Visions from Another Dimension - How A Sequential Optimizer Got More Efficiency from an Existing Boiler with Less NOX

Abstract

If there are a thousand ways to skin a cat then operating a boiler, with all the potential parameter settings involved and results - many of which have environmental, legal, or economic implications - to track, undoubtedly has more. Discovering which of these meet load, satisfy emissions requirements, and minimize heat rate, however, can pose a real Gordian's Knot. Key to unraveling it is isolating the means to "think in N-Space" (where N is the total number of settings and results involved) as well as actually see where you're going. The discovery technology for computerized boiler optimization, tuning plant settings to meet emissions and operational constraints (LOI, end gas temperature, etc.) while minimizing heat rate, has been in use since the early 1990s. Much more recently, however, the complementary ability to readily visualize all outputs as well as operational settings of a boiler simultaneously – in a single screen, so that one can both see and understand where the optimizer has been – has come into play. This provides tremendous synergy for any optimization project since all aspects of boiler operation can be viewed at a glance and, for once, the optimization engine itself not strictly operated as a software "black box." In effect, the human intelligence which has run a steam generator (frequently for many years) can be brought to bear in tandem with the artificial variety of an empirical optimization system. This paper will explore the benefits inherent in such a combination, using actual data from sequential optimization of a tangential unit.

Summary of Combined Technology Benefits

- NOx reduced by over 30% while boiler heat rate decreased
- Limits of boiler operability under non-adaptive control clear
- Visibility of, and confidence in, optimizer's actions gained

Tough Technology vs. Tight Standards

Engaging the strengths of computer technology and human intelligence has proven key to driving decades old, coal-fired generating units to emissions performance their designers never contemplated. This breaks down into two distinct phases for an effective project, which generally encompasses multiple measures of emissions output, boiler safety considerations, and economic effects on unit operation. The first employs sequential optimization, initially set-up and continuously utilized by those who run the unit (who knows its tricks better?), while the second incisively analyzes boiler behavior after as well as during optimization to assure that operational and controllability issues remain satisfied. This latter brings the human proclivity for pattern recognition to bear by explicitly engaging technology allowing project personnel to "see in N-Space." Both are dependent on relatively new, but well proven, mathematics now implemented in commercial software.

Sequential Optimization

Successfully in use since the mid-1980s, sequential optimization (as embodied in a system known as Ultramax®),[1] relies upon the notion of producing information about a system while optimizing that system's performance. Hence, trials concerning how to best run a boiler are conducted one (or at most, a handful) at a time, with results from those tests – whether trial recommendations were followed faithfully, modified, or completely ignored - fed back into the sequential optimizer. With that shot of fresh data Ultramax remodels, fitting local second order polynomials as an optimum is approached (less sophisticated and data hungry models are used when fewer observations are available). This is followed by generation of new advice designed to both step toward (or onto) the optimum – such as minimum heat rate with constrained NOx in the project discussed below – while generating more information about how to further tune the plant. Hence a jagged, but remarkably efficient, hill climbing approach is followed, with internal safeguards in play to preclude sticking on a local optimum or stepping into foul territory where excessively poor results are produced, as ideal performance is continuously sought. Consequently this system behaves as a multiple-input, multiple output adaptive controller, adjusting for changes to optimization requirements (such as a maximum acceptable level of NOx emissions) and external conditions (wet or off-quality coal) while learning to optimally control a plant in the act of optimizing it. Hence, such is a very powerful, albeit "black box" tool which, to run, one need only know process inputs, outputs, and external influences as illustrated in the "Optimization Plan Formulation" of Figure One, below. That formulation itself documents a 570MW, twin furnace, tangentially fired coal unit whose optimization and analysis will serve as the case study covered by the remainder of this paper. Moreover, the relatively new technology of Parallel Coordinates will be engaged to visualize and track the path of this relentless optimizer, allowing visibility of previously undetectable local optima and parameter regions where optimization goals can never be met.

ULTRAMAX OPTIMIZATION PLAN FORMULATION

Full Load Normal Operation. 14:01 14 OCT

VAR #	NAME	UNITS	TYPE	MODE	TRSF	PRIOR REGION LO	PRIOR REGION HI	CONSTRAINTS LO	CONSTRAINTS HI
2	O2_Setpoint	%	1	H	0	3.3	3.6	2.6	4.
3	F_Mill_Bias	%	1	H	0	-20.	20.	-20.	20.
4	Aux_Air_4_5	%	1	H	0	40.	75.	25.	100.
5	FA_Damp_E_F	%	1	H	0	80.	100.	40.	100.
7	Aux_Air_0_1	%	1	H	0	40.	55.	25.	100.
8	FA_Damp_A_B	%	1	H	0	80.	100.	40.	100.
9	Aux_Air_2_3	%	1	H	0	40.	50.	25.	50.
10	FA_Damp_C_D	%	1	H	0	80.	100.	40.	100.
11	HeatRate	Btu/kWh	6	H	0	9.13	9.37	0.	9.3
12	RH_Temp_A	Deg. F	5	H	0	.99	1.007	.98	1.01
13	RH_Temp_B	Deg. F	5	H	0	1.	1.029	.99	1.03
14	SH_Temp_A	Deg. F	5	H	0	.99	1.026	.99	1.03
15	SH_Temp_B	Deg. F	5	H	0	.987	1.014	.98	1.3
16	LOI	%	5	H	0	2.	3.	0.	4.
17	Tube_Temp	Deg.f	5	H	0	.85	1.04	.8	1.04
18	O2_Econ_A	%	5	H	0	3.3	3.6	2.5	4.
19	O2_Econ_B	%	5	H	0	3.3	3.6	2.5	4.
20	NOx	lbs/MMBt	5	H	0	.6	.8	0.	.58
21	CO	ppm	5	H	0	5.	15.	0.	100.
22	Load	MW	4	H	0	500.	575.	0.	0.
23	Opacity	%	5	H	0	2.	7.	0.	10.
24	WF_Diff_A	Inches H	5	H	0	3.	5.	2.5	5.
25	WF_Diff_B	Inches H	5	H	0	3.	5.	2.5	5.
26	Reheat_Spray	Klbs/hr	5	H	0	38.	118.	0.	195.
27	APH_Gas_In	Deg. F	4	H	0	693.	721.	0.	0.
28	APH_Air_Out	Deg. F	4	H	0	570.	592.	0.	0.
29	Steam_Flow	Mlbs/hr	4	H	0	3.8	4.	0.	0.

PAR(50) = -1 MINimizing variable # 11 HeatRate (Type 6)

Figure 1: Sequential Optimizer Problem Formulation. Note that "Type 1" variables are adjustable boiler controls, "Type 4" variables boiler outputs to be mathematically modeled (i.e., fit to a polynomial) but not considered in optimization, "Type 5" boiler outputs which are modeled and subject to optimization, and the lone "Type 6" variable to be minimized

Parallel Coordinates

Since knowing where a black box optimizer is taking your plant remains generically necessary – as opposed to simply nice – insight into behavior of the entire system as control adjustments are made becomes key. Developing such is crucial if operating personnel are to have confidence in the results and engineering staff successfully evaluate whether their equipment is being pushed too hard for consistent and sustainable operation… thus making the ability to "see in N-Space" nearly indispensable. This amounts to exploiting Parallel Coordinates visualization technology[2] recently commercialized as Curvaceous Visual Explorer™ to see all process variables (typically 20-50) simultaneously across the space of a single chart or computer screen. Such requires an N-D (N-dimensional) to 2-D (two-dimensional) coordinate transform:

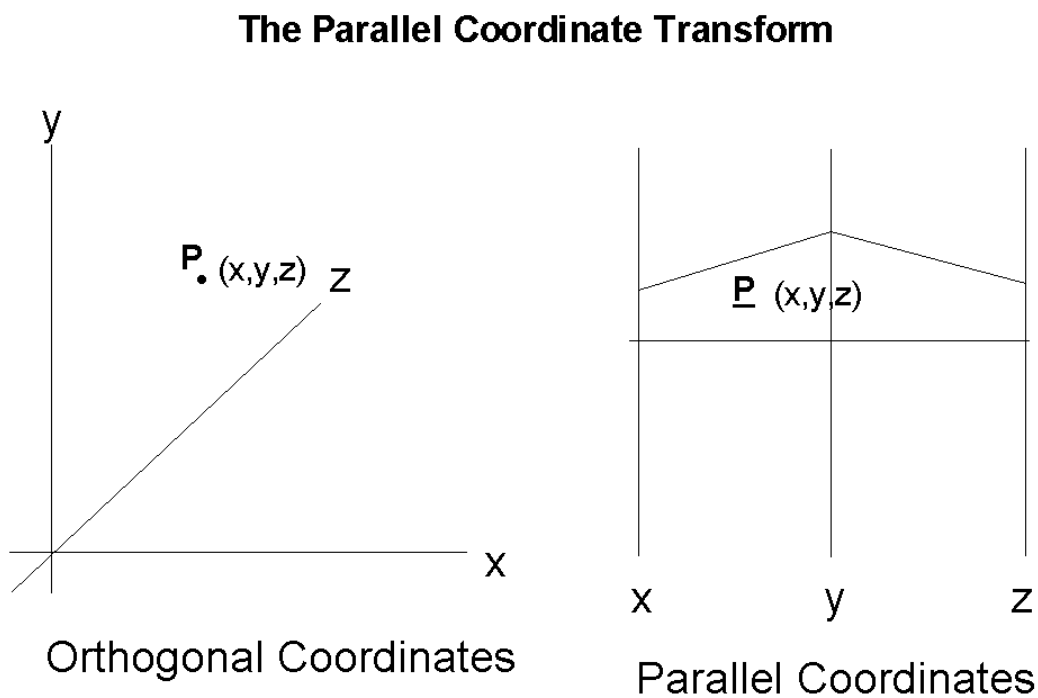

Figure 2: Orthogonal (Cartesian) to Parallel Coordinate Transform.

Hence a point (x,y,z) that plots in space on 3-Dimensional, Cartesian Coordinates (where axes are arranged at right angles to one another), displays as separate values joined by line segments on parallel X, Y, and Z coordinate axes. This replaces the point in space on Cartesian Coordinates – where three is the maximum number of axes that can be represented – with a polygonal line on Parallel Coordinate axes… where there is no limit to the number of axes that can be represented. Consequently, a single run from the optimization activity discussed below, with 18 separate variables including optimization run number, plots on Parallel Coordinate axes as Figure Three.

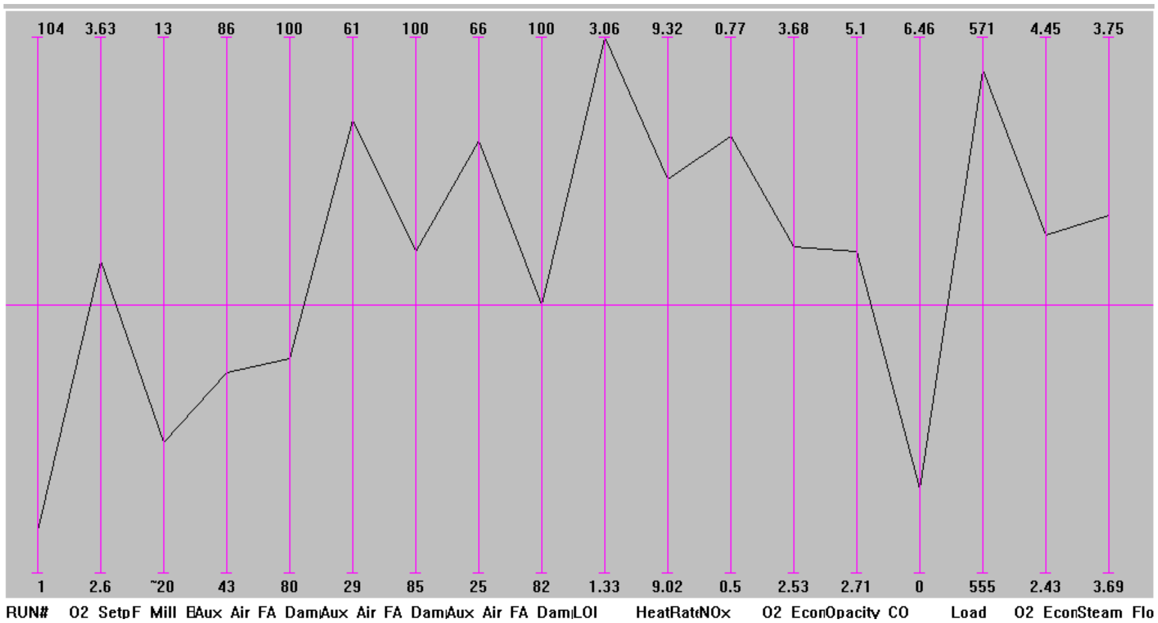

Figure 3: Parallel Coordinate Plot of a Single Optimization Run.

Optimization

The principal thrust of coal fired boiler optimization, at least initially, is to reduce NOx per MMBtu while holding other emissions, efficiency, and boiler safety measures – CO, stack gas opacity, loss on ignition (LOI), heat rate per KW-h generated, tube temperature, etc. – within acceptable bounds. With low NOx achieved, its performance can then be held at acceptable levels while economic measures, primarily heat rate, are minimized and other performance measures maintained within their constraints. Obviously such has to be accomplished at or near full load conditions, since most plant operation will be there, while additional optimization at conditions of interest (e.g., selected coal pulverization mills and their burners out of service) is also useful. However this optimization, where new settings for boiler operation will be empirically devised (often fairly quickly, particularly if closed loop integration has been employed) will almost invariably place the unit in uncharted territory… frequently unknown even to its manufacturer. Hence it is useful to both monitor plant behavior and analyze the optimization data record for assessment of operability issues, which frequently would not become apparent within the span of initial sequential optimization, associated with the new region. The NOx reduction and plant performance tuning project reported here followed almost exactly that pattern, exploiting sequential optimization with considerable success at a variety of conditions before switching to steady state operation for analysis of and reaction to longer term issues.

Boiler Tuning

After reviewing boiler characteristics, typical operational practices, and project objectives with plant stationary engineers as well as operators an optimization plan formulation typified by that of Figure One is input to the sequential optimizer (Ultramax). One will note that this situation is technically underdetermined – there are more results to be constrained or minimized than independent variables with which to drive them. Some, however, will prove related and hence controllable with the same set of independent parameters ("Type 1" variables) while others move easily to a position well within their constraints and become essentially inactive to the optimization problem. Hence there will be sufficient knobs to deliver the results desired. Consequently, conducting sequential optimization after an initial (baseline) run at standard operating conditions yielded the classic NOx Reduction Profile depicted below by Figure Four.

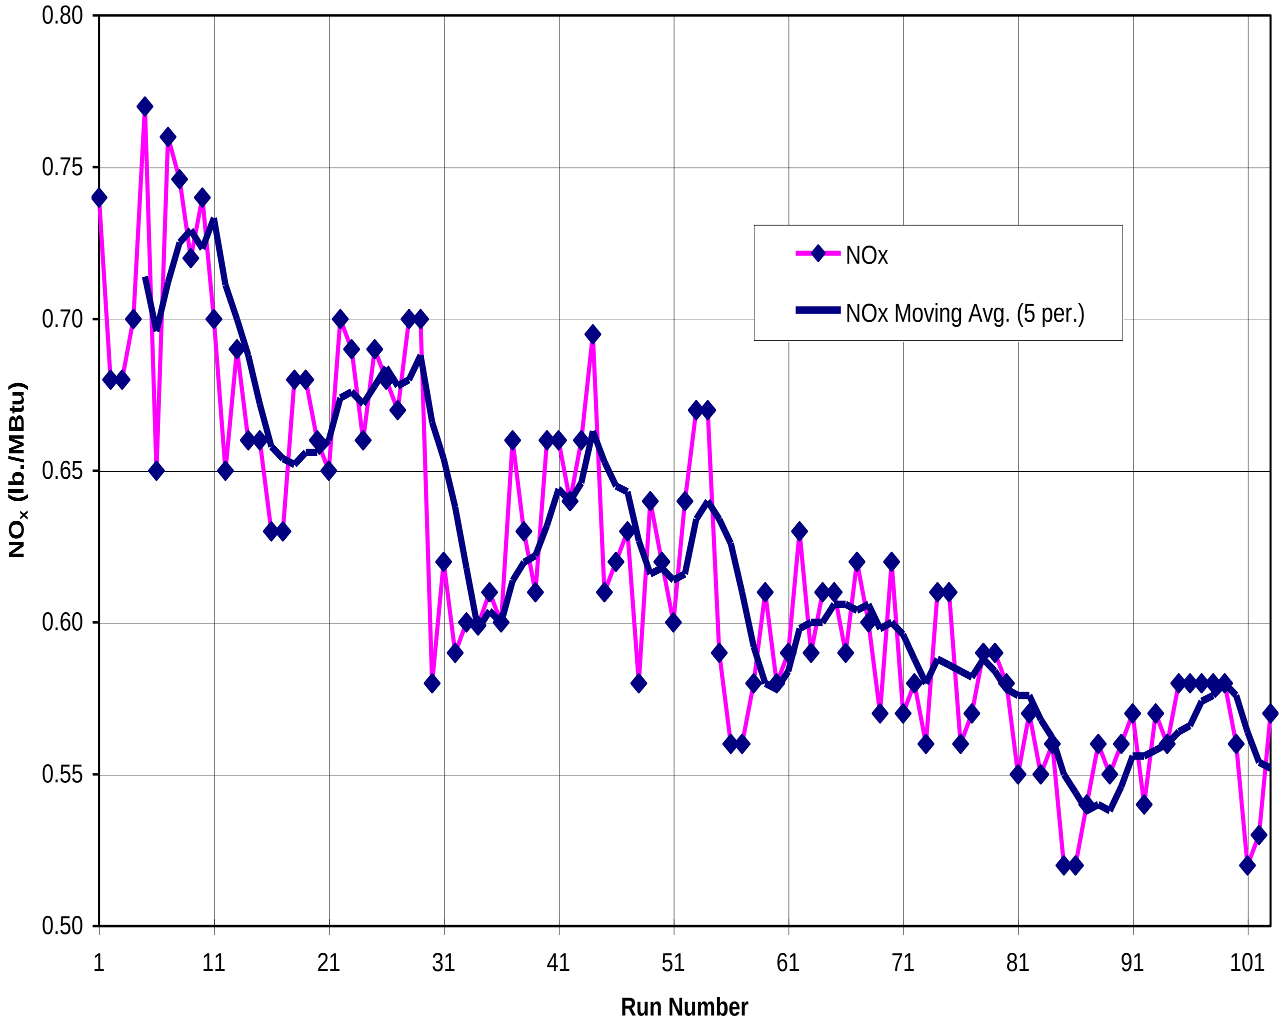

Figure 4: NOx Reduction Profile.

The steady decline in NOx is quite apparent, and typical for this character of project, with the only hitch being slagging noted around the burners as O2 setpoint was reduced. This led to the lower constraint of 2.6% on O2 setpoint reflected in Figure One being introduced mid-stream during optimization. Further, once it became recognized that very low levels of NOx (down to .56 lbs/MMBtu, with more clearly possible) had been achieved by run 85, that output was shifted to a constrained ("type 5") variable with an upper ceiling of .58 lbs/MMBtu and Heat rate subjected to minimization. By optimization run 104, a three hour test burn at best conditions, optimization was secured for full load conditions where all burners were in service with the following results:

- NO_x was reduced from a baseline of 0.74 lbs/MMBtu to 0.50 lbs/MMBtu, a reduction of 32%.
- Heat rate was reduced from 9,255 Btu/kWh to 9,023 Btu/kWh, a reduction of 2.5 %.
- All other performance and emissions constraints remained satisfied.

Similarly, conducting 37 further optimization runs at or near full load but with upper (F-mill) burners out of service, which simulates an over-fire air condition, yielded:

- NO_x reduction from 0.64 lbs/MMBtu to 0.43 lbs/MMBtu, amounting to 33%.
- Heat rate reduction from 9398 Btu/kWh to 9238 Btu/kWh, or 1.7%.
- Satisfaction of all other performance and emission constraints.

Finally, 46 more optimization runs with F-mill burners in service but lower (A-mill) burners off-line produced:

- NO_x reduction from 0.76 lbs/MMBtu to 0.60 lbs/MMBtu, a 22% decrease.
- Heat rate reduction from 9,082 Btu/kWh to 9,003 Btu/kWh, or 0.87%.
- Continued satisfaction of all other performance and emission constraints.

Since all burner configurations subjected to optimization, even those with inactive top or bottom rows (which actually did not have sufficient runs to confirm practical optimization), showed significant improvement over baseline conditions the next question was how durable such results would be.

Parallel Coordinate Analysis

Answering the durability question effectively requires a two-pronged approach. Simply operating the plant at optimized settings and watching the outcome – which in this case required a pullback on O2 setpoint to 3.0%, as further operational issues were detected – is the obvious first step. Additional to this, however, significant insight can be gained by bringing to bear the ability to "see in N-Space" through Parallel Coordinate analysis. Although this may have been engaged on historical data, typically to isolate arguments and initial ranges ("prior" regions) for Ultramax, or for observational purposes during optimization itself, its role in deducing plant stability at best settings from optimization data records is particularly key. The objective here is to determine whether a plant can be conventionally operated, with all boiler parameters independently set in ranges about the new operating point once adaptive control is switched off… or whether controllability issues will force other measures. Simple though optimization sounds, as boilers (or systems of any sort) are pushed beyond their design condition intricate pathways – which a computer's artificial intelligence can thread but a person's cannot – frequently appear in operational ranges which produce premium performance. These parameter ranges are then no longer simply connected, which effectively derails the notion of independently setting boiler parameters within a band around their operating point, and requires that sophisticated control intelligence be left in play while plant operators basically "take their hands off." As the latter can produce other effects, it is wise to clearly understand a plant's limits before implicitly taking such decisions.

For this project, using Parallel Coordinates to analyze optimization data associated with an A-mills and burners out of service condition through the capabilities of Curvaceous Visual Explorer (CVE) revealed that the boiler was well behaved at the load range and NOx levels tested. Hence, independent parameters were simply connected in their influence upon operational results, while unusual interactions as well as intricate behavior among variables were not to be expected. Black Holes [3], or voids in otherwise usable parameter ranges where desired performance could never be obtained were not apparent, and once optimized to low NOx, low heat rate performance the system should be capable of operating well in open loop. However, turning to optimization data with all burners in service and engaging CVE's cluster analysis facility, a means to discern numerically similar groupings of boiler performance data, brought further insight. Such revealed that Ultramax basically operated to minimize NOx in ten steps – ten clusters in NOx behavior (numbered 0-9) were found – refining its knowledge of the system at each step until enough data was acquired to take the next. In this instance, both the lowest and highest NOx levels recorded were unique, hence mapping them into the "no cluster" (value = -1) range on the NOxCls axis of the Parallel Coordinate plot in Figure Five.

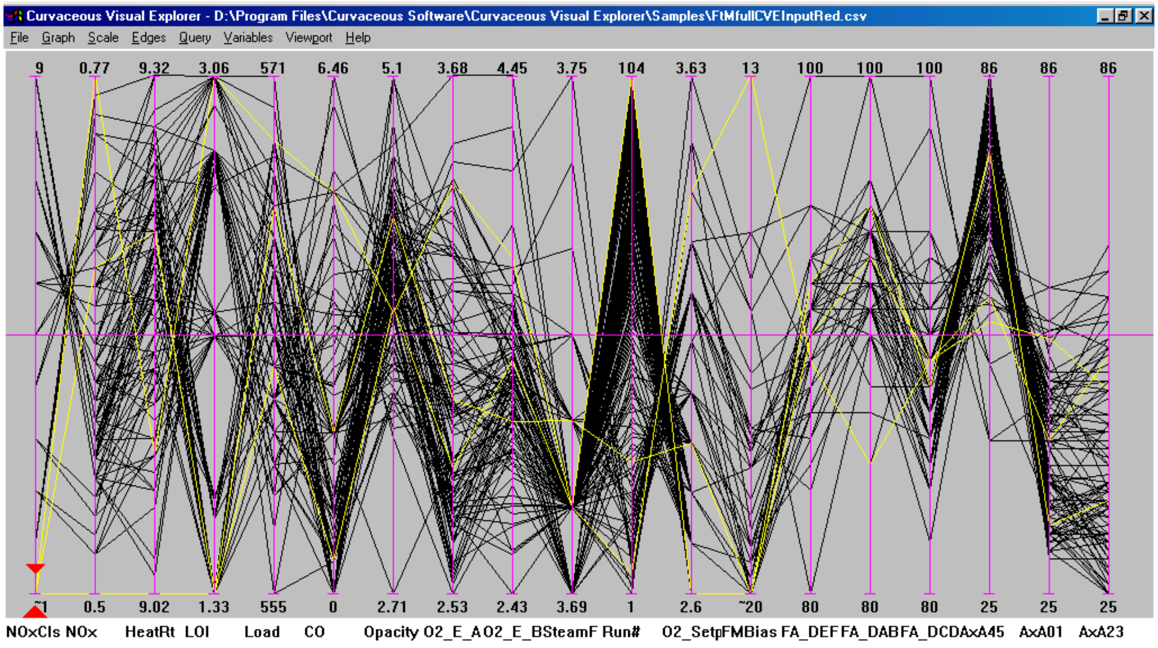

Figure 5: Parallel Coordinate Plot of Optimization Data Including Cluster Axis (leftmost).

Since this represents entirely normal Sequential Optimization behavior the next analysis step was simply to isolate runs of both low NOx "AND" low heat rate, as diagrammed by Figure Six:

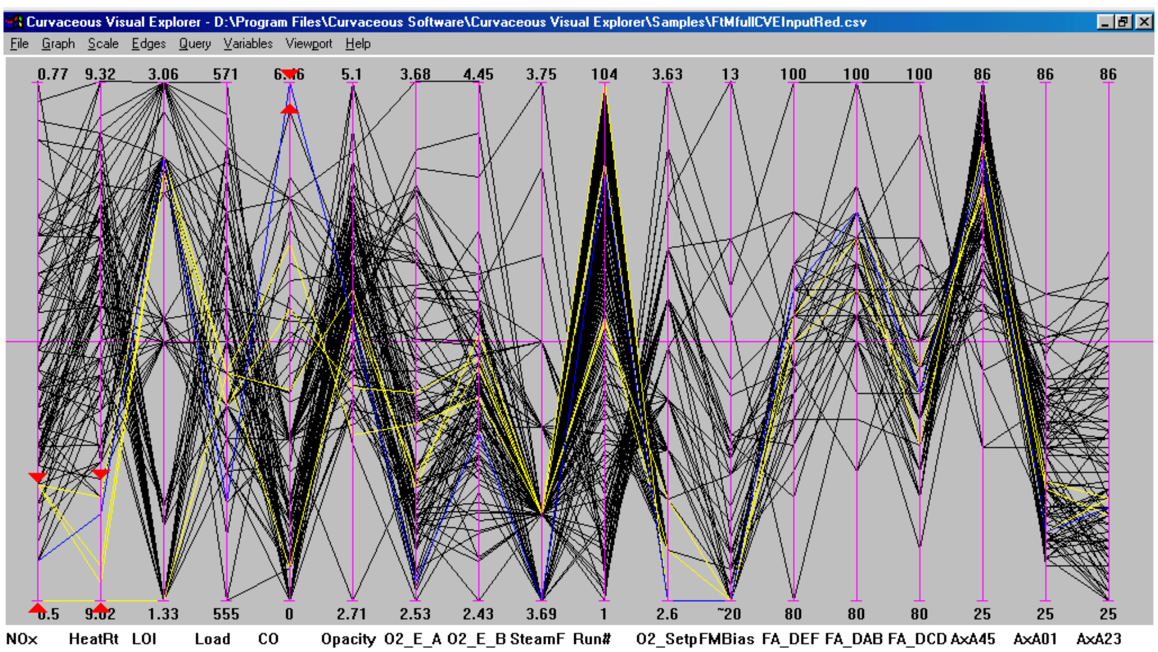

Figure 6: Parallel Coordinate Plot of Optimization Data With Low NOx "AND" Low Heat Rate Runs Selected. Note That High CO Observation Is Highlighted in Blue.

Unfortunately, of the five runs this selects, one (marked in blue) has the highest CO value in the dataset… so it's out as a potential operating point. Further, to obtain both low NOx and low Heat rate we must either accept the lowest NOx, lowest Heat rate, and lowest LOI observation in the dataset or tolerate being stuck with a relatively high LOI (roughly 2.69-2.81), as indicated in Figure Seven below.

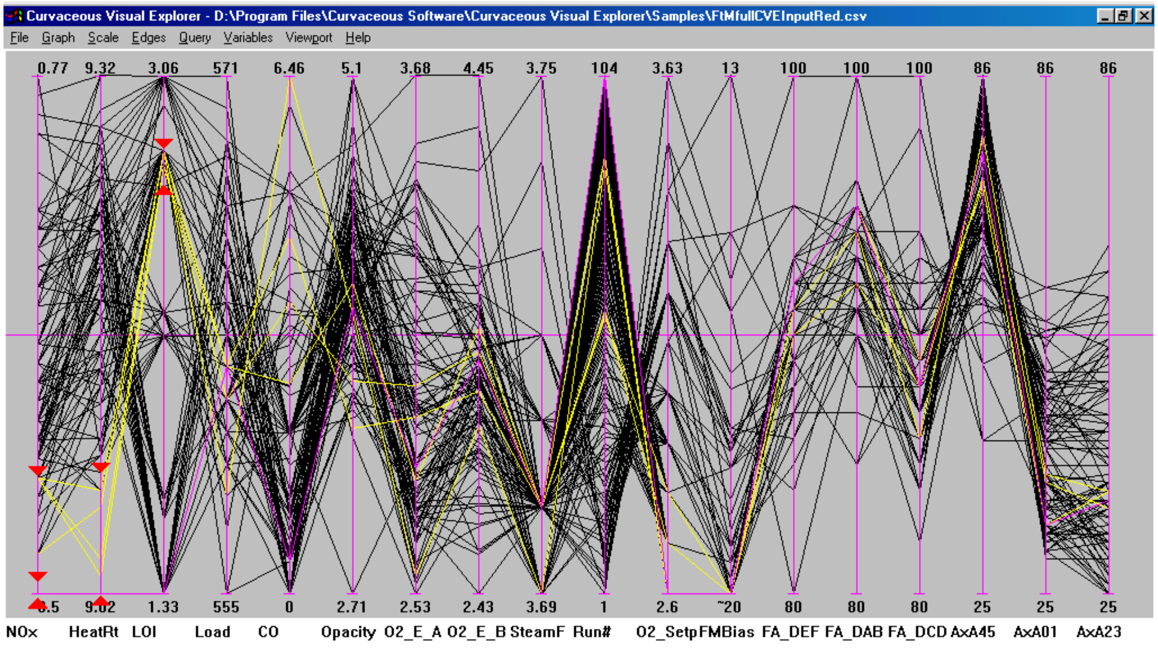

Figure 7: Parallel Coordinate Plot of Optimization Data With Low NOx, Low Heat Rate Runs Selected In Yellow But Best Point (which did not fall into a cluster and must be confirmed) Highlighted In Magenta.

Given the "wild duck" nature of the lowest NOx, lowest Heat rate, and lowest LOI run selected in the magenta query, that decision is not entirely obvious and either continued Sequential Optimization or a few more runs at that query's conditions are in order. Dismissing secondary queries (those in blue and magenta of previous screens) in Figure Eight and looking more closely at simply low NOx, low heat rate results, however, suggests that at such conditions we may be approaching this boiler's limits of stability.

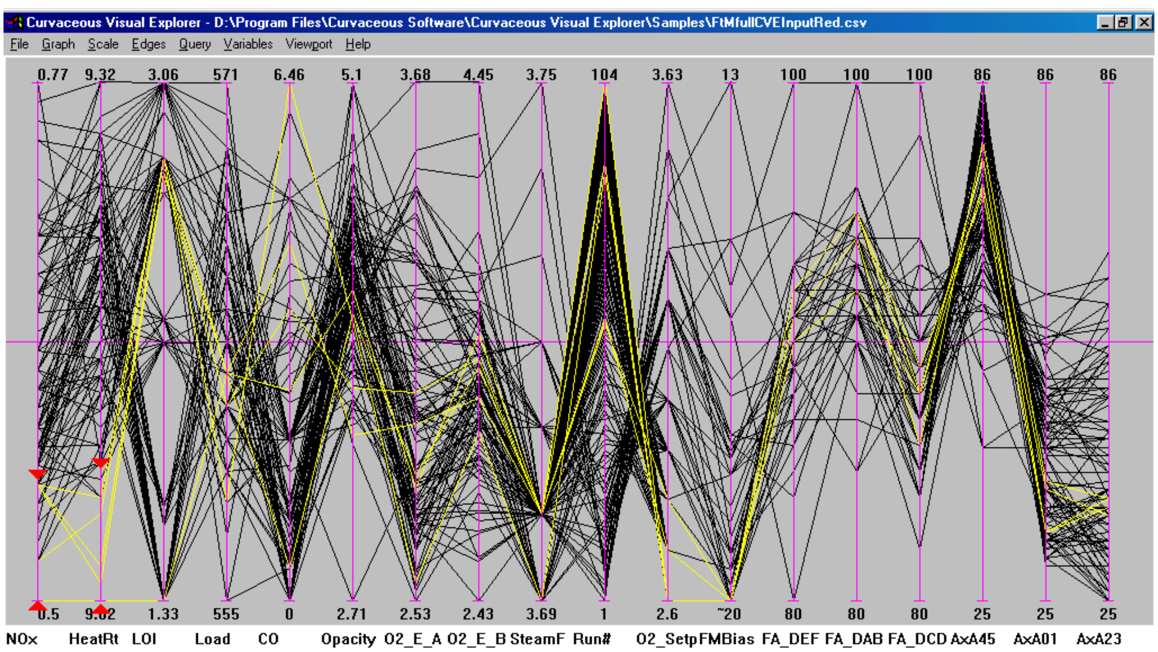

Figure 8: Parallel Coordinate Plot of Optimization Data With Low NOx, Low Heat Rate Runs Selected In Yellow.

Note the banding – yellow "good" zones interspersed with higher NOx or heat rate black - which is starting to develop between FA_DCD and AxA45 as well as AxA45 and AxA01, with the former actually displaying a lower Damper setting to higher Auxiliary Air and higher Damper setting to lower Auxiliary Air crossing pattern. Note also that FA_DCD, AxA45, and AxA01 are length compressed nomenclatures (to fit 18 variables across the space of one screen) for Fuel/Air Dampers C&D, Auxiliary Air Dampers 4&5, and Auxiliary Air Dampers 0&1, respectively. Since black represents poor performance and yellow good it's becoming clear that to maintain very low NOx and heat rate – notably lower than with A-mills and burners out of service – a control system will have to thread complex paths between at least some independent variables. Moreover, ranges for those independent variables which deliver optimum boiler performance will no longer be simply connected - the crossing pattern being one instance of this – but contain "Black Holes" where premium system performance can never be achieved. Thus, to maintain very low NOx, low heat rate, and potentially low LOI results in the presence of load, weather, and fuel (coal quality) changes this boiler will have to engage continuous, closed loop control… Since leaving the Sequential Optimizer, which is adaptive in nature, in place to run continuously is the simplest solution to that problem such is commonly engaged. Taking a less aggressive stance and tolerating higher NOx, heat rate and LOI, however, should allow open loop operation, as in the A-mills and burners out of service case. Boxing visible queries for low NOx, low heat rate performance, where CVE's boxing algorithm acts to deduce the lowest number of independent variables and their ranges necessary to include all selected observations, adds further evidence to this argument, depicted in Figure Nine below:

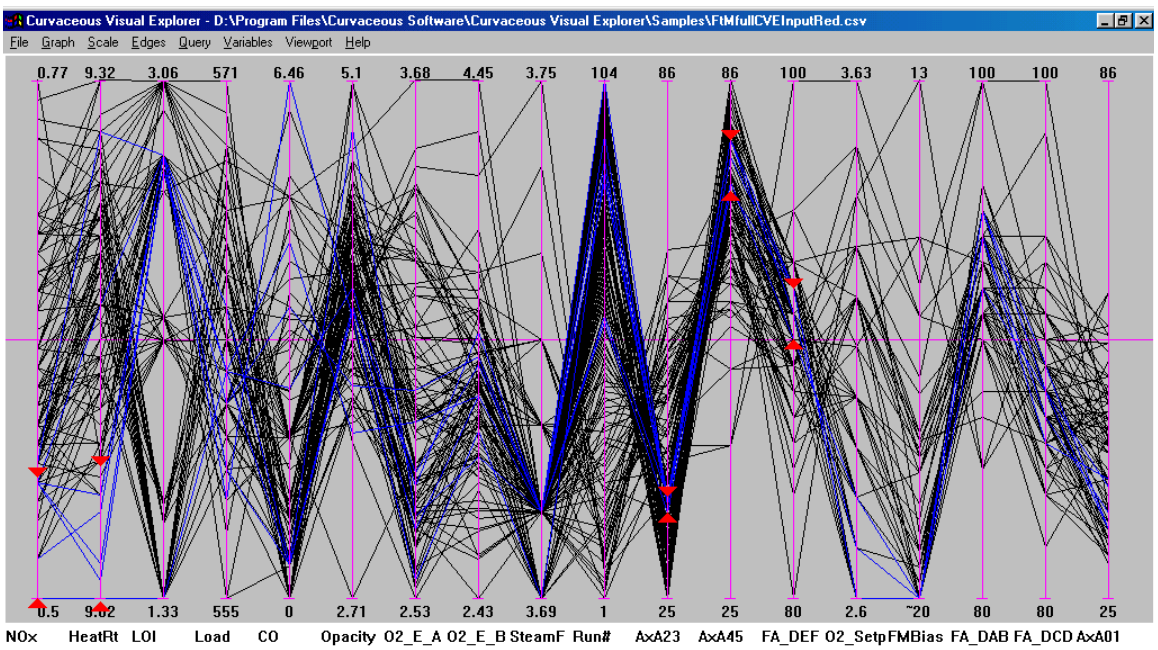

Figure 9: Parallel Coordinate Plot of Optimization Data Following "Boxing" Algorithm Usage.

At even a cursory inspection one will immediately note that the boxing algorithm was unable to deduce any set of independent and external variable ranges selecting only the five desired low NOx, low heat rate observations. The relationship between operational variables in producing desired boiler performance has become one that is no longer simply connected and extra runs (two, in this case) are inevitably included. Moreover, for producing low NOx, low heat rate results variables of significance have shifted from the less demanding, simply connected regime where lower pulverization mills and burners were out of service to AxA23, AxA45, and FA_DEF, in that order of influence.

Conclusions

As should be clear from the above, Sequential Optimization combined with Parallel Coordinate techniques amount to a 1-2 punch when teamed to reduce both stack emissions and the heat required to generate power. Although uniquely effective independently, and powerful process optimization technologies in their own right, together they successfully harness the numerical punch of artificial intelligence with the acumen and knowledge of a human observer. Even against thorny boiler optimization problems, where results must not simply be achieved on a transitory basis but held against the inherent variability of fuel, weather, load, and equipment, their relentless optimization and N-Space observation are effectively unstoppable. Burning fuel will always generate smoke, just as making an omelet requires breaking eggs… but with the right technology it will be much cleaner smoke.

References

1. Hoyte, D.W., "Computers - Optimization of Empirical Processes," Instrument Engineers' Handbook (3rd edition), Process Control Volume, 1995, pp. 835-839.
2. A. Inselberg, "The Plane With Parallel Coordinates," The Visual Computer, Volume 1, pp. 69-91, 1985 (published by Springer-Verlag).
3. R W Brooks, "New Dimensions for Improving Refinery Profitability," Petroleum Technology Quarterly 4, 1, 1999 (also available from www.curvaceous.com).