

FC07-94ID/332)

**Feedback Control of a Cupola –
Concepts and Experimental Results**

CONF-981006--1

Kevin L. Moore (corresponding author)
Measurement and Control Eng. Research Center
College of Engineering Idaho State University
Pocatello, ID 83209-8060
moorek@isu.edu

Mohamed A. Abdelrahman
Electrical and Computer Engineering Department
Tennessee Technological University
Cookeville, TN 38505-0001
maa0303@tnstate.edu

Eric Larsen
Lockheed Martin Idaho Technologies Company
Idaho National Eng. and Env. Laboratory
P.O. Box 1625, Idaho Falls, ID 83415

Paul King
U. S. Department of Energy
Albany Research Center
1450 Queen Ave SW, Albany, OR 97321

Denis Clark
Lockheed Martin Idaho Technologies Company
Idaho National Eng. and Env. Laboratory
P.O. Box 1625, Idaho Falls, ID 83415

**1998 AFS Cupola Conference
Cincinnati, OH**

October 1998

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

ph

MASTER

DISCLAIMER

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

DISCLAIMER

**Portions of this document may be illegible
electronic image products. Images are
produced from the best available original
document.**

Feedback Control of a Cupola – Concepts and Experimental Results

Abstract

In this paper we present some final results from a research project focused on introducing automatic control to the operation of cupola iron furnaces. The main aim of this research is to improve the operational efficiency and performance of the cupola furnace, an important foundry process used to melt iron. Previous papers have described the development of appropriate control system architectures for the cupola. These results are summarized. Then we describe the experimental results obtained with the U.S. Department of Energy Albany Research Center's research cupola. First, experimental data is used to calibrate the model, which is taken as a first-order multivariable system with time delay. Then relative gain analysis is used to select loop pairings to be used in a multi-loop controller. The resulting controller pairs meltrate with blast volume, iron temperature with oxygen addition, and carbon composition with percent coke. Special (nonlinear) filters are used to compute meltrate from actual scale readings of the amount of iron produced and to smooth the temperature measurement. The temperature and meltrate loops use single-loop PI control. The composition loop uses a Smith predictor to discount the deadtime associated with mass transport through the furnace. Experimental results validate the conceptual controller design and provide proof-of-concept of the idea of controlling a foundry cupola. Future research directions are discussed, including the concept of an integrated, intelligent industrial process controller, or I^3PC .

Feedback Control of a Cupola – Concepts and Experimental Results

Introduction

Overview

Key operational goals in running a cupola are to keep the iron properties within a prescribed range and, in some cases, to maintain a desired production rate. These goals are usually accomplished through judicious choice of the manipulated process variables, notably the blast properties (rate, temperature, and oxygen enrichment) and the charge composition (including percent coke, iron-to-steel ratio, and alloys). Cupola operation has always relied on the experience of the operator in deciding which process parameters to adjust to obtain the desired molten iron properties. Recently, two research projects funded by the Department of Energy (DOE) and the American Foundrymen's Society (AFS) have demonstrated the feasibility of using feedback control technology to help achieve better operation of the cupola furnace with less dependence on the experience and skills of a single operator. One thrust has been on improved understanding of the cupola process via a modeling effort. In that project a verified, first-principles, *analytical model* of the steady-state cupola was developed by Vladimir Stanek, with the assistance of the AFS Cupola Committee (described in [1]). The second project considered the use of process control to automatically select process inputs to force the cupola outputs to desired values [2].

In this paper we describe some results from the automatic control project. The key results are summarized. Then we describe the experimental results obtained with the U.S. Department of Energy (DOE) Albany Research Center's research cupola. First, experimental data is used to calibrate the model, which is taken as a first-order multivariable system with time delay. Then relative gain analysis is used to select loop pairings to be used in a multi-loop controller. The resulting controller

pairs melt rate with blast volume, iron temperature with oxygen addition, and carbon composition with percent coke. Special (nonlinear) filters are used to compute melt rate from actual scale readings of the amount of iron produced and to smooth the temperature measurement. The temperature and melt rate loops use single-loop PI control. The composition loop uses a Smith predictor to discount the deadtime associated with mass transport through the furnace. Experimental results validate the conceptual controller design and provide proof-of-concept of the idea of controlling a foundry cupola. Future research directions are discussed, including the concept of an integrated, intelligent industrial process controller, or I^3PC .

Project Team and Acknowledgements

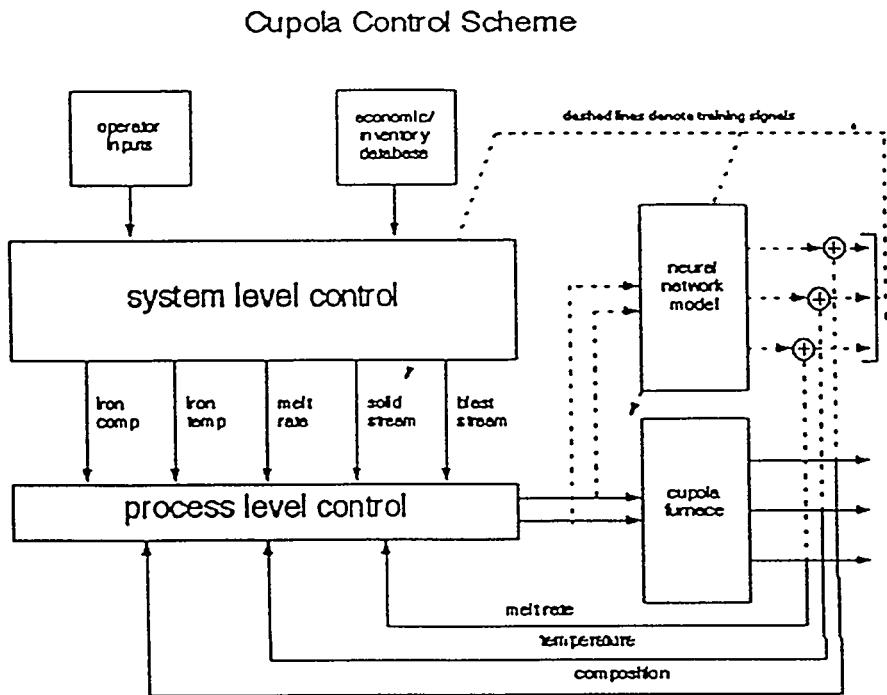
The project team included researchers from the Idaho National Engineering and Environmental Laboratory (INEEL), a DOE national lab, Idaho State University (ISU), researchers at the DOE Albany Research Center (ALRC), and an industrial oversight committee sponsored by AFS. An experimental research cupola (an eighteen-inch diameter furnace with a nominal melt rate of approximately one ton/hour) was designed, constructed, and instrumented at ALRC [3]. INEEL researchers developed a LabView-based computer instrument panel for data acquisition and control interfacing [4,5]. INEEL also developed a neural network model of the steady-state cupola [5,6]. ISU developed control concepts for the furnace [7,8,9]. The controller architecture has a hierarchical structure that includes system-level coordination for optimization and setpoint selection and process-level control for setpoint regulation [10].

This work was funded through Department of Energy Contract No. DE-FC07-94ID13321.

Summary of Results

The technical goal of the project was to develop and demonstrate concepts for applying intelligent control techniques, such as neural networks, to improve the operation of cupola

furnaces. Figure 1 shows a block diagram of the control architecture that was proposed in the early stages of the intelligent cupola control project. A key feature is that the controller has a two-level, hierarchical architecture. The *System-Level Control* block is intended to act as a supervisory controller to define operating points. The original ideas were that this would be done by using operator input, system databases (such as inventory, material costs, and cupola parameters), and artificial neural network models to perform setpoint selection for economic optimization. The *Process-Level Control* block acts as a feedback controller to maintain the cupola around the prescribed operating points. The original concept was that this would be done by using multivariable feedback control with predictive techniques (possibly based on neural network models) to deal with time delays associated with the charge movement down the cupola stack. A third aspect of the proposed architecture was that it included an *Adaptive/Learning Control* capability (indicated by the dashed lines shown in Figure 1) whereby the neural network model would be updated during operation to respond to changes in the process or in the operating objectives. The goal of the intelligent cupola control project was to develop appropriate algorithms for each of the blocks in the controller and to demonstrate the effectiveness of the control concepts on a real cupola.



5/31/95

Figure 1: Original intelligent cupola control architecture.

As the intelligent cupola control project has progressed, a number of significant accomplishments have been achieved. These are described against the backdrop of Figure 1:

1. A fully-instrumented 18-inch *Experimental Research Cupola* was designed and constructed at the Albany Research Center (ALRC), a DOE-operated facility, for verification and testing of models and control algorithms. This facility uses a PC-based, stand-alone data acquisition and control system developed using LabView by the Idaho National Engineering and Environmental Laboratory (INEEL).
2. A *Neural Network Model* of the cupola and *Techniques for Training* the neural network were developed at the INEEL, a DOE national lab operated by Lockheed Martin Idaho

Technologies Company (LMITCO). The neural network model gives steady-state relationships between selected cupola inputs and outputs, as the original AFS model does, but provides computational advantages that facilitate optimization and setpoint selection. It should be noted that the original project proposal did not explicitly propose a neural network model as a deliverable, but that the neural network model was developed as a logical part of the system-level control concept.

3. A *System-Level Control* strategy was developed by Idaho State University (ISU) and INEEL that uses the neural network model for economic optimization and "reverse" operation (for setpoint selection, for example). This methodology was demonstrated with software developed by INEEL. When combined with the AFS model, these results make it possible for cupola operators to determine lowest cost charging schedules and to determine proper operational setpoints. It is important to point out that these results can be used independently without necessarily trying to use them as part of a complete intelligent control system. Work is continuing to develop a mechanism to make the neural network, and its optimization capabilities, available to users as part of the commercial release of the AFS model.
4. *Proof-of-concept experiments at ALRC for Process-Level Control* algorithms have demonstrated the feasibility of using automatic control in cupola operations. A three-input, three-output multi-loop controller with dead-time compensation was successfully used to simultaneously maintain iron composition, temperature, and melt rate at desired setpoints by varying the percent coke, oxygen enrichment, and blast rate, respectively. As a part of these experiments, an empirical, lumped parameter, *transient model* of the cupola was also determined. These results show that it is possible to use feedback control to drive the cupola to a desired operating point, without "guesswork." Note that, although original controller

concepts suggested the use of neural networks for process-level control, subsequent work on the project indicated that, at least for the ALRC cupola, this is not necessary and more conventional techniques are adequate.

Feedback Control of the Cupola

The cupola furnace is one of the primary foundry processes used to melt iron. A cupola is usually constructed as a water-cooled vertical cylinder. The cupola is charged at the top with fuel (usually coke) and metal (pig iron, scrap metal, cast iron scrap, foundry return scrap, and ferro-alloys). Air is injected into the cupola through tuyeres located near the bottom of the furnace, above the molten iron. The blast air is often heated and enriched with oxygen. As the coke is consumed the charge drops and melts, producing a continuous flow of molten iron (large cupolas may produce up to 100 tons/hour of molten iron). As noted, typical operational goals are to keep properties such as composition and temperature within a prescribed range by varying manipulated variables such as the blast properties and the charge composition. The automatic manipulation of these variables for the purpose of regulating the cupola output variables is not commonly done in foundry operations. (Note: we do not mean to imply that control system technology is not used in foundries. Indeed, measurement and control instrumentation is essential to ensure that desired blast properties such as percent oxygen enrichment and blast temperature are achieved. What we refer to here is the automatic manipulation of the primary cupola inputs based on the observed variation of the primary cupola outputs of interest, such as temperature and composition. We call this process-level control as opposed to instrument-level control. Instrument-level control is also an essential part of the overall picture, but is in fact commonly found in foundries.)

Controller Development

Modeling for Controller Design: Based on preliminary analysis of the cupola process, information gathered from industrial cupola operators, and constraints placed by the actual instrumentation capabilities, manipulated and controlled variables were chosen for the ALRC cupola in the following way:

1. *Manipulated variables (process inputs):*

- (a) Percent coke relative to weight of metal (%Coke)
- (b) Oxygen enrichment (O_2)
- (c) Blast rate (B_R)

2. *Controlled variables (process outputs):*

- (a) Iron carbon content (%C)
- (b) Iron temperature (T_{FE})
- (c) Melt rate (MR)

Although other choices of inputs and outputs could be made, such as various types of metal input streams, concentrations of other elements such as S , Si , or Mn , or off-gas measurements, a decision was made to limit the scope of the proof-of-concept experiments to the fundamental signals of interest as listed above.

Next, data from a number of transient response tests was used to build a transient model of the cupola from an input-output perspective. A total of ten transient/control tests were conducted starting from the same nominal operating point (blast rate of 300 scfm, no oxygen enrichment, and 12 %Coke). From these tests a transient model was derived. The model is a first-order multivariable system with time-delay (expressed here using Laplace transforms, in terms of deviations from nominal):

$$\begin{bmatrix} \Delta\%C \\ \Delta T_{FE} \\ \Delta M_R \end{bmatrix} = \begin{bmatrix} \frac{0.04e^{-Ts}}{300s+1} & \frac{0.03}{300s+1} & 0 \\ \frac{4e^{-Ts}}{300s+1} & \frac{12}{300s+1} & 0 \\ \frac{0.04e^{-Ts}}{300s+1} & \frac{2}{300s+1} & 0.08 \\ \frac{300s+1}{300s+1} & \frac{300s+1}{300s+1} & \frac{60s+1}{60s+1} \end{bmatrix} \begin{bmatrix} \Delta\%Coke \\ \Delta O_2 \\ \Delta B_R \end{bmatrix}$$

The time delay T was determined to be one hour, or 3600 seconds. Notice that this is much longer than the 300 second time constant seen in most entries in the matrix.

Controller Design: The original long-term concept for the cupola controller architecture, shown in Figure 1, was a hierarchical structure. It included system-level coordination for optimization and setpoint selection, using the INEEL neural network, and process-level control for setpoint regulation, using conventional controllers with dead-time compensation to account for the transport delay associated with the solid charge stream. This has been described in various conference papers and in the annual technical reports prepared by ISU and will not be discussed further here. However, it became clear early in the project that we must “walk before we run” and thus the proof-of-concept experiments were limited to process-level control. Also, preliminary information about the cupola and its behavior suggested the need for multivariable controllers, with feedforward components and advanced robust control algorithms. These concepts are described in a conference paper [8]. However, after the experimental transient model was available a decision was made to use a simpler overall approach during the implementation. Using a technique called relative gain analysis, a multi-loop control strategy was developed.

To implement a multi-loop controller it is necessary to decide which input should be paired with which outputs. Although we might observe that percent coke is an obvious candidate for pairing with the percent carbon in the iron, it is useful to consider the issue more systematically. A tool commonly used in the process control community is the so-called relative gain analysis,

which is based on the steady-state gain matrix, which we denote K_{ss} . The relative gain matrix, R , is computed as:

$$R = K_{ss} * \left(K_{ss}^{-1} \right)^T$$

where “*” denotes entry-by-entry multiplication. The entries of the relative gain array matrix provide a measure of the effect of interaction in a multi-loop control system. It can be shown that one should use loop pairings that have relative gain array entries that are positive and close to unity.

For the ALRC cupola the steady-state gain matrix is defined by:

$$\begin{bmatrix} \Delta\%C \\ \Delta T_{FE} \\ \Delta MR \end{bmatrix} = \begin{bmatrix} 0.04 & 0.03 & 0 \\ 4 & 12 & 0 \\ -2 & 2 & .08 \end{bmatrix} \begin{bmatrix} \Delta CMR \\ \Delta O_2 \\ \Delta B_R \end{bmatrix}$$

From this we can compute the relative gain array matrix:

$$R = \begin{bmatrix} 1.3 & -.3 & 0 \\ -.3 & 1.3 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

This matrix makes it clear that, from the perspective of loop gain interactions, the following loop pairings should be used:

1. %Carbon controlled using percent coke.
2. Temperature controlled using oxygen.
3. Meltrate controlled using blast.

The resulting controller architecture is shown in Figure 2.

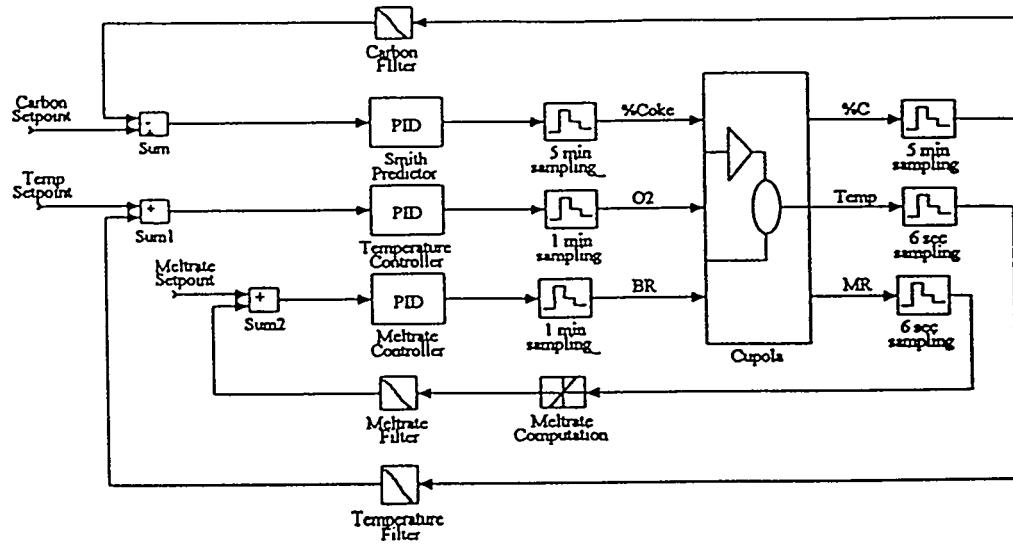


Figure 2: Control system architecture.

As can be seen in Figure 2, the controller pairs percent carbon with %Coke, temperature with oxygen, and meltrate with blast. All controllers in the system were standard proportional-integral-derivative (PID) algorithms with the derivative off. The carbon controller uses a PID algorithm combined with a Smith predictor to compensate for the delay time associated with the movement of the charge down the cupola. All controllers shown in the figure were implemented digitally, using LabView. Four points should be noted:

1. The control system is actually a cascade controller, where the controllers shown in Figure 2 are used to drive the setpoints for the instrument-level controllers. The one exception to this is the %Coke. This loop was implemented in a semi-automatic fashion as follows. The controller took measurements from the data acquisition portion of the LabView system and computed the appropriate changes to the percent coke. These changes were displayed on the monitor and were then relayed via two-way radio to personnel charging the cupola.

2. Due to hardware and data acquisition constraints there were a number of different sampling times in the actual implementation. These are also indicated in Figure 2.
3. All of the key output signals suffered from noise problems. As a result, it was necessary to use various filters in the control system:
 - a) *Carbon*: The percent carbon (and other chemical concentrations) in the iron was measured every five to ten minutes using a thermal arrest unit. This data was recorded and transmitted to the control system computer automatically. This measurement was the most “stable” of all our measurements, possibly due to the time between measurements. Consequently, only a simple three-point averaging filter was used to smooth the data.
 - b) *Temperature*: For all the tests we used a pyrometer. This instrument had to be focused very carefully on the stream of molten metal leaving the cupola. Difficulties with this measurement included the fact that the stream's position could vary, resulting in the pyrometer being focused on something besides the iron. It was also possible for the pyrometer to be bumped out of position, especially in the first several runs. In addition, sometimes the pyrometer just quit working. Later the pyrometer placement was stabilized, so that a more consistent reading of temperature was available. However, even so, all temperature data reported here must be taken as relative. This is because the pyrometer was found to be inaccurate by tests using a dip thermocouple. To deal with the variability in the temperature measurements we applied the same type of filtering algorithm as we did for melt rate (see below). That is, we applied hardlimits and standard deviation filters that threw out data points outside an acceptable range and we then did simple averaging (usually ten to thirty points, taken five to seven seconds apart).

c) *Meltrate*: Getting a good meltrate measurement was possibly the most challenging measurement problem we faced. This was because the measurement we based all our meltrate calculation on was the actual weight of iron. Thus it was necessary to differentiate the measurement of weight to get meltrate (weight per unit time). The technique used to do this was to compute a least-squares fit of a line to a fixed number of weight readings. The slope of this line was then passed through a hardlimiter filter that threw out suggested meltrates that were completely unreasonable. We also passed the meltrate through a standard deviation filter that rejected outliers. The algorithm used for producing meltrate is shown in Figure 3. Note that another way to compute meltrate is to use charge deck data. When the ALRC cupola is run the time between charges is recorded. This data can be used to determine meltrate and indeed, this approach is common in foundries. Figure compares the output of two different meltrate filters with meltrate computed from the charge deck data. The top plot in the figure is the meltrate computation used by the controller in the final two experiments. The middle plot is the meltrate computed from charge deck data. The lower plot is similar to the upper plot, but does not have exactly the same hardlimiting algorithms (note: all plots have the same zero point, but two have been offset by varying amounts for readability). It is clear that the meltrates are all quite similar. This implies that charge deck data could be used. This is an important point, because many foundries do not have the capability to measure meltrate based on the weight of output iron.

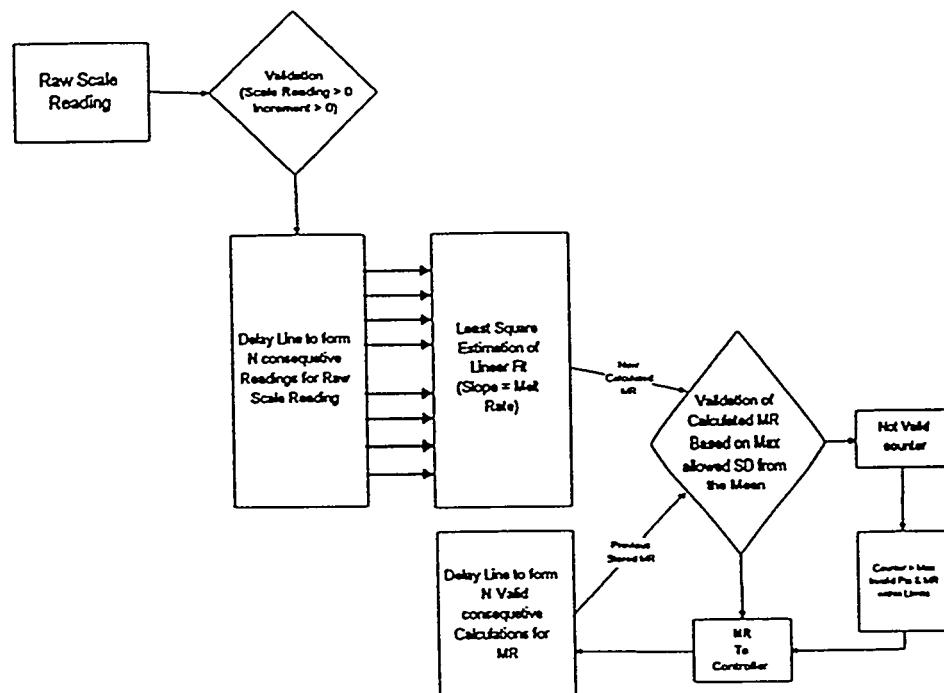


Figure 3: Flowchart for melt rate measurement.

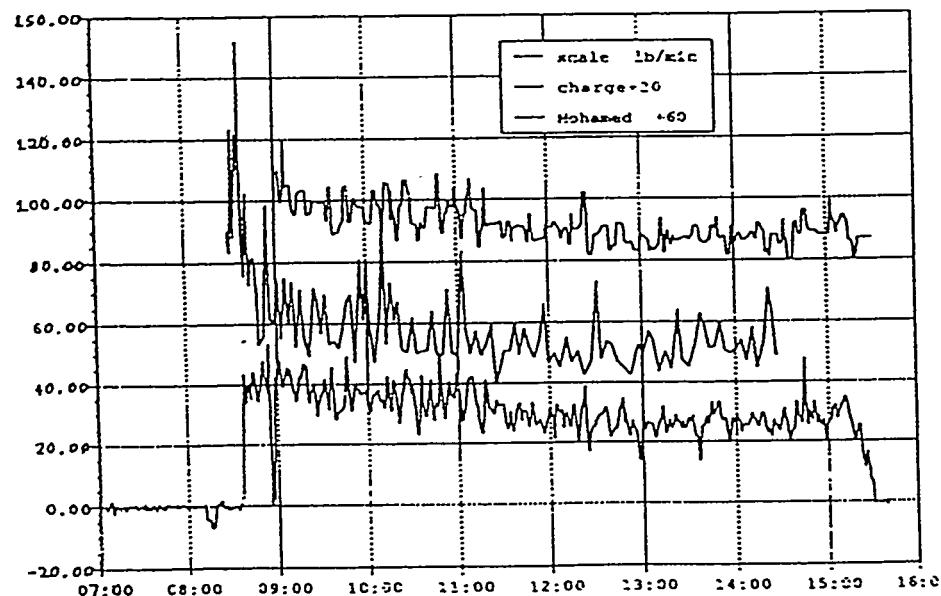


Figure 4: Melt rate computation from weight of iron versus charge deck data.

4. Actual controller gains were chosen via simulation. Closed-loop poles were chosen so that there was no overshoot in any signals in the simulated experiments. This was done using standard root locus-based design and then checked via simulation. The resulting controller had the form:

$$\begin{bmatrix} \Delta \% \text{ Coke} \\ \Delta O_2 \\ \Delta B_R \end{bmatrix} = \begin{bmatrix} C_1(s) & 0 & 0 \\ 0 & C_2(s) & 0 \\ 0 & 0 & C_3(s) \end{bmatrix} \begin{bmatrix} E \% C \\ E T_{FE} \\ E MR \end{bmatrix}$$

where E denotes the error signal. The Smith predictor used to regulate carbon concentration has the form:

$$C_1(s) = \frac{C(s)}{1 + C(s)G(s)(1 - e^{-Ts})}$$

where

$$C(s) = 0.1 + \frac{0.03}{s}, G(s) = \frac{0.04}{300s + 1}$$

The other two controllers are given by:

$$C_2(s) = 0.1 + \frac{2.8 \times 10^{-4}}{s}, C_3(s) = 3.0 + \frac{0.03}{s}$$

Experimental Results

A number of experiments were conducted, going gradually from single-loop operation to full multi-loop control. We began with single-loop control of melt rate and temperature, one at a time. A representative result is shown in Figure 5, which gives the output response of an experiment to control melt rate by adjusting the blast input. The setpoint in this experiment was 45 lbs/min. Notice that the melt rate shows significant variation about the setpoint. Analysis has shown that this variation is real and reflects how hard it is to control the cupola.

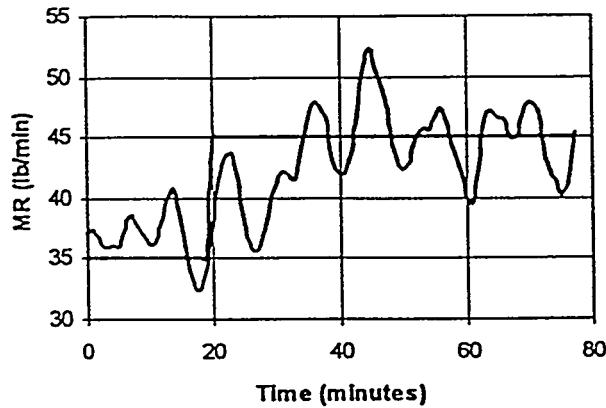


Figure 5: Single-loop control of meltrate.

After completing a number of single-loop experiments, we demonstrated multi-loop control of meltrate and temperature simultaneously. We then conducted a single-loop control experiment to regulate the carbon concentration using the Smith predictor. The final experiment consisted of demonstrating simultaneous control of all three outputs of interest: meltrate, temperature, and percent carbon. In the following we will discuss the final experiment. The sequence of events was as follows:

1. The furnace was started and brought to steady-state.
2. The controllers were turned on.

Meltrate setpoint was 40 lbs/min.

Iron temperature setpoint was 1400 degrees C.

%Carbon setpoint was 3.3%.

3. After about two hours the meltrate setpoint was changed to 35 lbs/min.

Figure 6 shows the result of the multi-loop control experiment. The plots on the left are the manipulated variables (i.e., the cupola inputs that are commanded by the controller) and the

plots on the right are the respective controlled variables (i.e., the cupola outputs that are being driven to a setpoint) that are paired with each manipulated variable. For approximately one and one-half hours the furnace was operated open-loop until it came to a steady-state. At this point the meltrate was approximately 36 lbs/min, the temperature was roughly 1340 degrees C, and the carbon composition was about 3%. When the controller was turned it began changing the blast rate, oxygen, and %Coke to drive the outputs to the desired setpoints. It can be seen from the figure that the controller action is effective.

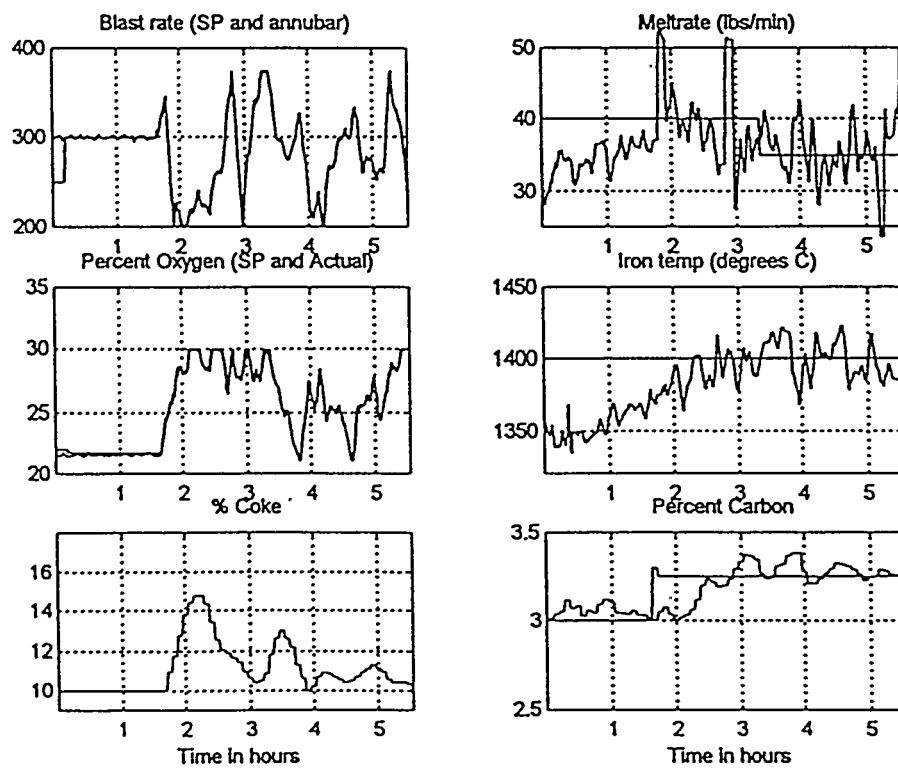


Figure 6: Simultaneous control of meltrate, temperature, and carbon composition; controller is turned on at approximately 1.5 hours; meltrate setpoint is changed at approximately 3.5 hours.

Summary and Future Work

Having described what has been accomplished in the intelligent cupola control project, we might also note what remains to be done. Referring again to Figure 1, though not described here, we have demonstrated software to perform *Neural Network Modeling* of the cupola, with appropriate *Neural Network Training* procedures based on the AFS model [5,6]. Also not described here, we have also demonstrated how to use the neural network model at the *System-Level* to optimize operating points based on economic considerations [10]. Finally, we have demonstrated experimentally the effectiveness of feedback control at the *Process-Level* to keep the cupola at a desired operating point. However, we must still integrate the algorithms from each level into a single software package and then demonstrate the use of the complete control scheme experimentally. In the experimental demonstrations carried out at ALRC, the setpoints for the process-level controller were not chosen by the system-level controller, but on the basis of operational considerations. Although we have successfully demonstrated the algorithms for each block in Figure 1, it remains to demonstrate the integration of these algorithms.

In addition to the complete integration of the system-level controller and the process-level controller, a number of other areas of further research and development work have been suggested by the intelligent cupola control project:

1. Our process-level control demonstrations were limited to only one metal stream (percent coke) and two blast streams (rate and oxygen). However, typically foundry operations also consider blast temperature and a metal stream composed of up to seven to ten different solid streams.
2. The ALRC demonstrations implemented a slightly less complicated version of the control schemes than were designed on paper (the implementation used multi-loop control; on paper

we had developed robust multivariable controllers utilizing feedforward elements for decoupling [8]). If we implement the multivariable controllers we will see improved results.

3. There is a strong need to develop more robust measurement techniques.
4. The availability of the AFS model and the neural network model provide an opportunity for developing advanced operational aids for diagnostics and troubleshooting.

These observations motivate additional project activity. Figure 7 shows a *generic architecture for an intelligent control system* that can be broadly applied to industrial automation and control

problems. We call this an intelligent, integrated industrial process controller, or I^3PC for short.

The concept of the I^3PC is that it is an enabling technology for advanced industrial process control that can be exploited to move iron melting furnace operations from an operator-intensive activity to a more modern, automated activity. The technical concept is to combine model-based intelligent control methodologies (including neural network, fuzzy logic, expert system, and advanced control system technologies) to define and develop an integrated approach to advanced industrial automation and control for ferrous metalcasting industries that also can be universally applied in other industrial process control environments. The next stage of the intelligent cupola control project is to develop and apply the I^3PC first to the ALRC cupola and then in an industrial foundry.

References

- [1] E.D. Larsen, D.E. Clark, H.B. Smartt, K.L. Moore, "Intelligent Control of Cupola Melting," *Transactions of the American Foundrymen's Society*, 95-67, pp 215-219, 1995.
- [2] K. Moore, "Intelligent Cupola Control," Funded Research Program, U.S. Dept. of Energy Metal Casting Competitiveness Initiative, Idaho Field Office, Idaho Falls, ID, October 1994.
- [3] P.E. King, L.G. Higgins, V. Stanek, E.D. Larsen, D.E. Clark, K. L. Moore, "AFS Cupola Model Verification, Part I - Initial Investigations," presented at the 1997 AFS Casting Congress, Seattle, WA, April 1997.
- [4] E.D. Larsen, D.E. Clark, H.B. Smartt, K.L. Moore, "Advanced User Interfaces and Intelligent Control," in *Proceedings of the 1994 INEL Technical Computing Symposium*,

Idaho Falls, ID, October 1994.

- [5] E.D. Larsen, D.E. Clark, K.L. Moore, V. Stanek, P.E. King, "Neural Network Applications for Cupola Melting" in *Materials Processing in the Computer Age, Proceedings of the 1995 TMS Annual Meeting*, Las Vegas, NV, February 1995.
- [6] E.D. Larsen, D.E. Clark, K.L. Moore, P.E. King, "Intelligent Control of a Cupola Furnace," in *Proceedings of Technology 2006, the 6th National Technology Transfer Conference and Expo*, Anaheim, CA, October 1996.
- [7] K. L. Moore, M. A. Abdelrahman, D. Clark, E. Larsen, H. Smartt, and C. Einerson, "A Multivariable Smith Predictor for Intelligent Control of a Cupola Furnace," In *Proceedings of the American Control Conference*, Seattle, Washington, June 21–23, 1995.
- [8] M.A. Abdelrahman, K.L. Moore, "Robust Control of Cupola Iron Furnaces," in *Proceedings of the 1996 American Control Conference*, Albuquerque, NM, June, 1997.
- [9] K. L. Moore, M. A. Abdelrahman, E.D. Larsen, D.E. Clark, P.E. King, "Experimental Control of a Cupola Furnace," in *Proceedings of the 1996 American Control Conference*, Philadelphia, PA, June, 1998.
- [10] M.A. Abdelrahman, K.L. Moore, "Optimal Choice of Cupola Furnace Nominal Operating Point," in *Proceedings of 30th IEEE Southeastern Symposium on Systems Theory*, University of West Virginia, Morgantown, WV, March 8-10, 1998.

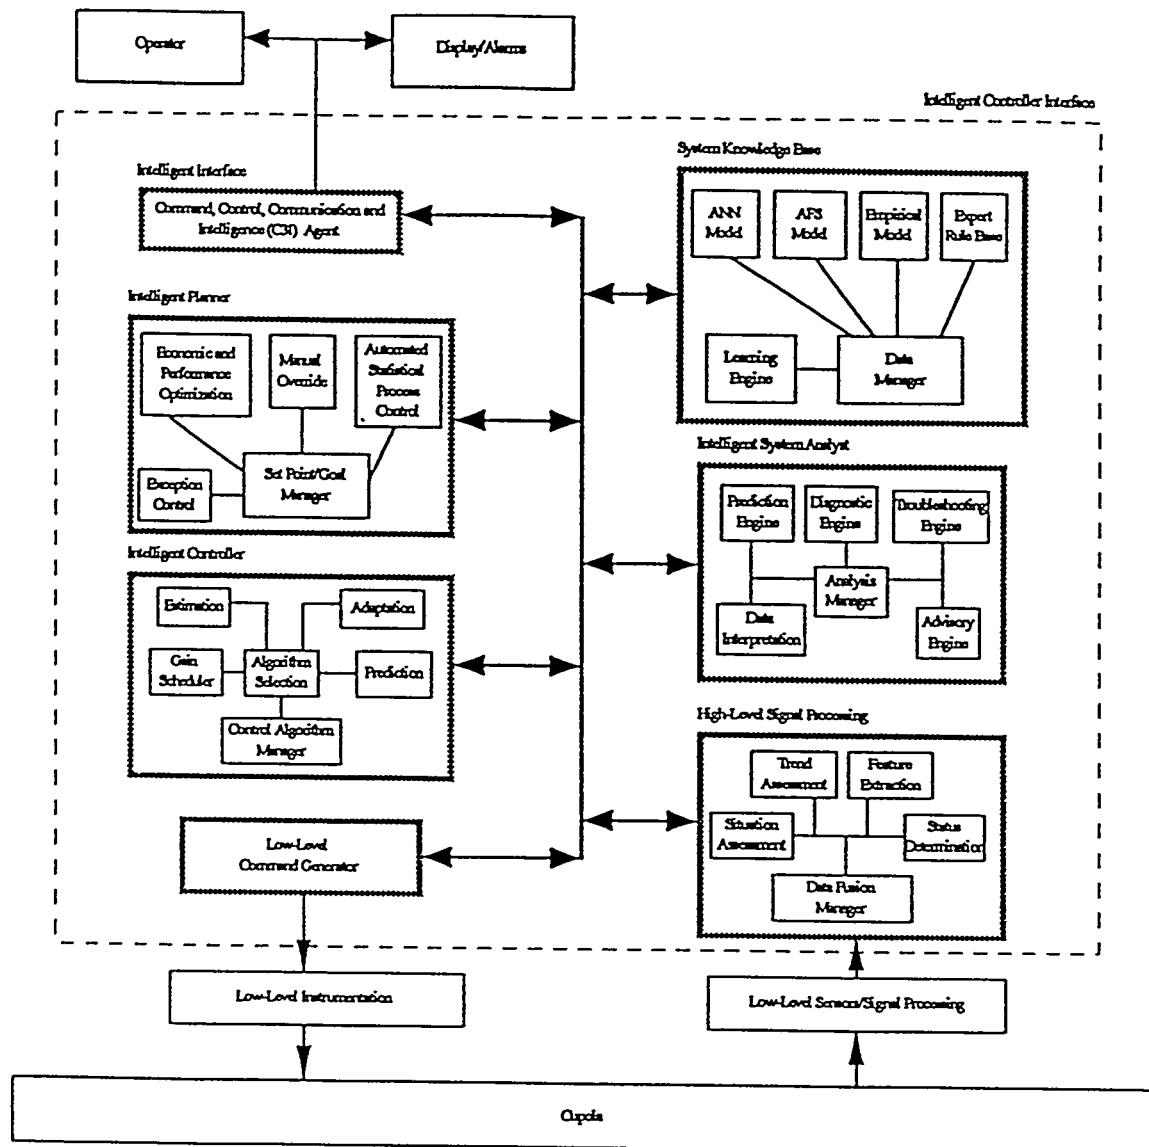


Figure 7: Integrated, intelligent industrial process controller.