

Final Technical Report

Title: “Intelligent Potroom Operation”

This material is based upon work supported by the U.S. Department of Energy under Cooperative Agreement Number DE-FC07-99ID13815

Project Period: September 30, 1999 to April 30, 2003

Date of Report: July 29, 2003

Report Authors: Jan Berkow & Larry Banta with guidance by Richard Love

Award Recipient: Applied Industrial Solutions, LLC
937 Southpoint Circle
Morgantown, WV 26501

Teaming Members:

- Jan Berkow, Project PI; Applied Industrial Solutions, LLC; Morgantown, West Virginia
- Larry Banta & graduate students; West Virginia University; Morgantown, West Virginia
- Richard Love & staff; Century Aluminum of West Virginia; Ravenswood, West Virginia
- Gensym Corporation; Burlington, Massachusetts; replaced in 3rd year by Mindmine, Inc., Sparks, Maryland and Intelligent Laboratory Solutions, Inc., Naperville, Illinois

Any opinions, findings, and conclusions or recommendations expressed in this material are based on those of the authors and do not necessarily reflect the views of the Department of Energy.

“Formal acknowledgement and a great deal of thanks must be given for the support, guidance, and technical vision provided by Richard Love, the Technical Manager for Century Aluminum. Without his involvement, the success of this project would not have been possible.”

Table of Contents

1	EXECUTIVE SUMMARY	1
2	PROJECT OBJECTIVES AND ACCOMPLISHMENTS.....	2
2.1	Cell Control Enhancement Module Objectives (CCEM).....	2
2.1.1	CCEM Project Activities	3
	Data Acquisition System	3
	Signal Processing	4
	FEM Development	4
2.1.2	CCEM Results	5
	Sampling Rates	6
	Filtering	6
	Cell Noise Discrimination and Metrics	8
	Parametric Studies Using the Finite Element Model.....	13
	Alumina Concentration Estimator.....	15
	Bath Ratio Estimation.....	15
2.1.3	CCEM Summary and Conclusions.....	15
2.2	Intelligent Potroom Advisor (Corrective Action Neural Network)	16
2.2.1	Intelligent Potroom Advisor Project Activities	17
	IPA Design Considerations	19
	IPA Development	21
	User “Buy-in”	22
	Cultural Challenges	23
	Predictive IPA	24
	Data Latency.....	27
2.2.2	IPA Results	27
2.2.3	IPA Summary and Conclusions.....	28
3	PRODUCTS.....	31

1 Executive Summary

In primary aluminum production plants, all potroom reduction cells suffer from occasional substandard performance in which cells operate at less than peak efficiency. The period when the cell operates in this manner can be referred to as an “Off-Peak Mode” or OPM. OPMs result in dramatic increases in energy consumption and, in the case when the OPM is an anode effect, dramatic increases in the emissions of perfluorocarbon gases (PFCs). The Intelligent Potroom Operation project focuses on maximizing the performance of an aluminum smelter by innovating technology to eliminate or continuously reduce both the duration and frequency of the OPMs that occur.

This project addresses two primary challenges in creating a continuous improvement strategy for OPM reduction in an aluminum smelter. The first is the ability to accurately assess the current operating state of a reduction cell. This difficulty is primarily due to the limited data set that can be obtained on a continuous basis due to extremely high temperatures and corrosive conditions that accompany the alumina reduction process. The second challenge is the ability to continually build and apply corporate knowledge for maximizing process performance. This is made more difficult by the limited data set and the lack of capability to apply a systems analysis perspective to obtain the synergistic benefit of analyzing multiple data sets.

The Intelligent Potroom Operation project has successfully addressed each of these challenges by innovating three key pieces of technology in order to continuously reduce reduction cell Off-Peak Modes. The first innovation was the development of Cell Control Enhancement Module (CCEM). The CCEM consists of a data acquisition system that derives additional understanding of reduction cell health by analyzing cell noise in new ways. The data acquisition system was used to refine and validate a finite element model-based cell state estimator to overcome the first challenge of understanding the current operating state of a reduction cell.

The potential to augment the existing process control system with more accurate cell state data was never realized during the scope of this project. The significance of the contributions made to the improved understanding of the alumina reduction operation was, however very favorably received by industry and academic peers at the Minerals, Metals, and Materials Society (TMS). Two peer-reviewed papers about the Cell Control Enhancement work have been presented in each of the last two TMS Light Metals conferences (Refs [3, 4, 5, 6]). Additional papers for this venue are in preparation (Refs [11, 12]). Future publication in archival journals is planned. Encouragement was also received to extend this work beyond the primary aluminum industry. Accordingly, Century Aluminum is continuing investment in this technology to further refine the data acquisition system in order to validate the benefits of cell state estimator approach. Century is presently using this technology as part of ongoing optimization studies.

The second innovation was the development of the formalism for both the methods and concepts involved in building an intelligent manufacturing system. One component of this system called the Intelligent Potroom Advisor (IPA), initially called the Corrective Action Neural Network (CANN), was successfully prototyped and deployed on one of the four production lines within the Century Aluminum plant. The IPA provides a platform to monitor process data in order to identify those cells exhibiting behaviors (defined as temporal data patterns) that require immediate attention. It serves as a resource

management tool for prioritizing where to focus lean operational resources. The IPA also provides a means of capturing heuristics of existing experienced staff or for purposes of facilitating discovery to determine the significance of unknown events that occur during OPMs. Based on the success of the IPA to demonstrate these benefits, Century is expanding portions of the IPA to all four production lines. These are the portions that offer known value within a reasonable response time per recognized data latency issues that are discussed later. Development of the other portions continue in the first production line.

The IPA provided Century Aluminum with the ability to manage heuristic-based corporate knowledge so that operational personnel could intervene when a reduction cell was exhibiting an OPM. The final goal of the Intelligent Potroom Operation project was to go beyond the “early intervention” of OPMs to an operational capability of “predictive avoidance” in order to maximize plant performance. It was recognized that a method was required to identify unknown behaviors that preceded periods in which the reduction cells were in OPMs. In addition, it was concluded that a method was required to overcome cognitive limitations in identifying these kinds of behaviors that may be in the form of complex variable relationships that are continually changing due to the dynamics that occur due to the natural life cycle of the reduction cell, changes due to modifications in operational practices, or a combination of both.

A third innovation was the development of a process to apply rule induction to identify statistically significant yet unknown complex behaviors that occur prior to OPMs. This process was developed and successfully performed manually during this project. Subsequent to this project and, with the encouragement of the Department of Energy’s Office of Industrial Technology, a proposal was submitted for the development of semi-automated solution for knowledge building, which addresses the “adaptive” requirement for an intelligent system.

2 Project Objectives and Accomplishments

The following section compares project accomplishments with the original objectives of the project. A discussion of the Cell Control Enhancement Module is provided in Section 2.1 followed by a discussion of the Corrective Action Neural Network, which was later renamed Intelligent Potroom Advisor in Section 2.2.

2.1 Cell Control Enhancement Module Objectives (CCEM)

The objectives of the Cell Control Enhancement Module were as follows:

1. To design, construct and install an enhanced data acquisition system (DAQ) on working aluminum reduction cells at an operating aluminum plant. The DAQ would allow the collection of many more channels of data than were available from the existing cell controller. The DAQ would also allow collection of data at sampling frequencies that could be controlled by the researchers.
2. To construct a cell “state estimator” by blending first principles models of the physical and chemical processes occurring in the cell with an enhanced sensor suite. The latter would measure available operating parameters and would not

require exotic new sensors, but would provide enough information to allow development of the state estimator.

3. To interface the CCEM with the Intelligent Potroom Advisor (IPA) and provide high-level information on cell state/status to the IPA.
4. To construct a cell controller that would “sit on top of” the existing cell control system and not require replacement of existing equipment.
5. To demonstrate the system in a working aluminum potline.

2.1.1 CCEM Project Activities

Development of the Cell Control Enhancement Module proceeded along two paths. One path involved the examination of data collected by a custom Data Acquisition System, designed and fabricated for this project at WVU. The focus of this work was to isolate signals that would provide insight into specific physical, electrical or chemical processes occurring in the cell. The objective was to derive new information from signals that are already available using the existing instrumentation in the typical aluminum reduction cell. Specifically, this offered the ability to derive value from existing cell noise. The second research path involved the construction of a finite element model of the cell, verification of the model, and use of the model to better understand and characterize the complex processes occurring in the cell. The objective of this path was to develop mathematical models that could be used for estimation of the state of the cell. The fundamental idea was to use the models and what sensor data was available, plus information derived from the signal processing task to deduce quantities that cannot yet be directly measured—alumina concentration and anode-cathode distance being the two most important ones. The two paths and their constituent parts are described below.

Data Acquisition System

It was necessary to design and fabricate a data acquisition system (DAQ) for installation on working reduction cells at an aluminum plant. This was necessary because it was required to collect data from a wide variety of sensors not supported by the existing cell controllers. It was also necessary to be able to adjust the sampling rate and to collect unfiltered data, neither of which was possible using the existing controllers. The DAQ was designed from standard components based on the Opto 22 (TM) system.

Four data acquisition systems were constructed and mounted on the duct end of four cells in a pot line at the test plant. Each system included compressed air cooling and fiber optic communications to a host computer in the main computer room of the pot room complex. Data from all four pots was uploaded at each sampling interval and stored on the hard drive of the host computer. Periodically, data was moved from the host computer to CD ROM and erased from the host hard drive to free up space.

Each DAQ can accommodate up to 32 channels of either analog or digital information. As configured, the systems continuously measured the parameters shown in Table 1 all the time. The 16 spare channels were used to measure additional parameters for special tests. The additional parameters included extra thermocouples and voltage taps on anode rods during a series of model verification tests described later in this document.

Signal Name	Input Type
Cell Voltage, limited	Analog 0-10V
Cell Voltage, attenuated	Analog 0-10V
Quadrant Voltage 1	Analog 0-75 mV
Quadrant Voltage 2	Analog 0-75 mV
Quadrant Voltage 3	Analog 0-75 mV
Quadrant Voltage 4	Analog 0-75 mV
Optical Encoder Channel A	Digital TTL
Optical Encoder Channel B	Digital TTL
Feed Switch Status	0-24VAC
Break Switch Status	0-24VAC
Duct Thermocouple	Type K T/C
Duct Pitot Tube (differential P)	0-50mV

Table 1: Data Acquisition System Inputs

Signal Processing

Data collected by the DAQ was analyzed in numerous ways in a search for signals that represented physical processes within the cell and which would provide information about the cell state. Numerous frequency domain methods were applied, including extensive experimentation with high-pass, low-pass and band-pass filters, Fourier Analysis, and wavelet analysis. Statistical measures were applied to the voltage signal and its components. Three important results emanated from this work. First, a filtering method was developed that allows clear distinction of the alumina feed cycle effects on the cell resistance signal. Second, a method was developed to make inexpensive measurements of cell current balance and to detect and quantify cell imbalances easily. Finally, a set of filters and metrics was developed to separate three components of what has heretofore been lumped together as “noise”. The metrics form the basis for a new anode control strategy that has the potential to increase cell power efficiency and stability. A thorough discussion of the signal processing methods is provided in [2].

FEM Development

The second major thrust in the CCEM development was the construction of a Finite Element Model (FEM) of the reduction cell. The FEM provides first-principles models of all of the major cell processes. The model has three components: an electrical model, a thermal model, and a chemical model. The electrical model is a 2-D steady-state finite element model, which is computed at each time- step. The thermal model is a 2-D transient finite element model. The chemical model is a lumped parameter model using scalar equations.

In developing models, a compromise must be made between the accuracy of the model and the computational speed. At one extreme is a lumped parameter model, which can be solved quickly, but has a limited accuracy. At the other extreme is a finite element model with a very fine mesh, resulting in very accurate spatial discretization, but requiring a significant amount of time for computation. For this work, a finite element model was chosen, but with a fairly coarse mesh. This enables the model to be solved in a reasonable amount of time, while providing adequate accuracy to characterize the major processes in the cell. A full 2-D slice of the cell, with two anodes, is modeled

using 292 elements and 1227 nodes. A schematic of the cell and corresponding mesh can be seen in Figure 1.

An extensive literature search was conducted to find the necessary governing equations for the FEM. The model was refined and verified by conducting parametric studies and comparing the results with both values from the literature and with data taken in a set of experiments performed on the instrumented cells at the test site.

Using the FEM, several types of parametric analyses have been run to better understand some of the interactions that occur in the cell. Extensive studies on the thermal balance in the cell and the impact of anode position, bath chemistry and alumina concentration on the cell have yielded some very useful insights. The FEM was also used as the basis for the construction of scalar models of alumina concentration that led to the development of an “estimator” for that important control parameter. The concentration estimator and parametric studies are discussed in more detail in the Results section of this report.




Figure 1: Schematic of cell with finite element mesh. Black lines are physical boundaries of cell elements like anodes, metal pad, ledge, etc. Blue lines show finite element mesh. Axis scales are in centimeters.

2.1.2 CCEM Results

Many of the early results of this work have already been published in the open literature. To date, the research has resulted in one Ph.D. Dissertation [1], one M.S.M.E. Thesis [2], and six conference papers [3-8]. Abstracts for two additional conference papers have been submitted [11, 12] and two manuscripts are in preparation for archival journals. Links to online copies of the Dissertation and M.S. Thesis are provided in the references.

Sampling Rates

Reduction cell voltage and current signals are extremely noisy, and low-pass filtering of the signals has long been practiced to eliminate some of the noise. Since cell voltage is a function of both the cell resistance and the current, voltage signals are generally converted to a “pseudo-resistance” R_p as a first step, and then filtering occurs on the R_p signal. Commercial controller manufacturers consider their filtering schemes proprietary, and details are difficult to learn. A common approach seems to be simple averaging filters. The characteristics of the filters are then functions of the sampling rate and the averaging period. The question of how fast to sample the cell voltage and current signals has been debated for decades. Fast sampling captures more information but generates huge amounts of data for storage. Still, there has always been a feeling that useful information might be hidden in the higher frequency portion of the R_p spectrum. Extensive studies of the frequency spectrum of the cell resistance signal were undertaken using data from the instrumented cells. In the data examined, little information content was found at frequencies higher than about 1 Hz. Broadband noise exists across a wide range of frequencies higher than this, but the power density of that noise did not seem to vary with any discernible cell conditions.

The examination of this data is extremely labor intensive and the data itself is probably site-specific, perhaps even cell-specific. The researchers do not claim to have made an exhaustive study of this topic, but report here only that no encouragement was found to pursue the detection of “high frequency” signals, i.e. those above 1 Hz. Consequently, the research team concluded that a sampling frequency of 5 Hz was adequate for the analyses described in the remainder of this document.

Filtering

Much of the work done in this part of the project focused on the detection of signals related to the alumina concentration and the anode-cathode distance, since these are two of the main control parameters. The objective of the alumina feed system is to maintain the concentration of dissolved alumina in the bath at a desired setpoint, usually between 2% and 4% by mass. Since alumina cannot be measured directly, many controllers operate by intentionally causing fluctuations in the alumina concentration and looking for changes in bath resistance that correspond to those fluctuations. At the test plant, cells feed 6-7 kg batches of alumina about every 6-7 minutes under “normal feed” conditions. “Overfeed” periods are about 2-3 minutes between feeds, and “Search” mode inhibits all alumina feeds in this specific cell and anode movements while the controller looks for a rise in the cell resistance curve. Figure 2 is a trace of cell resistance over an entire cycle of search, overfeed, normal feed and back to search mode.




Figure 2: Cell resistance signal over a complete cycle of feed modes for a typical cell.




Figure 3: Signals from Figure 2, filtered by 4-pole Butterworth low-pass filter with cutoff frequency 0.1 Hz. Vertical lines are feed events.

Various spectral techniques were examined for isolating the fluctuations in alumina concentration from feed events. Several FFT and discrete filter bank systems were investigated with little success, due to the extremely low signal-to-noise ratios in the data. The most effective approach was low-pass filtering of the cell resistance signal, with the objective of identifying variations through the entire break/feed cycle. Figure 3 is the same signal as Figure 2 after low pass filtering. In this case, a 4-pole Butterworth filter was used with a cutoff frequency of 0.1 Hz. Feed events are indicated by the vertical lines. The overall trend of the signal during the various phases of the feed/search cycle is apparent, although significant amounts of noise are still present. The filtering has made more obvious the small fluctuations in the cell resistance as feed events occur. Figure 4 is an expansion of a portion of the cycle, including both overfeed and normal feed periods. The rise and decay in the resistance during the individual feed events are clearly visible. It is worth noting that the fluctuations mentioned do not occur




Figure 4: Expanded section of Figure 3, showing distinctly the effect of the feed events on cell resistance. Fluctuations are due to temperature variations in the bath, not to variations in alumina concentration.

as a result of alumina concentration changes, but are driven by changes in the bath temperature. Bath conductivity increases with temperature, so when a mass of alumina and crust are dumped into the bath in a feed operation, the temperature decreases measurably—10C or more in Century’s cells. It is this sudden temperature decrease that drives the bath resistivity and thus the cell voltage up immediately after the feed.

Cell Noise Discrimination and Metrics

One of the most important results of this work was the development of a set of filters and metrics for discriminating “noise” from different cell processes. This is important because some types of noise are thought to be detrimental to cell current efficiency, and so all controllers employ some type of noise control strategy. Noise control is generally effected by raising the anodes. However, raising the anodes increases the power consumed by the cell, and may not be effective in controlling certain types of noise.

Three noise types were studied as the most important from the standpoint of pot control. Bubble noise is the fluctuation in cell resistance due to the buildup and release of gas bubbles on the bottom of the anodes. Bubble noise is a natural consequence of cell operation. It cannot be eliminated, but should be recognized as an essential background component. Some of the characteristics of bubble noise are illustrated in Figure 5 below. The intensity of bubble noise can be measured by the variance or standard deviation, with higher noise levels corresponding to higher values of sigma. From the data collected at Ravenswood, the value of sigma ranges from about 0.02 to about 0.04 .

A second and very important type of noise is anode short circuiting noise, which is indicative of a condition that should elicit control action. When the anodes are too close to the bath/metal pad interface, splashing or rolling of the metal can allow current to arc or short-circuit directly from the anode through a metallic path to the cathode. Short circuiting may also result from one or more improperly set anodes, from interactions between gas bubbles, the metal pad and sidewall freeze, or from other unknown factors. When this happens, energy is consumed, but no electrolysis of alumina takes place—

thus current efficiency and energy efficiency are both reduced. The anodes should be immediately raised to bring the short-circuiting under control.

This type of noise is characterized by large downward spikes in the pseudo resistance plot, as shown in Figure 6. Note in these plots that the upper envelope of the resistance signal is relatively uniform compared to the lower envelope which is punctuated by very large downward excursions.




Figure 5: Bubble Noise (left) and its Histogram (right), the enveloping curve on the histogram is a Gaussian distribution with mean zero and standard deviation 0.0234




Figure 6: Anode short-circuiting noise is characterized by downward spikes in the cell resistance

Another common phenomenon can also cause fluctuations or “noise” in the cell resistance signal. The strong electromagnetic fields present in the cell can cause waves to occur in the metal pad. This phenomenon is often called “metal pad roll”. Metal pad roll can look like high intensity bubble noise, but expansion of the signal reveals that it has a periodic nature rather than the white noise characteristic of bubble noise. Figure 7 illustrates this characteristic.

Metal pad roll can increase the likelihood of short circuiting, but is not the same thing, and in fact may require a different control strategy. Studies on the cells at an operating aluminum facility indicate that raising anodes is ineffective in quenching metal pad roll. In fact, anode movement often appears to instigate metal pad roll where none was evident before. The most effective strategy for quenching metal pad roll seems from these studies to be to avoid moving the anodes if possible—exactly the opposite of what the current controller does. Figure 8 shows that repeated efforts to quench metal pad roll by raising the anodes were ineffective.




Figure 7: Metal pad motion noise can look much like high-intensity bubble noise (left). Magnifying the time scale reveals the sinusoidal pattern characteristic of metal pad roll.

Magnifying the time scale reveals the sinusoidal pattern characteristic of metal pad roll. Algorithms were developed to discriminate the three types of noise from one another and metrics were developed to allow the relative levels of each type of noise to be quantified. The details of this work are described in references [2] and [3], however the results are shown in the figures below. There is no metric for the bubble noise beyond the signal sample variance, and it is not interesting as a control parameter. It may have significance towards determining modification or enhancement for a more effective anode shape. The metric for short circuiting noise is called the “splashing ratio”, and its range is from 1 to perhaps 3 or 4. Numbers higher than 1 indicate the presence of short circuiting noise. The metric for metal pad roll is called the “roll index”. Its range is from zero to 3 or 4. Numbers above about 1 indicate the presence of significant metal pad roll.




Figure 8: Metal pad motion noise, and the commercial controller's attempts to suppress it by anode raises.




Figure 9a: Composite signal exhibiting different types of noise. From 0-11 = bubble noise, 1100-2200 = short-circuiting, 2200-3300 = metal pad roll, and 3300-4400 = both roll and short-circuiting noise.

For control purposes, it is desirable to have the metrics be selective of the type of noise incident in the signal. Figure 9 demonstrates that the proposed metrics are highly selective for the chosen noise types. The top trace shows a composite signal consisting of periods of pure bubble noise, short-circuiting noise, metal pad roll and a combination




Figure 9b: Short-circuiting Ratio trace for signal in Figure 9a.




Figure 9c: Roll Index trace for signal shown in Figure 9a.

of short-circuiting and metal pad roll. The second trace shows the Short-circuiting Ratio (SR) and the third trace shows the Roll Index (RI). Note that for the initial bubble noise, both the SR and the RI are low. During the second part of the trace, where the short-circuiting noise is dominant, the SR is high, but the RI remains low. In the third segment, the RI jumps up to indicate high metal pad roll, but the SR drops back because there is not much short-circuiting noise present. In the last segment both types of noise are present, and both the SR and the RI exhibit elevated values.

Parametric Studies Using the Finite Element Model

The finite element model is a powerful tool for performing “what if” studies to understand how the physical environment, feed materials and control decisions affect the operation of the cell. Sometimes the interactions among the variables are so complicated that the results are not transparent.

An extensive set of parametric studies was performed on the cell, focused primarily on the energy/thermal balance and the effects of various upsets on the bath ratio. For example, the effect of a step change in anode-cathode distance (i.e. cell voltage) was examined, as were the effects of fluoride addition, soda ash addition, bath and metal pad depth changes, and improperly set anodes. Among the most important findings were the following:

- Time constants for thermal and chemical equilibrium are longer than expected, even in the relatively small cells at the test plant (92 kA). A step change in cell voltage required approximately 40 hours to establish new equilibrium condition (see Figure 10 below). Additions of soda ash and aluminum fluoride required similarly long periods to effect stable changes in bath ratio because of the interaction with sidewall freeze associated with changes in the liquidus temperature of the bath.
- Freezing and melting of sidewall ledge has a significant influence on bath ratio, which in turn affects the liquidus temperature of the bath. Fortunately the solid/liquid transitions have a damping effect rather than a destabilizing effect on the system.
- Increasing cell voltage to increase bath temperature is not effective as a long-term strategy, since the cell operating temperature is driven more by bath ratio than by voltage. Cell operating temperature should be regulated by monitoring and regulating the bath chemistry rather than cell voltage, which should be kept as low as possible to promote energy efficiency of the process.
- Bath depth has relatively little effect on the cell energy balance, but metal pad depth significantly influences the energy balance of the cell, since greater metal pad depth increases the amount of heat loss through the sidewalls adjacent to the liquid metal. The effect is to cool the bath, increasing the thickness of the sidewall ledge above the metal pad, as shown in Figure 11 below.




Figure 10: Ratio versus time for a step change in cell voltage from 4.50V to 4.40V

Details concerning these studies and the findings are presented in [1] and [4].




Figure 11: Sidewall thickness versus metal pad depth

Alumina Concentration Estimator

One of the most important parameters in the control of the reduction cell is the concentration of dissolved alumina in the bath. If the alumina concentration falls too low, the cell will go into anode effect: an unstable state in which the bath itself is disassociated. During anode effect, the cell produces perfluorocarbons, which are extremely persistent ozone-depleting chemicals. In addition, the cell power consumption increases by as much as ten times, and the thermal balance of the cell is completely disrupted. On the other hand, if the alumina concentration is too high, cell efficiency falls and undissolved alumina can be deposited on the cathode as “sludge” or “muck”. The alumina forms an electrically insulating layer initially creating shorting through low control voltage and then increasing power consumption and causing thermal problems in the cell.

Proper control of the alumina concentration is difficult because there is no alumina sensor that can withstand the hot, corrosive, abrasive environment of the bath. Most cell controllers force the cell to operate in a non-optimal range of alumina concentrations, for the reason that by doing so, it is possible to deduce alumina concentration from cell resistance changes. One of the major accomplishments of this research has been the development of an “alumina concentration estimator” that will enhance the ability of the cell controller to make correct and effective decisions concerning feed rate modifications.

The estimator is not yet completely refined. To date it has been impossible to test the estimator in an “active” mode, i.e. using the estimator output for control. In tests using real time data but with the estimator operating only in passive modes, it proved only modestly effective at providing an absolute concentration reading, e.g. “alumina concentration is 3.1%”. However, even in passive mode the estimator was quite effective in providing a relative concentration estimate, e.g. “too low”, “within working range”, or “too high”. In those cases the estimator provided correct advice more than 80% of the time. Even this level of information is a major step up from what has been available in the past. Given better test conditions and additional refinement, the algorithm should be able to improve its accuracy significantly. Details of the estimator design and performance can be found in [1], [5] and [8].

Bath Ratio Estimation

Attempts to estimate bath ratio from first-principles modeling were not as successful as those for the alumina estimator. No effective algorithm(s) could be derived to estimate bath ratio from the standard data set currently being taken. However, if a bath temperature measurement is taken, a very accurate estimate of the bath ratio can be computed. Since temperature is easy and relatively cheap to measure, the estimator is a major improvement over the existing method of taking a bath sample and sending it to the lab for chemical analysis. Periodic chemical analysis will still be necessary to recalibrate the estimator, but between analyses, this algorithm will provide continuous estimates of bath ratio. More information on the ratio estimator is given in [1].

2.1.3 CCEM Summary and Conclusions

All of the objectives of the CCEM task have been met with the exception of the last two. The DAQ was constructed and has collected more than 30 Gigabytes of data on the four

reduction cells in the two years it has been in operation. The system is still installed and operational. Data from the DAQ has been used in two ways. It was used to refine and validate a finite element model of the reduction cells, from which the cell state estimator algorithms were constructed. The data was also analyzed directly, and formed the data sets from which a new level of understanding of reduction cell noise was extracted. Although not identified explicitly as an objective in the original research plan, the work done on isolation and quantification of cell noise is one of the major contributions of this research.

The original concept for the DAQ called for it to work with rather than replace the existing cell control hardware. That approach has been found to be only partially realizable, due to a combination of hardware incompatibilities and intellectual property issues with the existing controller manufacturer. Each DAQ module has output capability designed into the hardware that would allow the individual boxes to serve as cell controllers, independent of the legacy system. However, it was felt that the software was not at a mature enough stage by the end of the project to trust with the safety of the four reduction cells. It is also impractical to try to patch the data from the CCEM into the data structures of the existing control system. Thus the demonstration goal was not achieved within the duration of the project. However, privately funded development and refinement of the alumina concentration estimator and bath ratio estimator algorithms continues beyond the project.

Significant contributions to the understanding of aluminum reduction cell operation and control have been generated by this work. The research has been well received by industry and academic peers in the Minerals, Metals and Materials Society (TMS). It has received favorable reviews in DOE Annual Project Review presentations in 2001 and 2002. Reviewers in the latter presentation suggested that the work may well be applicable to a host of industries beyond the Aluminum industry, and the team has in fact begun to pursue "crosscutting" applications of the ideas and approaches developed in this project.

Considerable development work still remains to bring the ideas generated in this project to commercialization. The entire DAQ must be redesigned to be smaller, cheaper and more capable. Software algorithms must be tested over larger data sets to prove stability. Interfaces with data archiving systems and Graphical User Interfaces must be developed. The effectiveness of the system in improving reduction cell energy efficiency and in reducing anode effect frequency must be quantified.

Despite these hurdles, the research team believes that the potential benefits available from this work merit the investment.

2.2 Intelligent Potroom Advisor (Corrective Action Neural Network)

The objectives of the Intelligent Potroom Advisor were as follows:

1. To create a knowledge repository for collecting and building corporate knowledge for the advisory function of the IPA;
2. To design, specify and construct the hardware and software interfaces between the Intelligent Potroom Operation (IPO) system and the existing distributed control system (DCS) at Century Aluminum;

3. To develop the models required to predict which reduction cells are in danger of entering Off-Peak Modes;
4. To develop the radio frequency operator dispatching and data entry system;
5. To expand the IPA to effectively use the higher level data coming from the CCEM; and
6. To integrate and test the CCEM and IPA functions of the IPO.

2.2.1 Intelligent Potroom Advisor Project Activities

The intended purpose of the Intelligent Potroom Advisor was to identify when a reduction cell is operating on the verge of instability or in a degraded operating state and to coach the operational personnel what action to perform for remediation of that condition in order to return the cell to an optimal operating state. Earlier evidence showed that reduction cells in the Century Aluminum Ravenswood, West Virginia plant operated in a bimodal state of around 89% current efficiency or up to 96% current efficiency. The purpose of the IPA was to continually shift those cells operating in the lower efficiency range into the higher range. The diagram in Figure 12 depicts the logic required of the Intelligent Potroom Advisor.

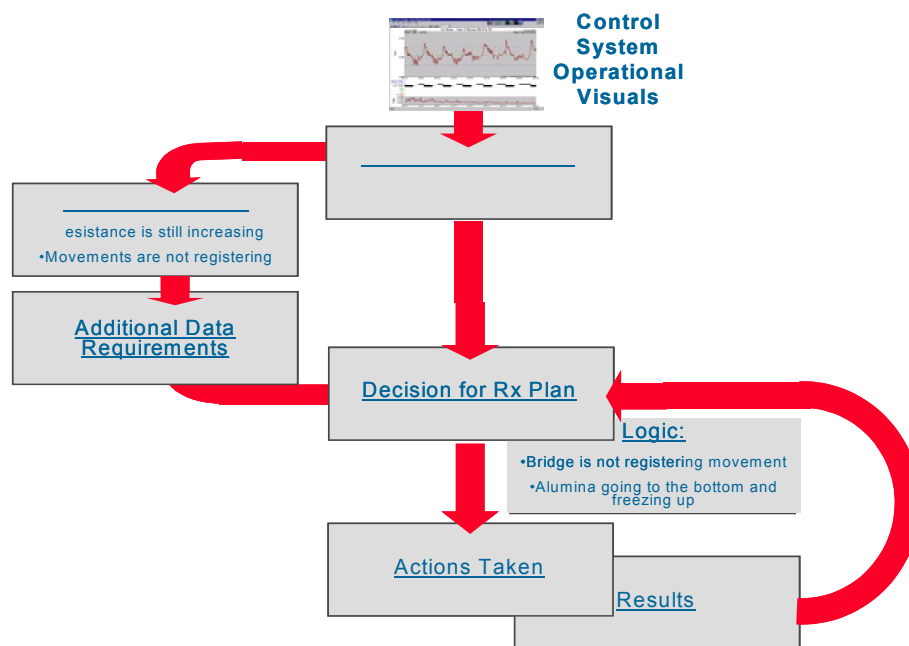


Figure 12: Intelligent Potroom Advisor Logic Diagram

Some of the key benefits anticipated early on in the project are represented here. First, and one of the primary issues that the IPA was anticipated to address was the problem of data overload through “exception” reporting or the ability to filter identified conditions of interest by the ability to focus lean operational resources on those operating conditions that are the most severe. Another benefit expected was the ability to “standardize” remediation strategies for given cell conditions. It was felt that a results-

based approach would provide the ability to develop a corporate knowledge repository of practices. In addition, the ability to link a rationale to specific remediation offered the opportunity to create an on-the-job training aid.

Within the first project year, much of the effort was focused upon developing both the correct system requirements to achieve the desired functionality and the construction of the correct systems infrastructure upon which to build the Intelligent Potroom Advisor. One of the problems in attempting to create the requirements for the systems infrastructure was the fact that the data requirements for the IPA had not been defined and yet the team needed to proceed on building an infrastructure to collect plant data. It was recognized because of this issue that the schema for the data historian and data warehouse would not be correct and may potentially be thrown away.

As described in the Figure 13 below, the potroom generates process control data from a proprietary distributed control system (Kaiser Aluminum's Celtrol DCS), which needed to be placed into a SQL database to allow for access and data manipulation. Additional process related data such as ad hoc measurements for data such as bath temperature or standard operating procedure measurements for items such as cell bath levels were collected at the start of the project on paper and collected via several entry points and eventually entered into a common set of spreadsheets. The replacement system allowed for remote data entry via a browser installed on Personal Digital Assistant (PDA) through a wireless local area network that was installed on line three comprising of two of the eight large potrooms within the plant. Many challenges were encountered in employing a PDA for this purpose due to the immaturity of the technology at that time. Some of these challenges included the awkwardly large form-factor of the devices available and the need to minimize negative impact on staff productivity due to either slow response times for completing each LAN handshake, extensive menus, or overcoming the need to use an operator's hands for entering or controlling the device.

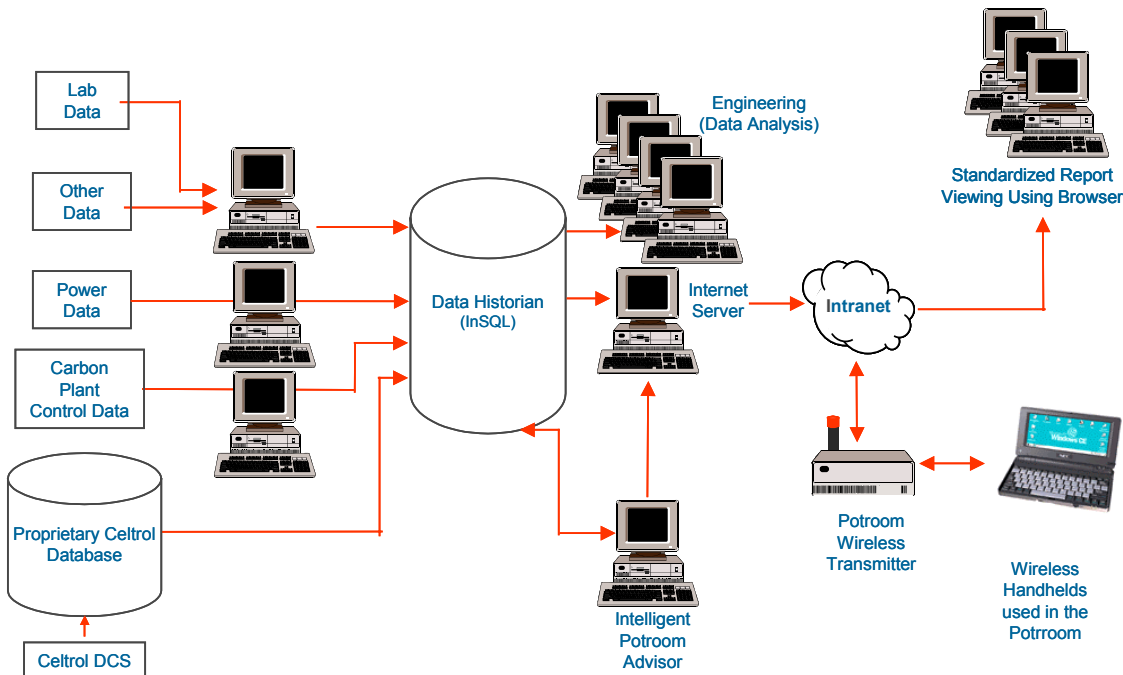


Figure 13: Overall Systems Architecture

IPA Design Considerations

As efforts focused on the development of the IPA, much thought was given to its primary objectives and functionality. The initial target usage for the IPA was to create a heuristic-based system that allowed operational personnel to optimize plant performance based upon standardizing on what they already knew. A secondary strategy and benefit was the ability to create a continuous improvement strategy based upon focus on the next set of problems considered to be “low hanging fruit”. It was felt that by continually eliminating the more severe operating states in this way, a significant plant improvement could be had by standardizing heuristic-based practices. The strategy agreed upon was to utilize a prioritization scheme that would assign a rating for the severity of the degraded operating condition such that operational staff would know where to place resources. A high level concept diagram for the IPA is provided below.

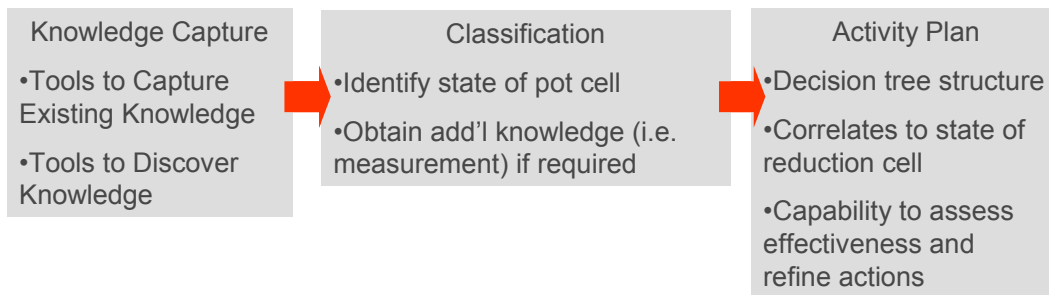


Figure 14: Intelligent Potroom Advisor Conceptual Overview

As the “exceptions” were studied, it was determined that an accurate way to describe a degraded operating condition would be in the form of a data “state” or pattern. It was also determined that a decision tree would be an effective way to catalog various data states corresponding to severity or stages of degradation. By linking agreed upon remediation strategies to that state in the form of a table look-up procedure, the IPA could provide the functionality of a decision-tree based advisor that was desired.

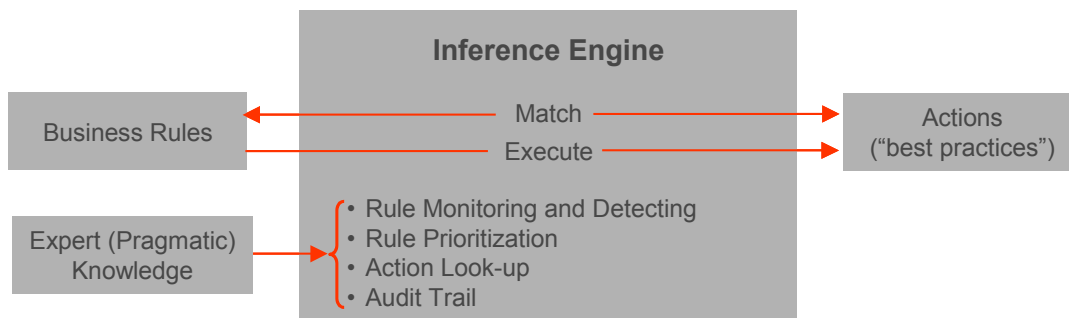


Figure 15: Expert System Functional Diagram

Further, it was recognized that degraded states could be represented in the form of rules (e.g. if there are five anode effects over a 24 hour period with a duration greater than 20 minutes, then set an alarm). Given that rules could be created to filter plant data to identify when a degraded state has occurred, it was determined that an effective application for performing this function was to use an Expert System type design. As

shown in Figure 15, an Expert System constantly fires rules at new operational data to determine if and when a match occurs. If there is a match, then an action is initiated to look-up the agreed upon remediation strategy or “best practice”.

The benefits of an Expert Systems approach are considered to be the following:

- Prioritization of critical events allows staff to focus on prevention or elimination of “low hanging fruit”
- Allows staff to focus limited resources on critical events
- Using agreed upon rules to identify critical events removes the burden on the staff to continuously analyze anomalies eliminating data “overload”
- Rules can include additional data beyond process control data to add context information (e.g. age of the cell, location in room, etc.)
- Provides capability to serve as an engineering search tool
 - Minimizes waiting for re-occurrence from a cell through the ability to search for a similar conditions throughout the plant
 - Allows engineer to more quickly assess known conditions
- System design of incorporating rationale for actions serves as an excellent knowledge repository and corporate training tool

It was recognized that Expert Systems have been around for 15 years or more but are no longer in use. It was recognized that a primary reason for this was the limited life expectancy. Once the business rules were developed and the company “playbook” of best practices was committed to memory by operational personnel, the system was seldom used. Expert Systems in the past could only be updated by a programmer, which made any maintenance, continual improvement, or discovery activity by the subject matter expert too difficult. This issue was dealt with by initially focusing on a rule editor that allowed any non-programmer to compose rules using Boolean logic combined with already defined operators and data fields (see Figure 16).

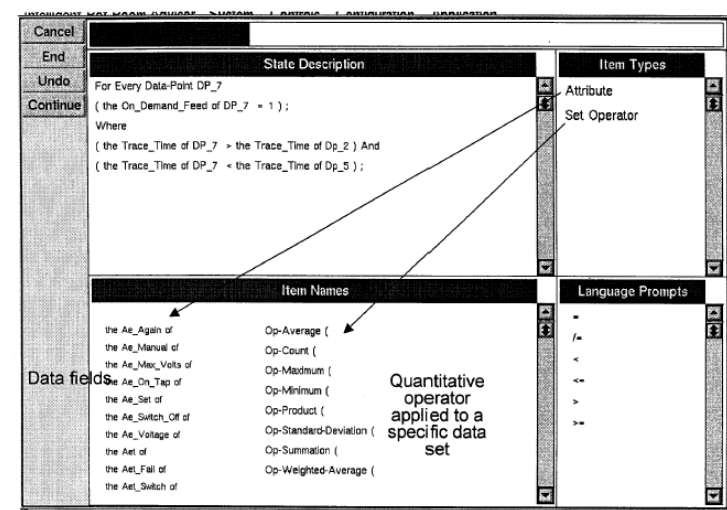


Figure 16: IPA Rule Editor

Another reason for the demise of Expert Systems was the economics. Now that PCs have large amounts of active RAM for dynamic storage of active programs and computing speeds have reached gigahertz ranges, the ability to leverage a single PC to filter data across many different production lines is possible. The leap in computing price performance has greatly increased the potential return-on-investment for these kinds of applications.

IPA Development

The IPA was prototyped using Gensym's G2 Expert System development platform. The Gensym G2 application offered the benefits of an object-oriented development environment based on a LISP-like programming language that allows for the inclusion of time accumulating constructs to accommodate rules that required calculating data changes over time. The object oriented aspect of G2 allowed for visuals to provide "user-friendly" alerts and management tools.

When constructing rules, the philosophy used was to look for some individual symptom or combination of symptoms, which then infers that particular problem or degraded state has occurred. In some cases, the root cause of the problem can be determined, while in other cases identification of the root cause may not be possible. The goal is to then bring the process back to a normal or optimal operating condition by suggesting one or more corrective action strategies. For simple diagnostics, there is often a straightforward mapping between symptoms (events), root causes, and corrective actions as shown in the Figure 17 below.




Figure 17: Rules that directly infer a degraded state

For more complex diagnostic cases, additional measurements may be required or the need to wait for effects of corrective actions before making final conclusions. Events may trigger intermediate conclusions, which may then be used to detect future states or events. These more complex diagnostic cases are shown in the figure below.




Figure 18: Rules that require additional diagnostics to determine correct diagnosis

The goal for the heuristic-based Expert System using these knowledge templates (Figures 17 and 18) to document and gather diagnostic knowledge, was to get as many of the simple cases as possible (i.e., "low hanging fruit"). Then, try to document as many of the more complicated diagnostic cases trying to avoid capturing the extremely rare cases, unless there was value in doing so. Approximately 30 rules were developed in this way. An example is provided in Figure 19 for repeat anode effects.

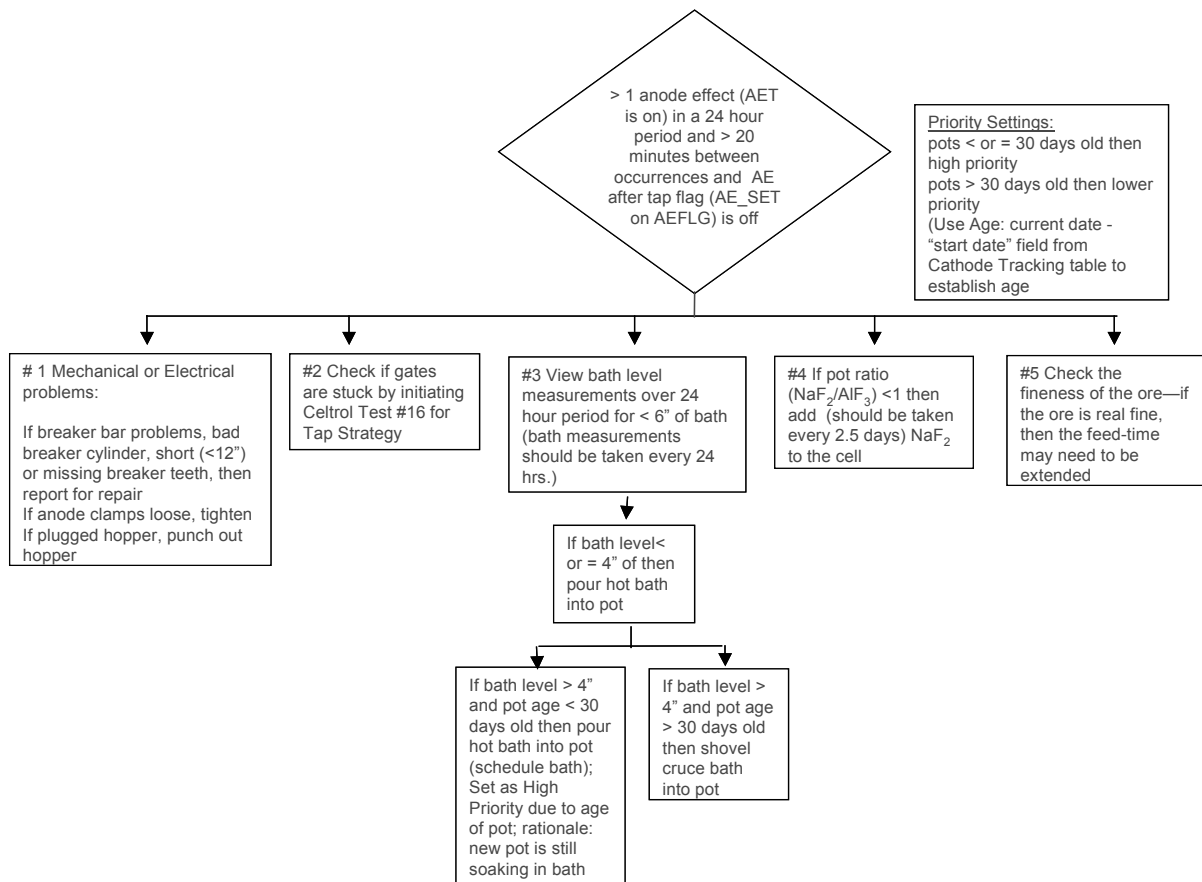


Figure 19: Repeat Anode Effect Rule

User “Buy-in”

Another activity considered critical to project success was obtaining user acceptance to ensure the IPA would be well received in the plant. It is strongly held belief that in making major technology innovations within an organization, perhaps as much as 30% of the success for innovation can be attributed to the technology and the remainder in the ability to address cultural changes that are incurred and user “buy-in” required. This was attempted by soliciting user input throughout the development cycle and by focusing on making the IPA an asset in enhancing the user’s ability to perform their respective job function.

The initial target user of the IPA was the potroom supervisor. The Supervisor’s role was, in part, to assess his room consisting of 84 reduction cells at the beginning and middle of each shift by walking the plant floor and determining the priority of tasks his team would perform. The goal was to gain support for the use of the PDA. The PDA was the target platform to interface with the supervisor in real-time for both data capture as well as provide real-time alerts and guidance. The PDA also provided the utility to investigate and validate process alarms that were suspect.

The project team became consumed with attempting to make the PDA an asset to the supervisor in performing his job in order to obtain buy-in. The project team performed extensive workflow analyses to better understand tasks performed and data and

information transfers. The PDA was turned into a platform where the user could both enter in data in real-time via developed entry forms and obtain a list of IPA alerts to determine which cells were acting in a degraded condition. The supervisor could also access other database applications such as the maintenance system to determine if there was an active work order to repair a mechanical problem. In the end, the size of these early PDAs were too large to be practical given the physical tasks that the supervisor needed to perform while walking the production line.

For this reason, the IPA's purpose was refocused on the creation of a "hot sheet" at the beginning of each shift. The hot sheet identified those exception conditions, which have occurred on the pot line and where that shift should focus their efforts. The exception conditions were typically time series-based trends that the process control system did not support. These rules were initiated automatically by the IPA with a hard copy report output and completed no later than 1 hour and 45 minutes prior to the change in shift (e.g. 6:15 AM for day shift and at 2:15 PM for swing shift). This allowed the potroom supervisor to take the hardcopy report with him when he performed his pre-shift and mid-shift assessment.

Each rule that was created had a priority level assigned to it. The priority allowed the follow-up action for that rule to be pre-assigned to some individual. Priorities were modifiable by the line supervisor who may wish to alter the priority of that task in order to effectively manage labor resources (e.g. a bathing task will be performed at that pot that shift anyway, no need to schedule a special activity to accommodate a requested IPA action). In addition to prioritizing IPA state alerts by importance and by individual, the priorities could also reflect those states that should be placed in the investigative category as part of a technical staff or line staff follow-up and not performed as part of the "hot sheet" of activities for operational staff.

Cultural Challenges

The project team stumbled onto the fact that by getting supervisor and plant experts to contribute to solving problems, the team was running up against a new cultural paradigm that required sharing subject matter expertise. The purpose of the Intelligent Potroom Operation project was in part, to create a method to create corporate knowledge that included the experience of subject matter experts who were rewarded in the past based on this expertise.

Because of the significance of the impact on culture in creating a knowledge-powered enterprise, much thought was given to more formal change management required at Century and that would need to accompany this activity if performed for others when commercializing this solution in the future. Once the impact on company culture was determined, training was started to identify these new roles to the operational staff. A chart was developed to convey the culture and values both currently and how it was expected to affect these areas when helping companies become knowledge-powered enterprises. This is provided in Figure 20. For a more complete explanation of the benefits and impact of the knowledge-powered enterprise see [9].

Personal Values	Cultural Values of Operational Personnel Within the Organization	
	Current	Future
	<i>[Labor viewed as an expense]</i>	<i>[Labor viewed as an asset or "knowledge worker"]</i>
Certainty (def. how life is today)	Continuing with status quo	Technology-based decision making
	Rewards are safety, security, and feedback that you are not the worst performer	Improvement is an expected part of performance
	No plan, do things as they were done before	Must meet a plan
	Autocratic supervisory roles	Leadership/management is telling others what to do
Uncertainty (def. "spice of life"; fears due to inability to control)	Being bottom of ladder in regards to performance in relation to others	
	Not knowing about support for technology	Trying to operate without technology
	Being respected for making changes	Not making changes for achieving goals
Significance (def. ability to contribute)	Permission to make decisions	Ability to make goals
	Experience	Ability to apply technology
	Treating knowledge as an art	Ability to create sustained improvement in a methodical manner
	Position/rank in the organization	Ability to make achievements
Connectivity (def. to belong)	Hierarchal	Team
	Experience	Leadership
Growth (def. ability to increase knowledge)	Years of experience	Technical understanding; ability to apply technology
	Functional role in the organization	Leadership

Figure 20: Cultural Changes Within the Knowledge-Powered Enterprise

Predictive IPA

The final goal of the Intelligent Potroom Operation project was to go beyond the "early intervention" of OPMs to an operational capability of "predictive avoidance" in order to maximize plant performance. It was recognized by the project team that the knowledge developed by experienced personnel was the ability to recognize existing degraded operating states and react. This final goal required identifying those operating states that occur prior to a degraded operating condition such that operational personnel could be proactive in avoiding this condition altogether--knowledge that was never developed.

Initially, it was thought that a neural network could provide this capability, which is why this portion of the project was called the Corrective Action Neural Network. The problem determined with this approach is that a neural network can be trained to recognize complex state changes but provides no visibility as to what the change represents. Neural networks must be treated, therefore, as "black boxes" due to this lack of visibility for these transformations. Constant retraining requirements for the neural network model would be required due to multitude of variances for each reduction cell, which was not considered practical. Visibility is also required for classification of the degraded operating "states" in order to correlate results into a decision tree.

The IPA with the expanded predictive rule discovery capability was envisioned to look like Figure 21.

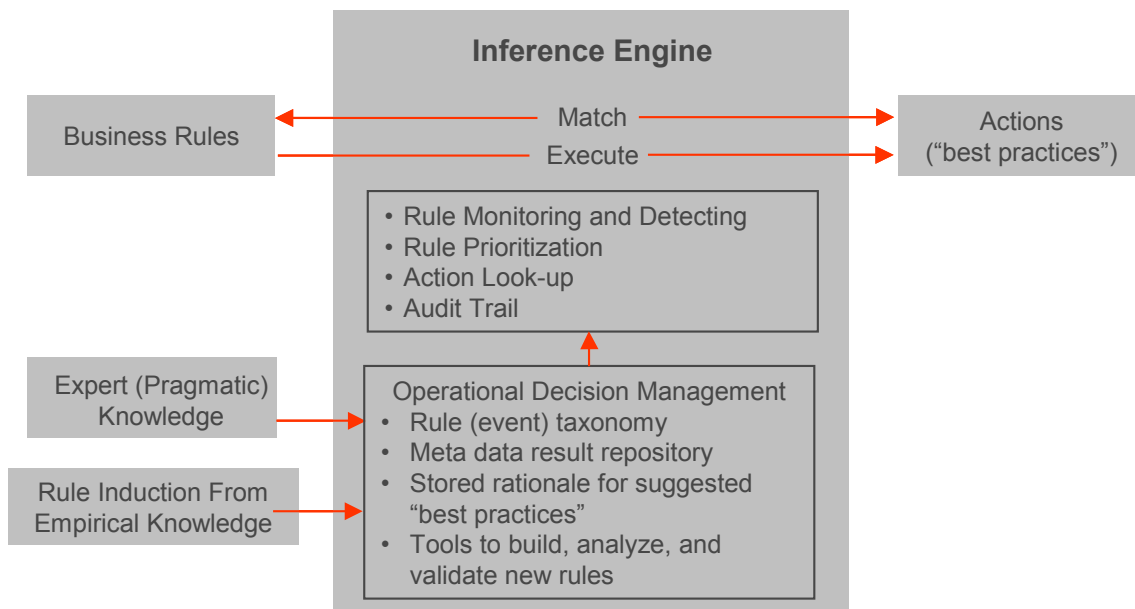


Figure 21: IPA with Inductive Knowledge Building Engine

During the third year of the project, in place of a neural network, advanced data mining technologies were applied to enhance existing alarms and to develop new predictive alarms with variable look-ahead time windows. Chi-Squared Automated Interaction Detection (CHAID) was the primary rule induction technique employed. It allowed the project team to create probabilistic decision trees with node splits that discriminate between various quantitative levels or classifications for each dependent variable. An example of a decision tree for a continuous dependent variable is provided in Figure 22.

In a decision tree, each “node” or “split” in the tree represents one possible probabilistic decision rule in a potentially multilevel and additive set of rules. Decision trees can be re-expressed as a set of probabilistic If-Then rules which is conducive to the Expert System rule design needed for the IPA. For example, there may be a 90% chance of an anode effect in the next 30 minutes if $w = \text{true}$, $x > 120$, $15 < y < 20$. Alternatively, the mean time to low pot level may be under 20 minutes when $c = 3$, $d < 30$, $e > 5$, and $f = \text{false}$. The independent variable distribution mined to find the rule is called the “split variable”. The statistics of each node provide the average or categorical value of the dependent variable as it becomes more finely discriminated at deeper split levels.

The level of predictive discrimination, i.e., the ability to accurately differentiate impending states, depends on the split level, sample size, derived variables used, the underlying phenomena, and noise present in the system as well as the tuning of the techniques. The CHAID trees were pruned to optimize statistical accuracy and robustness. Although finer discrimination may occur at higher split levels, this also yields greater complexity which may become harder for experts to follow, test, or verify. The model development process must receive iterative feedback from the subject matter experts in order to produce trees and rules that are more intuitive and actionable.

CHAID trees are also usually *degenerative*, i.e., multiple overlapping trees can be produced that achieve similar results but with different component variables and/or splits. This creates additional challenges and opportunities. The challenge is basically the greater time and complexity required to generate, present, categorize, and examine multiple rule trees. However, this can also provide wider choices and opportunities to search for and select more intuitive and actionable rules. It is therefore incumbent upon the modeler to categorize, prioritize, and present the results in a manner that best facilitates rule selection, a very time consuming pre-mining activity.

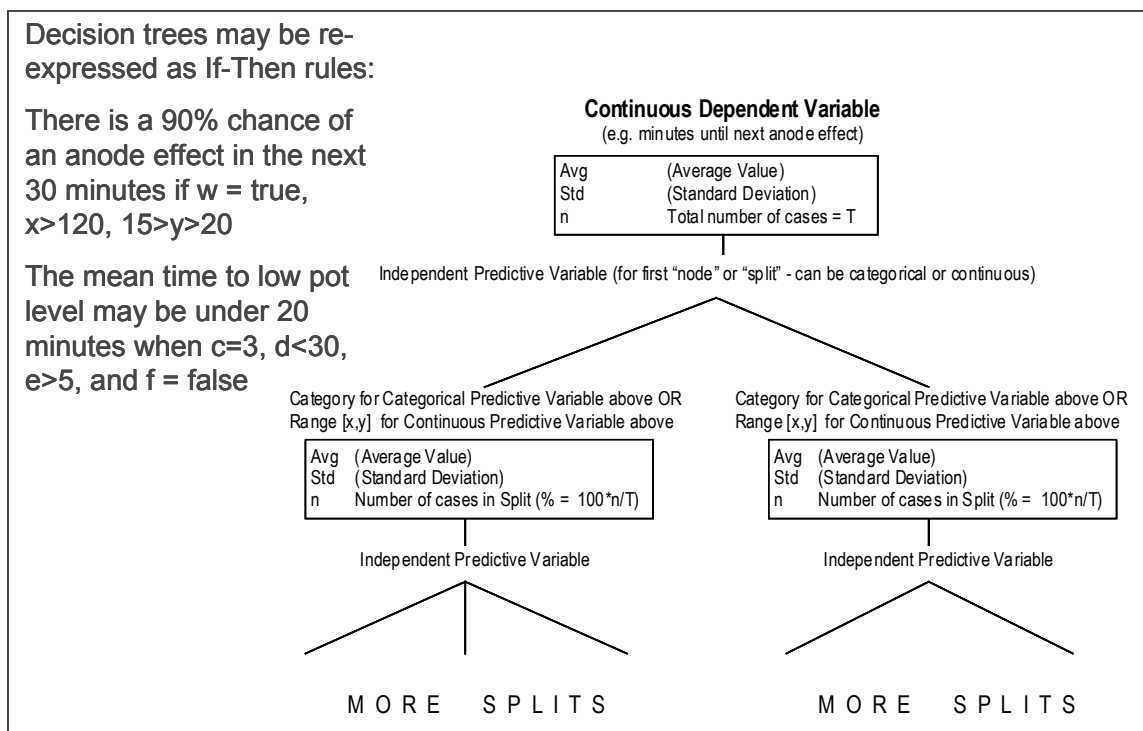


Figure 22: Schematic for CHAID Rule Induction Trees for Continuous Dependent Variables

Once a CHAID tree was created, it was then broken into branches and prioritized for subject matter expert analysis using a utility function. The utility function was a simple formula of sample size and discrimination power so that analysis could focus on those elements that yielded the most benefit. Each major tree branch or rule set was annotated by the expert to indicate:

1. Level of intuitive relevance
2. Usefulness
3. Actionability with prescribed or candidate reactions prioritized in a stepwise fashion
4. Relation to existing alarms, when applicable
5. Further information or refinement needed (if possible).

This expert feedback was used to:

1. Iteratively, enhance rule induction
2. Build an actionable look-up table
3. Implement new alarms or enhance existing ones.

The first set of CHAID trees were developed using 30 – 40 predictive input variables and at least one of four sets of dependent output variables for predicting anode effects, noise level, low bath level, and resistance error. From this starting point, the derived independent and dependent variables were iteratively enhanced and expanded based on statistical results and expert feedback.

One difficulty experienced during this process was the amount of data preparation time required to create the data set for CHAID analyses. Initially 2.5 months of data was collected which consisted of six gigabytes of analog and discrete bits of data used by the process control system for data flags. Once the data flags were expanded into fields within the SQL database, the file size increased by a factor of ten. Next, data needed to be synchronized for data granularity and normalized with time lags and by control protocols. Latter CHAID trees required data dictionaries to support 85 derived variables used to enrich the mined data set. For example, a derived variable may be a time horizon variable indicating 10 minutes before an anode effect, 15 minutes before an anode effect, and so on to determine which split in the CHAID tree had the most statistical significance. Every time a new derived variable was identified, the data pre-processing needed to be repeated prior to CHAID analyses. This effort minimized the amount of times the CHAID analyses could be performed to look for opportunities to identify more statistically significant rules.

Data Latency

Another issue recognized was the latency built into the data extraction and IPA rule firing process. The IPA application was set to re-fire each of the stored rules once per hour to accommodate the 25 minutes required for line three alone. The built-in data latency for the SQL data historian was up to 13 minutes as described in Figure 22 below. Latest set of rules included calculations for bridge raises/bridge lowers, time since last anode effect, and noise control in last 15 minutes. Viewing these variables once per hour combined with the method that the alarms were broadcast to operational personnel minimized the ability by operational staff to respond. The method utilized in the project did allow for the ability to induce rules that were predictive. However, the latency minimized the ability to utilize these rules in any beneficial manner.



Figure 23: Inherent data latency issues for data capture

1.1.1 IPA Results

1. The use of the Intelligent Potroom repository as a knowledge repository for standardizing corporate knowledge and practices has proven to be very successful. Prior to this project, the plant was successful in creating lists of practices and training programs that were time consuming to develop, difficult to maintain, and more difficult to ensure that practices were continually applied in a consistent manner. The IPA created a structure within which, degraded states could be articulated in the form of time-series based rules Energy, Office of Industrial Technology; April 2003.

and monitored to determine if these practices were consistently applied utilizing the “best practices” agreed upon.

The utility of the IPA in creating a low priority rule for purposes of determining significance of the event became a big advantage when compared to the alternative tools available. Without the IPA, the engineer would have to drill down through masses of data to determine whether an anomaly was merely noise or represented a significant statistical population worthy of institutionalizing as part of the set of standard operating practices. In order to take advantage of this utility, the engineer needed to become comfortable identifying opportunities in terms of statistical significance rather than in terms of absolutes. Acceptance of building operational practices in these terms was another challenge that needed to be addressed.

The strategy of creating the “hot sheet” for assisting the line supervisor to prioritize shift resources also proved to be very beneficial. The process control system had the ability to identify problem conditions for a given moment in time, which previously forced operational personnel into a reactive mode with the subjective interpretation for prioritizing shift resources left up to the line supervisor via his pre-shift line assessment. By the end of the project, line supervisors from the rest of the plant were interested in obtaining use of the IPA because of the “hot sheet” and degraded state validation capability.

The IPA was integrated with the existing distributed control system such that data could be extracted from the short-term proprietary historian into a longer-term non-proprietary environment for mining purposes. The latency issues minimized the benefit of new knowledge discovered. The ability to validate new knowledge also needed a more intimate coupling with the control system such that remediation strategies could be more effectively tested in a more expedited manner. In addition, the very time-consuming and tedious effort of pre-processing the extracted data for purposes of mining made the effort for improving accuracy and the ability to obtain a continuously updated result impractical.

The CCEM was never integrated with the IPA. It was recognized that an effective state estimator capability could significantly improve the quality of data that was being mined in order to generate more accurate predictive rules. Both the CCEM and IPA activities generated enough challenges that the ability to integrate the two was never realized.

2.2.3 IPA Summary and Conclusions

The primary aluminum industry is representative of many of the industries in the Office of Industrial Technologies portfolio. It is a very mature industry and as such, most of the significant performance gains have already been realized. The Intelligent Potroom Operation project and specifically, the experience derived through the development of the Intelligent Potroom Advisor, clarified that the goal of this project was to create a method upon which mature companies, like Century Aluminum could achieve continuous incremental gain. What was determined was that subject matter experts like those at Century, are very competent at identifying simple trends and process abnormalities. The challenge is that the next level of gains must come from the ability to identify those degraded states that are comprised of small statistical populations that are often very complex variable relationships and are cognitively beyond the capability of most individuals. It is the aggregated affect of continuously identifying these small complex

statistical populations and proactively addressing these degraded states where the next level of gain will be realized.

Defining appropriate technology to accommodate the methods needed to identify these gains also identified a need to change both the culture and practices that are needed to support the “knowledge enterprise” paradigm. The ability of engineers to institute new practices in terms of have a high probability of accuracy versus working in terms of absolutes is one obstacle that needed to be addressed. The change in management practices required for the knowledge enterprise both from a standpoint of how individuals are compensated to share and build knowledge to entrusting individuals at lower levels of the organization to employ this new knowledge is critical. The cultural challenges may be even more difficult to implement than the technical ones.

In performing literature searches and determining how others have attacked this problem, a summary of technologies was formulated. Figure 24 summarizes information technology solutions by how they contribute towards performance optimization. It was recognized early on in the project, that the heuristic-based form of the IPA offered similar capabilities to what several other vendors or researchers have termed Abnormal Situation Management applications. These tend to be Expert System type applications that apply time-series based rules, which then trigger a look-up to a corresponding action.

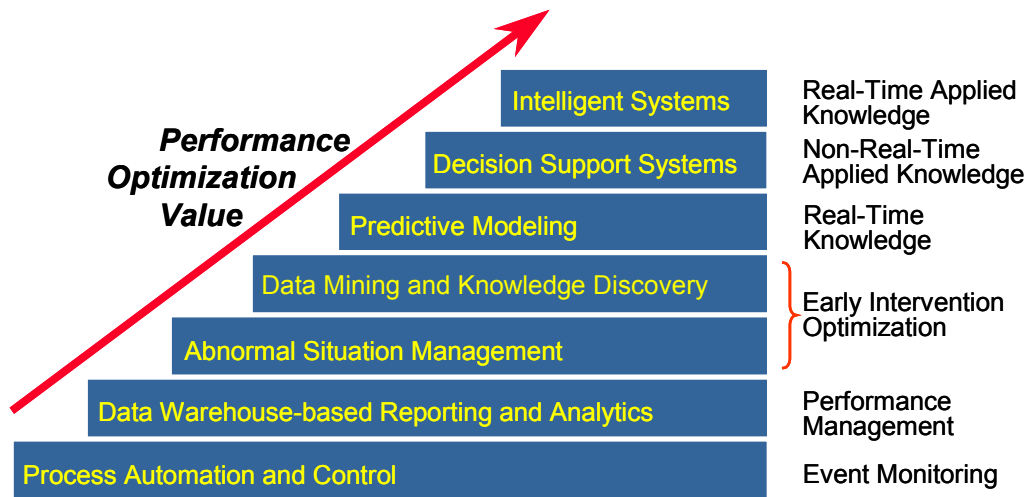


Figure 24: The Performance Optimization Value (POV) diagram

The difference between the IPA and these systems is the underlying concept of process knowledge. The building block of process knowledge is the concept of “behavior”. A behavior is a time-series based set (chronology) of data patterns or events. The distinction for behavior versus merely a time-series set of data patterns is the additional information that corresponds to behaviors in regards to context. Context refers to key indicators that allow for more accurate classification of behaviors such that when they occur again, we will be able to recognize what they mean. In an aluminum plant, a key indicator may be variables such as the age of the cell, its location within the plant, the mix of virgin and recycled alumina ore used in the cell, or combinations of these.

The process control system used at Century, like many other systems that are specific to a process, incorporates a first principles model that is applied to cells in the plant. It was evident that our challenge was to identify each unique instance where the more general first principle-based control system failed to perform. By building a catalog of degraded states and determining those significant variables that distinguish the context of those cells, we can customize the first principles model to accommodate the unique cell characteristics so that we can control them better. This cataloguing process still does not answer the question of why these behaviors occur, it is merely a recognition that they do.

Once the project team more fully grasped an understanding of process behaviors, parallels to other knowledge building initiatives became apparent. The understanding that an *Intelligent System* was a mechanism to create a recognition system for behaviors so that when it is recognized, we will know what to do about it the next time it is seen. NIST Senior Fellow James Albus and co-author Alexander Meystel define it in this way [“Intelligent Systems, Architecture, Design, and Control”, Willey & Sons, 2002]:

“Intelligence is the ability of a system to act appropriately in an uncertain environment, where an appropriate action is that which increases the probability of success, and success is the achievement of behavioral subgoals that supports the system’s ultimate goal.”

Albus and Meystel also indicated that the correct structure for an *Intelligent System* has never been fully developed.

There are many other realizations that became evident once the concept and requirements for an intelligent system were understood. These concepts also assisted in a better understanding of the correct system requirements, functionality, and design that an intelligent system requires in order to add an adaptive learning feature. These have all been included in [10] with the goal of using additional funding to automate the intelligence building (knowledge and action planning) processes to make systems adaptable and feasible for adoption and to turn the concept of the *Intelligent System* into a reality. Many of these concepts are defined in [9] and have been more fully expanded upon and incorporated into white papers and applied to a host of applications since the completion of this project.

3 Products

The following is a list publications resulting from this work

1. Biedler, Philip, *Modeling of an Aluminum Reduction Cell for the Development of a State Estimator*, Ph.D. Dissertation, West Virginia University, 2003.
http://etd.wvu.edu/ETDS/E2918/Biedler_P_dissertation.pdf
2. Dai, Congxia, *An Advanced Data Acquisition System and Noise Analysis on the Aluminum Reduction Process*, M.S.M.E. Thesis, West Virginia University, 2003.
http://etd.wvu.edu/ETDS/E2850/Dai_Congxia_Thesis.pdf
3. Banta, Larry, Dai, C., and Biedler, P., "Noise Classification in the Aluminum Reduction Process", *Light Metals, 2003: Proceedings of the TMS Annual Meeting*, March 2-6, 2003, San Diego, CA. pp 431-436.
4. Biedler, Philip and Banta, L., "Analysis and Correction of Heat Balance Issues in Aluminum Reduction Cells", *Light Metals, 2003: Proceedings of the TMS Annual Meeting*, March 2-6, 2003, San Diego, CA. pp 441-448.
5. Banta, Larry, Biedler, P., Dai, C., Love, R., Tommey, C. and Berkow, J. "Decomposition of Aluminum Cell Voltage Signals," *Light Metals, 2002: Proceedings of the TMS Annual Meeting*, 17-20 Feb, 2002, Seattle, WA, pp 365-370.
6. Biedler, Philip, Banta, L., Dai, C., Love, R., Tommey, C. and Berkow, J. "Development of a State Observer for an Aluminum Reduction Cell", *Light Metals, 2002: Proceedings of the TMS Annual Meeting*, , 17-20 Feb, 2002, Seattle, WA, pp 1091-1098.
7. Banta, Larry, Biedler, Philip and Dai, Congxia, "Processs signal classification for aluminum cells", *Proceedings of the SPIE International Symposium on Intelligent Systems and Advanced Manufacturing*, 28 Oct – 2 Nov, 2001, Boston MA. SPIE paper #4565-23.
8. Biedler, Philip, Banta, Larry and Dai, Congxia, "State Observer for the aluminum reduction process", *Proceedings of the SPIE International Symposium on Intelligent Systems and Advanced Manufacturing*, 28 Oct – 2 Nov, 2001, Boston MA. SPIE paper #4565-10.
9. Berkow, Jan, "Maximizing Operational Performance Through the Use Of Intelligent Systems", May 2003, White Paper submitted in behalf of *Sensors, Controls, and Automation Crosscutting Technologies Roadmap*, US Department of Energy, Office of Industrial Technology, April 2003.
10. Berkow, Jan, Banta, Larry, Oates, Tim, "Intelligent Manufacturing Systems for Plant Performance Optimization", proposal submitted in response to *Sensors, Controls, and Automation Technologies solicitation DE-PS07-03ID14442*, US Department of Energy, Office of Industrial Technology; April 2003.