interface and an optimization strategy. It should be noted that SootOpt has been developed to integrate with an existing ISB like the Diamond/ASI Sentry Series 1500, or be installed on-top of standard PLC or DCS based Sootblowing Control Systems that more typify the fossil fired fleet. The functionality it provides when layered on top of a relatively new ISB is a subset of what it provides when is connected directly to a SBCS, without any intermediate controls. In a context like Baldwin its function is to deliver direction to a locally intelligent system, using flux and cleanliness information, to direct the activity of that system toward global (i.e. higher-level plant) objectives. In a more typical context where it directs local actions (e.g. blower selection), the global optimization direction it provides is to its own propose-apply expert system.

Figure 32 - Figure 34 show the Master Enable for the CombustionOpt system (and SootOpt for U3). These charts show a 1450 day trend of MWs and the CombustionOpt Master Enable, as well as the same trend for the SootOpt Master Enable for Unit 3. For the Master Enable, 1 indicates "ON" and 0 indicates "OFF". The operator controls the Master Enable state of CombustionOpt, SCR-Opt and SootOpt from dedicated screens on the DCS. When master enable is "ON" CombustionOpt, SCR-Opt and SootOpt have permission to directly adjust the fuel and air distribution biases, or in the case of SootOpt direct the automated activity of the sootblowing system. In the case of CombustionOpt the operator can also enable or disable individual fuel and air distribution biases. SootOpt is either fully ON when its Master Enable is True (1, high) or fully OFF when it is False (0, low).

Observations:

- Unit 1 and Unit 2 were installed first, while Unit 3 continued on an extended high load pattern. This period of sustained high load operation delayed the installation of Unit 3 CombustionOpt and SootOpt because a significant outage was preferred for the installation of the logic to allow CombustionOpt to talk to the DCS in closed-loop, as well as the installation of some instrumentation needed for SootOpt. Since most units have scheduled or unscheduled outages on a semi-regular basis, this is not generally an expected delay. Unit 1 and 2 were in closed loop rapidly.
- Utilization in terms of percent of time enabled on all three units has improved over time.

- Unit 3 went through a period of tuning and adjustment in late 06 during which the interaction between CombustionOpt and SootOpt was coordinated through the use of shared model inputs and optimization objectives.
- SootOpt was installed just after CombustionOpt, once the integration work needed to optimize the Diamond ISB and the time needed for the vendor to install some necessary flue-gas temperature monitoring instrumentation.
- Utilization of SootOpt has been up and down in recent months due to issues with the Diamond ISB instrumentation array responsible for providing real-time measurements of density of soot buildup on interior furnace wall surfaces. SootOpt is currently functioning only in the convection region of the boiler while this is being remedied by the plant.

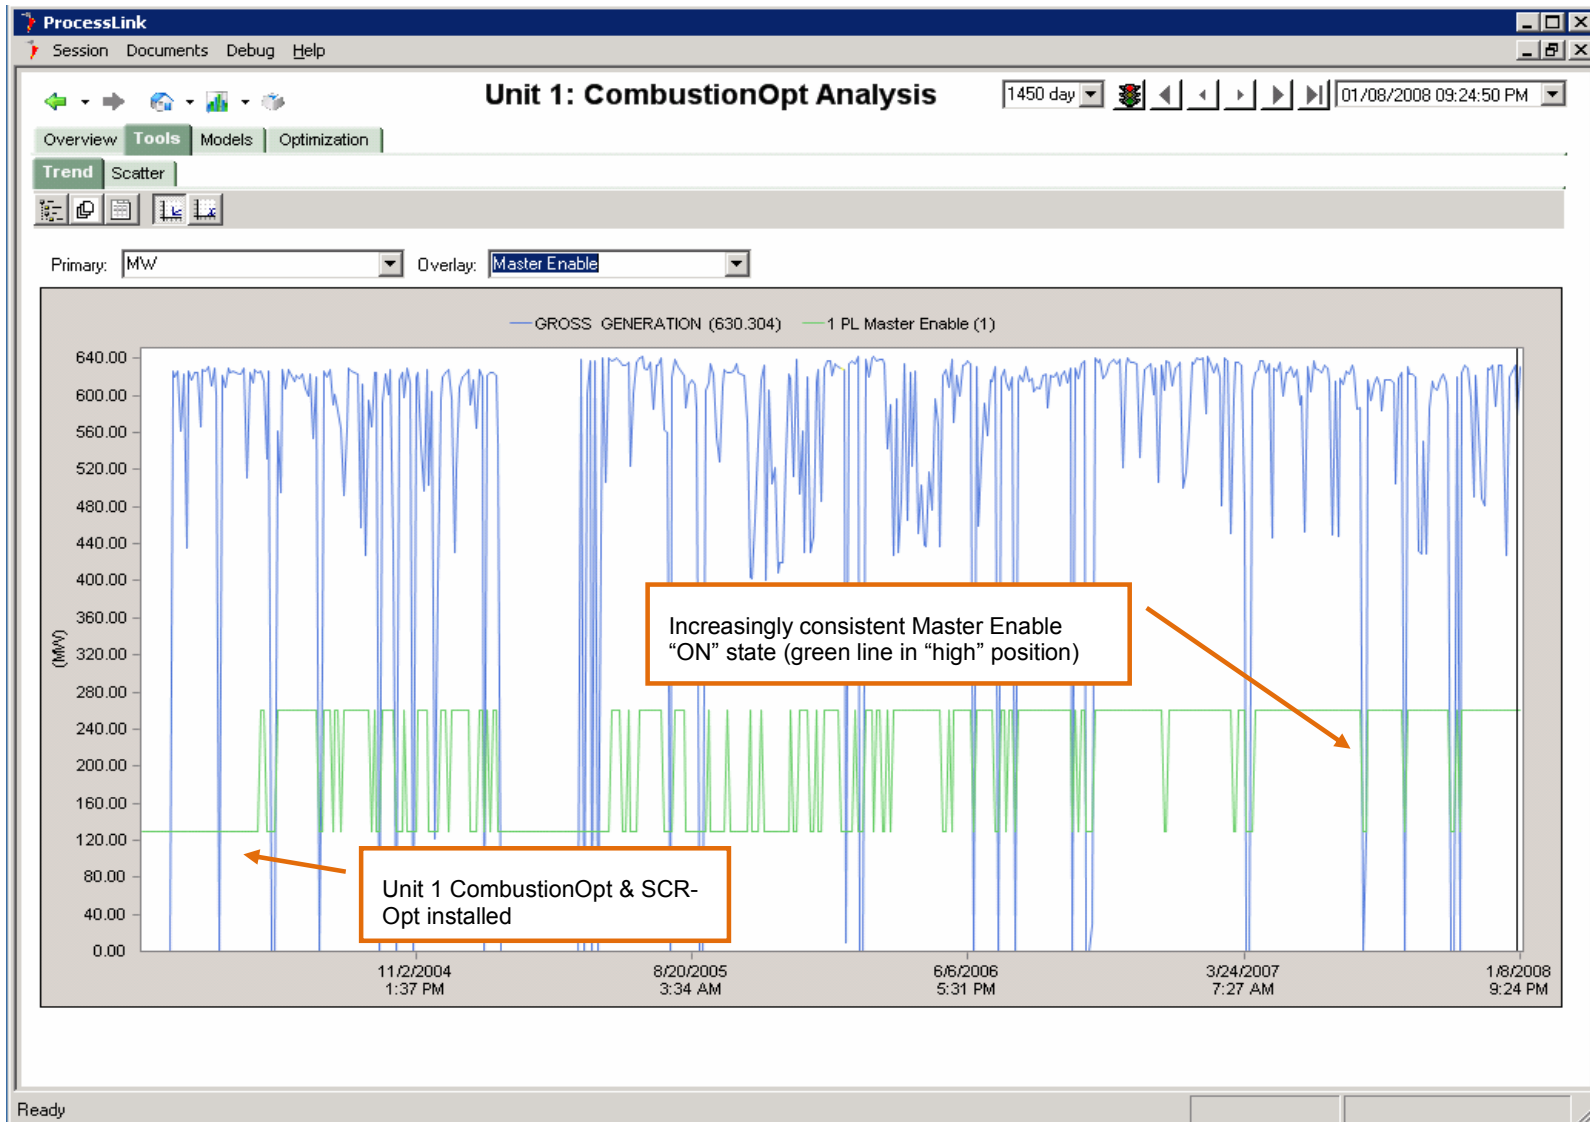


Figure 32 Unit 1 CombustionOpt Master Enable. Primary (left) y-axis is in units of Megawatts (blue trend). CombustionOpt Master Enable Signal is in green with low = Disabled, high = Enabled

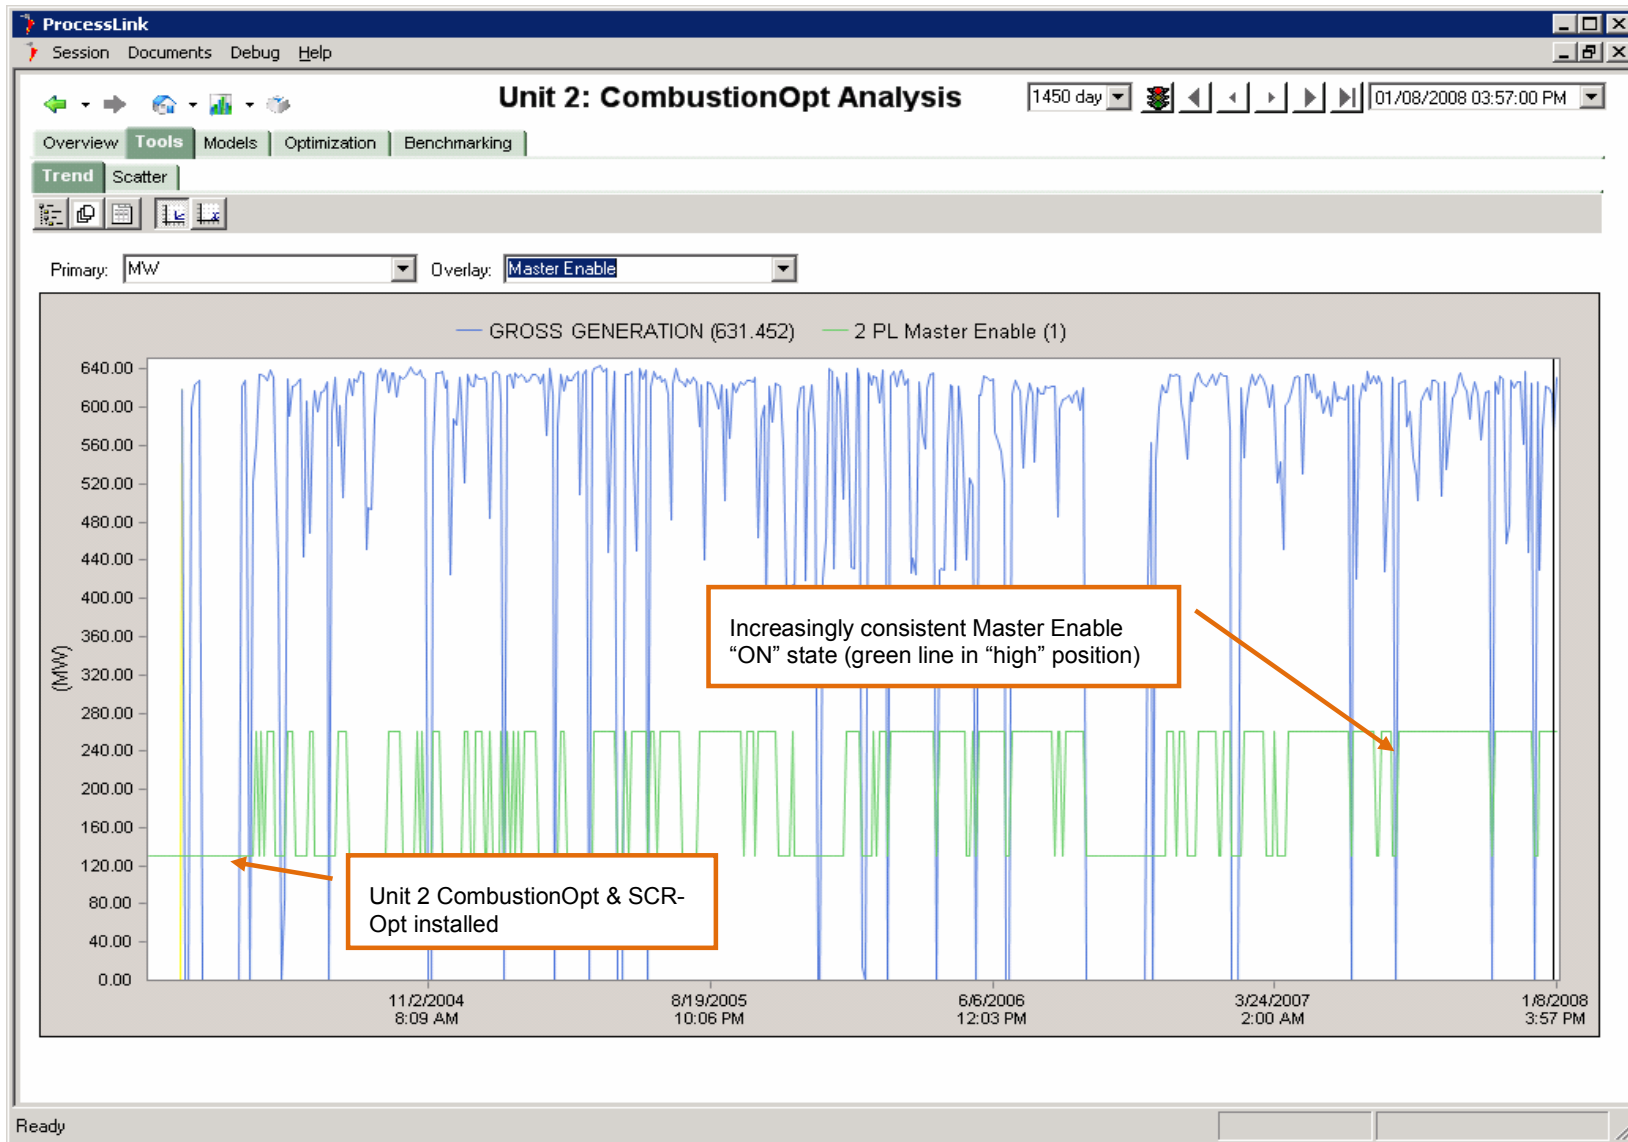


Figure 33 Unit 2 CombustionOpt Master Enable Left y-axis is in units of Megawatts (blue trend). CombustionOpt Master Enable Signal is in green with low position = Disabled, high position = Enabled

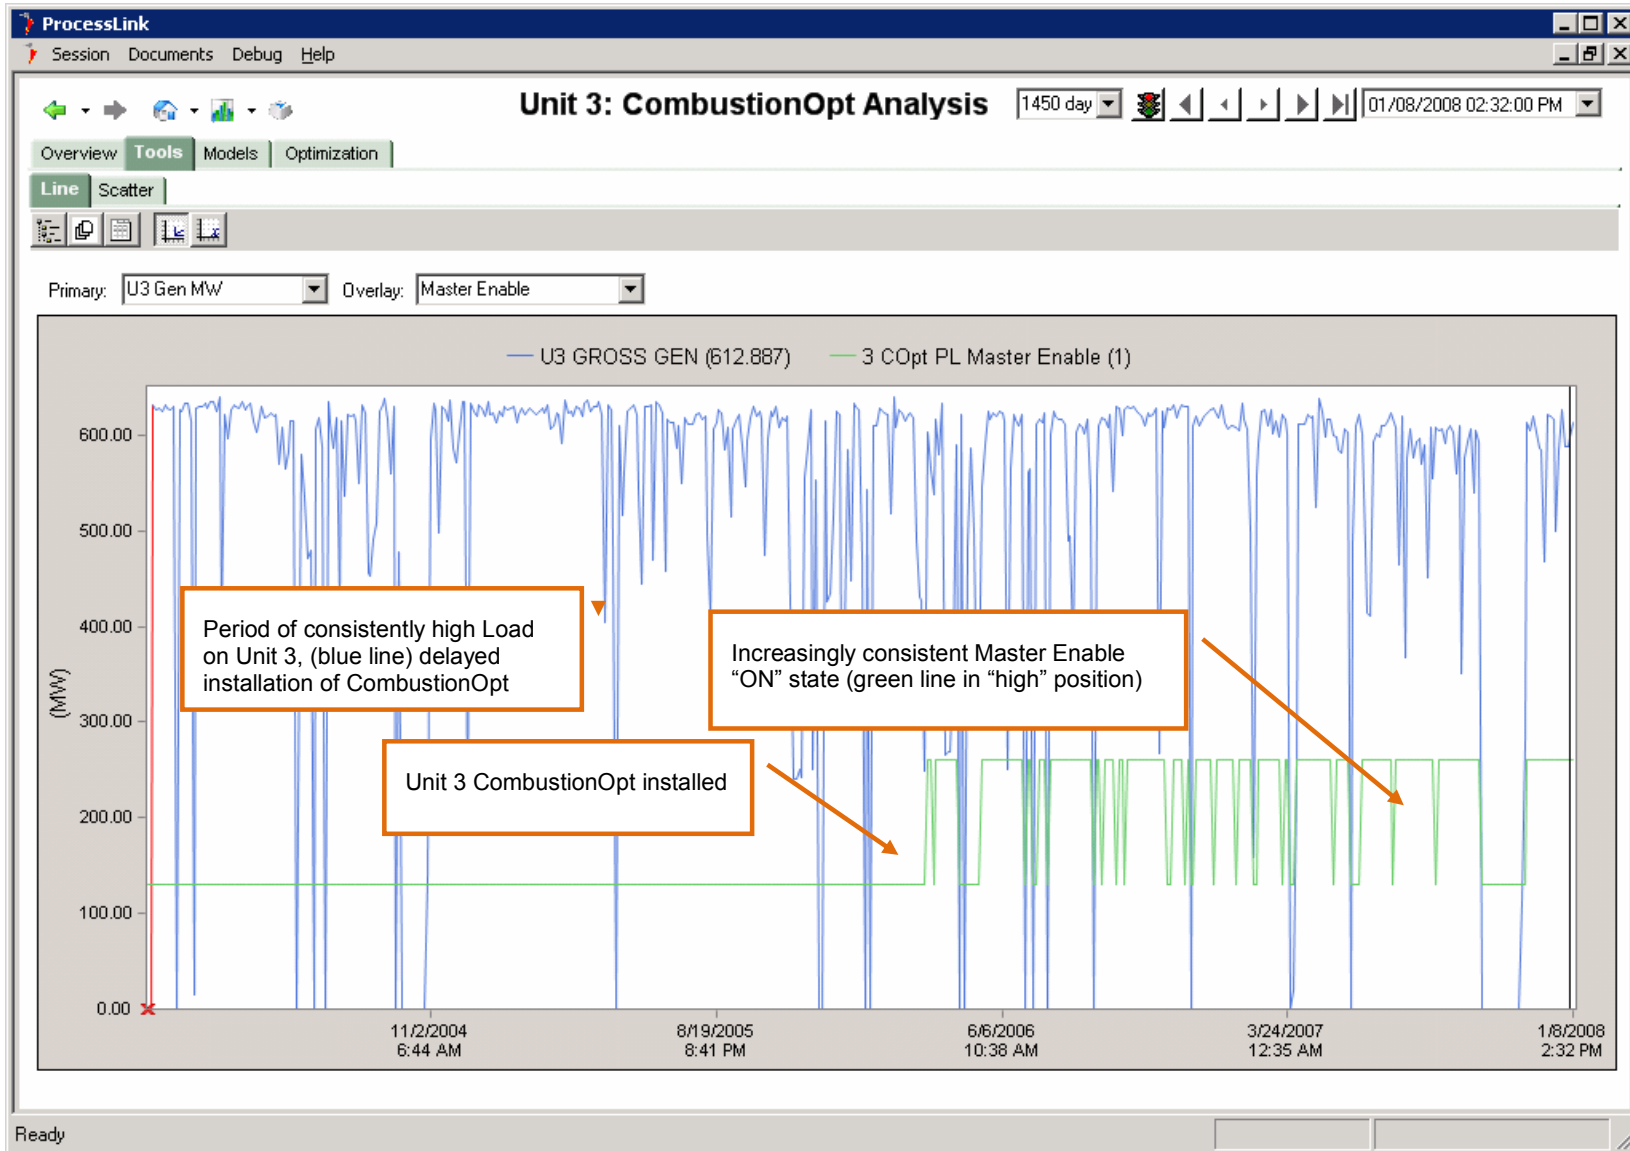


Figure 34 Unit 3 CombustionOpt Master Enable Left y-axis is in units of Megawatts (blue trend). CombustionOpt Master Enable Signal is in green with low position = Disabled, high position = Enabled

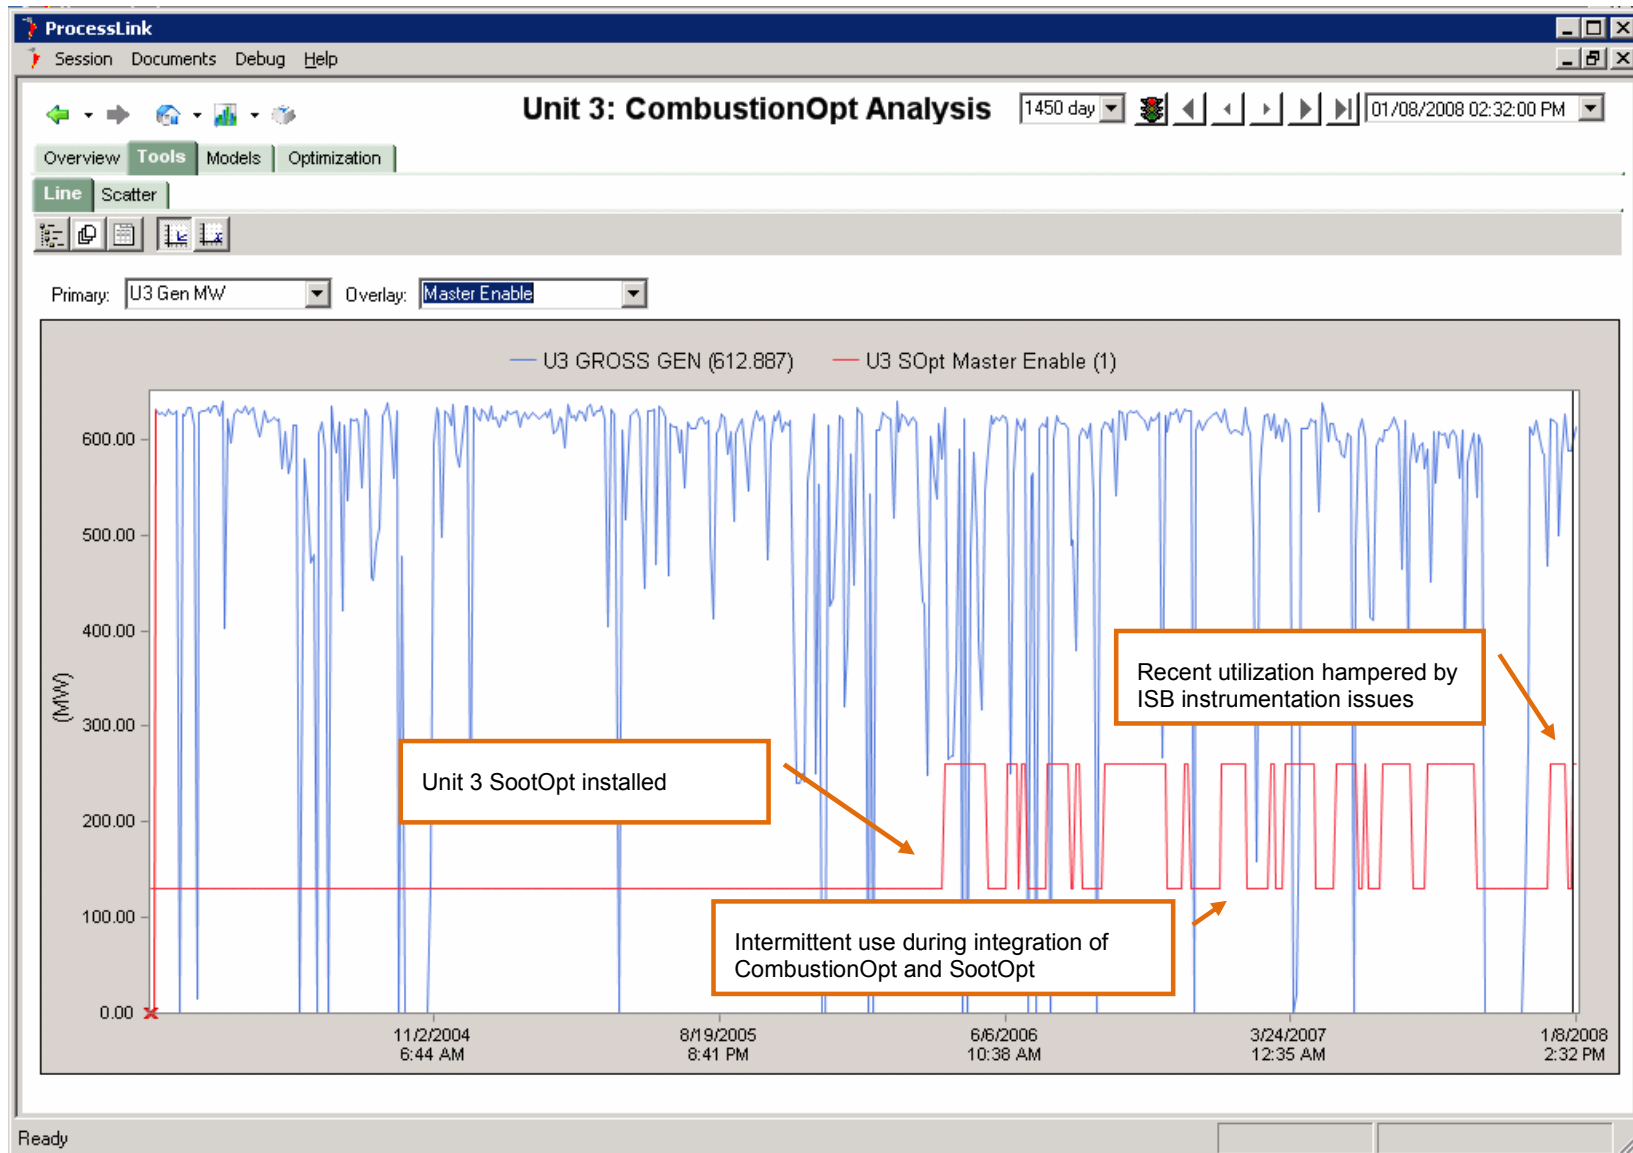


Figure 35 Unit 3 SootOpt Master Enable Left y-axis is in units of Megawatts (blue trend). SootOpt Master Enable Signal is in red with low position = Disabled, high position = Enabled

Figure 36 - Figure 38 shows the Percent of the total available Manipulated Variables Enabled by the operator for each Unit. For SootOpt there are no individual MV Enables. When master enable is "ON" CombustionOpt, SCR-Opt and SootOpt have permissive to directly adjust the fuel and air distribution biases, or in the case of SootOpt direct the automated activity of the sootblowing system. In the case of CombustionOpt the operator can also enable or disable individual fuel and air distribution biases. SootOpt is either fully ON when its Master Enable is True (1, high) or fully OFF when it is False (0, low)

Observations:

- The red in the Unit 3 percentage MV's trend (Figure 37) indicates that data is in error (prior to the installation of the system and the data points indicating the status of the MVs status.)
- For SootOpt there are no individual enables. It is either on or off and depends only on the health of the ISB and the operator master SootOpt Enable (Figure 35).
- This project included the installation of some 50+ hooks for CombustionOpt to potentially manipulate in the Control Systems for Unit 1 and Unit 2, with 30+ for Unit 3. The average for a commercial installation is between 20 and 40 installed MVs. In the future these MVs can be actuated in closed-loop with the flick of a switch, adjustment of constraints, and some close observation of initial moves.
- At BEC's request the introduction of optimization activity was structured to address the multivariable control problems they most wanted to improve and felt confident represented waiting opportunity, namely getting better control of the 28 cyclones on Units 1 and 2 and the numerous fuel-air and aux-air dampers on Unit 3 This has allowed them to minimize any risk to these high-value, high-capacity factor units. Although good results have been seen, many strong levers have yet to be employed for optimization, indicating that there is opportunity for even greater benefits.

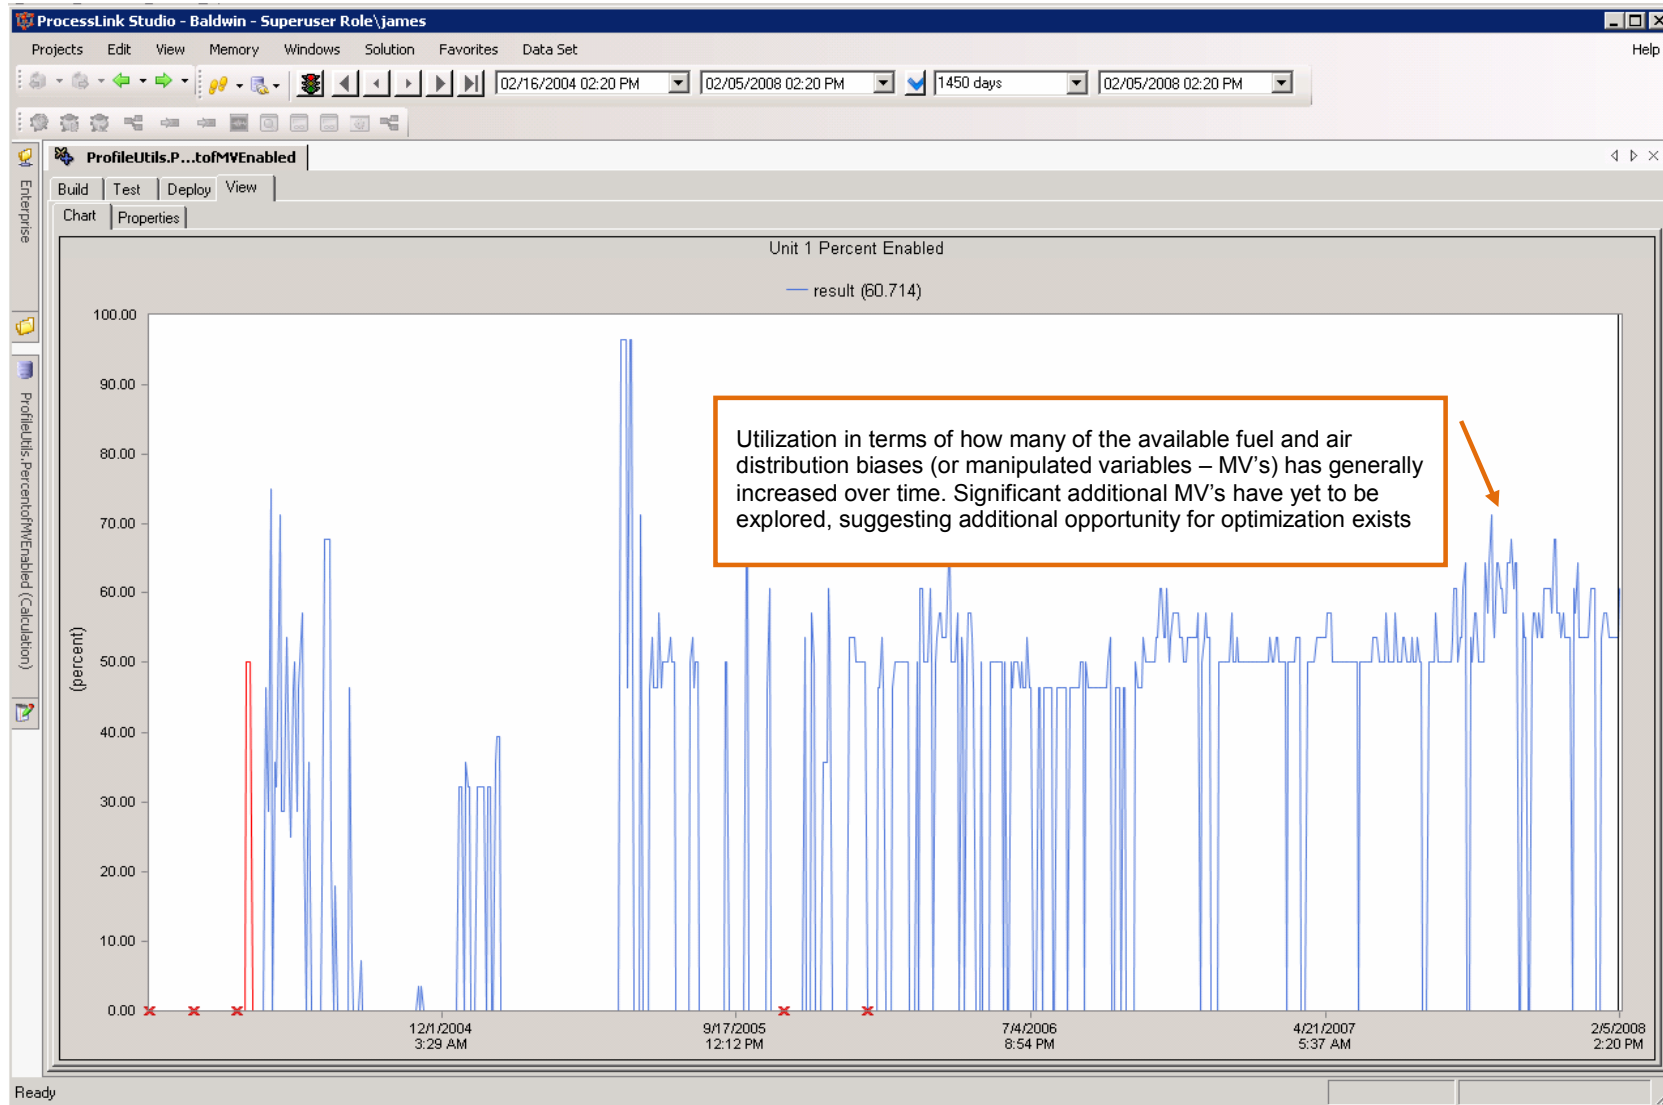


Figure 36 Unit 1 % Total MV's Enabled. This value equals the number of fuel and air distribution biases the operator has enabled, divided by the number of biases he has available to enable. For any individual biases to be enabled the Master Enable must be true or "ON"

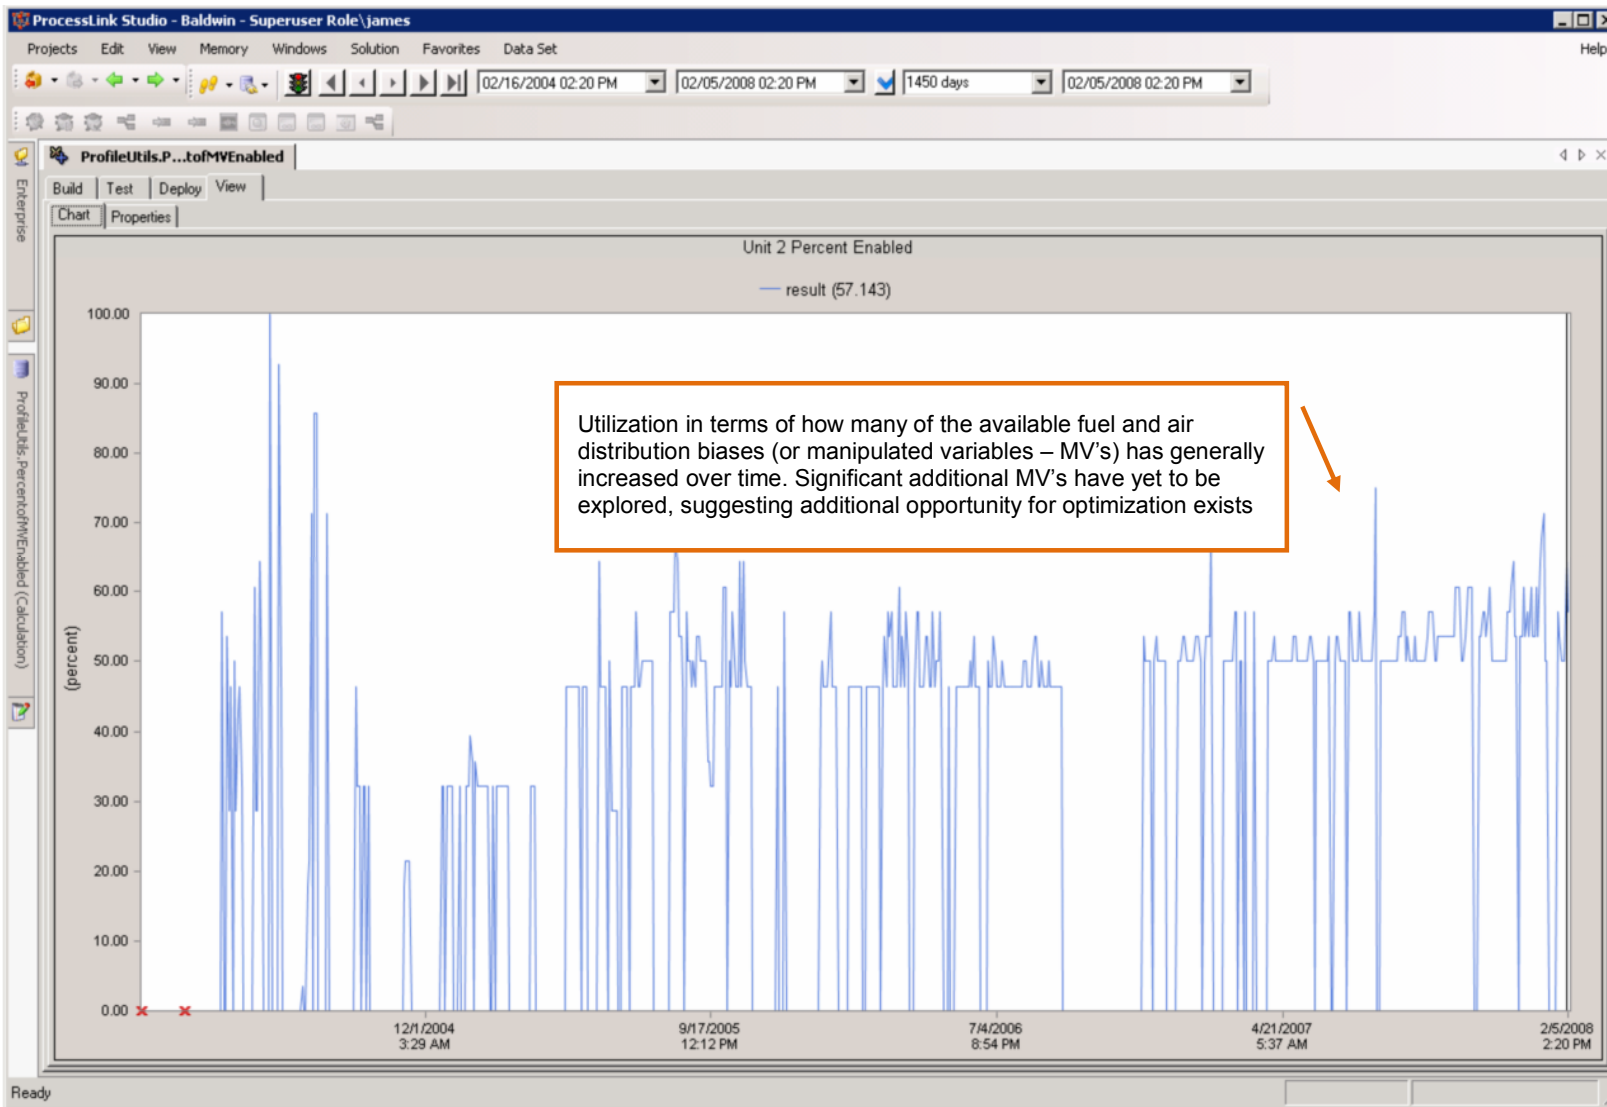


Figure 37 Unit 2 % Total MV's Enabled This value equals the number of fuel and air distribution biases the operator has enabled, divided by the number of biases he has available to enable. For any individual biases to be enabled the Master Enable must be true or "ON"

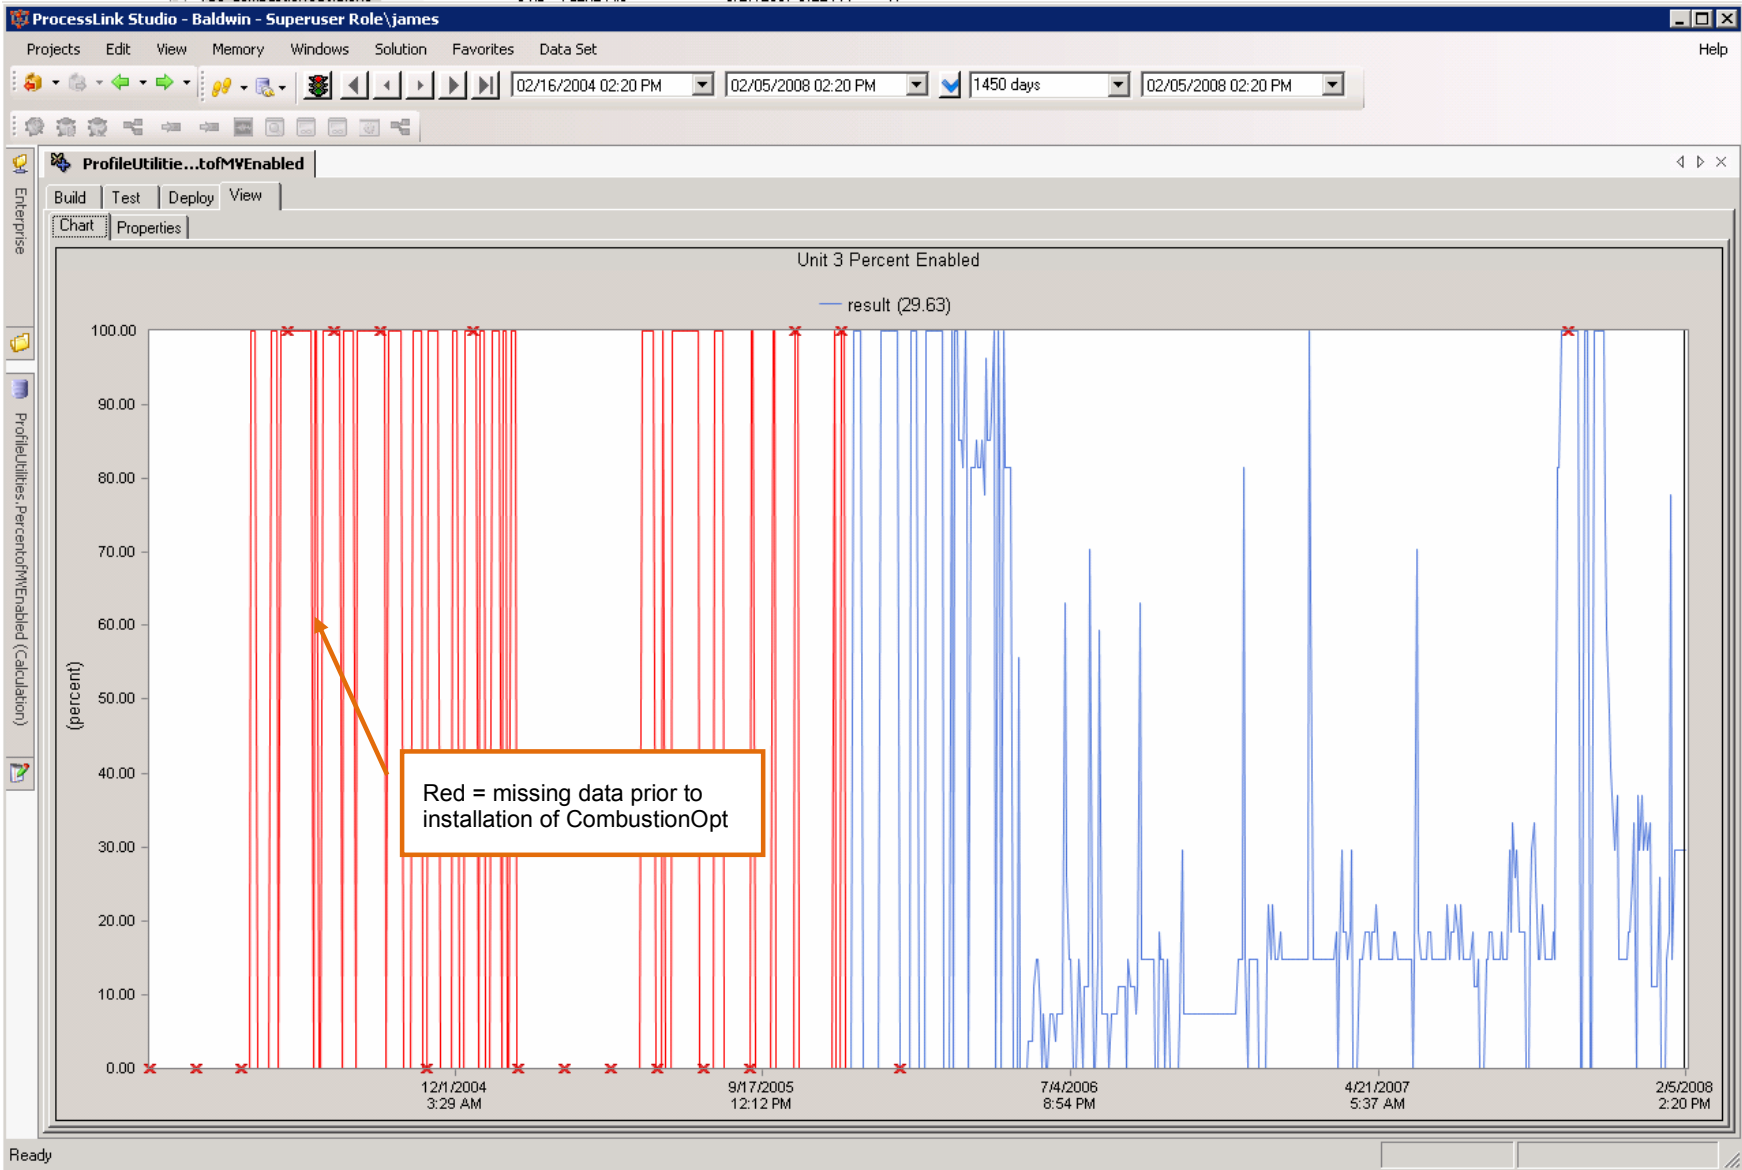


Figure 38 Unit 3 % Total MV's Enable

4.4 Closed-Loop Optimization Analysis

The methodologies used for estimating the optimization impacts are:

- Comparison of control population to experimental population
- Modeled response to disturbances and factors

4.4.1 Comparison of Populations Methodology

This analysis consists of comparing the properties of the populations, selected to represent the experiment and the control.

- Advantages
 - Easy to implement
- Disadvantages
 - It's hard to tell whether the experimental factor is what is causing the difference between the two populations in cases where there are disturbances that may not be equally present in both.

The time ranges that define the data selections for the control and experimental populations are 365 days prior to roughly 4/5/06 for the OFF or pre-optimization case (i.e. the control), and 365 days prior to roughly 2/8/08 for the ON or post-optimization case [Note these time range selection criteria are in addition to the other Optimizer ON, OFF criteria]. Note: All data from times before the last catalyst beds were installed is excluded from the experiment and control populations.

The categorized scatter plots shown in Figure 39 and in Section 4.4.1.1 (Measurements) will show:

- **[Left]** A population selected to represent the control (no optimization, CombustionOpt, SCR-Opt and SootOpt Master Enable OFF)
- **[Right]** A population selected to represent the experiment (CombustionOpt and SCR-Opt ON for Units 1&2 and CombustionOpt and SootOpt Master Enable ON for Unit 3),

In addition all populations are subject to filtering to 1) remove bad data (meaning bad status) and outliers (infeasible values 2) guarantee that the population is the one of interest, i.e. for Units 1 and 2 we are interested mostly in the effect at relatively high capacity (>600 MW) when the SCR is on, since that is where they spend most of their time; for Unit 3 the analysis is of the high load regime also (>600 MW) the SootOpt ON factor is a criterion for inclusion in the experimental population.

Tips for looking at the plots:

In most of the plots, the x-axis value is Megawatts, which is a major cofactor for all the things we are interested in looking at (namely NO_x, NH₃, and HR).

The Blue Line represents the Mean Value of each population.

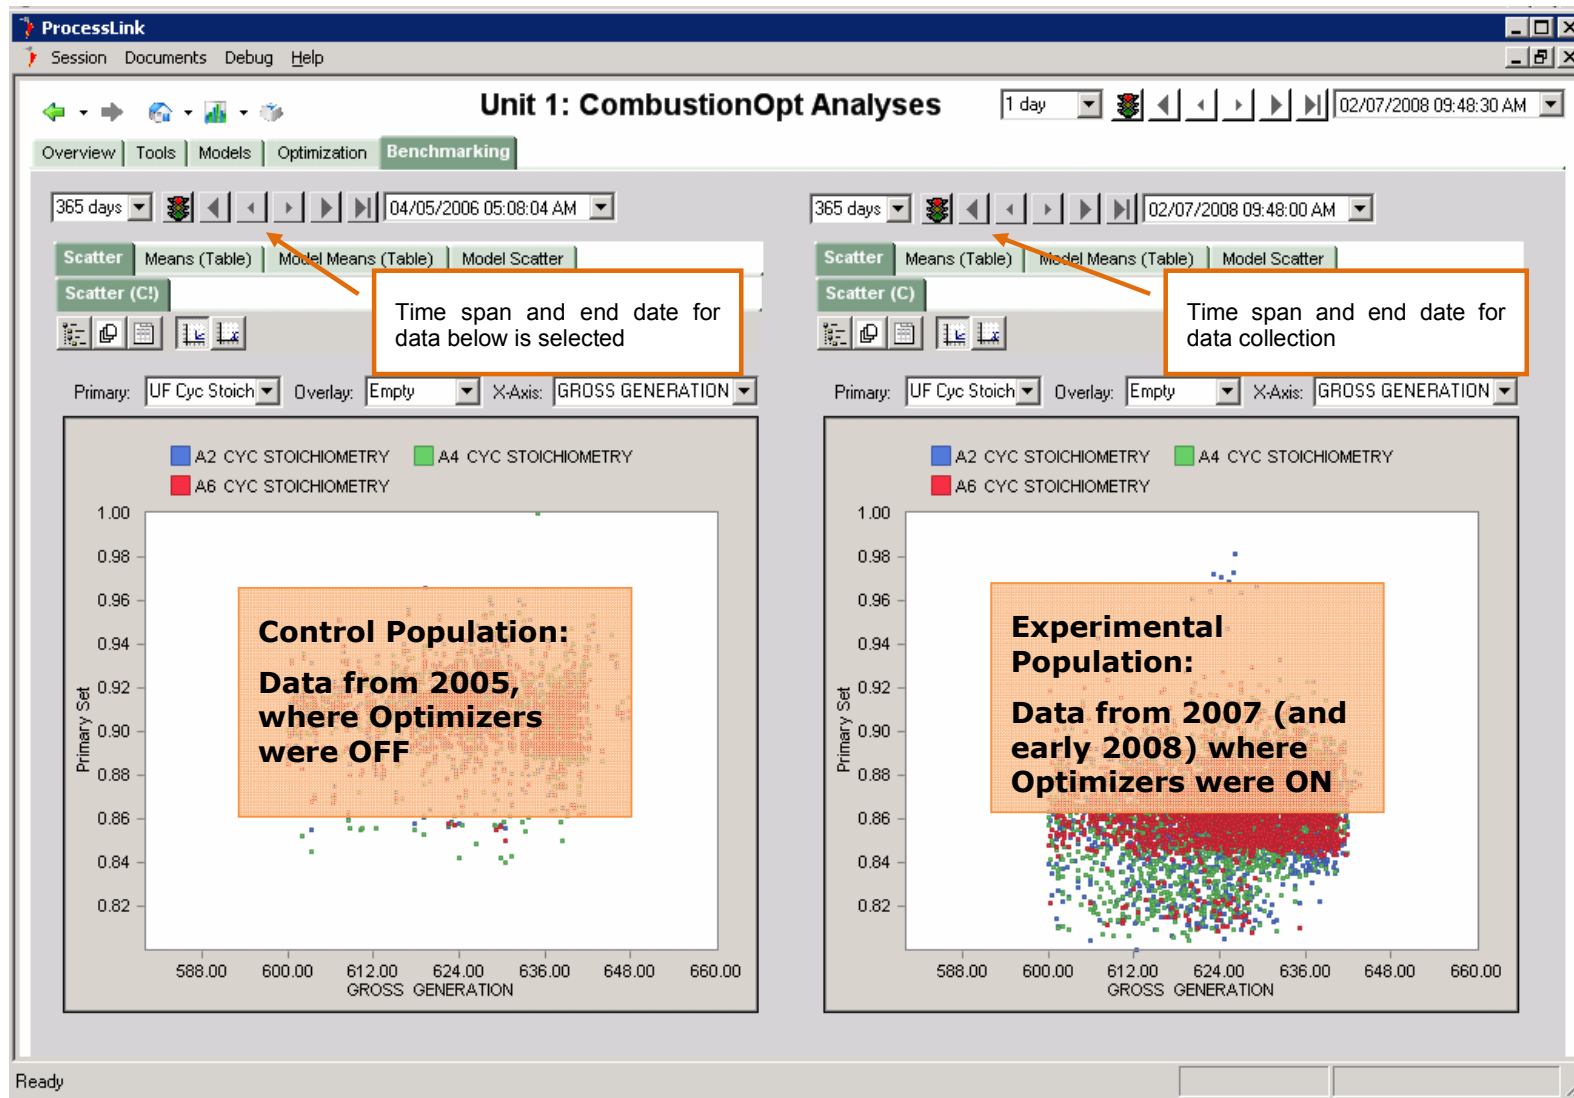


Figure 39 Layout of Categorized Scatter Plots

4.4.1.1 Measurements

4.4.1.1.1 NOx vs. MW

Figure 40 - Figure 41 show the relationship between CEMS NOx and Megawatts for Units 1, 2 and 3 for the control and experimental populations. Note: All data from times before the last catalyst beds were installed are excluded from the experiment and control populations.

Observations:

- For the SCR Units (1 & 2), CEMS NOx is controlled in PID by the injection of NH3. Some change in CEMS NOx is due to the optimization of the Furnace+SCR system reducing SCR NH3 controller error (hunting), but most difference is accounted for by changes to the NOx setpoint.
- Even though for Units 1 and 2 these are controlled changes (by the SCR CEMS NOx PID controller), any differences they represent do count in the assessment of impact, since without some kind of optimization, any reduction of the NOx setpoint must come with a corresponding decrease in SCR Inlet NOx (furnace NOx) conditions or an increase in NH3 flow (representing additional furnace NOx conversion). Any change seen without these offsets represents found optimization leverage (See section 5 "Managing Tradeoffs and Estimating Optimization Impact").
- The bimodal aspect of the distributions represents NOx season operations and non-NOx season operations, during which the NOx removal setpoints given to the SCR were different.
- For Unit 1 the experimental data is somewhat more tightly distributed than the control.
- For Unit 1 the density of samples at the lower setpoint (0.5 #/MBtu) appears to be greater in the experimental population
- For Unit 2 the experimental data may be somewhat more tightly distributed but it is less clear.
- For Unit 2 the density of samples at the lower setpoint (0.5 #/MBtu) may be greater in the experimental population
- The differences between SCR NOx control on Unit 2 and Unit 1 are widely different. Although Unit 2 received its 4th catalyst bed one year before Unit 1 (2004 rather than 2005) in all the data shown here is for both units have 4 catalyst beds. Although there is a lot to wonder about in this difference, the plant's expectation that these two superficially identical units should behave the same was long ago abandoned.
- For Unit 3 the distribution is less tightly controlled in the experimental data than in the control but also visibly lower.

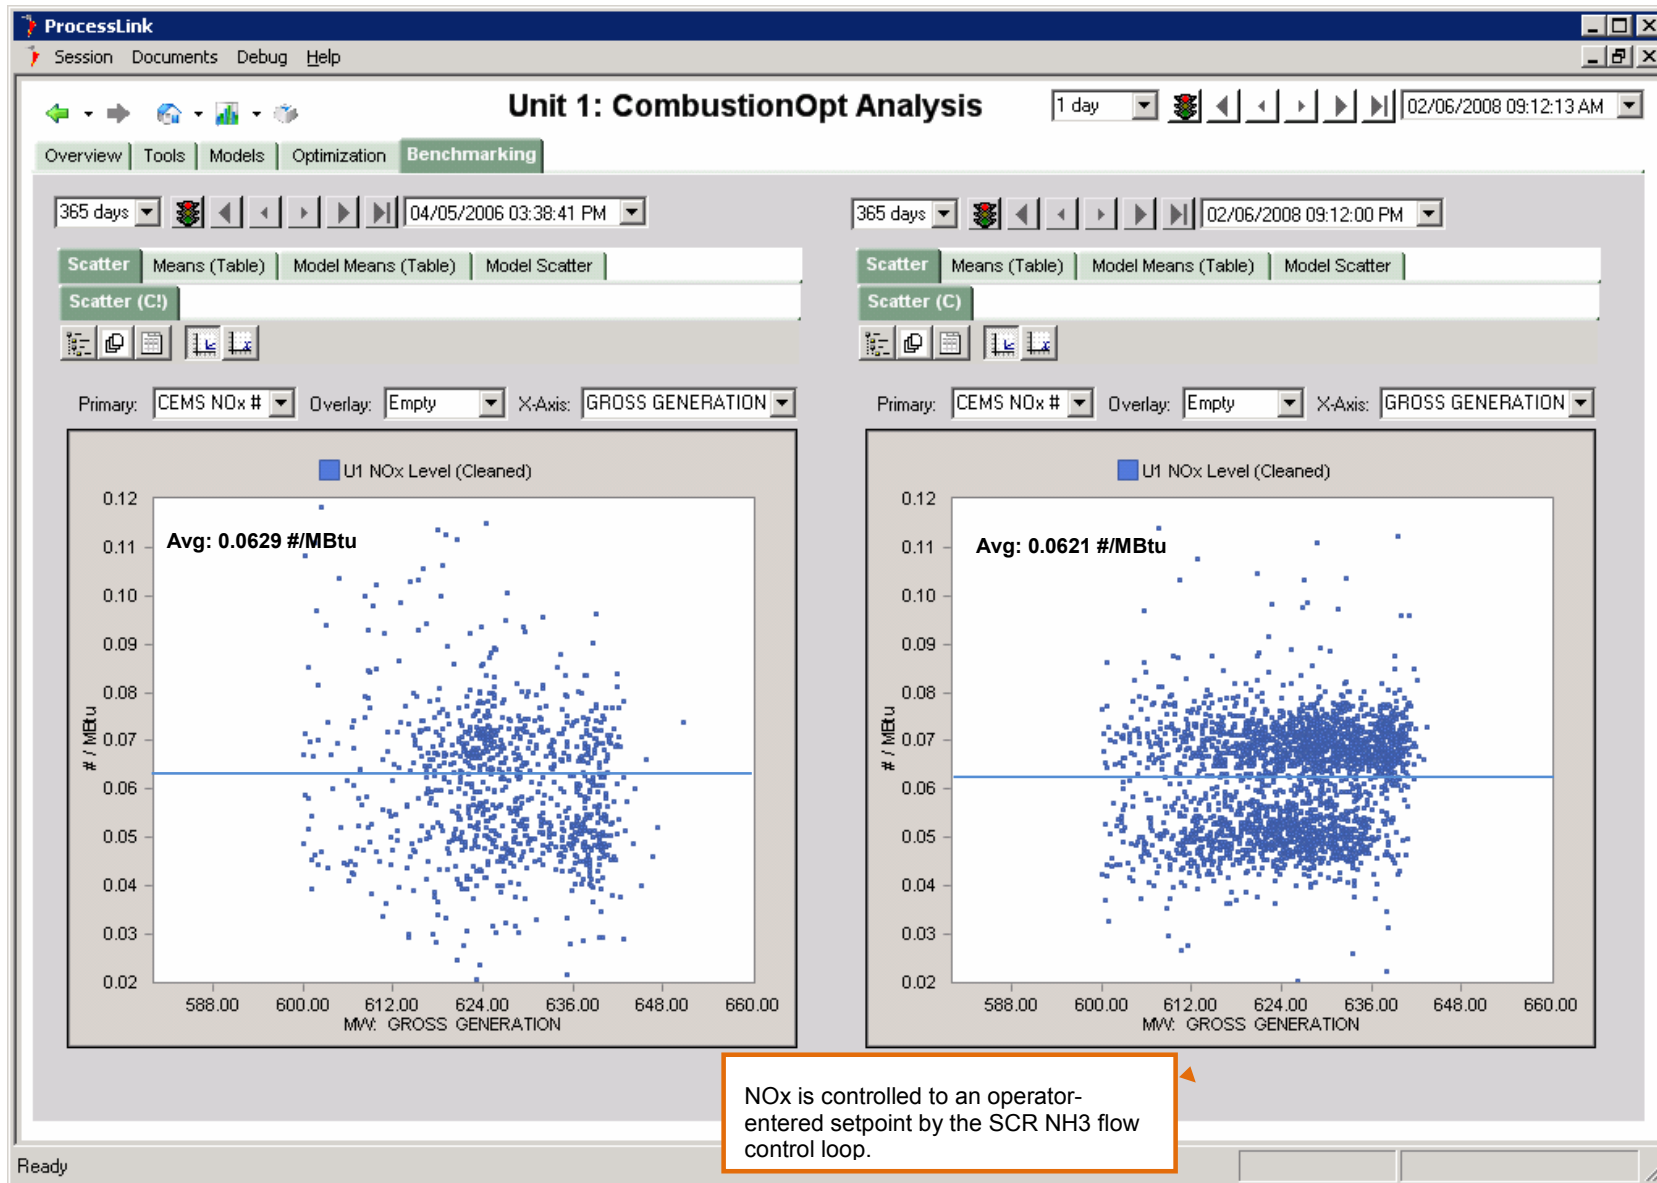


Figure 40 Unit 1 Measured NOx vs. MW

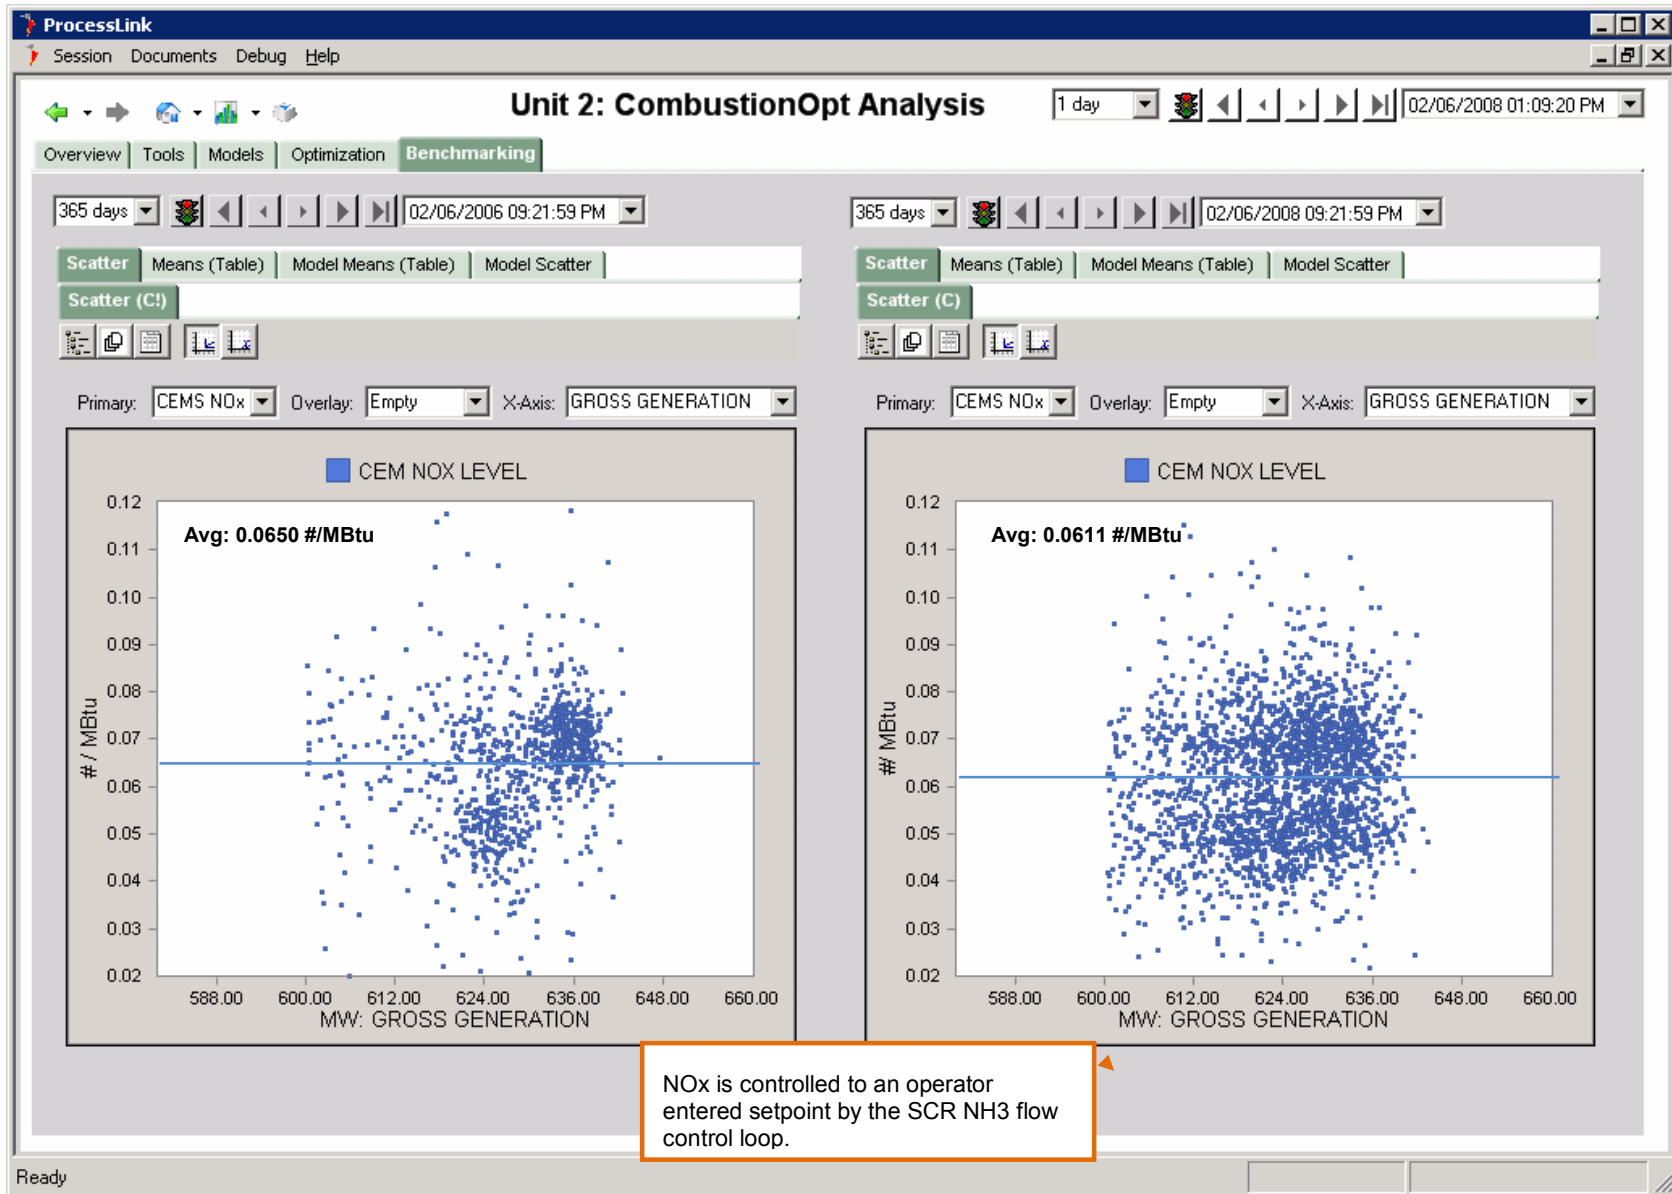


Figure 41 Unit 2 Measured NOx vs. MW

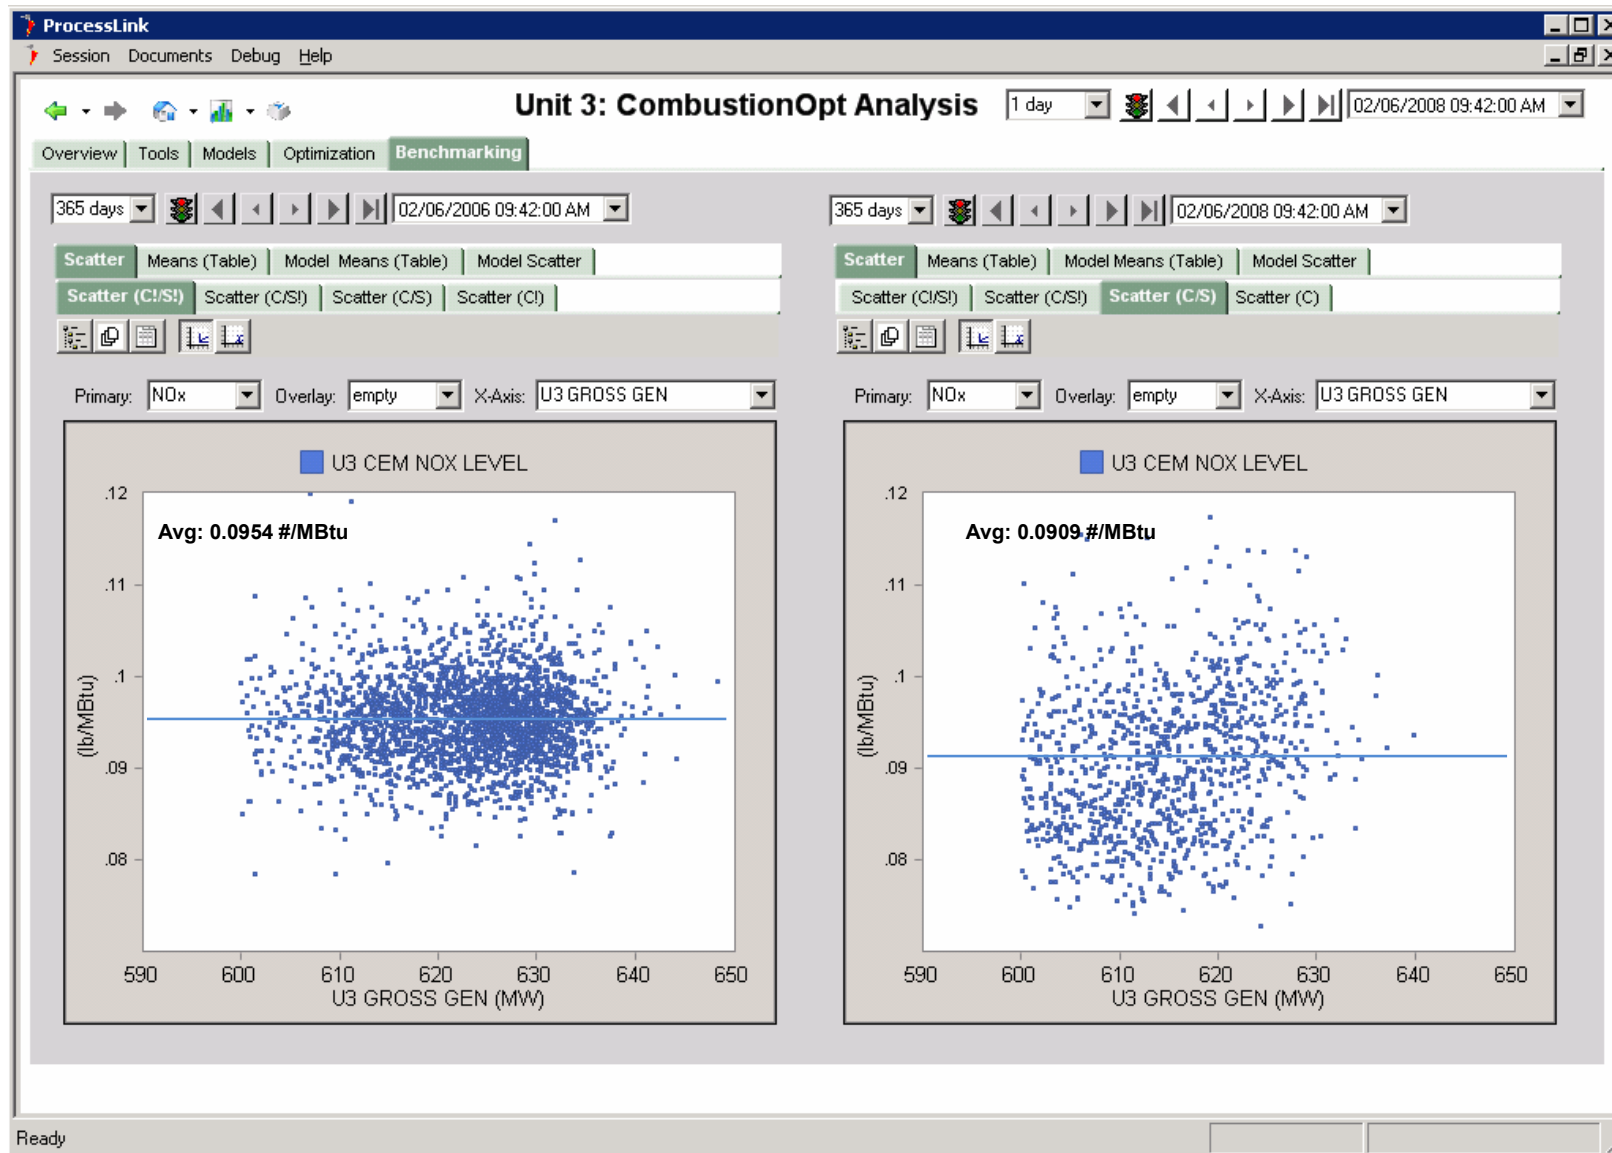


Figure 42 Unit 3 Measured NOx vs. MW

4.4.1.1.2 NH3 vs. MW

Figure 43 - Figure 44 show the relationship between SCR NH3 Flow and Megawatts for Units 1 and 2. Optimization of a unit with an SCR involves a net reduction in NH3 Flow (representing furnace NOx), or CEMS NOx. Any reduction in one, without increase in the other, represents reduced NOx production.

Observations:

- The difference in NH3 flow on Unit 1 is pronounced
- A difference on Unit 2 is less pronounced but significant.
- As mentioned above, these superficially identical units have distinct reputations. One obvious difference between them is that Unit 1 has a slightly newer 4th catalyst bed, however all data shown here is with 4 catalyst beds each in both units.

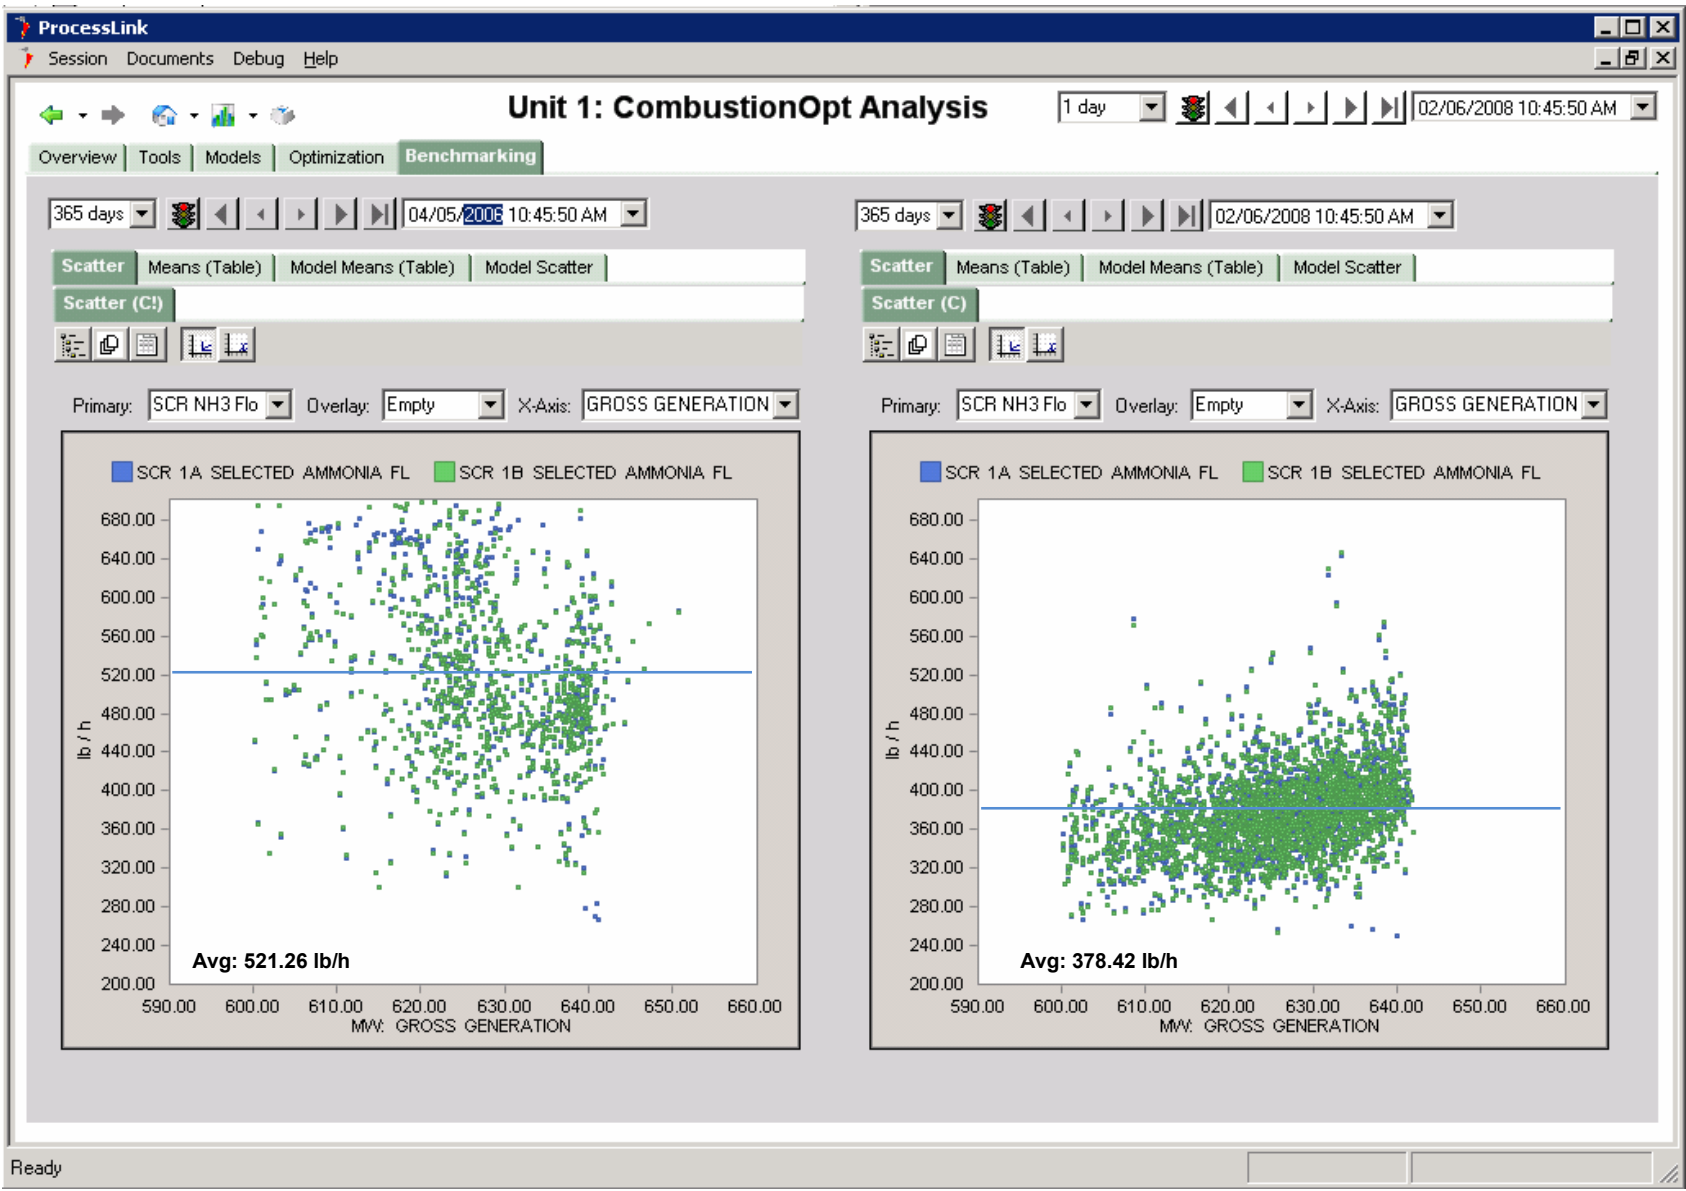


Figure 43 Unit 1 Measured NH3 vs. MW

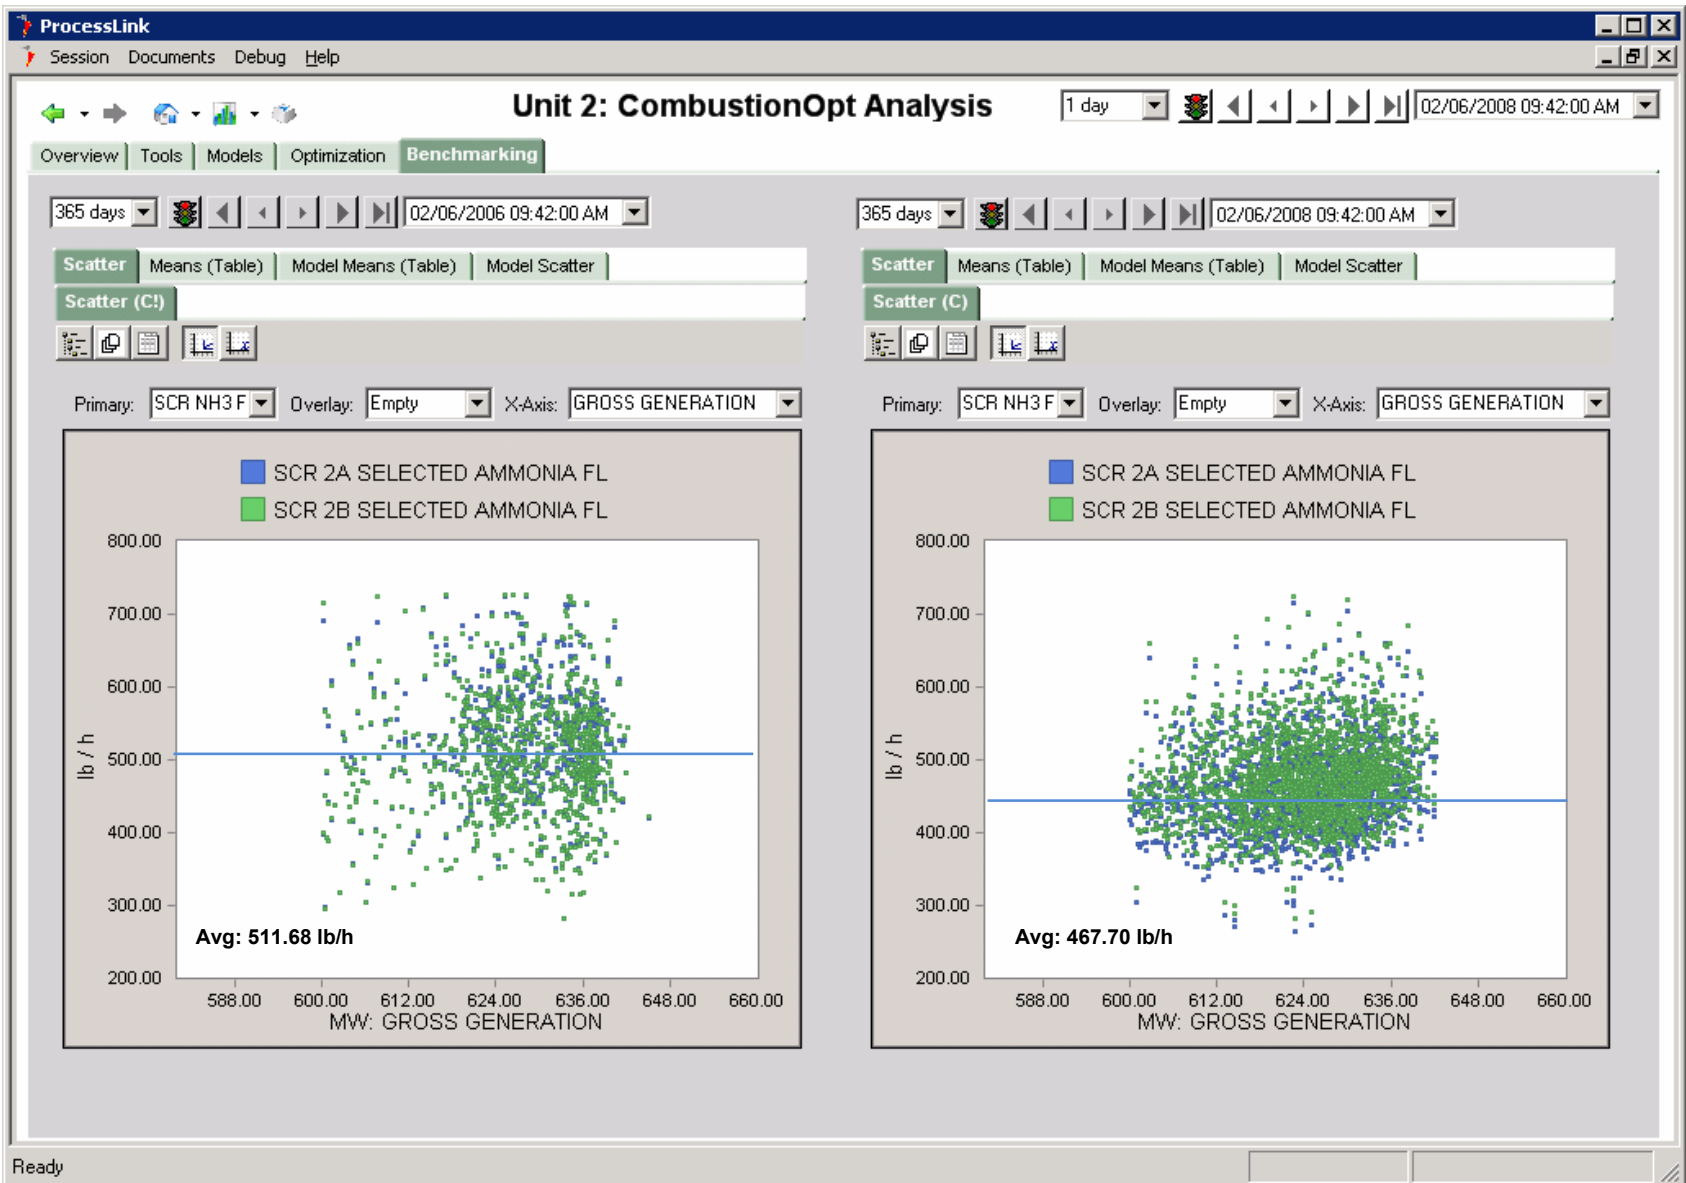


Figure 44 Unit 2 Measured NH3 vs. MW

4.4.1.1.3 HR vs. MW

Figure 45 - Figure 46 show the relationship between Heat Rate and Megawatts for Units 1, 2 and 3.

Observations:

- Heat Rate (especially) is a very noisy signal, and subject to significant disturbance from, among other things, ambient temperature and humidity, condenser cleanliness and of course MW.
- The Heat Rate signal being used here is from Eta Pro. Although PerformanceOpt also provides a very accurate Heat Rate measure, data for it does not go back as far and the use of the measure which was vetted historically by the plant prior to the project helps validate the experiment.
- Heat rate improvements are apparent in the measurements for both Unit 1 and 2, with Unit 3 showing a slight increase.
- Unit 1 Heat Rate is significantly lower than Unit 2
- Heat rate, NO_x and NH₃ flow are subject to disturbances from ambient conditions, condenser cleanliness, and especially load (among other things). In sections to follow, multi-variable regression modeling will be used to more accurately assess the differences between these populations and isolate out disturbances.

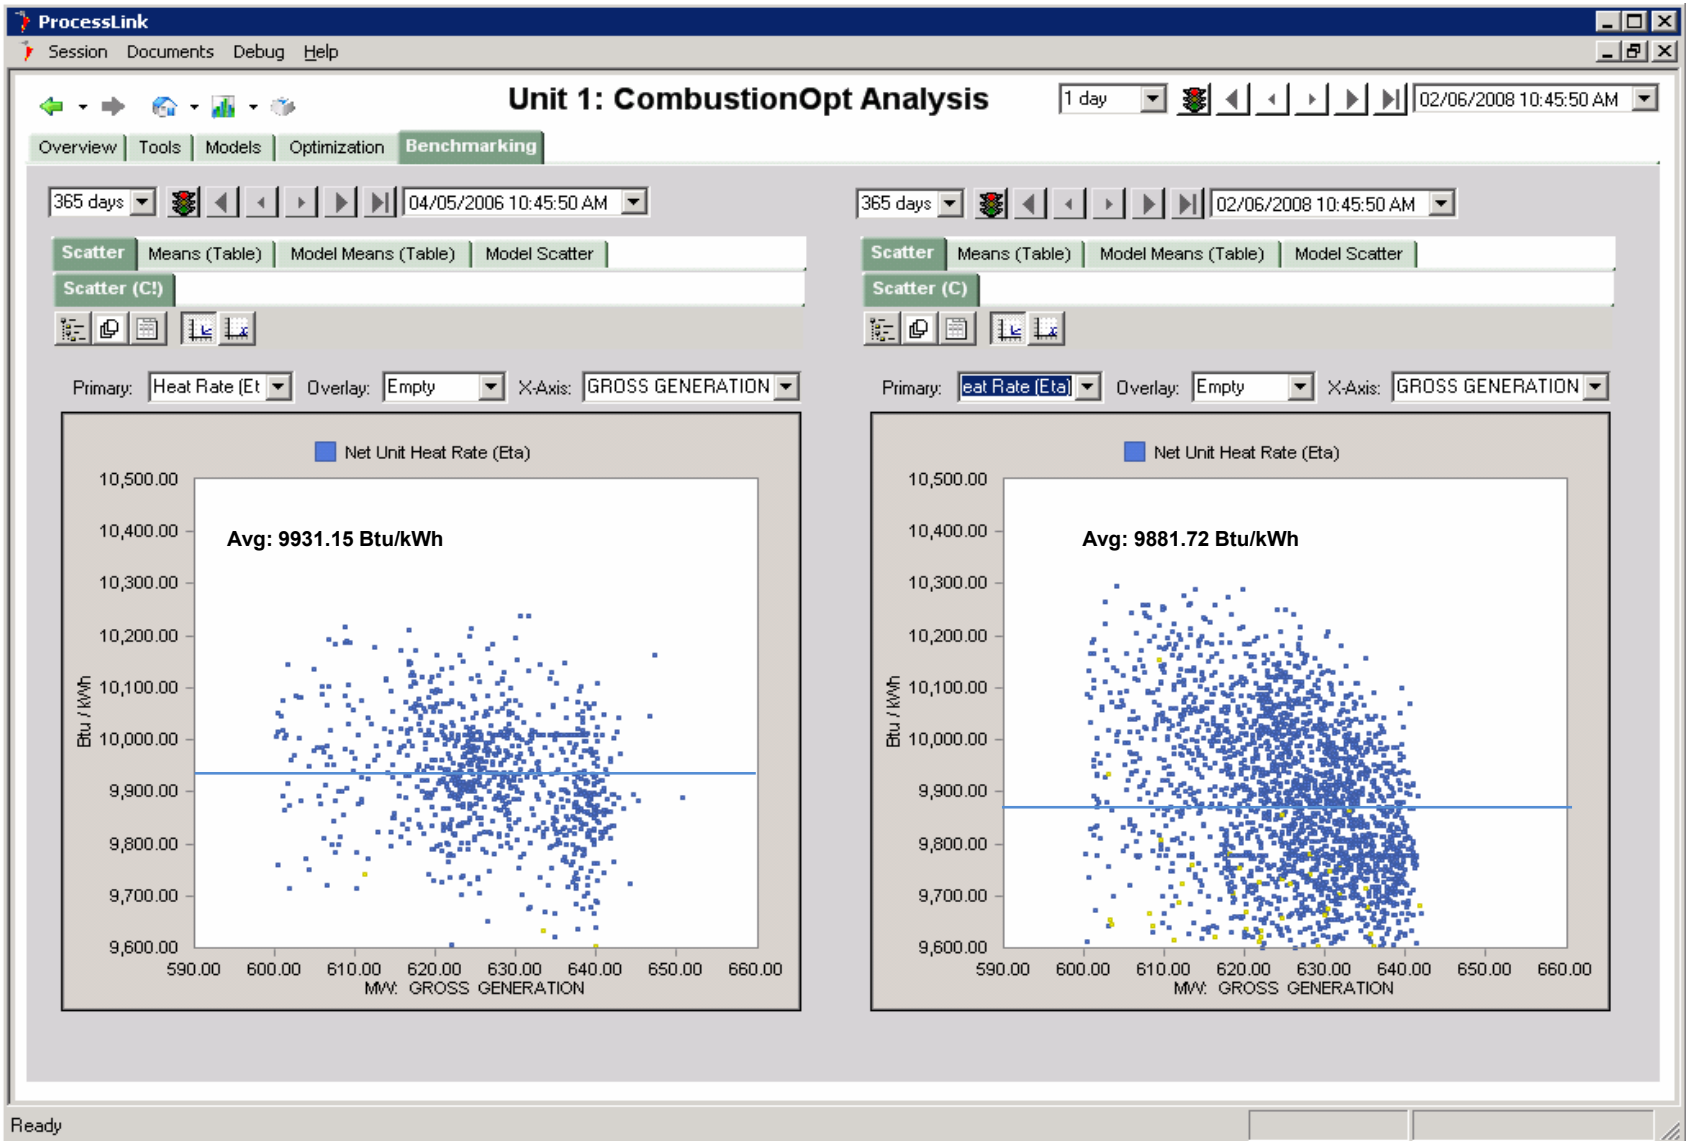


Figure 45 Unit 1 Measured Heat Rate vs. MW

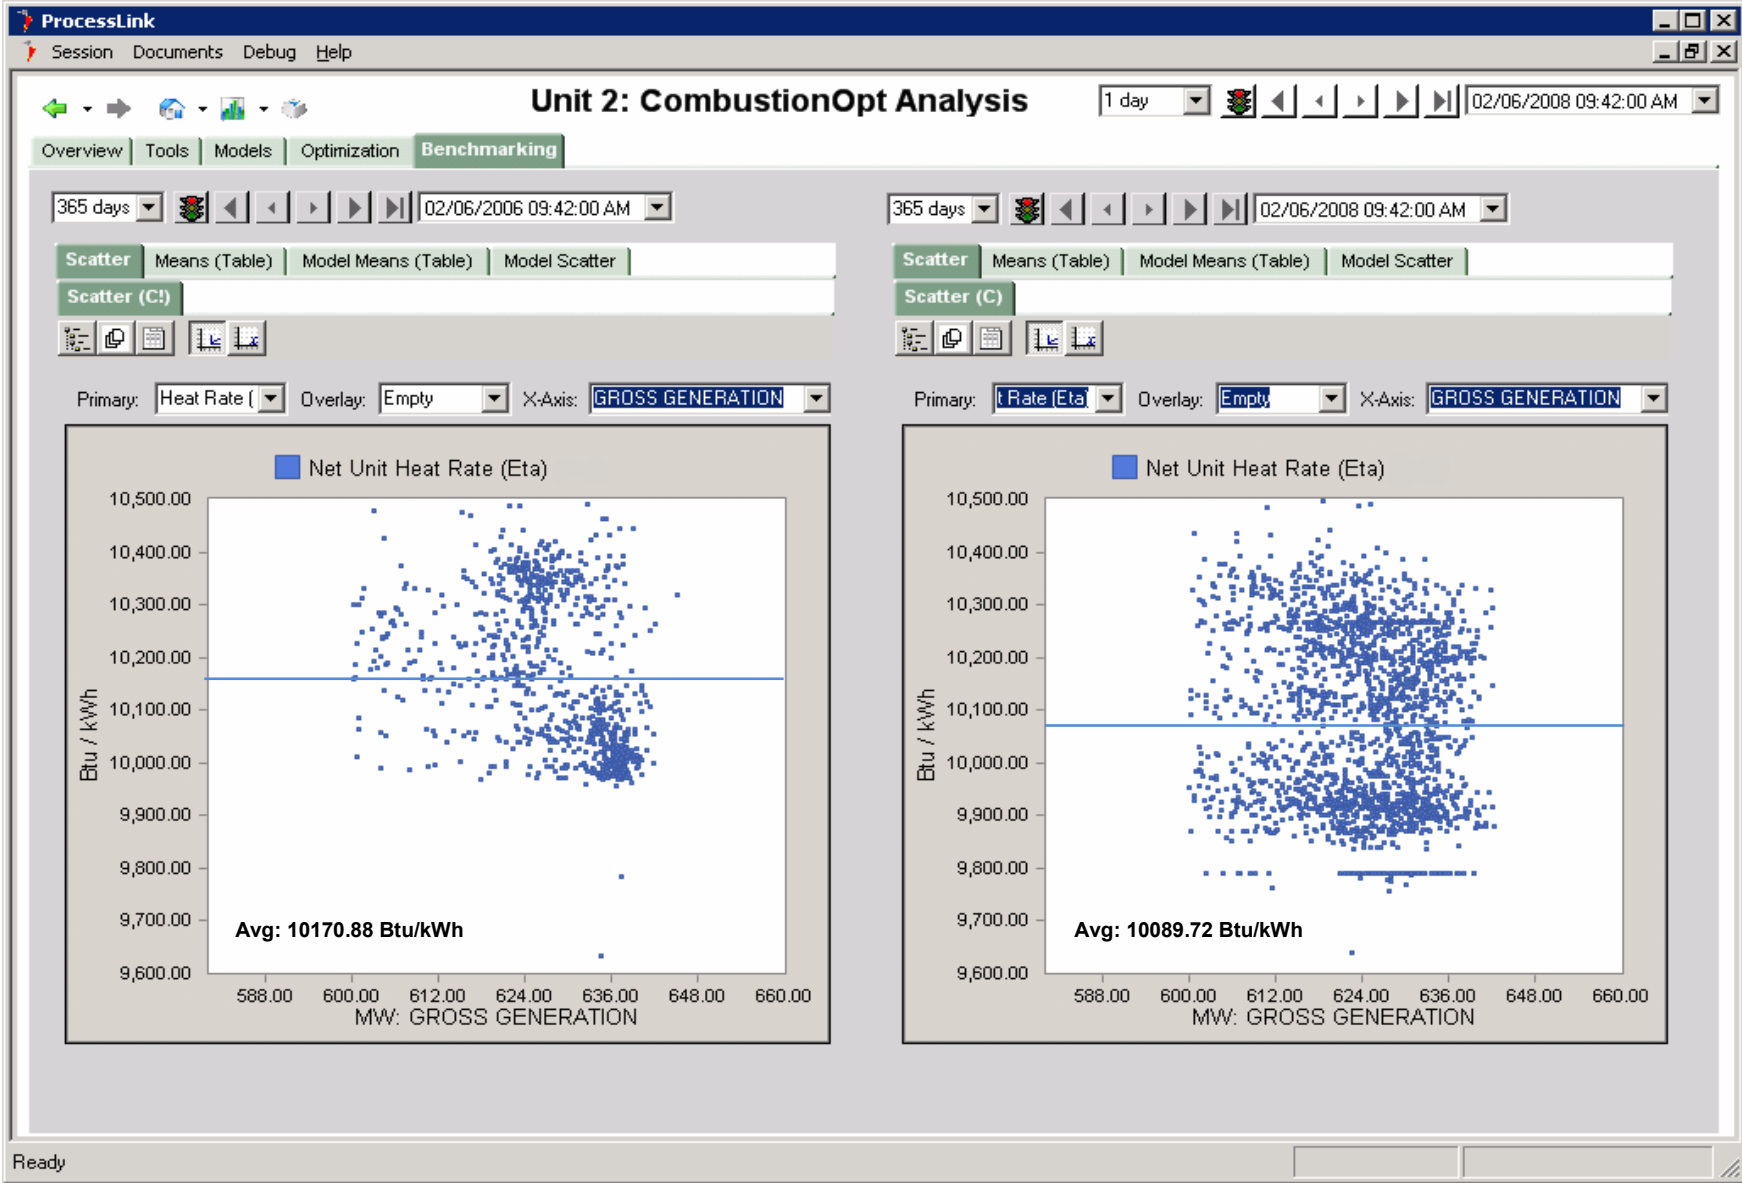


Figure 46 Unit 2 Measured Heat Rate vs. MW

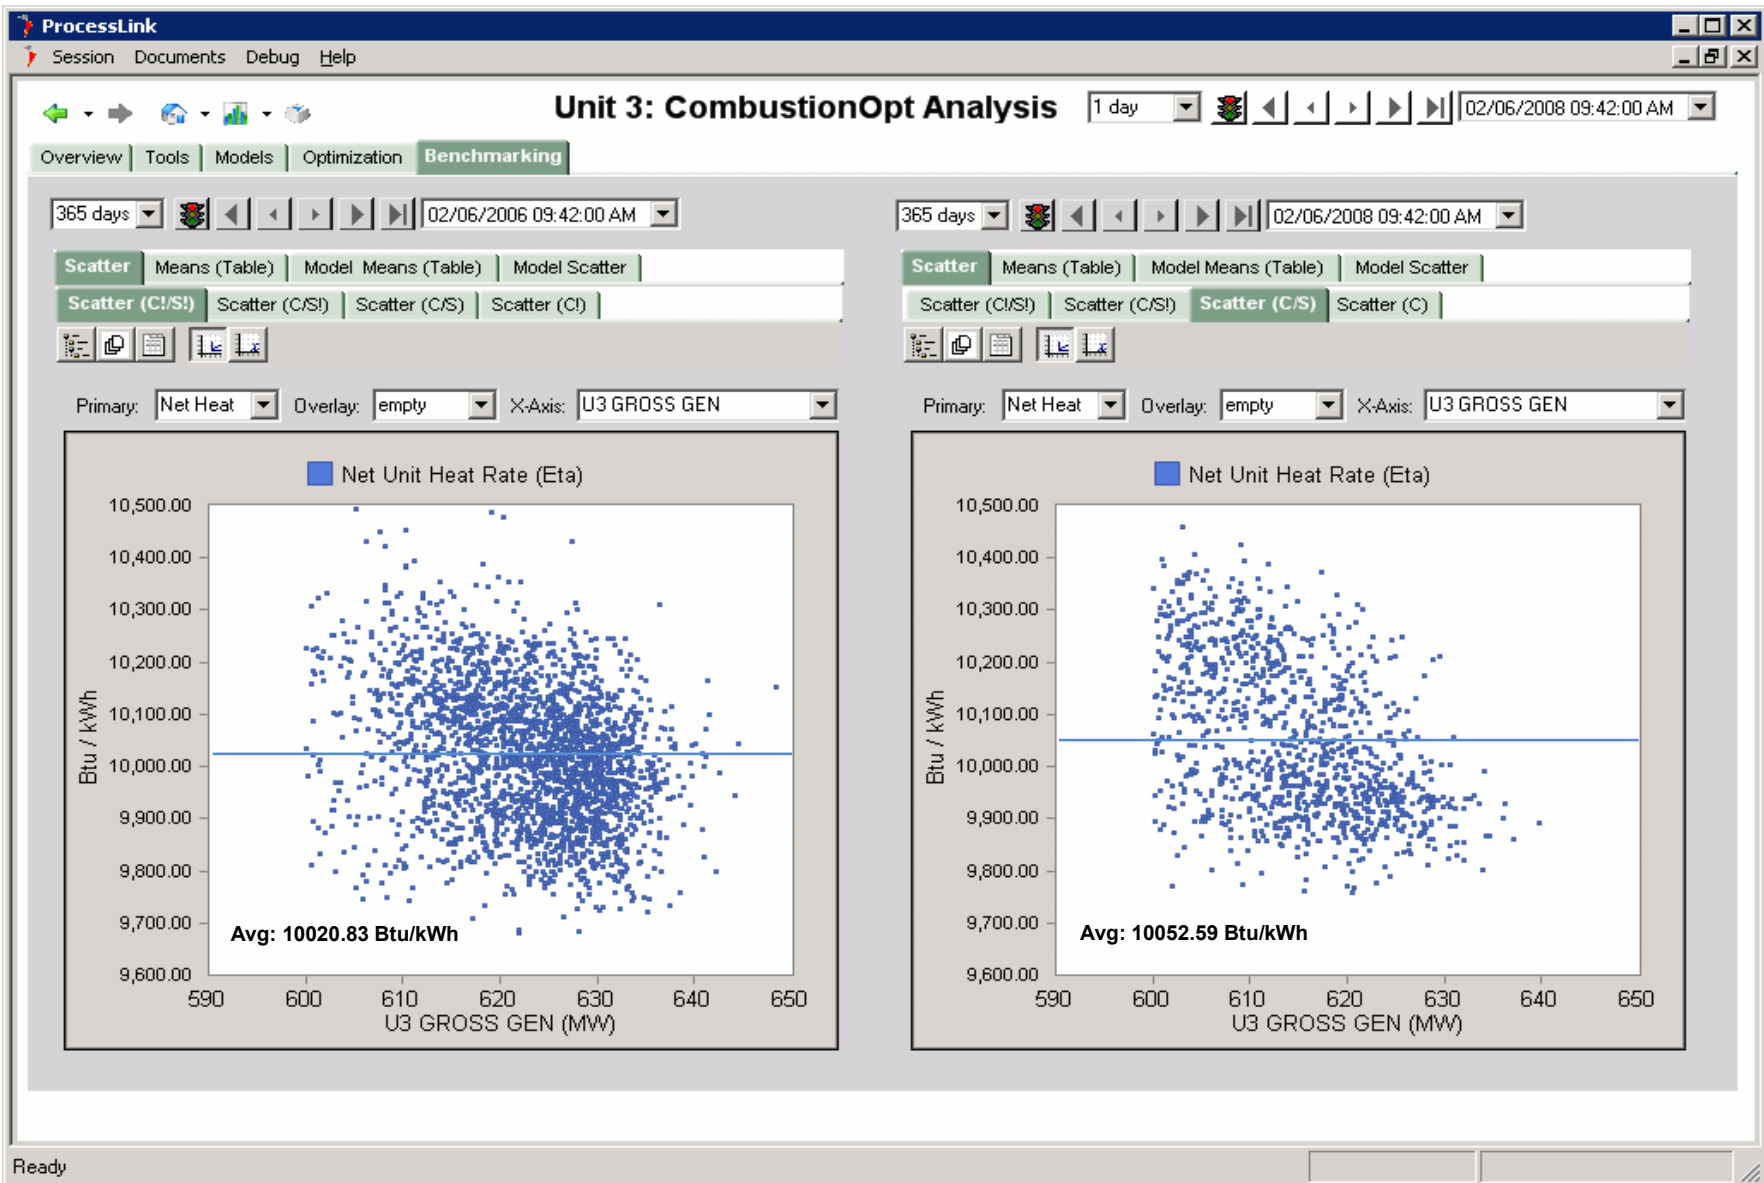


Figure 47 Unit 3 Measured Heat Rate vs. MW

4.4.1.1.4 NOx vs. NH3

Figure 48 - Figure 49 show the relationship between NOx and NH3 at Units 1 and 2. This comparison shows how things are different in two dimensions of interest but where one dimension is not the intuitively handy MW.

These views set the stage for the section to follow where fitting techniques are used to try and answer the questions they pose: NOx, NH3, MW, and HR are all varying together and are influenced by each other and at least two other major disturbances (Ambient temperature and humidity and Condenser Backpressure). Given that, how can we get a handle on how the total multivariable system is different between the two populations? Looking at two things at once (X and Y) is just too confusing when in reality there are multiple simultaneous X and Y relationships all changing at once.

Observations:

- Improvements in this plane are toward the origin (Lower NOx for Lower NH3). So a movement down and/or a movement to the left indicate improvement.
- The difference on Unit 1 is pronounced, as expected from the plots already shown.
- The difference on Unit 2 is less pronounced but still significant.

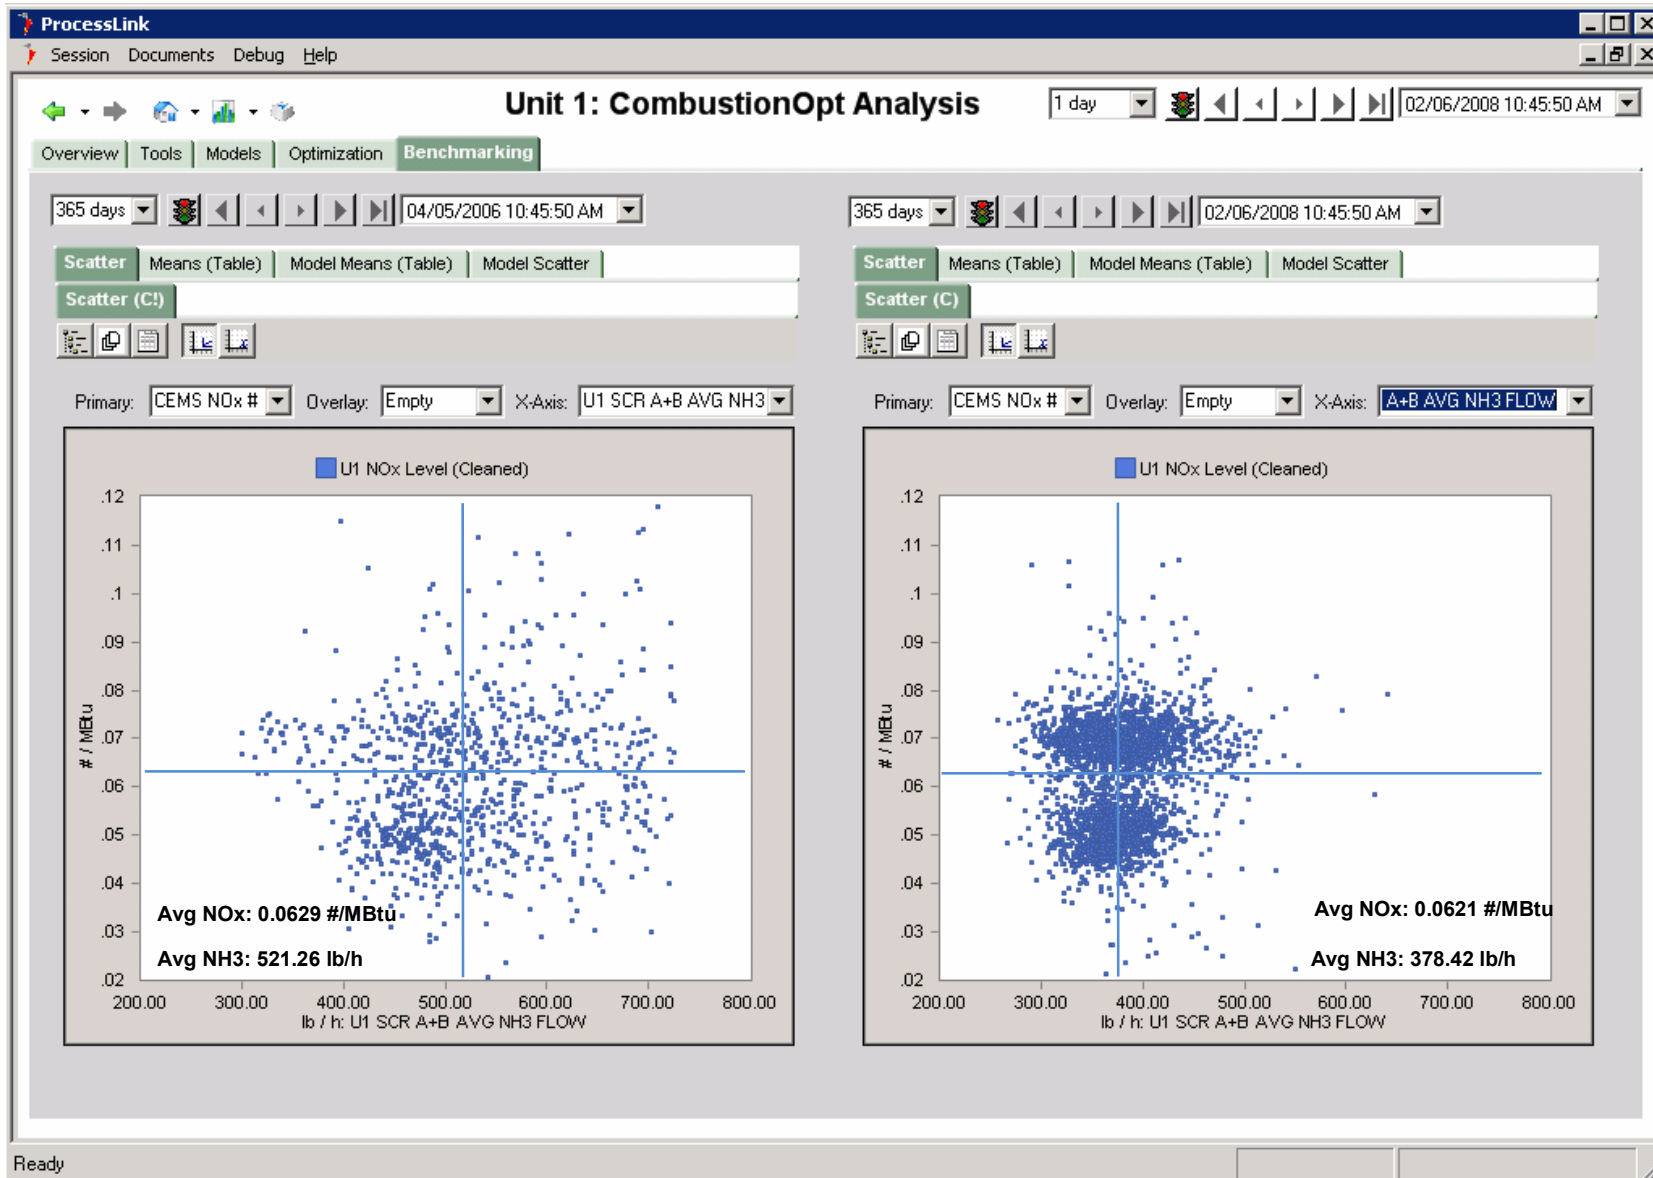


Figure 48 Unit 1 Measured NOx vs. NH3

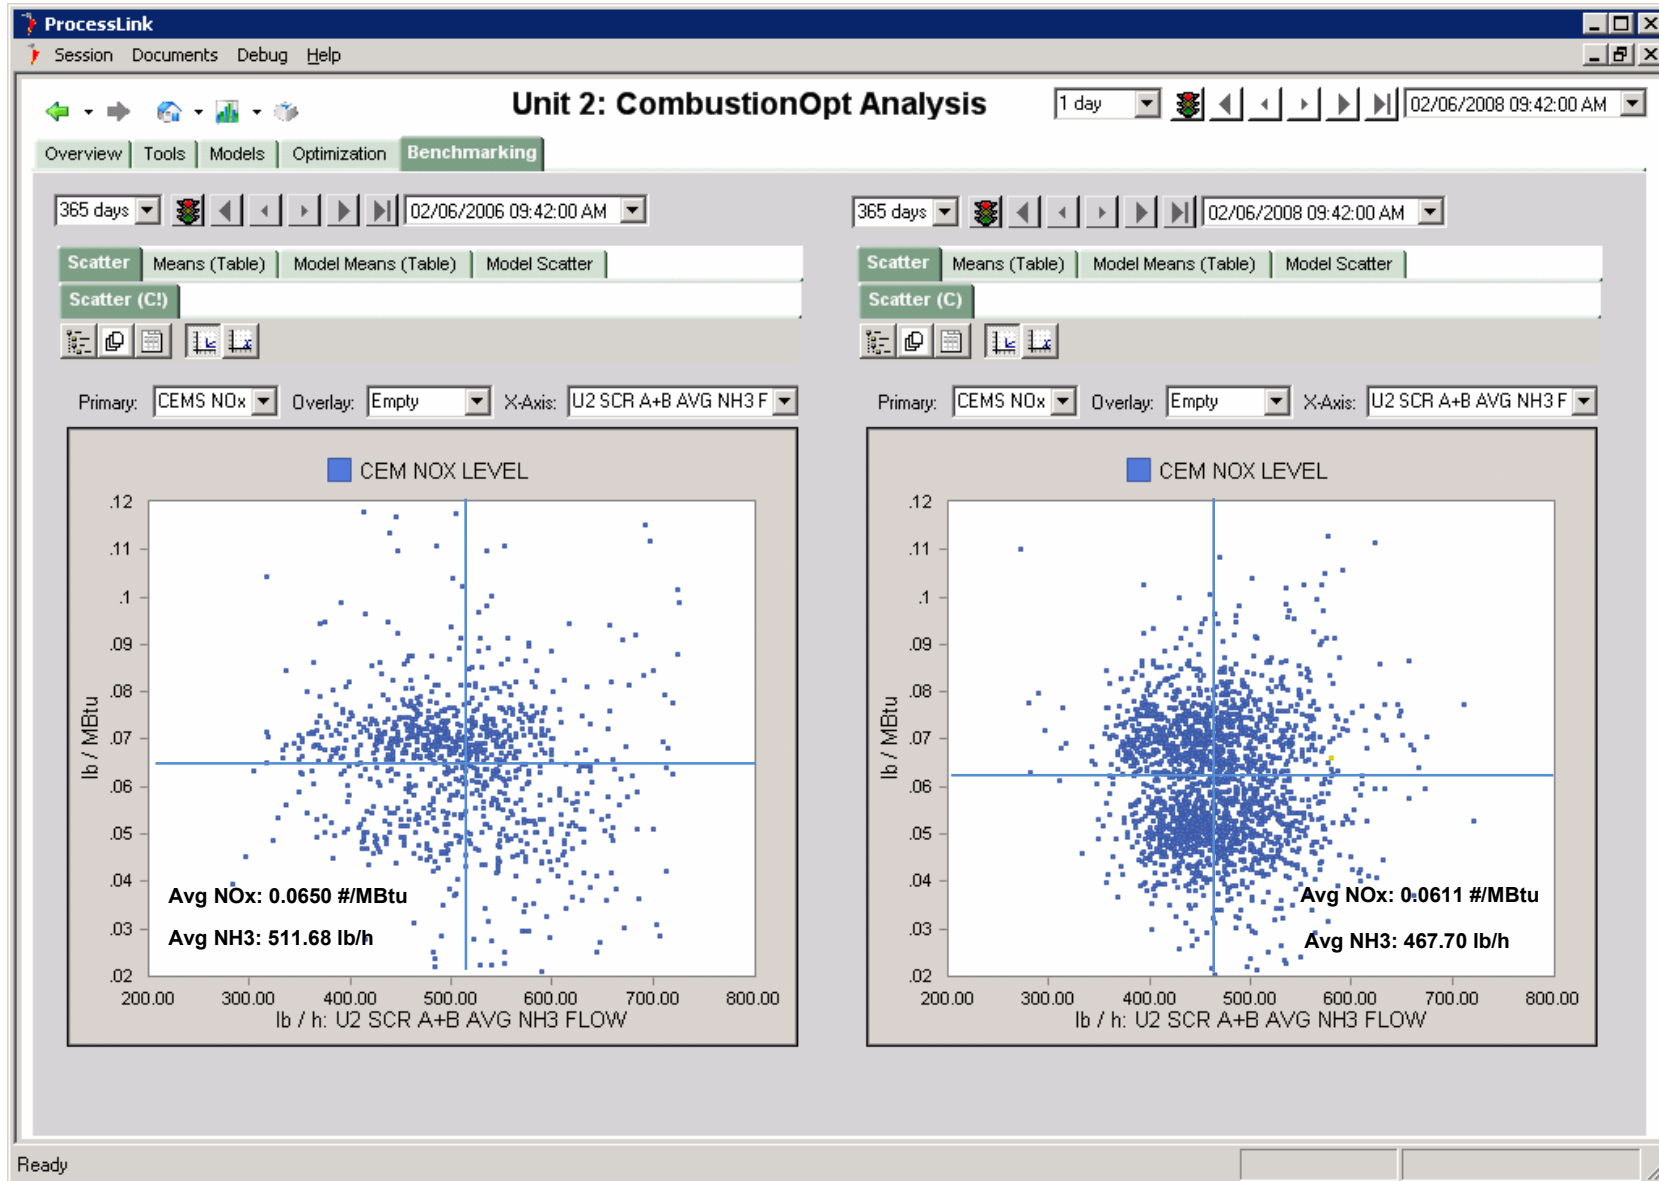


Figure 49 Unit 2 Measured NOx vs. NH3

4.4.1.1.5 NOx vs. HR

Figure 50 - Figure 51 show the relationship between NOx and Heat Rate for Units 1, 2 and 3.

This comparison shows how things are different in two dimensions of interest but where one dimension is not the intuitive MW.

Observations:

- Unit 1 shows the leftward movement of CEMS NOx vs. the Heat Rate coordinate, meaning lower Heat Rate was seen at the same CEMS NOx level in the experiment vs. the control.
- Unit 2 shows more pronounced improvement when viewed in this plane. Again the movement is largely to the left, which corresponds to the reduced Heat Rate seen in the previous plots.
- Unit 3 shows movement in both dimensions. The HR at higher NOx values is lower, while the Heat Rate at lower NOx values (which are missing from the control) is somewhat higher than the overall average of the control.

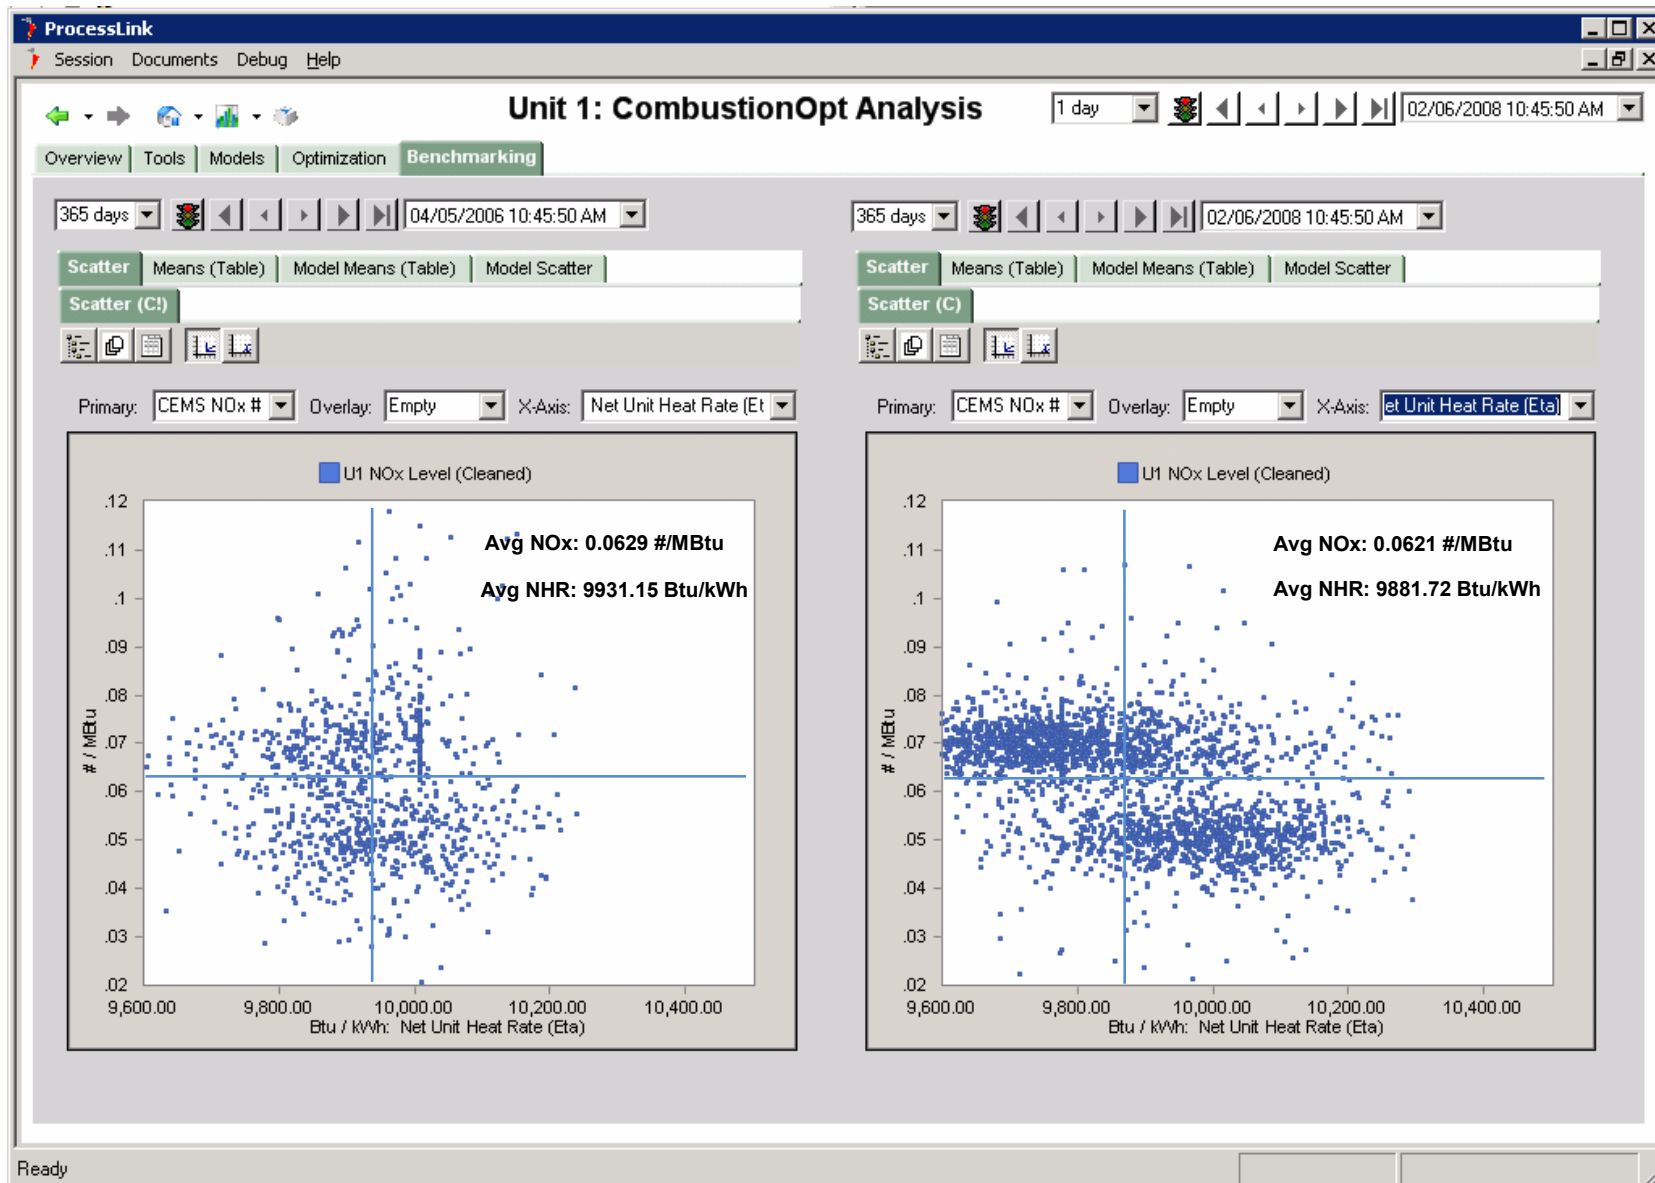


Figure 50 Unit 1 Measured NOx vs. Heat Rate

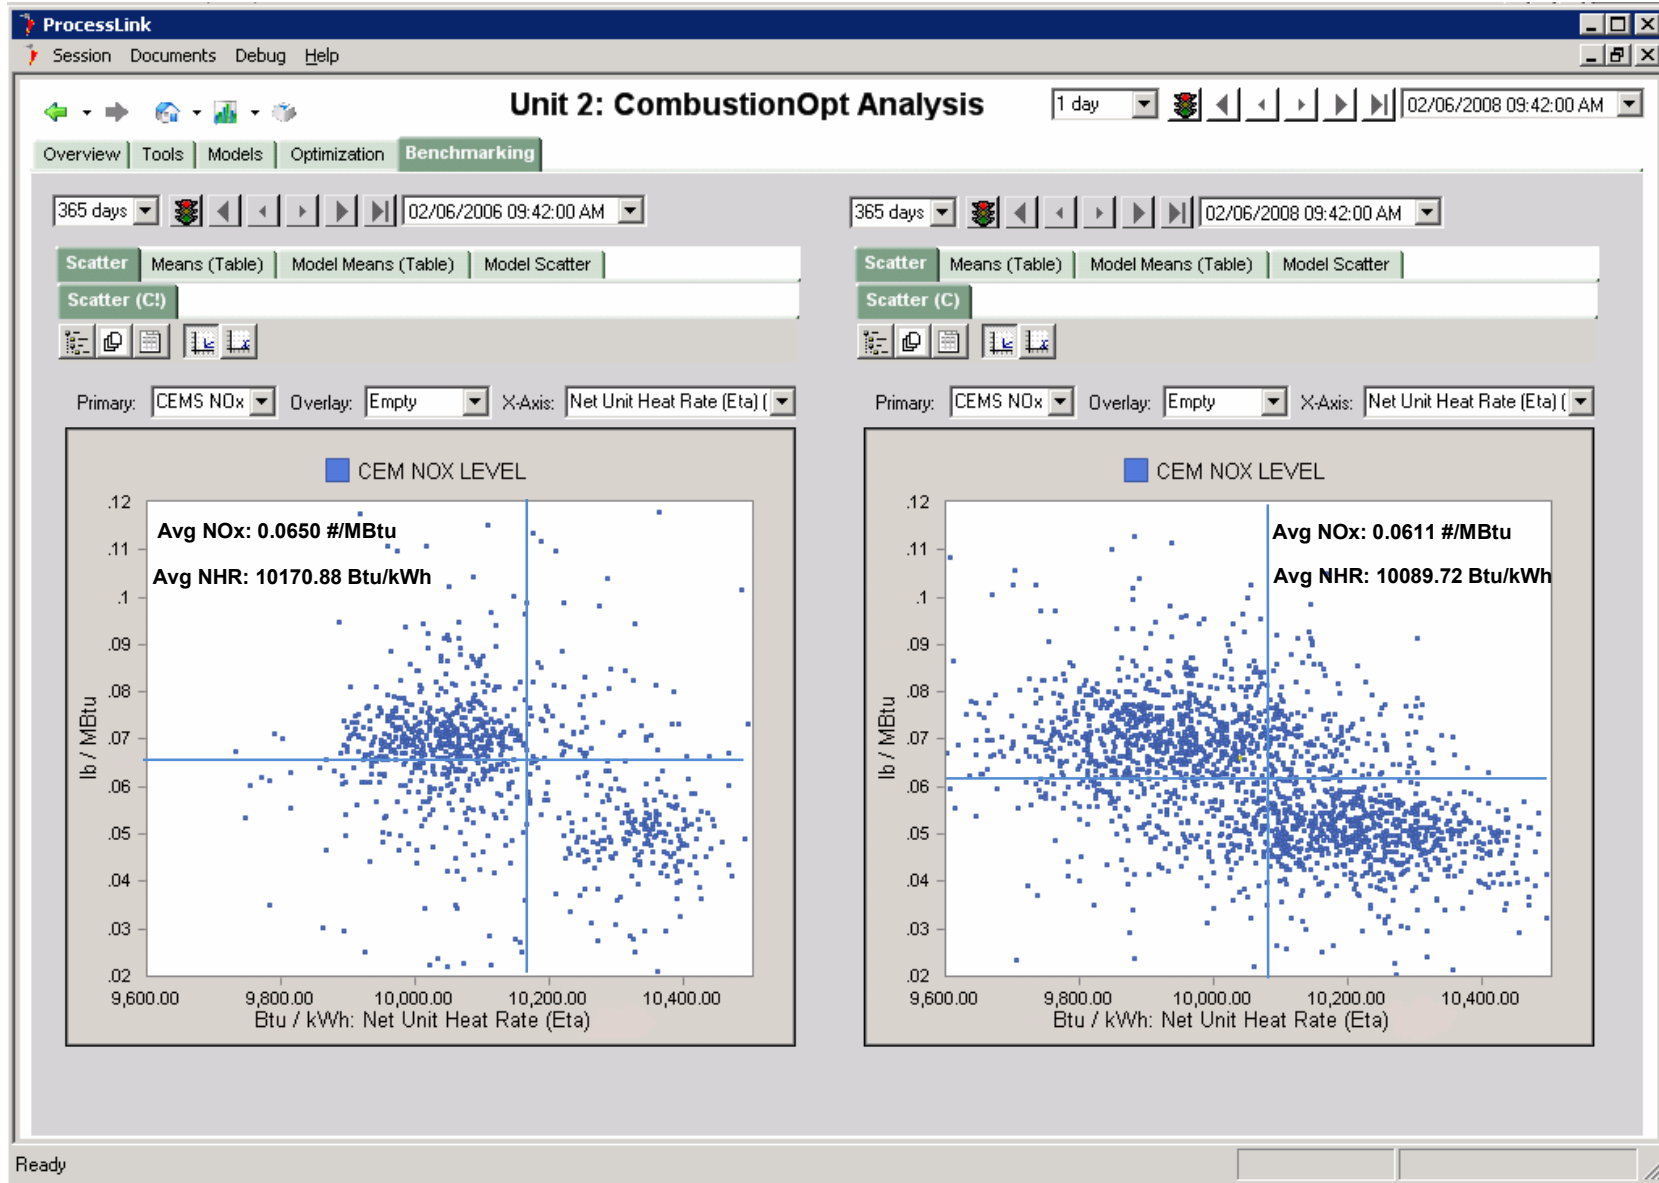


Figure 51 Unit 2 Measured NOx vs. Heat Rate

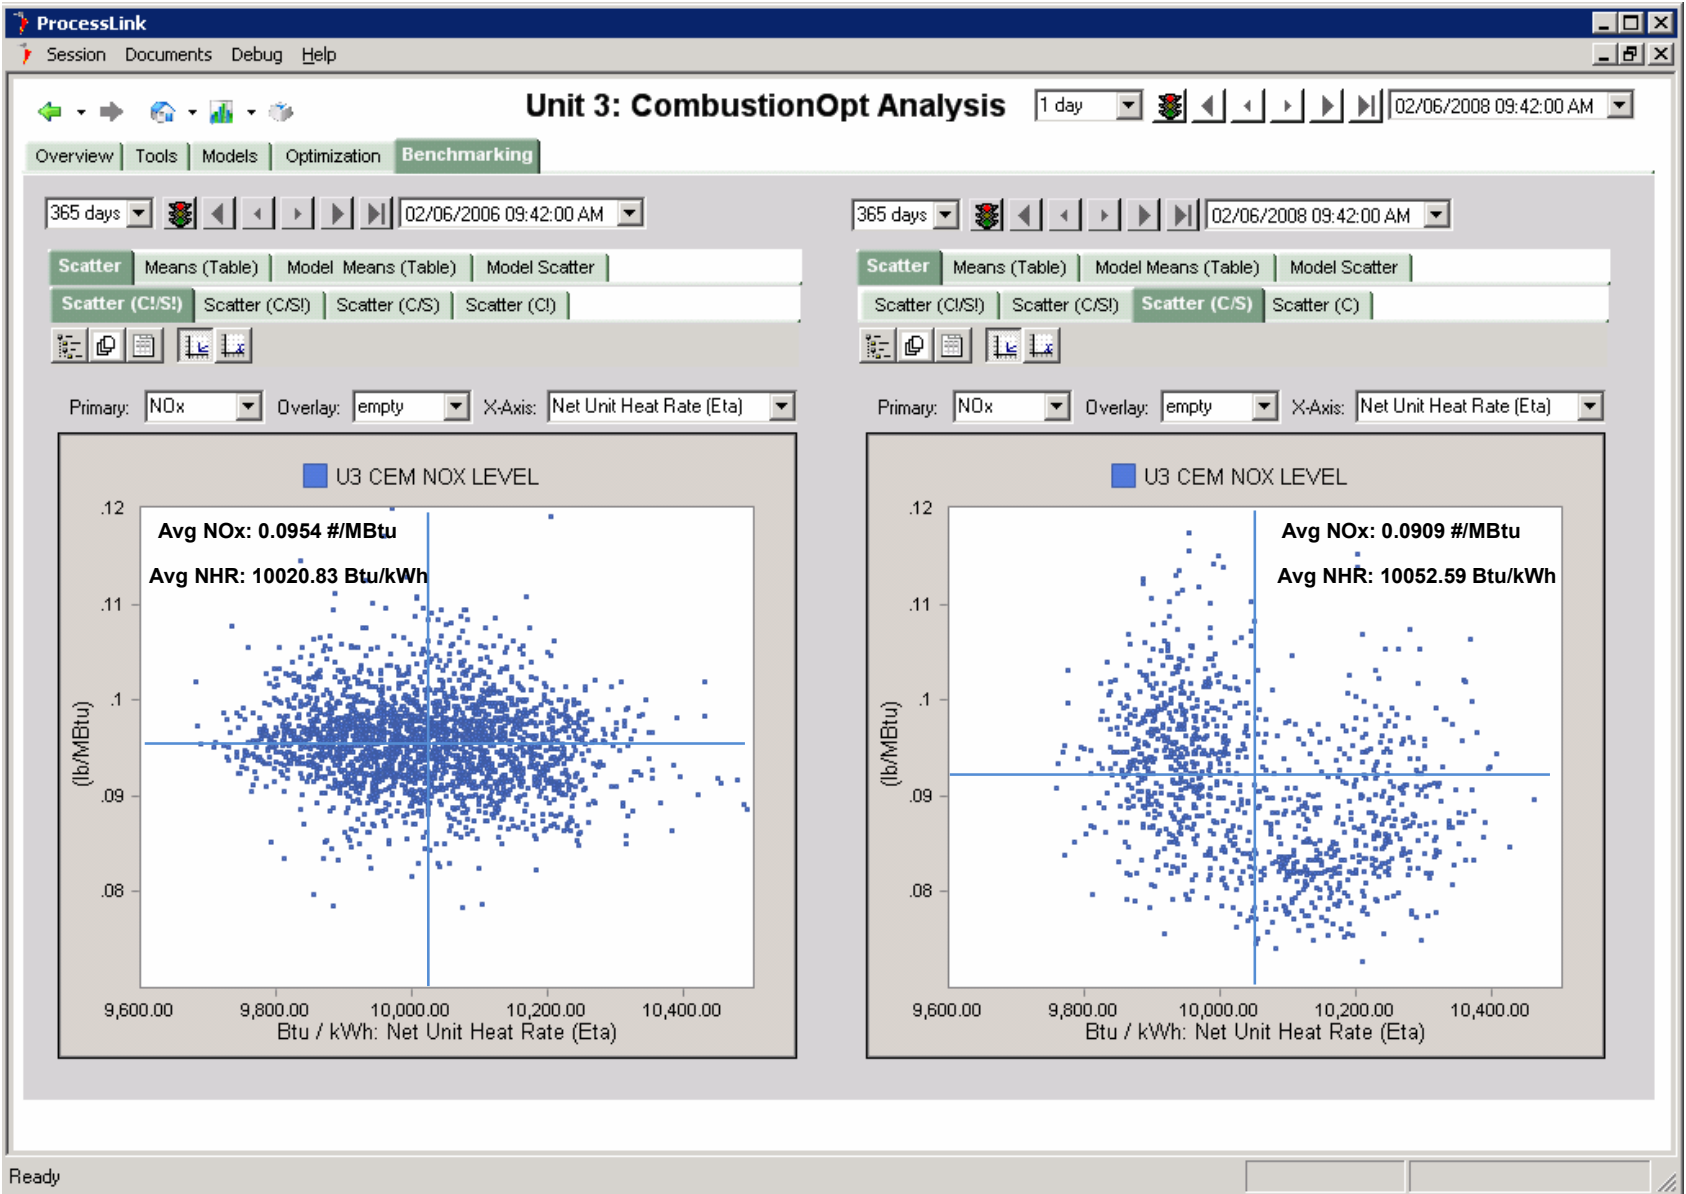


Figure 52 Unit 3 Measured NOx vs. Heat Rate

4.4.1.1.6 NH3 vs. HR

Figure 53 - Figure 54 show the relationship between NH3 and Heat Rate for Units 1 and 2.

This comparison shows how things are different in two dimensions of interest but where one dimension is not the intuitive MW.

Observations:

- For Unit 1 the change seen in this plane is both down and left, indicating that lower NH3 flow was seen at lower Heat Rates in the experiment.
- For Unit 2 the change is also in both dimensions but is less pronounced than Unit 1.

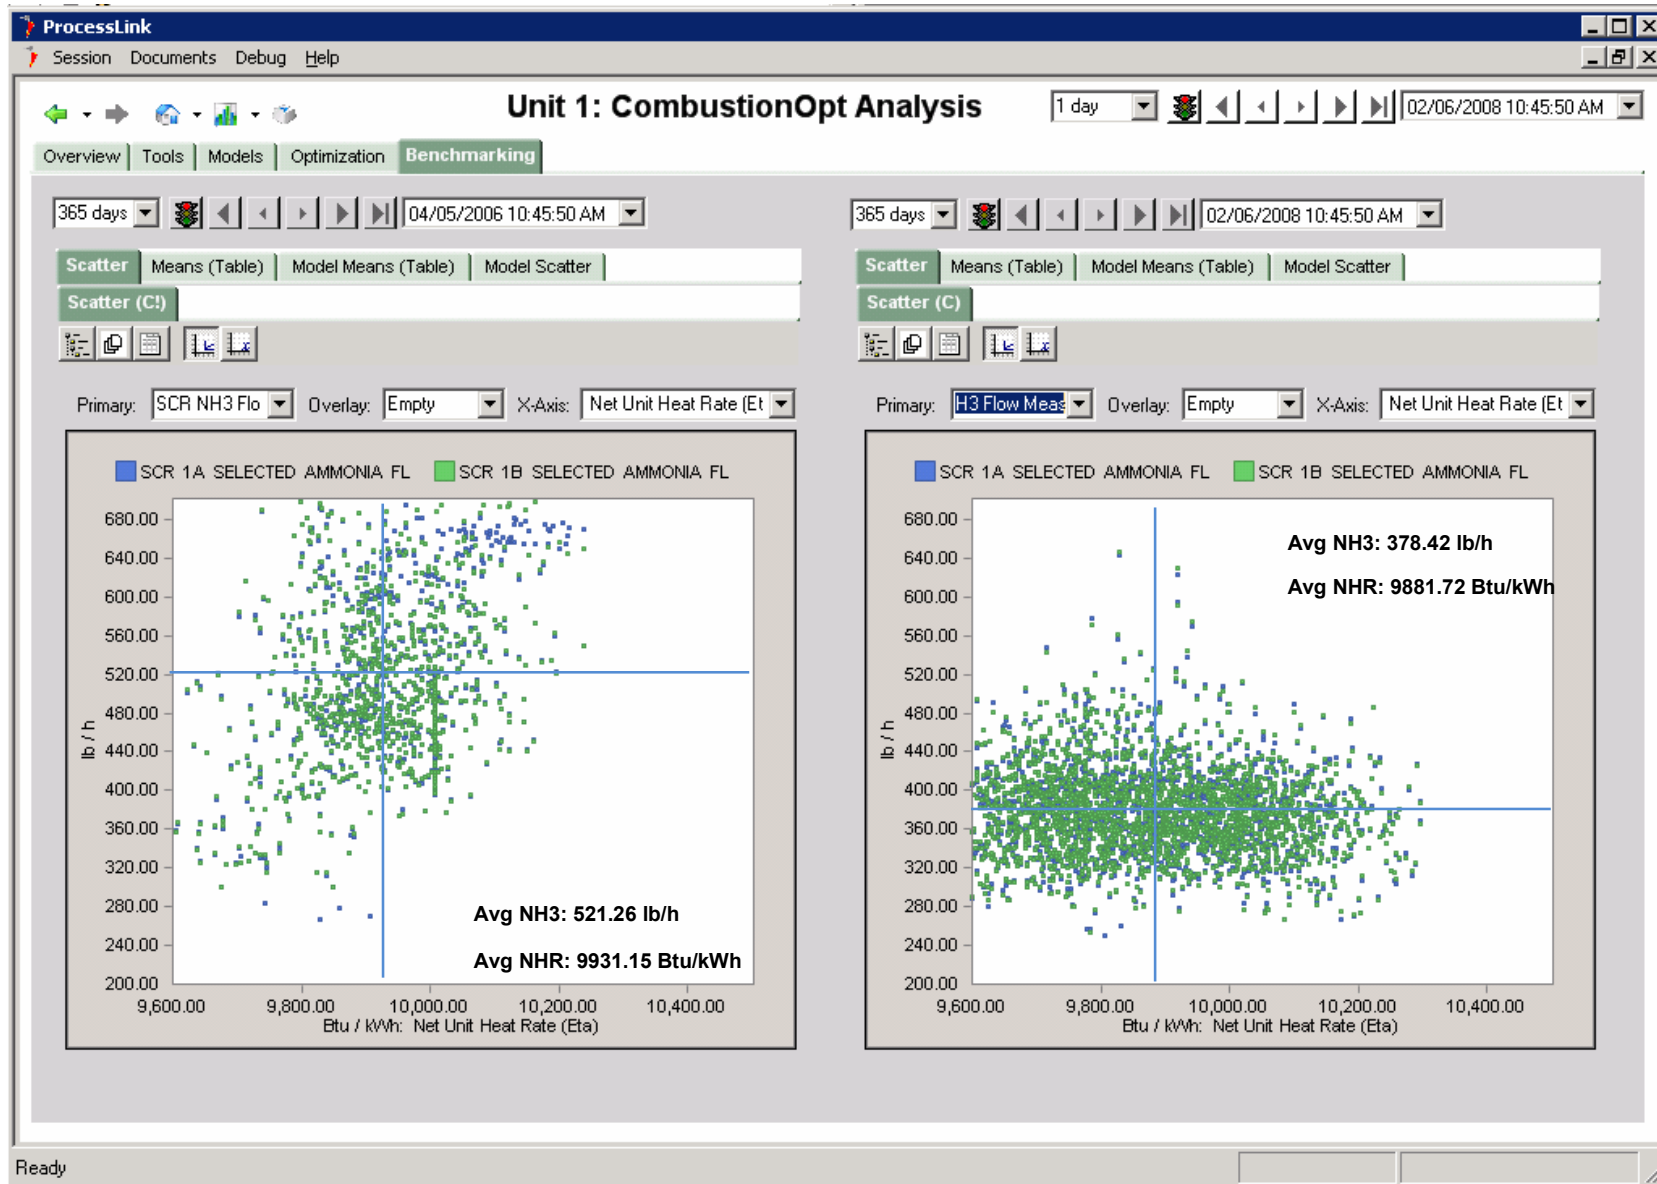


Figure 53 Unit 1 Measured NH3 vs. Heat Rate

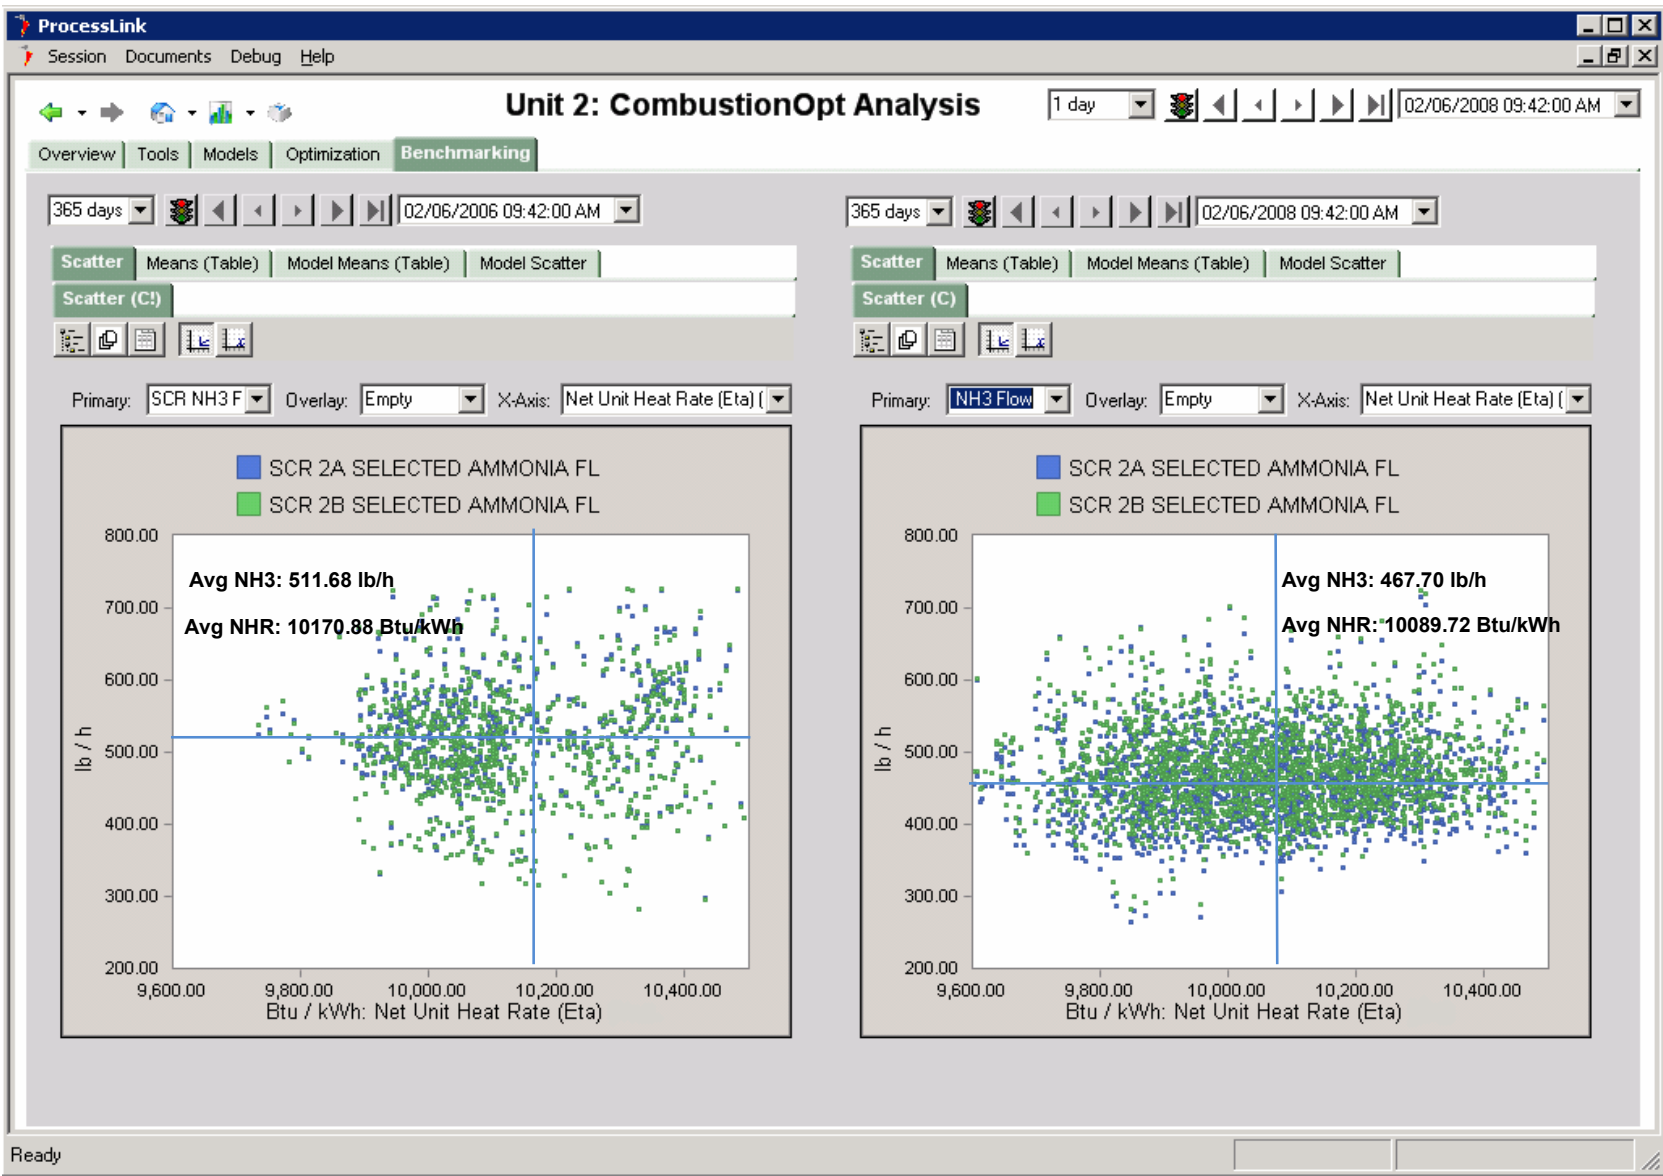


Figure 54 Unit 2 Measured NH3 vs. Heat Rate

4.4.2 Modeled Response to Disturbances

The second data analysis methodology used in this study is Modeled Response to Disturbances. This approach compares the response of models of the variables of interest (NO_x, HR, NH₃ usage), as functions of their disturbances..

- Advantages
 - Once the models are built they can provide an accurate picture of the effect of the experimental factor.
- Disadvantages
 - Non-linear multi-variable regression is more difficult to implement than the Comparisons of Populations method

The difficulty with looking at a collection of binary measured relationships is the fact the relationships between the variables we are looking at are not simply binary. In fact we know that NO_x, NH₃ and HR respond to each other in a coupled way and to factors other than just the experimental one. For this reason it is hard to get a clear picture of just what the multi-variable system's holistic response to the factor (optimization) is. The major cofactors or disturbance suspects for NO_x, Heat Rate and NH₃ include Ambient Temperature/Humidity and Condenser Cleanliness. We know this from theory and experience. Ambient temperature and Condenser cleanliness are especially strong disturbances for Heat Rate (as the analysis will indicate). For NO_x, NH₃ is of course a strong causal driver.

By using a fitting (or regression) approach we can model the response of the variables of interest to their disturbances and then use those models in a single disturbance scenario to isolate the effect of the factor of interest.

Because we have a powerful neural network training engine in ProcessLink (the same one used for CombustionOpt, SCR-Opt, SootOpt and MaintenanceOpt) and neural nets are excellent general multivariable regression machines (in fact that is precisely what they do), we overcome the challenges of using a modeling approach, which would typically involve the painstaking construction of a matrix of single variable regression models (and which would in the end be less accurate).

The following section presents the output of scenarios created using neural network (multivariable regression) models of NO_x, NH₃, and HR as functions of each other (for NO_x and NH₃) and their gross disturbances, namely MW, Wet Bulb Temperature and Condenser Backpressure.

The modeled Scenarios were developed in the following way:

1. *Selection of Control Data*: Selected data for NO_x, NH₃, and Heat Rate when Optimization was absent, meaning all data:
 - a. from 2005 (from roughly 4/05 to 4/06)
 - b. where MW > 600,
 - c. The SCR's were in service,
 - d. CombustionOpt Master Enable is OFF
 - e. And for Unit 3 SootOpt Master Enable is OFF
2. *Selection of Experimental Data*: Selected data for NO_x, NH₃ and Heat Rate when Optimization was present meaning all data
 - a. from 2007 (roughly 1/9/07 to 1/9/08)
 - b. where MW > 600, the SCR's were in service,
 - c. CombustionOpt Master Enable is ON.
 - d. And For Unit 3 SootOpt Master Enable is ON

3. Using the same neural modeling technology used to build models used for optimization in CombustionOpt and SootOpt and for equipment monitoring in MaintenanceOpt, we modeled the NO_x, NH₃ and HR values of the control and experimental population as functions of each other (where relevant) and their disturbances or cofactors (i.e. MW, Wet Bulb Temp, Condenser Back Pressure)
4. Used these models to compare the difference in the NO_x, NH₃ and HR performance predicted as a function of the real MW, Ambient Temp/Humidity, and Condenser Backpressure conditions found in 2007.
5. Based on the assumptions that the effects of the disturbances identified should be relatively constant between 2005/2006 and 2007/2008 periods. The difference in the predictions of the two models should represent the isolated effect of the factor (namely whether Optimization was in effect or not).

The categorized scatter plots shown in the figures in Section 4.4.2.1 (Modeled Scenarios) show (for Heat Rate as an example):

- **[Left]** The output of a modeled scenario where

$$y_{Control}(t) = HR(t) = g_{Control}[\alpha, \theta]$$

θ = weight vector adapted from control data,

α = input vector = MW, Wetbulb, CondBP for 2007.

- **[Right]** The output of a modeled scenario where

$$y_{Experiment}(t) = HR(t) = g_{Experiment}[\alpha, \theta]$$

θ = weight vector adapted from experimental data,

α = input vector = MW, Wetbulb, CondBP for 2007.

4.4.2.1 Modeled Scenarios

4.4.2.1.1 Modeled NO_x as f(Disturbances) vs. MW

Figure 55 - Figure 57 show the modeled relationship between CEMS NO_x and Megawatts for Units 1, 2 and 3.

Observations:

- For Unit 1 the NO_x distribution is more tightly controlled
- The density of samples in the lower NO_x mode is visibly greater in the experiment for both Unit 1 and 2.
- For Unit 3 the NO_x mode in the experiment is disrupted and shifted in the direction of lower values at the same MW (except for maximum rating) relative to the control.

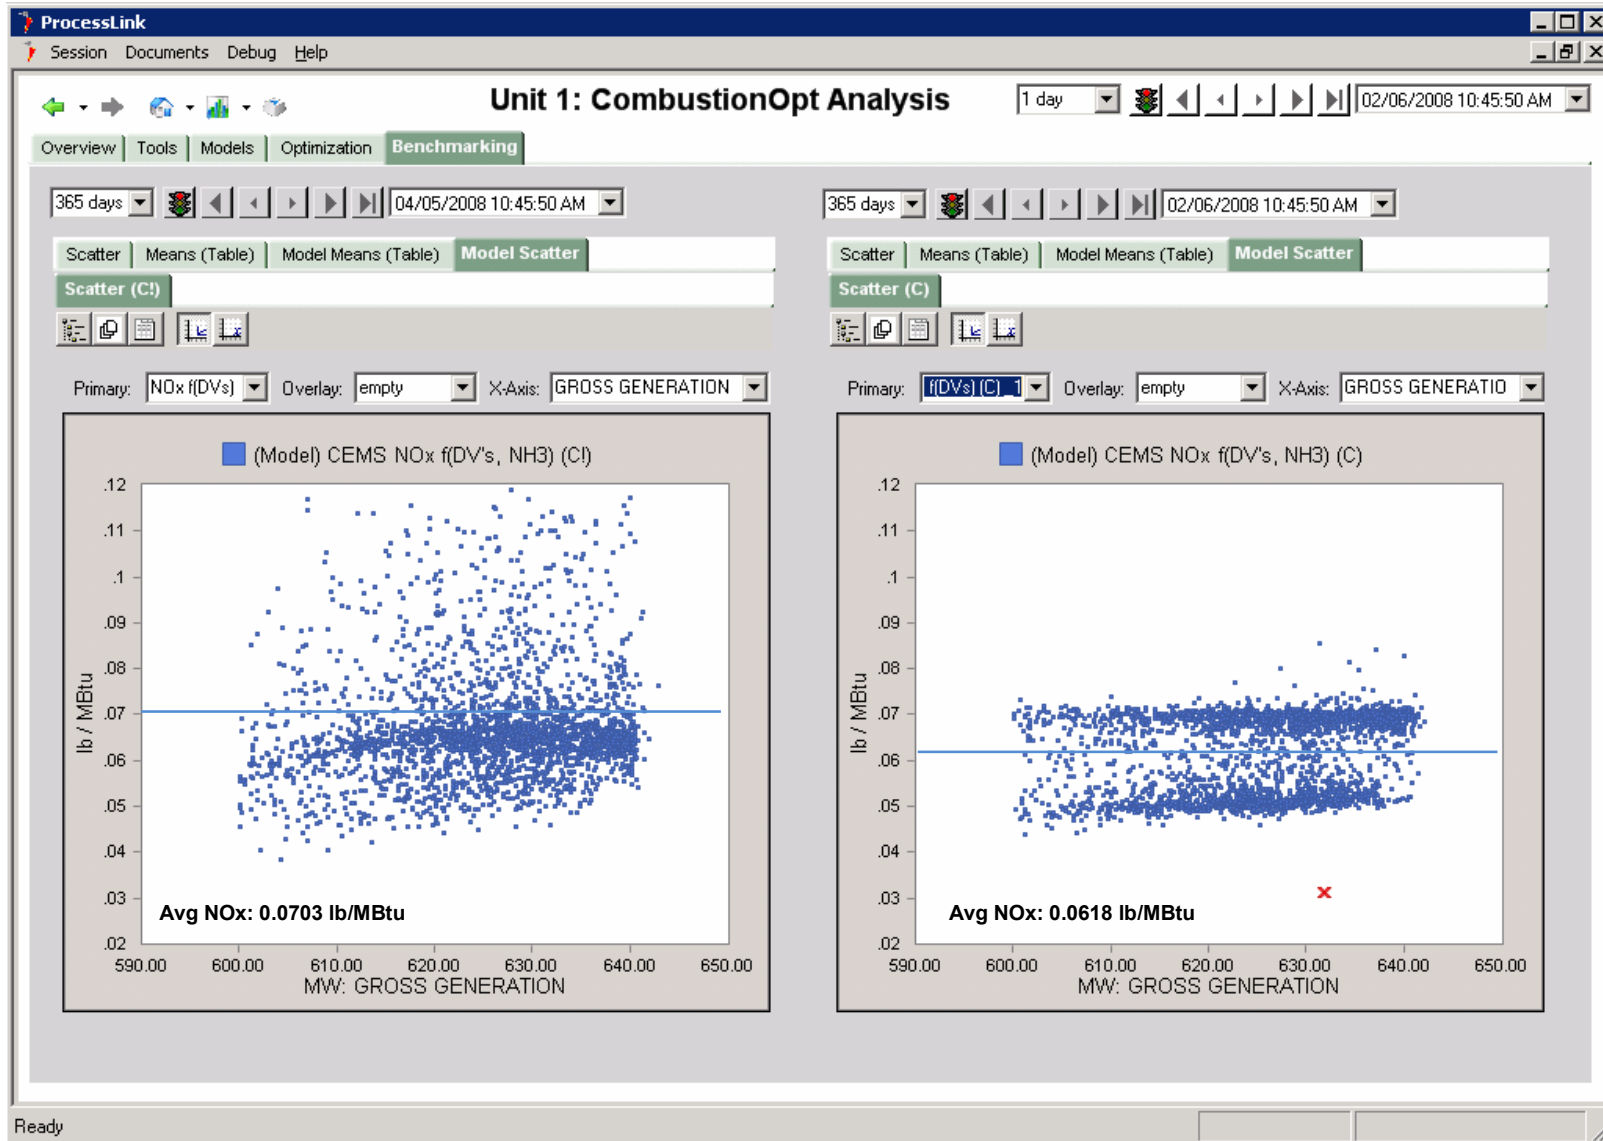


Figure 55 Unit 1 Modeled NOx as $f(\text{Disturbances})$ vs. MW

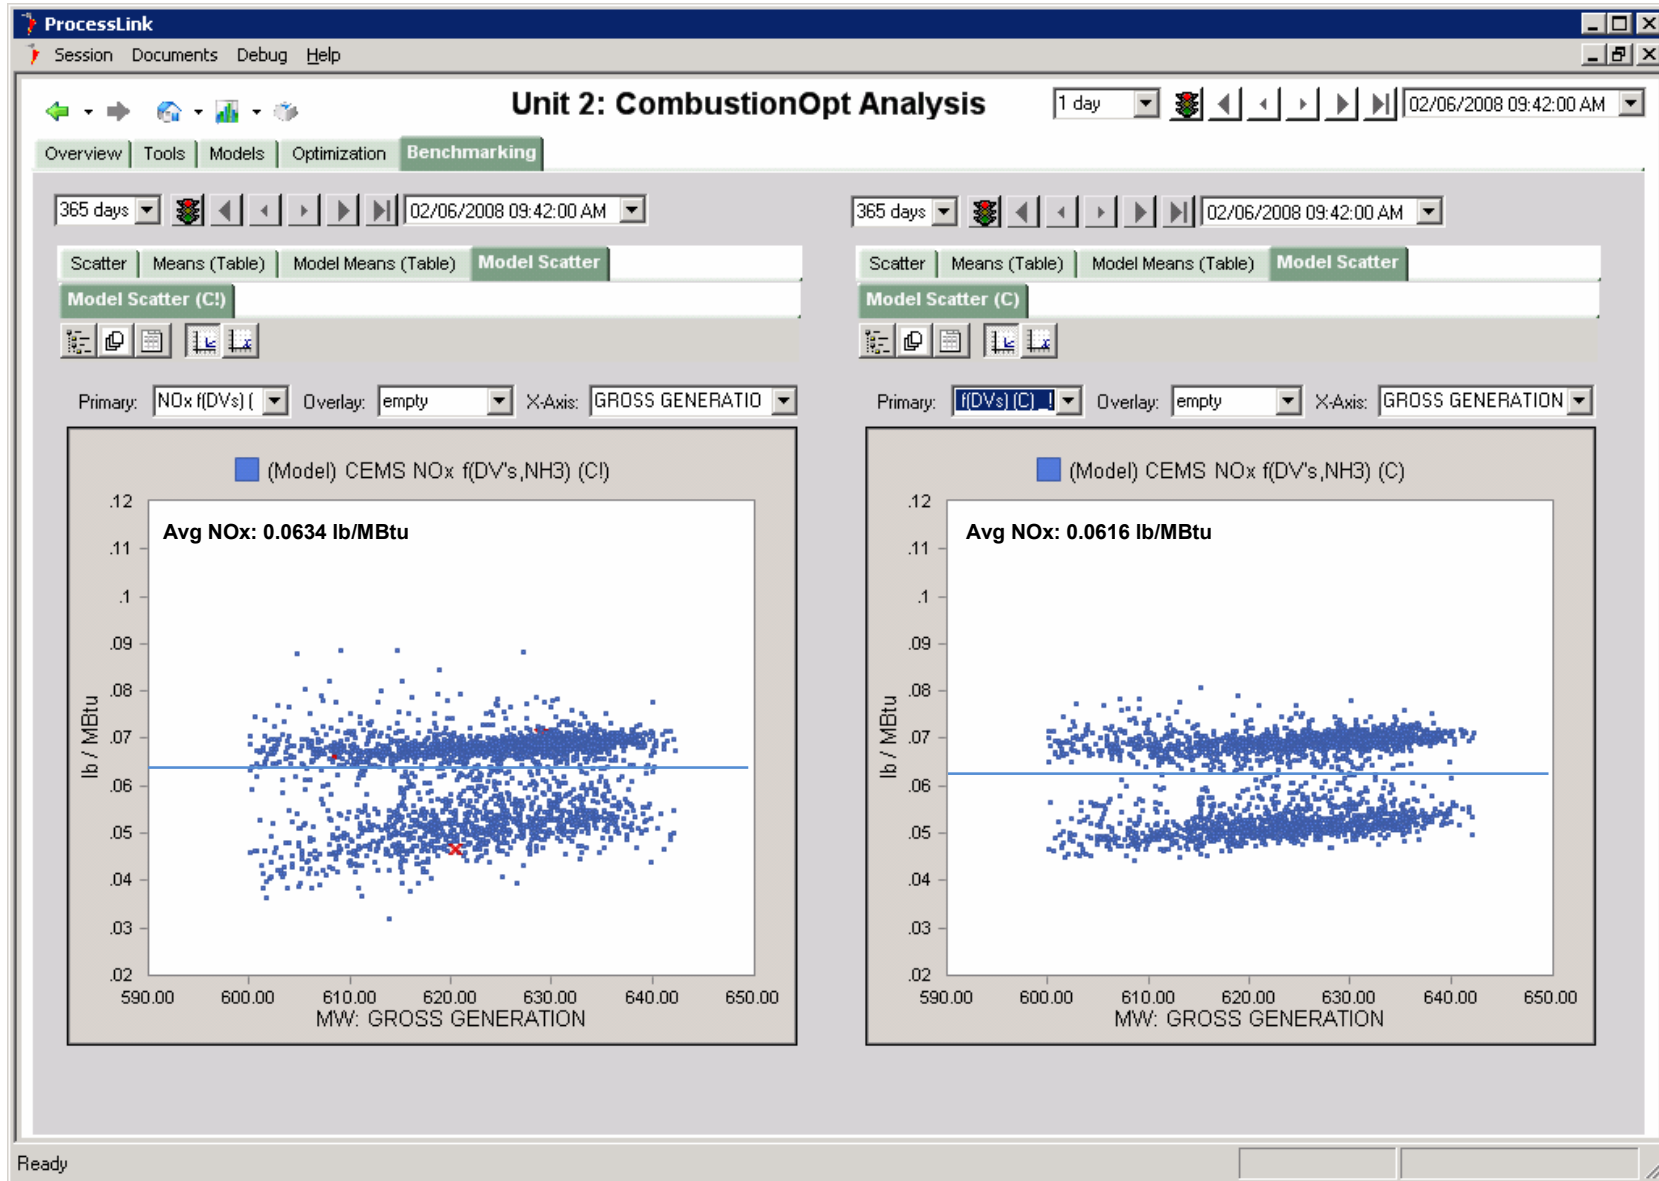


Figure 56 Unit 2 Modeled NOx as $f(\text{Disturbances})$ vs. MW

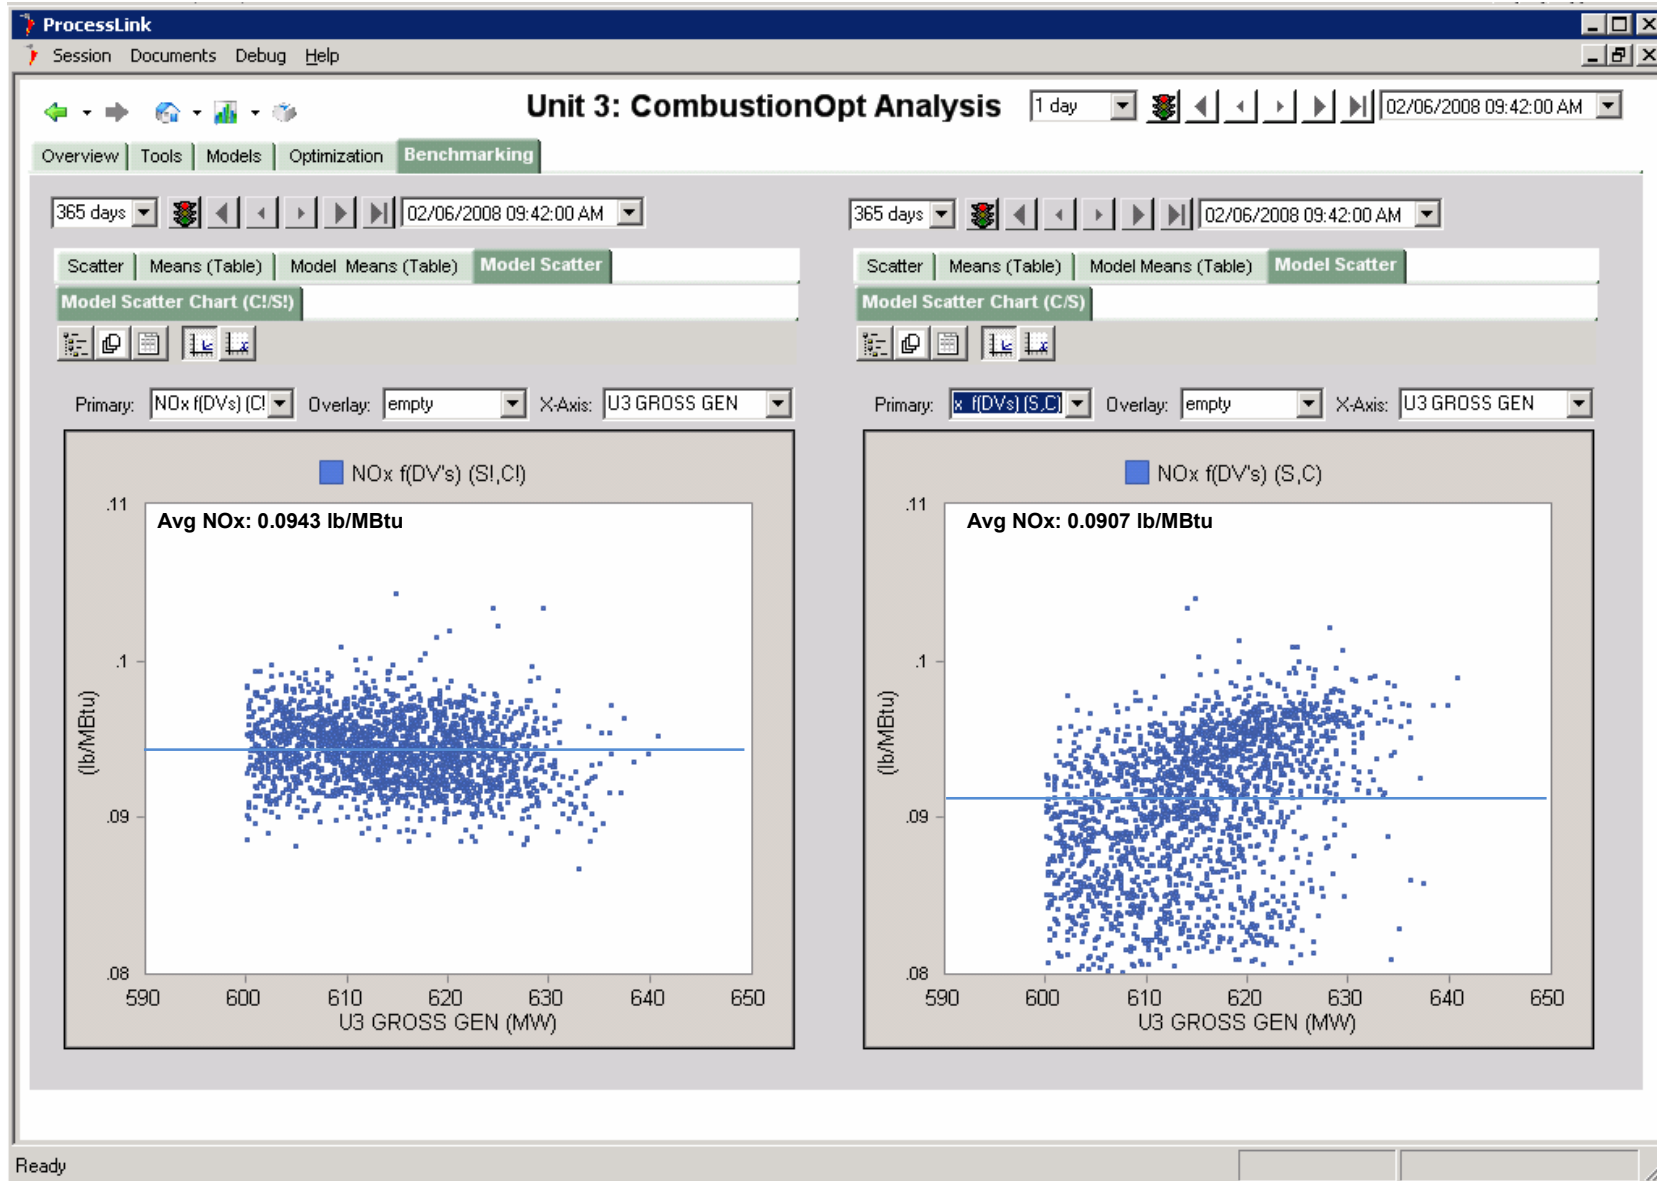


Figure 57 Unit 3 Modeled NOx as $f(\text{Disturbances})$ vs. MW

4.4.2.1.2 Modeled HR as f(Disturbances) vs. MW

Figure 58 - Figure 60 show the modeled relationship between HR and Megawatts for Units 1, 2 and 3.

Observations:

- The delta in Heat Rate for Units 1 and 2 is fairly pronounced and also more bi-modal.
- For Unit 3 there is a slight increase in the minimum Heat Rate, what might be a slight increase in the mean of the lower mode, but the relative absence of the upper mode.

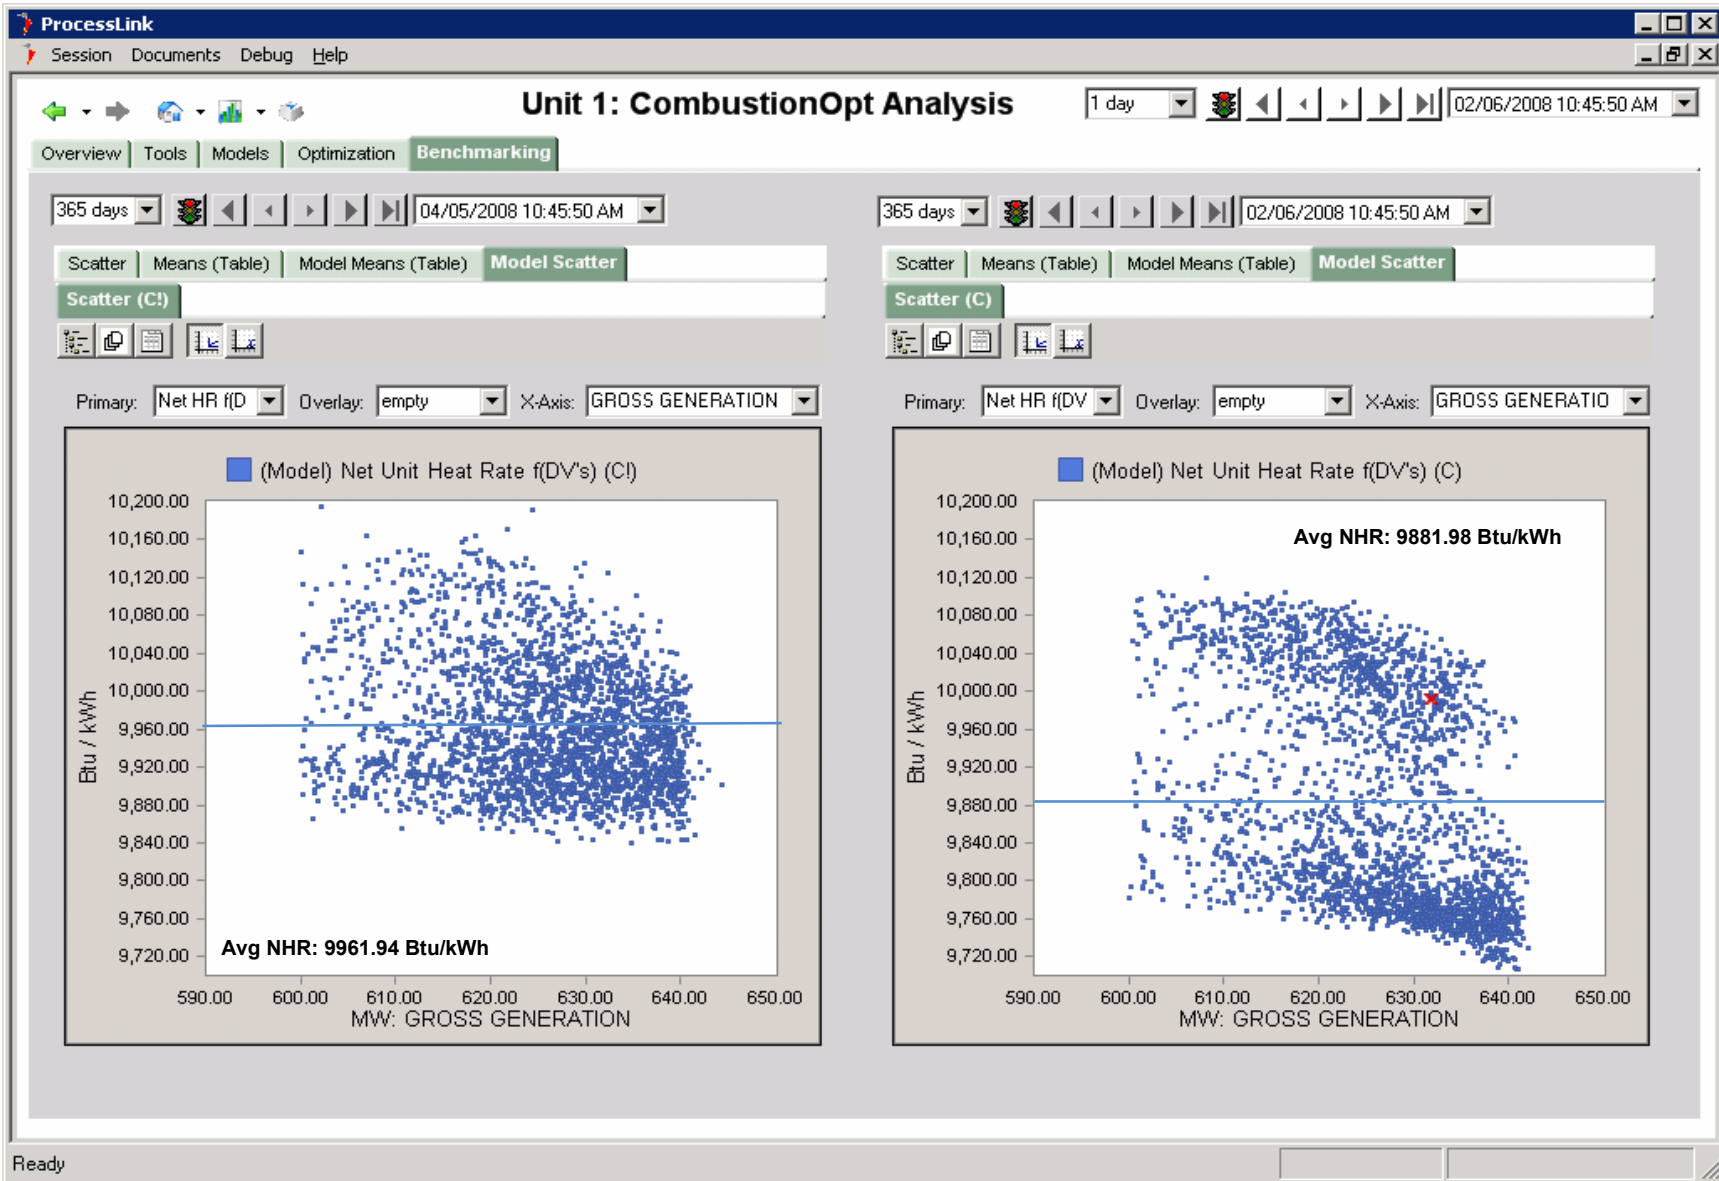


Figure 58 Unit 1 Modeled Heat Rate as $f(\text{Disturbances})$ vs. MW

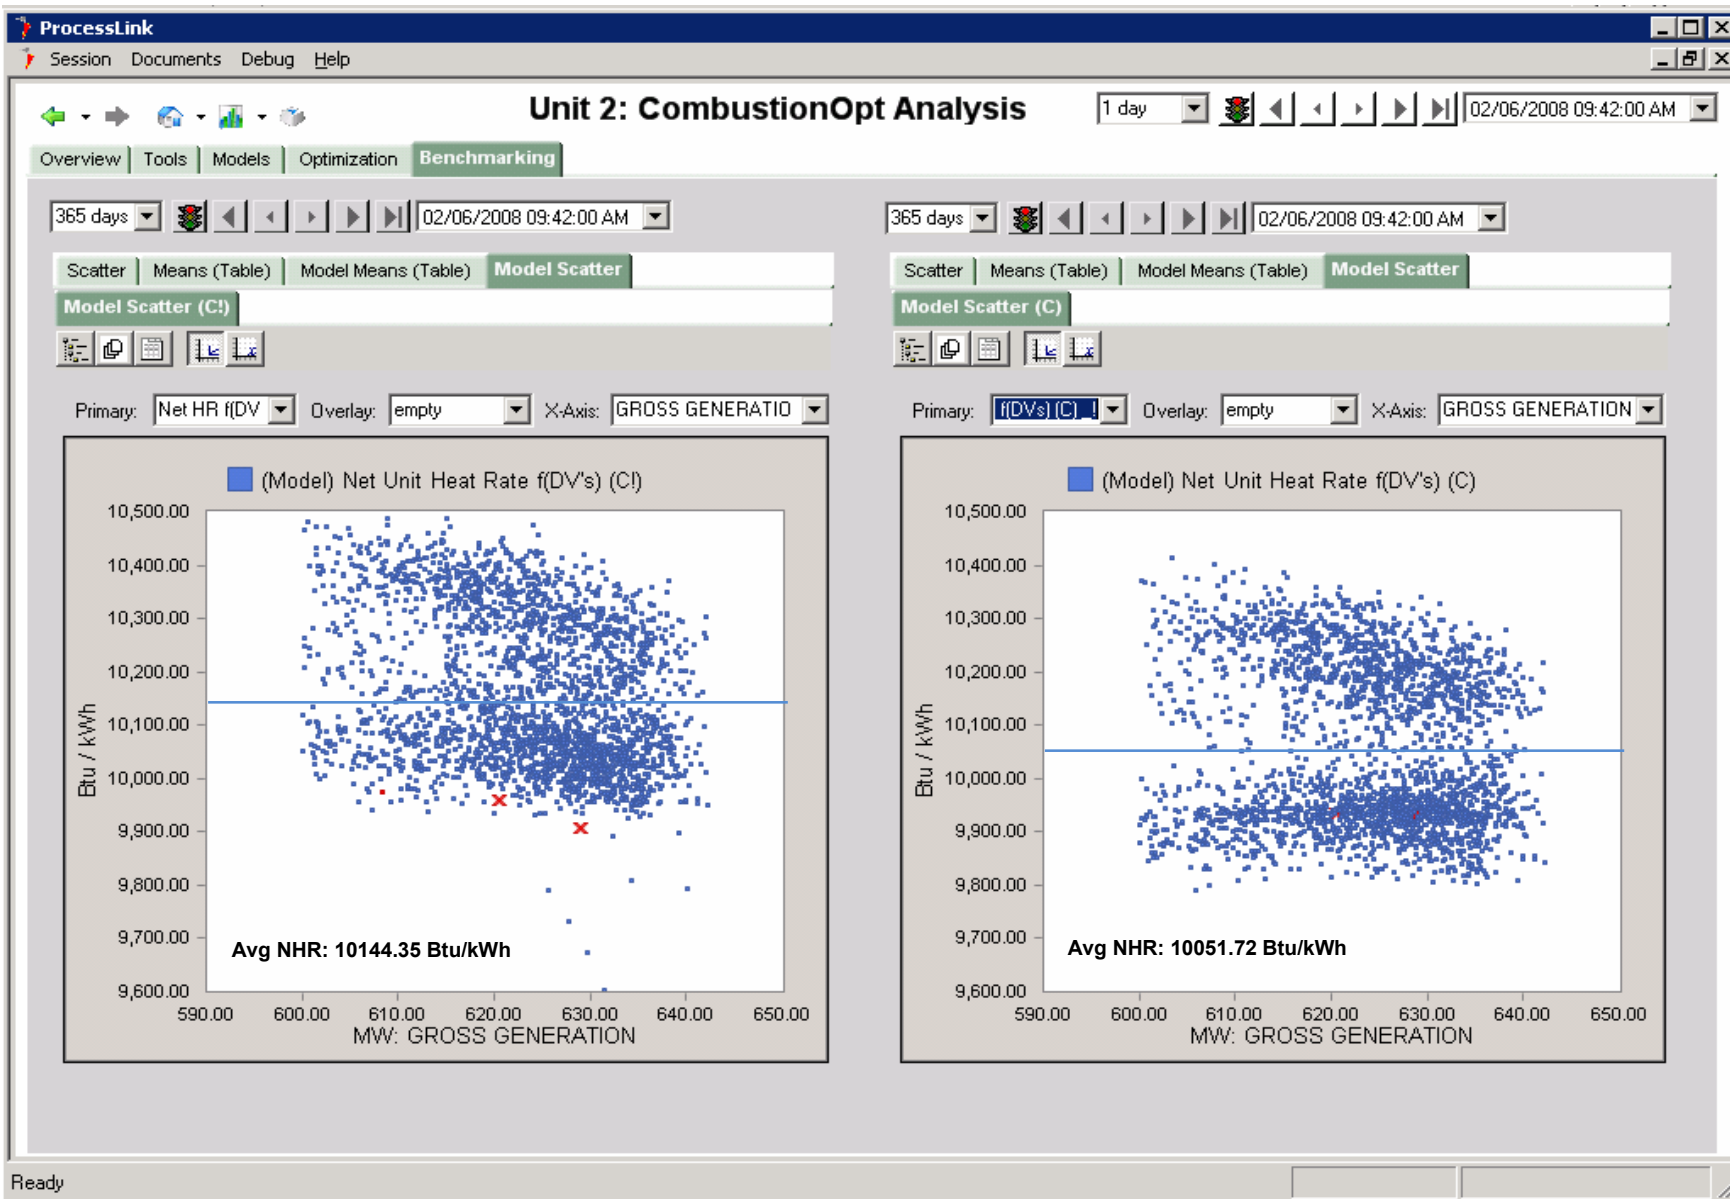


Figure 59 Unit 2 Modeled Heat Rate as $f(\text{Disturbances})$ vs. MW

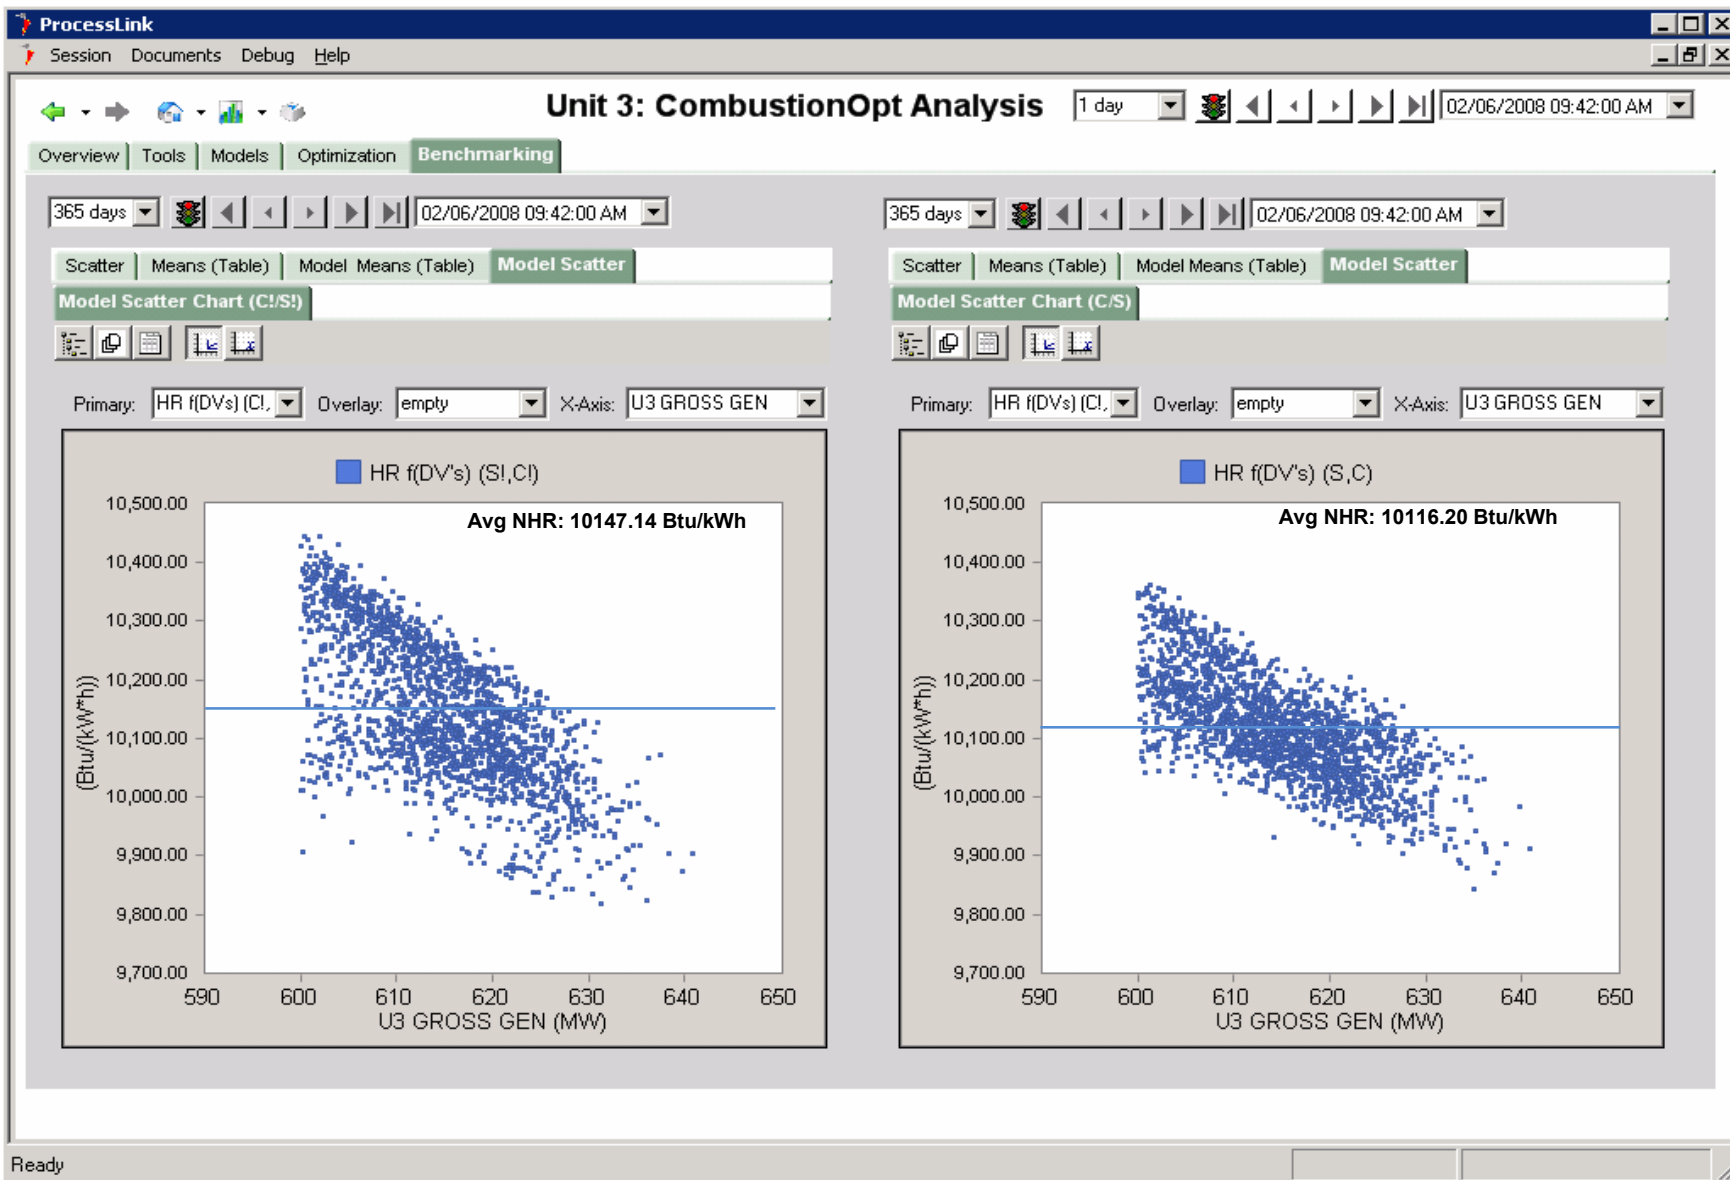


Figure 60 Unit 3 Modeled Heat Rate as $f(\text{Disturbances})$ vs. MW

4.4.2.1.3 Modeled NH3 as f(Disturbances) vs. MW

Figure 61 - Figure 62 show the modeled relationship between NH3 and Megawatts for Units 1 and 2.

Observations:

- For Unit 1 the difference is pronounced and positive.
- For Unit 2 the difference is hard to distinguish, though there are lower minimum values in the experiment.

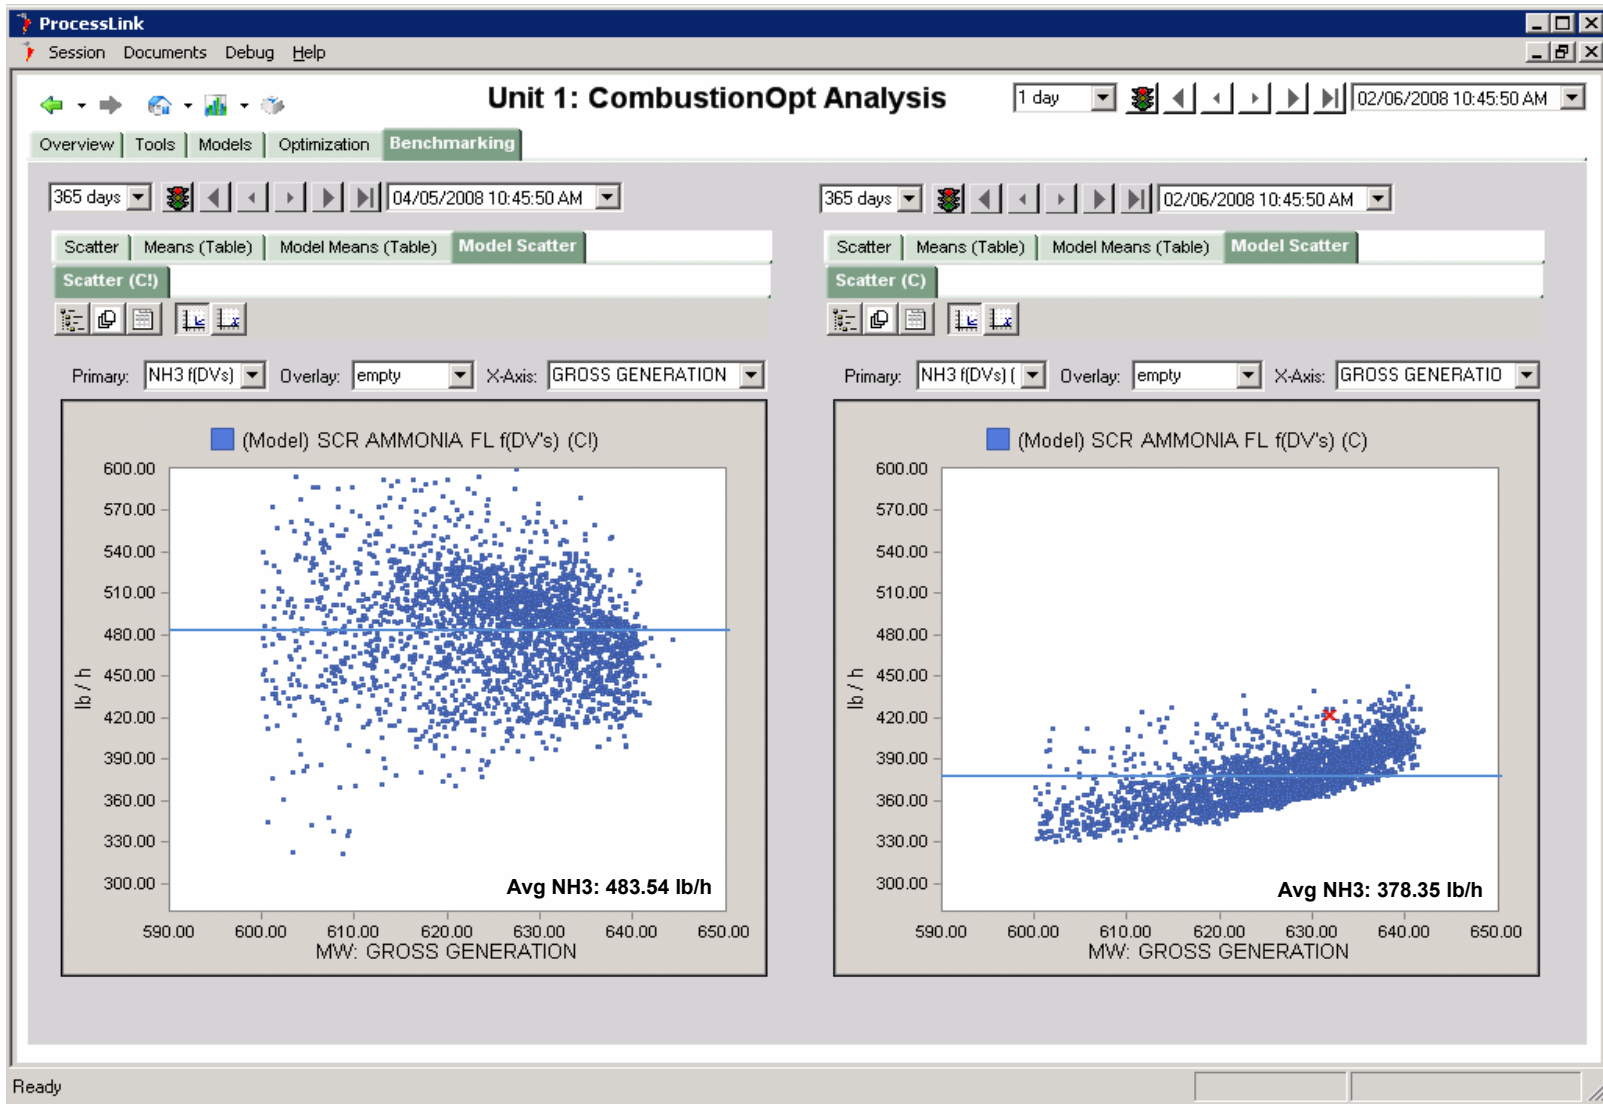


Figure 61 Unit 1 Modeled NH3 as $f(\text{Disturbances})$ vs. MW

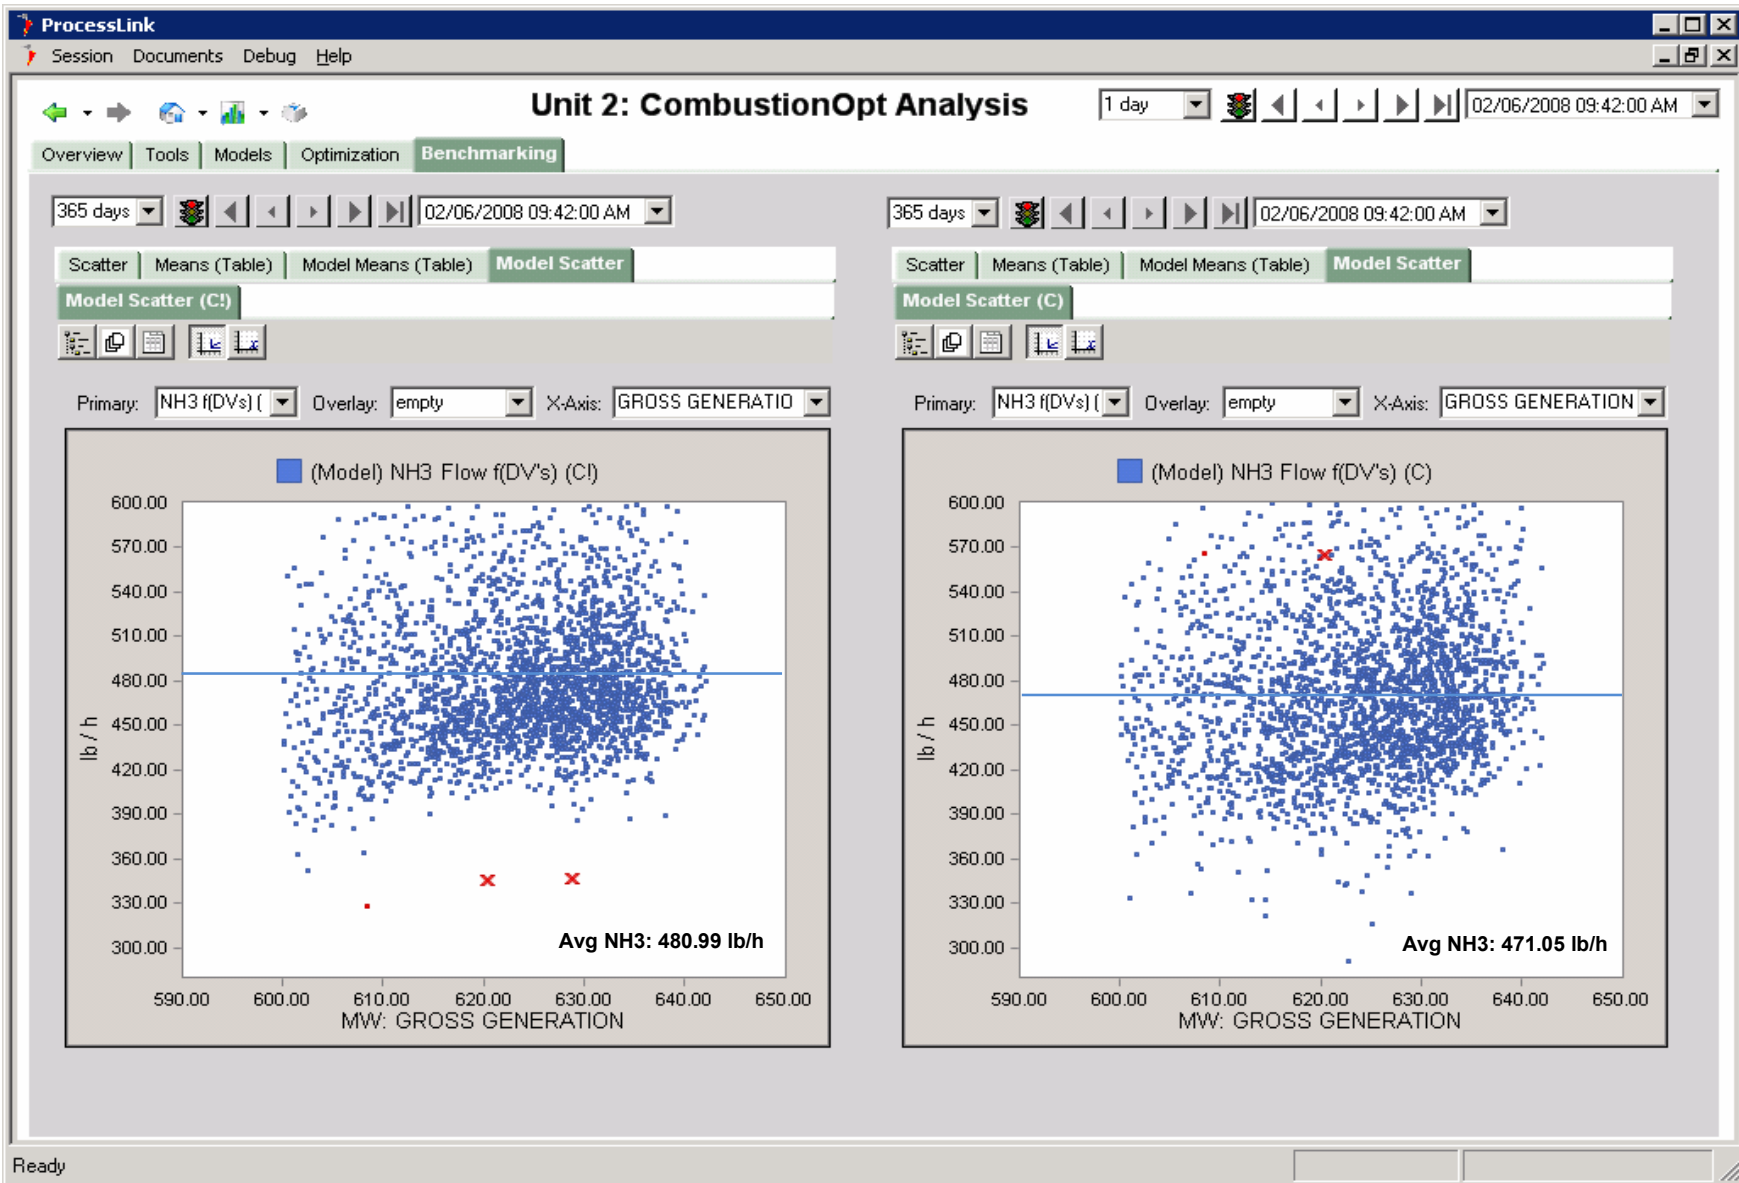


Figure 62 Unit 2 Modeled NH3 as $f(\text{Disturbances})$ vs. MW

4.4.2.2 Other Plots

4.4.2.2.1 Measured Stoichiometry vs. MW

Figure 63 - Figure 64 show the relationship between cyclone stoichiometry and Megawatts for the control and experiment. They illustrate how reduced stoichiometric operation of the cyclones is characteristic of the experimental population.

As previously mentioned, the major challenge for the cyclone units was moving toward reduced stoichiometry (lower NO_x production) without losing good cyclone slag flow. Cyclone function is largely based around the physics and chemistry of slag formation and flow.

PRB coal has higher ash content than high sulfur Illinois coal, with a different mineral make-up. This gives it a much narrower range of temperatures over which slag (molten ash) flows well. The proportion of O₂ present also affects the characteristics of the slag. Added to both of these issues, PRB coal has a lower overall energy density than high sulfur Illinois coal. This means more of it has to flow through the combustor to provide the same heat release. This is why plants (not just cyclones) typically have to take a 10-20MW derate when switching to 100% PRB coal. The fuel delivery system is just not designed to accommodate the higher mass flow and the different volatility characteristics of the fuel.

For BEC to switch to 100% PRB coal and avoid a de-rate, all three of these challenges, stoichiometry, temperature and mass flow, would have to be addressed. To complicate this further, there are no reliable temperature or O₂ measurements available to provide insight into conditions within the cyclone. Standard operating procedure has been for a very experienced combustion engineer to check the color of the flame through the sight glass on a twice-daily basis (more art than science). The consequence of failing to meet these combustion control challenges is filling up a cyclone with hardening slag, a condition so severe that it often requires an outage to dynamite the formation out. Cyclones in which slag flow is stalling also tend to become runaway NO_x generators.

Working with BEC Combustion Engineers, NeuCo configured CombustionOpt to meet these challenges by working to maintain the flame scanner quality signal, the closest proxy for temperature, above 65% while stoichiometry was reduced. It is worth noting that this signal is typically used only to determine whether or not flame is present (for safety interlocks).

Figure 65 - Figure 66 show the average Cyclone Main Flame Scanner Quality and demonstrate that strategy described above, along with advice and alerts fed to MaintenanceOpt (to relieve the requirement for vigilance by busy staff), improved average cyclone flame quality in the face of these challenges. This achievement provides direct input into the successful switch to 100% PRB and Low NO_x combustion.

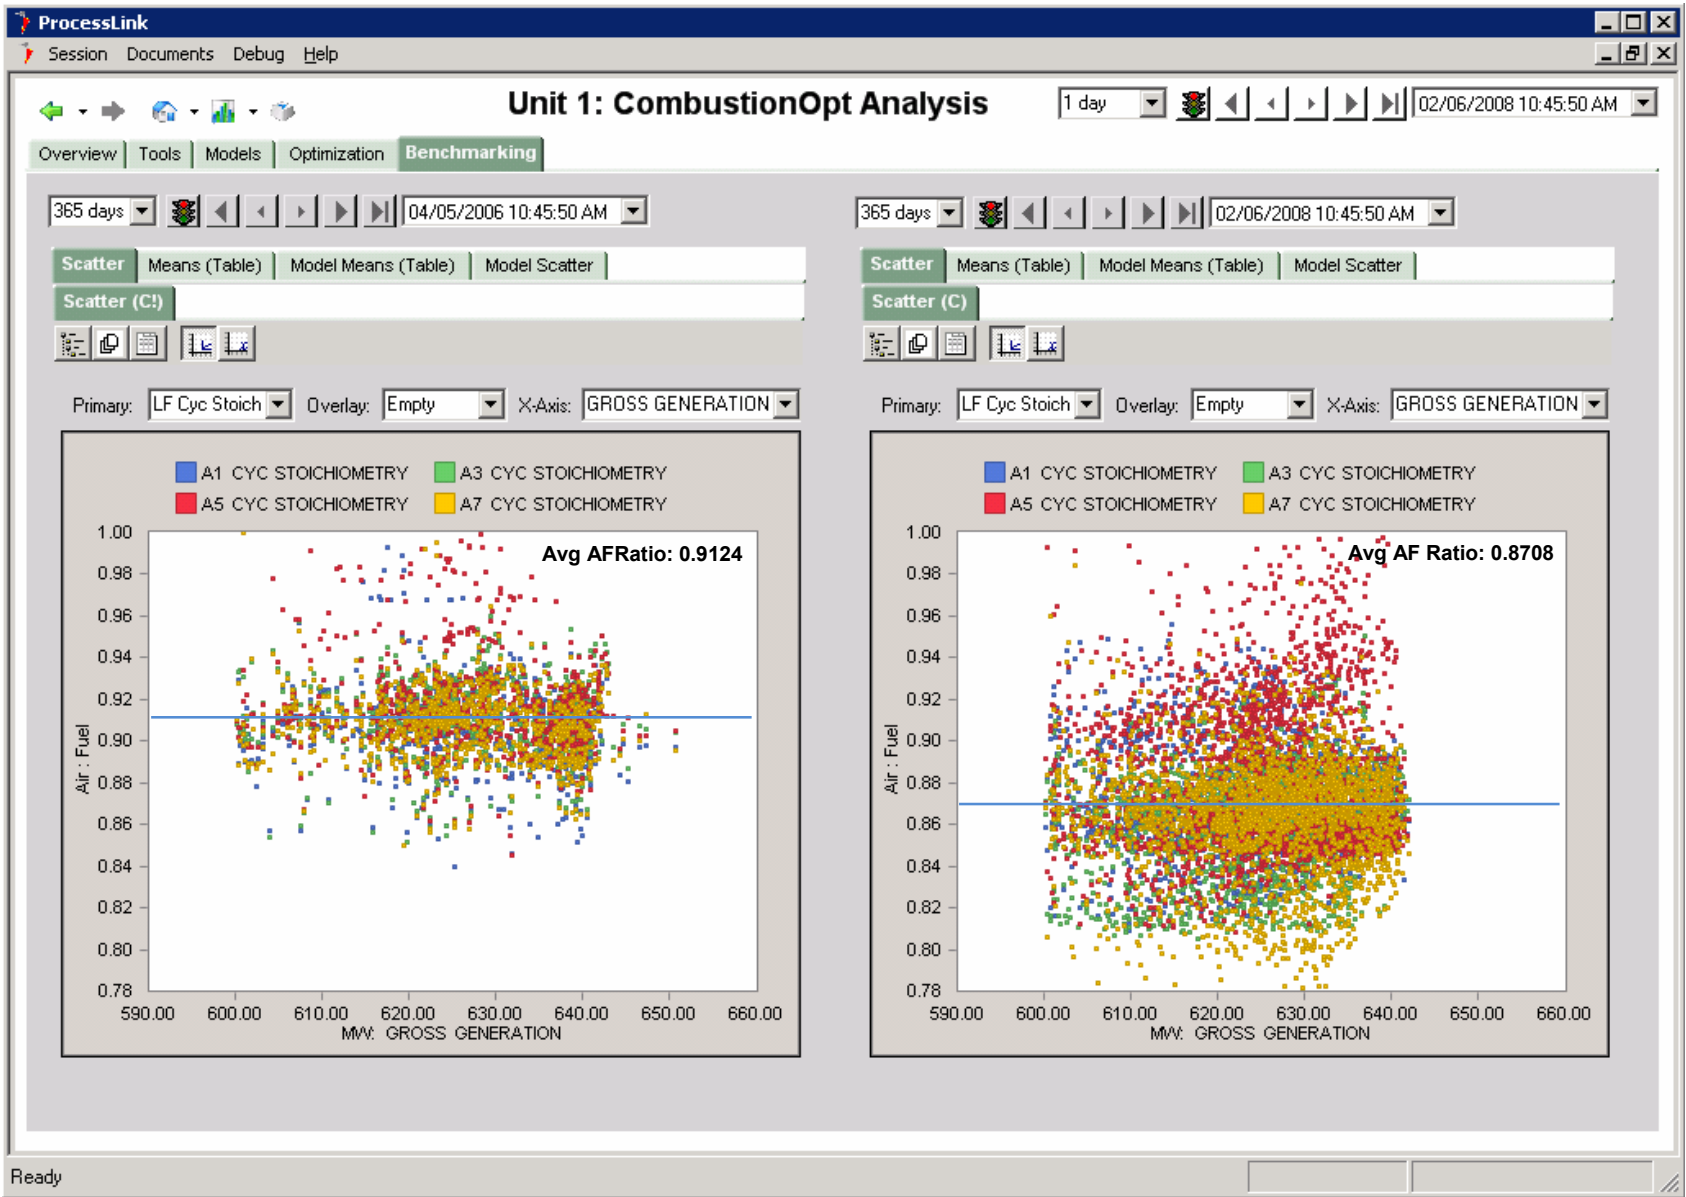


Figure 63 Unit 1 Cyclone Stoichiometry vs. MW

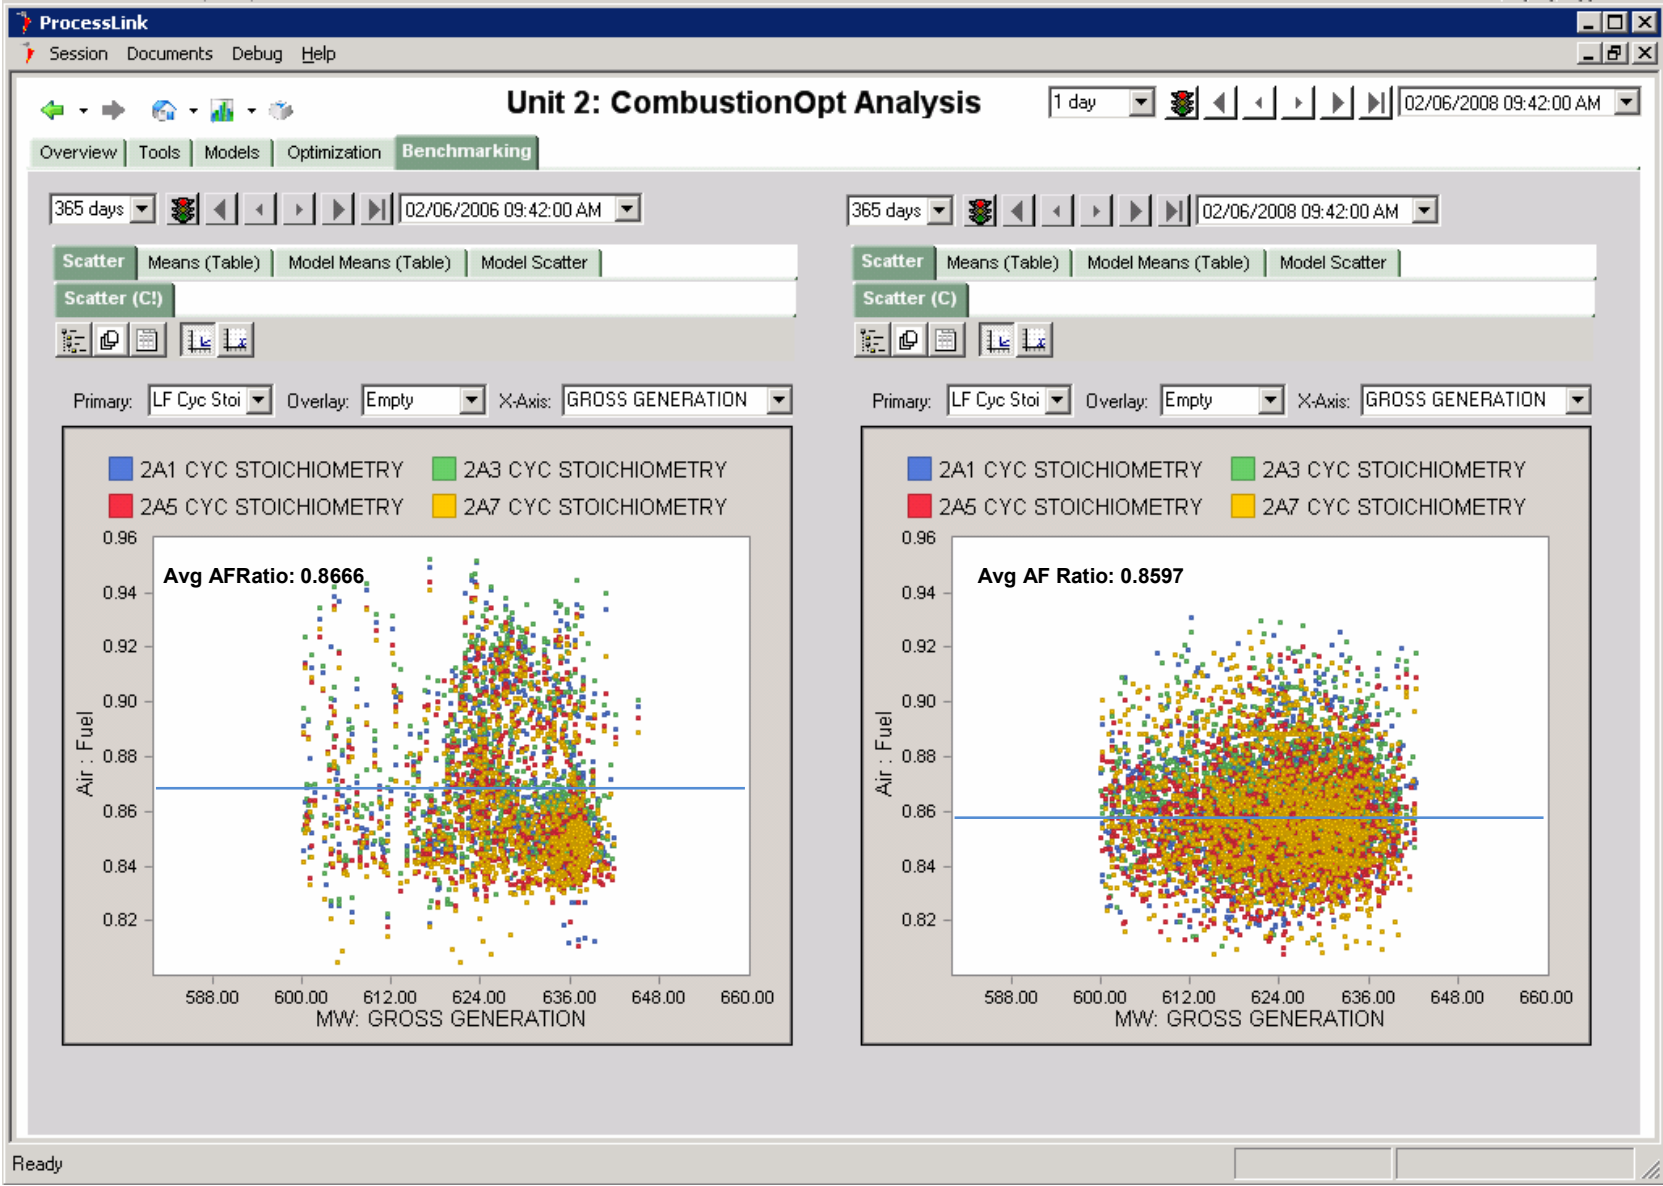


Figure 64 Unit 2 Cyclone Stoichiometry vs. MW

4.4.2.2.2 Average Main Flame Scanner Quality

Figure 65 - Figure 66 show the average of the Main Flame Scanner Quality signal over the course of the project. The goal of the optimizer has been to keep them above 65% while also minimizing furnace NOx and Heat Rate.

Observations:

- The cyclones on Unit 1 were more responsive than those on Unit 2.
- The flame scanners were installed and running on Unit 2 at the beginning of the project and were not added to Unit 1 until spring of 2005.
- Both units have seen improvements to average quality.

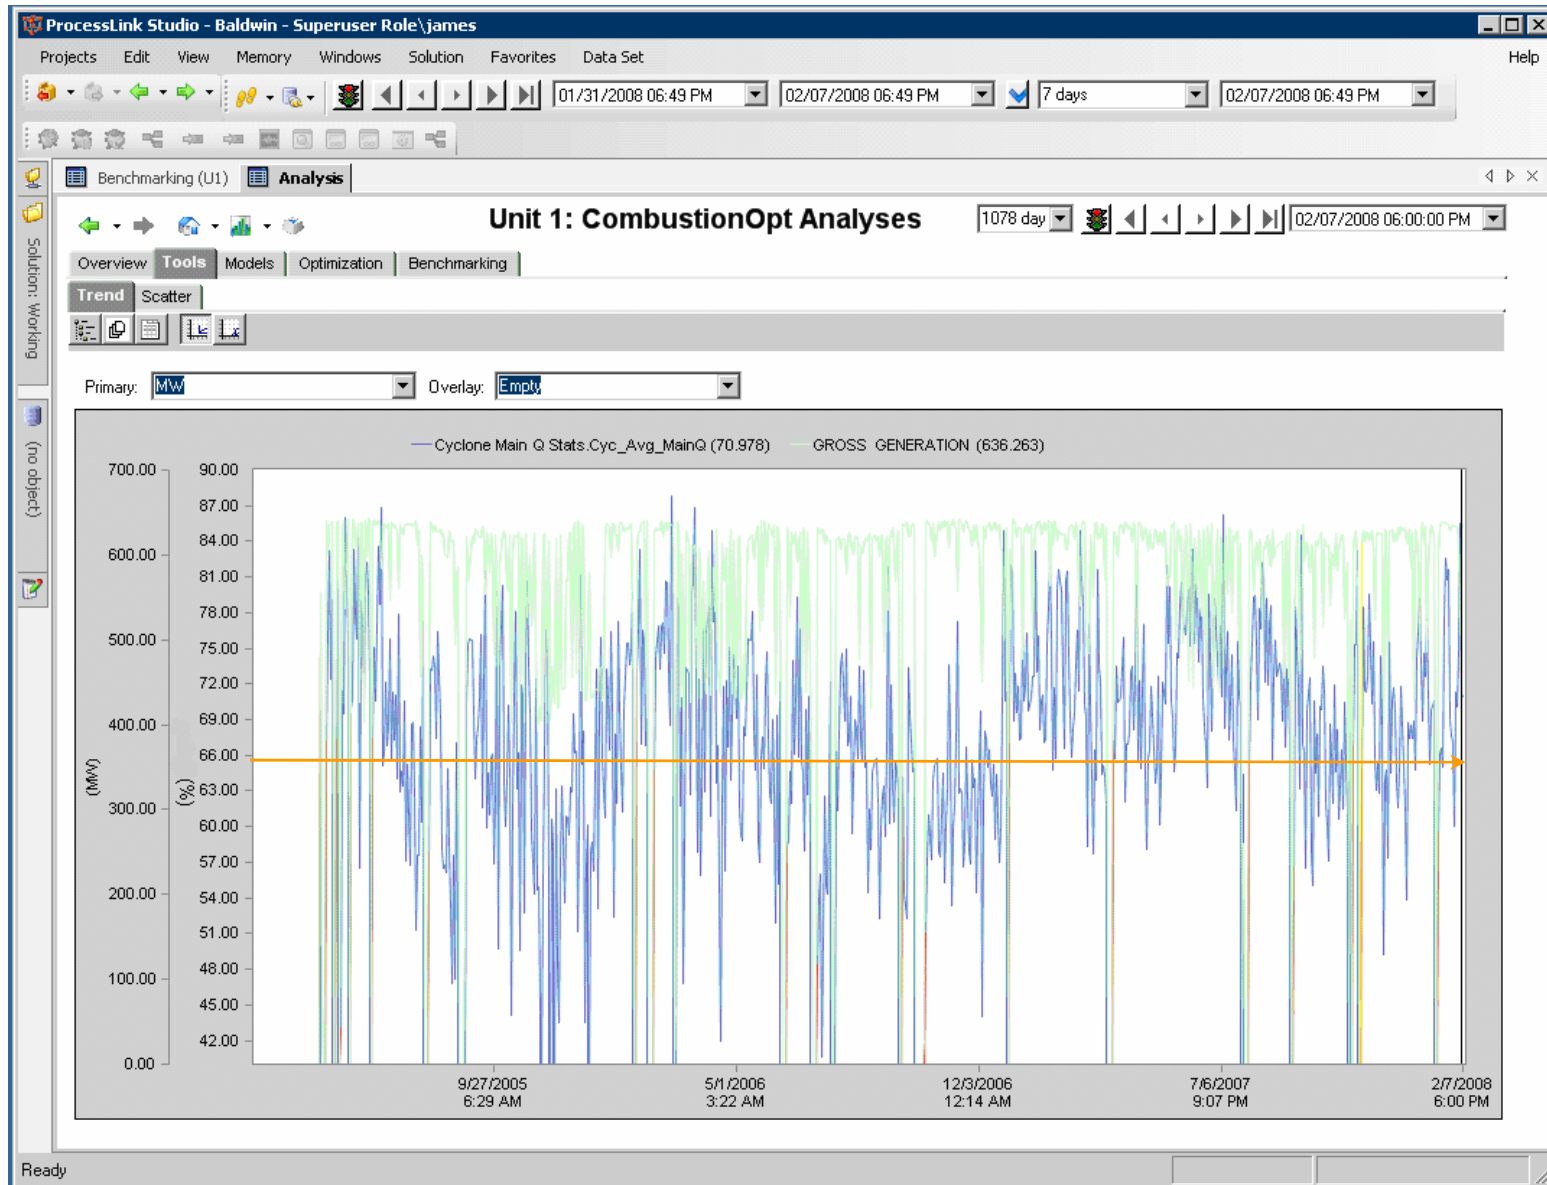


Figure 65 Unit 1 Average Cyclone Main Flame Scanner Quality (%) (2005 - present). Orange line shown for visual reference; Red lines indicate where data was not available

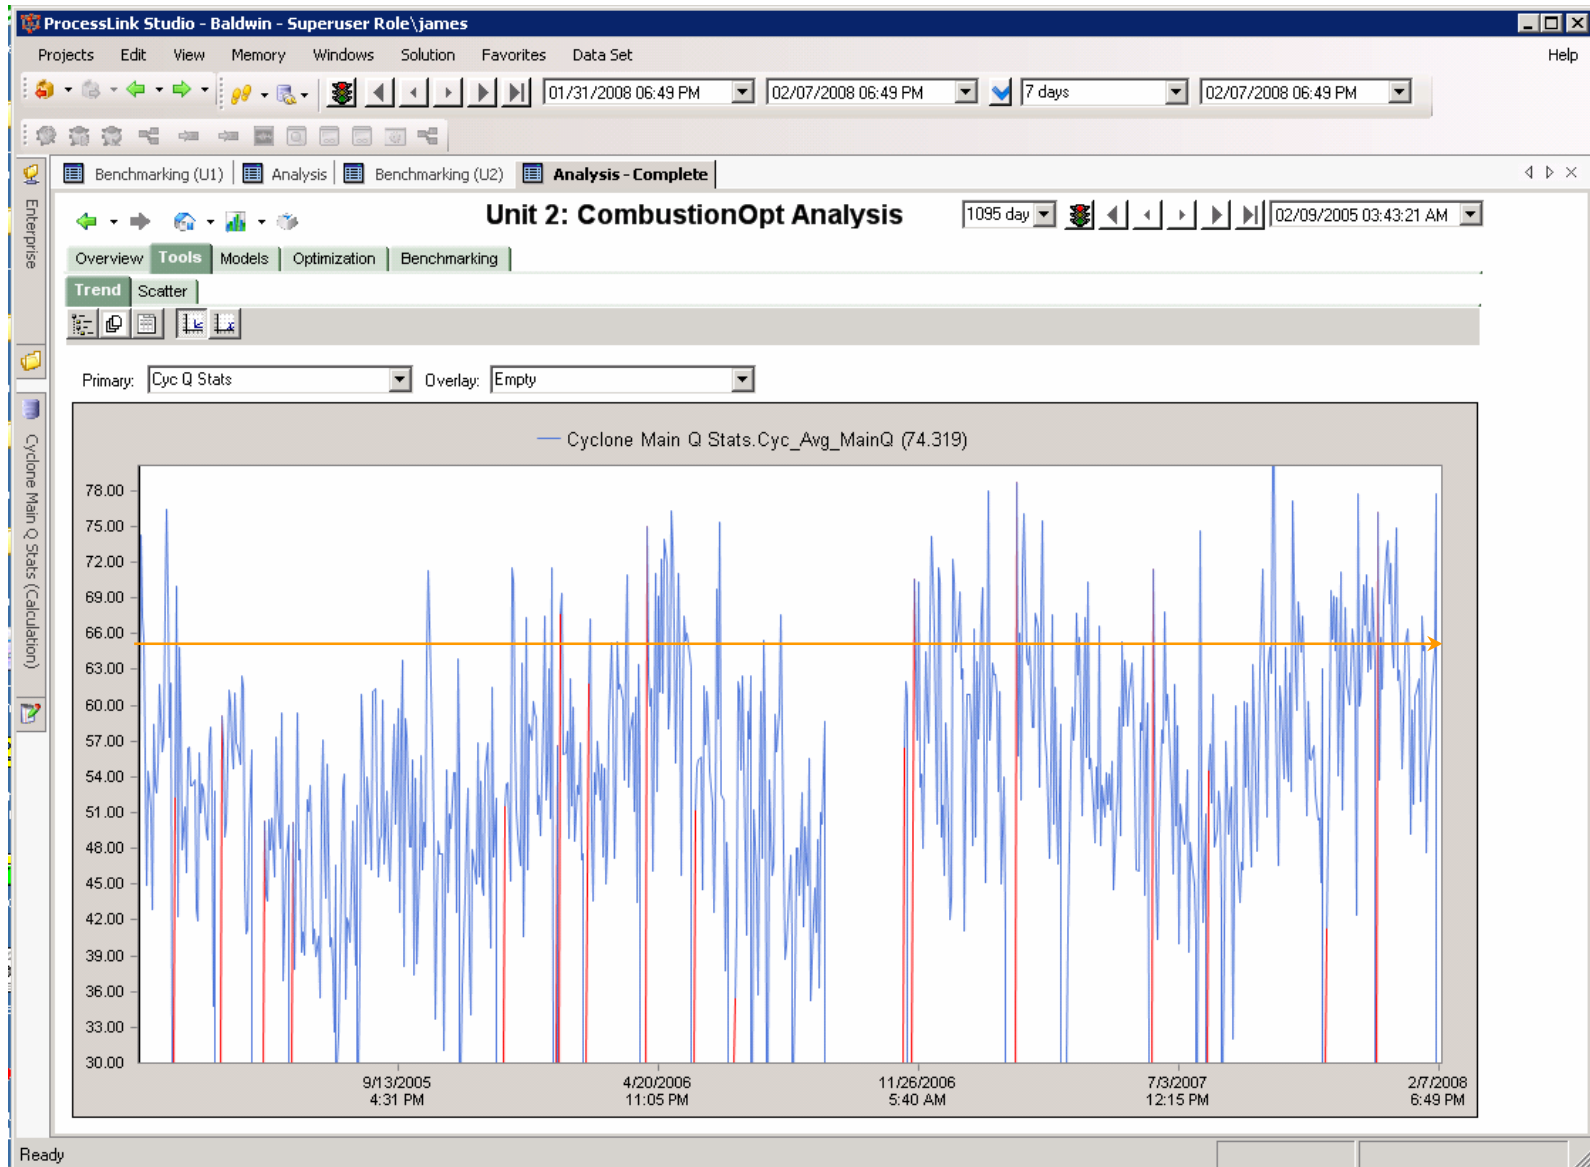


Figure 66 Unit 2 Average Cyclone Main Flame Scanner Quality (%) (2005 - present) Orange line shown for visual reference; Red lines indicate where data was not available

4.4.2.2.3 Measured Sootblowing Steam flow vs. MW (Unit 3)

Figure 67 shows the sootblowing steam flow in klb/h in the control and experiment populations (a delta of -1.58909 klb/h or 14.7%). This avoided steam usage is a direct contributor to improved Heat Rate. But more importantly it indicates that fewer sootblowing operations occurred because they were deemed unnecessary by SootOpt working in concert with the ISB.

Tube failures are the single most significant cause of forced outages and a significant portion of these failures are due to sootblower erosion and thermal shocking, which occurs when high pressure, relatively cool steam impinges clean tubes.

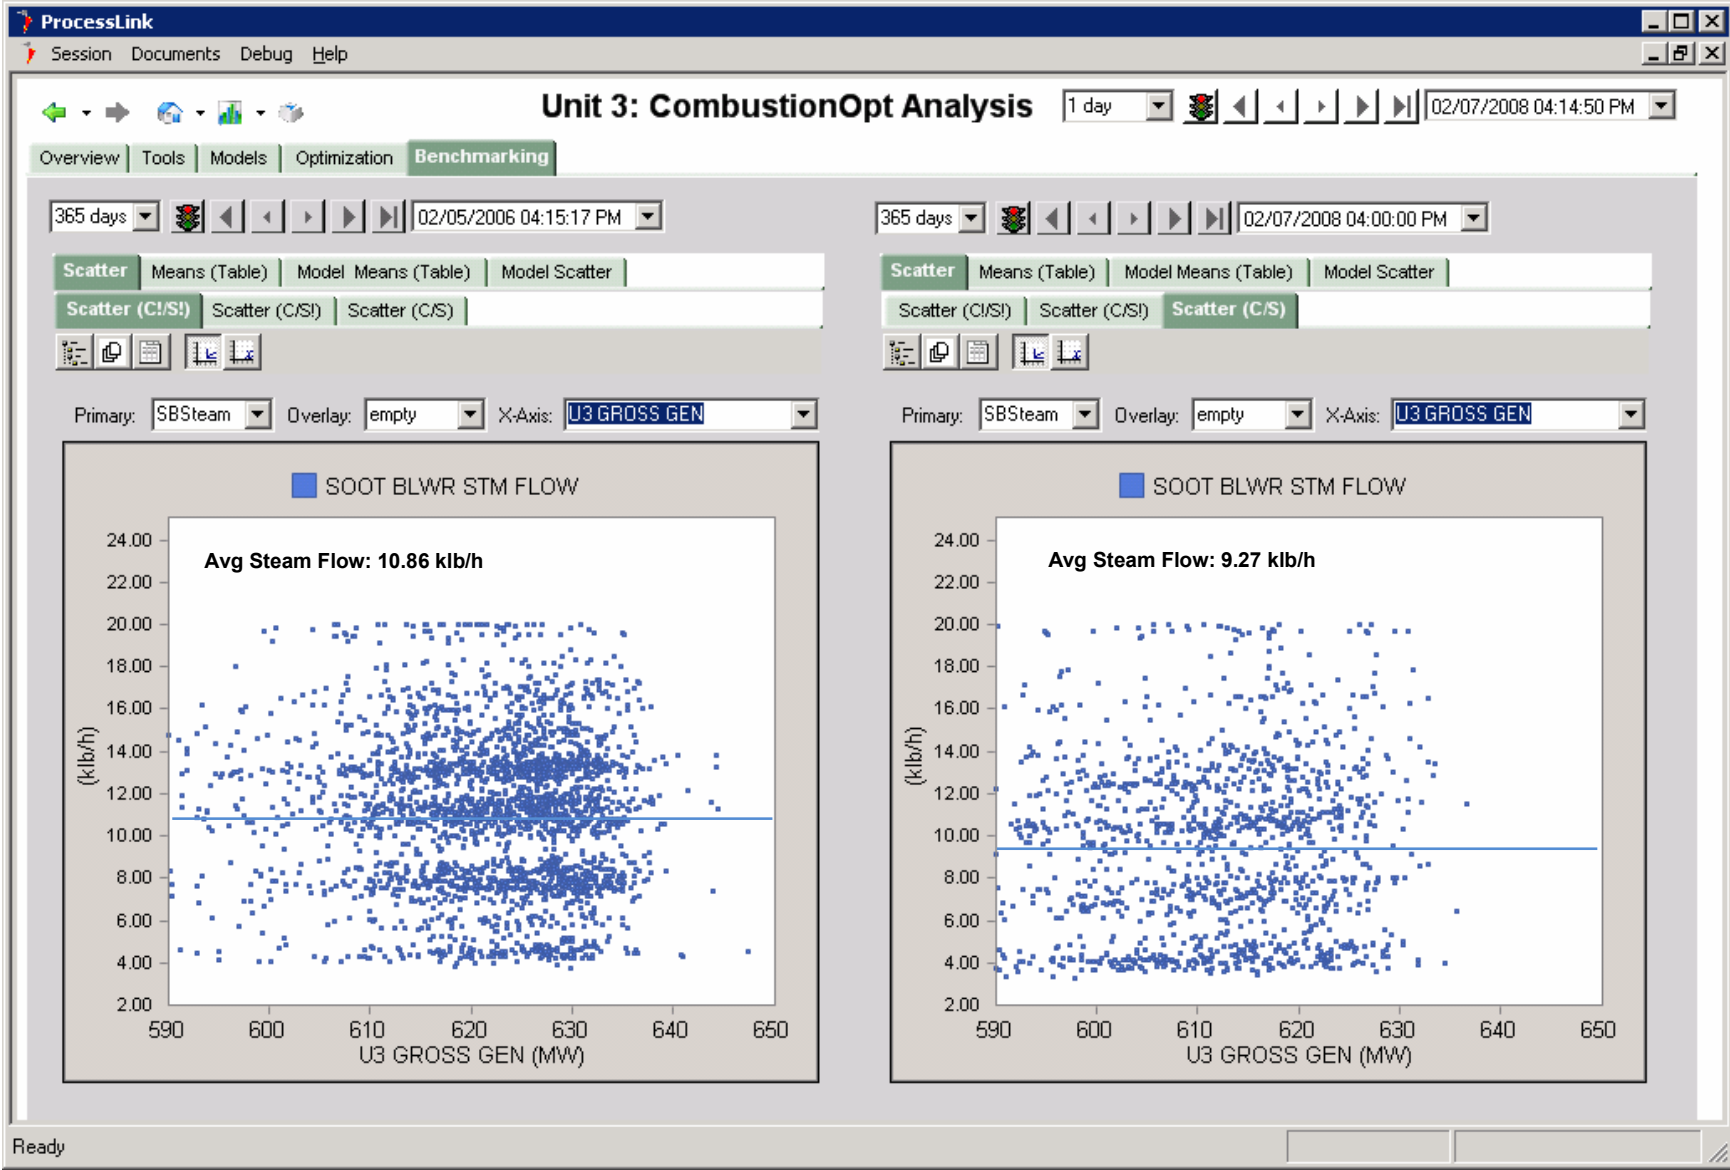


Figure 67 Unit 3 Sootblowing Steam Flow

4.5 Advisory Optimization

As stated in Section 4.1, it is necessary to evaluate the impacts of closed-loop and advisory optimization in different ways. It is much more difficult to quantify the benefits provided by advisory optimization because we often do not know how much faster a problem/opportunity was identified and resolved as a result of the advice, or what the impact would have been had the advice not been taken. This section will use examples of real alerts to demonstrate that advisory optimization is providing substantial value to BEC, both through the early identification/resolution of problems and opportunities, and through the time savings and peace of mind provided to plant personnel by knowing they have a second set of eyes constantly searching for performance issues and improvement opportunities.

While all of NeuCo's optimizers have an advisory component, when multiple optimizers are deployed MaintenanceOpt is the main interface for engineers and operations. MaintenanceOpt collects all of the alerts generated by all the products and presents them inside a framework for comparing their relative importance and for diagnosing and taking action. BEC has engineers assigned to use MaintenanceOpt and monitor each of the three units, and that is where they go to prioritize, analyze and take action as needed on problems that may have been identified by CombustionOpt, SCR-Opt, PerformanceOpt, SootOpt and MaintenanceOpt using their respective built-in detection engine and alerting rules.

One of the key features of MaintenanceOpt is its ability to propose the most likely known reasons for identified anomalies and provide the context data needed to systematically work through the diagnosis – all in one place allowing engineers to quickly identify the source of the problem and determine what action is needed. Another important aspect of MaintenanceOpt is its knowledge capture abilities so that as new things happen knowledge about what caused them and how they were handled can be reused in the future by other plant personnel. This can be in the form of annotations/logs through the detection-to-resolution workflow of the problem as well as in the form of electronic documents.

4.5.1 Advisory Optimization Examples

A complete list of the optimization alerts provided to BEC throughout 2007 can be found in Appendix A1. This section elaborates on five of those.

4.5.1.1 Air Duct Leakage

At the end of February, 2007, a MaintenanceOpt Alert triggered indicating an increase in total amps for the A-side forced draft fan, which is shown in the "Issues to Screen" display, below. As shown in the screen shot, the highlighted total fan amps symptom was accompanied by impact ranking, a risk estimate and context data relating to the most likely problems causing such an increase. This information was used in both the decision to escalate the alert into the "Problems to Diagnose" category (Top View Tabs in Figure 68 represent the sequential workflow steps).

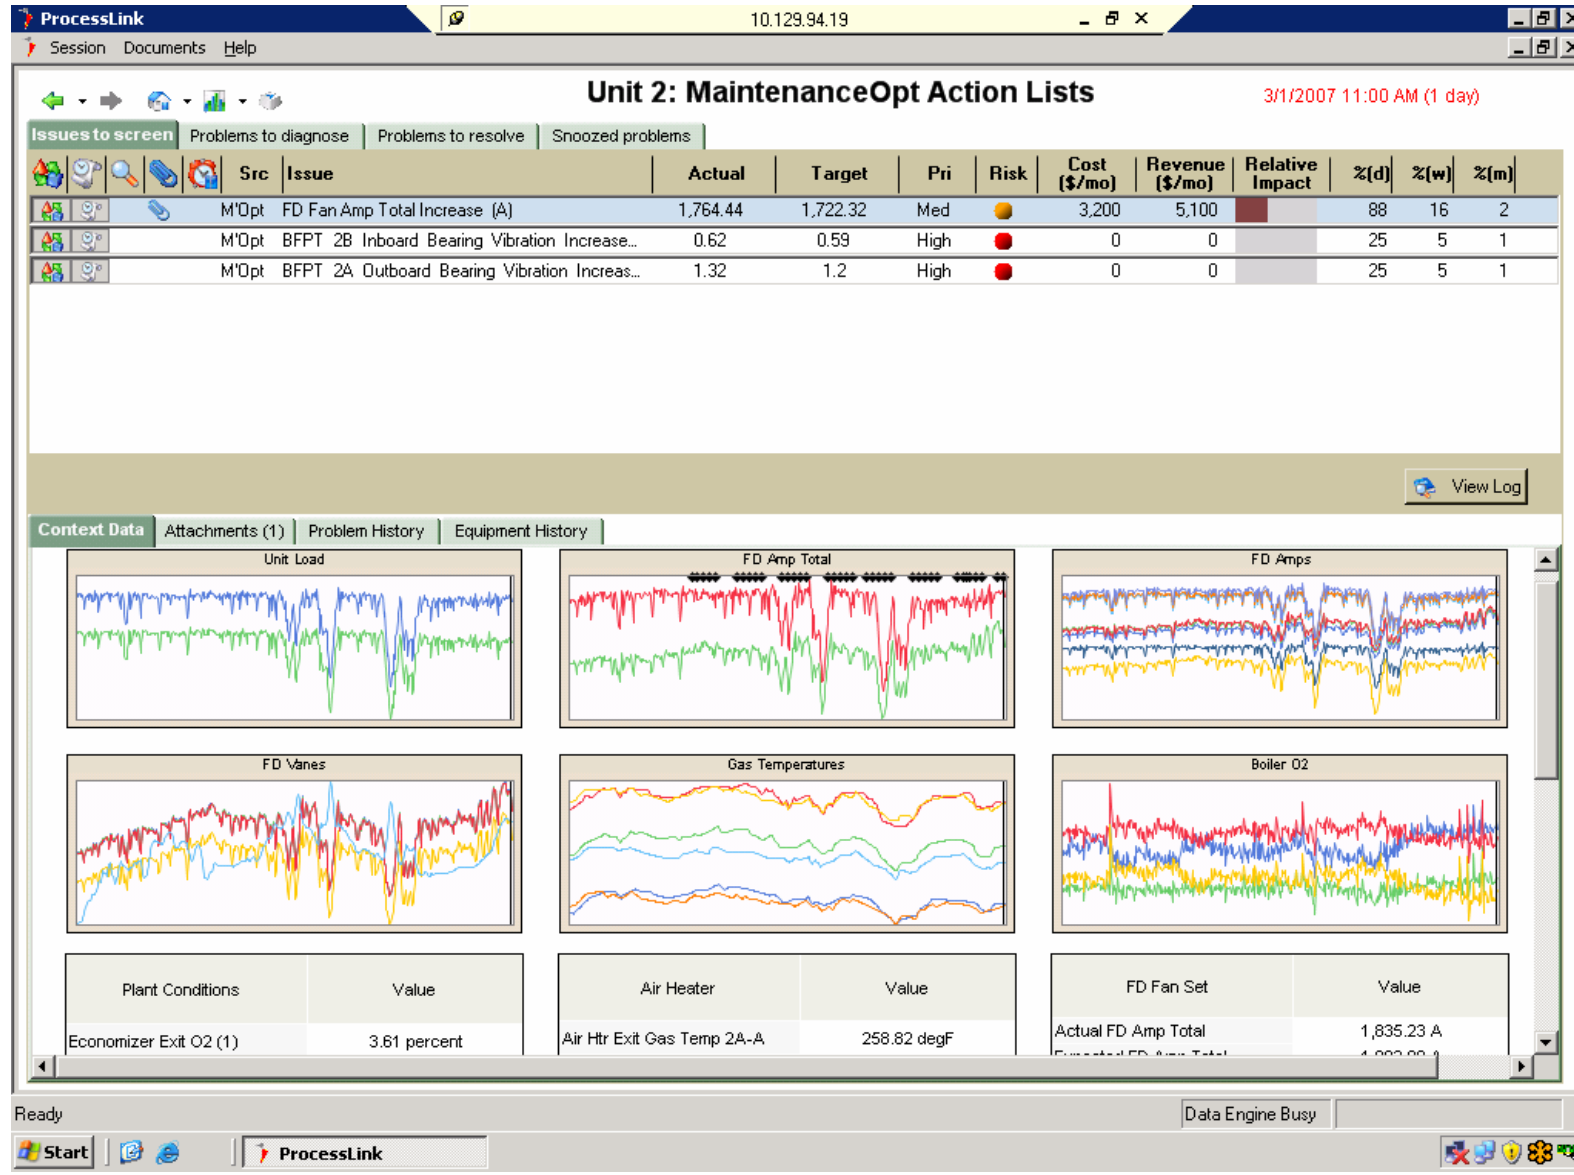


Figure 68 Issues to Screen View with FD Fan Amp Increase and Context Data

Once this problem was escalated, Baldwin engineering personnel could utilize MaintenanceOpt's diagnostics assistance to identify the probable cause as air duct leakage. The top half of Figure 69 below shows how MaintenanceOpt presents a list of potential root causes, ranked in order of the ease with which a definitive diagnosis can be made (e.g. causes which can be diagnosed with currently available online data first, then those requiring visual inspection or manual measurements, finally those requiring and outage and/or disassembly of equipment).

In the lower half of Figure 69 the graphical and tabular data required to determine the cause of the fan amp increase is displayed. Using that data, the Baldwin engineers determined that the probable cause of the problem was "air duct leakage" and made a note in MaintenanceOpt to perform a visual inspection of the potential culprit ducts at the next available opportunity.

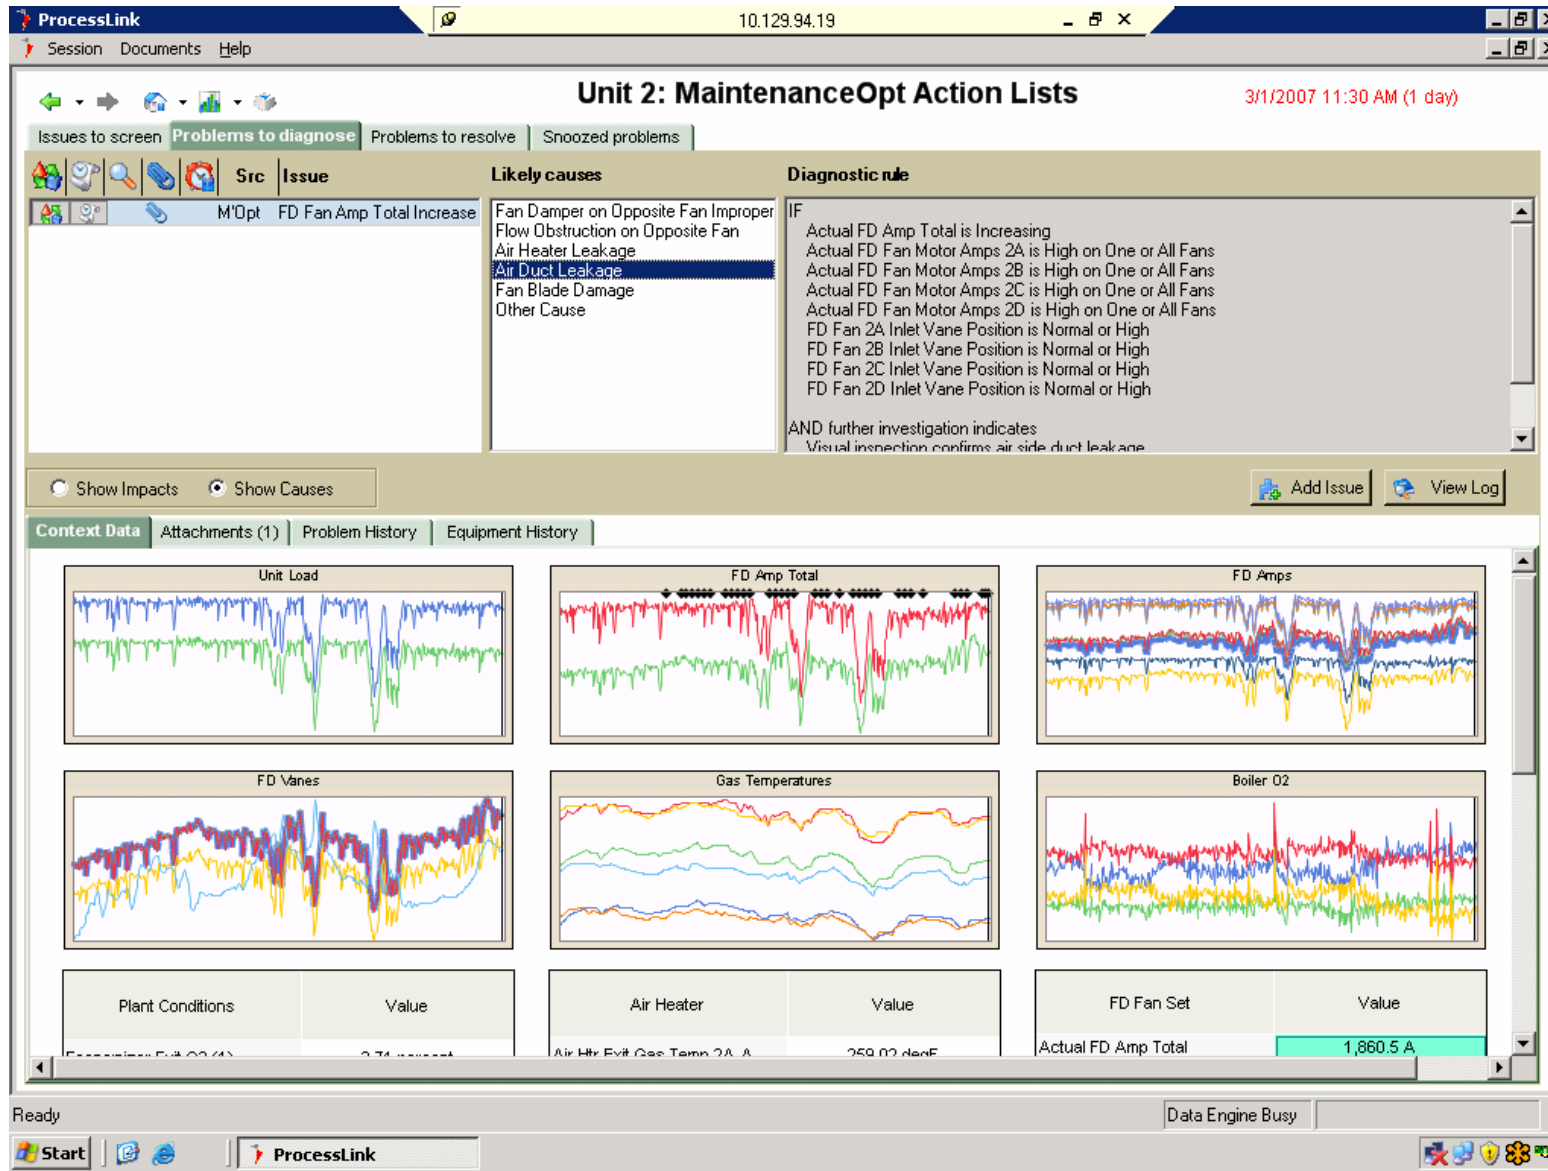


Figure 69 FD Fan Amp Increase Diagnostics View and Context Data

This visual inspection quickly revealed a large hole in one of the unit's over-fire air ducts. Two additional actions were specified at this point: 1) isolate the damaged duct until the next outage allowed it to be repaired; and 2) plan for the repair so that it could be efficiently addressed during this next outage (Figure 70).

One of the many ways that MaintenanceOpt facilitates problem resolution is by leveraging existing tools and technologies, allowing relevant electronic files or documents to be attached to the problem itself. These are then available throughout the problem lifecycle and into the future to aid in understanding similar problems when they occur. For this particular problem, Baldwin attached a photograph of the hole in the duct as a JPEG file, shown in Figure 71.

As described above, the history of any problem can always be accessed, whether categorized by type of problem or equipment. The history of this particular problem, from symptom alert to resolution is shown in Figure 72.

ProcessLink 10.129.94.19

Session Documents Help

Unit 2: MaintenanceOpt Action Lists

3/1/2007 11:30 AM (1 day)

Issues to screen Problems to diagnose Problems to resolve Snoozed problems

Src	Issue	Likely causes	Diagnostic rule
M'Opt	FD Fan Amp Total Increase	Fan Damper on Opposite Fan Improperly Opening Flow Obstruction on Opposite Fan	IF Actual FD Amp Total is Increasing

Change Status

Specify Action

Issue: FD Fan Amp Total Increase

Priority: **High**

Indicator: Total Current

FD Fan Set

Cause: **Fan Damper on Opposite Fan Improperly Opening**

Reliability Risk: **Medium**

Action

Timing: **Operations**

Action: **Inspect damper controls, verify position signal and visual position, fix online or schedule of**

Supporting Data

Diagnostic Information:

Visual inspection of damper controls shows discrepancy between actual position and position signal.

Actual FD Amp Total is Increasing

Actual FD Fan Motor Amps 2A is High on One or All Fans

Actual FD Fan Motor Amps 2B is High on One or All Fans

Actual FD Fan Motor Amps 2C is High on One or All Fans

Actual FD Fan Motor Amps 2D is High on One or All Fans

Expected FD Fan Motor Amps 2E is Normal or Lower than actual

Comments

This is an instrumentation problem

Close

Show Impacts

Context Data Attachment

Plant Conditions

Add Issue View Log

FD Amps

Boiler 02

FD Fan Set	Value
FD Amp Total	1.860 EA

Data Engine Busy

Ready

Start ProcessLink Change Status

Figure 70 Initial Actions Specified for Damaged Over-Fire Air Duct

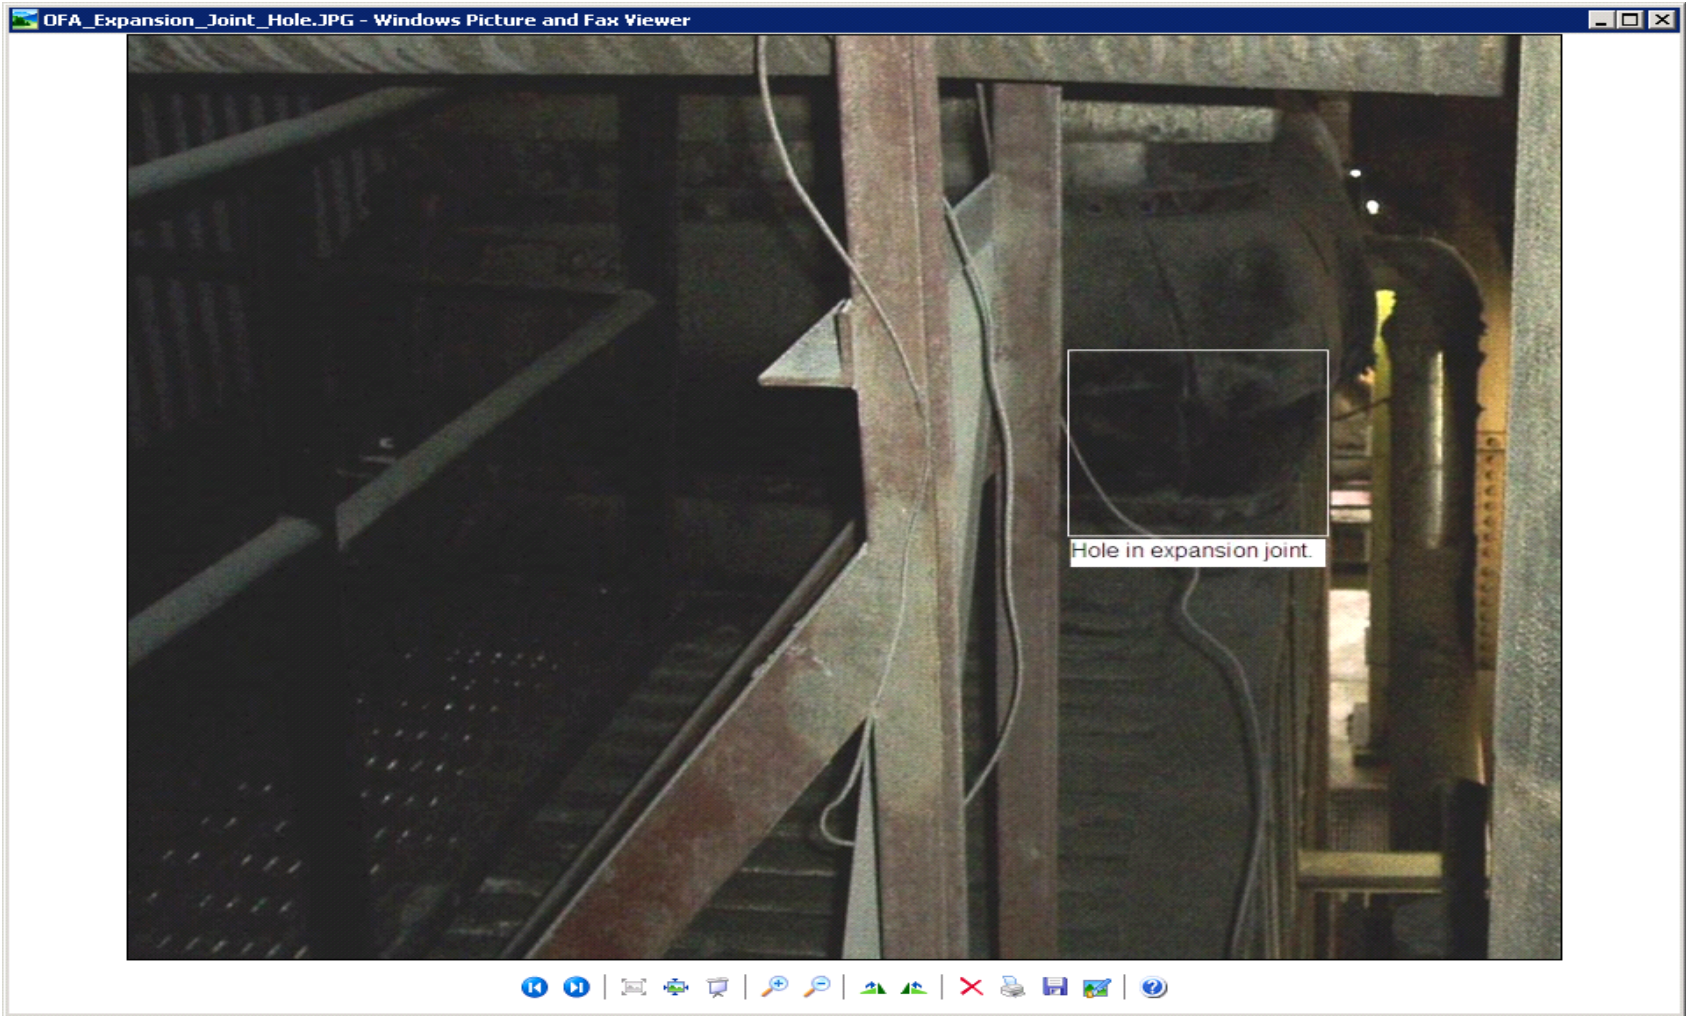


Figure 71 JPEG Image of Damaged Over-Fire Air Duct

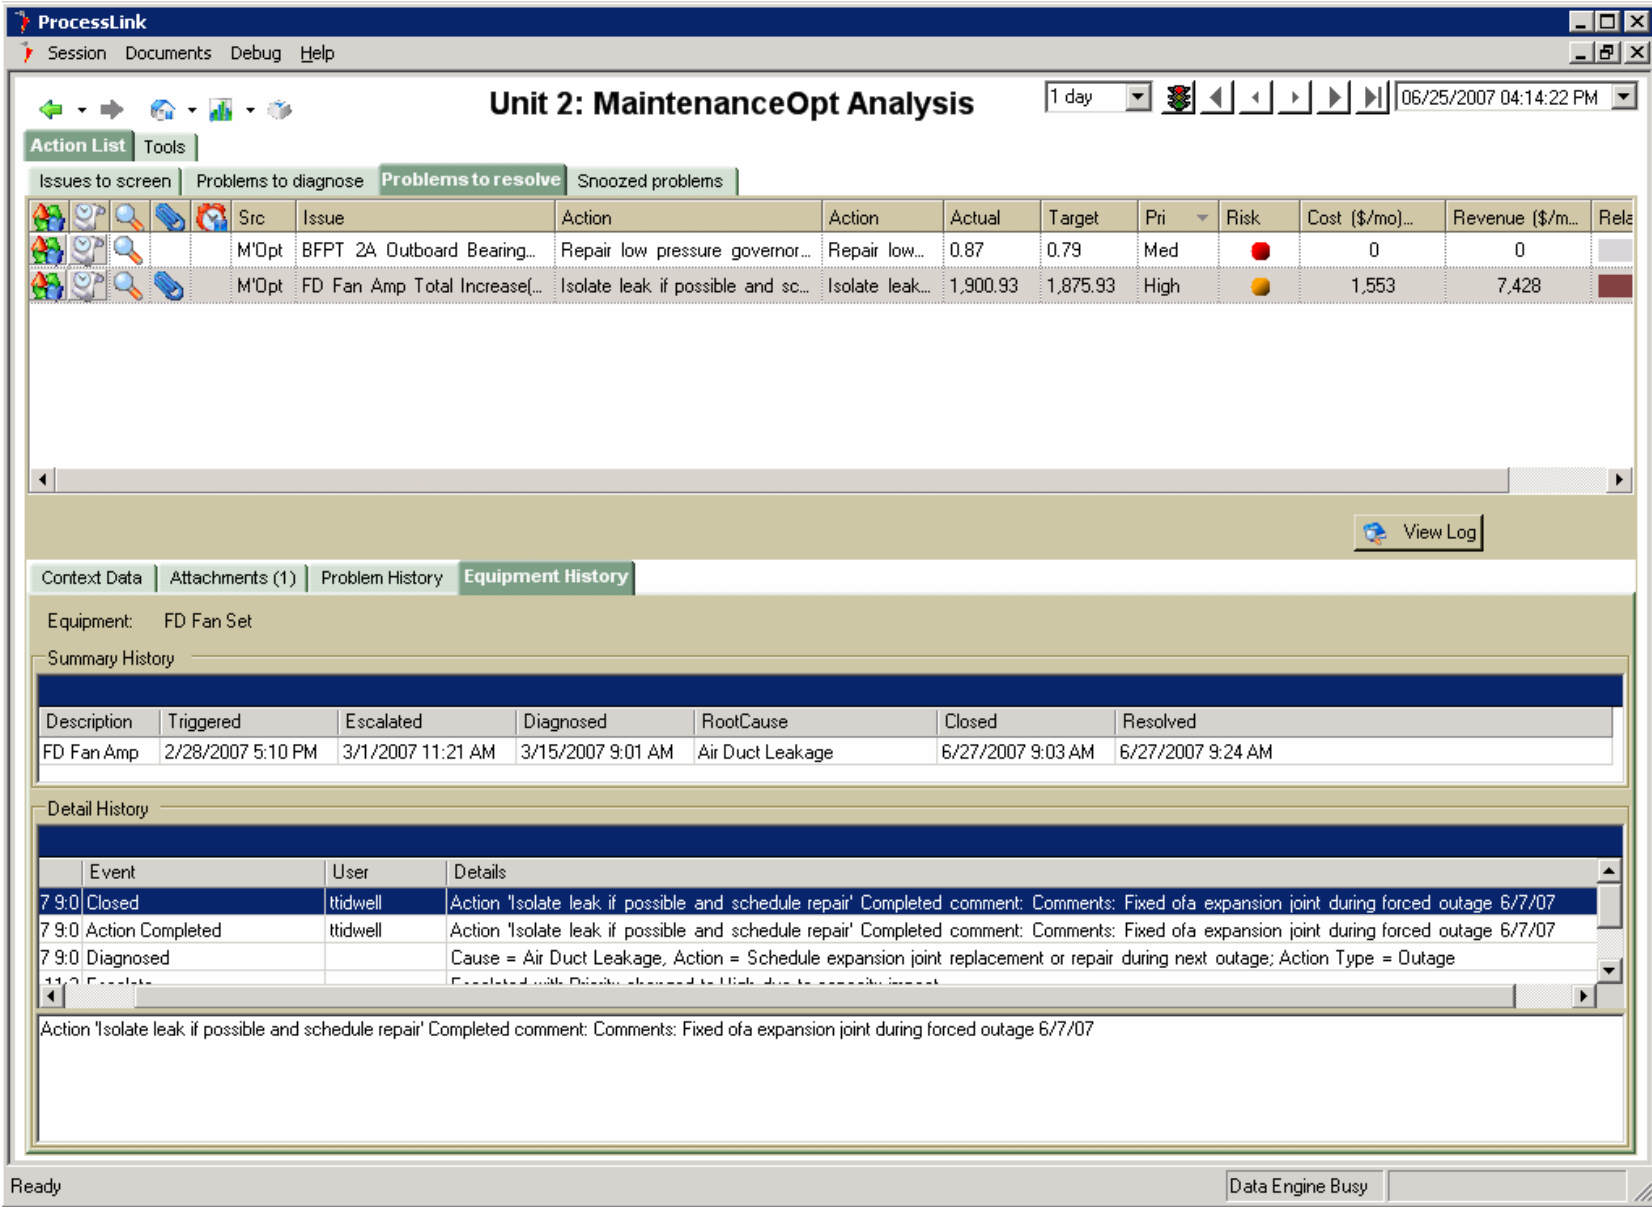


Figure 72 Problem History for Damaged Over-Fire Air Duct

4.5.1.2 Feedwater Heater/Tube Leak

In October, 2007, Baldwin engineers were alerted by MaintenanceOpt to the fact that one train of the high pressure feed water heaters were operating with higher than expected drain cooler approach (DCA). With the built-in knowledgebase of root causes and contextual information available through the MaintenanceOpt system, the plant was able to identify (or at least suspect) shell-side drain cooler tube leak as being the likely cause, and BEC continued to operate with the FWH train isolated until an outage in early November. Once the heat exchanger was opened, BEC Maintenance found that some nuts on the division plate had come loose and damaged the tubes, around 25 of which were plugged during the outage. All of these events and their details have been captured in the MaintenanceOpt archive and are available for reference in the future for similar or related issues; please refer to Figure 73.

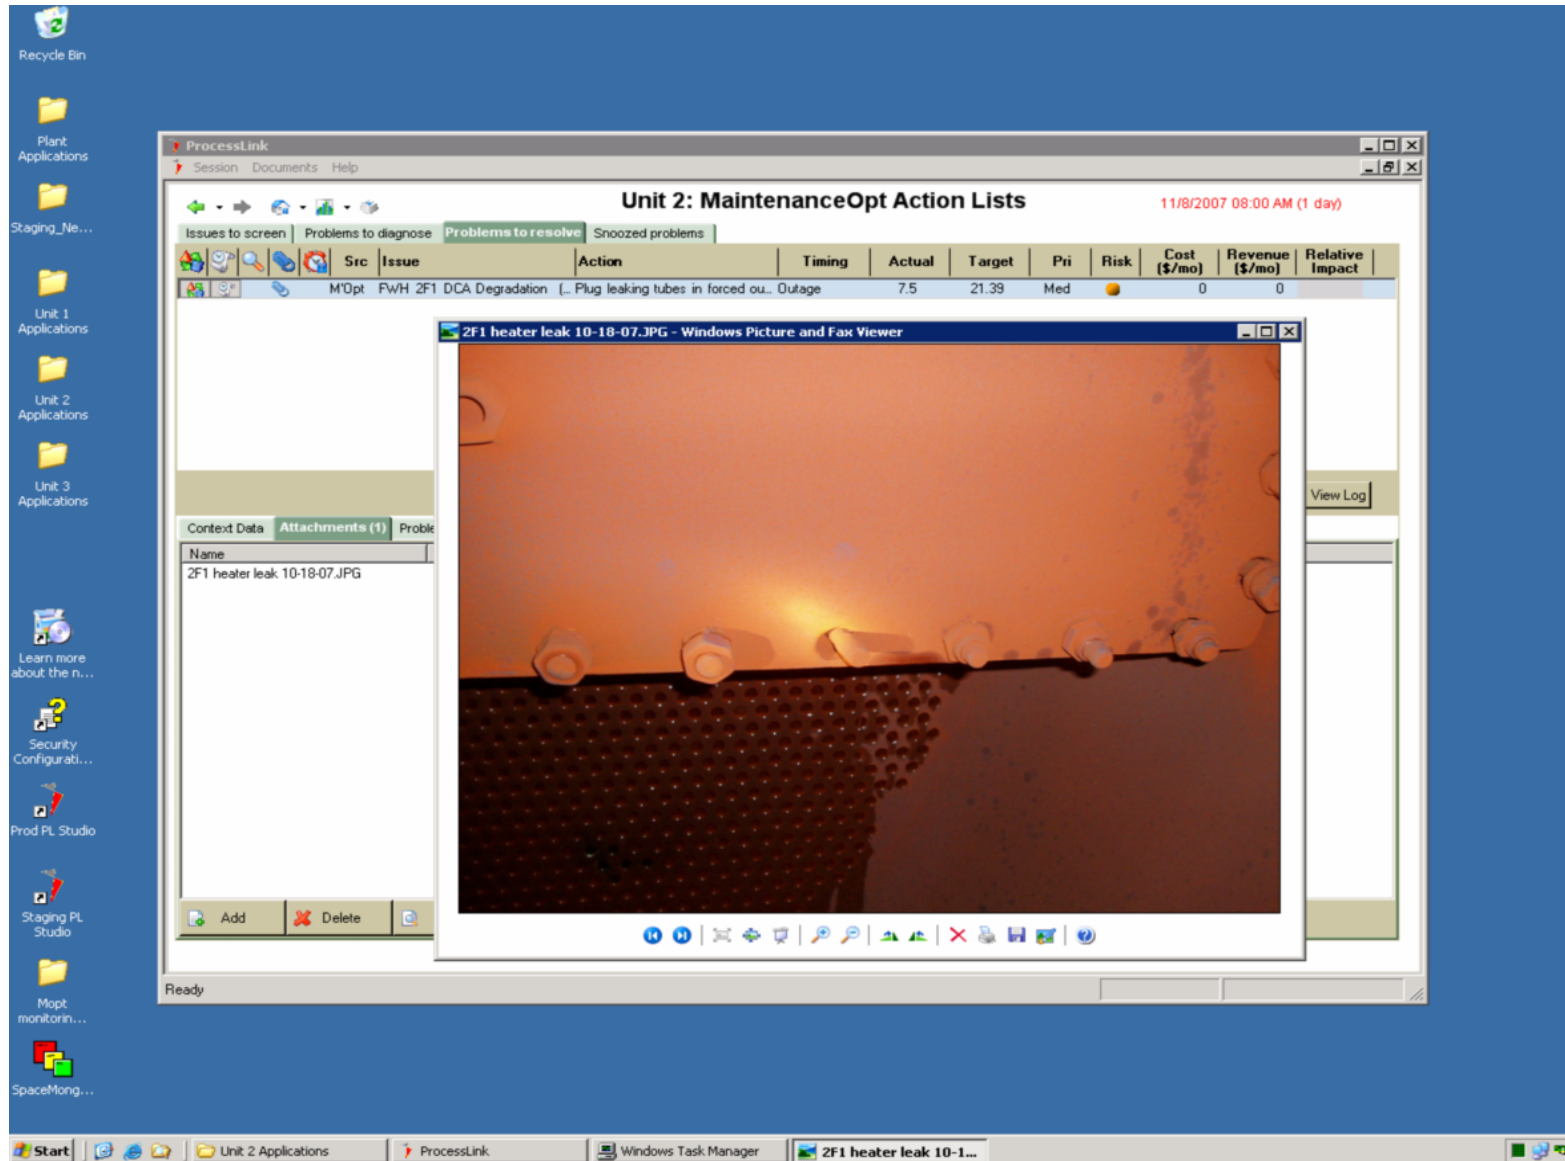


Figure 73 JPEG Image of Missing Feedwater Heater Nuts

4.5.1.3 Bad Thermocouple

SootOpt provides BEC with key insights into their soot cleaning operations that they did not previously have, facilitating the transfer of a process that had historically been done by operators based on a set schedule, to an automated one based on feedback from advanced instrumentation. An important aspect of SootOpt's expert system is the ability to define "holistic" rules and constraints that the soot cleaning controls need to operate within, rules that don't necessarily have to have an ASME performance metric (like DCS) as their foundation. In this example, SootOpt became aware that the frequency of soot cleaning was changing relative to its historical pattern. And more specifically, that very little sootblowing had occurred in the horizontal sections of the boiler in the past eight hours. The continuing deposition of ash and soot on these sections was constricting boiler airflow and causing differential pressures to increase. Although the knowledge of exactly why this was happening was not codified in a set of diagnostic rules at the time, knowing about it and having the opportunity to diagnose it before it became a real problem was critical. In the end the decrease in operating frequency of the Diamond Power ISB (which SootOpt is integrated with in closed-loop) was due to a bad thermocouple which had set its cleaning decision algorithm into alarm and stopped activity. As is typical of many control rooms, Operators had received training but had little experience with the possible alarm states of the ISB and the consequences these states represent. Because this alarm was not yet connected to a perceived consequence in their experience it had gone un-noticed. Had SootOpt's "second set of eyes" not been in place, this issue could have led to a derate or even tripped this base-loaded unit.

Figure 74 below shows MW (red line) and furnace to economizer differential gas pressure (blue line). The spike in differential pressure (1) indicates where plugging of the horizontal sections of the convection pass was beginning to restrict boiler airflow. The vertical cursor line (2) marks where a modeled objective for furnace economizer differential pressure was added. During the interval between, the IKs were exercised manually to make sure that any deposits that may have been built up were removed.

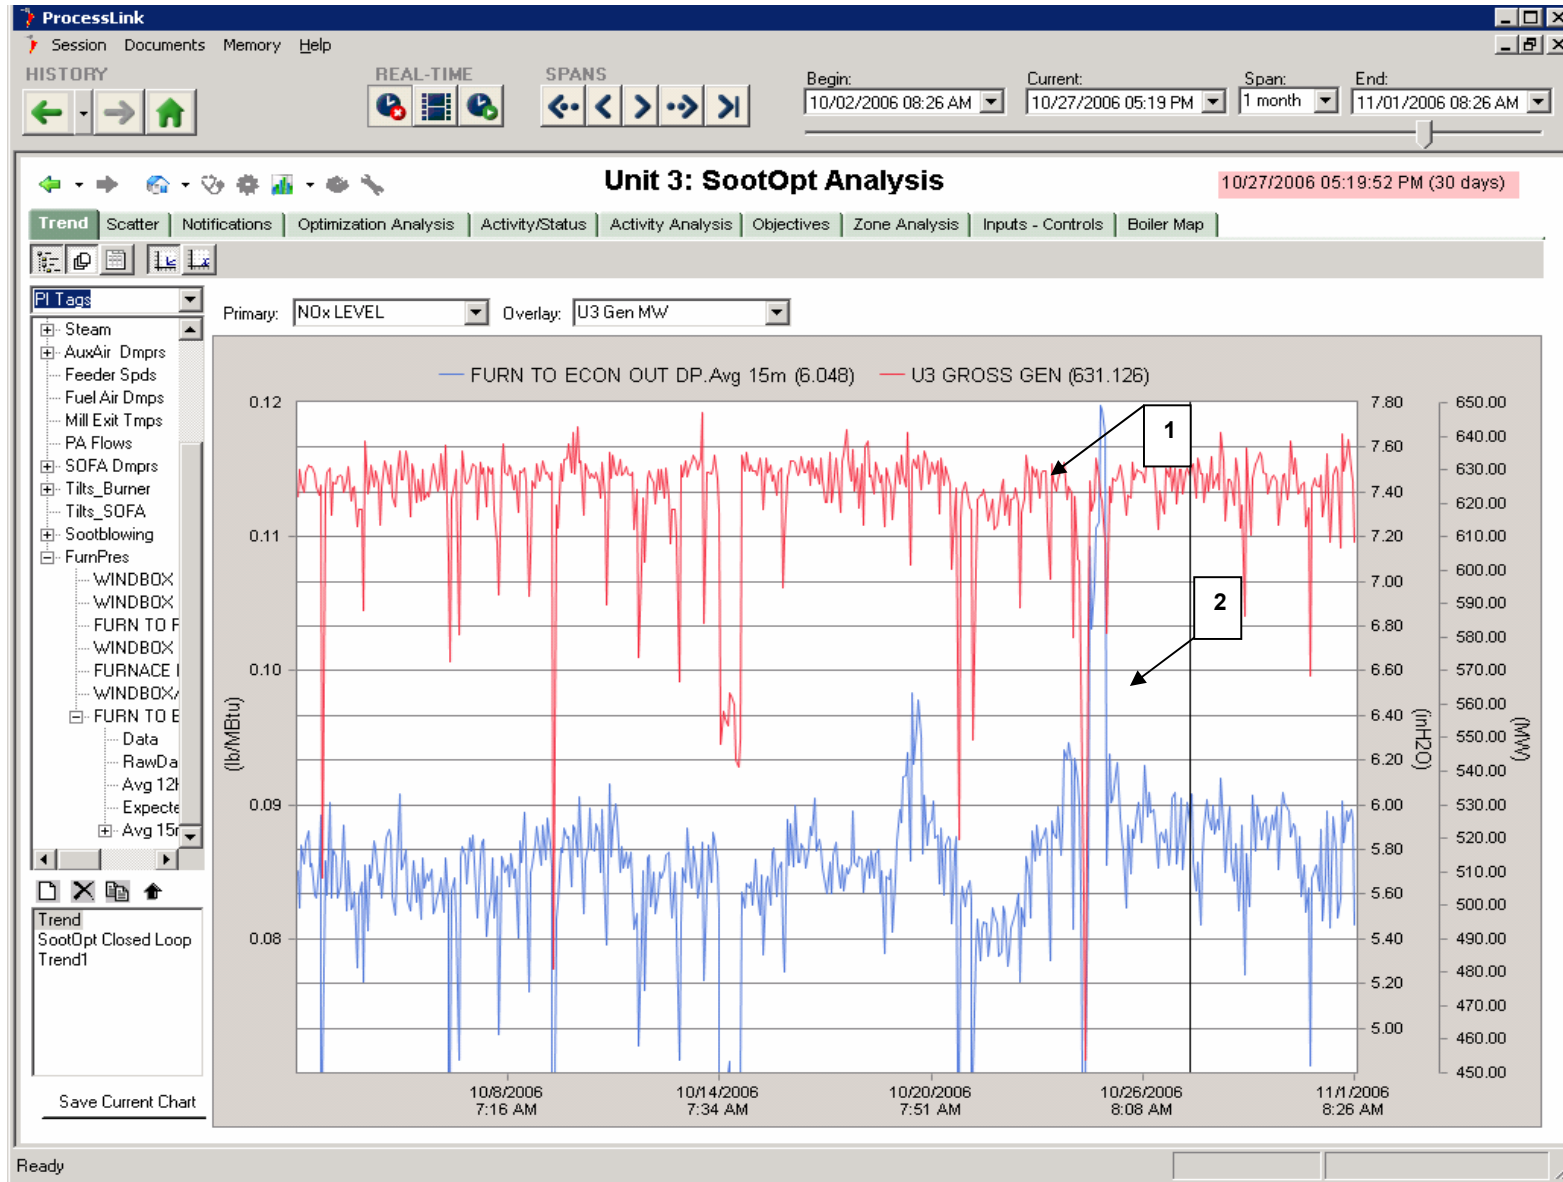


Figure 74 SootOpt Analysis Screen Showing Furnace-Economizer DPs

4.5.1.4 Condenser Cleanliness

PerformanceOpt's detailed online thermal plant model not only alerts plant personnel to discrepancies between actual and achievable efficiency and capacity; it also provides important analytical support, enabling the investigation of different operating scenarios. For instance, PerformanceOpt alerted BEC to a reduction in Unit 1 condenser efficiency, which was due to degrading condenser cleanliness (see Figure 75). This was beginning to have an impact on heat rate and potentially maximum unit load. BEC's approach has been to clean the condenser once a year, but with PerformanceOpt, NeuCo and BEC were able to examine the impact of heat rate and MW capacity in more depth and determine a cost/benefit analysis of cleaning once versus multiple times per year.

Two benefits methods were used:

- Efficiency Losses: Assumes improved cleanliness is used to reduce fuel input to achieve monthly loads
- Revenue Losses: Assumes improved cleanliness is used to generate additional capacity at full load and reduced fuel input at lower loads.

Ultimately it was determined that there was little total cost difference between cleaning one, two or three times per year, supporting BEC's existing cleaning strategy. For units without the benefit of sufficient experienced staff to undertake a detailed manual evaluation of these impacts using available historical data (as BEC had done), the correctness of this conclusion could not easily be verified. The economic consequences of being right here are significant. And for units with a different load and constraint profile the best decision may well have been different.

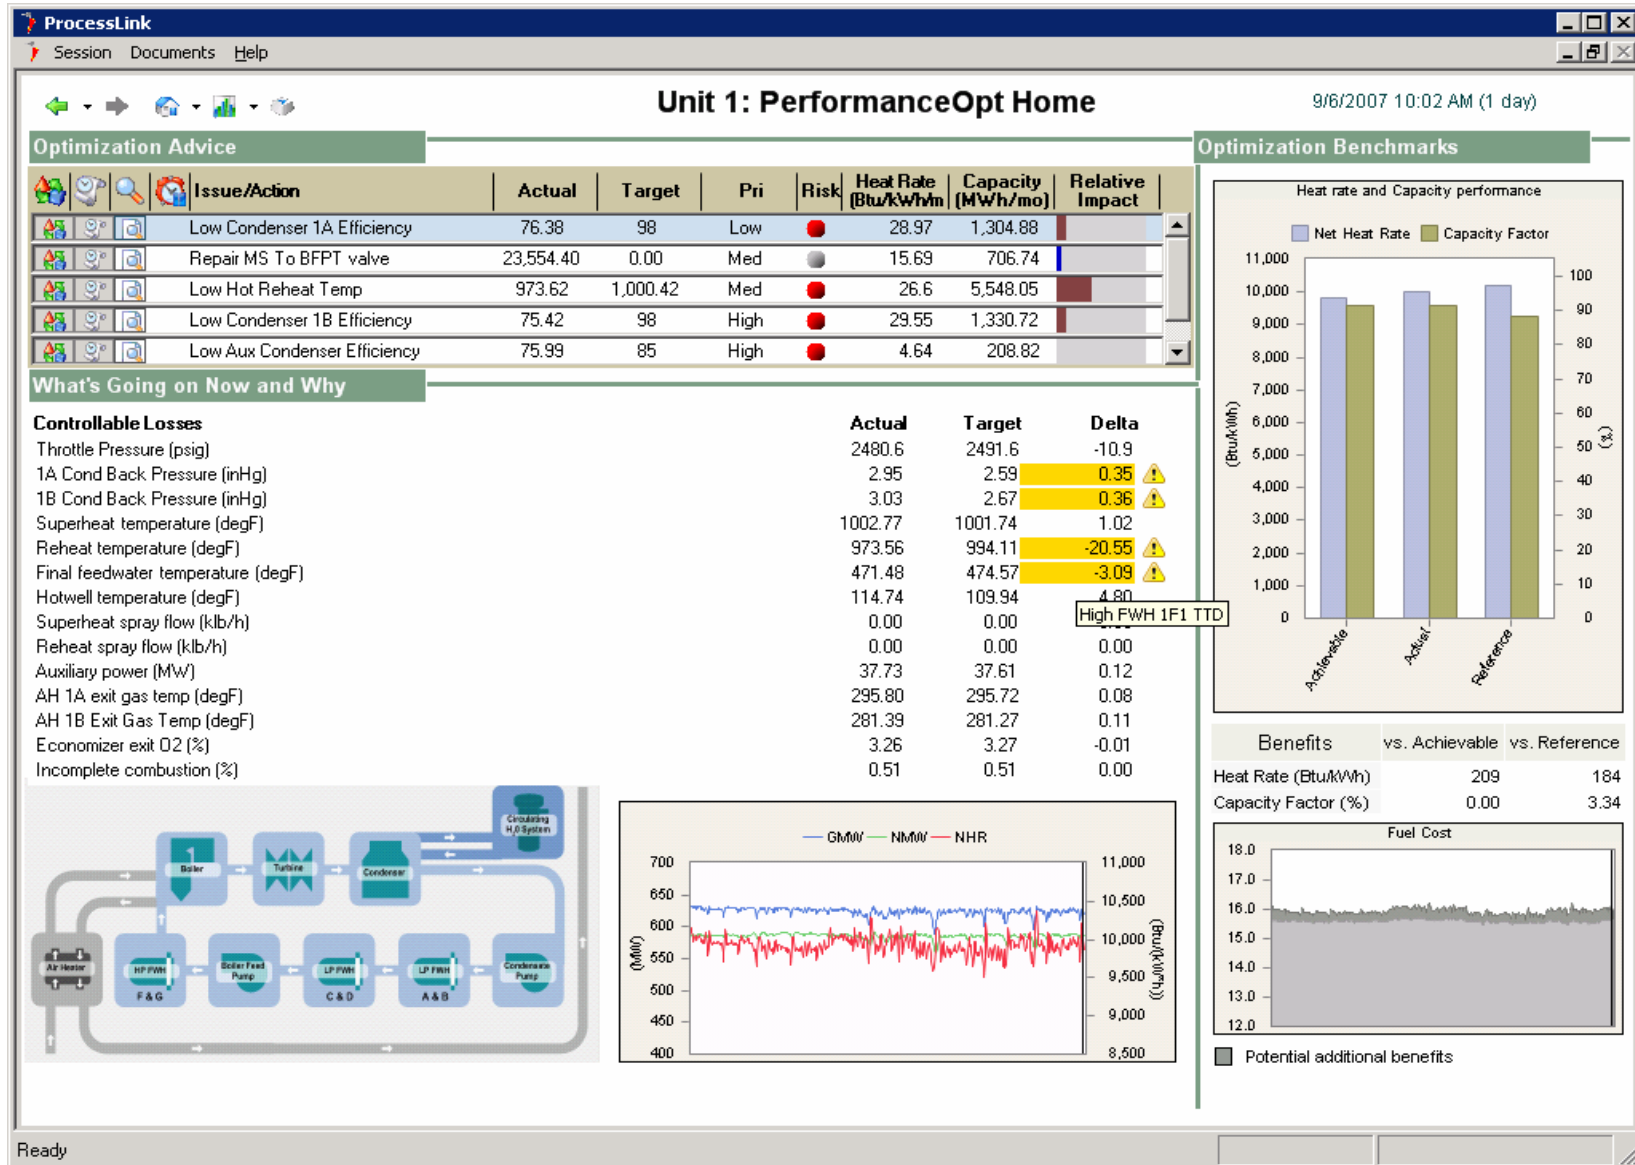


Figure 75 PerformanceOpt alerts BEC to low condenser cleanliness and the associated impacts

4.5.1.5 Cyclone Flame Scanner Quality

As discussed previously, the cyclone units at Baldwin are very challenging to operate in a low NO_x mode without losing good cyclone slag flow. One of the strategies NeuCo and BEC have adopted is to maintain the flame scanner quality signal, the closest proxy for temperature, above 65% while stoichiometry is reduced. To relieve some of the pressure on operations to monitor the health of the 28 cyclones on Units 1 and 2, NeuCo has incorporated triggers within the CombustionOpt advisory system to alert users when cyclone flame scanner qualities drop below this threshold, indicating that they could be in danger of slagging up. Warnings such as this are used to alert BEC that action needs to be taken in order to prevent the cyclones from filling up with hardening slag, a condition so severe that it often requires an outage to remedy. Other cyclone triggers let the Operations staff know when one or more of the coal feeders is either below its design rating low limit or above a safe high rating limit. Both of these conditions are operational inputs to cyclone health since the mass flow profile across these units is on the hairy edge and operating mistakes can be unforgiving.

Figure 76 shows CombustionOpt's Cyclone Flame Quality and Key Parameters analysis, which is designed to help operations staff maintain good situational awareness across the daunting array of information available for the 14 cyclones on Unit 1, providing a visual tool for observing patterns in the dimensions of cyclone cause and effect. Identical context data exists for Unit 2. It is also worth mentioning that this screen, though put together in a few hours by NeuCo, was conceived by plant staff, as being a good screen to have handy, to check in the morning, or pop up when things are going wrong.

In addition to providing tools an engineer can use to visualize possible patterns in cyclone cause and effect, CombustionOpt also provides neural network models of Main Flame Quality for all 14 cyclones. Cyclone Quality is seen as a potential proxy for good slag flow. Historically the 1A5 cyclone has been especially problematic. In Figure 77 the actual (green line) and model predicted (red line) Main Quality show that the model has knowledge about what is going on. This model can then be queried with respect to its sensitivity to cyclone control and state changes.

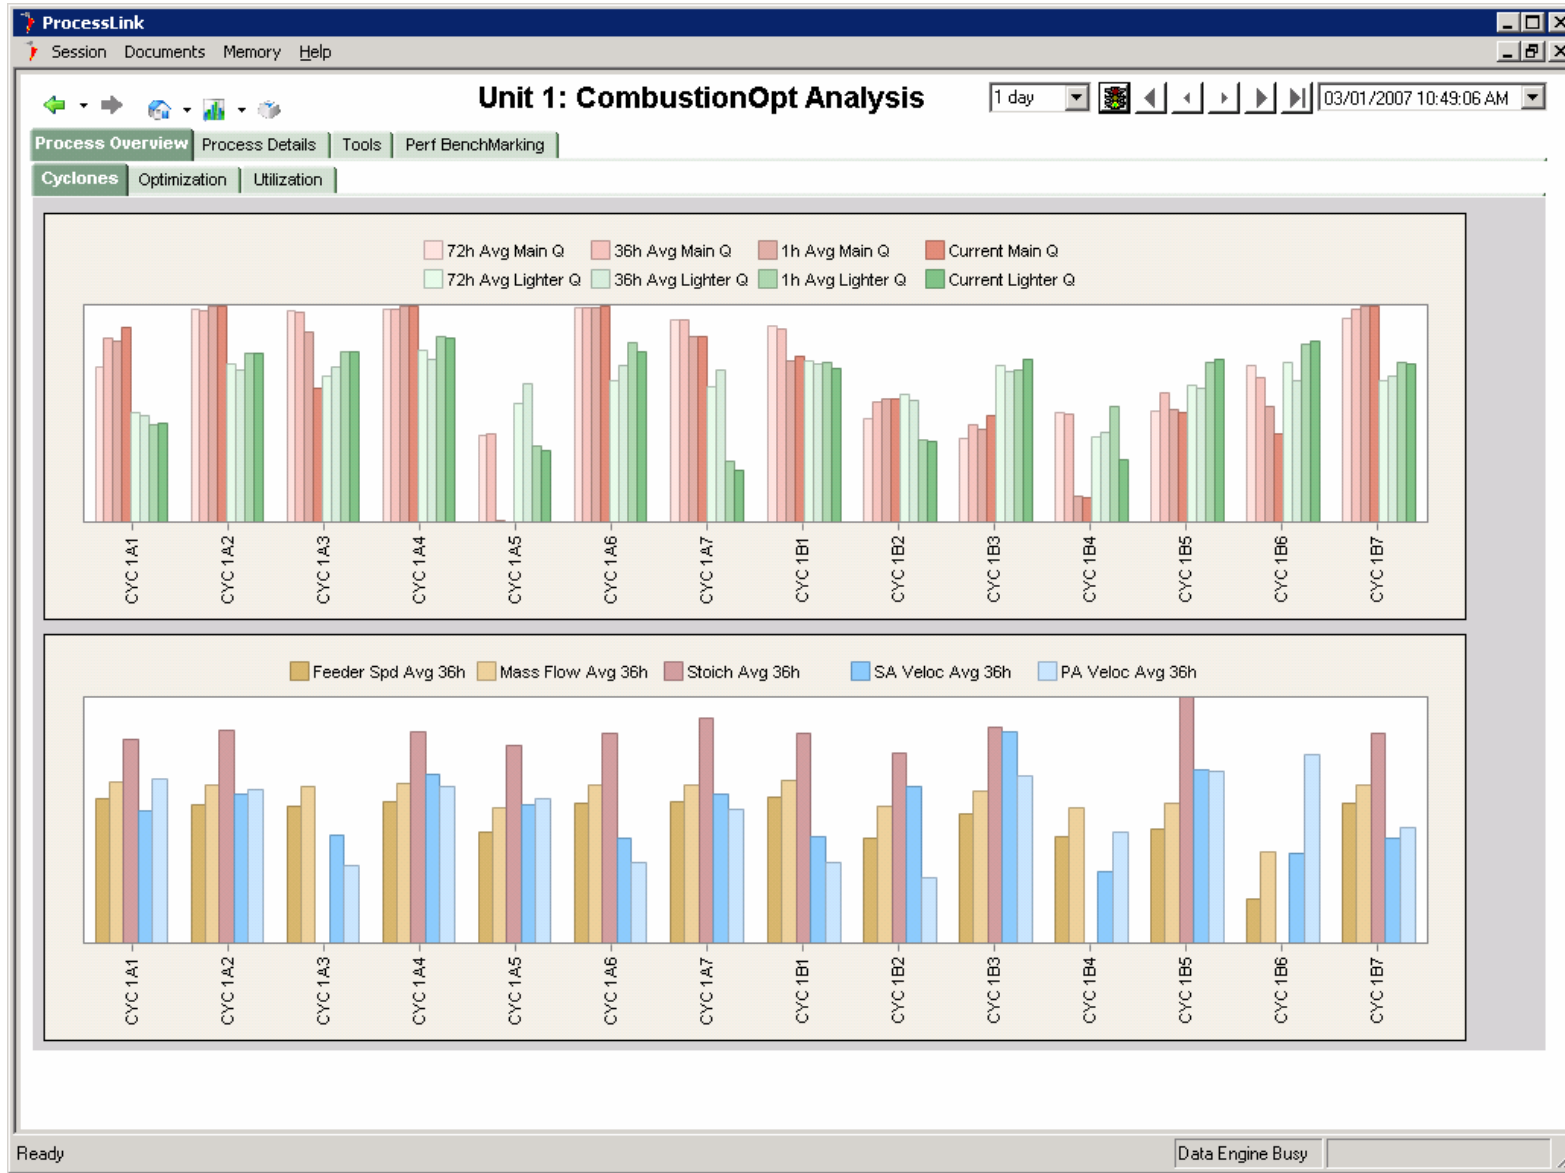


Figure 76 CombustionOpt Analysis Screen Showing Cyclone Flame Quality and Stoichiometries




Figure 77 CombustionOpt Analysis Screen Showing Modeled Flame Quality

4.5.2 Summary of Data Analysis

Unit 1		2005 vs 2007 MW>600 (Meas) 2007 vs 2007 MW>600 (Model)				2005 vs 2007 MW>550 Meas 2007 vs 2007 MW>550 (Model)			
		Control	Exp	Delta	% Delta	Control	Exp	Delta	% Delta
HR Meas	Btu/kWh	9931.15	9881.72	-49.42	-0.50	9938.90	9890.54	-48.36	-0.49
HR Model	Btu/kWh	9961.94	9881.98	-79.97	-0.80	9959.63	9893.57	-66.06	-0.66
NH3 Meas	lb/h	521.26	378.42	-142.84	-27.40	519.81	374.42	-145.39	-27.97
NH3 Model	lb/h	483.54	378.35	-105.18	-21.75	482.69	374.16	-108.53	-22.48
NOx Meas	lb/Mbtu	0.0629	0.0621	-0.0007	-1.1670	0.0624	0.0619	-0.0004	-0.7106
NOx Model	lb/Mbtu	0.0703	0.0618	-0.0086	-12.1675	0.0678	0.0619	-0.0059	-8.6423
Unit 2		2005 vs 2007 MW>600 (Meas) 2007 vs 2007 MW>600 (Model)				2005 vs 2007 MW>550 Meas 2007 vs 2007 MW>550 (Model)			
		Control	Exp	Delta	% Delta	Control	Exp	Delta	% Delta
HR Meas	Btu/kWh	10170.88	10089.72	-81.16	-0.80	10178.01	10082.11	-95.89	-0.94
HR Model	Btu/kWh	10144.35	10051.72	-92.62	-0.91	10176.82	10053.78	-123.05	-1.21
NH3 Meas	lb/h	511.68	467.70	-43.98	-8.60	505.77	468.72	-37.05	-7.33
NH3 Model	lb/h	480.99	471.05	-9.94	-2.07	473.32	468.98	-4.34	-0.92
NOx Meas	lb/Mbtu	0.0650	0.0611	-0.0039	-6.0146	0.0665	0.0613	-0.0052	-7.8164
NOx Model	lb/Mbtu	0.0634	0.0616	-0.0017	-2.7011	0.0617	0.0615	-0.0001	-0.1970
Unit 3		2005 vs 2007 MW>600 (Meas) 2007 vs 2007 MW>600 (Model)				2005 vs 2007 MW>550 Meas 2007 vs 2007 MW>550 (Model)			
		Control	Exp	Delta	% Delta	Control	Exp	Delta	% Delta
HR Meas	Btu/kWh	10020.83	10052.59	31.76	0.32	10029.73	10078.59	48.86	0.49
HR Model	Btu/kWh	10147.14	10116.20	-30.94	-0.30	10180.22	10155.78	-24.45	-0.24
NOx Meas	lb/Mbtu	0.0954	0.0909	-0.0045	-4.7224	0.0953	0.0912	-0.0042	-4.3559
NOx Model	lb/Mbtu	0.0943	0.0907	-0.0036	-3.7879	0.0944	0.0900	-0.0044	-4.6198
		MW>600				MW>550			
		% Delta				% Delta			
		Avg HR Delta Meas -0.32622				AvgHR Delta Meas -0.31385			
		AvgHR Delta Model -0.67356				AvgHR Delta Model -0.70415			
		Avg CEMS NOx Delta Meas -3.96802				Avg CEMS NOx Delta Meas -4.29428			
		Avg CEMS NOx Delta Model -6.21882				Avg CEMS NOx Delta Model -4.48637			
		Avg SCR NH3 Delta Meas -11.9994				Avg SCR NH3 Delta Meas -11.7653			
		Avg SCR NH3 Delta Model -7.94				Avg SCR NH3 Delta Model -7.80023			
		Avg Furn NOx Delta Meas -15.9674				Avg Furn NOx Delta Meas -16.0595			
		Avg Furn NOx Delta Model -14.1588				Avg Furn NOx Delta Model -12.2866			

Table 1 Summary of Data Analysis

5 Managing Tradeoffs and Estimating Optimization Impact

To quantify how much true optimization is being delivered where there are trade-offs between different objectives, it is useful to re-frame the measured effect in terms of overall optimization leverage, factoring out the direction the actual optimization may have been biased (for instance, toward more NO_x benefit or toward more HR for example).

The balancing and prioritizing of objectives is a key optimization challenge that CombustionOpt, SCR-Opt, and SootOpt address directly – through actual cost metrics where possible. In some cases however, concerns that are not easily quantifiable as costs (such as keeping a safe margin from a 30day average NO_x ceiling, given some uncertainty around how many high NO_x startups might have to be accommodated), may factor into the assignment of priority. At BEC the plant's desire to maintain a safe margin of error with respect to a 30 day average NO_x rate cap caused them to prioritize NO_x reduction over Heat Rate Improvements. The high risk penalty associated with violating the cap made this a logical tactic.

5.1 Visualizing Tradeoffs

As a way of visualizing this in action we can look at the following example (which shows a set of multivariable models (used in the sections above) of NO_x and HR, run through a scenario engine to get an estimate of what their trends would have looked like if SootOpt and/or CombustionOpt had not been used.

In Figure 78 and Figure 79 we see the model predicted NO_x & HR (blue) and actual NO_x & Heat Rate (green) for three multivariable neural network models. The models relate these objectives to their major causal factors and disturbances (MW, Ambient Conditions, and Condenser Backpressure).

For Both Figures:

The top trend is of a model that has seen data only where no optimization occurred.

- Predicted NO_x (based on this model's "experience") is higher than actual
- Predicted HR is fairly consistent with actual.
- This is a reasonable estimate of what the NO_x & Heat Rate trends would have looked like if neither CombustionOpt or SootOpt had been used on Unit 3 (in reality they were mostly enabled)

The middle trend shows the model that has only seen data with the effect of CombustionOpt only.

- Predicted NO_x is lower than the prediction in the top trend, but still higher than actual
- Predicted Heat Rate however is higher than actual
- This is a reasonable estimate of what the NO_x trend would have looked like if only CombustionOpt had been used.

The bottom trend shows the prediction of the model that only has seen data affected by both CombustionOpt and SootOpt

- Predicted NO_x is in line with actual, though in places the actual is higher (these represent conditions outside of its experience meaning one or both CombustionOpt and SootOpt are not enabled).

- Predicted Heat Rate is also back in line with actual, compared to the CombustionOpt only trend in the middle.
- Clearly this is the model best able to predict what NOx and Heat Rate.

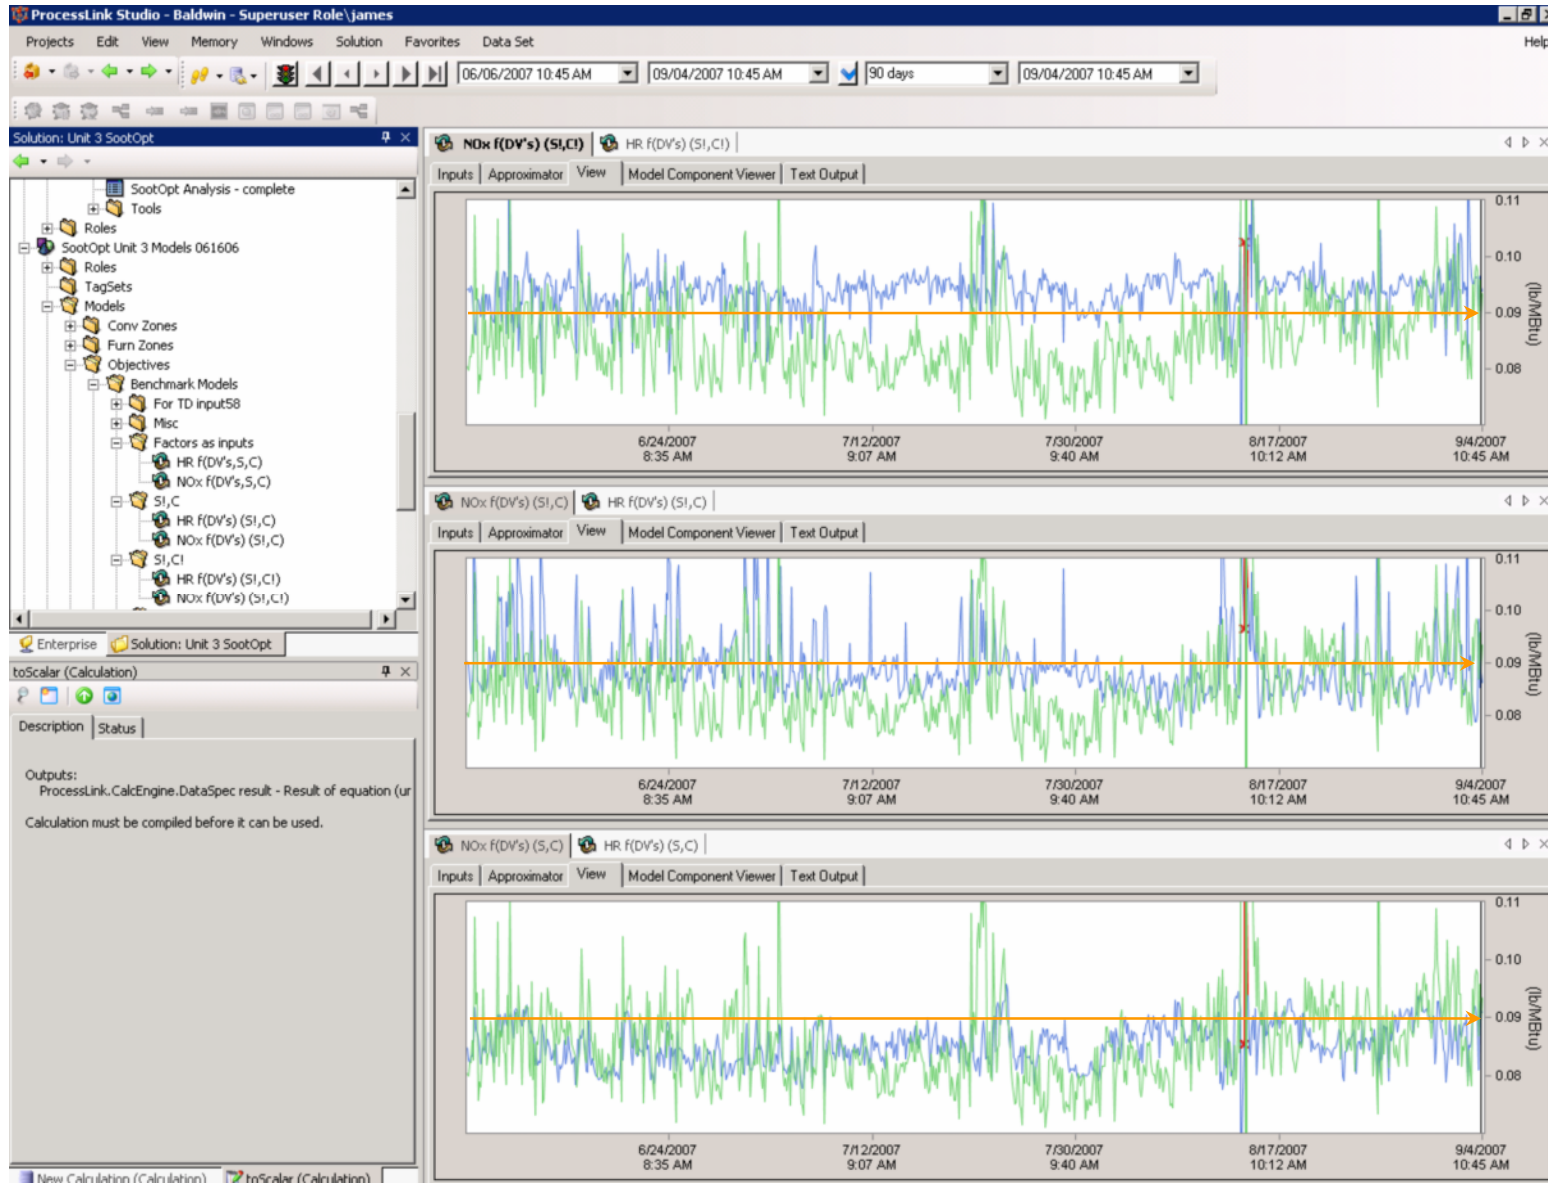


Figure 78 $NO_x=f(DV's); (SI,CI), (SI,C),(S,C)$

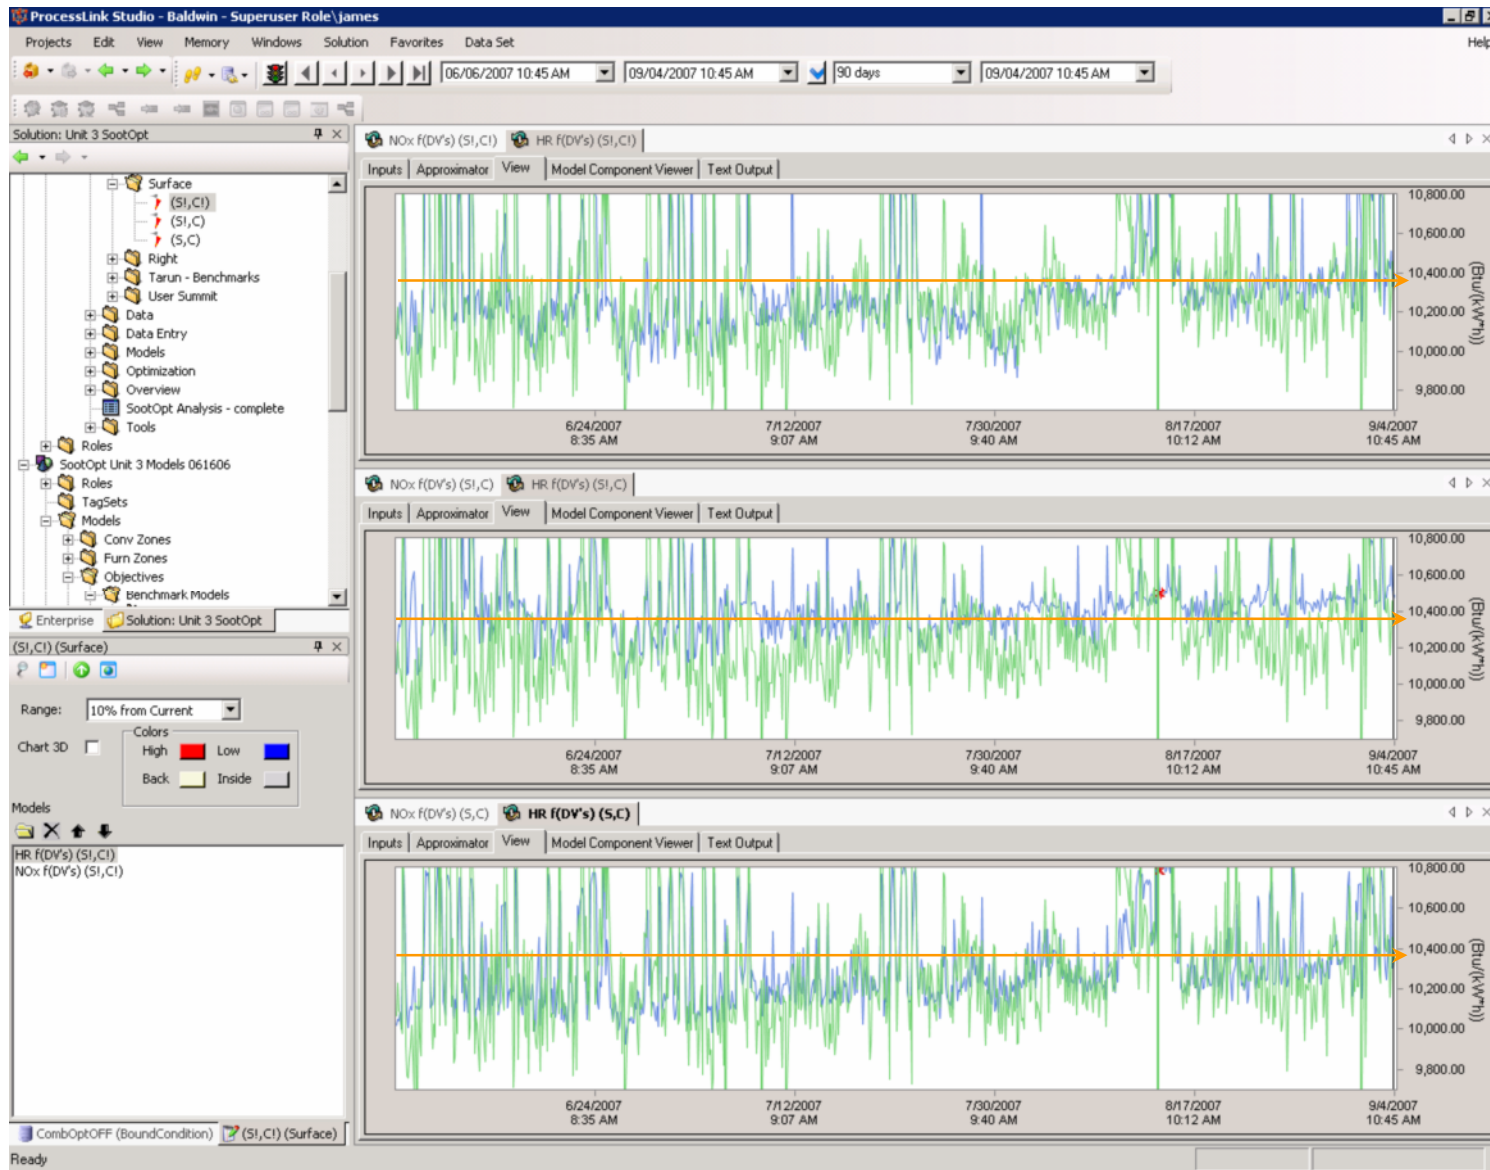
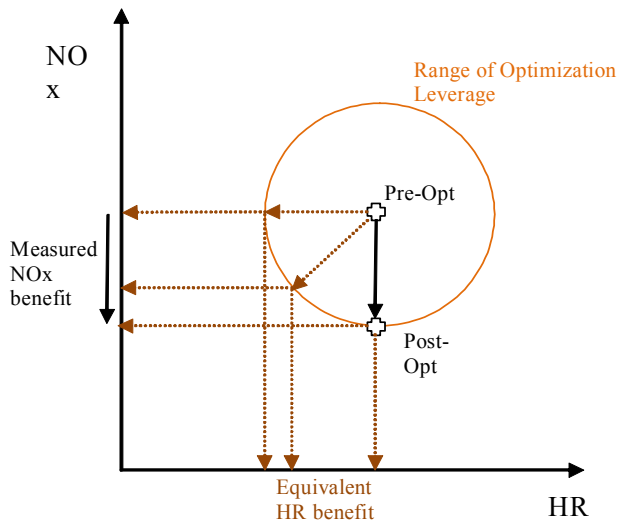


Figure 79 $R=f(DV's); (SI,C), (SI,C),(S,C)$

In the trends above it might appear at first that only NOx optimization occurred and that there was no net improvement in Heat Rate.

But with only the observed net optimization in NOx, without any final HR penalty and the assumption that all alternative scenarios are of a trade-off nature (NOx vs HR) we can visualize the benefit in the following way:



Based only on the assumption that no additional optimization benefit (other than what was measured) is available, and all further changes would involve a trade-off, if priority were altered, any combination of NOx and HR benefit on the circle could be achieved.

In summary the total Optimization leverage achieved is the net improvement in conserved dimensions. How the improvement was distributed between those two dimensions is a function of the priority placed on each – something that is fundamentally adjustable. At BEC it was the desire of plant staff that NOx reduction be placed above Heat Rate in terms of priority.

6 Conclusion

6.1 Estimate of Benefits

With respect to the goals set out for the project, and looking at the data for the period over which it was undertaken (using the multivariable regression-based estimation approach to adequately account for disturbance effects), NeuCo summarizes the integrated optimization project's benefits as follows:

MW>600	% Delta	MW>550	% Delta
Avg HR Delta Meas	-0.32622	AvgHR Delta Meas	-0.31385
AvgHR Delta Model	-0.67356	AvgHR Delta Model	-0.70415
Avg CEMS NOx Delta Meas	-3.96802	Avg CEMS NOx Delta Meas	-4.29428
Avg CEMS NOx Delta Model	-6.21882	Avg CEMS NOx Delta Model	-4.48637
Avg SCR NH3 Delta Meas	-11.9994	Avg SCR NH3 Delta Meas	-11.7653
Avg SCR NH3 Delta Model	-7.94	Avg SCR NH3 Delta Model	-7.80023
Avg Furn NOx Delta Meas	-15.9674	Avg Furn NOx Delta Meas	-16.0595
Avg Furn NOx Delta Model	-14.1588	Avg Furn NOx Delta Model	-12.2866

Table 2 Sub-set of Table 1, showing the average delta benefits for generated power >600 MW (left side) and for generated power >550 MW (right side)

- **NOx Reduction:**

- The 5% target for NOx reduction was exceeded with average CEM and SCR Inlet (furnace) NOx reduction of between 12% and 14%, as indicated by the green highlighted sections in Table 2 ("Average Furnace NOx Delta Model").

- **Heat Rate Improvement:**

- In addition to the heat rate improvement provided by the closed-loop operation of CombustionOpt, SCR-Opt and SootOpt, PerformanceOpt and MaintenanceOpt provided prioritized advice and diagnostic support. Together the systems delivered an average Heat Rate improvement of between 0.67% and 0.7%, as indicated by the green highlighted sections in Table 2 ("AvgHR Delta Model")
 - This falls short of the 1.5% heat rate improvement target largely because BEC prioritized Cyclone Stability (availability) and CEMS and SCR Inlet NOx over Heat Rate in the event they needed to be traded-off with one another, resulting in almost double the target NOx reduction but less than targeted Heat Rate improvement. A different prioritization of objectives could be used to drive a different balance between them with total combined leverage meeting the target of 1.5% improvement. See section 5 "Managing Tradeoffs and Estimating Optimization Impact".
 - Several factors may be masking greater Heat Rate improvements:
 - The decrease in fuel energy density (in addition to that seen with the switch to PRB) over the course of the project is a potentially significant contributor to degraded baseline Heat Rate, masking larger improvements. Fuel is one disturbance variable that is difficult to exclude because there are few indicators of its variability at this site.
 - "Actions" taken by the Optimizers in open loop, in the form of processing of the prioritized alerts, by plant and NeuCo engineers, deliver a benefit that is hard to quantify.

- **Increased Annual Available MWh:**
 - Although difficult to measure precisely, the target of increasing available MWh's by 1.5% was met in the following ways:
 - Providing prioritized alerts and knowledge-based diagnostics for a wide array of plant equipment and process anomalies
 - Helping the plant to move from high sulfur, high Btu Illinois coal to PRB and run that fuel at low stoichiometries. Typical permanent de-rates for units switching to PRB range from between 10 and 20 MWs. At a given capacity factor this works out to between 1 and 2% of unit production capacity.
 - Improved management of cyclone flame quality, and improved vigilance with respect to cyclone conditions avoided some degree of temporary de-rate due to cyclone slag build up. This improvement in was essential in facilitating the fuel switch sans derate.

- **Commensurate Reductions in Greenhouse Gases, Mercury, and Particulates:**
 - Reductions in all three of these indices can be associated directly with the optimization leverage observed in the Heat Rate and NOx reductions.

- **Commensurate Increases in Profitability from Lower Costs, Improved Reliability, and Greater Commercial Availability:**
 - Commensurate improvements in costs, reliability and availability project from the benefits detailed above.
 - Sustained operation of the cyclones using more available, less expensive but off-design fuel, at very low stoichiometry with correspondingly low NOx and SOx production levels.
 - More effective catalytic reduction of NOx, effectively increasing the reduction of NOx per unit of capital investment in SCR.
 - Reduced time required to discover, prioritize and diagnose plant equipment issues.

The gross improvement in the ability of this fossil-fired plant to deliver cost effective energy while steadily becoming cleaner and more efficient is a testament to the long term commitment the Dynegy team has made to the effort, and to the variety of techniques and technologies applied.

The data analyzed in this report show the important role that Optimization plays in this effort, delivering a measurably high rate of return per unit cost.

6.2 Economic Implications

This section summarizes the results of an engineering-economics benefits analysis applied to the BEC results and also to the entire US fleet of fossil-fired generating units, assuming the technology benefits demonstrated at BEC are broadly applicable.

6.2.1 BEC Economic Benefits

The operating and cost assumptions as well as the economic benefits achieved at the three BEC units are shown in Table 3. Note that while BEC itself is not yet participating in a liquid CO₂ trading market, many units are about to be affected by the initial auction for the Regional Greenhouse Gas Initiative (RGGI), some other

generators are participating in the Chicago Climate Exchange (CCE), and many others (including BEC) are in states that have or are in the process of forming other multi-state regional initiatives, such as the Western Climate Initiative and the Midwest Accord. Since there is not yet a mandatory federal cap and trade program for CO₂, the April 2008 average of current 2009 forward trading prices for the CCE (\$7.00/ton) and RGGI (\$6.00/ton) was used to derive the \$6.50/ton value of CO₂ reduction used for this analysis.

The current value of NO_x reduction at BEC is a complex calculation, and is determined by the New Source Review Consent Decree under which the plant is currently operating. Since the actual cost of NO_x factors is complex and depends on MISO wholesale power market conditions, the value of NO_x reduction at BEC, for the purposes of this report, is based on the sum of projected seasonal and annual NO_x allowance values for January of 2009.

BEC's fuel costs and actual unit heat rate values factor into Dynegy's operation of the plant within MISO and are thus considered proprietary. For the purposes of this report, a heat rate of 10,000 Btu/kWh was used for all three units, along with a \$1.50/mmBtu for delivered coal. These values are representative of pulverized coal-fired units burning Powder River Basin Coal in this general geographic area. The reported heat rate gains should be viewed in the context of coal quality (specifically heating value or Btu content), which was steadily declining over the entirety of the four-year project. A decrease in energy density is a potentially significant contributor to degrading baseline heat rate and is likely masking larger heat rate benefits provided by optimization.

As noted earlier in the report, the Optimizer benefits with respect to reliability and commercial availability are difficult to precisely quantify, however they clearly provided a substantial benefit as indicated in the following examples:

- SootOpt reduced the number of cleaning actions. We cannot directly relate this reduction in boiler cleaning actions to a reduction in water wall erosion and associated tube rupture outages. That said, there is no doubt that the 33 percent reduction of boiler cleaning actions (via water cannons and soot blowers) can be expected to help reduce such outages, which are the largest contributor to forced outage rates at all coal-fired plants.
- It is also very likely that the improved stoichiometry and flame quality control achieved at the two cyclone boilers resulted in fewer slagging events than would have occurred in the absence of CombustionOpt.

Given the difficulties in precisely quantifying the impacts of these operational benefits on reliability (EFOR) and commercial availability, we have conservatively estimated average availability improvements across the three units at 0.75%. While this is only half the 1.5% availability impact projected in NeuCo's CCPI proposal for the project, the associated annual benefits nonetheless range from \$1.3 to \$1.4 million dollars per-unit, representing more than a \$4 million annual benefit for the plant.

Dynegy	Baldwin Unit 1	Baldwin Unit 2	Baldwin Unit 3	Baldwin Total Plant	
1	Gross Capacity (MW)	640	640	630	1,910
2	Net Capacity (MW)	600	600	600	1,800
3	Capacity Factor (%)	85%	85%	90%	86.7%
4	Annual Output (MWh/yr)	4,467,600	4,467,600	4,730,400	13,665,600
5	Boiler Type	Cyclone	Cyclone	CE-T	Cyclones/CE-T
6	Baseline Heat Rate (Btu/kWh)	10,000	10,000	10,000	10,000
7	Annual Heat Input (mmBtu/yr)	44,676,000	44,676,000	47,304,000	136,656,000
8	Fuel Cost (\$/MMBtu)	\$1.50	\$1.50	\$1.50	\$1.50
10	PRB Coal (%)	100%	100%	100%	100%
13	Annual Coal Consumption (tons/yr)	2,628,000	2,628,000	2,782,588	8,038,588
14	Weighted Average Carbon Content	70.0%	70.0%	70.0%	70.0%
15	CO2 Output (tons/yr)	6,745,200	6,745,200	7,141,976	20,632,376
16	Annual Fuel Cost (\$/yr)	\$71,481,600	\$71,481,600	\$74,503,800	\$217,467,000
17	Heat Rate Improvement (-%)	-0.66%	-1.21%	-0.24%	-0.70%
18	Annual Fuel Savings	\$471,779	\$864,927	\$178,809	\$1,515,515
19	Value of CO2 reduction (\$/ton)	\$6.50	\$6.50	\$6.50	\$6.50
20	Annual CO2 Reduction (tons/year)	47,486	87,058	17,998	154,161
21	Annual CO2 Reduction Benefits	\$289,369	\$530,510	\$111,415	\$931,294
22	Baseline Average Boiler NOx (lb/MMBtu)	0.45	0.45	0.12	0.34
23	Baseline Annual NOx (tons/yr)	10,722	10,722	2,980	24,425
24	Average ProcessLink NOx Reduction, at boiler (-%)	-0.71%	-7.81%	-4.36%	-4.29%
25	SCR for Benefits calculations? (Yes/No)	Yes	Yes	No	Units 1-2
26	Nominal SCR-related NOx reduction (%)	90%	90%	0%	90%
27	NH3 Reduction (%)	-27.97%	-7.33%	0.00%	-17.65%
28	Average NOx Allowance Credit Value (\$/ton)	\$1,844	\$1,844	\$1,844	\$1,844
29	Ozone Season (Months)	12	12	12	12
30	NOx Reduction Allowance Benefits (\$/yr)	\$14,036	\$154,397	\$239,347	\$407,780
31	NH3 Cost (\$/ton NOx)	\$350	\$350	\$0	\$350
32	NH3 Reduction Value (\$/yr)	\$1,049,654	\$293,092	\$0	\$1,342,746
33	FGD for Benefits calculations? (Yes/No)	No	No	No	No
34	SOx Reduction Allowance Benefits (\$/yr)	\$35,855	\$65,734	\$13,589	\$115,179
35	Annual Availability Increase (%)	0.75%	0.75%	0.75%	0.75%
36	Increased Availability Value (\$/MWh)	\$55.00	\$55.00	\$55.00	\$55.00
37	Increased Availability Value (\$/yr)	\$1,306,773	\$1,306,773	\$1,392,512	\$4,006,058
	Total ProcessLink Suite Savings (\$/yr)	\$2,878,096	\$2,684,924	\$1,824,257	\$7,387,278
	Total ProcessLink Suite Savings w CO2 (\$/yr)	\$3,167,466	\$3,215,434	\$1,935,672	\$8,318,572

Table 3 Economic benefits achieved at BEC

As reflected by the numbers in Table 3, the total annual dollar value of the benefits associated with the products installed, refined, and commercialized at BEC are large, ranging from \$1.8 to \$3.2 million dollars per-unit; and \$7.2 to \$8.1 million dollars per year plant-wide depending on whether CO₂ benefits are included.

6.2.2 Economic Benefits as Applied to US Fossil Generation

The benefits achieved at BEC were extrapolated to the US fossil generation industry, as shown in Table 4. The table values come from a variety of sources: capacity and capacity factors from the 2005 UDI North American Fossil Generation data base; baseline NO_x values and SCR and FDG installations from McIlvaine Company; and baseline heat rate and fuel costs based on observations in the field.

It is worth noting that BEC is widely regarded as among the best operating fossil generating units in North America; All three of the units exhibit among the lowest boiler NO_x observed anywhere among boilers of similar design and fuel (i.e. cyclone and t-fired boilers burning 100% PRB coal). All three of the units were equipped with state-of-the-art instrumentation and controls (including a recent Emerson Ovation DCS) prior to the commencement of NeuCo's CCPI project at BEC. Extrapolating the results from units with such strong pre-project performance and state-of-the-art I&C equipment, to the current US industry, can be viewed as conservative given the preponderance of existing US units with more "low-hanging fruit."

Note that the oil and gas category includes both traditional steam turbine units and well as combined cycle plants. Neither SootOpt nor CombustionOpt were assumed to be included in the analysis of benefits as applied to combined cycle units, and SootOpt was not included in oil or gas-fired units. The 0.70 percent aggregate heat rate improvement gain demonstrated at BEC was used for all unit types but could be considered a conservative estimate for the following reasons:

- 1) NeuCo's experience applying CombustionOpt to oil and/or gas-fired units has consistently demonstrated larger benefits;
- 2) the complexities and interdependencies inherent to a combined cycle unit are such that NeuCo and its partners with domain expertise in combined cycle operations believe that heat rate gains for PerformanceOpt and MaintenanceOpt for these types of plants will likely be well in excess of one percent.

US Fossil Units	Typical PRB W/SCR	Typical PRB No/SCR	Typical Bitum W/SCR	Typical Bitum No/SCR	Oil/Gas (ST + CCCT)	Total (1950 Units) Industry Benefits
Gross Capacity (MW)	698	246	698	246	198	514,359
Net Capacity (MW)	645	227	645	228	192	475,782
Capacity Factor (%)	90%	80%	90%	80.0%	40.0%	82.7%
Annual Output (MWh/y)	5,089,038	1,592,075	5,089,038	1,594,670	672,768	3,446,769,199
Baseline Heat Rate (Btu/kWh)	10,000	10,000	10,000	10,000	10,000	10,000
Annual Heat Input (mmBtu/yr)	50,890,382	15,920,747	50,890,382	15,946,704	6,727,680	34,467,691,992
Fuel Cost (\$/MMBtu)	\$1.50	\$1.50	\$2.75	\$2.75	\$7.50	\$2.38
CO2 Output (tons/yr)	7,683,450	2,403,721	5,382,637	1,686,671	1,076,429	6,379,454,175
Annual Fuel Cost (\$/yr)	\$82,524,944	\$25,817,428	\$151,295,730	\$47,409,120	\$50,457,600	\$82,126,816,678
Heat Rate Improvement (-%)	-0.70%	-0.70%	-0.70%	-0.70%	-0.70%	-0.70%
Annual Fuel Savings	\$581,099	\$181,793	\$1,065,349	\$333,831	\$355,297	\$578,295,980
Value of CO2 reduction (\$/ton)	\$6.50	\$6.50	\$6.50	\$6.50	\$6.50	\$6.50
Annual CO2 Reduction (tons/year)	58,490	18,298	40,975	12,840	7,580	44,920,927
Annual CO2 Reduction Benefits	\$351,670	\$110,018	\$246,362	\$77,198	\$49,268	\$291,986,023
Baseline Average Boiler NOx (lb/MMBtu)	0.25	0.20	0.30	0.28	0.25	0.27
Baseline Annual NOx (tons/yr)	6,877	1,721	8,252	2,370	865	4,942,773
Average ProcessLink NOx Reduction, at boiler (-%)	-21.94%	-21.94%	-21.94%	-21.94%	-21.94%	-21.94%
NH3 Reduction (%)	-19.75%	0.00%	-19.75%	0.00%	0.00%	-19.75%
Average NOx Allowance Credit Value (\$/ton)	\$1,844	\$1,844	\$1,844	\$1,844	\$1,844	\$1,844
NOx Reduction Allowance Benefits (\$/yr)	\$52,809	\$696,242	\$63,371	\$958,894	\$349,999	\$251,515,127
NH3 Cost (\$/ton NOx)	\$350	\$350	\$350	\$350	\$0	\$350
NH3 Reduction Value (\$/yr)	\$475,282	\$0	\$633,709	\$0	\$0	\$248,197,651
SOx Reduction Allowance Benefits (\$/yr)	\$44,164	\$13,816	\$149,149	\$46,736	\$9,965	\$35,677,184
Annual Availability Increase (%)	0.75%	0.75%	0.75%	0.75%	0.75%	0.75%
Increased Availability Value (\$/MWh)	\$55.00	\$55.00	\$55.00	\$55.00	\$100.00	\$57.16
Increased Availability Value (\$/yr)	\$1,480,291	\$463,100	\$964,510	\$302,233	\$126,144	\$1,199,517,320
Total ProcessLink Suite Savings (\$/yr)	\$2,633,645	\$1,354,952	\$2,876,088	\$1,641,695	\$841,406	\$2,313,203,262
Total ProcessLink Suite Savings w CO2 (\$/yr)	\$2,985,315	\$1,464,970	\$3,122,450	\$1,718,893	\$890,674	\$2,605,189,284

Table 4 Economic Benefits as Applied to US Fossil Generation

As the numbers in Table 4 indicate, the benefits available to the industry based on the results achieved at BEC are between \$2.3 and \$2.6 billion dollars per year in annual savings across the full combination of unit types, fuel sources, and post-combustion controls characterizing the current US fossil generation fleet. These aggregate benefits are distributed across the categories of fuel efficiency, NOx reduction, reagent costs, CO₂ emissions, and commercial availability.

This estimate of available benefits represents a compelling financial incentive when viewed relative to the current costs of the four products that have been commercialized as part of this project. Table 5 shows the payback in months for the combination of CombustionOpt, SootOpt, MaintenanceOpt, and PerformanceOpt as they pertain to the categories of unit types and fuel sources in the US fleet. The product costs used include all software licenses, installation services, variable expenses (travel, living, computers, etc.) and one year of Annual Maintenance and Support.

Simple Payback for Commercial Products	PRB W/SCR	PRB No/SCR	Bitum W/SCR	Bitum No/SCR	(ST + CCCT)
Commercial Payback Excluding CO2 (Months)	4.38	8.51	4.01	7.02	9.45
Commercial Payback Including CO2 (Months)	3.86	7.87	3.69	6.71	8.93

Table 5 Estimated payback in months for an installation of CombustionOpt, SootOpt, MaintenanceOpt, and PerformanceOpt (a.k.a. Suite of optimizers)

The Suite of four integrated optimizers commercialized as part of this project are expected to yield well under a one-year payback for average sized units across all unit types and fuel categories comprising the US fossil power industry. This represents a highly cost-effective way of addressing some of the industry's most pressing challenges - one that complements and enhances a wide variety of other methods such as SCR, alternative fuels, low NOx systems, and modern control and instrumentation systems.

7 References (N/A)

8 List of Acronyms and Abbreviations

AI	Artificial Intelligence
APH	Air Pre-Heater
API	Application Programming Interface
ASI	Applied Synergistics Inc.
BEC	Baldwin Energy Complex
BFP	Boiler Feed Pump
BFPT	Boiler Feed Pump Turbine
BTU	British Thermal Unit
B&V	Black & Veatch
CCPI	Clean Coal Power Initiative
CEMS	Continuous Emissions Monitoring System
CMMS	Condition Monitoring Maintenance System
CO	Carbon Monoxide
CO ₂	Carbon Dioxide
DCA	Drain Cooler Approach
DCS	Distributed Control System
DMG	Dynegy Midwest Generation
DOE	Department of Energy, and Design of Experiments
EEGT	Economizer Exit Gas Temperature
EPRI	Electric Power Research Institute
ESP	Electro Static Precipitator
FCM	(ASI's) Furnace Cleanliness Module
FD	Forced Draft
FEGT	Furnace Exit Gas Temperature
FF	Functional Failure
FGD	Flue Gas Draft
FWH	Feedwater Heater
GUI	Graphical User Interface
HMI	Human Machine Interface
HR	Heat Rate
ID	Induced Draft
IP	Intermediate pressure

ISB	Intelligent Soot blowing
LOI	Loss on Ignition
mmBTU	Millions of BTUs
MW	Megawatt
mWh	Megawatt hour
M/year	Million per year
NH ₃	Ammonia
NO _x	Nitrogen Oxides
O ₂	Oxygen
OFA	Over Fire Air
OMU	Owensboro Municipal Utilities
PC	Personal Computer
PI	Plant Information Historian (from OSI Soft)
PID	Proportional, Integral, Derivative (control algorithm)
PLC	Programmable Logic Controller
ppm	parts-per-million
PRB	Powder River Basin
QA	Quality Assurance
RH	Re heater
S	Sulfur
SBCS	Soot Blowing Control System
SCE	ASI's Sootblower Control Expert
SCR	Selective Catalytic Reduction
SH	Super Heater
SNCR	Selective Non-Catalytic Reduction
SO ₂	Sulfur Dioxide
SO ₃	Sulfur Trioxide
SOFA	Separated Over Fire Air
T/C	Thermocouple
TTD	Terminal Temperature Difference
UI	User Interface
USD	U.S. Dollar (\$)
VPN	Virtual Private Network

9 Appendices

A.1 Optimization Advice given by the Optimizers

The following table lists the optimization advice that was given by the Optimizers during 2007 at the Baldwin Energy Complex. The table divides the advice by Unit, then by Optimizer, and then indicates the number of instances of advice for each piece of equipment on the Unit.

Unit 1

CombustionOpt

<u>Equipment</u>	<u># of Instances</u>
Cyclone 1A1	3
Cyclone 1A2	4
Cyclone 1A3	3
Cyclone 1A4	4
Cyclone 1A5	3
Cyclone 1A6	3
Cyclone 1A7	3
Cyclone 1B1	3
Cyclone 1B2	3
Cyclone 1B3	4
Cyclone 1B4	3
Cyclone 1B5	3
Cyclone 1B6	3
Cyclone 1B7	3
OFA System	1

PerformanceOpt

<u>Equipment</u>	<u># of Instances</u>
Aux Condenser	1
Boiler	2
FWH 1A1	1
FWH 1A2	1
FWH 1B1	1
FWH 1B2	1
FWH 1C	1
FWH 1D	1
FWH 1F1	2
FWH 1F2	2
FWH 1G1	2
FWH 1G2	1
HP Turbine	2
Main Condenser 1A	1
Main Condenser 1B	1

MaintenanceOpt

<u>Equipment</u>	<u># of Instances</u>
Air Heater 1A	1
Air Heater 1B	1
APH 1A	2
APH 1B	2
BFP 1A	3
BFP 1B	2
BFPT 1A	3
BFPT 1B	3
Boiler	6
CNDP 1A	2
CNDP 1B	2
CNDP 1C	2
CWP 1A	3
CWP 1B	3
CWP 1C	3
Exciter	2
FD Fan 1A	4
FD Fan 1B	3
FD Fan 1C	3
FD Fan Set	2
FWH 1B1	2
FWH 1B2	2
FWH 1C	2
FWH 1D	2
FWH 1F1	2
FWH 1F2	2
FWH 1G1	2
FWH 1G2	2
General	3
Generator	4
HP Turbine	6
ID Fan 1A	9
ID Fan 1B	9
ID Fan 1C	9
ID Fan Set	2
IP Turbine	9
LP Turbine 1A	5
LP Turbine 1B	5
Main Condenser 1A	1
Main Condenser 1B	1

Unit 2

MaintenanceOpt

<u>Equipment</u>	<u># of Instances</u>
Air Heater 2A	1
Air Heater 2B	1
APH 2A	2
APH 2B	2
Aux Condenser 2A	1
Aux Condenser 2B	1
BFP 2A	5
BFP 2B	5
BFPT 2A	3
BFPT 2B	2
Boiler	5
CNDP 2A	2
CNDP 2B	2
CNDP 2C	2
CWP 2A	3
CWP 2B	3
CWP 2C	3
Exciter	4
FD Fan 2A	4
FD Fan 2B	4
FD Fan 2C	3
FD Fan 2D	4
FD Fan Set	2
FWH 2B1	2
FWH 2B2	2
FWH 2C	2
FWH 2D	2
FWH 2F1	2
FWH 2F2	2
FWH 2G1	2
FWH 2G2	2
General	3
Generator	8
HP Turbine	10
ID Fan 2A	9
ID Fan 2B	9
ID Fan 2C	9
ID Fan Set	2
IP Turbine	12
LP Turbine 2A	8
LP Turbine 2B	8
Main Condenser 2A	1
Main Condenser 2B	1

Unit 3

SootOpt

<u>Equipment</u>	<u># of Instances</u>
Air Heater IKs	1
Convection IKs	14
Furnace H2O Cannons	14

MaintenanceOpt

<u>Equipment</u>	<u># of Instances</u>
Air Heater 3A	1
Air Heater 3B	1
APH 3A	2
APH 3B	2
Aux Condenser 3A	1
BFP 3C	8
Boiler	6
CNDP 3A	3
CNDP 3B	3
CNDP 3C	3
CWP 3A	3
CWP 3B	3
CWP 3C	3
FD Fan 3A	4
FD Fan 3B	4
FD Fan 3C	4
FD Fan Set	2
FWH 3B1	2
FWH 3B2	2
FWH 3C	2
FWH 3D	1
FWH 3F1	2
FWH 3F2	2
FWH 3G1	2
FWH 3G2	2
General	3
HP Turbine	3
ID Fan 3A	8
ID Fan 3B	8
ID Fan 3C	8
ID Fan Set	2
IP Turbine	4
Main Condenser 3A	1
Main Condenser 3B	1

Overall # of Instances

492

A.2 CombustionOpt/SCR-Opt Home Page Details

The Home Page provides an overview of the CombustionOpt system. It shows what the Optimizer is doing and why, provides advice, and displays performance benchmarks.

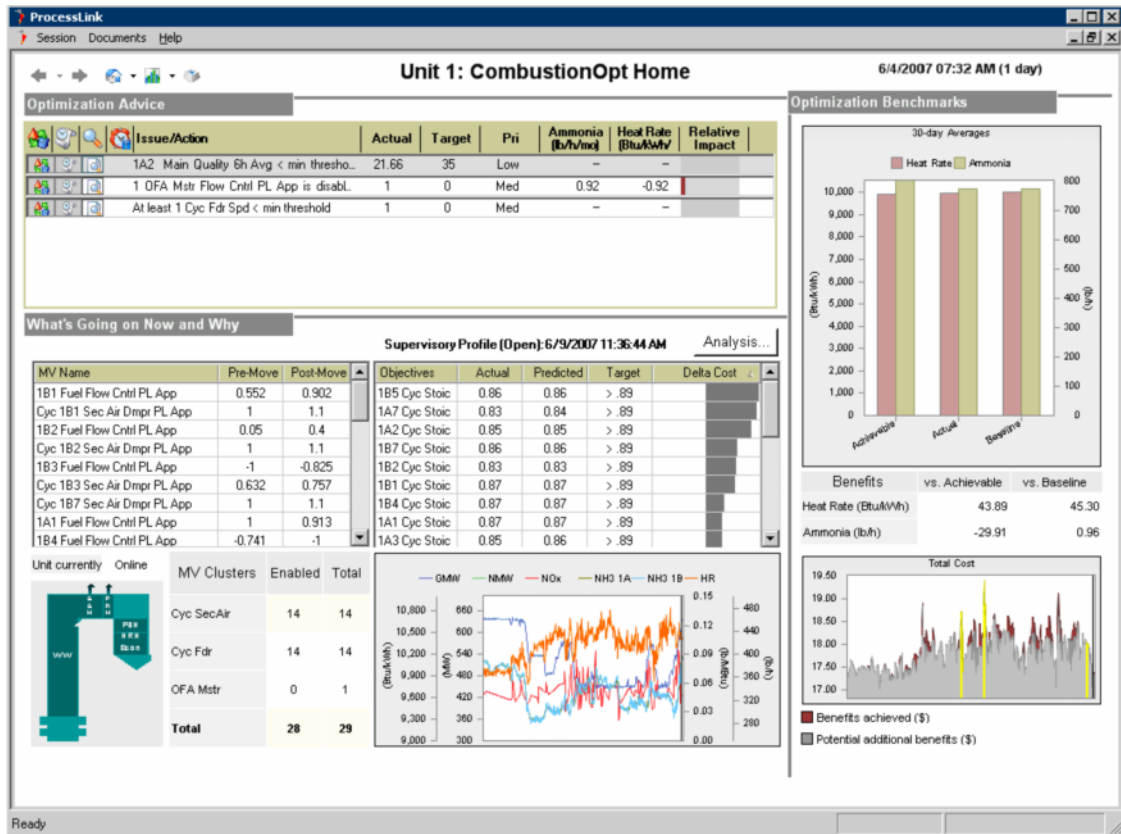


Figure 80 CombustionOpt's Home Page

CombustionOpt's Home Page is divided into three sections: What's Going on Now and Why (bottom left), Optimization Advice (top left), and Optimization Benchmarks (right side).

A.2.1 What's Going on Now and Why

The lower left section of the Home Page (Figure 80 above) provides information that explains why the closed-loop Optimizer made the move it did. It displays the optimization objectives in a way that gives insight into how much attention is being paid to each objective. It also displays what manipulated variables (MVs) were most recently moved by CombustionOpt, and how much each MV was moved. The iconized trend in the center contains key plant state variables and optimization objectives, to allow the user to quickly (by double-clicking) assess how they have been trending.

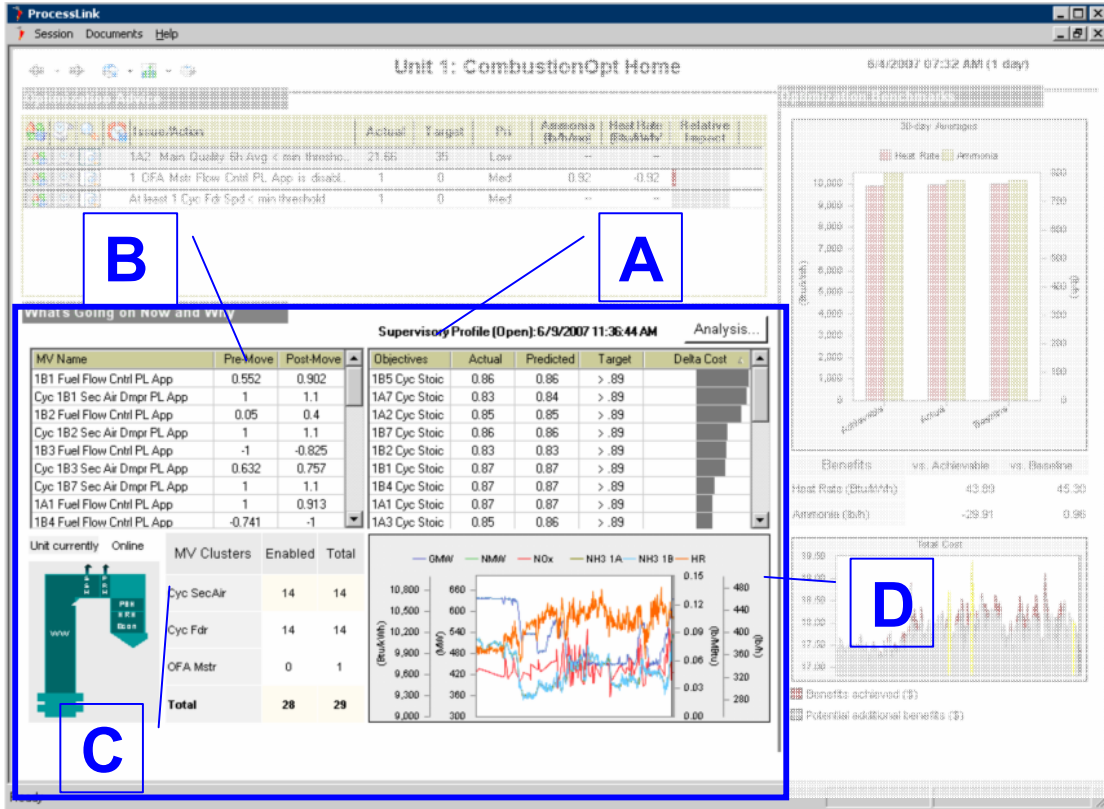


Figure 81 What's Going on Now and Why section on CombustionOpt's Home Page
CombustionOpt's What's Going on Now and Why section has four areas: Optimization Run Analysis table (A), Last MV Move table (B), MV Clusters Enabled table (C), and Trend Icon (D).

A.2.1.1 Optimization Run Analysis Table

The Optimization Run Analysis table (see "A" in Figure 81) shows all objectives and constraints in the selected optimization profile, for the selected optimization run.

The Optimization Run Analysis table has the following five columns:

Column	Description
Objectives	Name of optimization objective
Actual	Actual value of objective at time of optimization run
Predicted	Value optimizer predicts will result from changes to biases
Target	The optimization objective
Delta Cost	The relative importance of the objective in that run

Note that there is an Analysis button at the top right of this area. Clicking on this button takes you to the Optimization/Demystifier tab on the Analysis page.

A.2.1.2 Last MV Move table

The Last MV Move table (see "B" in Figure 81) shows the pre and post optimization values of each manipulated variable (MV) during that run. Note that the contents of the table can be sorted by clicking on any column header.

The Last MV Move table has the following three columns:

Column	Description
MV-Name	Name of manipulated variable (MV)
Pre-Move	Value of the MV before CombustionOpt's most recent move
Post-Move	Value of the MV after CombustionOpt's most recent move

A.2.1.3 MV Clusters Enable table

The MV Clusters Enabled table (see "C" in Figure 81) displays the enabled status of each manipulated variable (MV); the MVs are clustered by function.

The MV Clusters Enabled table has the following three columns:

Column	Description
MV-Clusters	Name we give to a group of similar MVs
Enabled	Number of MVs in the cluster that are enabled
Total	Total number of MVs in the cluster

A.2.1.4 Trend Icon

The iconized chart (see "D" in Figure 81) is a minimized version of a trend that shows a number of values relevant to combustion optimization; the values that are trended are specific to each site, but typically include gross MW, net MW, heat rate and NOx (or other key emissions parameters).

There are iconized charts on other screens as well. Double-clicking on any iconized chart maximizes it.

A.2.2 Optimization Advice

The upper left section of the Home Page shows current advice for how to further optimize the unit, along with tools and information to help prioritize and specify actions. The contents are generated by CombustionOpt's continuously running scenario evaluation system, which uses rules and models to look for opportunities to improve the benefits achieved by optimization.

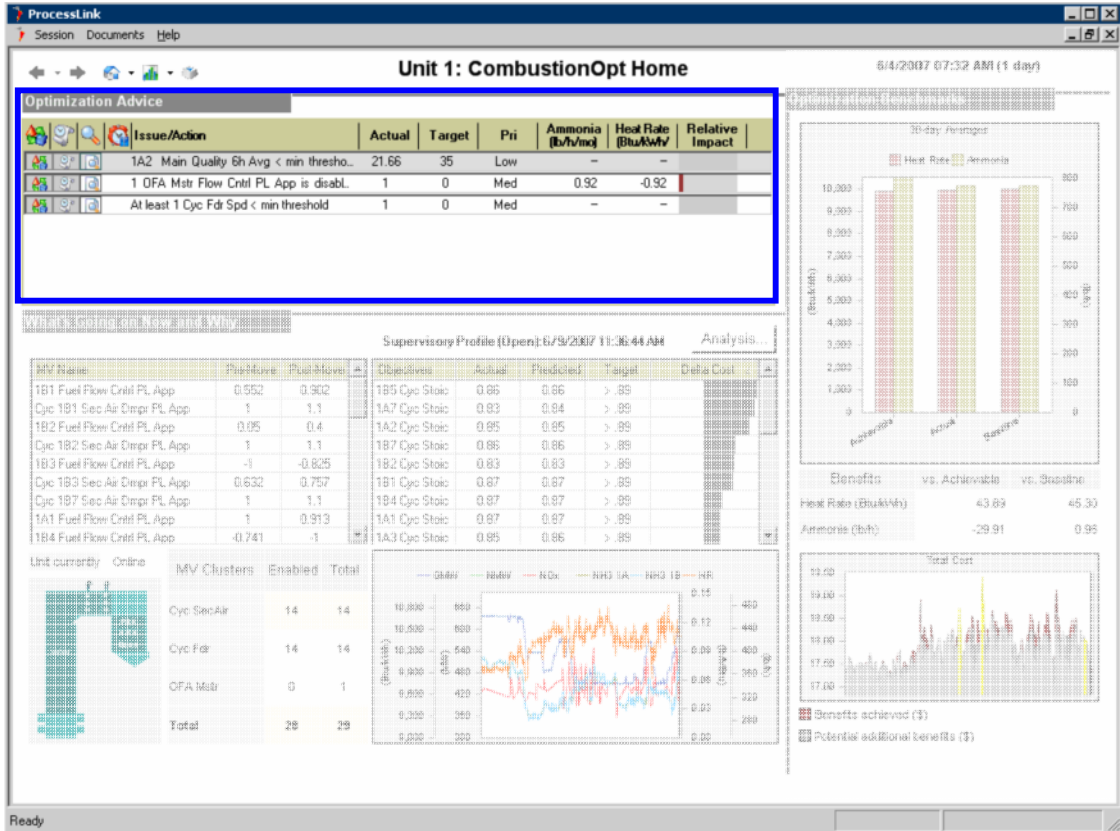


Figure 82 Optimization Advice section on CombustionOpt's Home Page

The Optimization Advice pane has the following ten columns:

Column	Description
	The Change Status button allows you to take action against each of the issues displayed
	The Problem History button displays the details of the history of past occurrences of the selected issue
	The Context Data button navigates you to a different view in the CombustionOpt application that provides drill-down information for analysis of the selected issue
	The Snooze icon is displayed if a user has placed an issue in the snoozed state. You can hide/display the snoozed issues using the right click menu on the Snooze column header
Issue/Action	Displays a description of the issue. If you have diagnosed an issue and specified an action (using the Change Status button), this column displays the description of the proposed corrective action. The tooltip (on mouse hover) shows both the issue and action descriptions
Actual	The actual value of the process/system variable that is used in the triggering rule for the alert

Column	Description
Target	The target value of the process/system variable that is used in the triggering rule for the alert
Pri	The priority for analysis and action, based on importance and urgency of the issue, assigned by a user using the Change Status button. The initial priority of each issue is a default value set during configuration
Svngs	For each objective that is dollarizeable, there is a column that displays the projected cost savings, for the next 28 days, that CombustionOpt believes will occur if its advice is taken. The savings are projected by summing/averaging the results of what-if optimizations executed for the last seven days and then multiplying by four
Relative Impact	Ranking based on the summation of all cost savings

A.2.2.1 Optimization Objectives

The impact on Objective X of taking a piece of advice is estimated using the formula below where the `achievable_impact_prediction` is the prediction of the model for objective X, under the scenario where the advice is taken and `actual_impact_prediction` is the impact based on the optimization that actually occurred.

$$\text{Impact on Objective X} = (\text{Projected Actual Cost}) - (\text{Projected Achievable Cost})$$

$$\text{Projected Actual Cost} = 4 \times \left(\sum_{\text{week}} \text{actual_impact_prediction}_{(\text{cost}/\text{MWh})} * \text{MW} \right)$$

$$\text{Projected Achievable Cost} = 4 \times \left(\sum_{\text{week}} \text{achievable_impact_prediction}_{(\text{cost}/\text{MWh})} * \text{MW} \right)$$

A.2.2.2 Change Status Information Collectors

You invoke the Change Status Information Collectors by clicking on the Change Status button. The various status change actions that can be taken are:

- Snooze
- Close
- Specify Action
- Change Risk
- Change Priority
- Escalate

The list of available actions will vary based on the current state (triggered, escalated, diagnosed, and snoozed) of the issue/problem that you are changing the status of. The contents of the Status Change dialog change, as detailed below, based on the status change action you have selected.

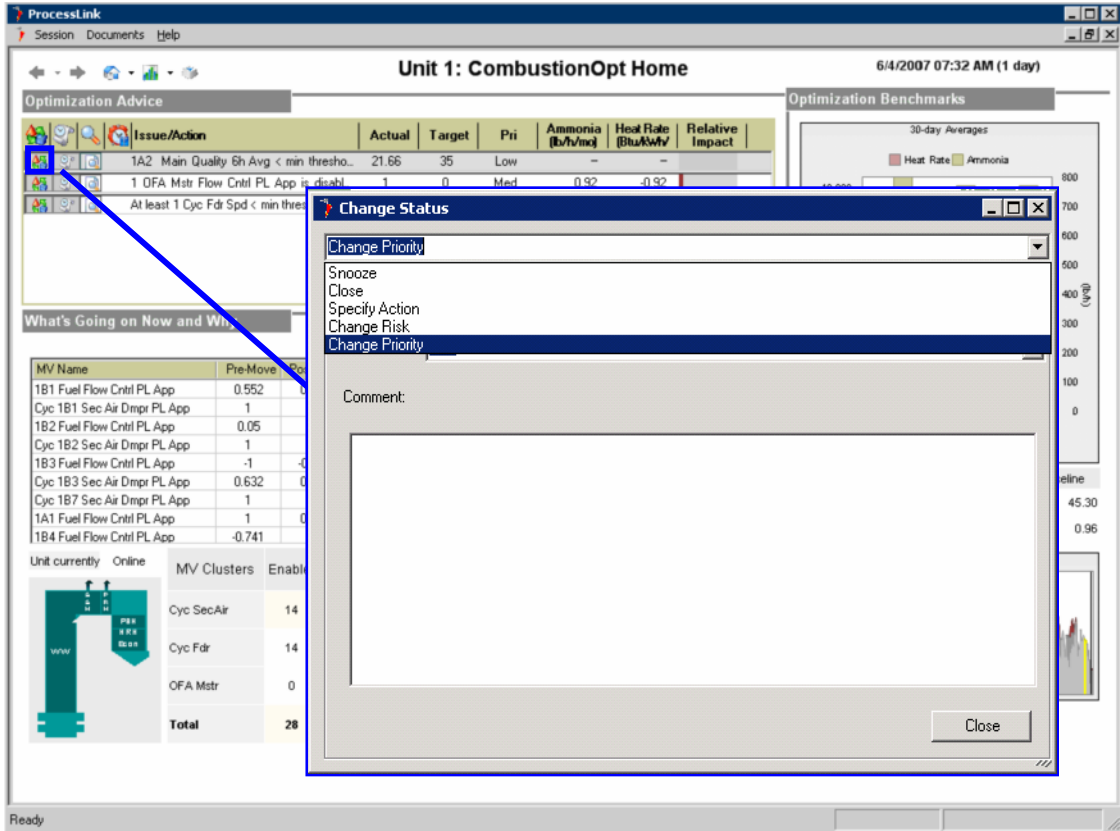


Figure 83 The Change Status Information Collector to Change Priority

Snooze

Puts the issue in a “snoozed” state for a specified time period during which you want it tagged as “snoozed”.

The Information Collector for snoozing/unsnoozing the selected issue asks for the following three pieces of information:

The Ask	Description
Time	The time period (in hours) that you would like the selected issue to be in a snoozed state
Priority	The priority (high, medium, or low) that you want to give the issue
Comment	Any comments you have that you want added to the history of the issue

Close

Dismisses the issue to indicate that it is not of concern or closes the issue to indicate that corrective action has been taken. This action moves the issue to the “closed” state; the issue is then removed from the Optimization Advice pane.

The Information Collector for closing the selected issue asks for the following piece of information:

The Ask	Description
Comment	Any comments you have that you want added to the history of the issue

Specify Action

Identifies a root cause and specifies a corrective action for the issue. This action moves the issue to the “diagnosed” state.

The Information Collector for specifying an action that will resolve the selected issue asks for the following six pieces of information:

The Ask	Description
Priority	The priority (high, medium, or low) that you want to give the issue
Cause	The root cause, selected from the list of likely causes available in the pull-down menu. If none of the causes listed apply, select “Other Cause” and provide specific details in the Comment section
Risk	Reliability risk associated with the issue – selected from the pull-down menu
Action Timing	The timing for the suggested corrective action – Outage, Derate, No Derate or Operations
Action	The corrective action; if the default action associated with the cause you specified is not the action that should be taken, type in the action that should be taken
Comment	Any comments you have that you want added to the history of the issue

Change Risk

Modifies the reliability risk associated with the issue.

The Information Collector for changing the reliability risk of the selected issue asks for the following two pieces of information:

The Ask	Description
Risk	The reliability risk associated with the issue – selected from the pull-down menu
Comment	Any comments you have that you want added to the history of the issue

Change Priority

Modifies the priority of the issue.

The Information Collector for changing the priority of the selected issue asks for the following two pieces of information:

The Ask	Description
Priority	The priority (high, medium, or low) that you want to give the issue
Comment	Any comments you have that you want added to the history of the issue

Escalate

Validates that the alert is not a false alarm and should be analyzed and resolved. Action available only for alerts currently in the "triggered" state – moves the issue from the "triggered" to the "escalated" state.

The Information Collector for escalating the selected issue from an issue to screen to a problem to diagnose asks for the following two pieces of information:

The Ask	Description
Priority	The priority (high, medium, or low) that you want to give the issue
Comment	Any comments you have that you want added to the history of the issue

A.2.3 Optimization Benchmarks

The section on the right of the Home Page shows how the unit has been doing with respect to its optimization objectives as well as how much better it might have done.

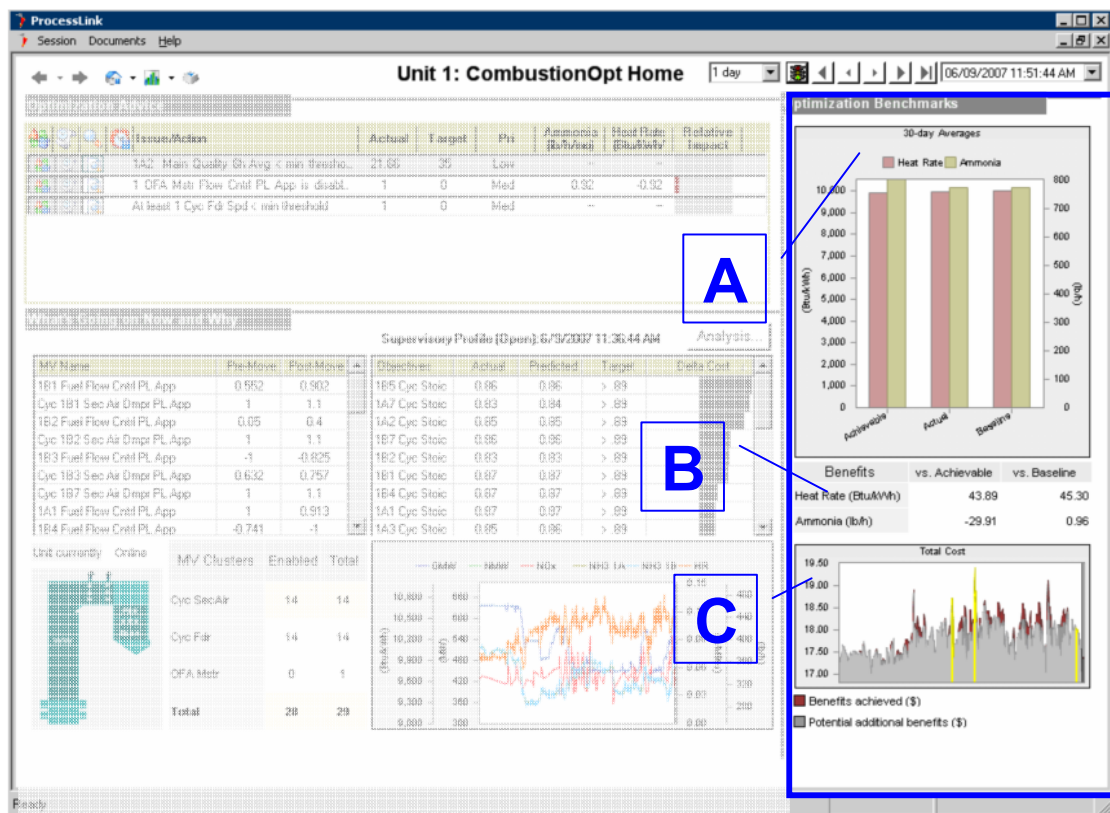


Figure 84 Optimization Benchmarks section on CombustionOpt's Home Page
CombustionOpt's Optimization Benchmarks section has three areas: Benchmark bar chart (A), Benefits table (B) and Triband chart (C)

A.2.3.1 Benchmark bar chart

The Benchmark bar chart (see "A" in Figure 84) shows you three values for each variable, where the variables represent optimization objectives such as NOx, Heat Rate, and NH3 Flow. The three values represent the Achievable, Actual, and Baseline benchmarks for that objective.

The Benchmark bar chart has the following three features:

Feature	Description
Achievable	An <i>achievable benchmark</i> represents the performance with respect to objectives and constraints that would have been achieved if all known optimization opportunities, both closed-loop and user actuated, had been taken. It is derived by summing, for each objective, at each point in time over the past 28 days, the cost that would have been incurred if the optimal actions had been taken
Actual	The <i>actual benchmark</i> value represents how the process being optimized has actually been doing, or in other words, the cost actually incurred
Baseline	The <i>baseline benchmark</i> represents the performance with respect to objectives and constraints that would have occurred if there had been no optimization. It is derived by summing, for each objective, at each point in time over the past 28 days, the cost that would have been incurred if all MVs were at their "neutral" or "zero-bias" position

A.2.3.1.1 Benchmarking

Benchmarking refers to using a "point of reference" to evaluate performance. There are many ways to develop these reference points. Past performance is one type of benchmark. Performance with respect to design is another. CombustionOpt uses statistical modeling to understand cause and effect relationships between manipulated variables and performance. Its models are used to run "what-if" scenarios. A useful analogy is using a curve fit to data containing historical performance and manipulated variables to "look up" the anticipated performance based on different values of the manipulated variables.

A.2.3.1.2 Benchmark calculations

The achievable benchmark is calculated using the formula below where the objective_pred is the prediction of the models for each naturally dollarizeable objective in question, under the "what if" scenario where all the known optimization actions are taken and objective_actual_value is the actual value of the objective.

$$Achievable = \left(\sum_{month} objective_pred_{(cost / MWh)} * MW \right)$$

The actual benchmark is calculated using the formula below

$$Actual = \left(\sum_{month} objective_actual_value_{(cost / MWh)} * MW \right)$$

A.2.3.2 Benefits table

The Benefits table (see "B" in Figure 84) shows, for the past 28 days for each optimization objective, the savings actually achieved through optimization, as well as the savings that were not achieved but could have been had all identified optimization actions been taken.

The Benefits table has the following three columns:

Column	Description
Benefits	The optimization objectives in the benchmark bar chart
vs. Achievable	(Achievable minus Actual) in the benchmark bar chart
vs. Baseline	(Actual minus Baseline) in the benchmark bar chart

Note that in "vs. Achievable", the values for some of the objectives may be negative (indicating that CombustionOpt needed to incur additional costs with respect to those objectives in order to achieve optimal savings overall.)

A.2.3.3 Triband chart

The Triband chart (see "C" in Figure 84) displays a single value for all objectives, in \$/MWhr, calculated at each moment. These values represent the benefit achieved through optimization (difference between actual and baseline), and the benefit that was not achieved but could have been had all identified optimization actions been taken (difference between achievable and actual.)

A.3 SootOpt Home Page Details

The SootOpt Home Page provides an overview of the SootOpt system. It shows what the Optimizer is doing and why, provides advice, and displays performance benchmarks.

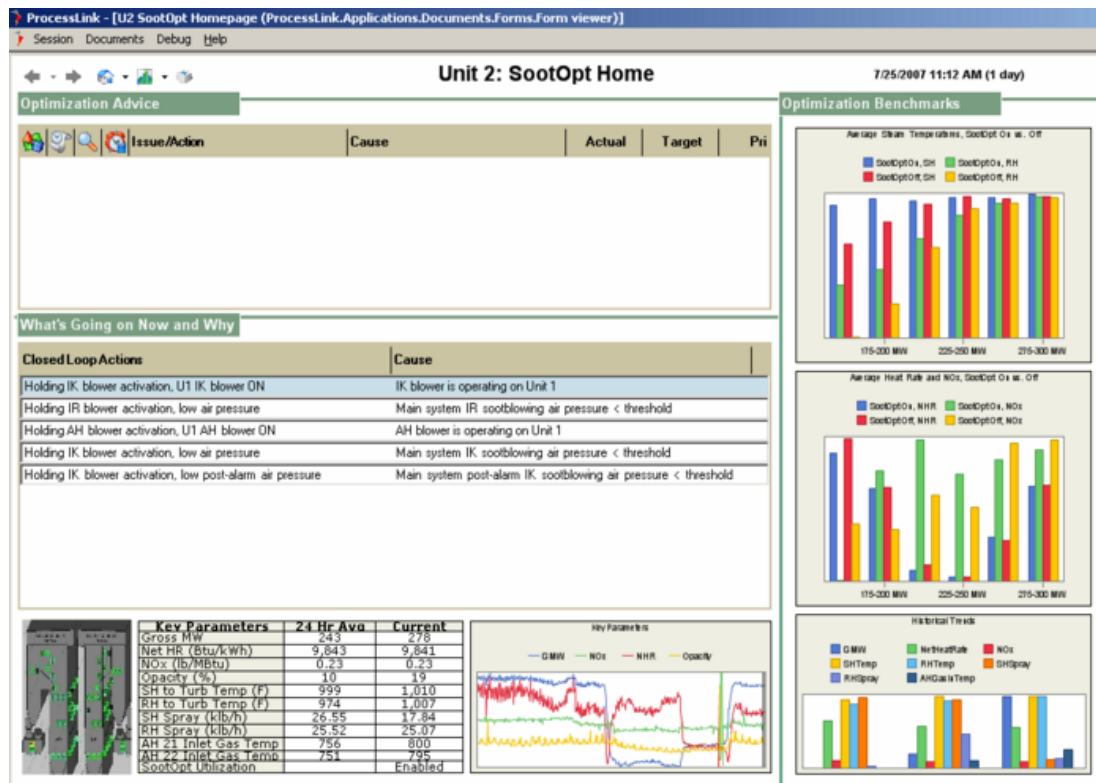


Figure 85 SootOpt's Home Page

SootOpt's Home Page is divided into three sections: What's Going on Now and Why (bottom left), Optimization Advice (top left), and Optimization Benchmarks (right side).

A.3.1 What's Going on Now and Why

The lower left section of the Home Page provides information that explains what the closed-loop Optimizer is doing and why. It displays the current sootblowing actions being taken by SootOpt's closed-loop optimizer, as well as any inhibiting conditions that may be affecting sootblowing activities. The trend and table display key plant state variables and optimization objectives, for quick access.

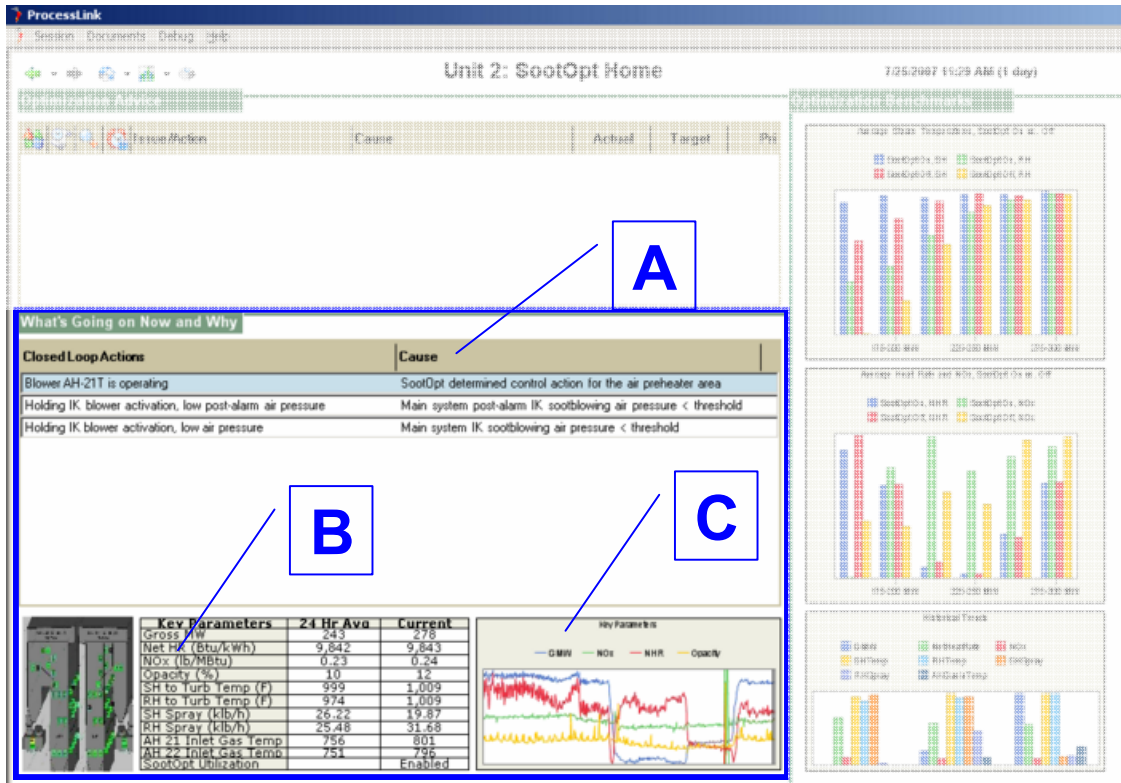


Figure 86 What's Going on Now and Why section on SootOpt's Home Page
 SootOpt's What's Going on Now and Why section has three areas: Optimization Run Analysis table (A), Key Parameters Table (B), and Trend Icon (C).

A.3.1.1 Optimization Run Analysis Table

The Optimization Run Analysis table (see "A" in Figure 86) shows the blowers currently operating (actuators determined by SootOpt) in each area, as well as any inhibiting conditions temporarily holding sootblowing activities in any of the areas.

The Optimization Run Analysis table has the following two columns:

Column	Description
Closed Loop Actions	What SootOpt is currently doing
Cause	Explanation of why

A.3.1.2 Key Parameters table

The Key Parameters table (see "B" in Figure 86) shows the current values as well as the 24-hour average values for a set of parameters related to sootblowing.

The Key Parameters table has the following three columns:

Column	Description
Key Parameters	Name of the parameter
24 Hr Avg	24-hour moving average value of the parameter
Current	Current value of the parameter

A.3.1.3 MV Clusters Enable table

The MV Clusters Enabled table (see "C" in Figure 81) displays the enabled status of each manipulated variable (MV); the MVs are clustered by function.

The MV Clusters Enabled table has the following three columns:

Column	Description
MV-Clusters	Name we give to a group of similar MVs
Enabled	Number of MVs in the cluster that are enabled
Total	Total number of MVs in the cluster

A.3.1.4 Trend Icon

The iconized chart (see "C" in Figure 86 above) is a minimized version of a trend that shows a number of parameters relevant to sootblowing optimization; the values that are trended are specific to each site, but typically include gross MW, net heat rate, NOx and Opacity (or other key emissions parameters).

There are iconized charts on other screens as well. Double-clicking on any iconized chart maximizes it.

A.3.2 Optimization Advice

The upper left section of the Home Page shows current advice for how to further optimize the unit, along with tools and information to help prioritize and specify actions. Alerts presented here represent anomalies that SootOpt has identified that have some impact on optimization objectives and/or are a risk to capacity. On sites where MaintenanceOpt also is installed, these items are forwarded for inclusion in its diagnostic workflow.

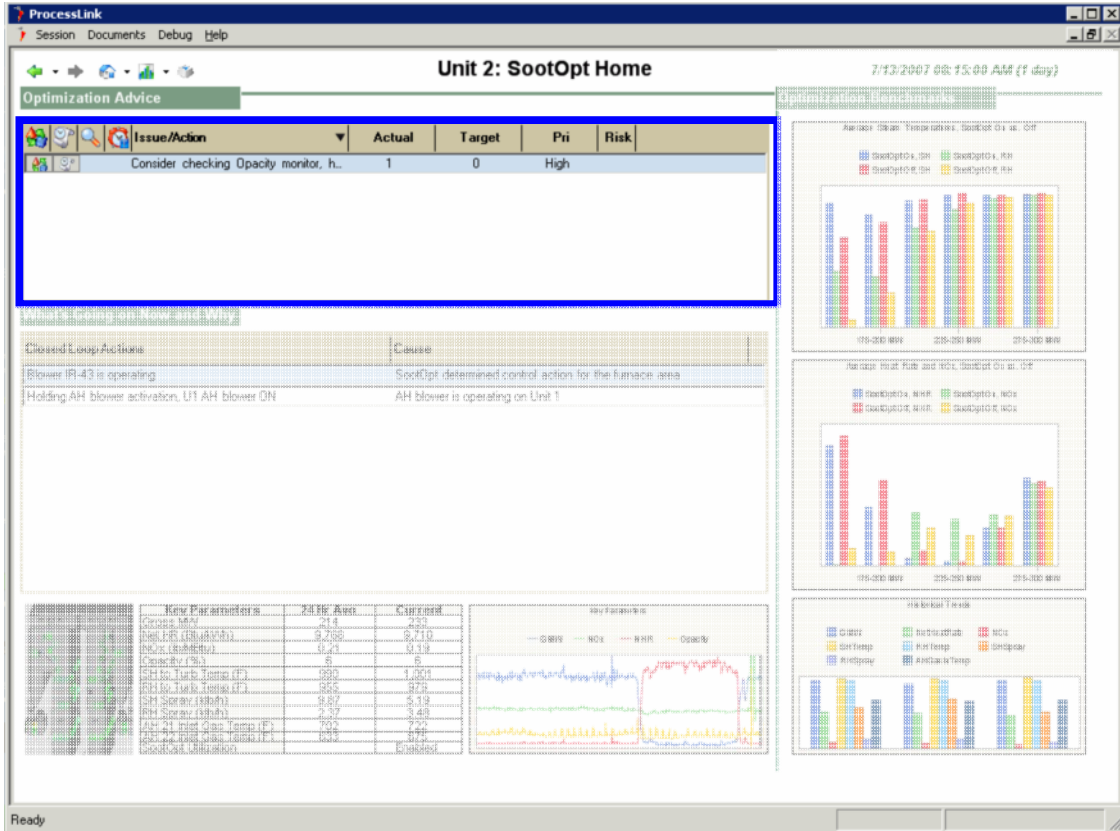






Figure 87 Optimization Advice section on SootOpt's Home Page

The Optimization Advice pane has the following columns:

Column	Description
	The Change Status button allows you to take actions against each of the issues displayed.
	The Problem History button displays the details of the history of past occurrences of the selected issue
	The Context Data button navigates you to a different view in the SootOpt application that provides drill-down information for analysis of the selected issue
	The Snooze icon is displayed if a user has placed an issue in the snoozed state. You can hide/display the snoozed issues using the right click menu on the Snooze column header
Issue/Action	Displays a description of the issue. If you have diagnosed an issue and specified an action (using the Change Status button), this column displays the description of the proposed corrective action. The tooltip (on mouse hover) shows both the issue and action descriptions
Actual	The actual value of the process/system variable that is used in the triggering rule for the alert

Column	Description
Target	The target value of the process/system variable that is used in the triggering rule for the alert
Pri	The priority for analysis and action, based on importance and urgency of the issue, assigned by a user using the Change Status button. The initial priority of each issue is a default value set during configuration
Risk	An estimate of the risk this issue poses to capacity (not always set)

A.3.2.1 Change Status Information Collectors

The Change Status Information Collector is invoked by clicking on the Change Status button. The various status change actions that can be taken are:

- Snooze
- Close
- Specify Action
- Change Risk
- Change Priority
- Escalate

The list of available actions will vary based on the current state (triggered, escalated, diagnosed, and snoozed) of the issue/problem selected for change. The contents of the Status Change dialog change, as detailed below, based on the status change action you have selected.

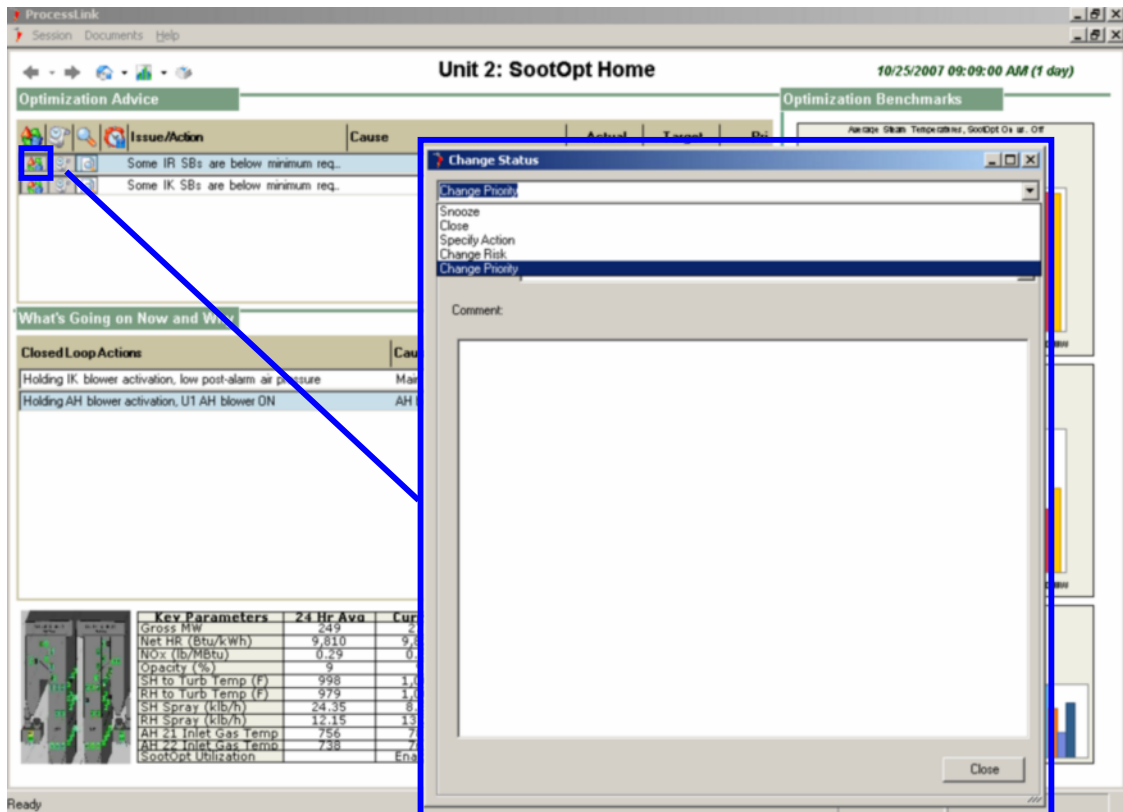


Figure 88 The Change Status Information Collector to Change Priority

Snooze

This action puts the issue in a “snoozed” state for a specified time period.

The Information Collector for snoozing/unsnoozing the selected issue asks for the following three pieces of information:

The Ask	Description
Time	The time period (in hours) the selected issue is to remain in a snoozed state
Priority	The priority (high, medium, or low) given to the issue
Comment	Any comments to be added to the history of the issue

Close

Dismisses the issue to indicate that it is not of concern or closes the issue to indicate that corrective action has been taken. This action moves the issue to the “closed” state; the issue is then removed from the Optimization Advice pane.

The Information Collector for closing the selected issue asks for the following piece of information:

The Ask	Description
Comment	Any comments to be added to the history of the issue

Specify Action

Identifies a root cause and specifies a corrective action for the issue. This action moves the issue to the “diagnosed” state.

The Information Collector for specifying an action that will resolve the selected issue asks for the following six pieces of information:

The Ask	Description
Priority	The priority (high, medium, or low) given the issue
Cause	The root cause, selected from the list of likely causes available in the pull-down menu. If none of the causes listed apply, select “Other Cause” and provide specific details in the Comment section
Risk	Reliability risk associated with the issue – selected from the pull-down menu
Action Timing	The timing for the suggested corrective action – Outage, Derate, No Derate or Operations
Action	The corrective action; if the default action associated with the cause specified is not the action that should be taken, a new action can be entered in
Comment	Any comments to be added to the history of the issue

Change Risk

Modifies the reliability risk associated with the issue.

The Information Collector for changing the reliability risk of the selected issue asks for the following two pieces of information:

The Ask	Description
Risk	The reliability risk associated with the issue – selected from the pull-down menu
Comment	Any comments to be added to the history of the issue

Change Priority

Modifies the priority of the issue.

The Information Collector for changing the priority of the selected issue asks for the following two pieces of information:

The Ask	Description
Priority	The priority (high, medium, or low) given the issue
Comment	Any comments to be added to the history of the issue

Escalate

Validates that the alert is not a false alarm and should be analyzed and resolved. Action available only for alerts currently in the “triggered” state – moves the issue from the “triggered” to the “escalated” state.

The Information Collector for escalating the selected issue from an issue to screen to a problem to diagnose asks for the following two pieces of information:

The Ask	Description
Priority	The priority (high, medium, or low) given the issue
Comment	Any comments to be added to the history of the issue

A.3.3 Optimization Benchmarks

The section on the right of the Home Page shows how the unit has been doing with respect to its optimization objectives as well as how much better it might have done.

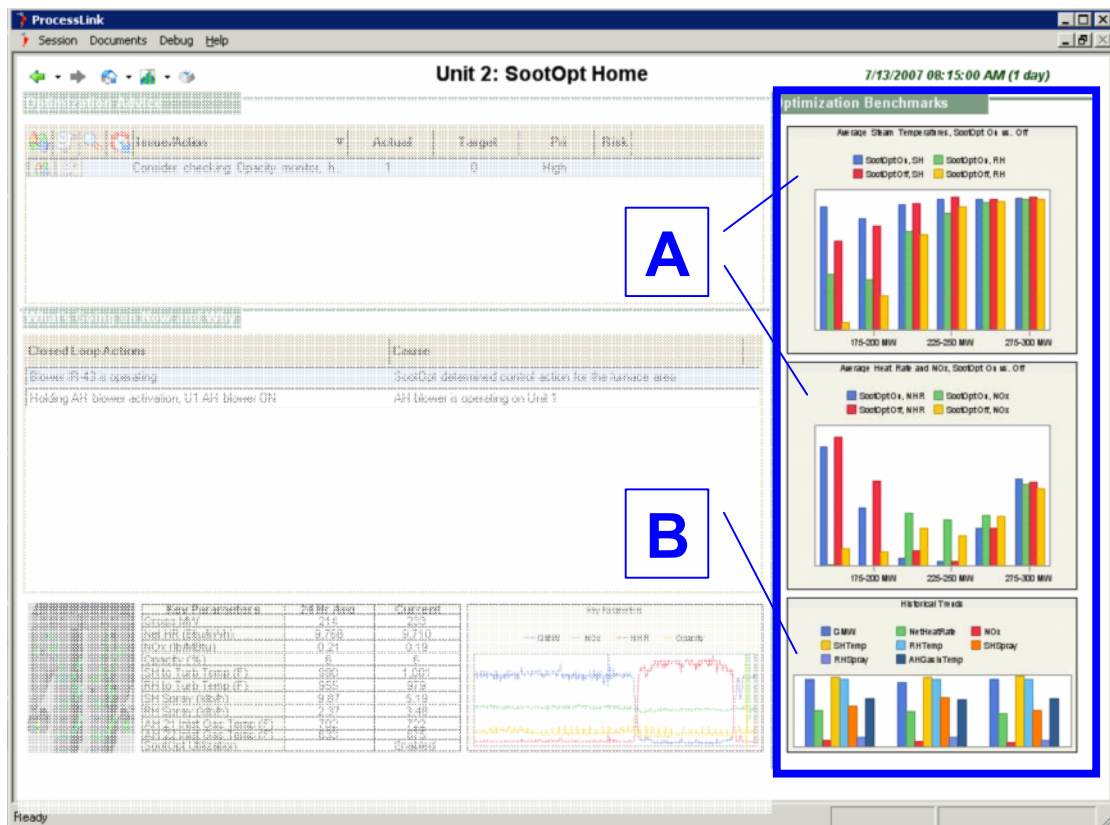


Figure 89 Optimization Benchmarks section on SootOpt’s Home Page

SootOpt’s Optimization Benchmarks section has two areas: Benchmark bar charts (A) and Historical trend chart (B)

A.3.3.1 Benchmark bar chart

The benchmark bar charts (see "A" in Figure 89) show 30-day average values of plant performance parameters such as superheat steam temperature, reheat steam temperature, net heat rate and NOx filtered for times when SootOpt is on and when it is off under different Unit Loads. Double-clicking on any iconized chart will open a full size view.

A.3.3.1.1 Benchmarking

Benchmarking refers to using a "point of reference" to evaluate performance. There are many ways to develop these reference points. Past performance is one type of benchmark. Performance with respect to design is another. Using statistical models that understand cause and effect relationships between manipulated variables and performance in hypothetical scenarios is another, more advanced technique.

A.3.3.2 Historical table

The bottom of the highlighted area (see "B" in Figure 89) shows various plant state and performance variables benchmarked against their historical values. The bar charts are categorized to show average values for a variety of parameters over three different time periods, Yesterday, Today and the past 7 days.

A.4 PerformanceOpt Home Page Details

The Home Page provides an overview of the PerformanceOpt system. It highlights the performance gaps that PerformanceOpt has identified, shows key performance indicators for current unit operation, and also benchmarks recent past performance.

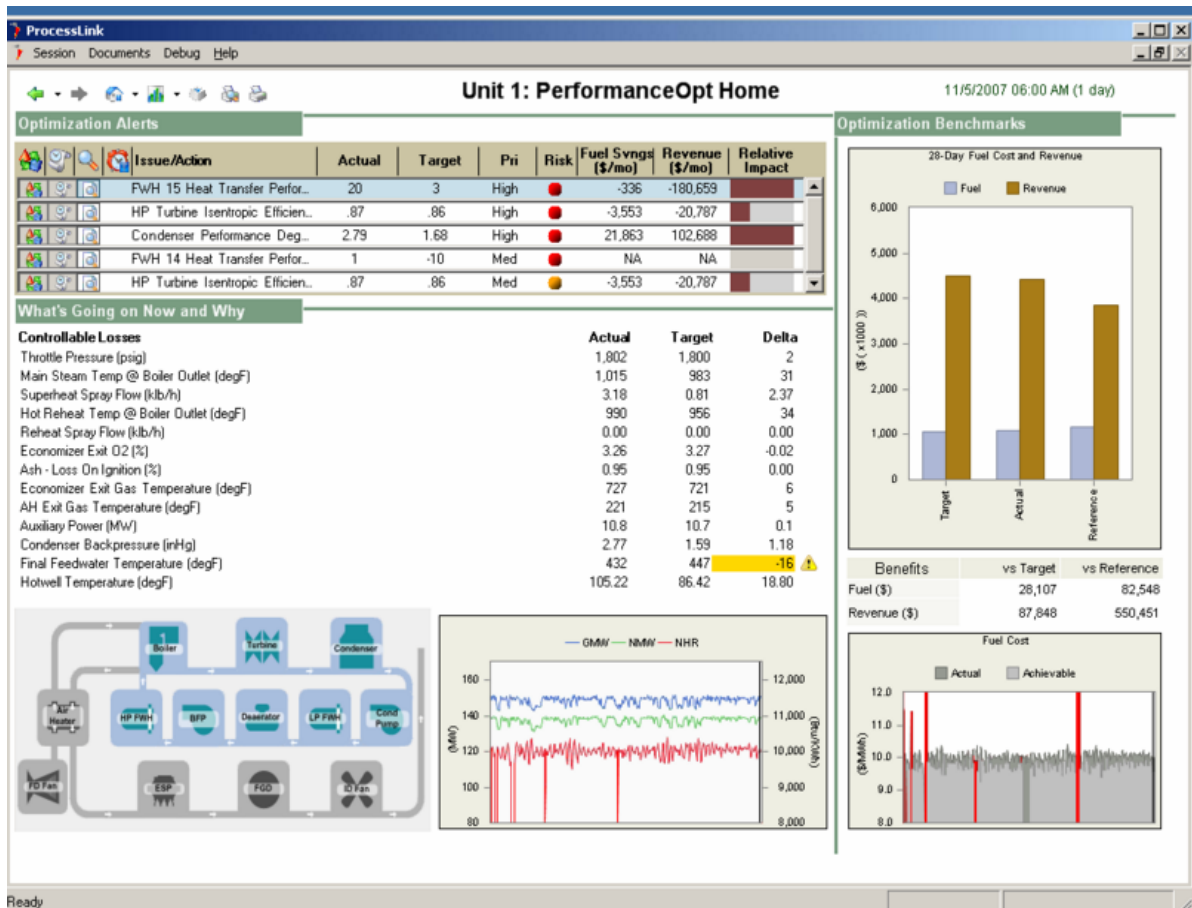


Figure 90 PerformanceOpt's Home Page

PerformanceOpt's Home Page is divided into three sections: What's Going on Now and Why (bottom left), Optimization Advice (top left), and Optimization Benchmarks (right side).

A.4.1 What's Going on Now and Why

The lower left section of the Home Page provides information on Controllable Losses, traditionally calculated by all performance monitoring systems, for the unit. The iconized trend in the center displays key unit-level performance measures, and allows the user to quickly (by double-clicking) assess how they have been trending. The process schematic in the bottom is a navigation aid to give users one-click access to the collection of pre-configured equipment-level data views.

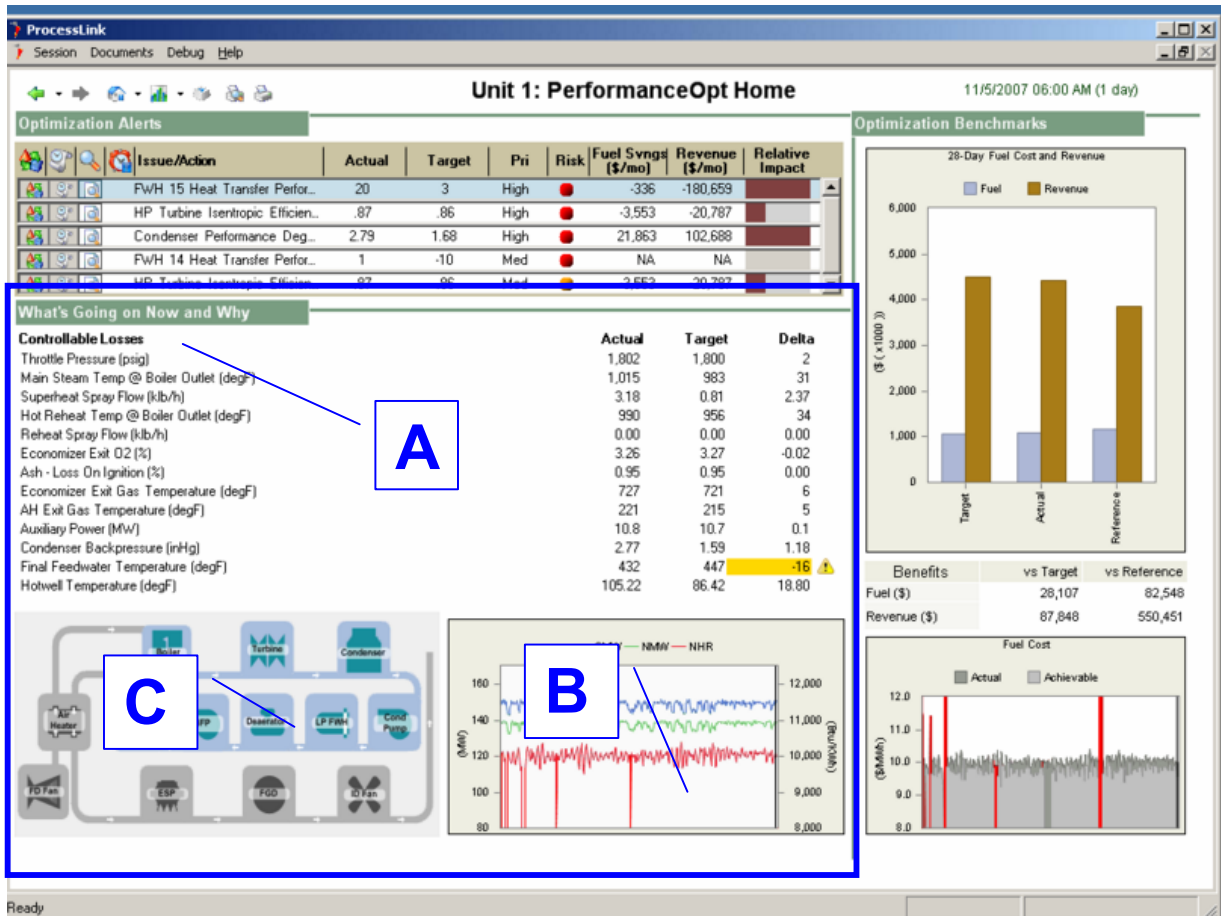


Figure 91 What's Going on Now and Why section on PerformanceOpt's Home Page
 PerformanceOpt's What's Going on Now and Why section has three areas: Controllable Losses table (A), Unit Performance Trend Icon (B) and Navigation Process Schematic (C).

A.4.1.1 Controllable Losses table

The Controllable Losses table (see "A" in Figure 91) lists the current actual and achievable values for the following controllable loss variables. You can generate a time trend of any value, row, or column of values in this table using the context menu available from the mouse right click.

Controllable Losses	Actual	Target	Delta
Throttle Pressure (psig)	1,802	1,800	2
Main Steam Temp @ Boiler Outlet (degF)	1,015	983	31
Superheat Spray Flow (klb/h)	3.18	0.81	2.37
Hot Reheat Temp @ Boiler Outlet (degF)	990	956	34
Reheat Spray Flow (klb/h)	0.00	0.00	0.00
Economizer Exit O2 (%)	3.26	3.27	-0.02
Ash - Loss On Ignition (%)	0.95	0.95	0.00
Economizer Exit Gas Temperature (degF)	727	721	6
AH Exit Gas Temperature (degF)	221	215	5
Auxiliary Power (MW)	10.8	10.7	0.1
Condenser Backpressure (inHg)	2.77	1.59	1.18
Final Feedwater Temperature (degF)	432	447	-15
Hotwell Temperature (degF)	105.22	96.42	18.80

FWH 14 Heat Transfer Performance Degradation
 FWH 15 Heat Transfer Performance Degradation

Figure 92 Controllable Loss parameters monitored on PerformanceOpt Home Page

In addition, controllable losses that are impacted by any of the alerts that are currently active and displayed in the Optimization Advice section, are highlighted in yellow as shown in

Figure 92. When you hover the mouse over the highlight, the tooltip shows the performance alert(s) which should be addressed in order to optimize the specific controllable loss impact.

The Controllable Losses table has the following three columns:

Column	Description
Actual	Value computed by the PerformanceOpt "as-is" simulation that models current unit operation
Target	Value computed by the PerformanceOpt "what-if" simulation that models unit operation under current conditions but with all equipment operating at "peak" performance
Delta	Difference between the actual and values for that controllable loss variable

A.4.1.2 Unit Performance Trend Icon

The iconized chart (see "B" in Figure 91) is a minimized version of a trend that shows a number of values relevant to combustion optimization; the values that are trended are specific to each site, but typically include gross MW, net MW, heat rate and NOx (or other key emissions parameters).

There are iconized charts on other screens as well. Double-clicking on any iconized chart maximizes it.

A.4.1.3 Navigation Process Schematic

The Navigation Process Schematic (see "C" in Figure 91) is a navigation aid to give users one-click access to the collection of pre-configured equipment-level data views. For PerformanceOpt configurations this typically consists of icons with navigation links for Boiler, Turbine, Condenser, Air Heater, High and Low pressure Feedwater Heaters and Deaerator, Boiler Feed Pumps, Condensate Pumps, and FD and ID Fans.

A.4.2 Optimization Alerts

The upper left section of the Home Page shows the current alerts for how to further optimize the unit, along with tools and information to help prioritize and specify actions. The contents are generated by PerformanceOpt's comparison of actual and achievable equipment-level performance computed by its continuously running as-is and what-if simulations using the rigorous heat and mass balance model of the unit.

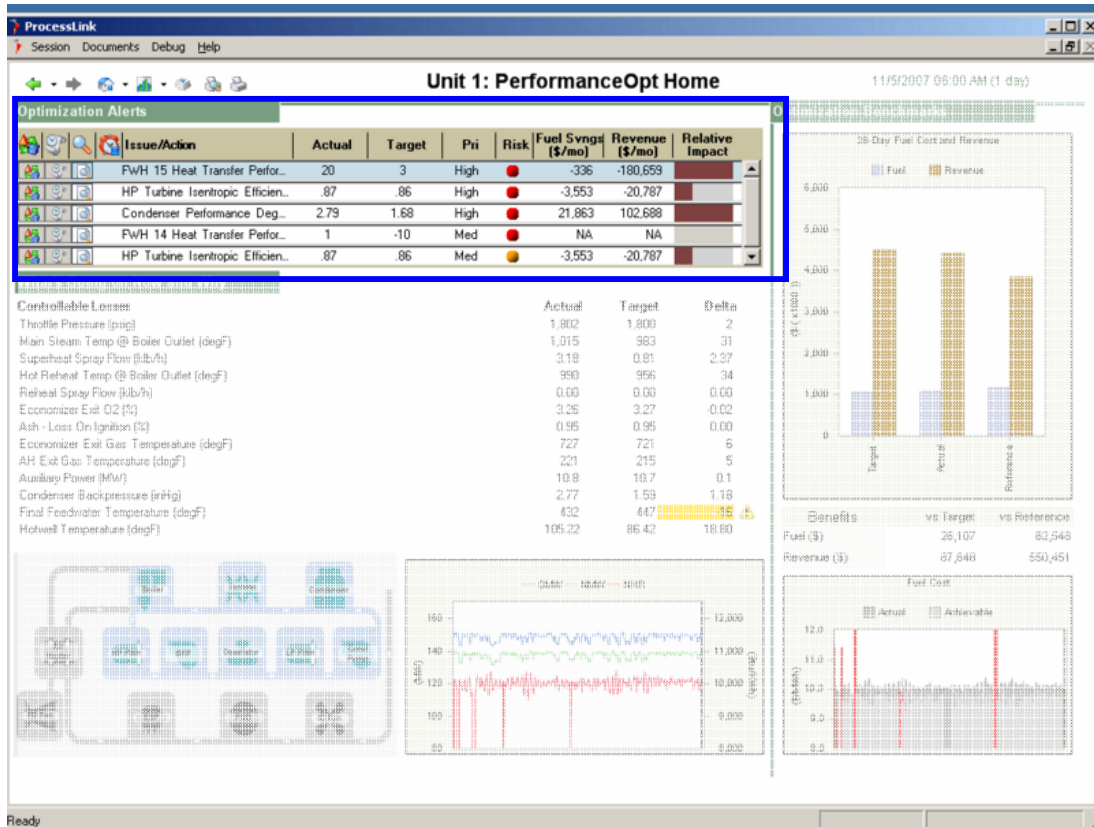


Figure 93 Optimization Alerts section on PerformanceOpt Home Page

The Optimization Alerts pane has the following twelve columns:

Column	Description
	The Change Status button allows you to take action against each of the issues displayed
	The Problem History button displays the details of the history of past occurrences of the selected issue
	The Context Data button navigates you to a different view in the CombustionOpt application that provides drill-down information for analysis of the selected issue
	The Snooze icon is displayed if a user has placed an issue in the snoozed state. You can hide/display the snoozed issues using the right click menu on the Snooze column header
Issue/Action	Displays a description of the issue. If you have diagnosed an issue

Column	Description
	and specified an action (using the Change Status button), this column displays the description of the proposed corrective action. The tooltip (on mouse hover) shows both the issue and action descriptions
Actual	The actual value of the process/system variable that is used in the triggering rule for the alert
Target	The target value of the process/system variable that is used in the triggering rule for the alert
Pri	The priority (High, Medium and Low) for analysis and action, based on importance and urgency of the issue, assigned by a user using the Change Status button. The initial priority of each issue is a default value set during configuration.
Risk	Qualitative probability (High, Medium, Low, and None) of the issue impacting the reliability and availability of the affected equipment.
Fuel Svngs	Shows the impact on Fuel Cost, based on change in unit heat rate, due to the identified performance gap. The Fuel Cost impact is shown as a cumulative projection for the next month, based on the average impact over the past week
Revenue	Shows the impact on Revenue, based on change in load generated, due to the identified performance gap. The Revenue impact is shown as a cumulative projection for the next month, based on the average impact over the past week
Relative Impact	Ranking based on the summation of Fuel Cost and Revenue impacts

A.4.2.1 Change Status Information Collectors

You invoke the Change Status Information Collectors by clicking on the Change Status button. The various status change actions that can be taken are:

- Snooze
- Close
- Specify Action
- Change Risk
- Change Priority
- Escalate

The list of available actions will vary based on the current state (triggered, escalated, diagnosed, and snoozed) of the issue/problem that you are changing the status of. The contents of the Status Change dialog change, as detailed below, based on the status change action you have selected.

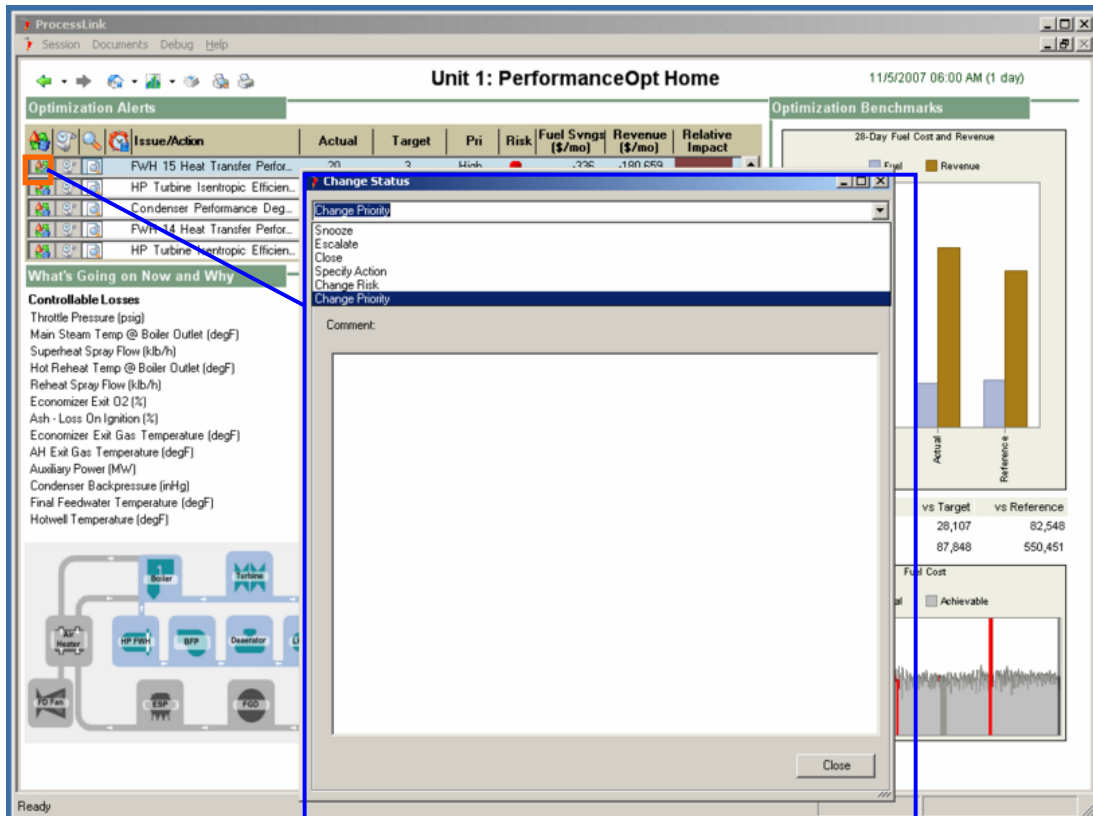


Figure 94 The Change Status Information Collector to Change Priority

Snooze

Puts the issue in a “snoozed” state for a specified time period during which you want it tagged as “snoozed”.

The Information Collector for snoozing/unsnoozing the selected issue asks for the following three pieces of information:

The Ask	Description
Time	The time period (in hours) that you would like the selected issue to be in a snoozed state
Priority	The priority (high, medium, or low) that you want to give the issue
Comment	Any comments you have that you want added to the history of the issue

Close

Dismisses the issue to indicate that it is not of concern or closes the issue to indicate that corrective action has been taken. This action moves the issue to the “closed” state; the issue is then removed from the Optimization Advice pane.

The Information Collector for closing the selected issue asks for the following piece of information:

The Ask	Description
Comment	Any comments you have that you want added to the history of the issue

Specify Action

Identifies a root cause and specifies a corrective action for the issue. This action moves the issue to the “diagnosed” state.

The Information Collector for specifying an action that will resolve the selected issue asks for the following six pieces of information:

The Ask	Description
Priority	The priority (high, medium, or low) that you want to give the issue
Cause	The root cause, selected from the list of likely causes available in the pull-down menu. If none of the causes listed apply, select “Other Cause” and provide specific details in the Comment section
Risk	Reliability risk associated with the issue – selected from the pull-down menu
Action Timing	The timing for the suggested corrective action – Outage, Derate, No Derate or Operations
Action	The corrective action; if the default action associated with the cause you specified is not the action that should be taken, type in the action that should be taken
Comment	Any comments you have that you want added to the history of the issue

Change Risk

Modifies the reliability risk associated with the issue.

The Information Collector for changing the reliability risk of the selected issue asks for the following two pieces of information:

The Ask	Description
Risk	The reliability risk associated with the issue – selected from the pull-down menu
Comment	Any comments you have that you want added to the history of the issue

Change Priority

Modifies the priority of the issue.

The Information Collector for changing the priority of the selected issue asks for the following two pieces of information:

The Ask	Description
Priority	The priority (high, medium, or low) that you want to give the issue
Comment	Any comments you have that you want added to the history of the issue

Escalate

Validates that the alert is not a false alarm and should be analyzed and resolved. Action available only for alerts currently in the "triggered" state – moves the issue from the "triggered" to the "escalated" state.

The Information Collector for escalating the selected issue from an issue to screen to a problem to diagnose asks for the following two pieces of information:

The Ask	Description
Priority	The priority (high, medium, or low) that you want to give the issue
Comment	Any comments you have that you want added to the history of the issue

A.4.3 Optimization Benchmarks

The section on the right of the Home Page shows how the unit has been performing based on Heat Rate and Capacity, in aggregated terms, over its recent operational history, typically one month.

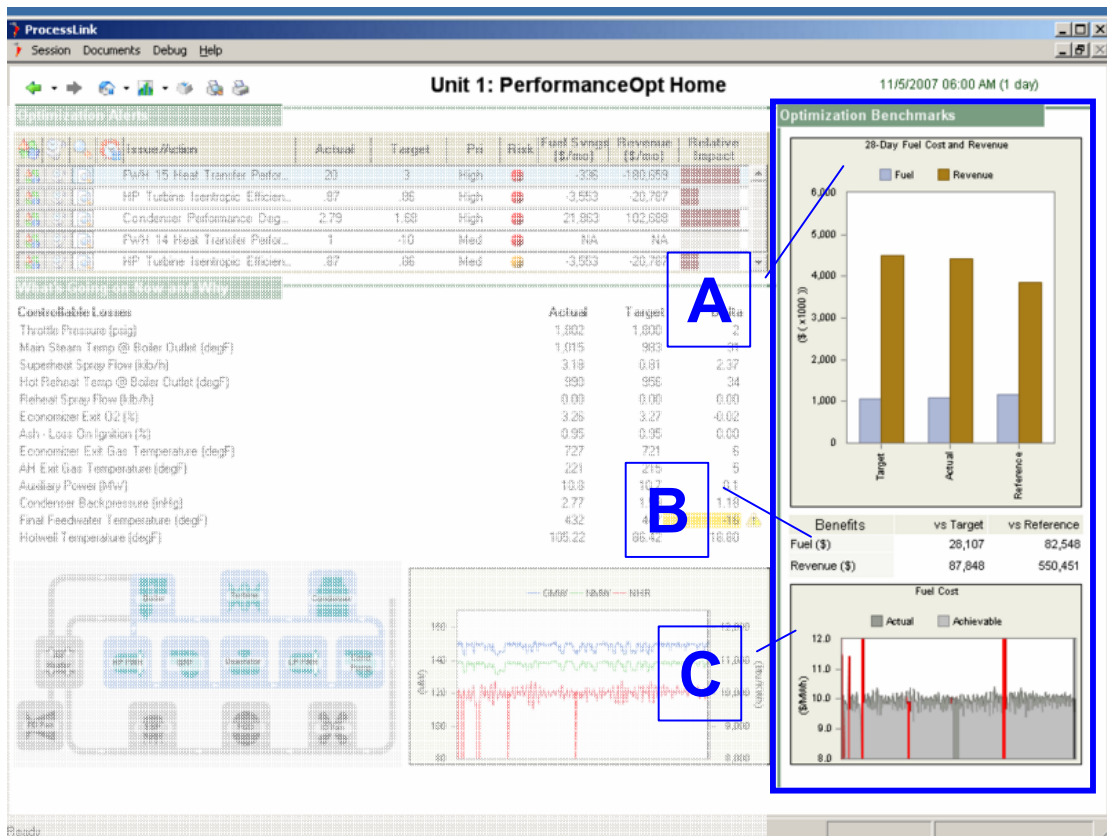


Figure 95 Optimization Benchmarks section on PerformanceOpt's Home Page
PerformanceOpt's Optimization Benchmarks section has three areas: Benchmark bar chart (A), Benefits table (B) and Unit Efficiency chart (C)

A.4.3.1 Benchmark bar chart

The Benchmark bar chart (see "A" in Figure 95) shows you fuel cost (measure of unit efficiency and heat rate performance) and revenue (measure of unit capacity performance), aggregated over the past month. The bar chart also compares the actual performance against two benchmarks, Achievable ("Target") and Baseline ("Reference"), to provide a view of additional opportunities that exist for unit performance improvement.

The Benchmark bar chart compares the following three scenarios of unit performance:

Feature	Description
Achievable ("Target")	The <i>achievable benchmark</i> represents how the unit process could be performing if the equipment were operating at peak capability. The comparison with the actual benchmark thus provides a measure of additional improvement opportunities available. The achievable Net Heat Rate comes from PerformanceOpt's Achievable HR what-if simulation where all equipment are performing at peak under actual conditions and generating the set load. The Achievable Net MW comes from PerformanceOpt's Achievable MaxMW what-if simulation where all equipment are performing at peak and the unit is being pushed to generate maximum power subject to load-limiting conditions in the process, as provided by the plant
Actual	The <i>actual benchmark</i> represents how the unit process has actually been performing, or in other words, the fuel cost actually incurred and the revenue actually generated over the past month. The fuel cost comes from the Net Unit Heat Rate calculated by PerformanceOpt's as-is simulation and the unit cost of coal in \$/mmBTu provided by the plant. The revenue comes from the actual Net MW generated by the unit and the Location Marginal Price (LMP) of power provided by the plant or available from other sources
Baseline ("Reference")	The <i>baseline benchmark</i> represents the unit performance with respect to net heat rate/fuel cost and capacity/revenue that is used by the plant as a floor and can be based on, for example, budgetary goals, history, etc. The comparison with the actual benchmark thus provides a measure of performance improvement that has been obtained

A.4.3.1.1 Benchmarking

Benchmarking refers to using a "point of reference" to evaluate performance. There are many ways to develop these reference points. Past performance is one type of benchmark. Performance with respect to design is another. PerformanceOpt uses first-principles modeling based on equations of thermodynamics to describe the cause and effect relationships between operating variables and performance. Its models are used to simulate "what-if" scenarios and provide benchmarks for measuring performance.

A.4.3.2 Benefits table

The Benefits table (see "B" in Figure 95) shows, for the past month the fuel cost savings and the revenue increase that occurred under actual operation (by comparison to baseline operation) as well as the additional benefits that could have been achieved (by comparison to achievable operation).

The Benefits table has the following three columns:

Column	Description
Benefits	The unit performance measures, Fuel Cost and Revenue, from the benchmark bar chart
vs. Achievable	(Achievable minus Actual) in the benchmark bar chart
vs. Baseline	(Actual minus Baseline) in the benchmark bar chart

A.4.3.3 Unit Efficiency chart

The Unit Efficiency chart (see "C" in Figure 95) provides a time trend view, over the past month, of actual and achievable Fuel Cost per unit of generated power (\$/MWh). This is based on Actual Net Unit Heat Rate, computed by PerformanceOpt's "as-is" simulation and a comparison with the Net Unit Heat Rate computed by PerformanceOpt's Achievable HR what-if simulation for the scenario of all equipment performing at peak.

A.5 MaintenanceOpt Home Page Details

The Home Page provides an overview of the current operation of the MaintenanceOpt system through a summary of the identified anomalies, either detected by MaintenanceOpt or detected by other systems and managed in MaintenanceOpt, along with their likely impacts on unit performance. In addition, the MaintenanceOpt Home page also presents benchmarks to compare unit performance, in terms of efficiency and capacity, from the recent past.

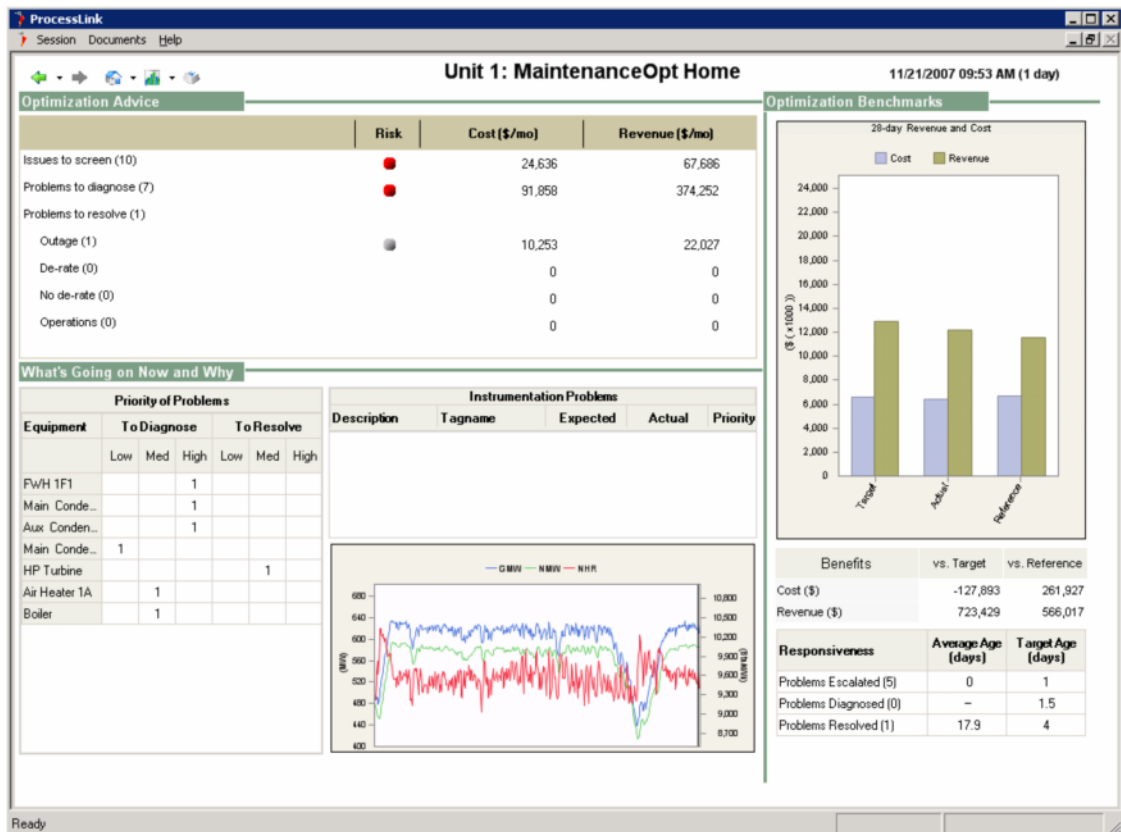


Figure 96 MaintenanceOpt's Home Page

MaintenanceOpt's Home Page is divided into three sections: What's Going on Now and Why (bottom left), Optimization Advice (top left), and Optimization Benchmarks (right side).

A.5.1 What's Going on Now and Why

The What's Going on Now and Why section of the MaintenanceOpt Home Page presents the following kinds of summary information:

- Distribution of current issues across equipment and their user-assigned priorities
- Instrumentation-related issues
- Trend of high-level unit performance measures – load and heat rate

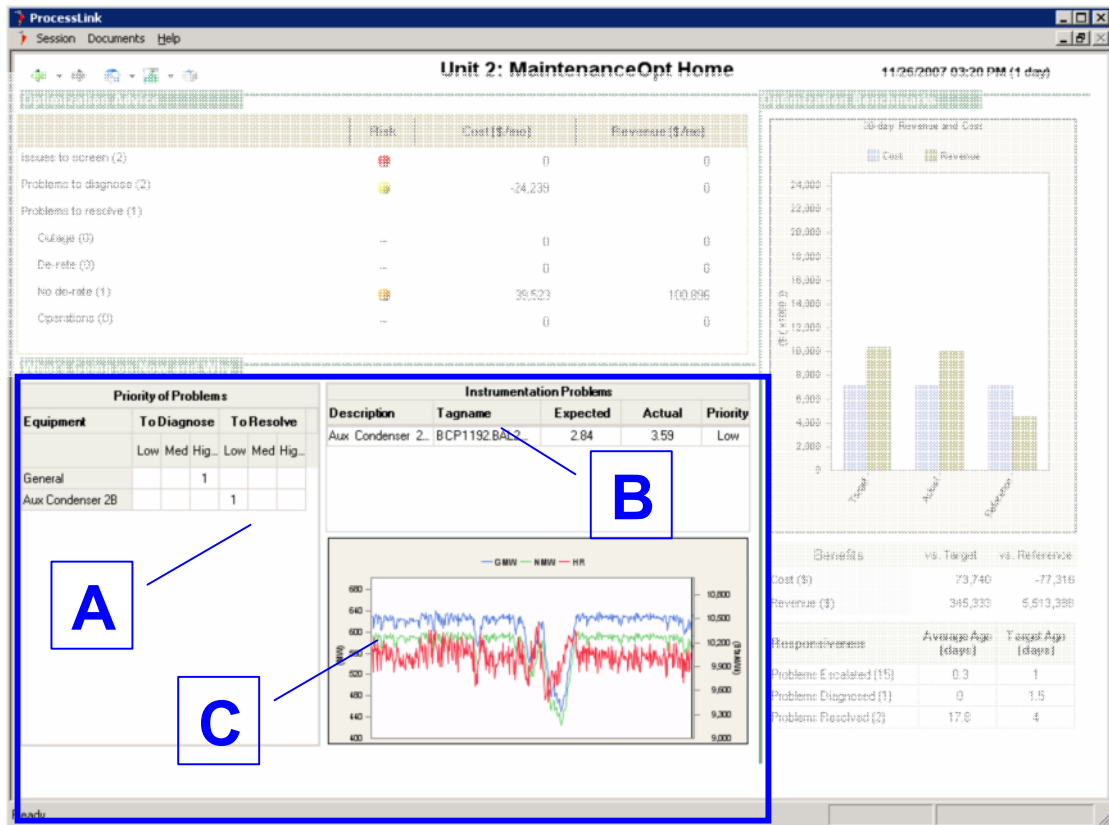


Figure 97 What's Going on Now and Why section on MaintenanceOpt's Home Page MaintenanceOpt's What's Going on Now and Why section has three sections: Priorities of Problems table (A), Instrumentation Problems table (B), and Unit Performance Trend Icon (C).

A.5.1.1 Priorities of Problems table

The Priorities of Problems table (see "A" in Figure 97) provides a summary view based on the current time snapshot of all MaintenanceOpt-managed issues grouped by:

- *Problem lifecycle state* – problems that have been validated/escalated by the user to the Problems to Diagnose category as well as problems that have further been diagnosed and are in the Problems to Resolve category.
- *Problem priority* – the user-assigned priority based on the importance and urgency of each problem. The priority classification is customizable but typical MaintenanceOpt implementations have three levels – High, Medium and Low.
- *Equipment* – equipment being monitored by one or more MaintenanceOpt triggers. The list of equipment displayed adjusts on the fly to only include specific equipment that currently has an active problem being tracked in the MaintenanceOpt application.

The information displayed in this table is the count of active issues in each group and aims to highlight if specific pieces of equipment require urgent attention. Users can get additional details about specific issues by using the toolbar to navigate to the Action Lists.

A.5.1.2 Instrumentation Problems table

The Instrumentation Problems table (see “B” Figure 97) provides a summary listing of all current MaintenanceOpt-managed issues that have gone through diagnosis and the user has determined to have a root cause related to instrumentation error. This table aims to highlight this class of problems for special attention by the appropriate plant staff.

The Instrumentation Problems table has the following five columns:

Column	Description
Description	MaintenanceOpt trigger that was diagnosed and ascribed to this instrumentation error
Tag name	ID in plant data historian of sensor that is monitored by the MaintenanceOpt trigger for anomaly detection
Expected	Value for the process variable predicted by the MaintenanceOpt models
Actual	Value for the process variable measured by the sensor
Priority	User-assigned priority for this problem

A.5.1.3 Unit Performance Trend Icon

The iconized chart in the bottom right of the What’s Going on Now and Why section (“C” in Figure 97) provides a quick trend view of unit-level performance measures, typically:

- Gross MV
- Net MW
- Net heat rate.

As with all other iconized charts in ProcessLink applications, the user can double-click on the chart to produce an expanded view that comes with many data manipulation/analysis capabilities.

A.5.2 Optimization Advice

The upper left section of the Home Page shows a summarized list of the issues currently being managed by the MaintenanceOpt system. The summary is presented in terms of the three stages of the MaintenanceOpt workflow:

- Issues to Screen
- Problems to Diagnose
- Problems to Resolve

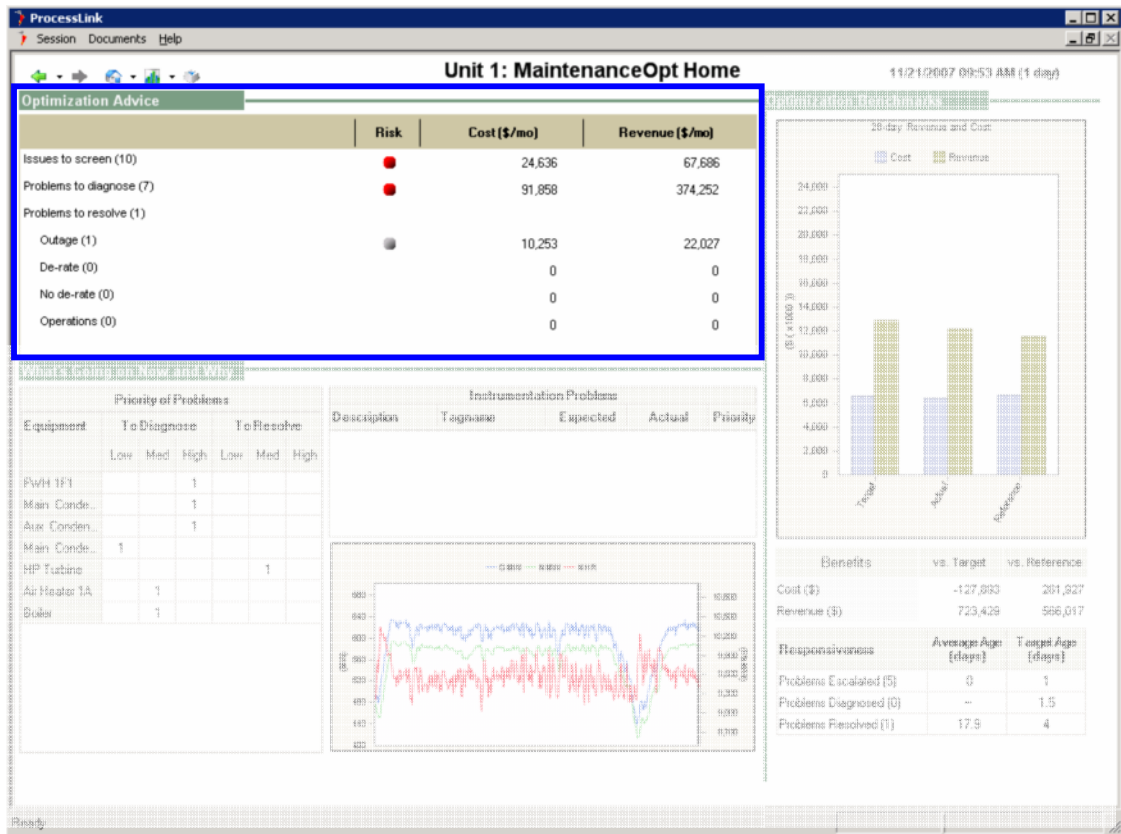


Figure 98 Optimization Advice section on MaintenanceOpt's Home Page

The summary information presented for each group of issues is the aggregation of:

Column	Description
Risk	Reliability risk, along a qualitative scale with four levels – None, Low, Medium and High
Cost	Heat rate impact, translated to fuel cost in \$/mo, using the unit fuel cost
Revenue	Capacity impact, translated to revenue in \$/mo, using the locational marginal price of power

In addition, Problems to Resolve are further sub-categorized based on the likely responsibility (Maintenance or Operations) and timing for the corrective action that has been identified:

- Operations
- No de-rate
- De-rate
- Outage

This tabular list is also hyper-linked to navigate via a single click to the specific Action List view for more details on the individual issues currently in that workflow step.

A.5.3 Optimization Benchmarks

The section on the right of the Home Page provides a high-level view of performance over the recent operational history of the unit and highlight opportunities that may be available for improvement.

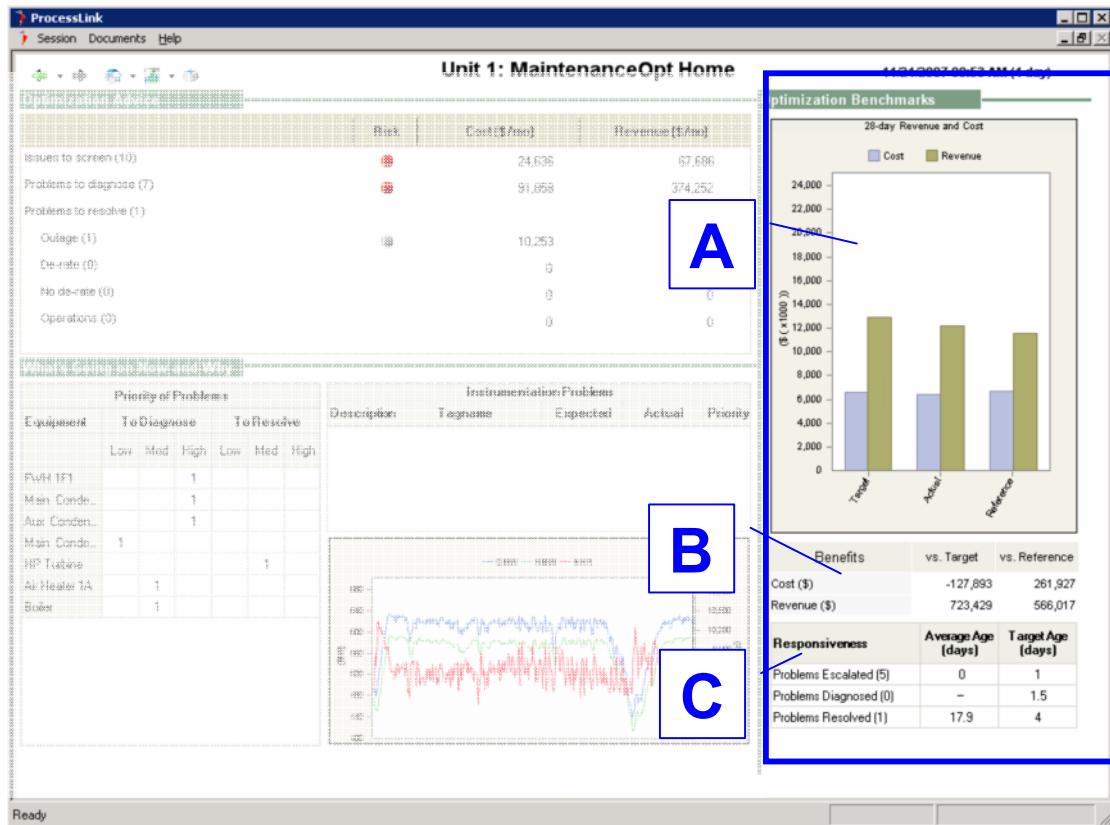


Figure 99 Optimization Benchmarks section on MaintenanceOpt's Home Page
MaintenanceOpt's Optimization Benchmarks section has three areas: Benchmark bar chart (A), Benefits table (B) and Workflow Responsiveness table (C)

A.5.3.1 Benchmark Bar Chart

The Benchmark bar chart (see "A" in Figure 99) shows you fuel cost (measure of unit efficiency and heat rate performance) and revenue (measure of unit capacity performance), aggregated over the past 28 days. The bar chart also compares the actual performance against two benchmarks, Achievable/Target and Baseline/Reference, to provide a view of additional opportunities that exist for unit performance improvement.

The Benchmark bar chart compares the following three scenarios of unit performance:

Feature	Description
Achievable / Target	The <i>achievable benchmark</i> represents the unit performance with respect to net heat rate/fuel cost and capacity/revenue that is used by the plant as a <i>target</i> and can be based on, for example, budgetary goals, history, etc. The comparison with the actual

Feature	Description
	benchmark thus provides a measure of additional performance improvement opportunities available
Actual	The <i>actual benchmark</i> represents how the unit process has actually been performing, or in other words, the fuel cost actually incurred and the revenue actually generated over the past 28 days. The fuel cost comes from the Net Unit Heat Rate calculated by the plant's performance monitoring system and the unit cost of coal in \$/mmBTu provided by the plant. The revenue comes from the actual Net MW generated by the unit and the Location Marginal Price (LMP) of power provided by the plant or available from other sources
Baseline / Reference	The <i>baseline benchmark</i> represents the unit performance with respect to net heat rate/fuel cost and capacity/revenue that is used by the plant as a <i>floor</i> and can be based on, for example, budgetary goals, history, etc. The comparison with the actual benchmark thus provides a measure of performance improvement that has been obtained

A.5.3.2 Benchmark table

The Benefits table (see "B" in Figure 99) shows, for the past 28 days the fuel cost savings and the revenue increase that occurred under actual operation (by comparison to baseline operation) as well as the additional benefits that could have been achieved (by comparison to achievable operation).

The Benefits table has the following three columns:

Column	Description
Benefits	The unit performance measures, Fuel Cost and Revenue, from the benchmark bar chart
vs. Achievable	(Achievable minus Actual) in the benchmark bar chart
vs. Baseline	(Actual minus Baseline) in the benchmark bar chart

A.5.3.3 Workflow Responsiveness table

The Workflow Responsiveness table (see "C" in Figure 99) provides a measure of efficiency, again based on a 28-day rear-looking view, of the problem lifecycle management process as operations, engineering, and maintenance groups take issues from identification to resolution.

The MaintenanceOpt workflow steps – escalation, diagnosis and resolution – provide the basis for the measures of process efficiency shown here:

Column	Description
Number of Issues	For each problem lifecycle stage, the number of problems processed over the past 28-day period
Average Age	For all the problems processed through each MaintenanceOpt workflow step in the past 28 days, the average time (in days) that the problem spent in that lifecycle stage.
Target Age	As a benchmark for this measure of efficiency, MaintenanceOpt

Column	Description
	also displays a corresponding Target Age, provided by the plant, as a goal for the detection-to-resolution lifecycle.