FOSTER-MILLER, INC

350 Second Avenue
Waltham, MA 02451-1003

781-684-4000
781-890-3489 fax

Date: 4/24/06

To: NETL AAD Document Control
U.S. Dept. of Energy
Laboratory - Bldg. 921
P.O. Box 10940
Pittsburgh PA 15236-0940

From: Roger Demler

Subject: Contract No: DE-FC26-03NT41846
Foster-Miller No: DDOE.040069

Report Type: Final Technical Report
Reporting Period: 10/1/03 to 9/30/05

Attention: Captain Erin Pettyjohn

Enclosed is the subject report for the above referenced contract. Please feel free to contact the person below if you require additional details.

Roger Demler
Senior Engineer

Email: rdemler@foster-miller.com
Work: 781 684-4673

Distribution:

Front Matter to S. Casey on 4/21/06.zg
NETL AAD -2 Copies
Contracts, Dill, Casey, Hill, Demler

FMI Report No: DOE-040069-1790

Sensor for Individual Burner Control of Coal Firing Rate, Fuel-Air Ratio and Coal Fineness Correlation

Final Technical Report

Reporting Period: 10/01/2003 - 09/30/2005

R. Demler

April 2006

Contract Number: DE-FC26-03NT41846

Foster-Miller, Inc.
350 Second Avenue
Waltham, MA 02451

Disclaimer

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

1. INTRODUCTION

1.1 Background

Accurate, cost-efficient monitoring instrumentation has long been considered essential to the operation of power plants. Nonetheless, for the monitoring of coal flow, such instrumentation has been sorely lacking and technically difficult to achieve. With more than half of the electrical power in the United States currently supplied by coal, energy generated by this resource is critical to the US economy. The demand for improvement in this area has only increased as a result of the following two situations: First, deregulation has produced a heightened demand for both reduced electrical cost and improved grid connectivity. Second, environmental concerns have simultaneously resulted in a need for both increased efficiency and reduced carbon and NO_x emissions.

A potential approach to addressing both these needs would be improvement in the area of combustion control. This would result in a better heat rate, reduced unburned carbon in ash, and reduced NO_x emissions. However, before feedback control can be implemented, the ability to monitor coal flow to the burners in real-time must be established. While there are several "commercially available" products for real-time coal flow measurement, power plant personnel are highly skeptical about the accuracy and longevity of these systems in their current state of development. In fact, following several demonstration projects of in-situ coal flow measurement systems in full scale utility boilers, it became obvious that there were still many unknown influences on these instruments during field applications. Due to the operational environment of the power plant, it has been difficult if not impossible to sort out what parameters could be influencing the various probe technologies.

Additionally, it has been recognized for some time that little is known regarding the performance of coal flow splitters, even where rifflers are employed. Often the coal flow distribution from these splitters remains mal-distributed. There have been mixed results in the field using variable orifices in coal pipes. Development of other coal flow control devices has been limited.

An underlying difficulty that, to date, has hindered the development of an accurate instrument for coal flow measurements is the fact that coal flow is characterized by irregular temporal and spatial variation. However, despite the inherent complexity of the dynamic system, the system is in fact deterministic. Therefore, in principle, the coal flow can be deduced from the dynamics it exhibits. Nonetheless, the interactions are highly nonlinear, rendering standard signal processing approaches, which rely on techniques such as frequency decomposition, to be of little value. Foster-Miller, Inc. has developed a methodology that relates the complex variation in such systems to the information of interest. This technology will be described in detail in Section 2.

A second concern regarding the current measurement systems is installation, which can be labor-intensive and cost-prohibitive. A process that does not require the pulverizer to be taken off line would be highly desirable. Most microwave and electrostatic methods require drilling up to 20 holes in the pipe, all with a high degree of precision so as to produce a proper alignment of the

probes. At least one electrostatic method requires a special spool piece to be fitted into each existing coal pipe. Overall, these procedures are both difficult and very expensive.

An alternative approach is pursued here, namely the development of an instrument that relies on an acoustic signal captured by way of a commercial accelerometer. The installation of this type of sensor is both simpler and less invasive than other techniques. An accelerometer installed in a pipe wall need not penetrate through the wall, which means that the system may be able to remain on line during the installation. Further, due to the fact that the Dynamical Instruments technology, unlike other systems, does not rely on uniformity of the air or coal profile, the installation location need not be on a long, straight run of pipe. In fact, an optimal signal is obtained near a pipe elbow. This is fortuitous, as bends are often more accessible on pipes in a power plant than straight sections. In contrast to measurement systems that rely on the uniformity of the air and coal profile, the accuracy of the system under development will not compromised by varying levels of flow uniformity.

Additionally, the ease of installation of an instrument relying on the use of an accelerometer is such that the resulting instrument can actually be implemented so that it can be portable. For those plants that do not require continuous online monitoring, the instrument could be moved from unit to unit, reducing costs even further. As an example, base-loaded plants may only be interested in optimizing flow balance to the burners under a single condition, full load. The use of a portable device, which could be moved from mill to mill or unit to unit, to achieve a balanced flow, would be highly appealing for plants with a large number of burners.

Overall then, the instrument under development should result in more accurate coal-flow monitoring, simpler installation and portability. This in turn would allow for an increased level of control for a broader base of plants and therefore serve to improve cost and energy efficiency.

1.2 Dynamical Instruments

Foster-Miller has developed a nonlinear technique that is very sensitive to subtle changes in system dynamics, enabling flow conditions to be extracted from time series measurements of the system dynamics. This technique, called Dynamical Instruments, involves the direct characterization of the dynamics of a time series signal as a set of statistics. The resulting vector of statistics comprises a “signature” of the dynamics, which can then be correlated to the system condition (in this case, coal flow).

Figure 1 provides an example of the type of results that have proven achievable using the Dynamical Instruments technique. This graph compares the actual and predicted flow rates of liquid in liquid-gas two-phase flows, in which the prediction was based on an ultrasonic measurement of the liquid film thickness. The points cluster near the Y=X line, indicating that the prediction is well correlated to the actual liquid flow. In fact, two standard deviations of error in this correlation comprise 2.9% of full scale. Since there is no first-principles instrument that can determine this liquid flow rate without first separating the liquid and gas into separate flows, this is an exceptionally good result.

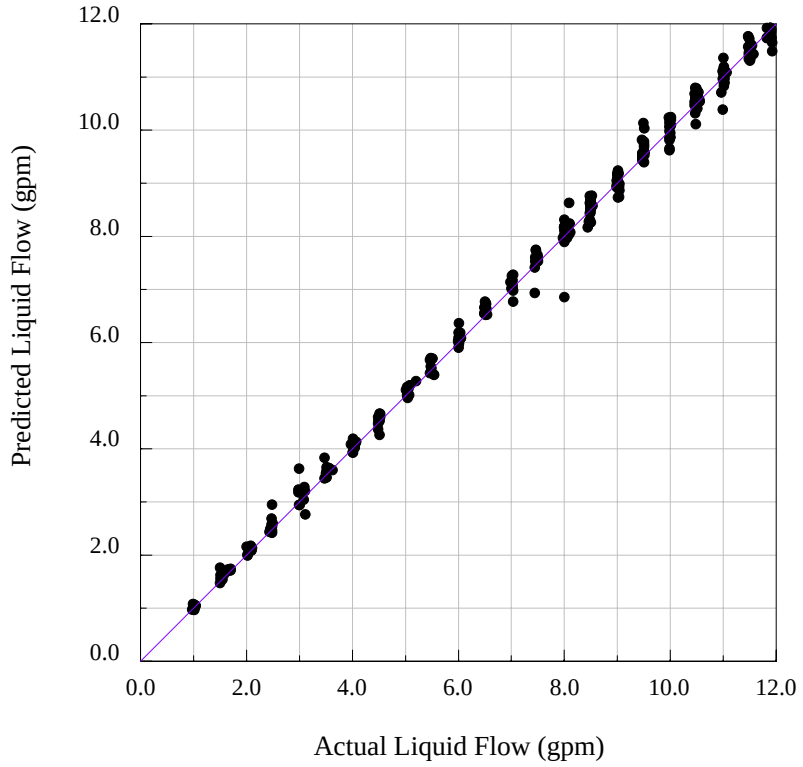

Figure 1. Sample Dynamical Instruments result: actual vs. predicted liquid flow from a wide range of liquid and gas flow conditions

1.3 Phase I Technical Objectives

The project's overall objective was the development of a commercially viable dynamic signature based sensing system that is used to infer the flow rate and fineness of pulverized coal. This effort focused on developments required to transfer the measurement system from the laboratory to a field ready prototype system. In order to achieve this goal, the program pursued the following specific objectives:

- Using an existing design, construct a portable instrument system. The instrument system includes an accelerometer to be attached externally to coal feeder pipes, data acquisition hardware and software, a signal conditioning and data processing module, and a neural network module. Raw data is collected from the impingement of the coal particles as well as the acoustic noise generated from the flow and is transformed into characteristic signatures through proper calibration to operational parameters that are meaningful to the operator.
- Install the system on the test rig located in the Coal Flow Measurement and Control Laboratory operated by Airflow Sciences Corporation in Livonia, MI, under contract to

EPRI. Conduct tests using a 12-inch pipe diameter, varying the following parameters: coal type, coal flow rate, air flow rate, pipe configuration and flow orientation. Coal particle characteristics will include fixed fineness, unless other tests provide data, and fixed, uncontrolled moisture. Perform an iterative process of data collection, data analysis and algorithm development in order to obtain a neural network that is representative of a range of flow conditions and produces a reasonably accurate determination of coal flow.

- Analyze the dynamic data to determine the combination of data acquisition rate and sensor characteristics that provides an optimal combination of instrument performance and final hardware cost.
- Using the instrumentation package developed, perform testing at two full-scale coal fuel power plants in order to evaluate the ability of the algorithm to accurately infer coal flow rates and determine if the measurement system can be used effectively in an active control loop for combustion diagnostics and burner balancing.

At the completion of this project, it was expected that prototype versions of both a portable system and a permanent installation would be available for final packaging and commercialization by one of the team members. Both types of systems would be marketed for conducting combustion diagnostics and balancing of individual flows to pulverized coal burners. The benefits sought through the use of this system include improvements to a plant's combustion feed utilization rate as well as the overall efficiency of a pulverized coal combustion system.

1.4 Program Results

A sensing system was designed as planned and is fully described in Section 3.2. The system is modular in structure, allowing for an easy installation that will not interfere with coal-plant operations. During the project, signal noise was more problematic than had been anticipated. The principal source of the noise ultimately was pinpointed to electrical grounding of the transducer, and was due to the particular design of this laboratory-grade transducer, although the component was originally selected as the top-of-the-line product among available high-frequency transducers. Our analysis showed that in fact we could use a lower-frequency transducer in the final instrument. This conclusion could only be reached by collecting and analyzing data as we have done in this program, and is a key result of the effort. In fact, this outcome is quite fortuitous for the instrument under development, as there are many more options for lower-frequency transducers, which as a rule tend to be both more robust and less expensive. Replacing the transducer with a lower-frequency transducer should eliminate the grounding-related noise issues, reduce the cost of the instrument system and permit the use of a much more rugged transducer.

Due to the grounding issues discussed above, far more laboratory data was required than had been planned. Fortunately, we were able to piggy-back on planned EPRI testing at the Coal Flow Test Facility and obtain a no-cost extension, which allowed us to obtain more laboratory data and to pinpoint and resolve the grounding issues. In the end, the collection of "clean" data was more limited than we had hoped at the start of the program. However, the correlation results were extremely good. For the 51 test cases from the laboratory, a neural network trained to

predict coal flow resulted in a correlation coefficient (r2-value) of 99.3%, providing a measurement accuracy of + 2.9% of full scale. This result is equivalent to a very good in-place instrument calibration based on a level of knowledge of actual flows that would not be possible in the field. These tests were based exclusively on a magnetically mounted sensor, which would be used for a portable instrument. The electrical grounding issues were more easily addressed with the magnetic mount than the stud mount that would be used for a permanent installation. Therefore, the final set of data used for the program effort was collected solely with the magnetic mount.

Field data was collected from two generation units at the St. Claire plant in Detroit, MI. As is detailed in Section 6.3, the flow conditions, pipe diameters and pipe configurations were all different from what had been examined in the laboratory. Further, the "actual" coal flows came from extractive sampling and were therefore not as accurate as the flow values determined in the laboratory, where the coal was weighed. The data was not as clean as the data that was ultimately obtained in the laboratory. After eliminating data files that showed significant evidence of noise, only 28 files remained. As shown in Figure 2, a neural net was trained to predict coal flow for the combined set of laboratory and plant data, obtaining a correlation coefficient of 96.1%, or a level of accuracy of +6.6% of full scale. In light of the fact that the "actual" coal flows from the field testing were obtained using extractive sampling and thus were probably no more accurate than this, the result was as good as one could hope for.

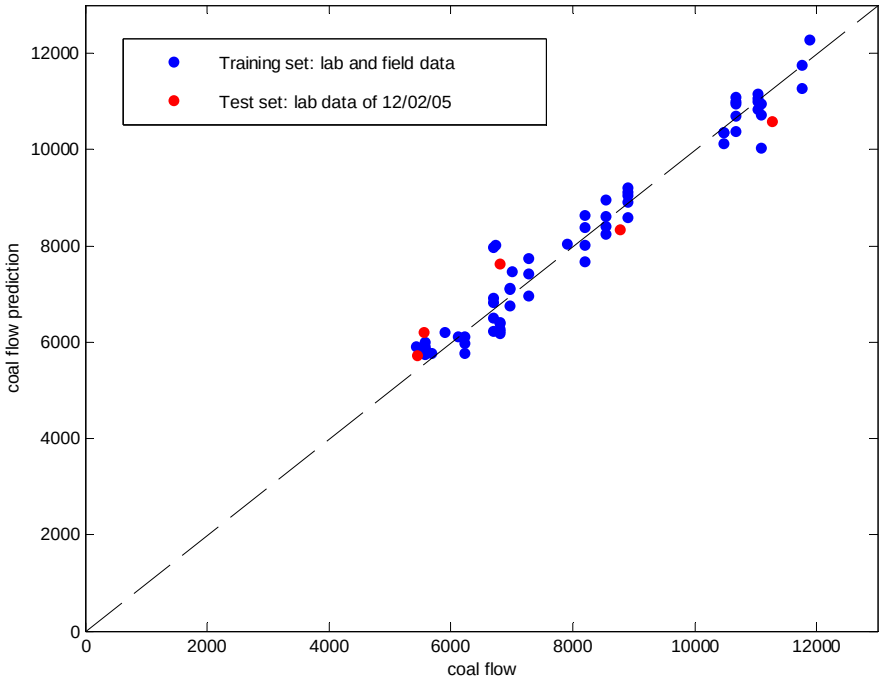

Figure 2. Correlation obtained for laboratory and field coal flow data

Due to the noise problems that had been present through most of the program, we felt it important to demonstrate the feasibility of including a real-time noise-detection algorithm as part of the instrumentation package. This would serve to provide automatic warnings when instrument readings are likely to be unreliable. Section 6.4 discusses this issue. The analysis presented there indicates that such an algorithm should be fairly easy to construct.

Overall, the results of this effort show that the operating envelope of the instrument can be expanded to accommodate quite different conditions. Availability of more data should both broaden the range of applicability and tighten the correlation significantly. Ultimately, with experience, application of an instrument in a new plant should not need to involve in-place calibration.

Information regarding particle fineness, moisture, and other coal characteristics were not available to us, but probably did not vary significantly in the laboratory testing. Therefore, we limited our efforts to the development of an algorithm to predict flow conditions.

The following sections describe the Dynamical Instruments technique, the test facilities and instrumentation developed, the data collected and the ensuing analysis, as well as the results and a recommended approach for future commercialization.

2. DYNAMICAL INSTRUMENTS METHOD

The Dynamical Instruments technique was invented in the course of a program for the U.S. Air Force involving the measurement of liquid-gas two-phase flows, and has since been used in numerous applications for fluid flow measurement, process characterization, and sensitive diagnostics and prognostics. To date, the technique is covered by 5 U.S. patents ([1]-[5]), with numerous complementary foreign patents either in force or pending.

In developing a flow meter, one is immediately faced with the problem that there is no single measurement that is directly related to the flow conditions. There is no simple gauge that reports the air and coal flow rates. Methods such as those relying on a turbine wheel or orifice, which are appropriate for use in a liquid-flow meter, cannot even be considered here as the coal particles are abrasive and the mechanical components would deteriorate rapidly. Further, the coal particles would tend to build up in the crevices of a turbine wheel and in any "dead" spots in the flow, both upstream and downstream of an orifice. This is before even considering the fact that two-phase flows, in this case mixtures of coal and air, are far more complex than single-phase flows, and therefore are much more difficult to monitor. In short, the possibility of a flow meter based on first principles is nonexistent for this application. In the absence of a suitable flow sensor, a different measurement approach is required. In what might be called a "conventional" approach, various features of the dynamics of the system are used, often relying on a characterization of the dynamics of the system in the frequency domain or via gross statistical measures:

- Frequency-domain analysis links the amplitudes of behaviors at particular frequencies, which are believed to reflect behaviors occurring at the fundamental or harmonics of frequencies of the behaviors of interest. One problem with this approach is that the flow dynamics are complex deterministic, but they are not periodic. Performing an analysis in the frequency domain has the effect of averaging out the "random" variations of the dynamics. Unfortunately, these seemingly random variations reflect the smallest scale of deterministic mechanisms that are the most sensitive indication of the system dynamics, so that removing them from the analysis limits the sensitivity of the diagnostic. Even large scale temporal variations of amplitudes, which may be clear indications of the flow conditions, are ignored. In addition, by considering only the amplitudes of frequency components, the phase information is discarded, so that the order of events is not taken into account.

- The alternative conventional approach, using statistics that characterize the time series, usually employs measures that are independent of the order of events. For example, the calculations of skewness, kurtosis, etc. do not depend on the "shape" of the evolution of the time series data, but instead on the distribution of values that are visited. Effectively, all of the time series values are jumbled together out of temporal order in computing these values. Thus, a great many time series can produce identical values of skewness and kurtosis. While conclusions based on the distribution associated with a the time series are valid, there is a loss of information in building a distribution from the time series, This renders statistical measures "blunt" instruments to apply to subtle data.

Overall, correlation of time-series measurements to flow rates should not be approached as a means to extract static measures. In fact, the variation inherent in the flow dynamics is in itself indicative of the flow conditions. The basic approach of the Dynamical Instruments method is illustrated below in Figure 3. It involves the following steps:

- Attach a suitable sensor to the system to obtain a signal that is representative of the system dynamics related to the operating conditions of interest.
- Apply amplification and signal conditioning to the sensor output.
- Digitize and buffer the data.
- Calculate a set of quantities that characterize the dynamic variation of the acquired data, comprising a "signature" of the system dynamics.
- Correlate the dynamic signature to the operating conditions of interest, in this case flow conditions, producing an output that can be used by the operator.

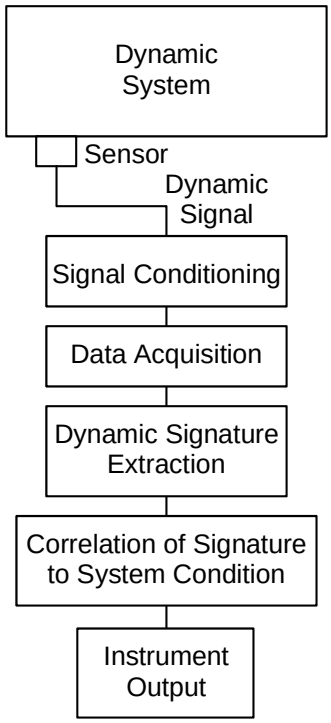

Figure 3. Dynamical Instruments concept

The approach described above is similar in appearance to the technique used in conventional machinery diagnostics, where a vibration or other machine behavior is related to the mechanical condition of the machine. In fact, however, Dynamical Instruments analysis is actually quite

different. The fundamental difference lies in the type of characteristics of the signal that are considered:

- Rather than time-averaging the sensor output, the raw time series data are used, including the seemingly "random noise" present in such signals.
- Rather than performing frequency-domain analysis, which is computationally intensive and commonly throws away half of the information of the raw signal (the phase), the statistics are calculated from the time series data themselves. No recursive computations are used, significantly reducing computational complexity.
- Rather than computing gross statistics of the data, the statistics take into account the detailed shape of the time series. This retains the ability to characterize the order of events, size, and "texture" of the data.

The net effect is a dynamic signature that is extremely sensitive to small changes in operating conditions. The following sections describe the fundamental rationale behind the Dynamical Instruments method and the issues that are involved in implementing it in practical applications.

2.1 Fundamental Reasoning

The only concept that is required to understand the Dynamical Instruments method that is new to most readers is the idea of an attractor. An attractor is a behavior to which a deterministic dynamic system is drawn whenever its global operating conditions are steady. A simple example of this is a damped driven oscillator. Figure 4 shows the attractor for this case, which happens to be a simple circle in velocity/position space. For all initial conditions not on this curve, the driving and damping forces of the system draw the behavior toward this circle. Thus, after a short period of time, the oscillator behavior settles out to this attractor. The radius of the circle increases as the driving force for the oscillator is increased (or the damping decreased), and conversely for decreases or increases in the driving force or damping. Thus, the attractor changes as the global operating condition of the system changes.

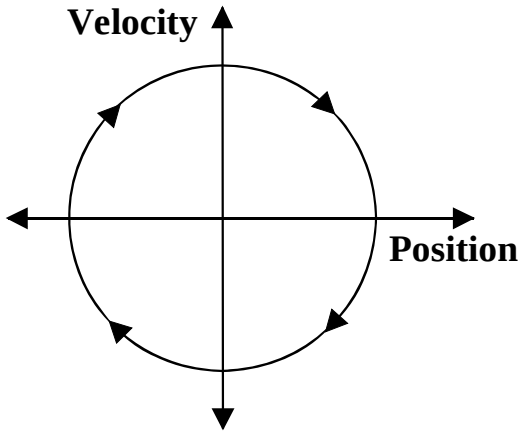

Figure 4. Attractor for a damped, driven oscillator

In a more complicated example, a fluid flow, the dynamics of the flow depend on parameters of the flow (e.g., Reynolds number, in a single-phase flow) and its initial condition. Even though the flow is probably turbulent, two different pipes of the same size, carrying identical flows, must have behaviors that are different instances of the same attractor. This assertion can be a bit hard to understand, since the detailed dynamics of a fluid flow are highly irregular. In fact, most engineers view the normal turbulent fluctuations in fluid flow as random. Although they are complicated (and in some cases extremely so), the dynamics of any macroscopic physical system are not random. They follow deterministic laws of physics that mandate their outcome. Mathematically, any deterministic system that has a steady global operating condition must approach an attractor. For relatively simple, periodic systems such as the damped driven oscillator, the attractor is quite simple. For highly complex systems such as fluid flow, this attractor happens to be highly complex.

In a sense, normal engineering practice already makes use of the pervasive existence of attractors. In turbulent fluid flow, detailed velocity fluctuations are responsible for the phenomena of pressure drop and heat transfer. Since the pressure drops or heat transfer coefficients in identical pipes with identical flow conditions are expected to be the same, there is an implicit assumption that they have the same attractor.

For systems somewhere between the simplicity of a damped driven oscillator and the complexity of a highly turbulent fluid flow, the attractor can look like the one shown below in Figure 5. This figure was constructed using experimental data collected for a nonlinear electronic oscillator. As the figure shows, the behavior of this circuit is very interesting, displaying significant new kinds of behaviors not seen in the attractor for the damped driven oscillator. For example, the attractor is not a simple closed curve, and in fact is not a closed curve at all. In addition, the attractor fills space, and involves a delicate interplay among several types of behavior. By the same token, this attractor has important similarities to the attractor of the damped driven oscillator. For example, both show that the long-term behaviors of both systems are limited to a small range of possible states, and that the time evolution of the system depends on the current state. The most important similarity for current purposes is that both behaviors are attractors, so they capture the dynamics of their respective systems for the given operating conditions: when the operating conditions of either of these systems return to the condition corresponding to the attractor, the dynamics return to the attractor.

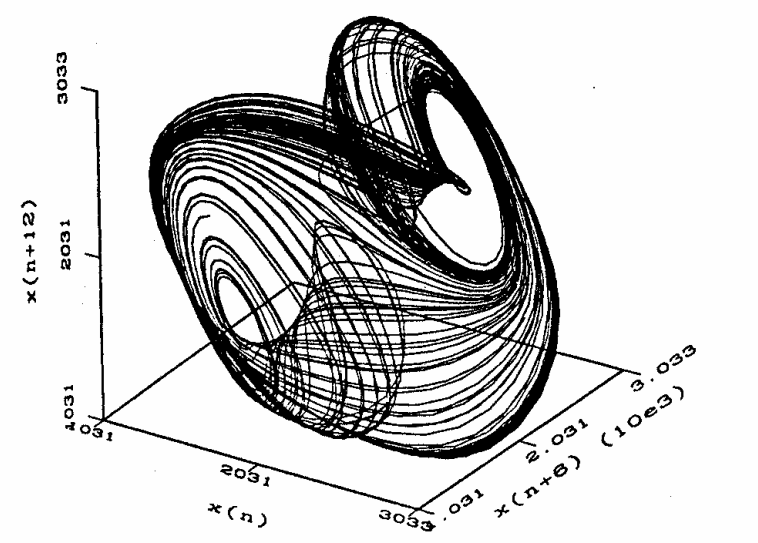

Figure 5. Attractor for a nonlinear oscillator

As the system operating conditions change, the attractor must also change. That is, any change in the system parameters must be reflected in a change in the system dynamics because, otherwise, the dynamics would not be a function of the parameter. Small changes in the operating conditions must cause continuous changes in the dynamics which, in many cases, might be expected to be small. Still, there are cases where a small change in operating conditions can produce a large change in the dynamics, such as the onset of new flow regime.

Given two attractors, it is always possible to differentiate between them if they are different. For example, one could measure the radius of the attractor, or find the average time between visits of given neighborhoods, or characterize the density of trajectories for different neighborhoods, etc. If a given value for a measure can be found each time the measure is applied to a given attractor, the measure is called an invariant of the attractor. By comparing the invariants for two attractors, one can always differentiate between attractors. If one set of invariants does not discriminate between the attractors, it is always possible to identify another invariant that does. In fact, it is always possible to define an infinite number of measures of an attractor that provides a given type of discrimination between attractors.

The observed, temporal behavior of a system is equivalent to the system's attractor. This is a proven mathematical result, but also one that can be understood intuitively. Mathematically, any temporal behavior of the system is a mapping of the attractor onto the time line. One can construct the simple attractor of Figure 4 from a time series history of the damped driven oscillator by estimating the velocity from sequential position measurements (or the position from velocity measurements). In fact, the much more complicated attractor of Figure 5 was reconstructed from experimental observations of the nonlinear oscillator. Thus, in a sense, the observed dynamics of the system is the attractor.

Reconstructing an attractor from time series data requires a sufficient quantity of data that the system would have visited a representative sampling of the neighborhoods of the attractor. For example, the simple attractor of Figure 4 could not be wholly reconstructed from less than a full period of the oscillator's motion. More to the point, the more complicated attractor of Figure 5 could not be reconstructed from a single oscillation, but instead requires many oscillations to obtain a representative sample of its dynamics.

The simplest way to find out how much data is needed to reconstruct an attractor is to calculate measures of the data and see if they vary with sample length. For example, the classic invariant measure of an attractor is its radius, which is usually taken as the standard deviation of the data. One can tell whether enough data is available by calculating standard deviations for subsets of the data sample and seeing whether the result is sensitive to sample length. When the values of a full set of such measures are relatively stable with sample length, the sample length can be declared sufficient. Interestingly, though, a wholly sufficient sample may not be needed to differentiate between two attractors. If a short sample displays a behavior that is peculiar to one attractor and not another, then it can be sufficient to differentiate between the attractors without being sufficient to reconstruct the whole attractor. Thus, it is not obvious what the minimum amount of data is that is required to estimate the system operating conditions to a given level of accuracy.

If two different sensors are used to observe a given dynamic behavior, the attractors associated with these two sensors are equivalent, even if the sensors are sensitive to different phenomena (e.g., acceleration and acoustic pressure). This is because all of the variables of a dynamic system are interrelated by the dynamics. Mathematically, these variables are simply different mappings of the attractor onto the time line. It is certainly possible, even likely, that one sensor's output will provide a better reflection of the dynamics of the system than another, but this is not a particular problem of the Dynamical Instruments method: one should always try to use the best instrument for an application. Still, there is not necessarily only one sensor that can be used as the basis for a Dynamical Instrument in a given application, and in many cases any of a variety of sensors could be used.

Although an attractor can be reconstructed from dynamic data, this is not to suggest that the attractor actually needs to be reconstructed before comparing one operating condition with another. This is simply not necessary, since the data are already equivalent to the attractor. Thus, one can calculate invariants of the data that are equivalent to invariants of the attractor. In addition, it is often difficult to reconstruct an attractor from experimental data, particularly in highly complex systems such as two-phase flows. The reasons for this are somewhat involved and beyond the scope of the current discussion. Still, in order for the Dynamical Instruments method to work, it is not necessary to be able to reconstruct the attractor, but simply that an attractor exists. Thus, the Dynamical Instruments method is implemented by correlating quantities calculated from the time series data to the operating condition of interest.

The following sections describe the detailed application of the Dynamical Instruments method, and the ways that individual technical issues are reflected in the method, its implementation, and the results that are obtained.

2.2 Data Acquisition

There is no longer much mystery to acquiring data digitally. From the standpoint of the Dynamical Instruments method, the key issues are the rate and duration of sampling. These issues have different implications for purposes of instrument development and implementation.

For purposes of both development and implementation, an insufficient sampling rate can miss important system dynamics, limiting the accuracy and even feasibility of a flow instrument. Thus, during development, one should always err on the side of a high sampling rate. The only cost of this approach during development is large data files. For implementation purposes, an excessive sampling rate would require faster data acquisition and analysis hardware, increasing the overall cost of the instrument, and may actually decrease the accuracy of the resulting instrument. Thus, one should optimize the sampling rate during development so that the resulting instrument achieves the best accuracy with the least expensive hardware. Fortunately, this optimization can be done in a relatively straightforward fashion. If the data are collected at an ample sampling rate, the effect of using lower sampling rates can be simulated accurately by skipping readings in the data files in calculating the candidate signature quantities. One can then determine the effect of sampling rate on instrument accuracy, providing a solid basis on which to make a design decision.

Just as in the case of sampling rate, during development the only cost of increasing the duration of sampling is large data files, while a short sampling duration potentially could cause important dynamics to be missed. For a diagnostic instrument implemented in the field, the ability to achieve a reliable flow estimate from a relatively short data snippet is desirable to provide a responsive instrument system. Thus, the size of the data sample used to update the instrument output should be selected to provide maximum accuracy subject to minimizing data requirements. As in the case of sampling rate, sampling duration can be optimized during development through analysis. During preliminary development, the system dynamics can be characterized using long data records to be certain of an accurate characterization. As the development process advances, the issue of sample size can be examined by analyzing the data in snippets of various sizes to determine how sample size affects update accuracy and responsiveness. The resulting strategy can then be implemented in prototype instruments with some confidence.

2.3 Dynamic Signature Extraction

This is the part of the method that most departs from conventional data analysis methods. It begins by calculating a broad array of quantities that characterize the signal dynamics, in hopes of finding a group of such measures that reliably discriminate one operating condition from another. There is no *a priori* method of defining which measures of a signal's dynamics will be related to system operating conditions, and no way of delineating which of any group of such measures will provide the best results. Thus, it is necessary to define a candidate set of measures

to apply to the signal, and then use the process of correlating these measures to the desired flow rates to help optimize the selection.

While no fundamental guidance is available to define a set of measures to apply to a dataset, this process can be pursued with confidence, because the previously described mathematical argument indicates that the method should work. If flow conditions are different, then the corresponding system dynamics are different, and an infinite number of measures of the dynamics will disclose this difference. Thus, if a given set of measures does not satisfactorily determine the system operating conditions from the observed data, there are other measures that will improve the instrument's performance. In addition, there is no shortage of possible dynamic measures. One could define many possible measures for any given phenomenon observable in the data. In some cases, these different measures may be effectively equivalent, while in other cases the information contained in one measure might be complementary to information contained in another, similar measure.

Thus, defining a set of measures to serve as a signature of a system's dynamics is essentially a black art. With that said, a set of guidelines has been developed over a series of past Dynamical Instruments programs to help provide a firm footing for a dynamic signature analysis:

- Use measures that characterize both the temporal and amplitude behaviors of the signal. Thus, in addition to measuring the "size" of events in the data via amplitude, one should measure the frequency or period of these events.
- Use measures that characterize a given phenomenon in different ways. For example, one measure might be the average of some phenomenon, and another might be the RMS (root-mean-squared) of the same phenomenon. The RMS value weighs larger events more heavily than the average does, so the two measures together provide information about the distribution of the sizes of the events.
- Use multiple similar measures. Some measures may prove to be surprisingly insensitive to the system operating conditions of interest while other, similar measures might incisively discriminate between operating conditions.
- Use measures that characterize phenomena at different scales within the data. For example, one set of measures might characterize the largest events occurring in the data, while another set might characterize the smallest, "texture" scale of events.
- Use measures that characterize both the actual time evolution of the signal and the passage of harmonic power in the signal. That is, since the signals occurring in flow dynamics reflect the passage of events associated to the flow, the resulting signal has aspects of both low-frequency variation (e.g., with the passage of each large event) and high frequency (with the detailed response to each event). By separately considering these aspects, one can gain further insight into the nature of the signal dynamics. This has led to the use of two separate groups of signature quantities, one calculated from the raw dataset and the other from a new dataset that characterizes the passage of harmonic power in the raw dataset.
- Use at least some measures that are independent of the signal amplitude *per se*. If such measures are successful, then two sensors with different sensitivity, or two installations of the same sensor with different coupling characteristics, have the potential of providing the same instrument output without special calibration.

Following this simple set of rules, a "standard" set of statistics has been developed, comprising 57 measures, which form the baseline when examining a new application of the Dynamical Instruments method. The details of these signature quantities are described in the Appendix.

In a previous effort, an ActiveX dynamic link library (DLL) was developed that can be called from Matlab to calculate each of these quantities. The raw data are imported into Matlab, preprocessed as necessary, and then sent to the DLL to calculate the 57-quantity signature. The resulting signature is then returned to Matlab, where it is used in further analysis.

2.4 Correlation of Signature to System Conditions or Parameters

The 57 quantities described in the previous section are intended to serve as candidates for inclusion in signatures of the flow dynamics. Not all of these measures would be used in the ultimate embodiment of an instrument, so each should be considered a candidate for inclusion in a dynamic signature. Having identified such a set of properties of the attractor, the job now comes to correlate these quantities to the flow condition.

2.4.1 Correlation Algorithms

There are quite a few different approaches that have been used with success in past Dynamical Instruments development programs, and different situations call for different methods. Each of these techniques is a form of state space analysis, in which the several statistics of the dynamic signature can be viewed as different dimensions in a state space, with the location of a given point being used to determine the flows conditions. Exactly how this analysis is performed can have a very strong influence on the results that are obtained. Some of the methods that might be considered include the following:

- Simple clustering methods – this approach involves geometrically separating different flow conditions according to their locations in the state space. Common techniques include principal components analysis and singular value decomposition, which linearly map the points for similar conditions into the best available estimate of a sphere. Whether or not a given point is similar to other points is determined by whether it lies within or near one sphere or another. This analysis can be performed readily in the Matlab environment using built-in function calls.
- Convex hull analysis – this technique is an extension of the simpler techniques. Rather than assuming that similar points comprise a sphere, they are assumed to form a simple convex shape, with points within the cluster being separated from other points by a geometrical boundary, called a hull. The location of a point within the hull, or distance of the point from the hull surface, determines the flow conditions. One major advantage of this approach over the simpler geometrical methods is that adding new points to the analysis simply shifts the hull, rather than rescaling the state space. Matlab routines were developed in a previous program to identify the convex hull and find a point's location relative to it.

- Local state space analysis – this is a technique in which several nearest neighbors from the experimental data determine the flow condition for a new data point. By considering neighborhoods separately, this approach accommodates fairly strong variation of the system behavior with location in state space. Among the highly capable nonlinear techniques, this approach has the advantage that it can be validated: it is easy to determine whether the experimental data are consistent, and whether a new data point falls within the range of the experimental data. In addition, this analysis can be updated easily when new experimental data become available. The algorithms required to perform this analysis are quite involved, and cannot be performed efficiently in the Matlab environment (because it does not support pointers). Consequently, in a previous program a special-purpose dynamic link library was developed to perform this analysis using calls from the Matlab environment.
- Neural networks – in this case, a randomized training technique is used to develop the algorithm that relates the location of a point in the state space to the flow condition. This renders the process difficult to verify, because the location of a given point relative to the experimental database is difficult to determine. Despite this disadvantage, neural networks are extremely adept at scaling the inputs differently in different neighborhoods. This capability often proves to be very useful, because the experimental points are usually distributed heterogeneously. Recently, Foster-Miller has funded the development of a proprietary analysis tool that allows very rapid training of neural networks. This tool, which uses Particle Swarm Optimization (PSO), has been implemented as an ActiveX DLL that can be called from the Matlab environment.

2.4.2 Selection of a Signature

Given a suitable selection of dynamic characteristics, one or more of the techniques described above can be used to correlate signature to the flow conditions. In many applications, it is the selection of the signature from among the many available statistics that is the principal barrier to the successful development of an analysis. A number of different techniques have been used to overcome this issue in past Dynamical Instruments programs:

- Graphing – viewing a scatter plot of the desired output (e.g., coal flow rate) as a function of a signature quantity can rapidly disclose whether a given statistic is a good candidate for inclusion in the signature. This process can be performed very rapidly in Matlab, with each promising signature quantity noted along the way. After this first round of selection, each of the promising quantities can then be plotted as a function of each of the other quantities, with the desired output represented by the symbol color. This rapidly indicates whether an added statistic improves the appearance of the correlation. The result is a relatively modest number of statistics that can be tested in exploratory algorithm development.
- Mutual Information – this approach uses a statistical calculation that relates the information that is known about a quantity Y if one knows the corresponding value of X. If Y is related to X, graphing Y as a function of X produces a curve, which could take on a variety of shapes depending on the functional relationship. In this case, X contains a lot

of information about Y. On the other hand, if X and Y are unrelated, the graph will produce a random distribution of points, indicating that X contains little or no information about Y. The mutual information calculation essentially characterizes the tendency for the X and Y values to correspond to one another in a pattern, by sorting the X and Y values into an NxN array of bins. This calculation can be performed very rapidly to find the mutual information between each signature quantity and the desired output, producing a short list of quantities that are good candidates.

- Numerical experimentation – if a simple signature is sufficient, then trying all possible combinations may be worthwhile. This approach works well when the algorithm is simple to automate, and when only a modest number of signature quantities are required (say, 4 or less). When larger numbers of quantities are used, the number of combinations becomes so large that other approaches tend to prove to be more efficient.
- Particle Swarm Optimization – this is an extension of the neural network training algorithm. Since PSO trains neural networks very rapidly, numerical experimentation can be used to try a large number of combinations of signature quantities as inputs to networks. It turns out that PSO itself can be used to rapidly identify a near-optimal signature of the dynamics, by optimizing the output of a neural network with any N inputs chosen from a list of, say 57. This capability that been implemented in the particle swarm ActiveX DLL described above, with great success. It has actually proven to be efficient to use PSO to identify a suitable signature, and then to use this signature as inputs to other analysis algorithms described above.

This section has described why and how the Dynamical Instruments technique works. Later sections will describe how this approach was used in the program at hand to infer flow conditions from accelerometer time series data.

3. TEST FACILITIES AND INSTRUMENTATION

Data was collected for this project in two stages. The first stage consisted of laboratory testing at the new Coal Flow Measurement and Control Laboratory, which is sponsored by the Electric Power Research Institute (EPRI), the cost share partner on this project, and has been built and is being operated by Airflow Sciences in Livonia, Michigan. The second stage consisted of field testing performed at Detroit Edison's St. Claire Power Plant.

3.1 Laboratory Facilities

The coal flow test facility is shown in the next four pictures taken at the project kickoff meeting in November 2003 at the facility in Livonia, Michigan. Figure 6 is the coal and air feed, metering, and measurement system, which produces known mass flows of air and pulverized coal to test sections. Figure 7 shows a transparent section of pipe with rotoprobe sampling port in an upflow orientation following two elbows. Figure 8 is a picture of the control loop display, showing the system schematic. Figure 9 is a closer view of this transparent test section showing a swirling flow of silicon dust used in shakedown tests of the facility.

Figure 6. EPRI & Airflow Sciences pulverized coal and air supply system

Figure 7. Transparent coal and air test section

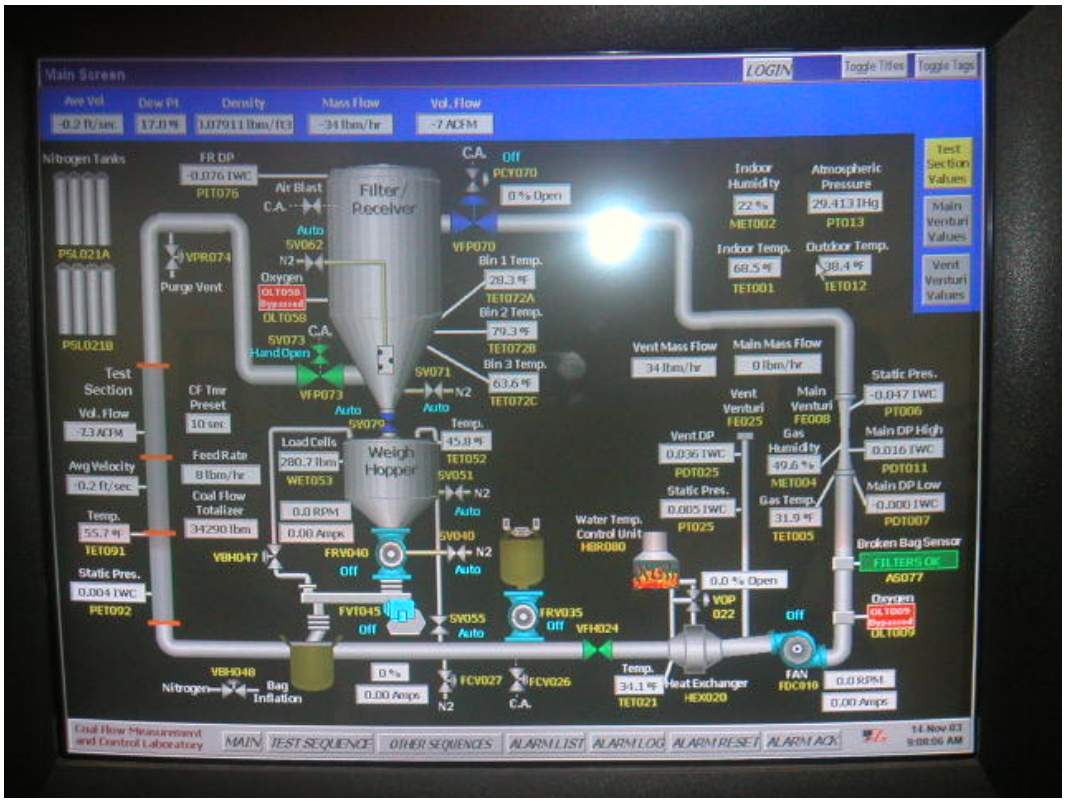

Figure 8. CRT display of Coal Flow Loop schematic, instrumentation and controls

Figure 9. Test section with silicon particles and air, swirl induced by offset elbows upstream

3.2 Instrument Package by Foster-Miller

The original coal flow measurement test package that was installed for use during the preliminary testing at the Coal Flow Test Loop consisted of four main components:

- An Endevco Model 7259 uniaxial accelerometer with high frequency response (its bandwidth very conservatively quoted as 30 kHz) and 10 mV/g sensitivity.
- A Kistler Model 504E isotronic signal conditioning amplifier.
- A compact filter/amplifier module featuring a Krohn-Hite 100 kHz low-pass filter and amplifier.
- A transportable computer with a Microstar Data Acquisition Processor board.

Initial shakedown testing revealed an unacceptable level of noise. In remedying this situation, the Kistler amplifier was replaced an Endevco 133 amplifier, and the Krohn-Hite filter/amplifier module was replaced with a Krohn-Hite Model 2284 filter. For a further discuss of these modifications, see Section 4. The revised package consisted of the following components:

- An Endevco Model 7259 uniaxial accelerometer with high frequency response (its bandwidth very conservatively quoted as 30 kHz) and 10 mV/g sensitivity.
- An Endevco Model 133 isotronic amplifier.
- A Krohn-Hite Model 2284 filter, a laboratory bench instrument.
- A transportable computer with a Microstar Data Acquisition Processor board

The components of the original instrumentation package and the revised package are illustrated schematically in Figure 10. An accelerometer mounted on the outside of the pipe senses the vibrations produced by the coal particles striking the pipe wall. The resulting raw signal is amplified using the accelerometer amplifier, producing a higher-level signal. This signal is then filtered to eliminate behaviors above 100 kHz to eliminate signal aliasing, and below 1 Hz to eliminate any DC offset. Originally, the signal amplitude was also increased in the same component, the Krohn-Hite filter/amplifier module, but the replacement module had gains set to unity. The resulting signal is then digitized and stored at 300 kHz using the data acquisition computer. Each data file is nominally 30 seconds in duration (with some variation due to details of data buffering in the data acquisition process).

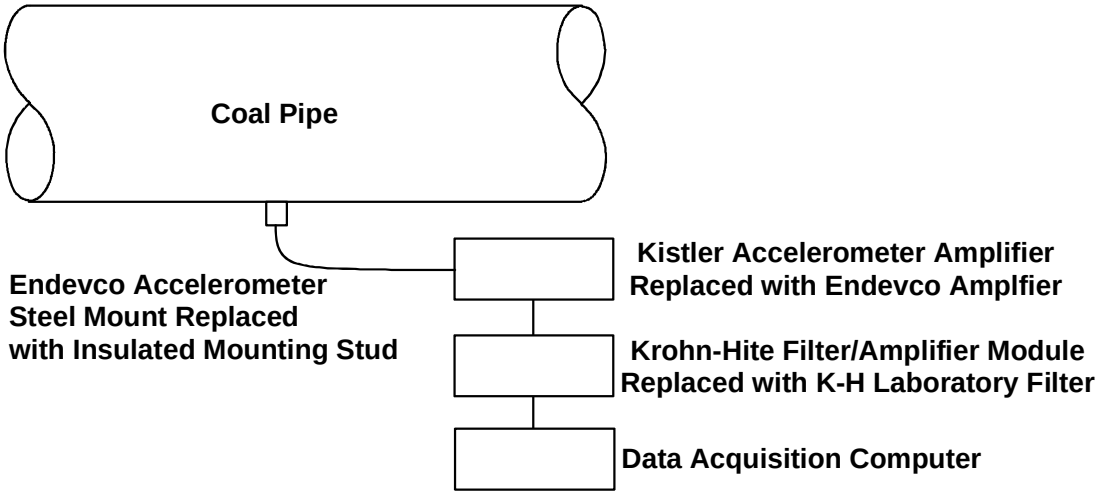

Figure 10. Instrumentation schematic

The Endevco accelerometer is a particularly fine instrument, with a frequency response that is as high as or higher than almost any transducer one might envision employing in the coal flow measurement application. By low-pass filtering the signal to eliminate behaviors above 100 kHz, and sampling the result at 300 kHz, the entire dynamics of the accelerometer response

can be captured. This allowed for extensive post-test analysis of the data to examine the influences of sensor response, filtering, and data sampling rate on the information content of the data. The intent was to select instrument configuration variables to be carried forward in further laboratory and plant testing.

Also provided with the test package was a choice of mounts for the accelerometer:

- An insulated mounting stud, to be threaded into a hole in the pipe wall, and
- A magnetic mount to be applied manually.

Our analysis indicated that the magnetic mount should provide very similar dynamic response to the direct threaded mount, because its size is small compared with the physical wavelength of a pipe wall vibration at the transducer resonant frequency of 90 kHz. The advantage of using a magnetic mount is that this provides a means for streamlining later testing both in the lab and in the field: transducers can be applied to the piping at will, without drilling holes (to which the utilities are understandably resistant, particularly for casual R&D purposes).

After much testing, we decided to focus exclusively on the magnetic mount. The stud mount resulted in ongoing grounding issues that were likely to persist in the field. The data collected using the magnetic mount was cleaner and resulted in a better correlation of the dynamic signature with coal flow. One additional modification to the instrumentation was made along the way. The Krohn-Hite laboratory filter was replaced by a custom built Krohn-Hite filter module, which functioned identically to the laboratory filter but could be plugged into the instrumentation without requiring operational settings, thereby facilitating field testing.

4. SUMMARY OF LABORATORY TESTS

The following summary details the various stages of testing performed by Airflow Sciences at the coal flow test facility. During the testing, there were a number of problems related to noise within the instrument system. As a result of these ongoing problems, more laboratory tests were needed than had been anticipated. Fortunately, many of these tests were performed in concert with EPRI-funded testing. Section 5.2 will contain a comprehensive discussion regarding the specific noise issues involved, as well as the implications for the instrumentation developed during this program.

4.1 Shakedown Testing of Instrument Package

The instrumentation package was installed at the Coal Flow Test Facility, located in an Airflow Sciences facility in Livonia, MI. Initially, several sample data files were collected to examine the behavior of the instrumentation:

- A no-flow test, with power applied to the inverter-based variable speed power supply for the blower.
- An air-only test with an air velocity of 80 fps (24.4 m/s).
- Two coal-air tests with an air velocity of 80 fps (24.4 m/s), a coal flow of 7000 lbs/hr (0.882 kg/s), and two different accelerometer amplifier gains.

The resulting data files were sent to Foster-Miller, where a preliminary analysis indicated the presence of strong frequency "lines" in the power spectrum, indicating that strong periodic oscillations were present in the data. As these lines were present even when no flow passed through the test section, the instrumentation must have been recording behaviors unrelated to flow. This issue needed to be resolved in order to proceed with the collection of calibration data.

The source of the noise could not be determined following long-distance consultation of Foster-Miller and Airflow Sciences personnel, so a Foster-Miller engineer, Bruce Barck, traveled to the test facility to perform hands-on debugging.

The most crucial issue proved to be the transducer mount: the accelerometer must be electrically isolated from the pipe using an insulated mounting stud, yet a plain steel stud had been used. Stray electrical currents are a common problem in coal piping, because of the tendency of the coal particles to become electrically charged. Consequently, coal piping systems are electrically bonded to reduce the likelihood of a dangerous buildup of electrical charge. With a plain steel mounting stud, any residual electrical currents would affect the output of the transducer directly. By replacing the steel stud with an insulated stud, the main source of noise was eliminated.

With the opportunity to test the instrumentation in the operating environment, Mr. Barck examined the performance of each of the other instrumentation subsystems. He found that the noise level of the Kistler Model 504E amplifier was higher than that of an Endevco 133 amplifier he had brought along, so he replaced the Kistler unit with the Endevco one. He also found that

the input gain of the compact Krohn-Hite filter/amplifier module was effectively amplifying the noise floor of the accelerometer signal. He remedied this by replacing the filter/amplifier module with a Krohn-Hite Model 2284 filter, a laboratory bench instrument. This instrument was set to have filtering characteristics identical to those of the compact module it replaced, but with input and output gains set to unity. This did not produce any problems, because the noise output of the Endevco amplifier appears to be essentially independent of its own gain. Thus, raising the gain of the Endevco amplifier compensates for the reduced gain of the filter box without raising the noise level. No changes to the data acquisition computer were needed.

The effect of these modifications is shown in the power spectrum graph of Figure 11. The red trace in this figure is the power spectrum of no-flow data from the initial shakedown testing. The blue trace is the power spectrum of no-flow data after the instrumentation was modified. Both traces reflect data collected with the blower inverter power supply turned on, so the principle likely source of electrical noise was present. As the two traces show, the signal was now quite a bit cleaner than previously. The strong variation in the red trace at low frequencies, which probably reflected amplifier noise, was now greatly diminished. In particular, there is essentially no component of 60 Hz and its harmonics. In addition, the strong lines at higher frequencies are nearly gone.

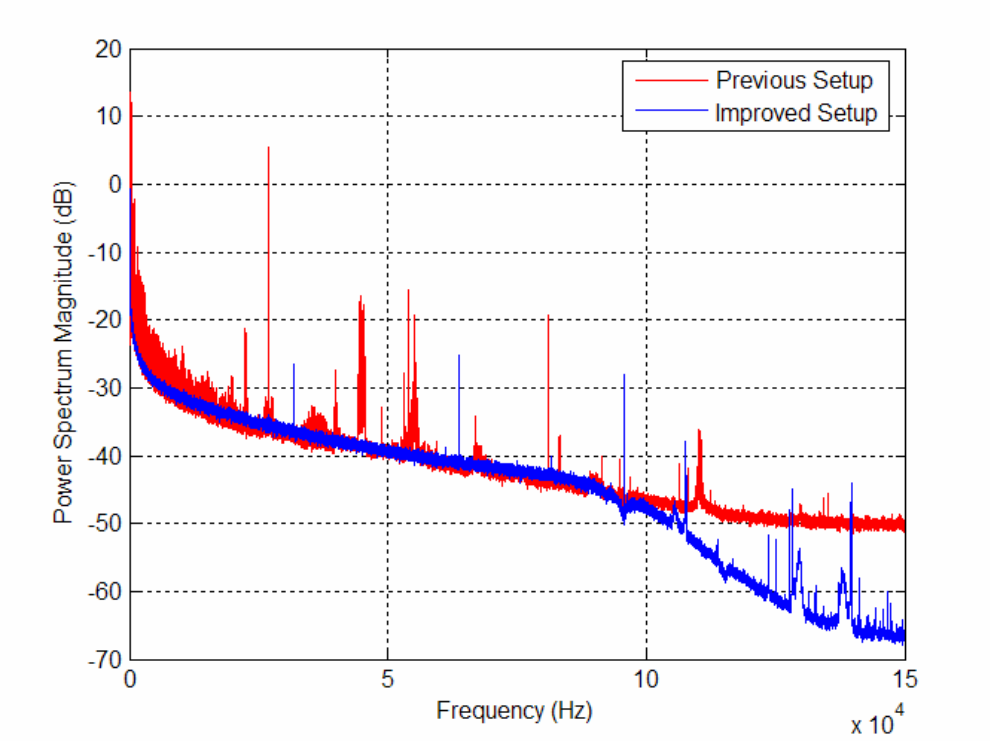

Figure 11. Power spectra of no-flow data before and after improvements

The improvement is just as clear when flow is present in the pipe, as is shown in Figure 12. In this case, the flow conditions for the two data files are similar, yet the characters of the two signals are quite different. With the improved instrument setup, the low-frequency noise seen with the previous setup is strongly suppressed, and the numerous strong frequency lines are essentially absent. With this major improvement to the instrumentation setup, collection of calibration data could begin in earnest.

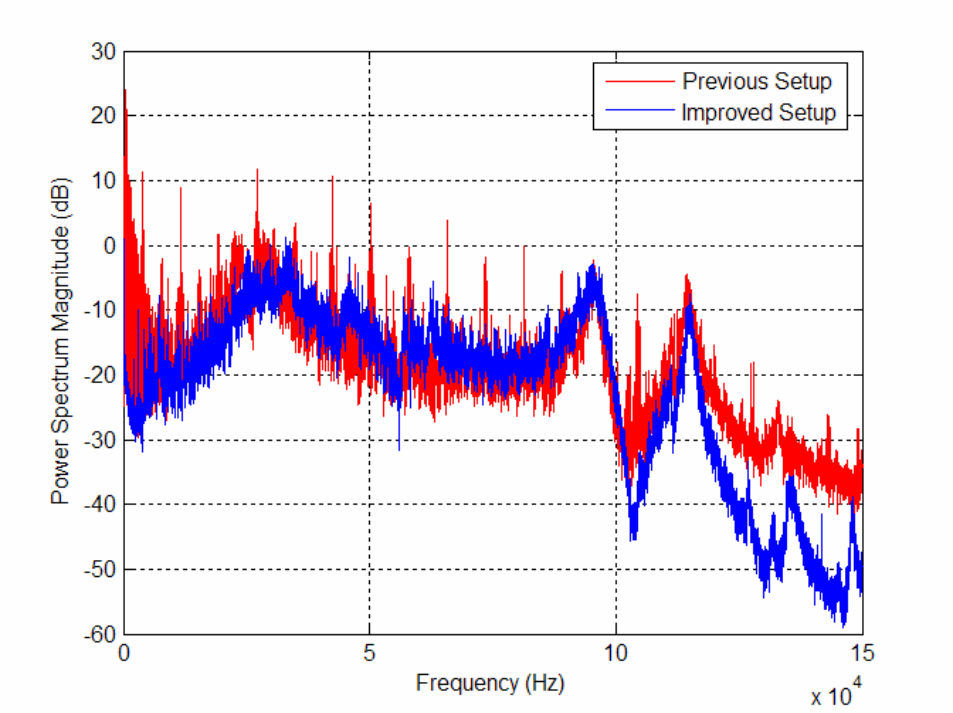

Figure 12. Power spectra of data with flow before and after improvements

4.2 Preliminary Data Collection

The first round of data collection was performed in August, 2004. Our focus for this test round was to gain insight into an issue that has remained unsettled for the ten year span of the development effort, namely how the type and location of the transducer mount affects the dynamics of the data and the extent to which the instrument calibration can accommodate this. The transducer has historically been mounted just downstream of an elbow, on the outside of the bend, using a stud threaded into the pipe wall. This has always worked well, but there are two reasons to consider alternative mounts:

- Although elbows are commonplace and generally accessible in coal piping systems, there are instances when an elbow is not situated in a readily accessible location. Thus, the

versatility of the instrument would be increased if the transducer could be mounted in other locations without loss of accuracy.

- Although a stud mount is suitable for a permanent instrument installation, one market of great interest is for portable instruments, either to be used by plant personnel or by contractors who perform plant balancing efforts. For a portable instrument, a magnetic transducer mount would be highly advantageous.

The availability of the Coal Flow Test Facility provided an excellent opportunity to address these issues. As a first step in developing a coal flow instrument calibration, data were collected with many mounting locations and two mount types:

- Fifteen mounting locations (illustrated schematically in Figure 13), beginning 5 pipe diameters upstream of the middle of an elbow, and moving one diameter at a time downstream, ultimately to a point nine diameters downstream of the elbow.
- Two mount types, including a stud threaded into the pipe wall and a magnetic mount.

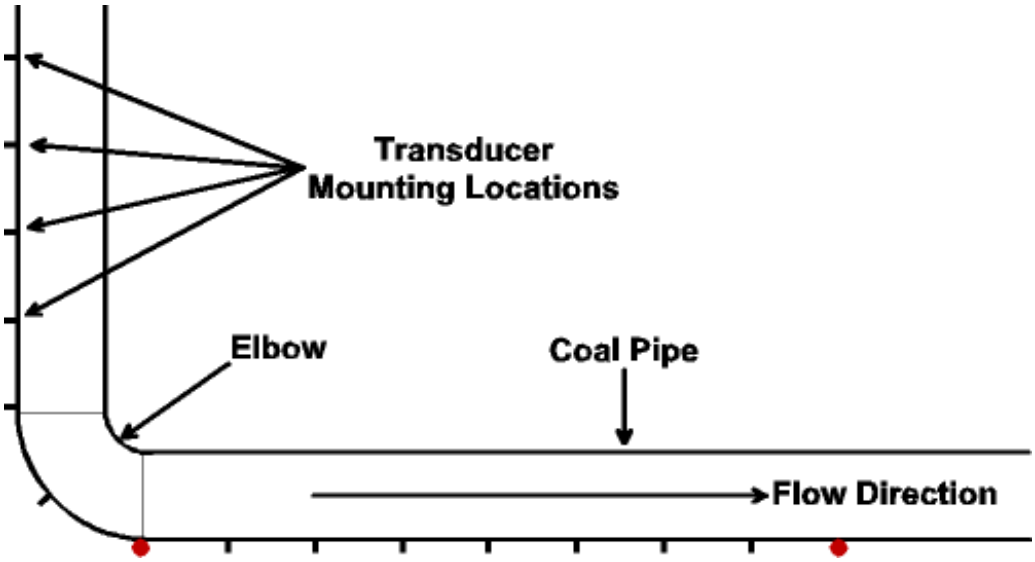

Figure 13. Schematic illustration of transducer mounting locations

Although only 7 flow conditions were visited in this testing, as outlined in Table 1, these conditions covered a broad range of air flow velocities and air/fuel ratios. Thus, although these tests do not fill the operating space, they visit conditions that should include the range of dynamics encountered in practice.

Table 1. Flow Conditions Visited in Scoping Testing

Flow Condition	Air Flow (lb/hr, (kg/s))	Air Velocity (ft/s, (m/s))	Coal Flow (lb/hr, (m/s))	Air/Fuel Ratio
1	11060 (1.394)	52.2 (15.9)	3690 (0.4649)	3.00
2	11060 (1.394)	52.2 (15.9)	11060 (1.394)	1.00
3	15670 (1.974)	73.9 (22.5)	0	inf
4	15670 (1.974)	73.9 (22.5)	5220 (0.6577)	3.00
5	15670 (1.974)	73.9 (22.5)	15670 (1.974)	1.00
6	20280 (2.555)	95.6 (29.2)	5790 (0.7295)	3.50
7	20280 (2.555)	95.6 (29.2)	13520 (1.704)	1.50

The resulting data files, filling 7 compact disks, were sent to Foster-Miller for analysis. Analysis of the first data set disclosed a very complicated variation of the signal dynamics with different types and locations of transducer mounts, as will be further discussed in Section 5.1. This does not imply that a correlation between signatures derived from data at different location mounts cannot be correlated. However, it does mean that in order to fully comprehend the impact of the mount location on the dynamic signature, far more data would be required than could be obtained within the constraints of this program. Therefore, it was decided to limit the transducer locations during further data collections for the duration of the program and to focus on an instrument calibration that works for these locations, with the expectation that experience gained in the field would eventually provide sufficient data to expand the calibration to other mount locations.

4.3 Data Collection of October 2004

The second round of testing was performed by our subcontractor, Airflow Sciences Corp., during October, 2004. At this time, it was decided that:

- It was still of great interest to be able to use either a stud mount, for permanent instrument installations, or a magnetic mount, for temporary installations.
- The outlet of an elbow offers a reliable means of obtaining a strong signal, but is not always readily available in a power plant. Thus, both the outlet of an elbow and a point well downstream of the elbow were selected as suitable measurement points (illustrated with red dots in Figure 13).

The calibration testing was performed using 2 elbows, one turning a vertical downward flow to horizontal and the other turning a horizontal flow upward. For the horizontal-to-vertical elbow, some tests were performed with an additional horizontal elbow installed immediately upstream. This induces swirl in the flow, which tends to cause the coal particles to “rope” together into a coherent structure in the middle of the pipe, an issue of considerable concern to plant operators.

Tables 2-4 below summarize the flow conditions that were visited during this testing. For the vertical-to-horizontal elbow, each flow condition typically represents four data files, including two mounting types and the two locations relative to the elbow. For the other two elbows, the mount type used is listed in the last column, so that each row in the table typically corresponds to two data files (for the two mounting locations). Overall, this testing produced 181 data files, including:

- Stud-mounted transducers in 95 tests.
- Magnetically-mounted transducers in 86 tests.
- Transducers mounted at the elbow outlet in 91 tests.
- Transducers mounted 8 diameters downstream of the elbow in 90 tests.

Table 2. Conditions Visited in Downflow-to-Horizontal Configuration

Air Flow (lb/hr, kg/sec)		Coal Flow (lb/hr, kg/sec)		Air/Fuel Ratio
16284	2.052	0	0	NA
12933	1.63	4263	0.537	3
16141	2.034	4736	0.597	3.4
19298	2.432	4802	0.605	4
16075	2.025	6314	0.796	2.5
19209	2.42	6339	0.799	3
12876	1.622	6370	0.803	2
19123	2.409	7481	0.943	2.6
16135	2.033	7722	0.973	2.1
12814	1.615	8428	1.062	1.5
18315	2.308	9351	1.178	2
15896	2.003	10223	1.288	1.6
12694	1.599	12167	1.533	1
15865	1.999	10335	1.302	1.5
16322	2.057	0	0	NA
19112	2.408	7566	0.953	2.5
16146	2.034	7747	0.976	2.1
16166	2.037	4675	0.589	3.5
16065	2.024	6391	0.805	2.5

Table 3. Conditions Visited for Horizontal-to-Upflow Configuration

Air Flow (lb/hr, kg/sec)		Coal Flow (lb/hr, kg/sec)		Air/Fuel Ratio	Transducer Mount
12767	1.609	12168	1.533	1	Magnetic
12769	1.609	21604	2.722	0.6	Stud
12844	1.618	8656	1.091	1.5	Stud
12858	1.62	8468	1.067	1.5	Magnetic

Air Flow (lb/hr, kg/sec)		Coal Flow (lb/hr, kg/sec)		Air/Fuel Ratio	Transducer Mount
12895	1.625	6352	0.8	2	Magnetic
12938	1.63	6595	0.831	2	Stud
12953	1.632	4353	0.548	3	Magnetic
13002	1.638	4434	0.559	2.9	Stud
15970	2.012	10069	1.269	1.6	Magnetic
16035	2.02	7799	0.983	2.1	Magnetic
16077	2.026	10421	1.313	1.5	Stud
16122	2.031	6326	0.797	2.5	Magnetic
16168	2.037	8230	1.037	2	Stud
16182	2.039	7827	0.986	2.1	Magnetic
16228	2.045	4618	0.582	3.5	Magnetic
16257	2.048	6136	0.773	2.6	Stud
16266	2.05	7409	0.934	2.2	Stud
16318	2.056	4453	0.561	3.7	Stud
16349	2.06	0	0	NA	Magnetic
16448	2.072	0	0	NA	Stud
18500	2.331	9282	1.17	2	Magnetic
18675	2.353	6306	0.795	3	Magnetic
18749	2.362	4869	0.613	3.9	Magnetic
18786	2.367	7513	0.947	2.5	Magnetic
18930	2.385	9505	1.198	2	Stud
19396	2.444	7752	0.977	2.5	Stud
19493	2.456	6747	0.85	2.9	Stud
19580	2.467	4833	0.609	4.1	Stud

Table 4. Conditions Visited for Horizontal-to-Upflow Elbow, Roping Configuration

Air Flow (lb/hr, kg/sec)		Coal Flow (lb/hr, kg/sec)		Air/Fuel Ratio	Transducer Mount
12788	1.611	11790	1.486	1.1	Stud
12851	1.619	7713	0.972	1.7	Stud
12928	1.629	6016	0.758	2.1	Stud
12990	1.637	4327	0.545	3	Stud
15965	2.012	6374	0.803	2.5	Magnetic
16064	2.024	10247	1.291	1.6	Magnetic
16072	2.025	10537	1.328	1.5	Stud
16149	2.035	7696	0.97	2.1	Stud
16216	2.043	8118	1.023	2	Stud
16238	2.046	6715	0.846	2.4	Stud
16294	2.053	7480	0.942	2.2	Magnetic
16298	2.054	4976	0.627	3.3	Stud
16399	2.066	0	0	NA	Stud
16400	2.066	0	0	NA	Magnetic
18622	2.346	9332	1.176	2	Magnetic
18781	2.366	7574	0.954	2.5	Magnetic
18817	2.371	6329	0.797	3	Magnetic
18913	2.383	4844	0.61	3.9	Magnetic

Air Flow (lb/hr, kg/sec)		Coal Flow (lb/hr, kg/sec)		Air/Fuel Ratio	Transducer Mount
18962	2.389	9128	1.15	2.1	Stud
19305	2.432	4690	0.591	4.1	Magnetic
19356	2.439	7900	0.995	2.5	Stud
19446	2.45	6753	0.851	2.9	Stud
19569	2.466	5058	0.637	3.9	Stud

4.4 Additional Laboratory Testing of July 2005

Analysis performed on the laboratory data collected prior to this point had revealed ongoing problems with noise, but had also produced some encouraging results, considering the limitations of the available data. Careful examination of each data file disclosed that various types of noise were present in many data files. In many cases, the noise was of a type that would allow for an identification routine to be built into the instrument system so that a warning could be issued for similar cases. However, through extensive analysis of the data sets that appeared to be noise-free, it was discovered that some data sets still suffer from some form of noise that could not be unidentified. The cases not suffering this noise produced extremely good flow correlations, while the noisy cases were essentially uncorrelated to coal flow. As a resulted, we concluded that additional laboratory testing was needed to identify potential sources of noise and to obtain additional data to serve as the basis of the instrument calibration. In order to reduce the cost to this program, we arranged for the additional testing to be done in concert with testing for an EPRI-funded program. We requested and received a no-cost time extension to accommodate this additional testing.

The testing was performed in the Coal Flow Test Facility during the week of July 25, 2005. Foster-Miller again sent Bruce Barck to the Coal Flow Test Facility to observe the testing and assist with the data collection. The key point of focus for Mr. Barck was to be to identify possible sources for the noise that had been evident in the laboratory data collected previously and, if possible, to find ways to reduce or eliminate the noise. Prior to this round of testing, we had Krohn-Hite build a filter module that would function identically to the general purpose Krohn-Hite laboratory filter box. The benefit of using such a module is that installation is accomplished by simply plugging the module into the system, which would facilitate field testing and ensure uniform results.

During the testing it became clear that the system still had many grounding issues that needed to be resolved. In order to ensure safe operation of the system, the coal pipes must be grounded throughout, mandating the use of insulated mounting studs, as was discussed in Section 4.1. Unfortunately, the mounting studs are easily damaged when they are repeatedly installed and removed, as was necessary in moving the transducer to enable data collection at multiple elbows. Once mounting studs become damaged, they are no longer non-grounding and collecting data with an acceptable level of noise is no longer possible. After several days of adjustments, it became clear that continued use of the stud mounts were likely to remain problematic in the field. While these issues might be resolved with the permanent installation of an instrument, it was not practical to continue further testing with the stud mounts for the balance of the program. Additionally, the analysis to date had produced far more encouraging results for data collected

with a magnetic mount of the transducer. Therefore, we decided it would be best to concentrate our efforts on obtaining a reliable instrument calibration using a magnetic mount. The previous test effort had clearly demonstrated that the sensor location one diameter downstream from the elbow generated an optimal signal as compared to the other locations that had previously been considered. Data collected at this location had more harmonic power and, as a result, had a higher signal to noise ratio. Given the delays that had resulted from time spent resolving noise issues, the decision was reached to focus exclusively, for the duration of the program, on data collected from a sensor at this position.

In sum, numerous adaptations were made to ensure uniformity of the data collection. After several days of modifications, a small quantity of very high quality data was collected using the magnetic mounts. In all, a total of ten data files were obtained using a vertical-down-to-horizontal pipe configuration. The air flow was approximately constant for this collection of case runs, and there were essentially three different coal flow conditions. The flow conditions for this data are summarized in Table 5 below.

Table 5. Conditions Visited in Downflow-to-Horizontal Testing using a Magnetic Mount

Air Flow (lb/hr, kg/sec)		Coal Flow (lb/hr, kg/sec)		Air/Fuel Ratio	Number of Files
14837	1.8695	7003	0.8824	2.12	1
14845	1.8705	6970	0.8782	2.13	3
14785	1.8629	8192	1.0322	1.80	3
14629	1.8433	11896	1.4989	1.23	1
14640	1.8446	11764	1.4823	1.24	2

4.5 Testing of August and October 2005

Though the noise issues had been thought to be resolved during the July test efforts, the process of determining the various noise sources had occupied the bulk of the test period. Therefore, the quantity of data collected was insufficient to develop the calibration, particularly as only three flow conditions were visited. Airflow Sciences agreed to collect more data to use in our analysis efforts.

In August 2005, Airflow Sciences collected and sent 60 data files to Foster-Miller. All data were collected with a sensor located one diameter downstream from the elbow and all with a magnetic mount. The set-up was reported to be unchanged from the final set-up in July. However, when the data was compared to the ten cases collected during July 2005, it was found that the signal in this latest collection had been extremely weak. In fact, the power of the signal had dropped 30-35 dB. Figure 14 below compares both the raw data and the power spectrum of one of the files collected in August with a file collected in July. While there is some evidence of peaks at similar frequencies in both data sets, clearly the harmonic power was very low for the data collected in August, which resulted in a signal that was dominated by noise.

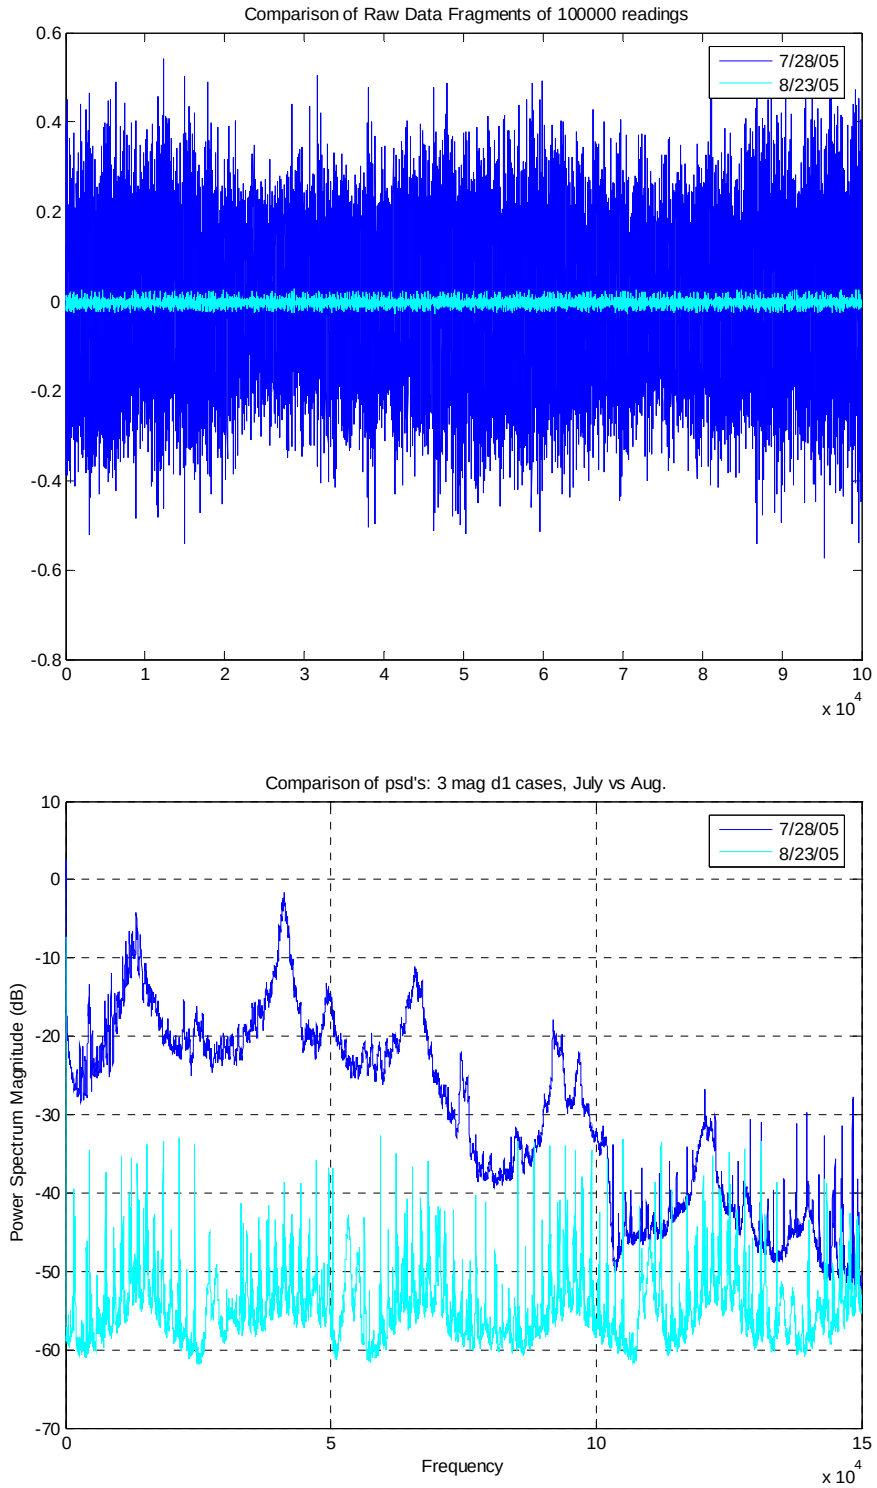

Figure 14. A comparison of two data files collected in July 2005 and August 2005

As a result of these problems, additional data was collected by Airflow Sciences in October 2005. The connections were checked to ensure that an adequate signal was generated. Unfortunately, the 39 files that were sent to Foster-Miller following this test run contained data with amplitude ten times that collected in July and, as a result, the data acquisition was saturating, as displayed in Figure 15, rendering it unsuitable for analysis. This was particularly surprising, given that ceramic particulate had been added to the coal, resulting in a mixture that was approximately 90% coal and 10% ceramic. One would expect the power of the acoustic signal to be somewhat lower than obtained previously when 100% coal was used, as had been the case in July, because the ceramic particles were lower in density than the coal. In sum, none of the data collected in August or October 2005 could be used to calibrate the instrument.

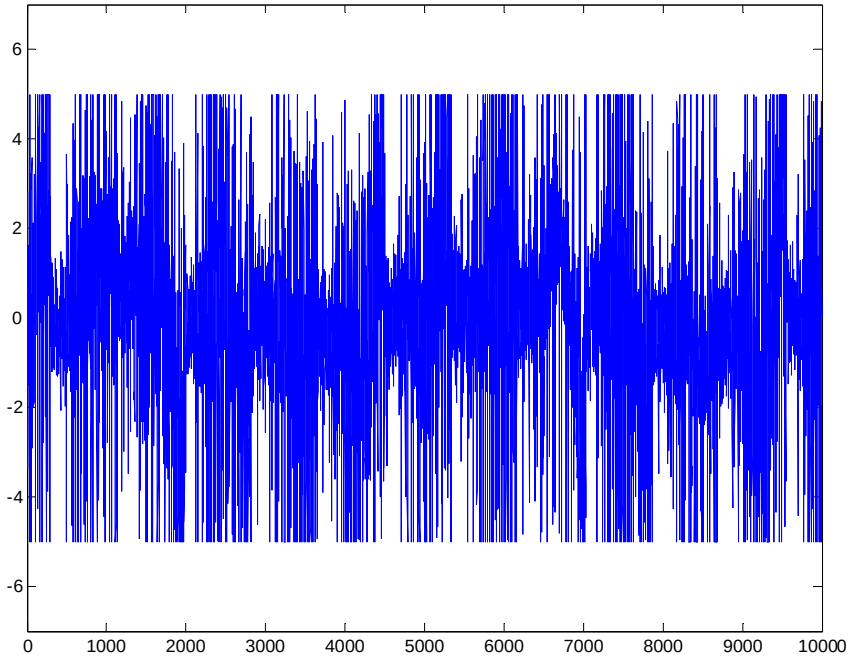

Figure 15. 10,000 readings of a sample data file from October 2005

4.6 Laboratory Testing of December 2005

Testing resumed at Airflow Sciences again in December, 2005. When it became clear that the problems experienced in August and October 2005 were persisting, Bruce Barck again visited the test facility in Livonia. Initially, the problem was resolved by tightening the accelerometer cable connector. This was a surprise to everyone from Airflow Sciences and Foster-Miller, as the connection did not appear to be loose and it was not obvious that this small change would

have such a large impact on the signal. Forty-one files containing high quality data were collected with this set-up. In all, 11 flow conditions were visited for which the airflow was near-constant and the coal flow varied from 5438 lb/hour to 11,098 lb/hour, as displayed in Table 6 below.

Table 6. Flow conditions for laboratory data with initial configuration, December 2005.

Air Flow (lb/hr, kg/sec)		Coal Flow (lb/hr, kg/sec)		Air/Fuel Ratio	Number of Files
16245	2.047	5438	0.685	2.990	1
16110	2.030	5584	0.704	2.890	5
16096	2.028	5898	0.743	2.730	1
16287	2.052	6229	0.785	2.610	3
16289	2.052	6694	0.843	2.430	5
16036	2.021	6799	0.857	2.360	5
16170	2.037	8541	1.076	1.890	4
15902	2.004	8913	1.123	1.780	5
15864	1.999	10670	1.344	1.490	5
15845	1.996	11037	1.391	1.440	4
16108	2.030	11098	1.398	1.450	3

The following day, the pipes were rearranged into a new configuration, as required for the EPRI testing, and the accelerometer was mounted once again, only to find that the previous noise problems had reappeared. At this point, Mr. Barck repeatedly disconnected and reconnected the accelerometer cable connector. In doing so, all the problematic behaviors that had been seen before resurfaced. Finally, following one of the reconnections, the problem appeared to be resolved. Five additional runs were recorded, four of which had flow conditions quite similar to conditions that were visited on the previous day. One of these had substantially lower air flow than all of the other December laboratory test runs. Table 7 below lists the flow conditions visited with the revised piping configuration. Figures 16-18 below contain photographs taken of the connector and its mount location during the first day of tests in December.

Table 7. Flow conditions following pipe reconfiguration, December 2005

Air Flow (lb/hr, kg/sec)		Coal Flow (lb/hr, kg/sec)		Air/Fuel Ratio	Number of Files
13086	1.649	5459	0.688	2.4	1
16351	2.06	5565	0.701	2.94	1
16303	2.054	6804	0.857	2.4	1
16286	2.052	8780	1.106	1.86	1
16150	2.035	11278	1.421	1.43	1

Figure 16. Close-up of accelerometer cable connector

Figure 17. Mount location

Figure 18. Piping configuration

4.7 Conclusions Regarding Instrumentation

The experience of the December test period reinforced the conclusion that the connector for the high-frequency accelerometer had been the source of many of the problems that have been experienced during this program. Disconnecting and reconnecting the connector produced dramatically different results despite the fact that there was no apparent difference in the connection from one try to the next. At this point, there was no way to tell if the flaw was due to a defective connector or whether the product itself is problematic, though its selection had been based on the fact that it is widely considered to be top quality among high-frequency accelerometers.

At first, the problems with the connector were a major concern as this particular part is quite costly and options for high-frequency connectors are extremely limited. Fortunately, as will be demonstrated in the analysis sections, we found that the best results were obtained when the data was filtered to pass a band of relatively low frequencies. The implications of this finding are that a lower-frequency accelerometer can be used, allowing a significantly wider range of sensors from which to choose. The lower-frequency accelerometers should be both more robust

and less expensive, which would permit a less expensive instrument system. Without the experimental evidence from this program, there was no way to be sure just how high a sampling rate was necessary.

4.8 Field Testing of December 2005

Field tests using the Foster-Miller instrument package were run during the week of December 12, 2005, at Detroit Edison's St. Claire Power Plant. The pulverizers were Babcox and Wilcox E-70 mills that underwent the conversion to EL-70 about 20 years ago. These are vertical spindle, air swept, ball mills.

Data was collected for Unit 1, Mills 1-5 and Unit 2, Mill 5. In both units, Mills 1 and 5 have three pipes per mill, each with pipe diameters of 13 inches. Mills 2-4 have two pipes per mill, each with pipe diameters of 14.5 inches. In contrast, recall that a pipe diameter of 12 inches was used in the laboratory testing. All tests were performed at locations where the flow was nominally upwards. For Mills 2-4, due to inaccessibility of positions downstream of an elbow, the sensor was placed approximately 1 diameter upstream of a 90 degree elbow. For Mill 1 and 5, the sensor was placed 1 diameter downstream of a 30-degree bend.

In comparison with the laboratory data taken earlier that month, the data from the field tests showed evidence of far more noise contamination and also displayed significantly less power in the frequency band under consideration. While it is not at all surprising that the signal transmitted to a sensor upstream of an elbow (as it was for three of the mills) would have less power, when combined with the higher noise level, the result was a far less favorable signal to noise ratio. Additionally, some of the raw data displayed erratic behavior, including unusual variation of amplitude. We learned from Airflow Sciences that some of these cases had been collected during extractive sampling, an invasive process that proved highly disruptive to the signal. Each of the 73 files collected was examined carefully, both by way of the raw data and the power spectrum. In this way, we weeded out data that would not be of use for our analysis. Of the 73 files, 28 were selected for further analysis. The result of this analysis will be discussed in detail in Section 6, as will the process of data selection. Tables 8 and 9 below summarize the flow conditions of the data collected in the field and also indicate both the number of files collected for each flow condition and the number of files selected.

Table 8. Summary of Flow Conditions for 13" pipe (Mills 1 and 5)

Air Flow (lb/hr, kg/sec)		Coal Flow (lb/hr, kg/sec)		Air/Fuel Ratio	Number of Files	Number of Files Used	Unit and Mill
18947	2.387	5592	0.705	3.388	4	0	U2 M5
22119	2.787	6120	0.771	3.614	4	2	U1 M1
21445	2.702	6733	0.848	3.185	4	1	U1 M1
20402	2.571	7284	0.918	2.801	5	4	U2 M5
18107	2.281	7423	0.935	2.439	6	0	U1 M5
20508	2.584	8194	1.032	2.503	8	7	U2 M5
21105	2.659	10224	1.288	2.064	5	0	U1 M1

Table 9. Summary of Flow Conditions for 14.5” pipe (Mills 2-4)

Air Flow (lb/hr, kg/sec)		Coal Flow (lb/hr, kg/sec)		Air/Fuel Ratio	Number of Files	Number of Files Used	Unit and Mill
27803	3.503	5683	0.716	4.893	7	1	U1 M2
26005	3.277	6040	0.761	4.305	4	0	U1 M3
25425	3.204	6749	0.850	3.767	6	0	U1 M4
30022	3.783	7911	0.997	3.795	7	7	U1 M4
26816	3.379	8266	1.041	3.244	7	0	U1 M2
29369	3.701	10474	1.320	2.804	6	6	U1 M3

5. PRELIMINARY ANALYSIS

Each of the data files collected in the program effort consisted of 9 million readings, acquired at 300 kHz for 30 seconds. This produced files of 36 MB, large enough to require planning and attention in their storage and transport, but small enough not to be completely unmanageable. With 30 seconds of flow dynamics in each file, any irregular variation of the dynamics with time (such as wandering of a central "rope", or passage of lower and higher particle concentrations) would be expected to occur quite a few times within a given data file. Thus, any reasonably sized subset of the data could be considered a different instance of the flow dynamics, so that a given data file could be considered multiple data sets that could be compared against one another. This proved to be very useful in the data analysis.

As was mentioned previously, the signal was hardware filtered prior to acquisition to limit the influence of behaviors above 100 kHz. This was done to eliminate aliasing, a classic issue in signal processing. This phenomenon arises because the Fourier transform of a time series data set is symmetrical about the Nyquist frequency, which is half the sampling rate. Basically, it is not possible to know whether a given behavior represents a low-frequency behavior or a high-frequency behavior that is "aliased" to appear as a low-frequency behavior. By applying the hardware filter to eliminate the high frequency behaviors, digital signal processing can be applied to the data without concern: what appears to be low-frequency really is low-frequency.

All of the data analysis was performed in the Matlab environment using software tools developed in previous projects. The first step in pre-processing the data was to digitally filter each data file to eliminate behaviors above 100 kHz, just to be certain that extraneous dynamics had not entered the data. For example, noise could be introduced in the analog-to-digital conversion, although this was not expected. Basically, it is difficult to check each of these large data files for unusual behaviors, so that it is easier to apply the filter first and not and eliminate one potential source of contamination. Each of the resulting data sets was then partitioned into 30 one-second snippets of data. A dynamic signature was calculated for each snippet using an ActiveX DLL, and then the median value was found for each of the signature quantities over the 30 snippets of a single file. This procedure greatly reduces the influence of outliers, a necessary precaution given that noise spikes have been found to be common in both laboratory and plant data.

5.1 Analysis of the First Round of Test Data

In analyzing the preliminary round of data collected in scoping tests, we were only concerned with the effects of the type and location of the transducer mount on the signal dynamics, as reflected by the signature quantities discussed in Section 2.3. Thus, we were not interested at the initial stage in relating the signature quantities for each data file to the flow conditions, but instead were studying the trends exhibited in the signature quantities calculated for a given flow condition as the type and location of the transducer mount varied. If the statistics had varied in a relatively simple manner with transducer location, this would have suggested that a universal calibration could be found for prediction of flow parameters (coal flow, air flow, and coal fineness) irrespective of transducer mounting location. Similarly, if the effect of the transducer

mounting type (stud or magnetic mount) on the statistics was relatively simple, then the application of the instrument could be broadened using both types of mounts. The primary goal was to determine the type and location of transducer mounts to be pursued in further testing.

The initial analysis of the scoping data indicated that the type and location of the transducer mount had strong effects on the signature quantities. The results presented below led to our decision to limit mount locations for future testing for the remainder of the program. In fact, later on, following a more comprehensive analysis, it became clear that relationships between dynamic signatures based on data collected at different locations of the transducer mount were obscured by various types of noise, as will be further discussed in Section 5.2.2. This in turn implied that the relationships between the statistics for data collected at different transducer-mount locations might not be as complicated as they first appeared. In the long run, a universal correlation, independent of mount location, might be more within reach than it first appeared. While this development held promise for our long-range goal, the decision to limit mount locations for this program was upheld as the noise issues limited the quantity of "good data" available.

Sample results from the preliminary analysis are presented in Figures 19 and 20, from flow condition 5 (equal air and coal flows of 15,670 lb/hr). Figure 19 displays the graph of the standard deviation of the signal as a function of position for both the stud and magnetic mounts. One striking feature of this graph is how different this variation with position is for the two different mounts. For the stud mount, position 6 (corresponding to the outlet of the elbow) produces the strongest signal, with the signal becoming weaker in a manner that is essentially symmetrical for positions upstream and downstream of this location. This behavior can be seen in quite a few of the signature quantities that measure the amplitude of the signal. By contrast, the behavior of the magnetic mount is quite different, with the signal at position 6 being weaker than any other location. Both mounts produce similar results for positions well downstream of the elbow.

The signature quantity shown in Figure 20 is a measure of the period for the passage of the largest events in the signal. Larger values for this statistic reflect a longer interval, on average, for the passage of large disturbances. The behaviors for the stud and magnetic mounts are somewhat similar for locations downstream of the elbow, but markedly different for locations upstream of the elbow.

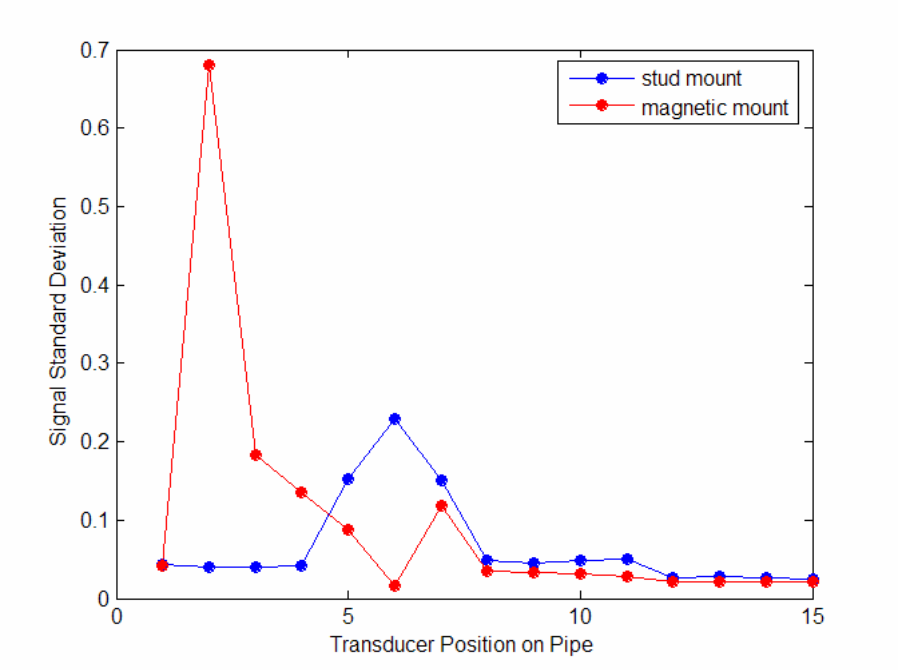

Figure 19. Standard deviation of signals as functions of transducer position and mount

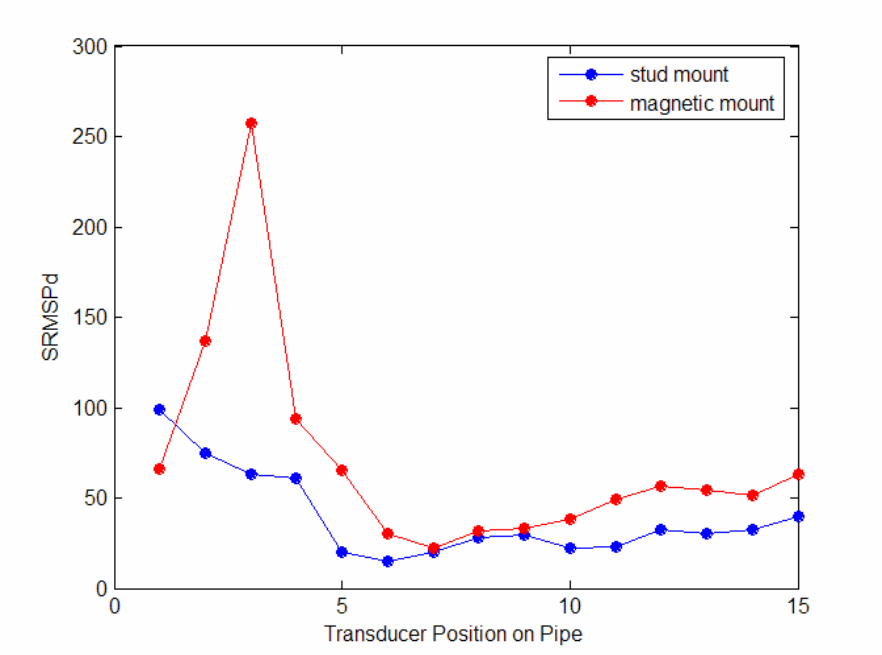

Figure 20. A time measure of the signals as functions of transducer mount and position

The results for other flow conditions were generally similar to the results shown in Figures 19 and 20, and revealed quite a bit of variation in the behaviors of some signature quantities with position and mount type. The net result was that there did not appear to be a simple means of identifying the effects of the mount type and position of the transducer on the signal dynamics.

These results did not indicate that a universal correlation cannot be achieved. In fact, there must be identifiable effects of mounting type and location on the signal dynamics, but there is no requirement that these effects be simple. In principle, given enough experimental data, these effects can be "learned" by a suitable analysis, so that the flow conditions can be determined for any reasonable transducer mounting type and location. The problem is that developing such an analysis could potentially require a very large quantity of data, which could not be obtained within the constraints of this program. Whether the actual effects of the mounting type and location are complicated or, as we now believe, noise made the effects appear more complicated, the objective of this program was to develop commercially viable instrument systems for both fixed and portable applications, and limiting the transducer installation to specific locations was a reasonable compromise at this point. Experience gained through field use may provide the data required to expand the instrument calibration to different transducer locations, but it appeared that the best approach for the present program was to concentrate on the issues that are of greatest interest:

- Given a standard transducer mounting position (e.g., at the outlet of an elbow), can we generate an instrument calibration that works favorably for a variety of pipe sizes, orientations, and coal flows?
- Can both magnetic and stud mounts be accommodated?

As was noted in Section 4.3, the test matrix for the following round of tests, performed in October, 2004, was designed with the view towards answering these questions.

5.2 The Impact of Noise

Once the October, 2004, data collection had been completed, we could begin to focus our efforts on the actual data calibration. Frequently, in performing a Dynamical Instruments analysis, the computation of the dynamical signature itself effectively filters out much of the noise. The signature selection process, and subsequent training of neural nets, generally results in a dynamical signature consisting of signature quantities that are less sensitive to noise and more sensitive to the key dynamics. Therefore, originally we had hoped that digital data filtering would provide sufficient noise reduction.

The first stage of the analysis, presented below in Sect 5.2.1, produced disappointing results from our point of view. However, surprisingly, when these results were presented at an EPRI meeting at the Coal Flow Test Facility in Livonia, MI, this initial attempt fared surprisingly well in comparison to two commercial coal flow meters they had tested. Nonetheless, in comparison to other Dynamical Instruments projects, we did not feel that the results were satisfactory.

At this point, we began to concentrate more closely on characterizing the impact of noise. Section 5.2.2 will describe the results of a case-by-case examination of the data, which ultimately resulted in the elimination of 28 files of the 209 cases collected during the first round of testing, and 38 files of the 181 cases collected during the second round. Sections 5.2.3 and 5.2.4 will then describe further efforts to develop an analysis that would be less noise-sensitive. While the results of these efforts resulted in a much improved calibration, they still were not yet what we would call "instrument-quality". There were many indications of multiple noise sources that were effectively obscuring the flow dynamics. Therefore, it became clear that a high priority of the next test effort must be locating the noise sources and, where possible, revising the data collection techniques in order to eliminate much of the noise, so that the analysis efforts could focus more on calibrating an instrument with flow conditions and less on noise reduction.

5.2.1 A First Pass at Calibration

In analyzing the combined data sets of August and October, 2004, a pre-processing algorithm similar to that described earlier was followed. In order to reduce the impact of pipe resonance on the observed dynamics, we decided to digitally low-pass filter the data, eliminating behaviors above 75 kHz, rather than 100 kHz as had been used previously. Signature quantities were then computed based on the filtered data. The signature quantities were organized into a number of different arrays, which considered data from various combinations of the following variables:

- Transducer mounting types
- Transducer mounting locations
- Different pipe elbows

The various arrays included data from both the earlier scoping tests and the more recent calibration tests. The signature arrays were then analyzed using Particle Swarm Optimization to identify correlations between the flow dynamics and the coal flow rate.

Figure 21 is a typical result that was obtained for a stud-mounted transducer located at the outlet of the vertical-to-horizontal elbow. An ideal result would place all points on the diagonal line, producing a correlation coefficient, r^2, of one. The actual correlation coefficient was 0.946 in this case, which is not a bad result but does not represent a high-precision instrument.

The results for the magnetic mount were not as favorable, as illustrated in Figure 22. In this case, a correlation coefficient of 0.875 was obtained, a significant reduction in accuracy compared to the result in Figure 21.

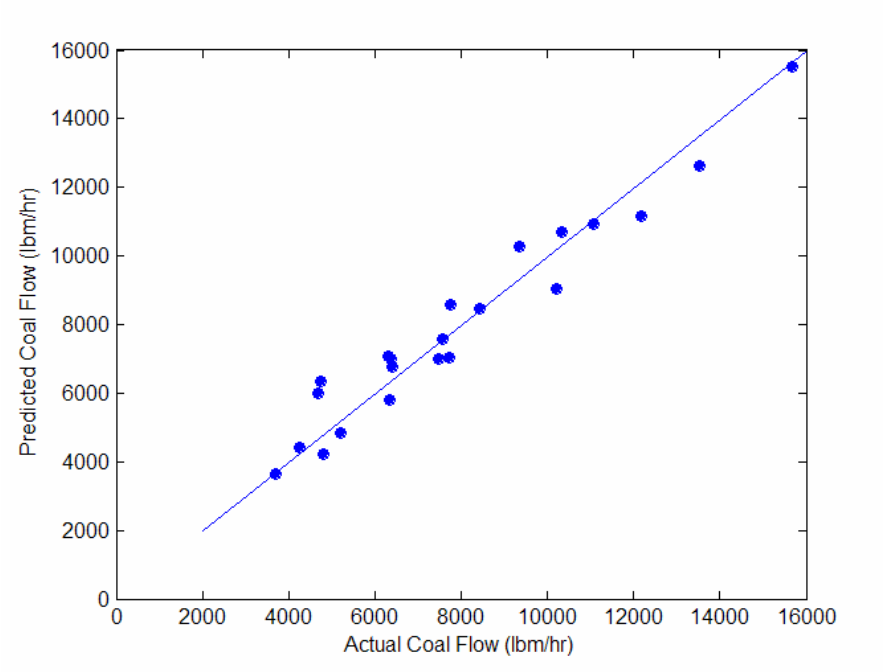

Figure 21. Typical result for stud-mounted transducer for a specific elbow mount

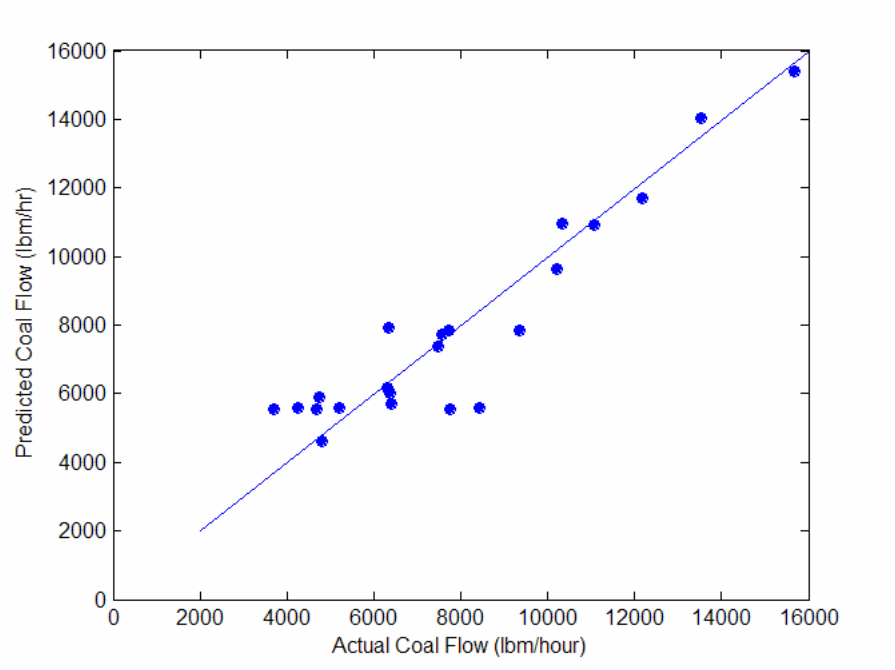

Figure 22. Typical result for magnetically-mounted transducer for a specific mounting location

The results degraded even more markedly when data for the different elbows are combined in one correlation, illustrated in Figure 23. In this case, the correlation was based on data for stud-mounted transducers in all 3 mounting situations, including the vertical-to-horizontal elbow, the horizontal-to-vertical elbow, and the horizontal-to-vertical elbow with a roping configuration. The correlation coefficient here is 0.677, far below the level one would consider acceptable in a commercial instrument. Nonetheless, as previously noted, none of the results obtained from the commercial instruments were a significant improvement over the Dynamical Instruments analysis.

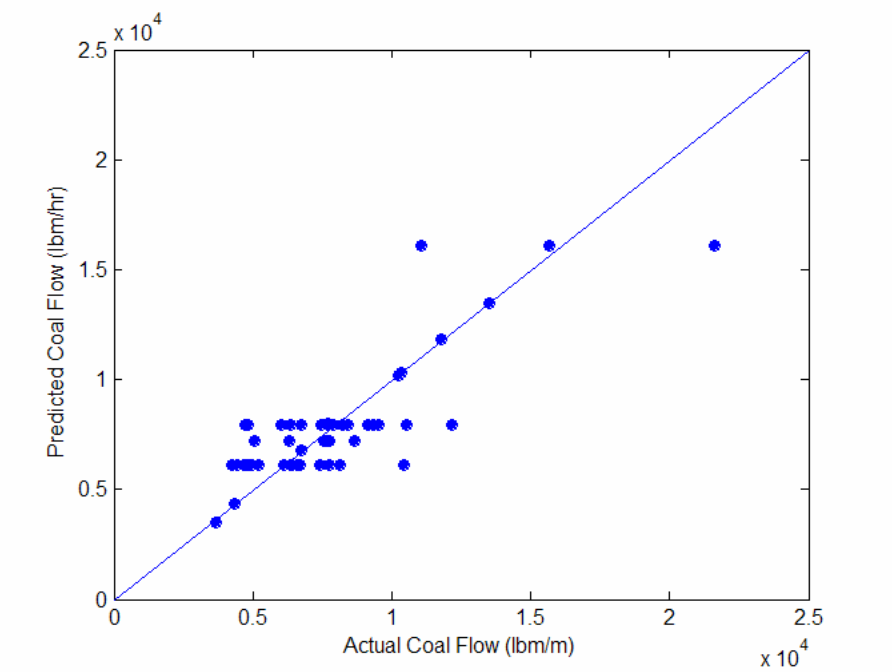

Figure 23. Typical results for a stud-mounted transducer for all elbow mounts

5.2.2 A Case-by-Case Examination for Noise

It was discovered in analyzing one of the data files that not all of the noise had been eliminated from the data through the debugging effort early in testing. The first type that was identified was a relatively large-amplitude, low-frequency behavior, such as that shown in Figure 24. Assuming that this noise is simply additive to the otherwise correctly sampled dynamics, this type of noise can be eliminated by digitally high-pass filtering the data. Consequently, it was decided to high-pass filter data files to suppress behaviors below 20 kHz in calculating dynamic signatures for the remaining analysis.

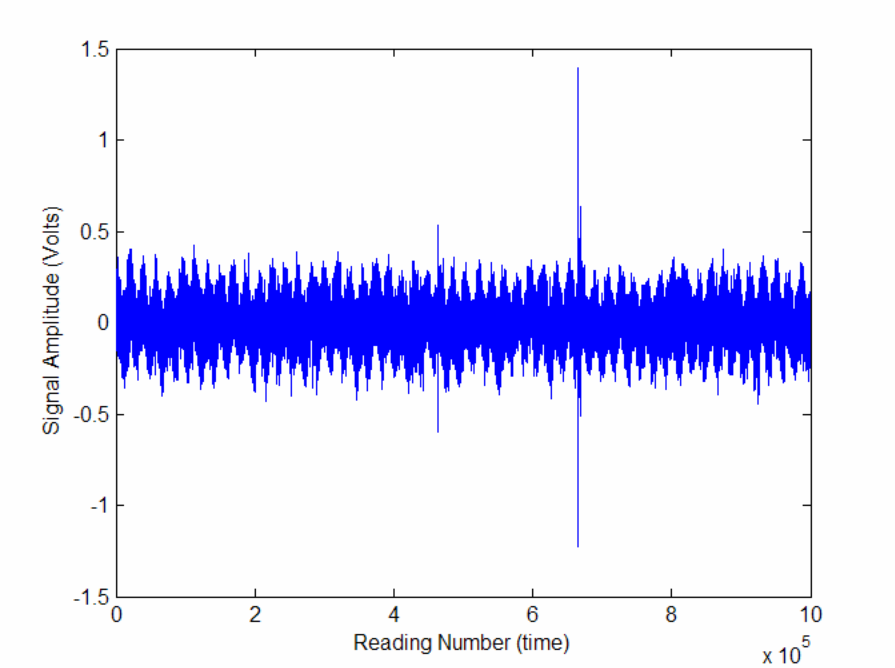

Figure 24. Typical data with low-frequency noise

The example above led us to perform an exhaustive examination of all of the data obtained to date. We had not performed this screening earlier because of the sheer volume of data. In this process, each data file was opened, graphed at two different time resolutions, and its power spectrum graphed. Figure 24 above is the result of graphing the first 100,000 data points, comprising a mere 1/3 of a second of data. A second kind of noise is found in graphing the entire data set, shown in Figure 25. The central section of the data has an amplitude that has proven to be typical for the bulk of the "good" data sets. The earlier and later sections presumably suffer from additive noise. When a data file like the one shown in Figure 24 was encountered, the range of data that displayed "normal" amplitude was noted, and signatures calculated for that limited window only.

Another form of noise visible in full-file graphs is shown in Figure 26. It is not at all clear what combination of instrumentation, data acquisition, and/or system conditions could produce such a bizarre behavior, but it clearly is not representative of flow dynamics. Such files were discarded out of hand.

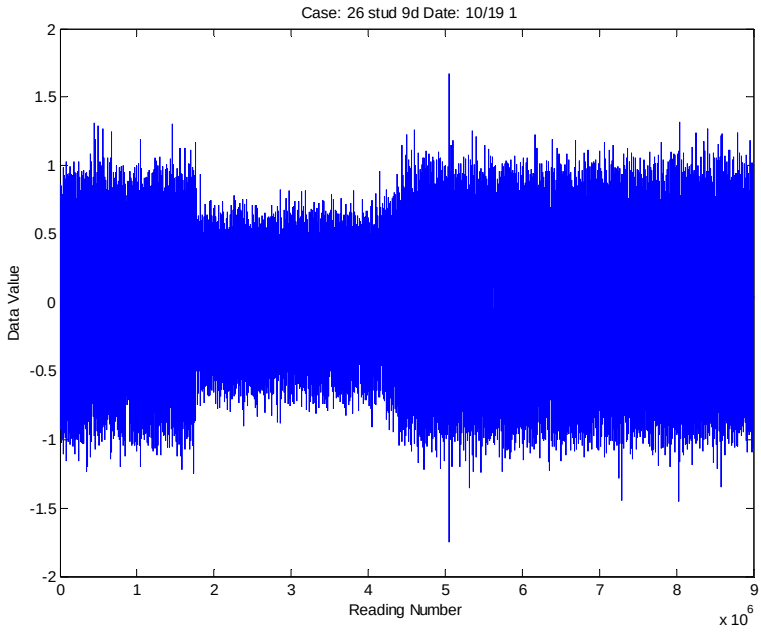

Figure 25. Noise producing a varying signal amplitude

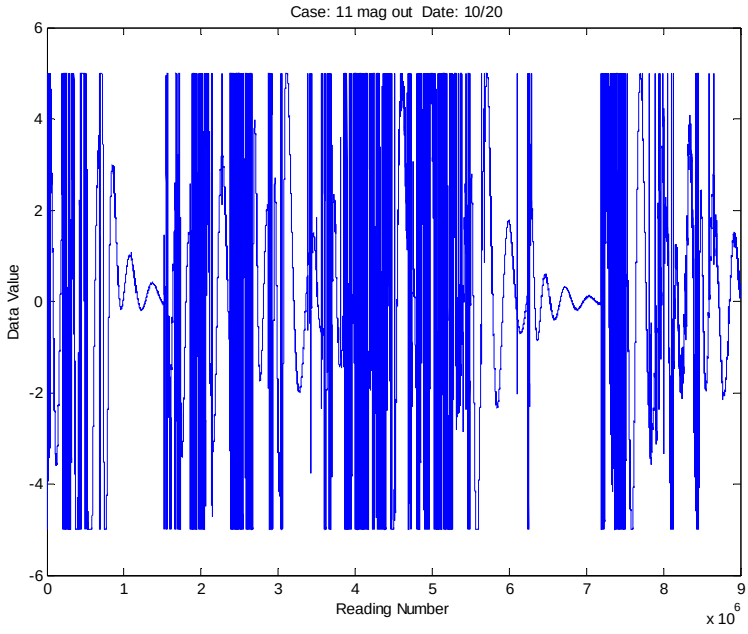

Figure 26. Over-the-top noise

Yet another form of noise can be seen in a power spectrum of a data file, as shown in Figure 27, In this case, the power spectrum is essentially flat from DC up to the roll-off produced by the hardware low-pass filter in the instrumentation package. Such a behavior represents pure white noise, a behavior never encountered in previous Dynamical Instruments development efforts. Again, it is not clear what combination of instrumentation, data acquisition, and/or system conditions could produce such a behavior, so these cases were discarded as well.

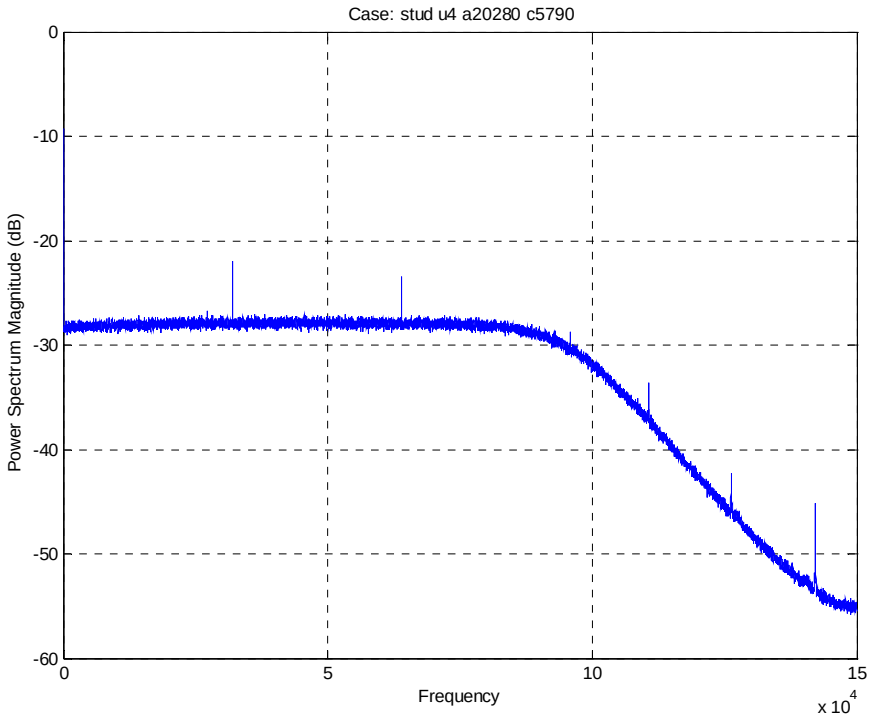

Figure 27. Noise producing a flat power spectrum

A significant number of files were discarded through this screening process. Of the 209 data files in the original scoping data, 28 cases were discarded. Of the 181 files in the more recent calibration set, 38 were discarded. In general, the cases with magnetic mounts tended to be noisier than with the stud mounts: of the 95 stud mount cases in the calibration data, only 12 were discarded, while 26 of the 86 magnetic mount cases were discarded.

Figures 28 and 29 provide an indication that this screening process was on the right track. Figure 28 plots the signal standard deviation calculated for tests performed at a single flow condition with both mount types and various transducer mounting locations. The data had been filtered to pass frequencies between 20 kHz and 75 kHz (the upper limit being chosen to avoid over-emphasizing behaviors at the transducer's resonance in the neighborhood of 85 kHz). A

similar graph was presented as Figure 19 in Section 5.1. The current graph looks the same, but identifies the first four points for the magnetic mount as noisy cases. Ignoring these cases, the variation with position is seen to be much smoother.

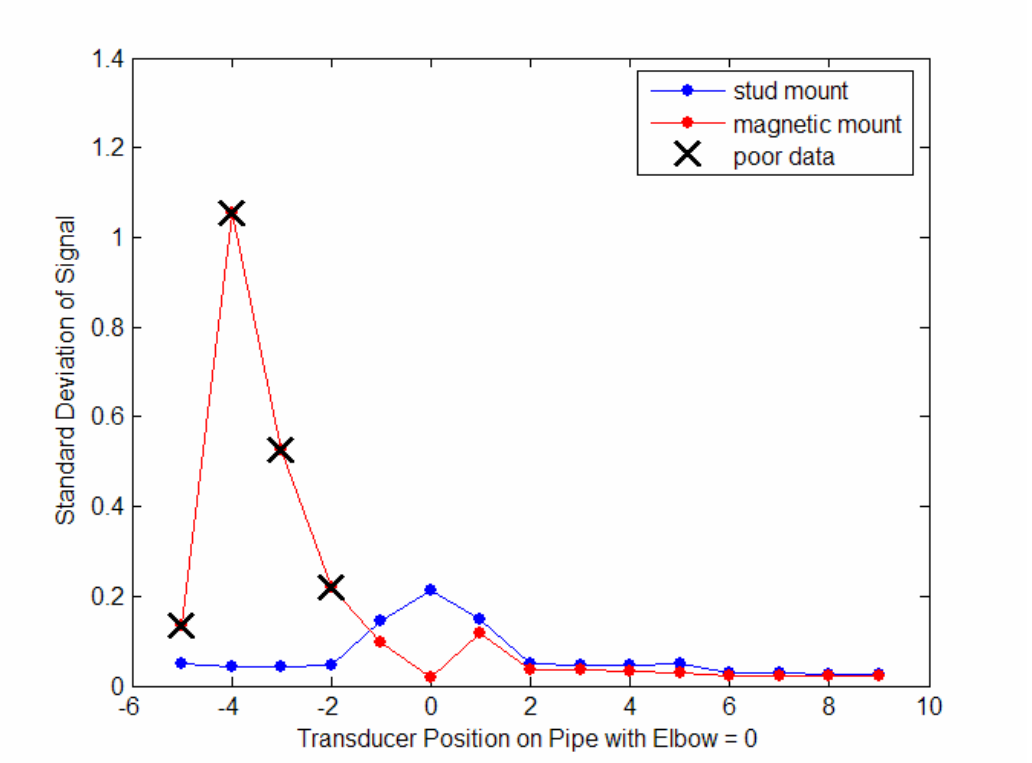

Figure 28. Signal standard deviation as a function of position and mount type

Figure 29 plots the value of another signature quantity as a function of transducer position and mounting type. This figure is similar to Figure 20 in Section 5.1, but the values of the statistic are now quite different. The reason for this is that filtering the data to eliminate low-frequency behaviors changes the character of the signal. The net result, ignoring the four noisy cases, is that the variation with position is quite smooth. Overall, the behavior of the signatures, after eliminating the noisy cases, is more in line with what had been anticipated based on our experience in earlier programs.

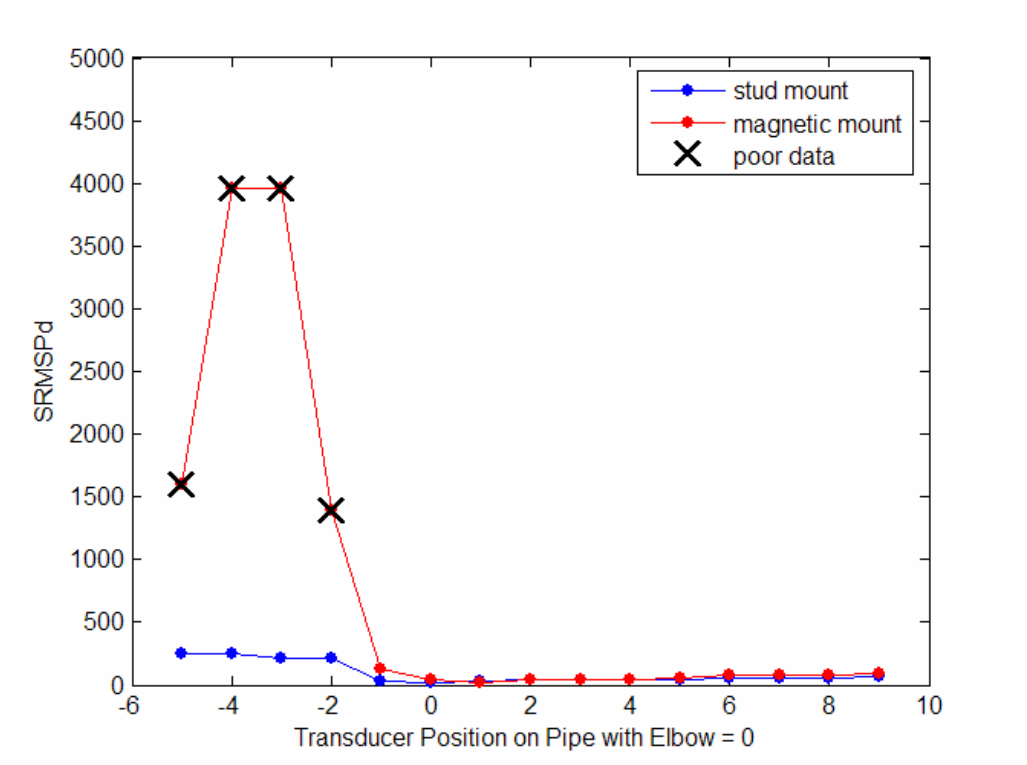

Figure 29. A time measure of the signal dynamics as a function of position and mount type

5.2.3 Bimodal Correlation of Coal Flow to Dynamic Signature

After eliminating the cases that had been identified as noisy, an attempt was made to correlate all of the magnetically-mounted cases where the transducer was mounted at an elbow outlet. The procedure was repeated numerous times to ensure that the best result was obtained. The analysis typically produced correlation coefficients in the general range of 0.75. This was better than the results of less than 0.7 obtained in earlier analysis, but not to the extent that was hoped for. The outcome was surprising, as there was no particular reason to believe that any of the remaining cases were noisy. It was by sheer accident that one particular correlation was graphed, with the extraordinary result shown in Figure 30. In this case, the overall correlation coefficient was 0.78, but the graph clearly shows that there are 2 separate behaviors. One subset of the cases has an excellent correlation coefficient of 0.994, while the other is essentially uncorrelated (r^2 = 0.014). The well-correlated cases, approximately half of the total, indicate that at least some cases sensitively disclose dynamics that are related to coal flow. The fact that roughly half of the other cases did not share the same dynamics suggested that there might be yet another, as yet unidentified form of noise in some cases. Fortunately, since the uncorrelated cases all fall within a narrow band of predicted flow, it should be possible to discern cases with normal dynamics from those with noisy dynamics.

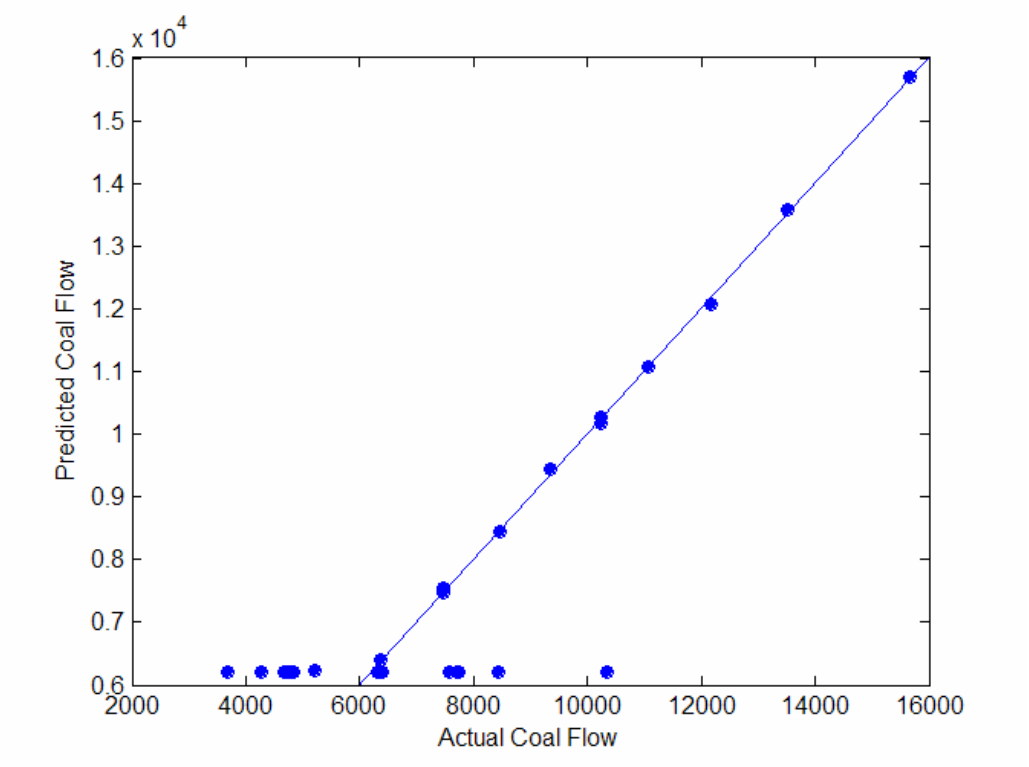

Figure 30. Bimodal correlation obtained after removing identifiably noisy cases

The well-correlated data points in Figure 30 comprised the best result we had encountered with coal flow data in the many data sets we had collected over numerous test efforts, confirming that observed dynamics should be sensitively related to flow. This correlation accuracy is typical of the results we have obtained with the Dynamical Instruments technique in other flow measurement applications.

As these results suggested that an additional noise source was obscuring the dynamics, the focus turned to the identification of noisy data, with an eye towards locating measures less sensitive to noise. Using Particle Swarm Optimization, neural networks were trained on Dynamical Instruments signature quantities to distinguish the two sets of cases, those that were well-correlated with coal flow and those that were not. Using just two inputs, the noisy cases could easily be distinguished from the others with a 100% success rate. However, examining pairs of signature quantities that were used as inputs to achieve this separation, little insight was gained towards qualitatively distinguishing the two sets of cases or identifying the noise source. Figure 31 shows a typical behavior of a signature-quantity pair that distinguishes those cases where coal flow is well-predicted from those where it is poorly predicted. The data does not fall cleanly into two distinct clusters. Rather, there is a cluster of cases with intermediate signature-quantity values for which information related to coal-flow could not be extracted.

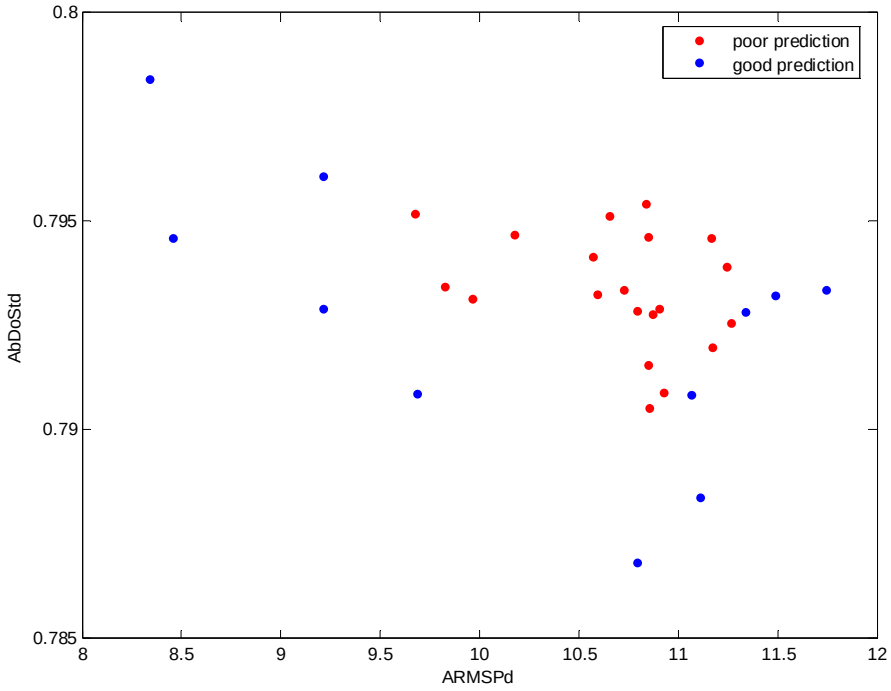

Figure 31. Separation of cases with good and bad predictions

The noise identified in the data provided an explanation for much of the trouble that had been encountered in analyzing the data collected to this point. The noise could be identified, either by direct analysis of the time series or by comparison of the signal dynamics to non-noisy cases. Still, this did not explain where the noise originated. In addition, the data remaining after eliminating noisy cases was of such a modest quantity that it would be insufficient for the development of a reliable instrument calibration.

Consequently, we felt that additional data should be collected in the Coal Flow Test Facility to complete the instrument calibration before plant testing was undertaken. By piggy-backing on planned EPRI testing, this testing could be performed at very little cost to the program, but necessitated an extension to the program schedule. A no-cost time extension was requested and approved.

Until new data would be made available to us, there was still information to be gleaned from the current data set. While our hope was that noise issues would be substantially reduced during the next round of testing, there would be a considerable advantage to fine-tuning the analysis techniques in order to reduce the noise sensitivity of the results. We first performed a comparative analysis to optimize the frequency range used when digitally band-pass filtering the data during the preprocessing stage. The results of this analysis are described in Section 5.2.4. We then examined the length of the data windows used in computing the dynamic signature and

were able to significantly improve the results by comparing dynamic signatures for windows of different lengths. This will be discussed in Section 5.2.5.

5.2.4 Comparison of Ranges for Band-Pass Filtering

Based on the earlier observation that band-pass filtering had produced better results than merely low-pass filtering, the noise problem was now approached from the point of view of optimizing the window of frequency values that were passed in filtering. Focusing on the elbow-outlet, magnetic-mount cases, again excluding those that were identifiably noisy, each raw data file was filtered six different times to allow a comparison of the results obtained when each of the following frequency ranges were passed:

- 20-40 kHz.
- 30-50 kHz.
- 40-60 kHz.
- 50-70 kHz.
- 20-60 kHz.
- 40-80 kHz.

As before, each filtered data file was broken up into 30 one-second windows of data, with signature quantities computed for each window. Again, only the median value of the signature quantities for each data file was used, in an effort to reduce the inaccuracies due to noise spikes. The resulting median signatures from each frequency window were then trained thirty times using Particle Swarm Optimization with five inputs and three hidden nodes. Repeated training is needed to optimize the performance of the neural net due to the fact that success is highly dependent on initial training conditions. Certainly, the most important result is the highest r^2 value obtained. However, from a practical point of view, accessibility of information can ultimately have a major impact on the best value achieved. The results of this analysis are summarized by Table 10 below.

Table 10. Comparison of Results in Passing Different Frequency Ranges

Frequency Range in kHz	Highest r^2 value obtained	Number of runs with $r^2 > 0.70$
20-40	0.77	6 out of 30
30-50	0.75	1 out of 30
40-60	0.76	2 out of 30
50-70	0.78	4 out of 30
20-60	0.76	4 out of 30
40-80	0.69	0 out of 40

Overall, the highest r^2 value obtained was not all that different for the different frequency ranges, with the exception of the 40-80 kHz range. Beyond that, for some of the frequency ranges the

information was more accessible, with more runs resulting in the higher r^2 values. However, overall most frequency windows contained comparable information and comparable noise interference as well, as is indicated by Figure 32 below. Further analysis was restricted to the data filtered to band pass 50-70 KHz, the frequency window that yielded the highest correlation value.

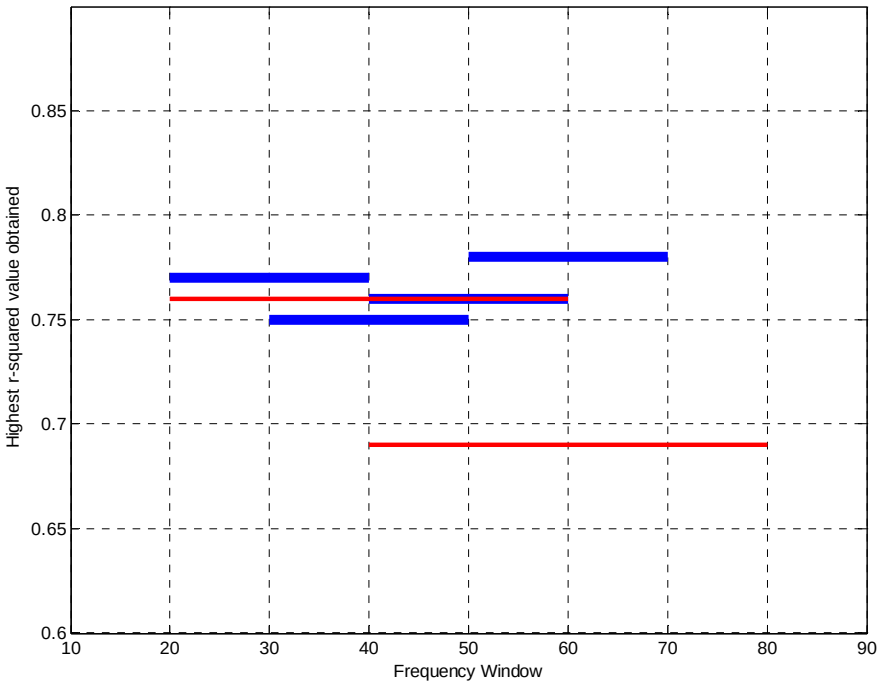

Figure 32. Comparison of r-squared values achieved for different frequency windows

5.2.5 Coal Flow Predictions from Temporal Variation of Dynamic Signature

An examination of the signature-quantity values for individual one-second windows revealed significant differences in the temporal variation from case to case. This is illustrated in Figure 33 below for a representative signature quantity, sdADuSize. That observation in turn raised the question of whether the temporal variation might itself be related to coal flow, or whether it was a reflection of the instrument noise.

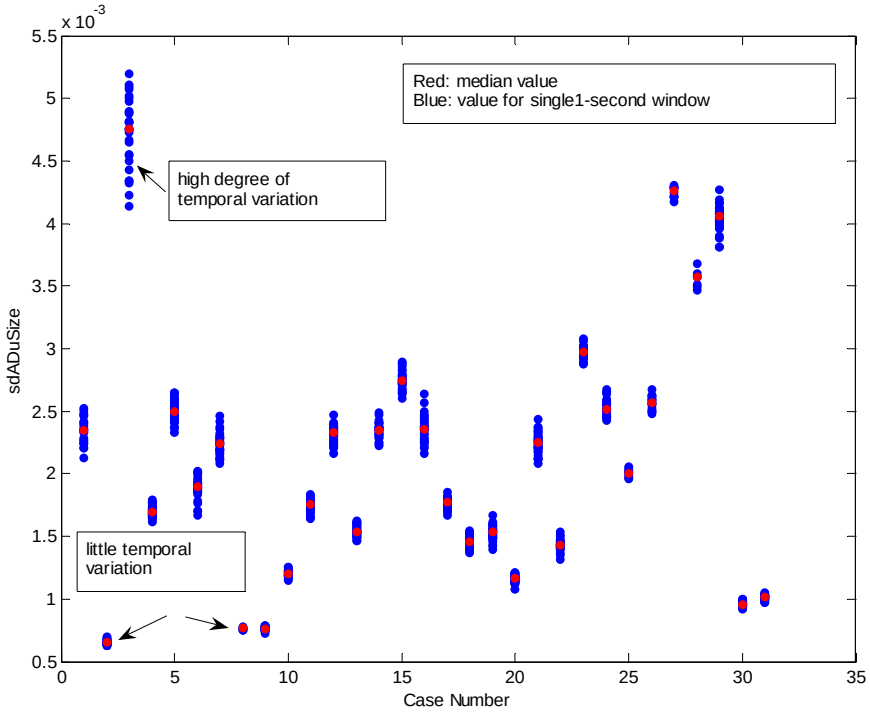

Figure 33. Temporal variation of sdADuSize over each 30-second file

A simple measure of the magnitude of temporal variation is the standard deviation. In order to address the above question, the training procedure that had been applied previously was now repeated with the median of each signature quantity replaced by the standard deviation of each signature quantity over the 30 one-second windows for each case. This training resulted in a significant improvement of results, attaining a maximum r^2 value of 0.85, as compared to a previous maximum of 0.78. Figure 34 below displays two histograms, providing a comparison for the correlation results based on the median of each signature quantity with those based on the standard deviation. Clearly, more information is extracted when the analysis is based on the standard deviation. Combining both dynamic signatures did not improve the analysis further.

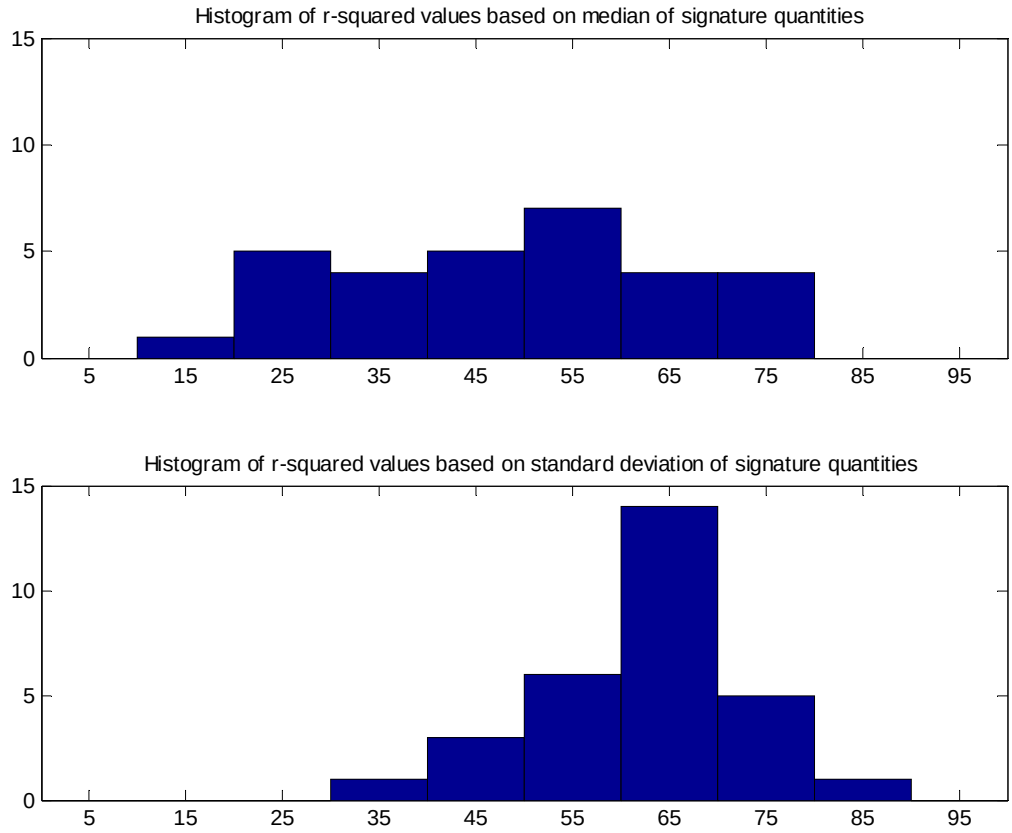

Figure 34. Histograms comparing correlation results based on median of each signature quantity to standard deviation of same

The fact that the temporal variation of the dynamic signature over a test run provided more information concerning coal flow than the median value suggested that one-second windows of data were inadequate for extracting coal flow values and that the use of longer windows was warranted. Therefore, the one-second windows of data were replaced by five-second windows. Consecutive windows had starting points separated by one second and therefore overlapped by four seconds, once again resulting in a total of 30 windows for each data run. Training was now attempted as before, but this time using the median of the five-second windows of each signature quantity over each file. The results were not significantly different from those obtained based on the one-second windows. A maximum r^2 value of 0.79 was achieved as compared to that of 0.78 using one-second snippets. Both these results were based on the data filtered to pass a frequency range of 50-70 kHz. Figure 35 below compares the histograms for 30 trainings based on the one-second windows to 30 trainings using five-second windows. At best, a modest improvement is attained with the longer windows.

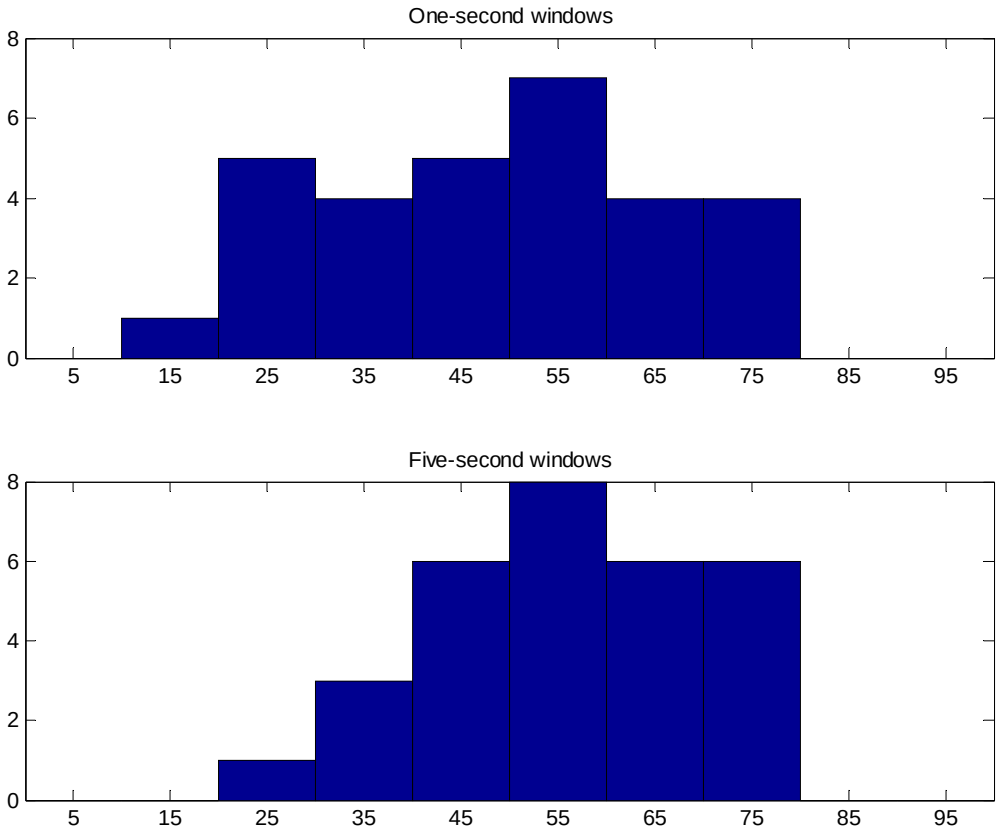

Figure 35. Results based on one-second windows vs. those based on five-second windows

In short, the signatures of longer snippets of data produced lesser results than the standard deviation of each signature quantity over a longer period of time. A comparison of the medians of the one-second signature quantities to the median of the five-second signature quantities revealed that there were two types of behaviors. While all of the median signature quantities displayed similar behavior for the one-second windows and the five-second windows, for some signature quantities the values were essentially independent of the window size while other signature quantities displayed subtle variations in value between the two different window sizes. Figures 36 and 37 display representative examples of each of these two types of signatures.

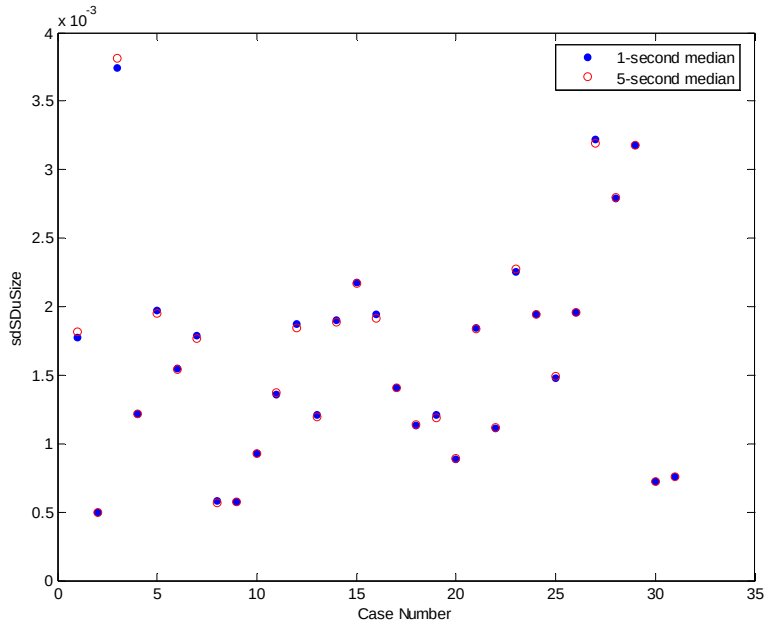

Figure 36. A signature quantity for which the median value is essentially independent of window length

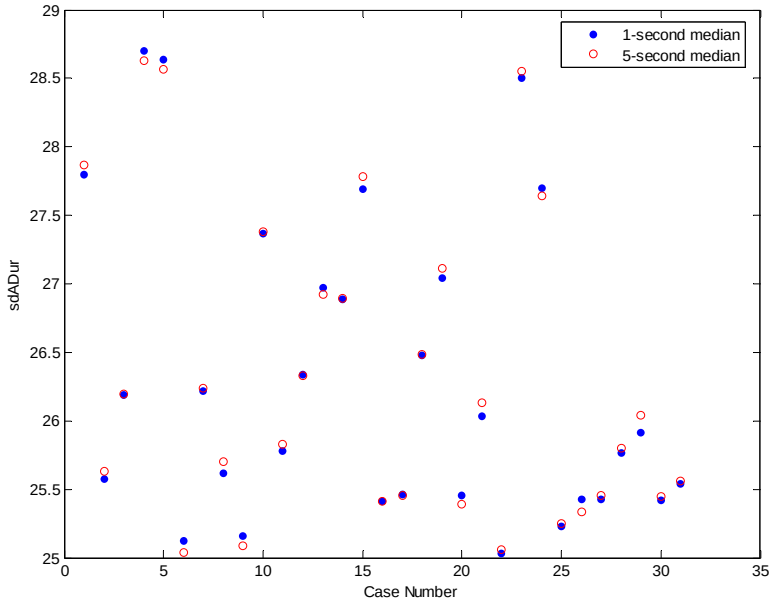

Figure 37. A signature quantity for which the median value varies with window length

Given the fact that the standard deviation of the signature quantities had better captured temporal variation than the signature quantities over longer intervals, the question arose whether the difference between the median values of the one-second signatures and the five-second signatures might in fact be more sensitive to coal flow than the values of the signature quantities themselves. Training was repeated once again, this time based on the differences between the medians of the five-second signatures and the medians of the one-second signatures. This improved the results dramatically, achieving a maximum r^2 value of 0.97. These results are shown in Figure 38 below. Note that a "perfect" result would place all data points on the diagonal. Referring back to Figure 30, the best results obtained based on the median signatures for the 1-second snippets, the level of improvement is quite striking. Previously, the best neural net yielded no information at all about the coal flow for a substantive percentage of the data, namely, those cases mapped to the horizontal line at height 6212 in Figure 30. In contrast, the current analysis performs reasonably well on the entire collection of data.

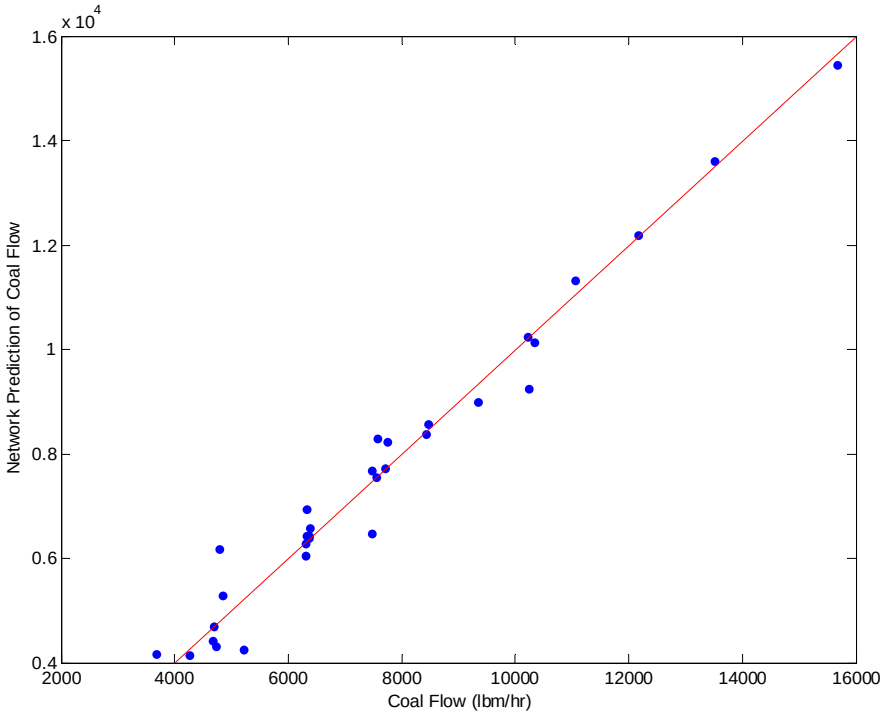

Figure 38. Best neural net prediction of coal flow, based on difference between median of one-second signatures and median of five-second signatures

The results displayed in Figure 38, although not yet of precision instrument quality, gave rise to a fair amount of confidence that coal flow information could be extracted with the techniques at hand. It should be noted that the number of cases involved, 31, was still quite small and the quantity of data was insufficient to permit a test set. Still, the fact that the product of the number of inputs and the number of hidden nodes was less than 50% of the number of cases precludes

the possibility of the neural net simply "memorizing" the data. These cases included all pipe arrangements described Section 4.3.

A closer examination of the signature quantities used as inputs for the optimal neural net revealed some unusual trends in the values over time. For some test runs, the values of the one-second and five-second signature quantities showed little variation with time. For other signature quantities, there was a distinct trend, with the value either increasing or decreasing over time. Figure 39 below contrasts the autocorrelation signature quantity for two different test runs with similar coal flow and air flow. The top half of the figure illustrates a case in which the autocorrelation clearly decreases over time while the bottom half of the figure illustrates a case with no such trend. This distinction could not be correlated with coal flow, air flow, pipe arrangement or test date. It remains unclear whether these trends were related to the dynamics or noise.

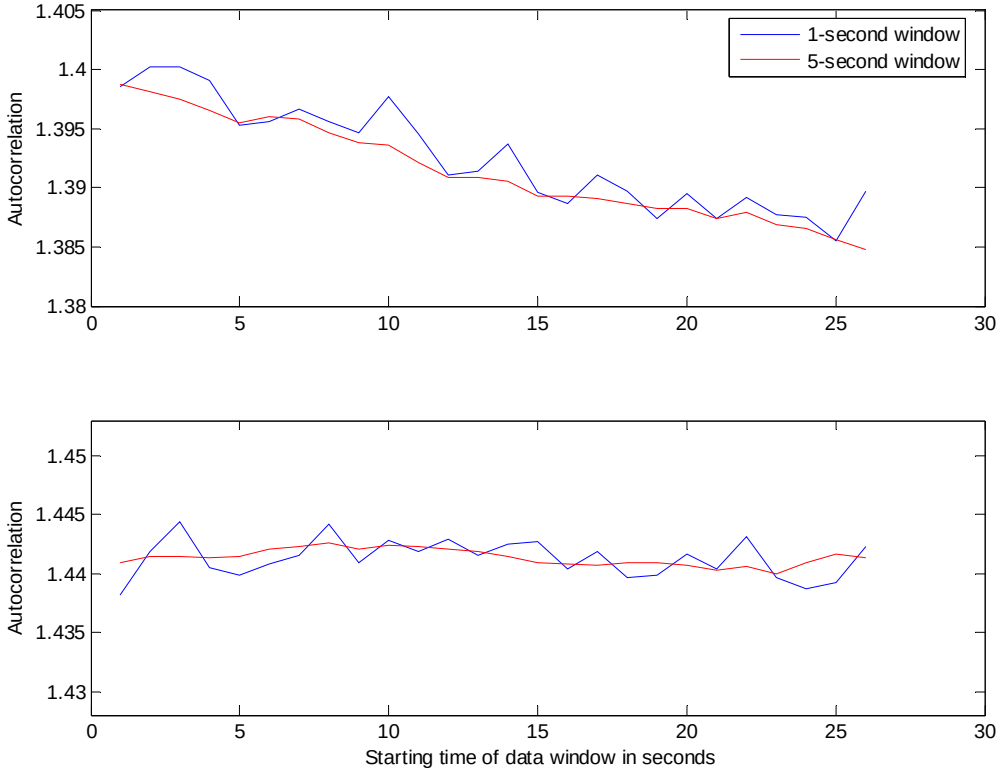

Figure 39. Trends of autocorrelation over time for two test runs with similar coal flow

Two possibilities came to mind to explain the improved predictions in using the difference between median signatures for the longer and shorter windows over using the median signatures for a single window length,. We already knew that noise has been a major problem with this data. It was possible that subtracting one median from the other effectively subtracts effects of

the noise, with the remaining measurement a better reflection of the dynamics of the flow system. A second possibility is that the degree to which a one-second window of data effectively populates a state space as well as the five-second window may in some way be dependent on the coal flow. Our view at this point was that the analysis of additional data to be collected at the Coal Flow Facility would allow us to determine the extent to which the improvement in the analysis was due to reduced noise sensitivity and the extent to which the current algorithm had better pinpointed those aspects of the dynamics that vary with coal flow. If the temporal variation turned out to be related to the dynamics, the signature quantities themselves could be fine-tuned to better capture the temporal variation of the dynamics to create a more incisive algorithm. However, as we will see in Section 5.3 and Section 6, based on the later data collections, it now appears that the improved predictions being discussed here were due to noise reduction by way of the subtraction process.

5.3 Ten Noise-free Cases

As noted earlier, major efforts were directed toward noise reduction during the July, 2005, test round. While this effort occupied most of the test period, once the modifications were put into effect, ten files of remarkably consistent data were obtained. Although the quantity of data and flow conditions represented were not sufficient to produce an instrumentation calibration, they could be analyzed to ensure that we were moving in the right direction.

Both the raw data and power spectral density were examined for each individual case, as had been done with the earlier data sets. The aberrant behaviors evident in the earlier raw data sets had vanished. Further, the power spectral density displayed a level of consistency far beyond what had been observed earlier, particularly in the 50-70 kHz frequency range that we had been using for band-pass filtering. Figure 40 below illustrates the power spectral densities of all ten cases simultaneously.

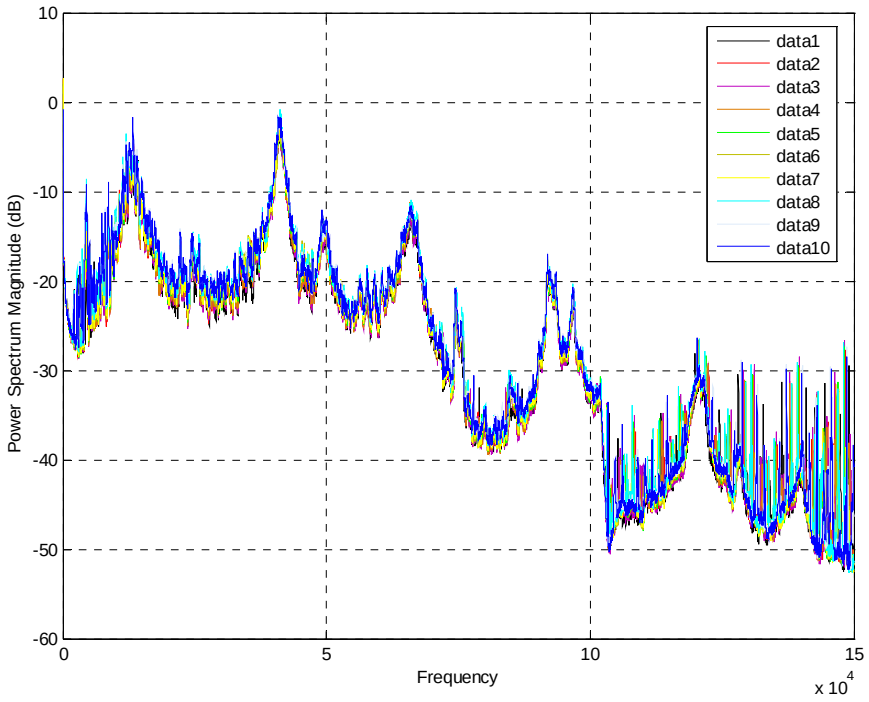

Figure 40. Comparison of power spectral densities for the 10 cases of July, 2005

In examining these ten magnetic-mount cases, we addressed several questions:

- How do the dynamic signatures of the old "good" magnetic-mount cases compare to the new?
- How well does the previous best network perform on the new data?
- What is the best correlation that can be achieved between the dynamic signatures and coal flow when training on the new data?

Any comparison of the dynamic signatures between old and new cases is limited by the reduced set of flow conditions covered by the ten cases of July 2005. A comparison of the flow conditions covered is shown in Figure 41 below.

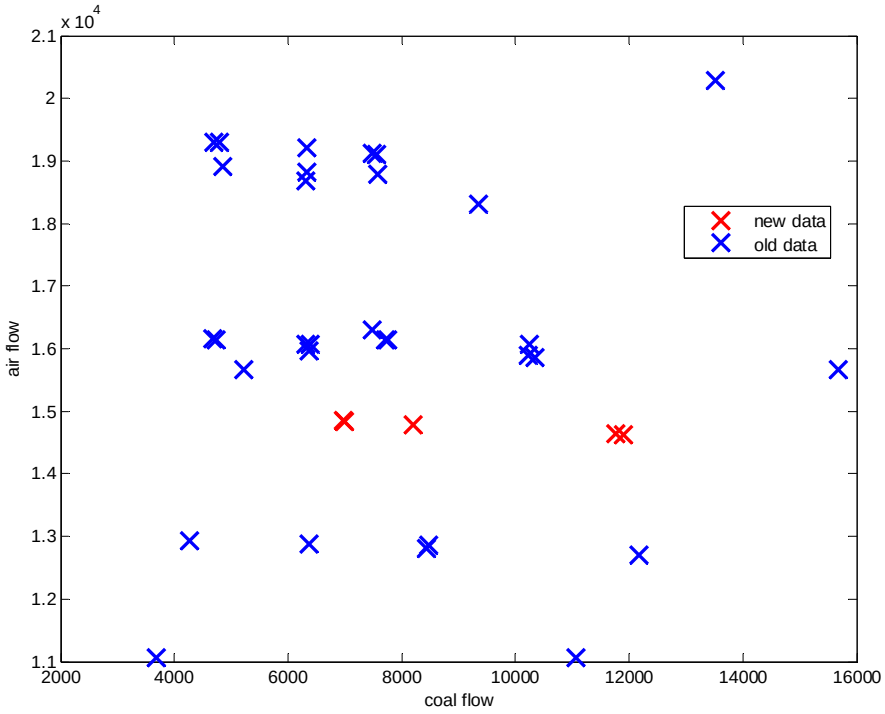

Figure 41. Comparison of Flow Conditions between Data of October 2004 and July 2005

The most dramatic differences between signature quantities in the two sets of data occur for the size measures. As an example, Figure 42 below shows the standard deviation as a function of coal flow. This behavior is representative of the behavior exhibited for all size measures. Namely, the standard deviation is larger but more consistent for July data when compared to the data collected earlier. If this increased magnitude was solely due to a difference in amplification, there would be more variation, for the July data, in the standard deviation for any fixed coal flow level. In other words, there would be less consistency in the July data rather than more. The ten cases of July suggest that standard deviation may increase with increasing coal flow when airflow is held constant. However, ten cases are not a sufficient number to confirm this behavior. Two sets of cases from the original data, each corresponding to approximately constant airflow levels, are circled, one in green and the other in cyan. These cases demonstrate that the original data did not exhibit a similar pattern.

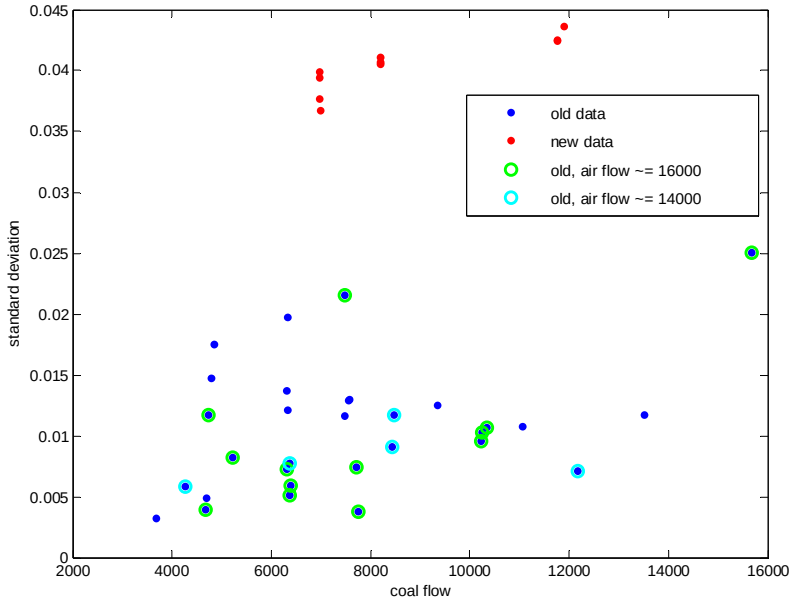

Figure 42. Comparison of standard deviation for original data and data of October 2004 and July 2005

Figure 43 below shows a comparison of a normalized time measure between the two sets of data. Although the two sets exhibit a similar order of magnitude for the signature quantity, the July data tends to be more consistent. The legend for Figure 42 is valid for Figure 43 as well.

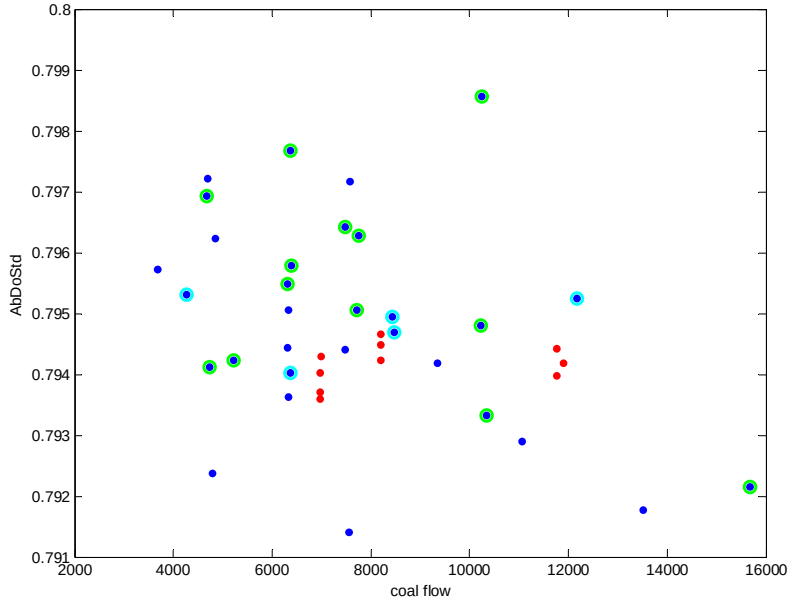

Figure 43. Comparison of AbDoStd for original data and data of October 2004 and July 2005

The dynamic signature that had exhibited highest correlation with coal flow to date had been based on the difference between the signatures of the 5-second windows of data and those of the 1-second windows. The signature quantities employed for this analysis were of the same order of magnitude for the old and new data. However, when the "best" neural net previously developed was applied to the new data set, the prediction results were essentially uncorrelated with actual coal flow, with a correlation coefficient of only 9.71%. The results are displayed in Figure 44 below. This outcome was not really surprising considering all the adjustments that had been made to the instrumentation over the test period. A number of these adjustments would ensure uniformity of future data collections, whether in the lab or the field. However, the implications of such a poor prediction are that consistency in the installation of the sensor mount is critical to the success of a coal-flow prediction instrument.

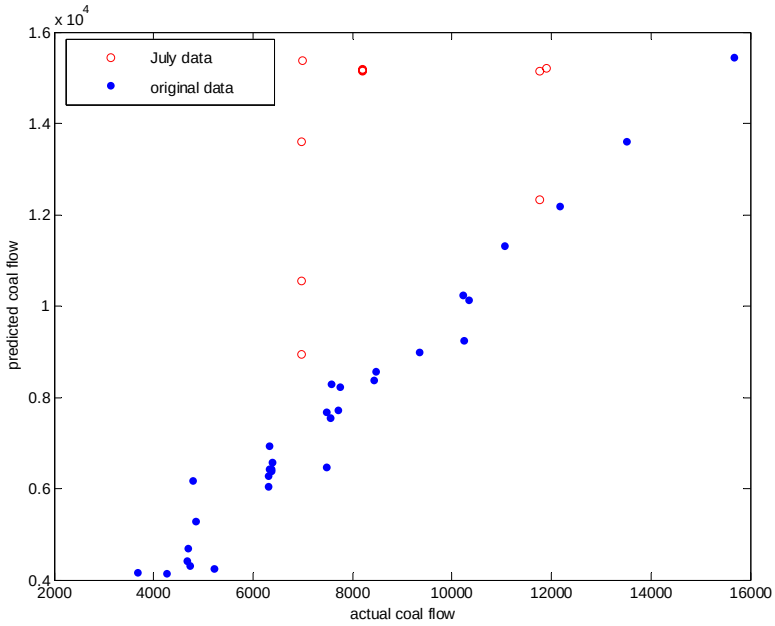

Figure 44. Comparison of coal flow predictions for original data and data of July 2005

When neural nets were trained using only the ten cases of July 2005, the result was quite different. Correlation coefficients above 99% were easily obtained, regardless of whether the original median signatures for the 1-second windows were used, or the median signatures for the 5-second windows, or the difference between the signatures of the 5-second windows and the 1-second windows, or the standard deviation of the signatures of the 1-second windows. In fact, using either the median signatures of the 1-second windows or those of the 5-second signatures, correlation coefficients above 99% could be achieved with just one input and one hidden node. While it is important to remember that this data was limited in quantity and visited only three flow conditions, reducing coal-flow prediction to a sorting problem, these results were highly encouraging. Although any predictions based on such limited data could not achieve statistical

significance, this data did display a variation in individual signature-quantity values that appeared to depend on the flow conditions. The flow predictions for the single-input, single-hidden-node network are shown below in Figure 45.

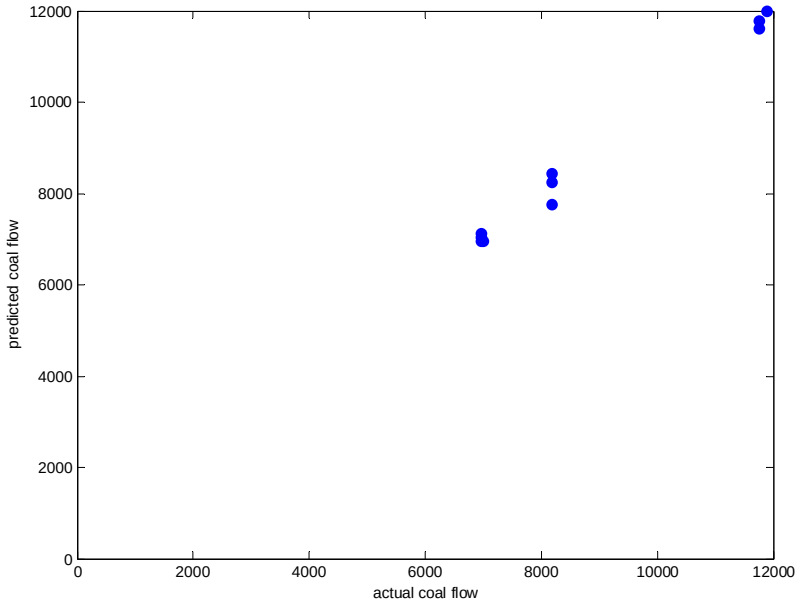

Figure 45. Prediction of coal flow with 1 input and 1 hidden node

6. ANALYSIS OF LAB AND FIELD DATA OF DECEMBER 2005

At last, the noise issues that had been plaguing the laboratory tests had been resolved. A total of 56 data files had been collected in the lab and an additional 73 data files had been collected in the field. Our goal at this stage was as follows:

- Examine all data files and weed out any that showed clear evidence of noise corruption.
- Apply the Dynamical Instruments techniques to both the laboratory data and the field time-series data for the purposes of coal-flow prediction.
- Evaluate the success of this approach and compare the laboratory and field results to see if information gained through the analysis of laboratory data could be used to develop a field instrument.
- Evaluate the possibility of developing a Dynamical Instruments analysis that could identify noisy data in real time.

6.1 Weeding out Noisy Cases

The greatest challenge throughout the program had been that of obtaining clean data. Further, our current state of knowledge was that every time the sensor was relocated, the only way to guarantee continued reliability of the output quality was by trial and error adjustments to the sensor electrical connector. This situation mandated that the analysis begin with the examination of each individual data file to ensure that the prediction algorithm would be based on reliable data. Once the prediction algorithm is developed, the easier problem of automating this weeding-out process could be tackled.

Examination of the laboratory data revealed that, once the connector was appropriately connected, the data was of high quality, similar to the ten files collected in July. The raw data of the 41 test runs in the first piping configuration was consistent and the power spectrum of any one case could be used as a gold standard. The five laboratory test files collected after reconfiguring the pipes, as described in Section 4.4, were also quite good. Any major sources of noise were on the high frequency end and therefore beyond any frequency band that would be passed by the digital filtering process. Figures 46 and 47 below show a representative power spectrum for each of the laboratory piping configurations.

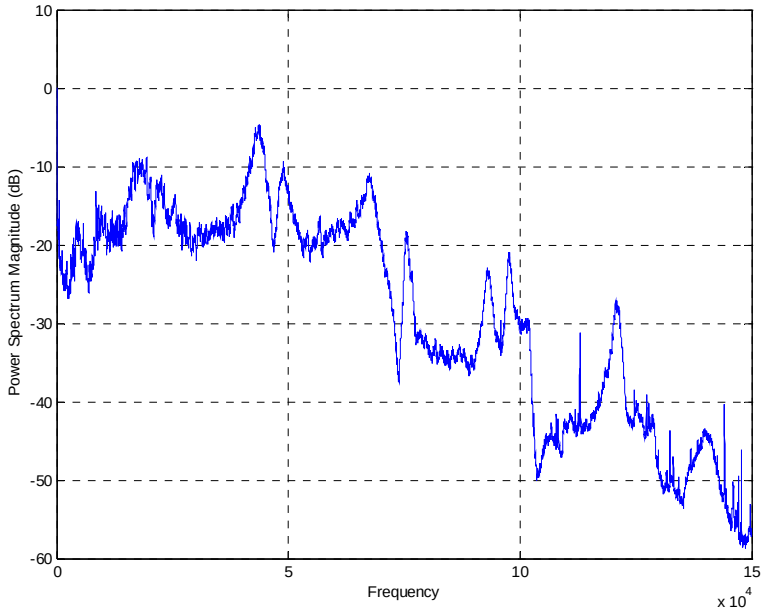

Figure 46. Sample power spectrum for lab data collected on December 1, 2005

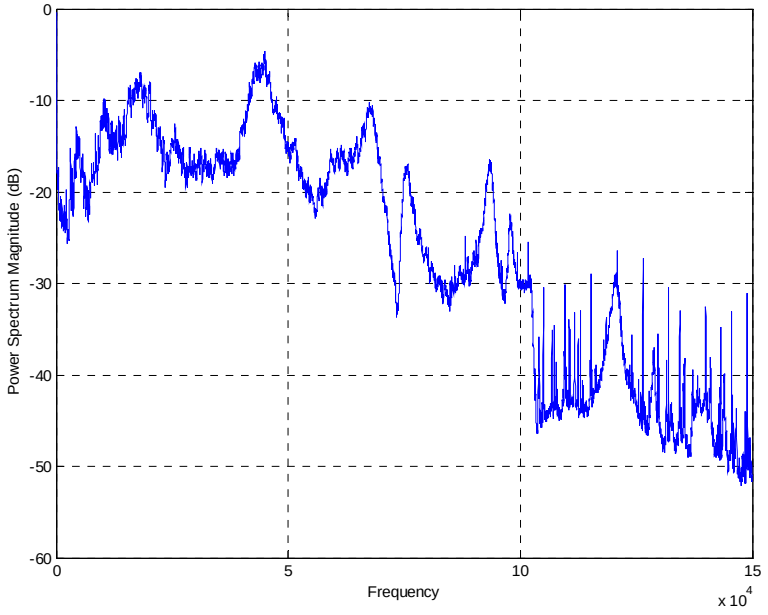

Figure 47. Sample power spectrum following reconfiguration of pipes on December 2, 2005

At first pass, examination of the field data was far less encouraging. Some of the files displayed erratic variations in the amplitude of the raw data, an indication of severe noise. Figure 48 below shows one such case from the data collected at Unit 1, Mill 5. Such output may have resulted from extractive flow sampling being performed concurrently with data collection. In any case, this behavior would obscure the coal flow dynamics, so files displaying such amplitude variation were not included in the analysis. Note that typically the files displayed an amplitude similar to that present at the end of this file, i.e. the smallest of the local amplitudes that are present in this file.

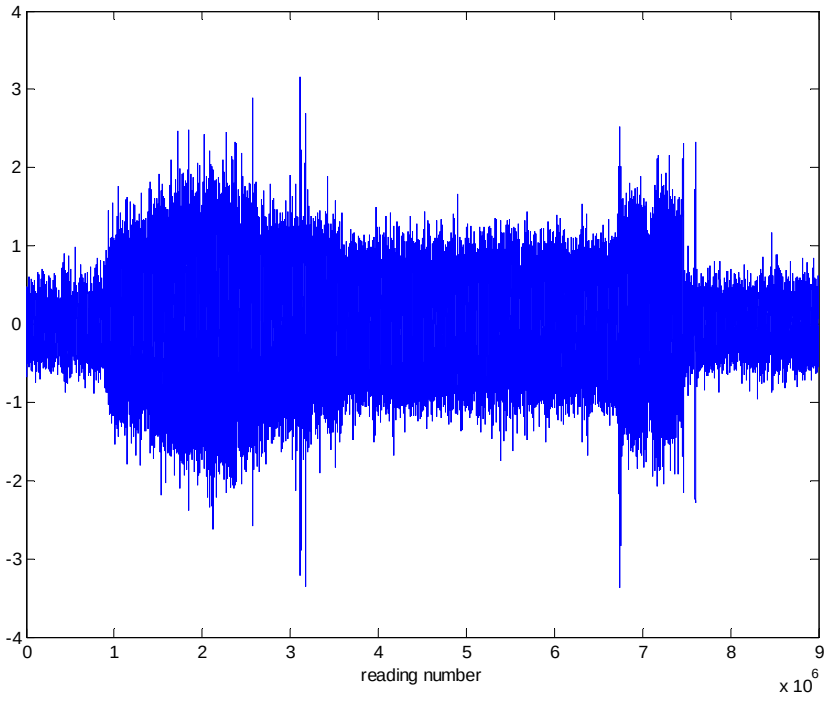

Figure 48. Example of raw data displaying amplitude variation

When the power spectra of the field data were examined, it was found that for some of the files, the power spectrum essentially dropped off as the frequency increased, with barely a hint of the characteristic peak generally present at approximately 40 kHz. These power spectra tended to be unusually spiky, which can be another indication of noise, because fluid flow behaviors cannot produce sharp spectral lines. An example of such behavior is illustrated in Figure 49. Further examination of the entire collection of power spectra for the field data revealed that, while many of the power spectra possessed some degree of spikiness, many appeared cleaner, i.e. less spiky, in the region between 35 and 50 kHz, and displayed a peak around 40 kHz that was far more prominent than that present in the example illustrated in Figure 48. However, during the analysis of earlier data, we had ultimately opted to digitally filter the data to pass a band of 50-70 kHz. Fortunately, earlier in the program, as was described in Section 5.2.4, we had analyzed the effect of changing the band-pass filter range on the resulting coal flow prediction. At that time, we had

concluded that neither the specific frequency range nor the width of this range had much impact on the results, provided that the components of the signal below 20 kHz and above 70 kHz were filtered out. The range 50-70 kHz had been selected earlier by virtue of the fact that, in attempting to predict coal flow, the best correlation coefficient achieved for this range in was marginally better than for the other frequency bands examined. Based on this earlier comparative analysis, we were confident that filtering the data to pass 35-50 kHz, the range where the field data appeared to be less corrupted by noise, would not hinder the analysis in any way. Therefore, the decision was made to filter the data in this way. For consistency, this decision was applied to both the field data and the laboratory data. The latter appeared to be clean throughout the entire 20-70 kHz. At this point, in our efforts to weed out any noisy field cases, we focused exclusively on the power spectra for the field between 35 and 50 kHz.

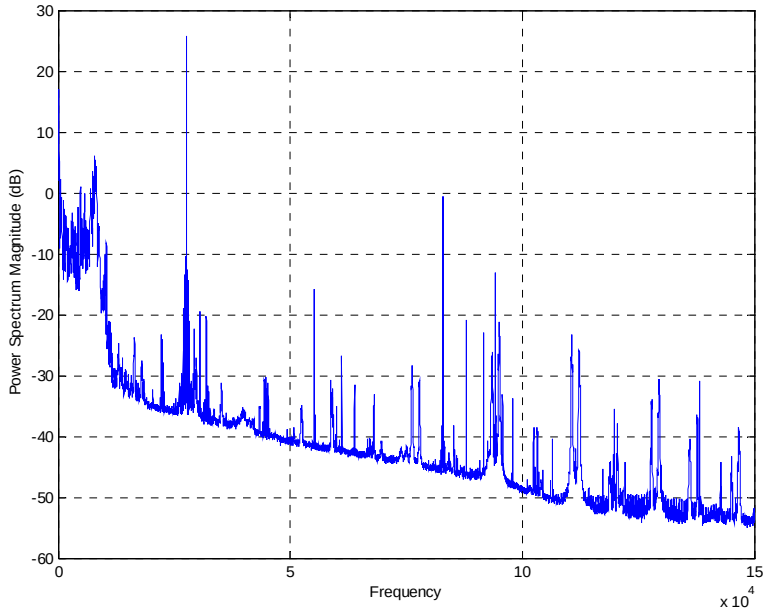

Figure 49. Sample behavior of power spectrum deemed unacceptable for analysis, collected at Unit 1, Mill 4

After eliminating cases with behavior similar to that of Figure 49, we found that some of the other cases displayed power spectra that contained the characteristic peak at 40 KHz but with relatively low power compared to other field cases. Some of these reduced-power cases were quite spiky while others were not. Figures 50 and 51 below exhibit these behaviors, with each graph displaying a comparison of such a power spectrum to that of other data collected at the same site. It should be noted that Figure 50 compares two data files collected at Unit 1, Mill 2, where the sensor was placed upstream of a 90° elbow with a pipe diameter of 13", while Figure 51 compares two data files collected at Unit 2, Mill 5, where the sensor was placed downstream of a 30° elbow with a pipe diameter of 14.5". Therefore, this reduced power is not the result of the piping configuration. It should also be noted that, for all of the field data files,

the power displayed at the 40 kHz peak was significantly lower than the power present at 40 kHz for the files collected in the laboratory, where that power was consistently greater than -5 dB.

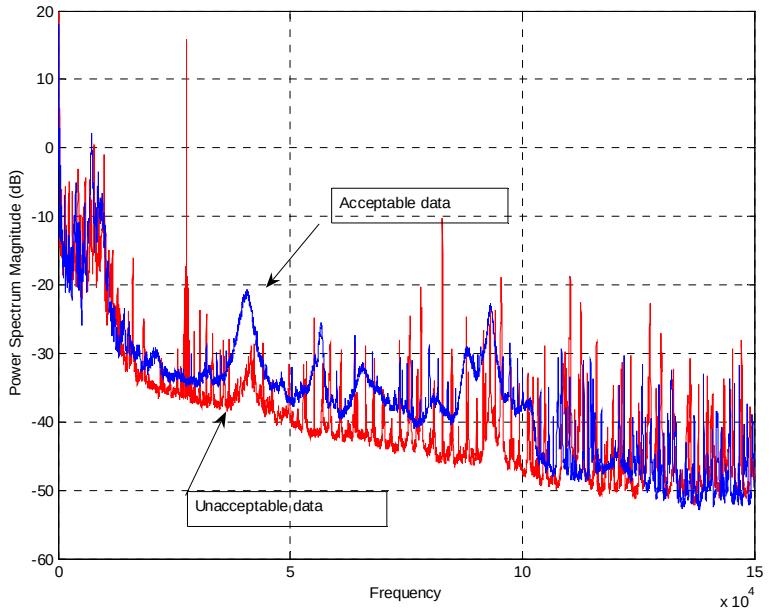

Figure 50. A comparison of two power spectra for data collected at Unit1, Mill 2

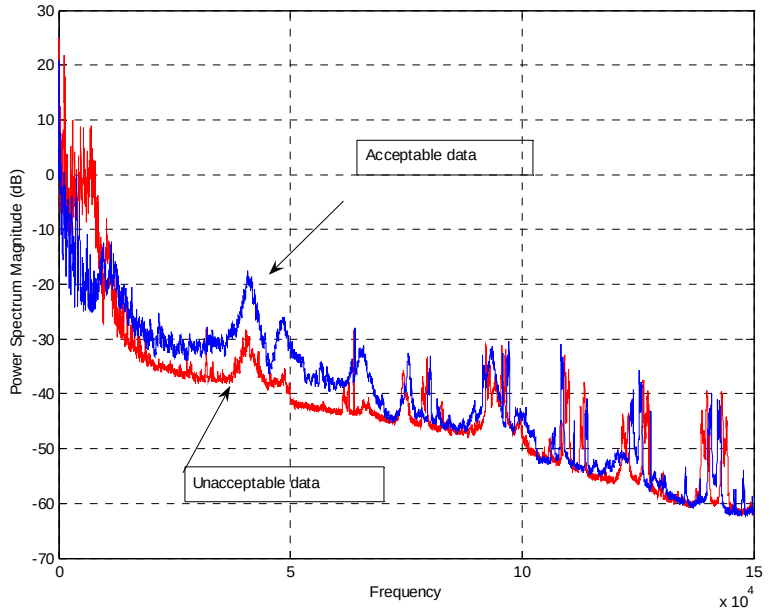

Figure 51. A comparison of two power spectra for data collected at Unit2, Mill 5

Admittedly, any decision to select data on the basis of case-by-case examination of raw data and power spectrum graphs is somewhat subjective. However, in order to retain as much objectivity as possible, we established certain criteria for the acceptable power spectra. There appeared to be essentially two levels of power at 40 kHz. We decided to eliminate those cases where the power was approximately -30 dB or lower and retained those cases where the power was in the -25 to -20 dB range. We also eliminated cases that were particularly spiky in the 35-50 kHz frequency range. However, for the most part, the latter cases also had reduced power at 40 kHz, and so would have been eliminated anyway. Our overall goal, in establishing this procedure, was to balance the fact that more data would enable improved neural net learning and ultimately lead to more accurate predictions, with the fact that noisy data would tend to result in neural-net learning more dependant on memorization of exceptions than on pattern discovery, ultimately hindering generalization to other data sets. We will return to the question of establishing more objective criterion for weeding out noisy data in Section 6.4.

When all was said and done, only 28 of the 73 files of field data collected could be used for the analysis. Each mill was represented in this reduced data set, with the exception of data for Unit 1, Mill 5, which needed to be discarded in its entirety.

6.2 Analysis of Laboratory Data

At this stage, the analysis of laboratory data was restricted to the 56 cases collected during July and December 2005. Too many changes had been made in the instrumentation during the July test effort to consider combining the earlier noisy data sets with the recent, high quality data sets. A preliminary review of the new data revealed no advantage to using the difference between the medians of the 5-second signatures and the medians of the 1-second signatures over simply using the medians of the 1-second signatures. This was anticipated based on the analysis of the 10 cases from July and was quickly confirmed with the additional data. In fact, the median signatures alone performed far better than that difference. Additionally, there was no advantage to using the 5-second signatures over the 1-second signatures. Unlike the earlier data sets, the latest laboratory tests showed no tendency for the variation of the signature-quantity values over time to increase with coal flow. Looking back, two possible explanations had been offered at the time for the higher level of performance of neural nets trained on the difference. The additional data from the latest test sets tended to support the noise-reduction hypothesis as the explanation. With reduced noise in the actual signal, the effect of subtracting the signatures is to simply damp the signal, which provides no benefit whatsoever.

The analysis of the laboratory data had begun with digitally filtering the data to pass the 50-70 kHz frequency band, as had been done earlier in the project. A Dynamical Instruments analysis based on the filtered data performed reasonably well. However, following the arrival of the field data, the decision was made to adjust the frequency range passed to 35-50 kHz, for the reasons specified above. Surprisingly, in contrast to the analysis of the earlier noisy data sets, this adjustment resulted in a clear and immediate improvement in the quality of the results.

The immediate question was how the July signatures compared with the December signatures, and how the December signatures for the five cases obtained after the piping reconfiguration compared with the signatures from the previous day. A very encouraging sign was the fact that measures of size, such as standard deviation, exhibited the trend of increasing monotonically with coal flow, as had been noted in the July data set (see Figure 52 below.) There were variations in the magnitude of these measures that depended on the individual data sets. Some were subtle, as with standard deviation, while others more marked. Still, other signature quantities involving time measures, such as the duration of standard deviation-crossing waves, known as SDur, (graphed in Figure 53) clearly distinguished between the three data sets, while exhibiting levels that appeared to be independent of coal flow. Such signature quantities can actually be quite important to the accuracy of a neural-net prediction as they facilitate the fine-tuning of a correlation by linking an individual data set to the flow conditions that it came from. For example, with the data at hand, there was no obvious way to adjust the somewhat different levels of the standard deviation obtained for the three data sets, having ruled out a simple explanation such as a change in the gain. However, a neural net can learn to recognize and accommodate these differences, a task which is made easier by a signature quantity such as SDur, in this case, which can be used to point the way.

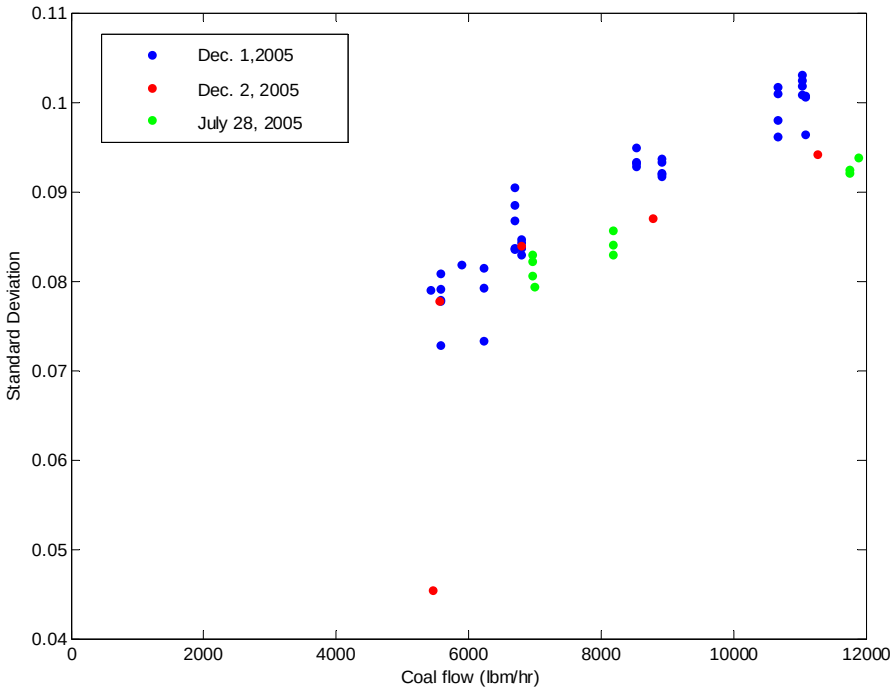

Figure 52. Standard deviation vs. coal flow

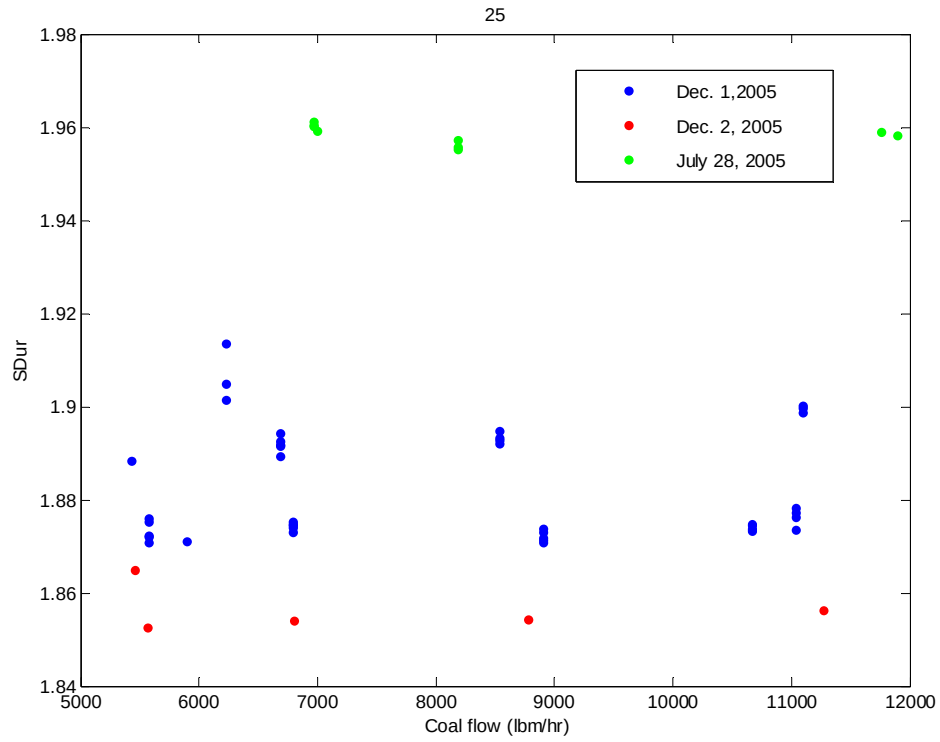

Figure 53. SDur vs. coal flow

Neural nets were trained to predict coal flow using particle swarm optimization applied to a training set consisting of the 46 cases collected in December. Using just two inputs, standard deviation and SDur (graphed in Figures 52 and 53 above), and four hidden nodes, we were able to obtain a correlation coefficient of 99.6% between actual coal flow and predicted coal flow for the December laboratory data. Given that the product of the number of inputs (2) and the number of hidden nodes (4) is much smaller than the number of data files (46), these results are quite good and of the level that had been anticipated at the start of the program. However, applying the network to the laboratory data from July revealed that the results did not generalize well. This was not particularly surprising, as a neural net must learn to identify the types of differences that we knew existed between the data sets. Unfortunately, there simply was not enough data to put together a training set of files that was representative of all data sets and still have enough cases remaining for a test. The results discussed above are illustrated in Figure 54

A second attempt, illustrated in Figure 55 also based on particle swarm optimization, used a training set consisting of 51cases, namely, the July data and the December data, excluding the five cases collected following the piping reconfiguration. Using 4 hidden nodes, a neural net, trained to predict flow based on 3 inputs, achieved a correlation coefficient of 99.3% between actual and predicted coal flow, an accuracy level of $\pm$ 2.9% of full scale. Applying this network to the five December cases that had been excluded showed mixed success, with prediction capabilities being compromised at the high end. Once again, this was not really surprising as many of the signature-quantity values for these five cases were out of the range of values for which the neural net had been trained.

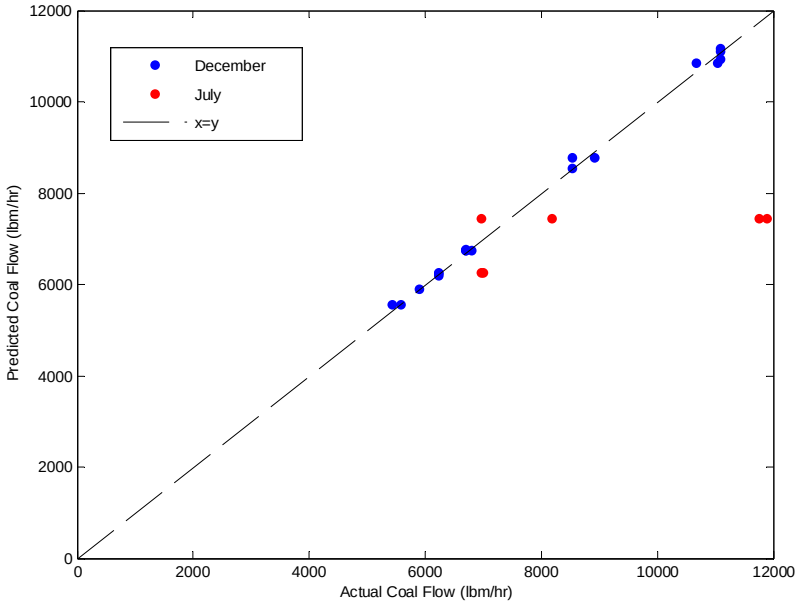

Figure 54 Sample results, accuracy of ±2.2 % of full scale on training set using 2 signature quantities trained on December data sets

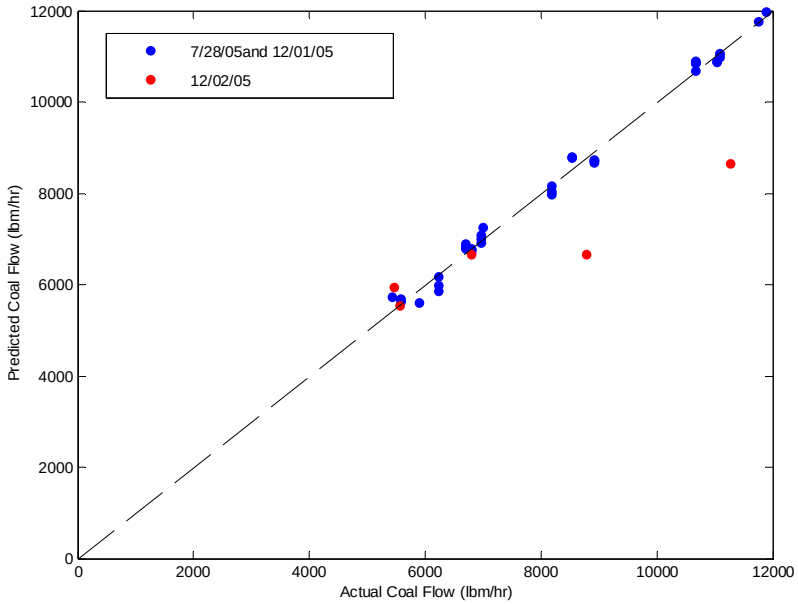

Figure 55 Sample results, accuracy of ± 2.9% of scale on training set using 3 signature quantities trained on July 28, 2005 and December 1, 2005 data sets

In the absence of a suitably large collection of data that would provide for both a *representative* training set and a test set, the best results for which one could reasonably hope would be a neural net that performs well on the training set using a suitably small number of inputs and hidden nodes, as was found here. In that case, the neural net's performance cannot simply be based on memorizing or overtraining. Therefore, the only way that the neural net can attain such a high correlation factor is by learning to identify patterns in the data. While these results cannot supply the level of validation that could be achieved with a larger data set, considering the quantity of available data, they are very promising.

6.3 Joint Analysis of Laboratory and Field Data

The Dynamical Instruments approach is a data driven method that relies on the neural net learning to recognize the flow conditions from the dynamic signature. The immediate problem in expanding the analysis to include the field data as well as the laboratory data was the fact that the flow conditions in the field were quite different from those in the lab, with the airflow in the field both higher and more varied than in the laboratory. Additionally, both pipe diameters used in the field were larger than the diameter of the pipes used in the laboratory (see Figure 56). While one would generally expect that higher airflow levels would be needed for the larger pipes used in the field, there was no direct way to translate the dynamics of the small diameter pipe to the dynamics for the larger pipe diameter. Additionally, there were a number of other differences between the laboratory and field data, as summarized in Table 11. Some would have affected the actual dynamics of the flow while others would have had an impact on the signal being used to identify the flow dynamics. As a result, no neural net trained exclusively on the laboratory data had been given the opportunity to learn the conditions that were visited in the field. Given all of the above, it would have been truly surprising if the networks exhibited in Figures 54 and 55 of Section 6.2 performed well when applied to the data collected in the field. When these networks were applied to the field data, the results were as anticipated. The neural nets developed to date did not generalize well to the field data. In fact, there had been no reason to expect otherwise.

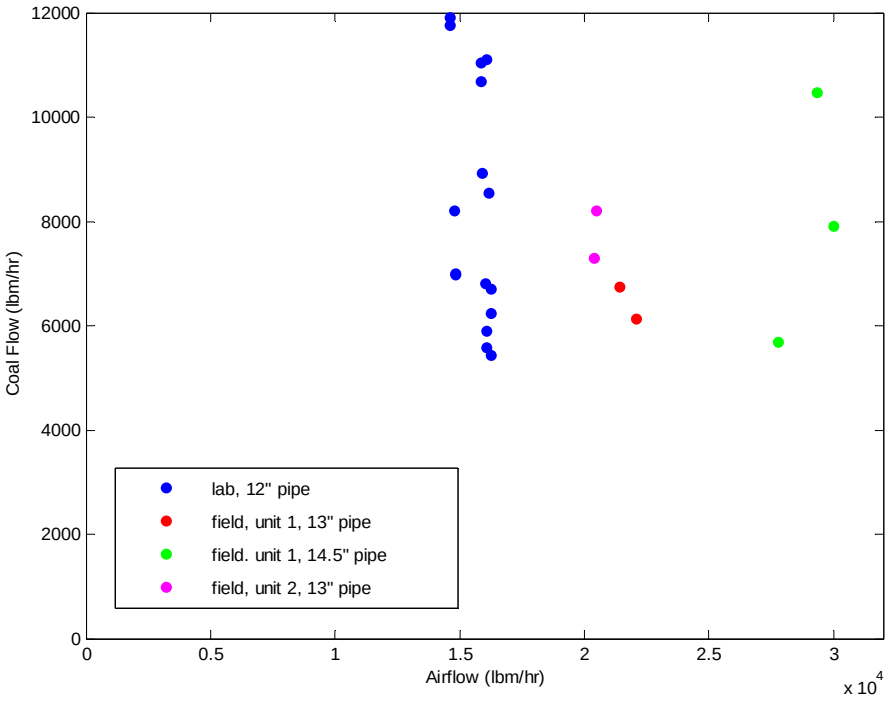

Figure 56. Summary of flow conditions

Table 11. Comparison of Conditions for Laboratory and Field Data

	Laboratory	Field
Pipe Size	12 inch pipe diameter.	13 and 14.5 inch pipe diameters.
Sensor Position	1 diameter downstream of 90° elbow.	1 diameter downstream of 30° elbow or 1 diameter upstream of 90° elbow.
Flow Conditions	Airflow ~ = 16,000 lbm/hr.	Airflow range: approximately 19,000-30,000 lbm/hr.
"Actual" Flow Measures	Obtained by weighing the coal, high level of accuracy.	Obtained through extractive sampling, accuracy ~ $\pm$ 6% of scale.
Signal	Peak in power spectrum at ~ 40 kHz is ~ -5 Db.	Peak in power spectrum at ~ 40 kHz is ~ -25 to -20 Db.
Noise	Extra measures taken to insure clean data, little evidence of noise in data.	Noisy environment, noise apparent in data.

Although the calibration developed from laboratory data could not predict the flow conditions in the field (keeping in mind that the data were collected under different conditions), the laboratory data would be useful for the analysis of the field data. As discussed in Section 6.1, after weeding out the noisy cases of field data, only 28 field cases remained. With such a small quantity of data, an analysis limited to the field data alone would not be optimal. A better alternative would be to train a neural net with a training set that included laboratory and field data. In a best-case scenario, the neural net would not only learn to distinguish between the data sets but also learn the common patterns that existed between the two and adjust those patterns as necessary to the individual data sets. While attempting these goals with a model-based analysis would be quite a challenge, with a neural net, the training process amounts to a search for a best fit to the data. If the number of inputs and hidden nodes are suitably small, the best fit must (from an information-theoretical standpoint) depend on general patterns, whether or not those patterns can be identified explicitly by a model.

Once again, the training process was implemented using particle swarm optimization. The total available data comprised 85 cases:

- 10 from the lab data of July,
- 41 from the lab data of December with the initial piping configuration
- 5 more from the lab in December following the reconfiguration, and
- 28 from the various locations in the field.

Of these, the only cases that could reasonably be withheld without compromising the data representation of the different conditions were the 5 cases collected in the lab following reconfiguration. While 5 cases are not a representative test set, at least there would be some indication of whether the efforts were pointed in the right direction. The best correlation achieved involved a neural net trained using 4 inputs and 4 hidden nodes and resulted in a correlation coefficient of 96.1% between actual and predicted coal flow, providing an instrument accuracy of $\pm$6.6% of full scale. The results of this neural net prediction are shown below in Figure 57. These results gain added strength when one considers that under the best of circumstances, extractive sampling results in an accuracy level on the order of 5%. In less than ideal circumstances, as were reported by Airflow Sciences for the field data set considered here, the accuracy of the extractive sampling is further compromised. In sum, the accuracy level of $\pm$6.6% is about as good a result as one could hope to achieve based on extractive sampling. Note that this neural net also performs well for our limited test set of 5 cases, whose predictions are also included in Figure 57. The correlation coefficient for actual and predicted coal flow for the set of 5 test cases is comparable at 96.8%.

Further, limiting the number of inputs to 3, and still using 4 hidden nodes, had only a limited impact on the correlation coefficient reducing the best fit result to 95.8%. Even a neural net with 3 inputs and 3 hidden nodes achieved a correlation coefficient of 94.4%.

These results suggest that an instrument could be developed to accommodate an expanded range of conditions and pipe-layouts. As the availability of data expands so that the data set visits a wide range of flow and operating conditions, the level of accuracy should improve.

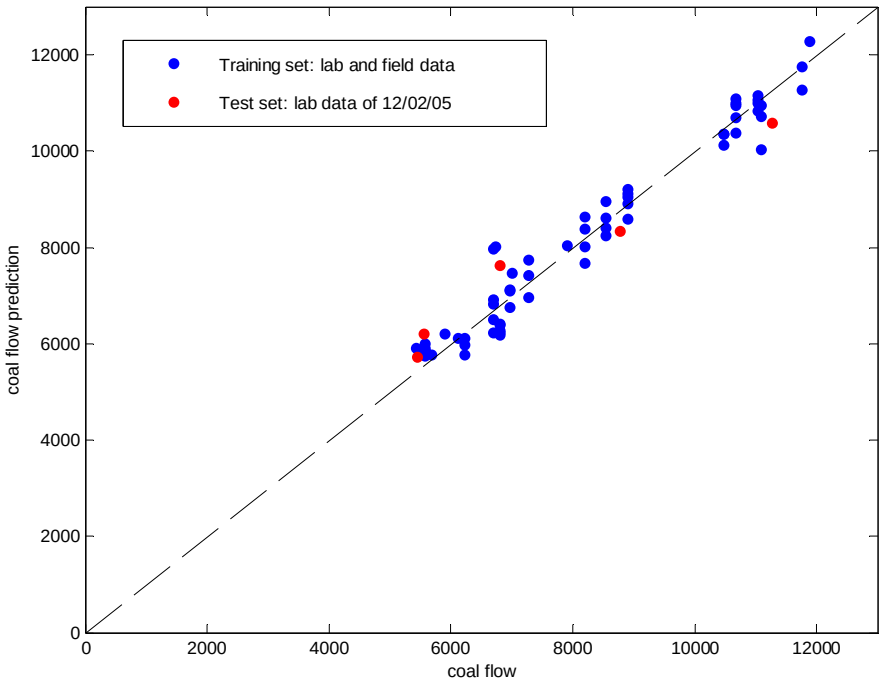

Figure 57. Coal-flow prediction for combined data set of lab and field data

6.4 Development of a Noise-Recognition Algorithm

The results obtained using the 35-50 KHz frequency range indicate that a more robust instrument can be developed through the use of a lower-frequency sensor than has been used throughout this program. Nonetheless, a coal plant is a noisy environment and it is neither realistic nor practical to expect to eliminate every source of noise. Further, unlike current instrumentation, the instrument under development is intended to monitor the coal flow without interruption of normal plant operations and with minimal human supervision. In the long run, the process of weeding out noisy data cannot be accomplished by personnel in the field arduously examining raw data and power spectrum plots as has been done during the development stage, as we discussed in Section 6.2. Rather, noise recognition must be automated so that the instrument can learn to recognize "noise signatures" and warn the user when predictive capabilities of the instrument are compromised.

A key aspect of the Dynamical Instruments approach is that an estimate or prediction of the coal flow level can be made in real time based on only 30 seconds of data. Data could be collected on a continuous basis, if desired, so that in the space of minutes, a sequence of estimates for the flow conditions would be obtained. Averaging consecutive predictions would serve to increase

the accuracy as the noise would be effectively averaged out. Even with the small quantity of data available here, the effects of this averaging process are noticeable. Many of the data files collected by Airflow Sciences were collected consecutively, spaced by intervals of 10-20 minutes in time. During these time periods, the flow conditions were nominally constant and the "actual" coal flow assigned to such cases was obtained by a single measurement, calculated either by weighing the coal in the lab or by extractive sampling in the field. The number of consecutive files available to us ranged from 1-7 files, with the median number of files equal to four and the mean nearly four as well. We identified those flow conditions for which a minimum of four relatively noise-free files were collected during a single test run. We then averaged the predictions obtained for a single flow condition. Figure 58 below shows that, even after averaging just 4-7 predictions, we see evidence of improved accuracy, as the average of the predictions clusters more closely to the line than many of the individual data points.

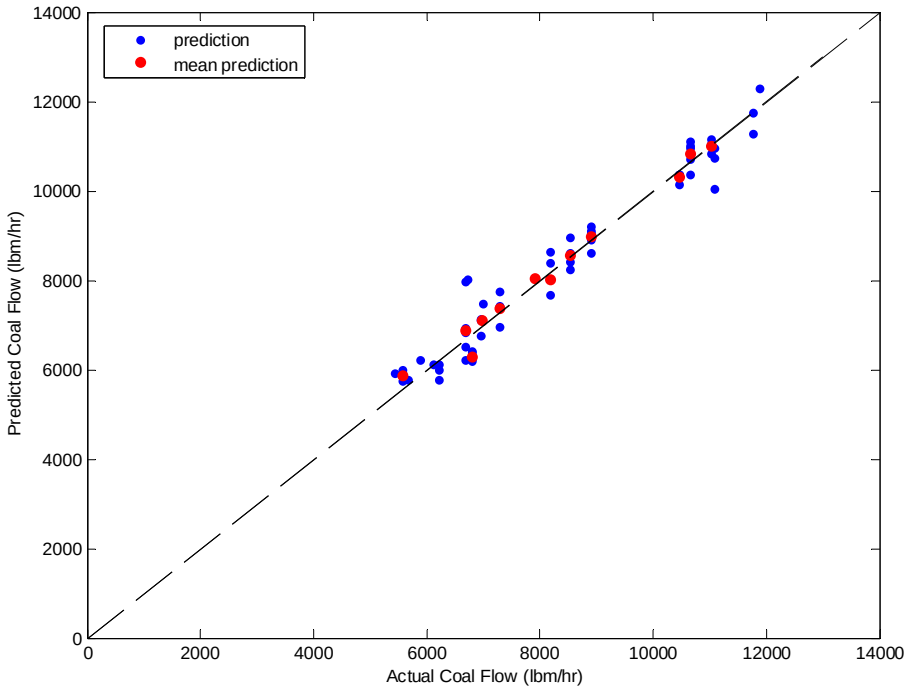

Figure 58. Mean predicted coal flow over 4-7 files compared to individual predictions

The lesson is not that noise will not have an impact on the accuracy of the prediction, since we have seen repeatedly during this program that a poor signal-to-noise ratio is highly problematic. However, once a reasonable signal has been obtained, averaging the prediction will tend to be forgiving of at least low-level noise. This statement is not unique to a Dynamical Instruments analysis, but rather depends on basic statistical principles. However, the statement gains strength with the Dynamical Instruments approach, due to the fact that many consecutive predictions can be determined within a short space of time, and in particular during a time period where the flow conditions are stable.

While low-level noise can be damped through averaging, high-level noise must be identified so that inaccurate predictions are automatically discarded. In order to approach the problem of noise identification, we considered the entire collection of data that had been obtained in the field during December 2005, noisy cases and all, along with the recent laboratory data from July and December 2005 obtained after all changes to the instrumentation had been finalized. We expected that the neural nets trained on the selected data would not perform well when applied to the identifiably noisy data and confirmed that this was in fact the case, as is illustrated in Figure 59 below. Although some of the predictions for the noisy data do appear to be reasonably accurate, it is difficult to conclude that there is any significance in that observation, as the entire set of files corresponding to the noisy data is mapped to essentially two levels. Given that the noisy data covers a range of flow conditions, the more probable explanation is that on occasion one of those two levels "happens" to actually be the correct.

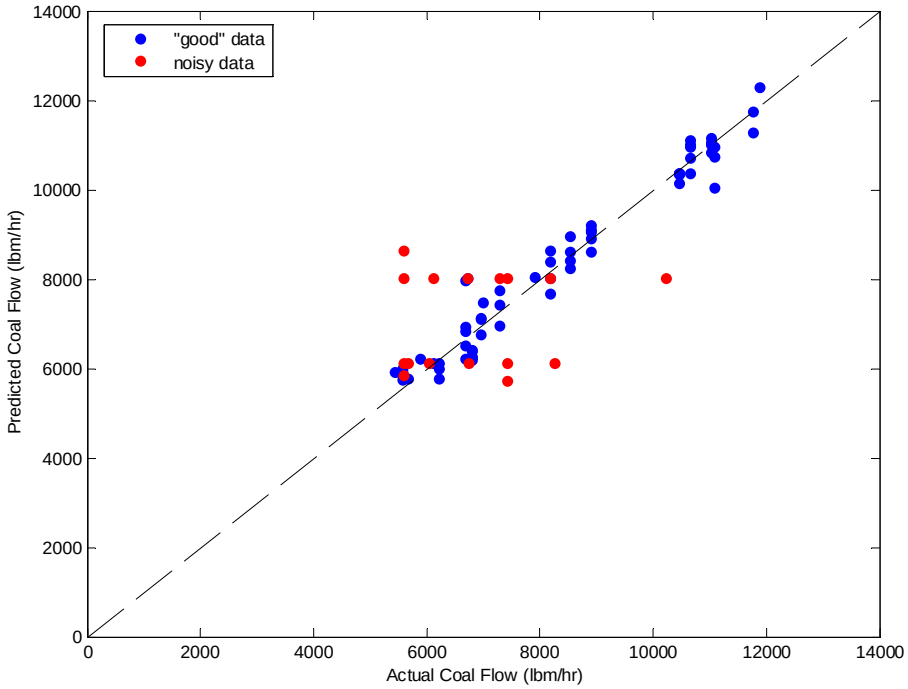

Figure 59. Coal flow predictions for noisy data as compared to those for clean data

Nonetheless, the development of a noise identification algorithm should be based solely on quantifiable observations. Rather than try to identify those cases that we had classified as noisy based on a process that was admittedly somewhat subjective, we attempted to identify those cases for which the neural net performed poorly. This meant that some of the cases that were considered to be noisy would actually be considered "good" cases for the present. The entire collection of data was partitioned into two subsets, those for which the predicted coal flow

differed from the actual coal flow by less than two standard errors of measurement, i.e. 6.6% of scale or 785 lbm/hr (as determined in Section 6.2) and those for which the difference between predicted coal flow and actual coal flow was above that threshold. In fact, there was a clean break between the error levels of the two subsets and the "good" cases were mostly well below that threshold, as is illustrated in Figure 60. Despite our view that the occasional accurate prediction obtained for the noisy data occurred by chance, this approach ensured that the analysis focused directly on the question of whether we could predict those cases for which the coal-flow prediction would be inaccurate and depended, in this way, solely on measured quantities.

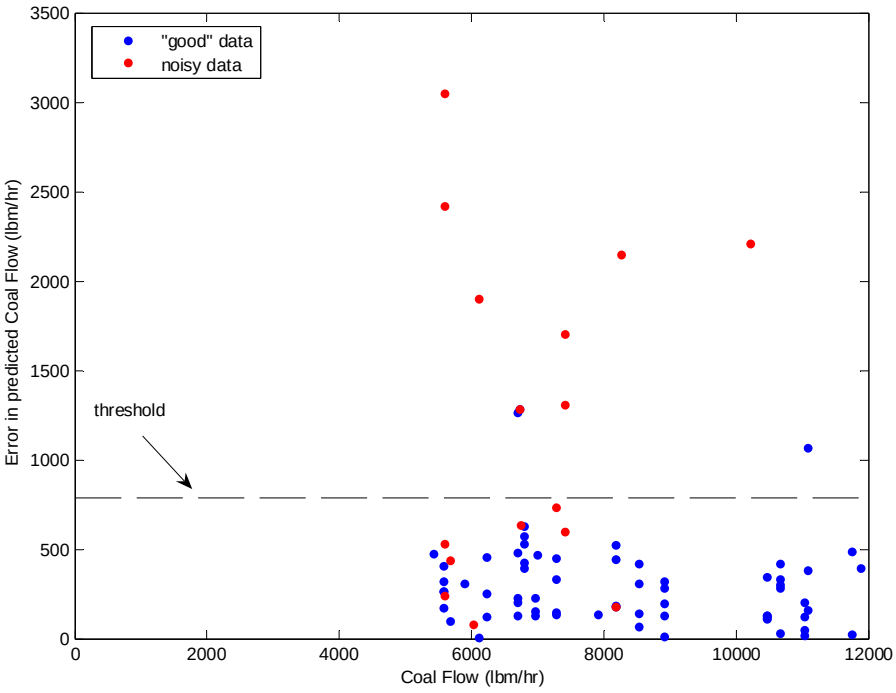

Figure 60. Error levels in coal flow prediction for both clean and noisy data

In all, of the 45 files that had been rejected as noisy by way of the process described in Section 6.2, 27 resulted in predictions that were greater than the two standard errors of measurement and only 3 of the files selected as non-noisy resulted in such poor predictions. (Some of the readings were sufficiently similar to one another that they cannot be distinguished from one another in Figure 60.)

The initial approach to noise identification was similar to that taken for predicting coal flow. Particle swarm optimization was applied to the field cases to train a neural net that could successfully predict the Boolean output, 0 for the "good" cases and 1 for the noisy cases, where in this case "good" was defined by a prediction error less than two standard errors of measurement, as discussed above. As it turned out, constructing such a neural net was

accomplished easily. Using three inputs and three hidden nodes, a neural net was trained that produced a correlation coefficient of 94% for actual accuracy classification vs. predicted accuracy classification. Such a correlation coefficient actually produces an extremely accurate Boolean classification, as correct classification does not require the neural net to produce exactly 0 for the "good" cases or exactly 1 for the noisy ones. All that is needed is for the outputs to be "close enough" to their true Boolean values so that the classification can be deduced. The results of this neural net prediction are illustrated below in Figure 61. In fact, out of 73 field cases, only one would be misclassified by this neural net.

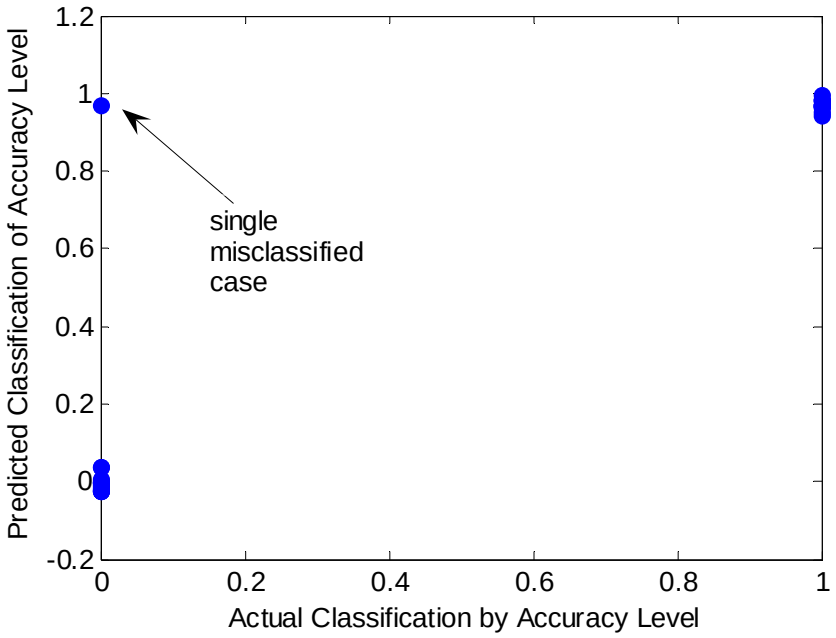

Figure 61. Neural net classification of field data as either "good"or noisy

While this was an excellent starting point, we have found that in contrast to algorithms that predict flow rates (numbers in a continuum), classification algorithms tend to be more robust if they do not rely on a neural net. Rather, we use the neural net to locate a collection of signatures quantities that will enable the desired classification. Then, using other techniques, such as cluster analysis, an algorithm that functions less as a black box can be developed. The importance of this approach for discrete classification lies in the fact that, when the training set is limited and the output is Boolean, it is particularly easy to overtrain a neural net so that it effectively memorizes any anomalous cases. As with a prediction along a continuum, the risks of overtraining for a discrete prediction decrease as size of the training set increases. In other words, with enough data, either technique is likely to be robust.

Prior to proceeding with a cluster analysis, the above neural net was applied to the lab data. Recall that this neural net had been trained exclusively on field data. Only two cases in the lab data had resulted in coal flow predictions whose errors were greater than two standard errors of measurement. The neural net above predicted that all of the December data was "good," missing the two cases that resulted in a larger prediction error, and erroneously predicted that all ten July laboratory cases were noisy. This outcome is best understood in the context of cluster analysis.

Figure 62 below gives an indication of the potential success of cluster analysis for noise identification. The graph contains a scatter plot of the 3-dimensional signature selected during particle swarm optimization of the neural net. The field cases that are colored green are those selected for use in the original analysis, based on the original selection process described in Section 6.2. The points in red represent the field data that was rejected during that selection process. The lab data, not included in the neural-net training set is colored blue, the lighter blue corresponding to the 10 good cases from July. Those points that would be classified as noisy based on the level of accuracy achieved in the coal flow prediction are circled. Focusing first on field data, the clusters of red and green points demonstrate that the visual classification applied in Section 6.2 would be nearly duplicated by a cluster analysis of dynamic signatures. Of the 73 field cases, it appears that approximately five might change their classification through a clustering routine. Considering the subjective nature of the original procedure used to identify noise and the fuzzy nature of the acceptance boundary, this is extremely encouraging. The field cases whose coal flow was well-predicted despite the original identification of noise (red, but not circled in black) would likely be classified as noisy in a cluster analysis, reinforcing our earlier conjecture, based on Figure 59, that the accuracy did not result from true prediction capabilities for this particular set. Three field cases that had originally been selected by us as "good" appear based on Figure 62 to have been misclassified. One of those resulted in a poor prediction of coal flow. The other two seem to lie near the cases that we suspect were predicted well by chance.

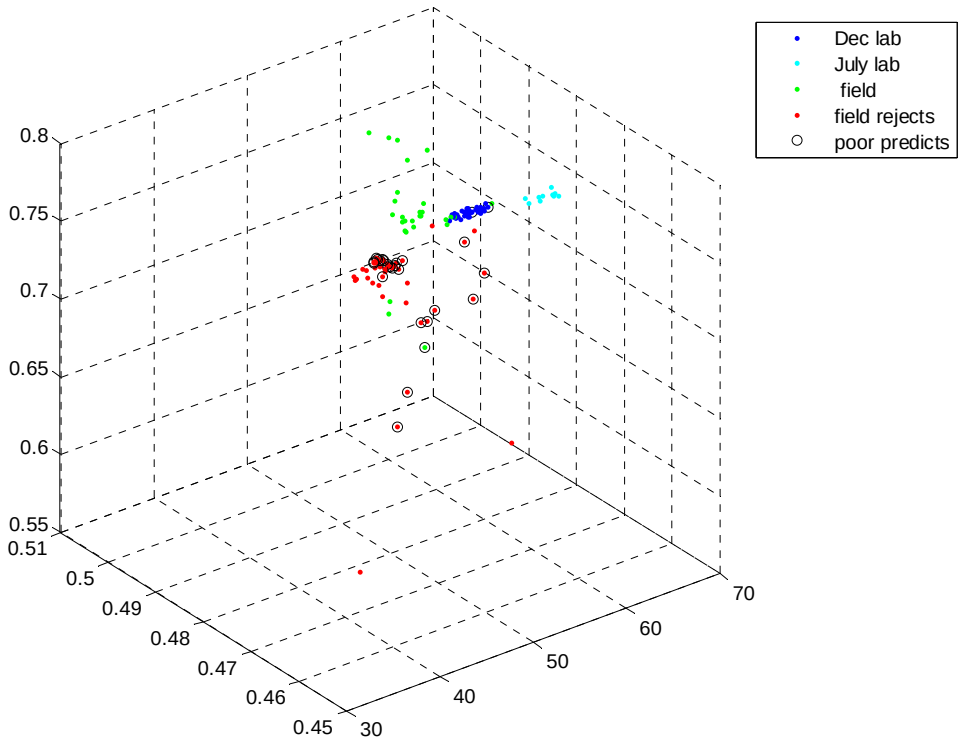

Figure 62. Scatter plot representing 3-dimensional dynamic signatures used to identify noise

For the laboratory data, we see that a cluster analysis based on these signature quantities would do well overall, but again miss the two cases that were misclassified by the neural net. As for the 10 July cases, it is clear that they did in fact fall in a region of space not visited by the training set of the neural net, which explains the poor performance of the neural net on this set. Visually, these cases appear closer to the "good" set, but a more dependable analysis, whether depending on clusters or neural nets, would include sufficient data to cover all conditions during the learning process. In fact, the location of the July data points serves to illustrate a fact that has been emphasized throughout this report. Specifically, during the training phase, the neural net learns the classification scheme. However, only regions that are visited by the training set can be expected to be well-predicted by the resulting neural net.

A key point in addressing the feasibility of designing an effective noise-identification algorithm for the instrument under development is that the algorithm must perform well with regards to eliminating noisy data, but need not perform nearly as well in evaluation of "good" data. In other words, false positives are far less costly than false negatives, where positive indicates classification as noisy. Data can be collected on a near-continuous basis, so if some good data is rejected, there will be no shortage of good data to serve the same purpose. Even the occasional false negative, i.e. acceptance of noisy data as clean, will not be particularly problematic after the implementation of the averaging procedure discussed at the start of this section. In short, based

on the results of this section, there is every reason to believe that a robust, real-time instrument can be developed that will perform the two-fold process of eliminating noisy data and predicting coal flow on the remaining data.

7. COMMERCIALIZATION PLAN

Any discussion of the commercial development of an instrument based on the Dynamical Instruments technique must take into account its inherently empirical nature. Since the technique works by learning to relate observed system dynamics to the flow conditions that produced them, it is only as applicable, useful, or accurate as the database on which the analysis is developed. Given this basic starting point, plus the very broad range of conditions encountered in various power plants, two essentially opposite approaches to commercialization can be identified:

- Develop special-purpose instruments for individual power plants. This would involve installing an instrument and performing an intensive in-place calibration. Through experience with data collected at a number of such installations, the amount of in-place calibration required would gradually decrease, hopefully ultimately to be eliminated.
- Develop a general-purpose instrument that can be applied in a broad range of power plants. This would require developing an extensive database of flow data for a broad range of conditions, including varying pipe sizes, piping configurations and flow ranges.

In principle, either of these approaches should be achievable. The first approach could be pursued with any interested utility, requiring only flexibility on the part of the utility operators to permit testing the plant over the full range of coal and air flows likely to be encountered in operation. The calibration data would be obtained using extractive sampling, which probably represents the largest component of the instrument cost. In return, the utility would be essentially assured of obtaining an on-line instrument that is at least as accurate as the extractive sampling that would be used to calibrate it.

On the other hand, coal plant balancing efforts are the stock in trade for Airflow Sciences and several other companies, offering the opportunity to obtain an extremely large quantity of calibration data for a broad variety of installation conditions. The marginal labor required to collect Dynamical Instruments data in the course of a plant balancing effort would be quite minimal, especially with a well-designed portable instrument package. Our experience in this program indicates that a magnetic sensor mount provides a handy and accurate means of sampling the flow dynamics, so developing such a portable instrument appears completely straightforward. As the database grows with each power plant balancing effort, the calibration analysis would be performed on an ongoing basis, permitting the calibration of the portable instrument system to be upgraded over time. Thus, the portable instrument would become applicable to more and more operating conditions and, we believe, more accurate as well. Thus, the personnel performing the plant balancing efforts could gradually gain confidence in the portable instrument, so that ultimately the extractive sampling involved in plant balancing would be obviated.

Between these two approaches to the market, the first appears more attractive at the outset, because it would be expected to produce an income stream that would, at least partly, offset the cost of instrument development. To pursue the latter would require deferring the income stream, and in fact would involve some cost to cover the data collection in plant balancing efforts. By

the same token, the latter approach would also be expected to produce a generally applicable instrument in a shorter period of time, increasing the market and likely profit margin.

Ultimately, the market would decide which approach was more practical. Both approaches were discussed in meetings with EPRI's Pulverizer Interest Group, the EPRI activity that funded the Coal Flow Test Facility. Although it was agreed that the special-purpose instrument approach would be workable for a sufficiently interested utility, none could think of a utility that would be willing to undergo the intensive calibration process. Considering that this group of utilities is among the strongest boosters of coal flow measurement in the US power industry, this is a strong statement about the efficacy of this approach.

The group was more interested in the development of the general-purpose instrument, even if it delayed the instrument's commercial availability. Presumably, there was an element of letting someone else pay for the detailed calibration effort in this viewpoint, but their interest in the technique appears genuine: this group is more likely to purchase such systems for their plants than most other utilities. Many found it particularly attractive to consider that an instrument could be installed with no need for in-place calibration, although this was viewed somewhat skeptically: the currently available commercial instruments typically require some baseline calibration, and the utilities simply do not believe that any is accurate "out of the box". Thus, in marketing an instrument, being open to the challenge of an in-place calibration or demonstration would be viewed positively by most utilities. Success of such demonstrations would improve the customer perception of the instrument's quality over time, solidifying the market position of the instrument.

The process of developing a commercial instrument in this approach would involve the following steps:

- Identify available accelerometers that would be suitable for use in a commercial instrument. This involves considerations of performance characteristics, packaging, and cost, with the first two criteria weighted more heavily. Since the part cost of the Dynamical Instruments coal flowmeter should be small compared with that of other commercial offerings, sensor cost is relatively unimportant. The ideal outcome of this effort would be the identification of essentially interchangeable sensors from multiple vendors, so that the failure of any single vendor to perform would not limit availability.
- Develop a portable instrument package that can be used both as a field data collection platform and prototype operational instrument. In the program effort, we found that the package used in the lab and field testing was workable, but inconvenient. The computer was bulky, and the signal conditioning components needed to be connected on-site. For a practical field test program, all components should be arranged in a single, compact and convenient package with all wiring pre-connected. This could be accomplished with the use of a laptop computer, PC-card based data acquisition, and an external signal conditioning module. These can all be housed in a suitcase-sized package, so that setting up for a test would be quick and simple. For software, it would be desirable to use a flexible data acquisition and analysis program that permits collecting raw data, calculating dynamic signatures, and implementing and updating instrument calibrations. It appears that the Matlab programming environment, from the MathWorks, would be a

good choice for this, especially since we perform all of our instrument analysis with Matlab. Overall, the development of this instrument package appears straightforward.

- If possible, perform instrument shake-out testing using the Coal Flow Test Facility to develop baseline data that can be compared to the data collected in the program effort, and to confirm noise free operation. This testing could also compare the dynamics found with magnetic and permanent mounts (now with a transducer immune to the grounding problem of the previous transducer). From this testing, and/or the data collected in the program effort, develop and implement a preliminary instrument calibration, including a noise-identification algorithm to confirm to the operator the quality of the data being collected.
- Collect large quantities of data in plant balancing efforts. This effort would become a routine part of the plant balancing effort, providing substantial quantities of data for new conditions at little cost. The prediction provided by the current instrument calibration would hopefully assist in the plant balancing effort, with this assistance hopefully improving as the effort continues.
- After each plant balancing effort, the data would be downloaded to permanent storage media and sent to Foster-Miller for analysis. This analysis would include comparing the new data to the existing instrument calibration and the creation of a new calibration that includes the new data in its basis. The resulting calibration would then be sent to implement in the field test system.

As the calibration of the field test system expands to include a broad range of plant configurations and operating conditions, the potential rapidly arises to sell instruments that can be installed permanently in operating plants. In many cases, the plants that contract for balancing efforts would be interested in permanently installed instruments if an accurate, inexpensive system were available, so the management of each plant will be approached about their interest in purchasing a system. The sale, installation, and checkout of an instrument system would be pursued wherever this interest is found.

The recommended route to commercialization thus appears to be through continued support by DoE and EPRI. This provides an opportunity to develop both portable and permanently installed instrument systems. With EPRI support and Airflow Sciences involvement, the portable version would add to the software data base during their plant balancing work with the aim of providing them with a much quicker and recordable balancing process. With DoE support, there would be opportunities to incorporate prototype instruments in larger DoE supported programs in advanced powerplant and gasification projects, where the Dynamical Instrument would be used in control loops and diagnostic systems.

8. CONCLUSIONS

At the conclusion of the program effort, we are ready to move forward and develop a package that can serve as a prototype commercial instrument. The key challenge throughout the program has been elimination of noise contamination during data collection. However, the source of the noise has now been linked to electrical grounding issues specific to the particular high-frequency transducer used in this program. Fortunately, we have demonstrated, as discussed in Section 6, that a lower-frequency transducer will not limit the analysis in any way. Substituting a lower-frequency transducer will do more than solve the grounding problems. Many lower-frequency transducers are both less costly and more robust. Both these qualities will serve the instrument well. Further, as there are many more low-frequency transducers from which to choose, there will be more options available in selecting a transducer that is well suited to the final product.

The only real challenge in analyzing the data during this program has been the development of techniques to compensate for the noise that was present in the data. Once the noise was reduced to an acceptable level, correlation results between predicted and actual coal flow were as good as had been expected at the start of the program. However, the data that has been collected during this program needs to be supplemented in order to fully develop an instrument that can be used in the field.

It is our view that the best way forward is to develop a prototype that can be readily installed temporarily in interested plants and to then perform calibration based on extractive sampling. As the data collection expands to include a wide range of flow conditions, pipe sizes and piping configurations, the calibration will be improved and the range of applicability will be expanded. We will continue to gain additional experience by monitoring a greater number of plants until the algorithm has been fully developed and fine-tuned to the point where no in-plant calibration is needed.

9. REFERENCES

1. Hill, W.S., “Method and System for Analyzing a Two-Phase Flow,” US Patent 5,600,073, February 1997.
2. Hill, W.S., “Method and System for Analyzing a Two-Phase Flow,” US Patent 5,714,691, February 1998.
3. Hill, W.S. and Barck, B.N., “Flow Analysis System and Method,” US Patent 5,741,980, April 1998.
4. Hill, W.S., and Heirtzler, F.J., “Object Classification and Identification System,” US Patent 6,014,520, January 2000.
5. Heirtzler, F.J., and Hill, W.S., “Tire Defect Detection System and Method,” US Patent 6,381,547, April 2002.
6. Bliss, L.A., “Selecting Artificial Neural Network Inputs Using Particle Swarm Optimization,” Doctoral Dissertation, School of Computer Science and Information Systems, Pace University, June 2003.

APPENDIX: DESCRIPTION OF DYNAMICAL INSTRUMENTS SIGNATURE QUANTITIES

The following discussion describes the several classes of measures used to characterize the evolution of a dynamic signature, followed by a description of each of the specific signature quantities. In the following discussion, terms in parentheses refer to the short-hand notation used to denote the different measures.

A.1 Characterization of Harmonic Power

As mentioned above, two separate sets of signature quantities are calculated for each dataset, one calculated for the raw data and the other for a new vector that characterizes the passage of harmonic power in the raw data. The single most common measure of harmonic power for dynamic data is the standard deviation of the data. It is more correct to say that the variance of the signal (the square of standard deviation) characterizes power, but the standard deviation is proportional to the amplitude of the signal. To characterize the passage of harmonic power, a new vector of moving standard deviations of the raw data is created (illustrated in Figure A1). This vector is constructed by calculating the standard deviation of a block of 20 successive readings. The raw data block is then updated by dropping the "oldest" value from one end of the block and adding a "new" value to the other end of the block, and the process repeated until the moving standard deviation is computed for the entire raw data vector. This produces a standard deviation vector that is 19 readings shorter than the original raw data vector. This procedure can be performed very quickly when implemented in an efficient calculation. In shorthand notation, signature quantities calculated for this vector are indicated by prefixing the letters **sd** to the signature quantity name. The signature quantities calculated on the raw data lack this prefix.

Whenever one devises a given measure for dynamic data, one can expect to find an application where the measure falls short. A limitation of the moving standard deviation vector is that it has a single time scale (20 readings) for characterizing the power in the signal. For very rapidly varying signals, such a short window is quite reasonable, because it can capture several periods of the signal variation. For a slowly varying signal, such a window could prove to be a fraction of a period, and thus will under-report the amplitude of the variation. Numerical experiments have demonstrated that, for data that are rapidly varying, the results of Dynamical Instruments analysis are insensitive to the standard deviation window size: although the signature quantity values change with window size, the correlation to the system conditions of interest proves to be equivalent. Consequently, the data sampling rate should be selected to ensure that the signal varies significantly, on average, from reading to reading. In particular, this means that the data should not be greatly over-sampled, producing a very smooth variation from one reading to the next. By and large, a signal that is sampled at roughly twice the maximum frequency of interest (or less) has proven to be suitable. For the current program, this requirement was met by collecting the data at a sampling rate of 300 KHz and lowpass filtering to pass frequencies below 100 KHz.

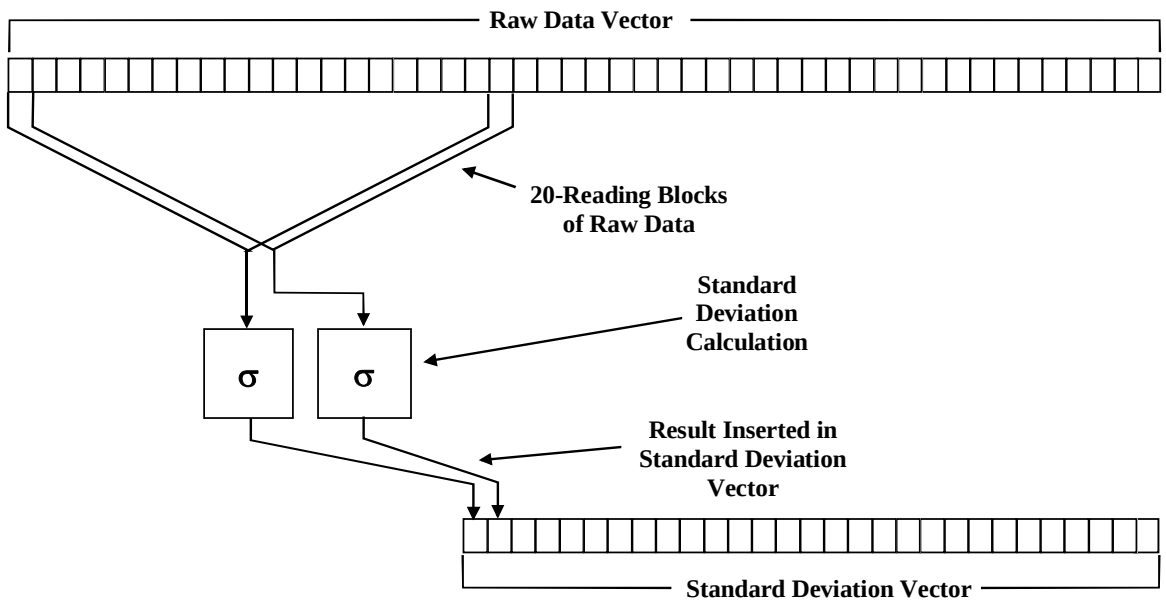

Figure A1. Generation of moving standard deviation vector

A.2 Amplitude Measures

Average (Avg) – This is the simplest of all amplitude measures. For a zero-centered signal such as those provided by accelerometers and acoustic emission detectors, the average value should nominally be zero. Thus, this measure is not used to characterize the raw data of such signals. However, when the passage of harmonic power is characterized, the average value is significant, so the average is used as one measure for the standard deviation vector. This is the only amplitude measure that was calculated for the standard deviation vector and not for the raw data.

Standard Deviation (StDev) - The simplest measure of amplitude of an irregular signal is the standard deviation, which reflects the harmonic power of the signal. In addition, in nonlinear dynamics, the radius of an attractor is normally taken to be the standard deviation. Thus, this measure characterizes the size of the attractor.

Absolute Deviation (AbsDev) – The standard deviation weighs large variations in the signal more strongly than small variations, so we identified this measure to weigh all events evenly. The absolute deviation is the average absolute value difference between each value in the time series and the average value of the time series.

RMS Difference (RMSDif) – This measure is the RMS difference in sequential readings, a measure of the small-scale "texture" of the data.

Average Difference (AvgDif) – The RMS difference weighs larger differences more strongly than smaller differences, while this measure weighs all differences evenly. The average difference is the average absolute value difference between sequential values in the time series.

Since the average value is not used with the raw data, applying the preceding measures to the raw signal and standard deviation vector accounts for 9 of the 57 signature quantities.

A.3 Integral Measures

This group of measures characterizes the size of large events by summing the "height" of each value in the time series above a reference line. Two different reference lines are used for these signature quantities: the average value of the data, and the average plus the standard deviation.

Size of Duration Events Above the Average Line (ADuSize) – The average sum of each value in a duration event minus the average value. A duration event occurs between an upward crossing of the average line and the next downward crossing, as illustrated in Figure A2, below. Thus, this measure is an estimate of the integral of this event.

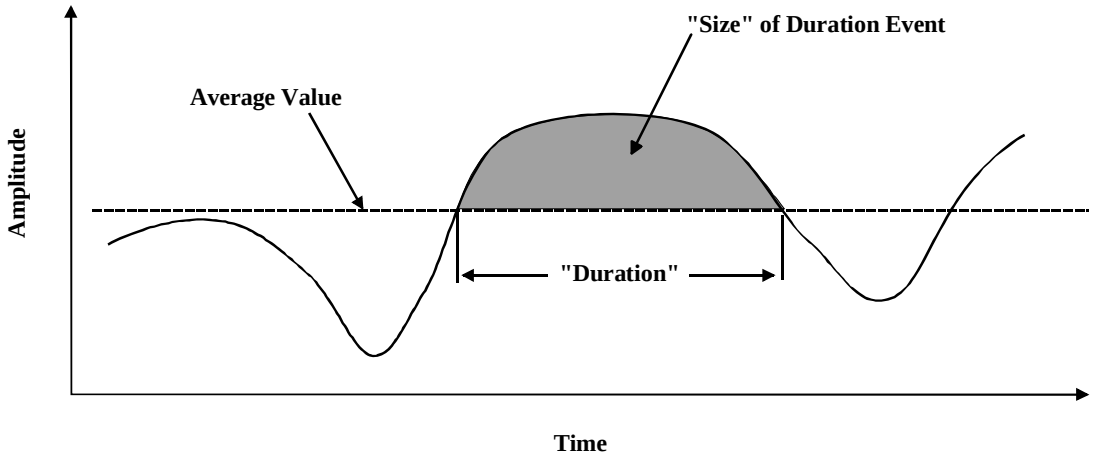

Figure A2. Size of duration events for average-crossing waves

RMS Size of Duration Events Above the Average Line (ADRMSiz) – Similar to the previous measure, but the RMS of the event sizes. This measure weighs larger events more strongly than smaller events. Together with the previous measure, this measure offers information about the distribution of the sizes of events.

Size of Duration Events Above the Average-Plus-Standard-Deviation Line (SDuSize) – Similar to ADuSize, but considering a higher "bar", and thus examining only the largest events in the data.

RMS Size of Duration Events Above the Average-Plus-Standard-Deviation Line (SDRMSiz) – The RMS version of the above measure, weighing the largest events more heavily than smaller events.

Used in both the raw data and standard deviation vectors, the preceding measures comprise 8 of the total of 57 measures.

A.4 Normalized Amplitude Measures

Each of the measures above, including the amplitude and integral measures, is directly sensitive to the amplitude of the signal. Thus, if one sensor is 10% more sensitive than another, with everything else held steady, the measures above will be 10% larger for that sensor. In an attempt to reduce or eliminate the sensitivity of the results to sensor sensitivity, additional measures were created by computing the ratio of each amplitude measure to the average and/or standard deviation of the corresponding vector. Thus, new signature quantities were produced that comprise AbsDev/StDev, AvgDif/Avg, etc. In principle, one might expect that only one normalization parameter would prove to be necessary (e.g., dividing by average or standard deviation, but not both). However, it often occurs that one or the other normalization parameter is strongly favored for different amplitude-sensitive measures, so both sets are retained. In the shorthand notation, values normalized by dividing by average were signified by appending **oAve** to the signature quantity descriptor, and those normalize by dividing by standard deviation by appending **oStd**. Of course, since the average value of the raw data is a meaningless parameter, the signature quantities calculated for the raw data used only the standard deviation as the normalization parameter. This produced a total of 22 new measures (7 for the raw data and 15 for the standard deviation vector).

A.5 Time Measures

Of course, the amplitude of the signal is rendered totally irrelevant in signature quantities that are measures of time. The following measures were developed to characterize the actual time evolution of the signal.

Characteristic Autocorrelation Time (Auto) – This is the time scale during which the signal loses linear correlation, and has been suggested as a measure of the rate of mixing in a process. The autocorrelation $c_{xx}(\tau)$ is the integral over all time of the signal x(t) times a time-delayed version of the same signal, x(t-τ). For no time delay (τ=0), the autocorrelation thus equals Σx^2. It turns out that, for long time delays (large τ), the autocorrelation of aperiodic signals approaches $n\bar{x}^2$, where n is the number of readings in the record. As shown in Figure A3, the characteristic autocorrelation time is the intercept of a line at the initial downward slope of the autocorrelation curve with the runout value. Measured in sample intervals, it can be calculated using

$$\Delta t_{auto} = \frac{\Sigma x^2 - n\bar{x}^2}{\Sigma x^2 - \Sigma x_i x_{i+1}}$$

The last term in the denominator is the sum of the product of each value with it immediate successor in the time series.

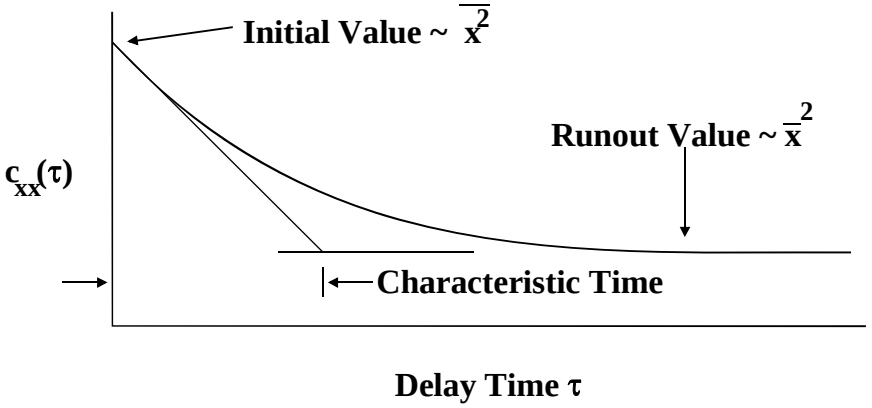

Figure A3. Characteristic autocorrelation time

Period of Average-Crossing Waves (APd) – This is the average time between positive crossings of the average line, as illustrated in Figure A4, and characterizes the time period for first-order signal variation.

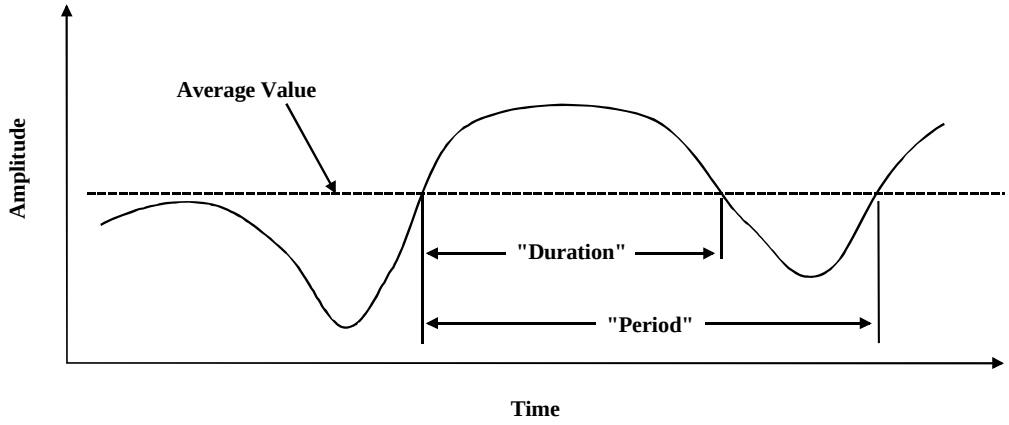

Figure A4. Measures of period and duration of average-crossing waves

RMS Period of Average-Crossing Waves (ARMSPd) – This is the RMS value of the preceding measure, and emphasizes the influence of waves with longer periods.

Duration of Average-Crossing Waves (ADur) – This is the average number of readings that occur between an upward and downward crossing of the average line. It is a measure similar to APd, but is generally not half its value.

RMS Duration of Average-Crossing Waves (ARMSDur) – This is the RMS value of the preceding measure, and emphasizes the influence of large-duration waves.

Period of Standard Deviation-Crossing Waves (SPd) – This measure is similar to the period of average-crossing waves (APd), but in this case the "bar" is raised somewhat to look at the period of waves that cross the average plus standard deviation (Figure A5). This measure characterizes the timing of large-scale waves in the data.

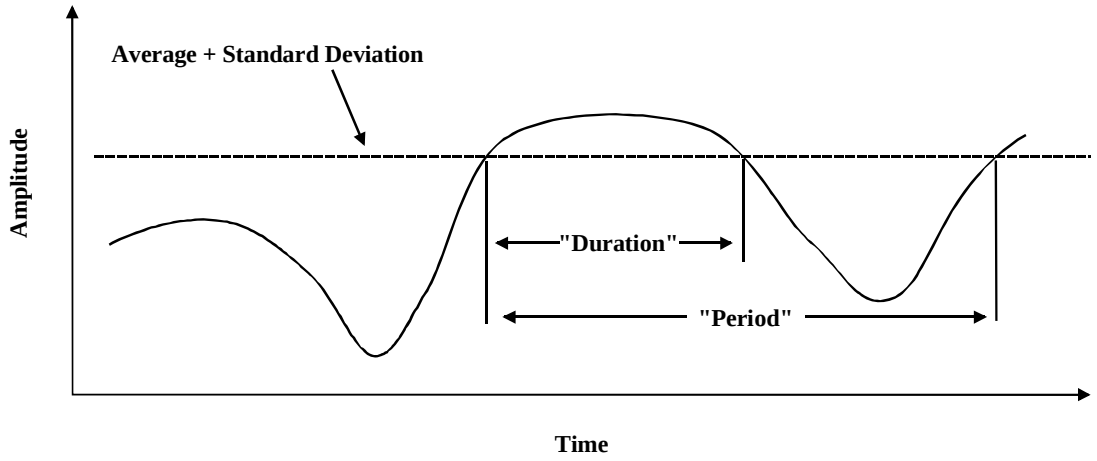

Figure A5. Period and duration of standard deviation-crossing waves

RMS Period of Standard Deviation-Crossing Waves (SRMSPd) – The RMS value of the previous measure, again emphasizing the influence of the largest waves.

Duration of Standard Deviation-Crossing Waves (SDur) – The average duration of these large waves.

RMS Duration of Standard Deviation-Crossing Waves (SDur) – The RMS value of the preceding measure, again emphasizing the influence of the largest waves.

These time measures were applied to both the raw data and standard deviation vector, comprising 18 of the 57 signature quantities.

A.6 Summary of Signature Quantities

Table A1 below summarizes the signature quantities described in the preceding sections. This set of quantities has proven surprisingly capable at relating the dynamics of sensor signals to the system conditions in numerous Dynamical Instruments development programs.

Table A1. Summary of Dynamic Signature Quantities

Sig Qty Name	Raw Data or Standard Deviation	Amplitude	Integral	Normalized	Time
StDev	Raw data	X			
AbsDev	Raw data	X			
AbDoStD	Raw data	X		X	
RMSDif	Raw data	X			
RDioStD	Raw data	X		X	
AvgDif	Raw data	X			
AvDoStD	Raw data	X		X	
Auto	Raw data				X
APd	Raw data				X
ARMSPd	Raw data				X
ADur	Raw data				X
ARMSDur	Raw data				X
ADuSize	Raw data		X		
ADSoStD	Raw data		X	X	
ADRMSiz	Raw data		X		
ARSoStD	Raw data		X	X	
SPd	Raw data				X
SRMSPd	Raw data				X
SDur	Raw data				X
SRMSDur	Raw data				X
SDuSize	Raw data		X		
SDSoStD	Raw data		X	X	
SDRMSiz	Raw data		X		
SRSoStD	Raw data		X	X	
sdAvg	Std dev	X			
sdStDev	Std dev	X			
sdSDoAve	Std dev	X		X	
sdAbsDev	Std dev	X			
sdAbDoAve	Std dev	X		X	
sdAbDoStD	Std dev	X		X	
sdRMSDif	Std dev	X			
sdRDioAve	Std dev	X		X	

Sig Qty Name	Raw Data or Standard Deviation	Amplitude	Integral	Normalized	Time
sdRDioStD	Std dev	X		X	
sdAvgDif	Std dev	X			
sdAvDoAve	Std dev	X		X	
sdAvDoStD	Std dev	X		X	
sdAuto	Std dev				X
sdAPd	Std dev				X
sdARMSPd	Std dev				X
sdADur	Std dev				X
sdARMSDur	Std dev				X
sdADuSize	Std dev		X		
sdADSoAve	Std dev		X	X	
sdADSoStd	Std dev		X	X	
sdADRMSiz	Std dev		X		
sdARSoAve	Std dev		X	X	
sdARSoStD	Std dev		X	X	
sdSPd	Std dev				X
sdSRMSPd	Std dev				X
sdSDur	Std dev				X
sdSRMSDur	Std dev				X
sdSDuSize	Std dev		X		
sdSDSoAve	Std dev		X	X	
sdSDSoStD	Std dev		X	X	
sdSDRMSiz	Std dev		X		
sdSRSoAve	Std dev		X	X	
sdSRSoStD	Std dev		X	X	