

Development of a Safeguards Monitoring System for Special Nuclear Facilities

A Dissertation Presented for the
Doctor of Philosophy
Degree

The University of Tennessee, Knoxville

James Joseph Henkel

August 2011

Copyright © 2011 by James J. Henkel
All rights reserved.

DEDICATION

To my Father and Mother; thank you for your unwavering support.

ACKNOWLEDGEMENTS

I would like to thank the members of my dissertation committee for taking time to review this work: Dr. J. Wesley Hines, Dr. Belle Upadhyaya, Dr. Jason Hayward, and Dr. Mary Leitnaker. I would especially like to thank Dr. Hines for his guidance, dedication, and patience during my four years as his research assistant.

The Oak Ridge National Laboratory International Safeguards group has been an invaluable resource during these past years. Aside from providing a test facility, their guidance and support was instrumental in successfully completing this research.

During the course of this research I was fortunate enough to work with David Hooper. His help and insight were invaluable and this research could not have been completed without him.

As an undergraduate, I was fortunate to work with the Oak Ridge National Laboratory's Nuclear Materials Detection and Characterization group under Dr. John Mihalcz. This work continued throughout my graduate education. My positive undergraduate experience with Dr. John Mihalcz and the Oak Ridge National Laboratory convinced me to pursue a graduate degree.

Finally, a large debt of gratitude is owed my family. I would not have been able to complete my education without their support and encouragement. A special thanks is owed my parents, who at an early age began teaching me what they coined "Real Life 101" – those lessons have been more valuable than any class I have taken.

ABSTRACT

Two important issues related to nuclear materials safeguards are the continuous monitoring of nuclear processing facilities to verify that undeclared uranium is not processed or enriched and to verify that declared uranium is accounted for. The International Atomic Energy Agency (IAEA) is tasked with ensuring special nuclear facilities are operating as declared and that proper material safeguards have been followed. Traditional safeguards measures have relied on IAEA personnel inspecting each facility and verifying material with authenticated instrumentation.

In newer facilities most plant instrumentation data are collected electronically and stored in a central computer. Facilities collect this information for a variety of reasons, most notably for process optimization and monitoring. The field of process monitoring has grown significantly over the past decades, and techniques have been developed to detect and identify changes and to improve reliability and safety. Several of these techniques can also be applied to international and domestic safeguards.

This dissertation introduces a safeguards monitoring system developed for both a simulated Uranium blend down facility, and a water-processing facility at the Oak Ridge National Laboratory. For the simulated facility, a safeguards monitoring system is developed using an Auto-Associative Kernel Regression model, and the effects of incorporating facility specific radiation sensors and preprocessing the data are examined. The best safeguards model was able to detect diversions as small as 1.1%. For the ORNL facility, a load cell monitoring system was developed. This monitoring system provides an inspector with an efficient way to identify undeclared activity and to identify atypical facility operation, included diversions as small as 0.1 kg. The system also provides a foundation for an on-line safeguards monitoring approach where inspectors remotely facility data to draw safeguards conclusion, possibly reducing the needed frequency and duration of a traditional inspection.

TABLES OF CONTENT

1	INTRODUCTION	1
1.1	Problem Statement.....	4
1.2	Original Contributions	6
1.3	Organization	7
2	LITERATURE SURVEY	8
2.1	Process Monitoring (PM)	8
2.1.1	Parametric Models.....	12
2.1.2	Kernel Regression	14
2.1.3	Data Reconciliation	19
2.1.4	Sequential Probability Ratio Test.....	24
2.2	Nuclear Safeguards.....	25
2.2.1	Process Monitoring for Nuclear Safeguards	29
3	METHODOLOGY.....	41
3.1	Safeguards Monitoring of a Simulated Uranium Blend-Down Facility.....	41
3.1.1	Building an AAKR Model.....	46
3.2	Load Cell Monitoring at a Mock Feed and Withdrawal Facility	51
4	APPLICATIONS AND RESULTS.....	60
4.1	The Uranium Blend-Down Facility Simulation	60
4.2	The ORNL Mock Feed and Withdrawal Facility	70
4.2.1	Types of Operation at the ORNL Facility	76
4.2.2	Development of a Load Cell Monitoring System	78
4.2.3	Addition of a PI Controller	91
4.2.4	GUI Development	100
5	CONCLUSIONS.....	120
5.1	Summary of Contributions	122
5.2	Future Work.....	122
	LIST OF REFERENCES	125
	APPENDIX A: Relevant Equations to the Uranium Blend-Down Model	133

APPENDIX B: Relevant Equations to ORNL Facility Model	136
APPENDIX C: Descriptions of Algorithms used by PlotEvents	138
APPENDIX D: Sample Input File For PlotEvents	142
APPENDIX E: Process Summary Report.....	143
VITA	146

LIST OF TABLES

Table 2-1. An application of data reconciliation to a hypothetical flow splitter. The physical system is constrained by $M1 = M2+M3$, but due to measurement error the measured values do not exactly obey the constraining law. Data reconciliation is used to calculate a correction factor for each measured value (weighted by both the magnitude of the measured value and the standard deviation of the measurement), to force the measured values to follow the conservation law.	23
Table 3-1. Summary of the simulated instrumentation of the Blend Down Facility.	45
Table 3-2. Summary of monitoring models for the Simulated Blend Down Facility. While each model was based on an empirical AAKR model, there were a total of 9 different parameters available to each model – 8 different sensors and the possibility of preprocessing the measured signal with data reconciliation. An “X” indicates that the sensor or technique was included in the AAKR model.	47
Table 3-3. A material declaration sheet from the ORNL mock feed and withdrawal facility. The sheet list the initial and final tank weight, as read by the accountability scales, before and after processing.	57
Table 4-1. Summary of AAKR model performance for unfaulted data from the blend-down monitoring system simulation. The accuracy is expressed in terms of percent error, so a lower value is indicative of less error. The uncertainty is expressed as a percent of the mean sensor value.	63
Table 4-2. Model performance metrics for model 3 (No Data Reconciliation) and Model 4 (With Data Reconciliation). With each metric, a lower value is desired. The accuracy is expressed in terms of percent error, so a lower value is indicative of less error.	71
Table 4-3. Summary of thresholds for identifying station states.	85
Table 4-4. Material Declaration Sheet for October 14 and 15, 2010	111

Table 4-5 Process Declaration Sheet for October 14 and 15, 2010	112
---	-----

Table 4-6. Results of comparing the Material Declaration Sheet and Process Declaration Sheet. Reconciled tanks are shown first, then un-matched tanks listed on the material declaration sheet, and finally undeclared tanks.	113
--	-----

LIST OF FIGURES

Figure 2-1. The basic structure of a process monitoring model. Measured values are used to create a prediction of the measured signals, and then the residuals between the measured and predicted values are used in decision logic.	9
Figure 2-2. Outline of a parametric and non-parametric empirical model. A parametric model optimizes parameters of a mathematical model once and then uses those parameters for each subsequent prediction. The model training data set is not saved for future use. A non-parametric model uses historical operation data to build a parametric model for each prediction. Each new prediction requires a new parametric model, and the optimized parameters are not saved. The historical operation data is saved, but the set of optimized parametric parameters is discarded after predictions are made.	15
Figure 2-3. The effect of incorrect Gaussian Kernel bandwidth selection. Too large of a bandwidth (top) results in a larger range of distances receiving a non-trivial weight, which can introduce a large prediction bias, where the peaks and values of the data are not well fit. Too small of a bandwidth (bottom) results in large weights only when the distance is small, and can introduce a large variance [26]. An optimal bandwidth minimizes the model uncertainty.	17
Figure 2-4. A properly sized bandwidth balances the effect of prediction bias and model variance. The result is that the peaks and values of the function are well fit and that the overall model uncertainty is minimized [26].	18
Figure 2-5. Diagram of a simple flow splitter. In this system the mass flow out of M1 must equal the mass flow into M2 and M3. However, due to random measurement error, the measured mass flow rate may not exactly represent the relationship $M1 = M2 + M3$. Data Reconciliation is used to calculate a correction factor to the measured values so that they obey the physical relationship.	21

Figure 2-6. Block diagram of material flow through a hypothetical gas enrichment plant. As material is received, its contents are verified and it is put into storage. When material is ready to be processed, it is moved from storage to an input accountability tank, thereby crossing from MBA 1 to MBA 2. After the material is processed, it is removed from the process line and placed in storage, thereby crossing from MBA 2 to MBA 3, where it stays until it is shipped offsite. At each MBA crossing, the MCA is done through destructive assay and key flow measurements. 27

Figure 2-7. A hypothetical load cell profile as a feed tank is loaded, processed, and removed from a GCEP feed station. At an actual plant, the profile would be contaminated with noise spikes, most significantly when a tank is first loaded or removed. 32

Figure 2-8. Changes in a GCEP feed station profile under different operating conditions. The spikes when a tank is added or removed from the scale represent expected fluctuations. 34

Figure 2-9. A general feed template for a GCEP feed station. Section A marks the period when a cylinder is emptying. Section B marks the period when an empty cylinder still resides on the station load cells. Section C marks the delay before a new feed cylinder is placed on the station. Section D marks the period when a new cylinder is first placed on the station, and the weight value is still fluctuating. Section E marks the period when a new cylinder is ready to be emptied [49]. 36

Figure 2-10. A generic feed, accountability, and Buffer tank connected in series [50]. A material balance area surrounds the three tanks. The flow rates are denoted by F , the temperature by C^0 , the density by ρ , and the volume by V . Measurements of the nitric acid concentration, plutonium concentration, and uranium concentration are denoted by H , Pu , and U respectively [50]. 38

Figure 3-1. A schematic of the Fissile Mass Flow Meter (FMFM) used in the blend-down monitoring system. When the shutter is opened a Californium-252 neutron source induces fissions in a uranium hexafluoride stream. Downstream of the neutron source

is a shielded detector array used to detect delayed gammas from Uranium-235 fission. The number of delayed gammas detected is proportional to the amount of Uranium-235 in the stream. The time delay between the source shutter opening and the initial detection of fission gammas can be used to calculate the flow velocity. 42

Figure 3-2. Schematic of the simulated uranium blend-down facility. All legs are instrumented with flow sensors and tank weight sensors. Additionally, the HEUF₆ and PLEUF₆ legs are instrumented with FMFMs. Diversions were simulated by removing material from the HEUF₆ just before the blending tee. 44

Figure 3-3. Block diagram of AAKR model with data reconciliation. The data selection, model building and predictions, and computing the residuals are the exact same as an AAKR model without data reconciliation. The purpose of including data reconciliation is to better constrain the measurement values using known physical laws about the system. 48

Figure 3-4. Schematic of the ORNL mock feed and withdrawal facility. Tanks of water are placed on the feed station and the water is pumped into a surge tank. Water is then partially routed to a tail and production station. In this setup, the pump, surge tank, and stream splitter serve as a mock cascade area. 52

Figure 3-5. The actual ORNL mock feed and withdrawal facility. Three feed stations are pumped into a Surge tank, which then drains into the product and tail stations. Before and after a tank is processed at each station, it is weighed on an accountability scale. 53

Figure 3-6. The life cycle of a tank at the ORNL mock feed and withdrawal facility and the type of data generated during each stage of the tank life cycle. The black boxes and arrows represent how the physical material is moved through the facility. The blue boxes and dashed arrows represent the data collected from each physical location. ... 55

Figure 3-7. A sample of the different data types collected at the ORNL mock feed and withdrawal facility. The accountancy scale data (a) is a step function of a tank's weight before and after processing. The process scale data (b) is continuous and shows the

instantaneous weight of a tank on a station as material is being processed. The feed pump power (c) data is a type of facility control data and may either be a step function or continuous depending upon the instrumentation..... 56

Figure 4-1. Simulated data from the Uranium blend-down facility Simulink model. The simulated data represents unfaulted operation. Initially, the LEU process scale is vacant, then a tank is added and its material processed, and finally the tank is removed. The LEU Flow and FMFM show a step change corresponding to when a tank begins processing. The LEU mass flow is not a direct output of the model but is approximated by taking the derivative of the LEU tank weight..... 62

Figure 4-2. Residuals and SPRT Alarm for the P-LEU velocity sensor (model 2) 10% diversion scenario. Because there is no statistical change between the predicted and measured values, a 10% diversion could not be detected with this model. Note that in the top graph the residuals show no change and the SPRT does not alarm..... 64

Figure 4-3. Cumulative weight difference between the three tanks during a diversion. Under normal operation, the cumulative weight difference should be zero, since everything that leaves the HEUF₆ and LEUF₆ tanks enter the PLEUF₆ tank. During a 10% diversion, there is a growing residual as the cumulative amount of material diverted increases..... 66

Figure 4-4. Predicted and measured PLEUF₆ flow rates for the 10% diversion scenario for the 8 sensor model (models 3 and 4). In the top plot the measured values are in blue and the prediction values are in red. There is a large statistical change between the predicted and measured values. A 10% diversion could be detected with this model. Note that in the top graph, the predicted values are indistinguishable from the measured values before the diversion takes place. 67

Figure 4-5. Predicted and measured PLEUF₆ flow rates for the 2% diversion scenario for the 8 sensor model (models 3 and 4). In the top plot, the measured values are in blue and the prediction values are in red. The SPRT identifies statistical change in the residuals, indicating that a 2% diversion could be detected with this model..... 69

Figure 4-6. An essential function of the MATLAB GUI is to determine if the amount of material processed according to the accountancy scale data (and listed in the declaration sheet) agrees with the amount of material processed according to the process scale data. 73

Figure 4-7. Original Facility control setup. The yellow-handled valve behind the throttling valves acted as a master on/off valve for the surge tank. During operation, the master valve was opened and the two throttling valves were used to route the proper amount of material to the product and tail legs. Proper operation required that the operator be aware of the surge tank level and the ratio of material routed to the product and tail legs. In practice, balancing material flow into and out of the surge tank while simultaneously routing the proper amount of material to each leg was difficult to achieve by simply adjusting the throttling valves. The control system was very sensitive to operator attentiveness and to changes in operator personnel..... 74

Figure 4-8. Facility data from March 27, 2009 when the facility was still under manual control. The black line at 125 minutes represents when operation should have finished. There is no more feed material entering the mock cascade area, and the product tank is completely filled. However, because the operator was unable to balance the material flow within the surge tank, approximately 30 kg of water accumulated in the surge tank which then had to be drained into a tail tank. Additionally the blue circles highlight times when the tail tank flows were adjusted, but without any corresponding adjustments to the feed or tail flow. This control scheme led to high variability in the system, where inflow to the mock cascade area was no longer an accurate predictor of outflow of the mock cascade area. 75

Figure 4-9. The life cycle of a tank during undeclared activity. This scenario looks like normal operation except no authenticated accountability data is generated. This scenario represents a facility processing undeclared material in an attempt to make HEUF₆ for their own, possible clandestine, purpose. 77

Figure 4-10. Diversion path at the ORNL facility. 79

Figure 4-11. Description of an enrichment masking scenario at the ORNL facility. 80

Figure 4-12. Raw tank station profiles. (a) shows the profile from a Feed Station, while (b) shows the profile from a tail and product station. Both (a) and (b) show the large spikes when a tank is added or removed (circled in red). In (b) the product station (blue) spikes are not as significant because the product tanks are much smaller. (c) Shows the profiles from all 8 stations (3 feed, 3 product, and 2 tail) at the ORNL facility. 81

Figure 4-13. The effect of using a median filter on a feed scale profile. The raw data shown in blue is noise free except for spikes when a tank is unloaded and loaded – at approximately 600 and 650 minutes respectively. The dotted red line shows the smoothed data. It follows the same profile as the raw data, only without the noise spikes. 83

Figure 4-14. A feed station profile and its approximated mass flow rates. The top plot shows a feed station profile with the 0.5 kg limit plotted as a dotted line. When the weight reading is below the dotted line, the station is said to be in the Empty state. The bottom plot shows the approximated mass flow rate for the same profile. The thresholds for the Static, Filling, and Draining state are shown as dotted lines. In addition to the mass flow rate falling within the specified region, the magnitude of the weight must be above 0.5. 86

Figure 4-15. The top plot shows the raw data from a feed station. The bottom plot shows the same station only the data has been filtered and the states identified. The profile portions are correctly identified as Empty, Static, or Draining based on the derivative and the absolute value of the data. 87

Figure 4-16. The life cycle of a process station with respect to 4 possible states. 88

Figure 4-17. A sample accountancy scale profile and process scale profile. The dark blue bars represent data from the accountancy scales. These values are step functions that represent the authenticated tank weights. The accountancy scale values are then

used to generate the declaration report. The light blue line represents a station profile as recorded by the process scales. This information is used to generate a process summary report which can be compared directly to the declaration report, providing the inspector with more assurance that all the material processed by the facility was recorded in the declaration report. 90

Figure 4-18. Simulated data from incorporating a PI controller instead of a manual control scheme. In the simulated data, the PI controller keeps a constant level in the surge tank, so the mass flow into and out of the surge tank are kept equal. When the mass flow into and out of the surge tank is equal, the correlations between each individual flow rate increase, and more sensitive monitoring techniques can be employed to search for atypical facility operation. 92

Figure 4-19. The updated control scheme at the ORNL Facility. The installation of a pressure transducer with a PI control valve removed a human element from control of the facility, removing the largest source of variability in facility operation. The pressure transducer gave a measure of the water level in the surge tank, and the PI controller adjusted the valve to maintain a set level. 94

Figure 4-20. Plots of the cumulative inventory difference from the March 27, 2009 run, when the facility was still under a manual control scheme. 95

Figure 4-21. Cumulative inventory difference from July 19, 2010, after installation of the PI controller. Plot (a) is the CID for each station, and plot (b) is the total CID calculated by summing the data in plot (a). 97

Figure 4-22. The cumulative inventory difference and SPRT results for the July 19, 2009 run. The cumulative inventory difference was smoothed with a median filter, and the mean and variance was taken from a window of data points right after the startup transient. The bottom plot shows the result of the SPRT. The red circles represent when the data failed the SPRT, and blue represents normal operation. 99

Figure 4-23. Block Diagram of the MATLAB GUI. The inputs to the GUI are shown as blue lines and the outputs are shown as dotted lines. The solid black lines represent different sections of the GUI. 101

Figure 4-24. The opening PlotEvents window after data has been loaded. Clicking the Load Data button prompts the user to load a data file. Initially every station is plotted, but the user can select difference stations and then use the Replot button to update the plot. 102

Figure 4-25. The window opened when the Analyze Events button was pressed. Four different types of analysis are available and can be selected using the four toggle buttons in the top right corner. The initial analysis is the System States. 104

Figure 4-26. The AnalyzeEvents window when the Cumulative ID analysis is selected. Different stations can be selected and plotted together. Additionally, the total cumulative inventory difference can be plotted. 105

Figure 4-27. The AnalyzeEvents window when ID analysis is selected. The default top plot is the raw and smoothed cumulative ID difference. The bottom plot is the SPRT alarm. Red indicates an atypical state; blue represents a normal state. Startup and shutdown transients are recognized as atypical operation. The bottom plot is the Sequential Probability Ratio Test (SPRT) alarm. When the facility is under normal steady state operation, the CID should be stable. The SPRT alarm is a statistical test to determine if the CID is remaining stable. The red portion indicates that the CID is not stable, while the blue portion means it is stable. The SPRT alarm provides indication of changes to facility operation and is meant as an inspector tool to recognize atypical facility operations. 107

Figure 4-28. The AnalyzeEvents window when ID analysis is selected and the Gaussian ID plot is selected. Instead of the CID vs. time, the top plot is a histogram of the CID fitted to a normal curve. The bottom plot remains unchanged. 108

Figure 4-29. The AnalyzeEvents window after the Write Summary Report button has been executed. A new button becomes available: Compare to Declaration, which allows the user to select a material declaration sheet and process declaration sheet 109

Figure 4-30. The ID analysis of normal operation. The CID remains stable with the exception of startup and shutdown transients. The SPRT alarm triggers during the transients, but not steady state operation..... 115

Figure 4-31. The ID analysis of a diversion operation. The cumulative inventory difference does not remain stable. The SPRT alarm identifies the deviation from steady state operation. 116

Figure 4-32. SPRT alarms during a masked enrichment run. 117

Figure 4-33. Simulated masked enrichment detection. Because the feed flow was higher than expected, the PI voltage was larger than predicted. Using an SPRT alarm on the residuals clearly identified that atypical operation had taken place. 119

ABBREVIATIONS

NNSA	National Nuclear Security Administration
IAEA	International Atomic Energy Agency
NPT	Nuclear Non-Proliferation Treaty
MBA	Material Balance Area
PM	Process Monitoring
CM	Condition Monitoring
OLM	On-Line Monitoring
MCA	Material Control and Accountability
PDI	Person Days of Inspection
GCEP	Gas Centrifuge Enrichment Plant
HEU	Highly Enriched Uranium
BDMS	Blend-Down Monitoring System
ORNL	Oak Ridge National Laboratory
FMFM	Fissile Mass Flow Meter
HEUF ₆	Highly Enriched Uranium Hexafluoride
LEUF ₆	Low Enriched Uranium Hexafluoride
PLEUF ₆	Product Low Enriched Uranium Hexafluoride
GUI	Graphical User Interface
SPRT	Sequential Probability Ratio Test
EULM	Error Uncertainty Limit Monitoring
CID	Cumulative Inventory Difference

1 INTRODUCTION

Two important issues related to nuclear materials safeguards are: 1) the continuous monitoring of uranium enrichment facilities to verify that undeclared special nuclear material is not processed and 2) that declared special nuclear material is accounted for. For both cases, the material enrichment and flow need to be monitored. This is of importance to both the National Nuclear Security Administration (NNSA) and the International Atomic Energy Agency (IAEA) in enforcing the elements of the Nuclear Non-Proliferation Treaty (NPT) and other agreements. The signatories of the NPT have agreed not to develop, manufacture or otherwise acquire nuclear weapons or other nuclear explosive devices and have accepted safeguards on all nuclear material used in peaceful nuclear activities [1, 2]. The basic measure by which the IAEA verifies the fulfillment of these obligations is through nuclear material accountancy [3]. Therefore, one goal of the IAEA is to detect diversions of significant quantities of nuclear material from peaceful activities in a timely manner [4, 5]. The significant quantity of nuclear material and time period depends on the type of material. For example a significant quantity of nuclear material is defined as 8 kg of plutonium, 8 kg of ^{233}U , 25 kg of ^{235}U in highly enriched uranium ($\geq 20\%$ ^{235}U), 75 kg of ^{235}U in low enriched uranium ($< 20\%$ ^{235}U), or 20 tons of thorium. Time periods may range from one month (for pure plutonium metal) to three months (for plutonium in irradiated fuel) [6]. A second goal of the IAEA is to ensure that a plant declares all of the material it processes. Both are of importance to the IAEA because diverted material and undeclared nuclear material could be used in illegal manufacture of nuclear weapons. Nuclear safeguards are the protective measures by which the IAEA ensures that nuclear material is accounted for.

Traditional safeguards include designating material balance areas (MBA) throughout a facility. Several MBAs are common within bulk handling facilities. For some nuclear facilities, because of the itemized nature of most nuclear material, simple material control and accountability (MCA) is a sufficient safeguard. For example, a nuclear power plant receives and ships nuclear fuel assemblies. The assemblies are uniquely labeled and the nuclear components are not subdivided or easily removed. For these assemblies, tracking each container and ensuring its integrity ensures that the

nuclear material is kept accountable. In this sense, the entire plant could be viewed as a large MBA. Other facilities, namely fuel reprocessing plants and gas centrifuge enrichment plants (GCEPs), separate, divide, mix, or redistribute special nuclear material. In such facilities, MCA is no longer as simple as tracking a fuel bundle and ensuring its integrity, so multiple MBAs are defined within the facility. For example, reprocessing facilities have one MBA at the front end, a second MBA is typically defined that encompasses the process area, and a third MBA is usually defined around the product storage area at the tail end of the process. Material flow into and out of the MBAs is monitored by collecting samples from various locations within the MBA for destructive assay and measuring flows at predetermined locations at the boundaries of the MBAs [7, 8]. The IAEA will inspect, either during scheduled or random short notice inspections, equipment in each facility's MBA and the calibration of the instrumentation used to monitor the material flow between MBAs. Additionally, full inventories are conducted periodically to verify the amount of material contained within a MBA, rather than just confirming the material flow between MBAs.

All MBA measurements have a nominal measurement uncertainty. As facility throughput increases, the absolute value of the measurement error uncertainty can grow to larger than a significant quantity, thus providing no assurance that a significant quantity of material has not been diverted within the required time period. Additionally, the IAEA only verifies instrumentation calibration during an inspection. As nuclear facilities continue to grow in number, size, and complexity, the manpower required to monitor these facilities will continue to tax the limited resources of the IAEA. However, modern facilities will incorporate digital instrumentation and data acquisition systems. These data systems can be used to develop on-line, continuous or near-continuous monitoring systems, allowing IAEA inspectors to, possibly in a remote manner, efficiently verify equipment and instrument calibrations, as well as the total amount of material processed. While historically used for maintenance scheduling and process optimization, the general field of Process Monitoring encompasses the methods and techniques of on-line continuous monitoring, much of which is directly applicable to nuclear safeguards monitoring.

Process monitoring (PM) was developed to assist condition monitoring (CM) of plant equipment. CM is the practice of identifying the operating status of system components and using the current component condition to determine the optimal maintenance schedule and to optimize operations. Previous maintenance strategies relied on preventive or periodic maintenance and reactive maintenance. Preventive maintenance means performing maintenance on a set time schedule, regardless of the equipment's current condition. Preventive maintenance has two disadvantages. First, maintenance resources are wasted on systems that do not require maintenance, leading to expensive and unnecessary maintenance schedules. In addition, performing unnecessary maintenance on healthy components can introduce fault catalysts into previously properly working systems. Conversely, reactive maintenance is performing maintenance when a system component fails. This maintenance strategy leads to unplanned and expensive system downtimes [9].

One aspect of PM is to identify sensors that may require maintenance due to sensor drift or other malfunctions. In PM systems, data collected from plant sensors are evaluated with empirical and/or physics-based models to obtain an independent estimate of the actual, un-faulted plant parameter. This estimate is compared to the sensor signal in order to determine if the sensor has faulted. With this technology, continuous or near-continuous sensor surveillance is possible. These techniques have been approved in a general way for nuclear power plants for monitoring and calibration of safety-critical components and are used extensively in the nuclear power industry [10].

Adjusting standard PM techniques to monitoring for safeguards relevant information and combining those techniques with traditional nuclear safeguards offers the possibility to do real-time or near real time material inventory verification and instrumentation calibration verification. All available data can be integrated into a central model. This model can monitor changes in the system configuration, identify suspect instrumentation, validate operations and monitor equipment, and provide information for authentication and other decision-making tasks. While research has examined the potential of PM for nuclear safeguards in a general way, one key

component missing from published safeguards research is a method to integrate traditional nuclear safeguards with modern process monitoring techniques for facility safeguards conclusions. This work attempts to fill that need by developing a PM system specifically for the nuclear safeguards and for the verification of material processed. This work presents the research completed to develop a central monitoring system incorporating traditional nuclear safeguards with modern process monitoring techniques.

1.1 Problem Statement

The ultimate goal of nuclear safeguards is to validate the amount of material processed at a facility and to ensure that no material was diverted from its intended peaceful use. Traditional nuclear safeguards rely heavily on facility and equipment inspections. Facility inspections require large amounts of manpower, time, and money. The possibility of remote, continuous or near-continuous monitoring of a facility provides an efficient way to reduce the burden of facility inspections by taking advantage of the large amount of data already collected by a facility. Facility data is generally large in volume but comes from unauthenticated instruments it is not used for safeguards monitoring. On the other hand, inspector data comes from authenticated instruments placed at key locations between within the facility and is used in safeguards monitoring. While several papers have mentioned the possibility of incorporating facility data with inspector data for safeguards monitoring in a general way, the actual methodology and quantitative analysis of doing so has been overlooked in published literature.

A major hurdle to incorporating facility data for nuclear safeguards is the volume of data and its authentication. Most facilities already incorporate different degrees of PM to efficiently monitor signals from a wide variety of sensors, incorporate correlated signals into a central monitoring model, and use that model to look for subtle process changes that may arise from natural events such as sensor drift. For example, consider a gas flowing through a pipe. Standard instrumentation would include pressure, temperature, and flow rates. Since these measurements would be correlated with each other it is possible to build a monitoring model that relates two pressure and

temperature measurements to the flow rate measurement. By comparing the predicted flow rate values to the measurement values and tracking the residuals, an operator can quickly be alerted to changes in the process flow. For example, if the temperature and pressure measurements remain constant, then the predicted flow rate should remain constant. If the measured flow rate does not remain constant due to natural sensor drift, then the residual between the measured flow rate and predicted flow rate would be non-zero and would alert an operator. In this sense, PM can be viewed as somewhat self-authenticating, because two independent measurements are used to predict a third independent measurement. If the predicted value matches the measured value, then some assurance is provided that all three measurements are correct. As the facility grows in complexity, the relationships between measurements grow in complexity, and the effort required to fool the monitoring system grows exponentially.

Traditional safeguards methods are time consuming, tedious, and expensive. Additionally they ignore the potential information available in the unauthenticated facility data. Since the large amounts of facility data are already collected, incorporating this data into safeguards monitoring incurs little financial burden with huge potential gains in the effectiveness of the monitoring system. The development of automated PM routines coupled with the possibility of remote data access further reduces the required manpower to monitor a facility. The focus of this research is developing the methodology to create a PM system specific to nuclear safeguards incorporating both unauthenticated facility data with authenticated inspector data. While several different PM models are available, the auto-associative kernel regression model is used here. Initially, PM models are developed for a simulated uranium blend-down facility to determine the sensitivity of the model to detect nuclear material diversions. Then these algorithms are incorporated into a monitoring system developed for the Oak Ridge National laboratory (ORNL) mock feed and withdrawal facility. The ORNL facility mimics a continuous batch feed process, similar to a GCEP, and serves as a test-bed for the development of a nuclear safeguards monitoring systems. A continuous batch feed process is a process where containers of feed material, such as a feed cylinder at a GCEP, are brought into the processing area, the contents of the container are

processed, the empty container is removed, and then a new container is brought into the processing area. The monitoring system for the ORNL facility checks the unauthenticated data for sensor masking and abnormal process operation, and creates a log of the material processed which can be coupled with traditional inspections to check for undeclared material processing.

1.2 Original Contributions

The research described herein culminates in several original contributions to the field of nuclear safeguards monitoring. These contributions lie mainly in the development of PM algorithms to monitor and validate the material processed at a nuclear facility. Additionally a graphical user interface (GUI) was developed to aid in the prototyping of algorithms for safeguards monitoring and material validation. The main original contributions:

1. Application of an Auto Associative Kernel Regression (AAKR) model for continuous or near continuous safeguards monitoring. These models quantitatively characterize the ability of a monitoring system to detector sensor masking
2. Integration of data reconciliation and an AAKR model for safeguards monitoring.
3. Development of an automated routine to validate the amount of material processed at the ORNL facility, including flags alerting an inspector to abnormal process changes.
4. Development of a MATLAB-based toolbox and GUI for safeguards monitoring and material validation at the ORNL facility. The toolbox and GUI is designed as a tool for potential inspectors to quickly, and possibly remotely, use facility data to draw safeguards conclusions.

1.3 Organization

Chapter 2 reviews relevant literature in process monitoring algorithms and applications, as well as current methods in nuclear material safeguards monitoring. Chapter 3 presents the specific problem addressed by this work and the methodology proposed to solve this problem. Section 3.1 presents a MATLAB Simulink model of a uranium blend-down monitoring facility. This section also outlines the development of the safeguards monitoring system using an AAKR model and data reconciliation. Section 3.2 describes the Oak Ridge National Laboratory mock feed and withdrawal facility and the components of a load cell monitoring system.

Chapter 4 presents the application and results of the safeguards monitoring system to the two facilities. Section 4.1 presents the results of incorporating radiation sensors and data reconciliation into the AAKR architecture for the simulated uranium blend down facility. Section 4.2 presents the results of the load cell monitoring system for the ORNL mock feed and withdrawal facility. Section 4.2.4 describes a graphical user interface implementing the load cell monitoring system developed for the ORNL facility. Chapter 5 presents the conclusions from this dissertation, a summary of the contributions, and proposed areas of future work.

2 LITERATURE SURVEY

Traditional safeguards techniques rely on inspectors performing material balances between MBAs. MBAs are designated boundaries within a facility, and all special nuclear material is recorded as it enters and leaves a MBA. The role of inspectors is to validate MBA measurements and the amount of material within a MBA as reported by the facility. These safeguards methods are time consuming, tedious, labor intensive, and expensive. Additionally, these methods are limited in the sensitivity and speed at which diverted material may be detected. However, incorporating remote on-line monitoring of a plant has the potential to reduce the burden of IAEA inspectors. Most facilities incorporate some form of process monitoring already, and while the idea extending PM techniques for safeguards is not new, little research has been done develop to explore and quantitatively analyze the benefits of doing so. The basic theory of process monitoring is to collect measured signal values, use them as input to a prediction module to create signal predictions, compare the measured signal values to predicted signal values, and then make decisions based on the residual between the measured and predicted value, as shown in Figure 2-1. Based on the residual pattern the decision logic would decide if a sensor had drifted or if there was an abnormal process happening. Depending on the prediction model and implementation, PM is very sensitive to small sensor drifts in near real time. In an application to safeguards, the sensitivity of PM may be useful for detecting small diversion or sensor masking. The following sections outline the development of PM, including relevant monitoring models, and then the development of nuclear safeguards, including traditional safeguards methods and current research in the integration of PM to nuclear safeguards.

2.1 Process Monitoring (PM)

For successful operation of any process, it is important to detect process upsets, equipment malfunctions or other special events as early as possible and then to find or remove the factors causing those events [11]. The earliest forms of process monitoring (PM) were developed for chemical plants and were referred to as solution monitoring.




Figure 2-1. The basic structure of a process monitoring model. Measured values are used to create a prediction of the measured signals, and then the residuals between the measured and predicted values are used in decision logic.

Prior to the advent of computer, chemical plants relied on destructive analysis techniques, detailed and complicated physics models, and highly trained operators to detect and analysis process upsets. As computers became mainstream and increasingly exploited for scientific investigation, the field of chemometrics emerged. Chemometrics is the science of relating measurements made on a chemical system or process to the state of the system via application of mathematical or statistical methods. As solution monitoring and chemometrics became more widespread, mostly due to advances made in both plant instrumentation and in the computer technology that evaluates it, these techniques were applied to other industries and were more generically referred to as process monitoring. Many chemical problems and early applications of chemometrics involve calibration. Using multivariate statistical techniques, fast, cheap, and non-destructive analytical measurements can be used to estimate sample properties that would otherwise require time-consuming, expensive, destructive testing. As plants become more heavily instrumented, the ability to identify sensors that have fallen out of calibration is increasingly important. As computer technology continued to evolve, newer PM techniques were developed to assist condition monitoring (CM) of plant equipment in real time or near-real time. These techniques are commonly referred to as On-Line Monitoring (OLM) methods [15].

CM can be divided into two major tasks: state estimation and state monitoring. State estimation refers to estimating the current condition, or state, of a system component. Three techniques are generally used in state estimation: redundant signal monitoring, reference signal monitoring, and diverse signal monitoring. State estimation using redundant signals takes the average sensor reading when a large number of redundant signals are available and uses that average sensor reading as an estimate of the system state. State estimation using reference signals involves comparing a sensor's response to a calibrated reference signal to distinguish between sensor drive and process drive. The third technique, state estimation using diverse signals, is useful when a large number of redundant sensors are not available. Instead of determining the system state from a set of redundant signals or a reference signal, the system state is determined from other system sensors that are correlated with the sensor being

monitored. This technique is generally more practical than installing a large number of redundant sensors and relies heavily on multivariate modeling.

Multivariate modeling can be further divided into physical modeling and empirical modeling. Physical modeling relies on knowledge of the physics of the system to predict system states. For instance, by measuring the temperature of a gas in a known volume, the pressure of the gas can be predicted based on established physics relations of gas pressure and temperature. While physical modeling has the advantage that it allows one to predict future system states for a new process operation, it has the disadvantage that it requires detailed knowledge of the physics particular to the system. Alternatively, empirical modeling does not require knowledge of the system's physics, and it does not provide any analyzable relationships between system parameters. Instead, empirical modeling compares current operating conditions to past operating conditions to determine the expected current condition. As chemical facilities grew in size and complexity, empirical models became more widespread for process monitoring due to the difficulty of developing detailed physics models. The disadvantage of empirical modeling is that its predictive range is limited by the past operating conditions. If a new process condition is encountered that is significantly different from the past operating history, the empirical model will be unable to predict the new operating condition. However, as data exemplifying the new operating condition is appended to the known past operating condition, model predictions become more reliable.

Historically, these techniques have been applied specifically for condition monitoring in maintenance scheduling [9, 13]. Recently, a three volume NUREG/CR-6895 series [14-16] was developed for the Nuclear Regulatory Commission (NRC) to provide background, technical guidance and explore implementation issues related to the use of PM for the extension of safety critical sensor calibration intervals. One application of implementing OLM methods is summarized in [14]. Nuclear power utilities began investigating and developing OLM methods that would allow them to move away from periodic maintenance strategies towards condition-based maintenance philosophies. Traditionally manual calibration of nuclear instruments is required through regulation and is performed during each nuclear power plant refueling outage. There

are several weaknesses associated with manual calibration that could be addressed with OLM methods. Manual calibration is the practice of having personnel periodically recalibrate instrumentation regardless of their condition. Previous studies have shown that between calibration periods less than 5% of the 50 to 150 calibrated instruments have fallen out of calibration. Therefore, during a refueling outage around 95% of the calibrations are unnecessary, and performing unnecessary maintenance on properly working equipment provides opportunities for faults to enter the system. Additionally, if a sensor falls out of calibration it remains undetected until the next calibration period. An OLM method could continuously monitor the instrument channel and identify which sensors have degraded. The identified sensors could be identified as needing calibration at the next outage or as entirely inoperable depending on the amount of degradation. Identifying only the instruments that need calibration would reduce the amount of maintenance required during an outage and would limit the opportunity to introduce faults by performing unnecessary maintenance on working equipment. Implementing OLM methods would allow a plant to move develop a condition-based maintenance schedule rather than a periodic maintenance schedule, providing advantages such as reduced radiation exposure, increased instrument reliability, and increased equipment availability.

2.1.1 Parametric Models

There are two common empirical model architectures: parametric and non-parametric. Parametric models identify a set of parameters to define the functional relationship of the system. An example of a common parametric model is linear regression:

$$y = \beta_1 \cdot x_1 + \beta_2 \cdot x_2 + \dots + \beta_n \cdot x_n + \epsilon \quad (2.1)$$

Where y is the prediction, β_n is the n^{th} parameter, x_n is the n^{th} input, and ϵ is an error term. The error term is not a fitted parameter of the model, rather it identifies the error between the prediction, y , and the measurement. For several linear regression models, this equation can be written more compactly in matrix form as:

$$Y = X \cdot \beta + \epsilon \quad (2.2)$$

Where

$$Y = \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix}, X = \begin{bmatrix} x_{11} & \cdots & x_{1p} \\ \vdots & \ddots & \vdots \\ x_{n1} & \cdots & x_{np} \end{bmatrix}, \beta = \begin{pmatrix} \beta_1 \\ \vdots \\ \beta_n \end{pmatrix}, \epsilon = \begin{pmatrix} \epsilon_1 \\ \vdots \\ \epsilon_n \end{pmatrix}, \quad (2.3)$$

Parametric models use historical plant data to create a model training data set and then optimize the parameters (in the case of linear regression, each β) of a predefined mathematical function. The mathematical function used as the prediction model does not need to be linear and will depend on the system being modeled. For instance, population growth is generally modeled with the Malthusian Growth Model, which assumes a mathematical model of the form:

$$Y = X_o \cdot e^{\theta \cdot t} + \epsilon \quad (2.4)$$

Where Y is the predicted population, X_o is the initial population, θ is the parameter to be fit, t is the time, and ϵ is an error term [18]. θ is a measure of growth rate with respect to time. The parameters are normally optimized by minimizing an objective function. For linear regression, the objective function is the sum of the squared error. Once the parameters have been optimized, the model training set is no longer useful, and only the optimized parameters and the predefined mathematical function are saved. When a new query is made, the measured sensor values are fit to the mathematical function, using the previously optimized parameters, and the residual between the new prediction and the measured value are fed to the decision logic. For example, in equation 2.4, once θ is optimized (X_o , the initial parameter is known and the error term not included when making predictions), all future predictions rely only on X_o , θ , and t and not the data that was used to optimize θ . Once a parametric model has been defined, the model is inflexible unless a new training data set becomes available and either the mathematical function is updated and/or the parameters are re-optimized based on the new training set. More information about different parametric models, including models of higher complexity, is available in [19-22].

2.1.2 Kernel Regression


Kernel Regression is an empirical, non-parametric modeling technique. A non-parametric model uses historical data to build a local model. It differs from a parametric model in that every time a new estimation is made, a new model is built from the historical data. Essentially a non-parametric model creates a temporary parametric model, optimizes the temporary parametric model parameters using historical data similar to the input signals, uses the optimized temporary parametric model to perform predictions, and then discards the temporary parametric model. When a new set of input signals is available, the process repeats. The process of using a parametric and non-parametric prediction model is outlined in Figure 2-2. Instead of saving a set of parameters, a non-parametric model saves the historical data and parameters are re-optimized for every prediction and then discarded after every prediction. In the case of kernel regression, instead of creating a temporary parametric model, a prediction is performed using a weighted sum of historical values. Kernel regression can be formally defined as the process of estimation using a weighted average of historical exemplar observations [23]:

An important first step creating a kernel regression model is selecting which data should become part of the model training data set. Only historical data that is a close representation of the query data becomes part of the model training data set. Determining the similarity between the query data and the historical data is performed through a distance function. The most common distance function is the Euclidean distance, also known as the L^2 -norm, and is given by:

$$d_i(X_i, x) = \sqrt{(X_i - x)^2} \quad (2.5)$$

The distance between every historical data record and the query vector is calculated, and the historical data records with the minimum distance are used in the model training set. While the L^2 -norm is the most common distance measure, several more robust distance measures have been explored, including the L^1 -norm [24].

Parametric Model



Non Parametric Model




Figure 2-2. Outline of a parametric and non-parametric empirical model. A parametric model optimizes parameters of a mathematical model once and then uses those parameters for each subsequent prediction. The model training data set is not saved for future use. A non-parametric model uses historical operation data to build a parametric model for each prediction. Each new prediction requires a new parametric model, and the optimized parameters are not saved. The historical operation data is saved, but the set of optimized parametric parameters is discarded after predictions are made.

In Kernel regression, estimation is performed based on a weighted average of similar historical data. Once the model training data set has been selected, each record in the model training data set is assigned a weight. Records that are very similar to the query vector (and, therefore, have a small distance) should receive a very high weight, while less similar records should receive a smaller weight. The most commonly used kernel function is the Gaussian Kernel, given by [25]:

$$K_h(d) = \frac{1}{\sqrt{2\pi}h} e^{\frac{-d^2}{2h^2}} \quad (2.6)$$

Where h is the kernel's bandwidth. The bandwidth determines how small the distance, d , must be in order to generate a large weight. A smaller bandwidth will only generate high weights when the distance is close to zero. The bandwidth is usually chosen to minimize model uncertainty, given by:

$$\text{uncertainty} = \sqrt{\text{bias}^2 + \text{variance}^2}. \quad (2.7)$$

The effects of over or under estimating the bandwidth can be seen in Figure 2-3, and a properly sized bandwidth is shown in Figure 2-4. Too large of a bandwidth can result in a larger prediction bias, and too small of a bandwidth will increase the prediction variance. The bandwidth is generally optimized so that the quadratic sum of the bias and prediction uncertainty is minimized. The Gaussian kernel is only one example of a kernel function. Other kernel functions include inverse distance, exponential, absolute exponential, uniform weighting, triangular, biquadratic, and tricube [23]. While other kernel functions may perform better in specific applications, the Gaussian kernel is generally adequate.

Because the kernel regression, and empirical modeling in general, does not require detailed knowledge of a facility can be easily applied to several different types of facilities. Historically, it was used in large chemical facilities. These facilities are highly instrumented and the relationships between measurements are often non-linear and too complex for the development a physical model. However, kernel regression can be used in any situation where a physical model isn't available. For instance, reference




Figure 2-3. The effect of incorrect Gaussian Kernel bandwidth selection. Too large of a bandwidth (top) results in a larger range of distances receiving a non-trivial weight, which can introduce a large prediction bias, where the peaks and values of the data are not well fit. Too small of a bandwidth (bottom) results in large weights only when the distance is small, and can introduce a large variance [26]. An optimal bandwidth minimizes the model uncertainty.




Figure 2-4. A properly sized bandwidth balances the effect of prediction bias and model variance. The result is that the peaks and values of the function are well fit and that the overall model uncertainty is minimized [26].

[17] describes the use of various ecological variables, such as temperature, wind speed, precipitation, and relative humidity to predict the area burned in wildfires.

2.1.3 Data Reconciliation

Data reconciliation (DR) is a measurement correction technique based on the known physical laws of system and the instrument uncertainties. Known material and energy balances serve as additional constraints to improve measurement accuracy by reducing random errors associated with measurements. DR produces reconciled (or corrected) estimates that are consistent with the known material and energy conservation laws and are inherently more accurate when no systematic error exists [14]. An example would be liquid transfer between two tanks. Under normal operation the liquid leaving tank 1 must enter tank 2; however, because of random measurement error, the measured amount of liquid leaving tank 1 and entering tank 2 may not be exactly equal. DR would reconcile the two measurement values so that the measured values would be equal.

The first application of DR was by Kuehn and Davidson [27] of the IMB Corporation. Kuehn used DR on a linear material balance system in which all parameters were measured. DR has since been applied to dynamic processes [28], non-linear models [29], and recently to processes that are both dynamic and non-linear [30].

DR can generally be written as a constrained, weighted, least squares optimization problem:

$$\text{Min } \sum_i \left(\frac{y_i^* - y_i}{\sigma_i} \right)^2 \quad (2.8)$$

where y_i^* is the reconciled value, y_i is the measurement, and σ_i is the measurement standard deviation. The weights are proportional to the inverse of each measurement's accuracy. Most DR techniques assume that measurements contain only a random error component with a Gaussian distribution. If systematic errors exist, DR may reconcile

the measurements in such a way that the reconciled values are more inaccurate than the raw measurements [31]. Therefore, systematic error detection techniques are generally coupled with DR to remove and eliminate systematic errors.

As an example of DR, consider a simple flow splitter, shown in Figure 2-5. In this system, under normal operation, the mass flow out of M1 must equal the mass flow into M2 and M3. However, because every measurement has some random noise, the measured values will not exactly obey the physical relationship that $M1 = M2 + M3$. DR is used to calculate a correlation factor that is applied to the measured value so that they obey the physical relationship. The corrected values are referred to as reconciled values.

Data reconciliation calculates a correction factor, ν , that is applied to measured values to force the system to follow known conservation laws. The relationship between the reconciled values and the measured values is given by:

$$\bar{x} = x + \nu \quad (2.9)$$

where $\bar{x}$ is the reconciled values, x is the measured values, and ν is the corrective factor to be determined. The corrective factor is applied to the measured values to force the measured values to follow conservation laws, such as $M1 = M2 + M3$. The solution to equation 2.8 requires ν to be minimized. In matrix notation, and substituting ν for $y_i^* - y_i$, and $\frac{1}{\sigma^2}$ for S_x^{-1} , the objective function can be expressed as:

$$\nu^t \cdot S_x^{-1} \cdot \nu \quad (2.10)$$

where ν is the vector of corrective factors, S_x^{-1} is a diagonal matrix of σ_x^{-2} . For a first order, linear system the minimizing solution, through the application of Lagrange multipliers, is :

$$\nu_{min} = -\left(\frac{\partial f}{\partial \bar{x}} S_x\right)^T \cdot \left(\frac{\partial f}{\partial \bar{x}} S_x \cdot \left(\frac{\partial f}{\partial \bar{x}}\right)^T\right)^{-1} \cdot f(x) \quad (2.11)$$




Figure 2-5. Diagram of a simple flow splitter. In this system the mass flow out of M1 must equal the mass flow into M2 and M3. However, due to random measurement error, the measured mass flow rate may not exactly represent the relationship $M1 = M2 + M3$. Data Reconciliation is used to calculate a correction factor to the measured values so that they obey the physical relationship.

The correction factors are weighted by both the standard deviation and the magnitude of the measurement value. So measurements with higher precision, and therefore a lower σ , have a lower ν . Equation 2.11 is substituted into equation 2.8, which is then used to reconcile the measured values. A summary of applying this correction factor to the hypothetical system in Figure 2-5 is shown in Table 2-1.

The data in Table 2-1 shows that the measured values of M1, M2, and M3 do not obey the conservation law initially, but after applying data reconciliation the reconciled values do obey the mass conservation law.. Also the size of the correction factor is related to both the standard deviation and the magnitude of the measured values, which explains why M1 had a much larger correction factor than either M2 or M3. It is important to note that DR assumes the data is under normal operation, the sensors are not faulted, and the conservation laws are actually true (which would not be the case in a diversion scenario). If these assumptions are not met, DR can actually increase the error and uncertainty of the measured values [31]. In this sense, DR can be thought of as a de-noising technique which relates the measured values to their most likely true values (i.e. removes the random noise component).

Data reconciliation was first applied only to linear processes, such as mass balance problems seen in almost all types of nuclear facilities. As the techniques matured, the areas of application increased. For instance, data reconciliation is used in several types of power plants (coal, nuclear, natural gas, etc.) where both non-linear and dynamic processes take place, such as plant start-up, load changing, and energy balances within a combustion chamber. Specific to nuclear power plants, data reconciliation has successfully been used to reconcile flow rates into and out of the reactor core to provide a better measurement of reactor power, which enables the plant to operate closer to its maximum capacity throughout the entire year [14].

Table 2-1. An application of data reconciliation to a hypothetical flow splitter. The physical system is constrained by $M1 = M2 + M3$, but due to measurement error the measured values do not exactly obey the constraining law. Data reconciliation is used to calculate a correction factor for each measured value (weighted by both the magnitude of the measured value and the standard deviation of the measurement), to force the measured values to follow the conservation law.

Sensor	Measured Value, $f(x)$	Standard Deviation	Correction Factor (v)	Reconciled Value, $f(\bar{x})$
M1	500	25	-3.36	496.64
M2	245	12.25	0.81	245.81
M3	250	12.5	0.84	250.84
M1-M2-M3	5			0

2.1.4 Sequential Probability Ratio Test

The sequential probability ratio test (SPRT) is a sequential hypothesis test used as a fault detection technique that was developed by Abraham Wald [36]. The SPRT is useful for determining if a sequence of residuals indicates the system is in a degraded state. The sequences of residuals being generated by a random process (such as measurement noise) should have a mean of zero, but if the system is degraded there will be a shift in the mean or variance of the residuals. The SPRT determines whether a sequence of residuals indicates a faulted condition by calculating the log likelihood ratio. According to the central limit theorem, the measurement residuals can be assumed to be normally distributed. Using this assumption, the SPRT determines whether or not a sequence of numbers comes from a specified normal distribution or not. In the case of fault detection, the unfaulted normal distribution's mean and variance would be determined by training the model. In this case the SPRT determines whether or not the sequence of residuals comes from unfaulted normal distribution. If the sequence fails the SPRT, then the data is in a faulted state. Complete derivation of the log likelihood ratios is outside the scope of this work, but it can be found in [35]. The SPRT determines how small of a drift can be detected with a specified false alarm rate by calculating:

$$\lambda_m = \lambda_{m-1} + \frac{m_i}{\sigma^2} \cdot \left(s_m - \frac{m_i}{2} \right) \quad (2.12)$$

where m_i is the mean, s_m is the residual at time m , and σ^2 is the variance. If λ is within certain bounds, then the residuals are said to come from unfaulted data. The bounds are calculated by specifying a false alarm rate, α , and missed alarm rate, β .

$$A = \ln \left(\frac{\beta}{1-\alpha} \right) \quad (2.13)$$

$$B = \ln \left(\frac{1-\beta}{\alpha} \right) \quad (2.14)$$

If λ is less than A , then the residuals are said to come from unfaulted data, if λ is greater than B , the data is assumed faulted. If λ is between A and B the test is inconclusive.

2.2 Nuclear Safeguards

The initial development of nuclear technology was for military purposes. After World War II, the attention of nuclear technology turned to civil applications. Nuclear safeguards refers to the development of technology and policy to assure that special nuclear material is used for its intended peaceful applications. In 1953, in his "Atoms for Peace" address to the UN General Assembly, Dwight D. Eisenhower proposed the creation of an international body to both regulate and promote the peaceful use of atomic power [18]. In the following year, the United States proposed the creation of an international agency to take control of all fissile material and serve as a nuclear bank [36]. From this proposal the International Atomic Energy Agency (IAEA) was born. However, international disagreements prevented the IAEA from becoming the nuclear bank originally proposed. Instead, the IAEA became tasked with ensuring that nuclear material is not diverted from its intended peaceful uses. To this end, the IAEA is charged with: the inspection of facilities to ensure their peaceful use, providing information and developing standards to ensure the safety and security of nuclear facilities, and as a scientific center for fields involving the peaceful application of nuclear technology.

Parties who have signed the Nuclear Non-Proliferation Treaty (NPT) are required to enter a safeguards agreement with the IAEA. The NPT is a treaty to limit the spread, or proliferation, of nuclear weapons. The treaty first came into effect in March of 1970. The NPT is generally interpreted as a three-pillar system: non-proliferation, disarmament, and the right to peacefully use nuclear technology. The first pillar states that the recognized nuclear states will not transfer nuclear weapons technology to non-nuclear weapon states or encourage non-nuclear weapon states to acquire nuclear weapons and that non-nuclear weapon states will not seek to acquire or manufacture nuclear weapons. The second pillar is to reduce the worldwide stockpile of nuclear weapons. The third pillar recognizes that non-nuclear weapon states have the right to peacefully use nuclear technology.

The basic measure by which the IAEA verifies the fulfillment of the NPT obligations is through nuclear material accountancy [3]. The availability of fissile material has long been considered the principal obstacle to a country's nuclear weapons development effort. The spread of enrichment and reprocessing technology allows countries to develop their own nuclear technology for the peaceful purposes stated in the NPT, but it also inherently gives those countries the capabilities of producing fissile material for weapons development. The IAEA routinely performs safeguards activities at over 900 facilities in over 71 countries, including power reactors, research reactors, conversion plants, fuel fabrication plants, enrichment plants, storage facilities and reprocessing plants, in addition to other facilities [37].

Traditional safeguards rely on defining MBAs within a plant and tracking the material flow between MBAs. Periodically, full inventories are conducted to verify the amount of material within a MBA rather than just the amount of material entering and leaving a MBA. Since most facilities deal with uniquely labeled components, such as fuel assemblies, and the nuclear components are not easily subdivided, simply ensuring tracking the component as a whole is a sufficient accountability measure. However, GCEPs and reprocessing plants are unique in that they don't deal with uniquely labeled components. Special nuclear material is mixed and divided throughout the plants and proper MCA requires taking measurements at key locations at the boundaries of MBAs [7, 8]. A sample MBA for a GCEP is shown in Figure 2-6. Typical key measurements include concentrations, isotopic compositions, and volume, density and/or weight. Instrumentation and measurement techniques including alpha spectrometry, calorimetry, chemical titration, K-edge absorption densitometry, manometers and vibrating-tube densitometers, neutron techniques, spectrophotometry, uranium gravimetry, mass spectrometry, x-ray fluorescence, and gamma-ray spectrometry [7]. Material inventory must be accounted for to within one significant quantity of material within an appropriate time period [6].

The Rokkasho plant in Japan is a commercial-sized reprocessing plant. It will have the capacity to reprocess 800 metric tons of spent fuel a year. The average spent




Figure 2-6. Block diagram of material flow through a hypothetical gas enrichment plant. As material is received, its contents are verified and it is put into storage. When material is ready to be processed, it is moved from storage to an input accountability tank, thereby crossing from MBA 1 to MBA 2. After the material is processed, it is removed from the process line and placed in storage, thereby crossing from MBA 2 to MBA 3, where it stays until it is shipped offsite. At each MBA crossing, the MCA is done through destructive assay and key flow measurements.

fuel is approximately 0.9% plutonium by weight, meaning that the plutonium throughput is 7.2 metric tons per year. Assuming a realistic measurement uncertainty of 1.0%, 72 kg of plutonium falls within the annual 1-sigma error bounds [7]. A significant quantity of plutonium is defined as 8 kg, so the Rokkasho plant must perform a monthly plutonium inventory balance to keep the 1-sigma error bound below a significant quantity. Regardless of measurement uncertainty, a high enough material throughputs could force a facility to perform more frequent inventory balances to ensure the 1-sigma error bound stays below a significant quantity of material.

As commercial plants continue to spread and grow in size, eventually plant throughput will reach a point that even monthly inventory balances will not have the precision to account for one significant quantity of material. Also, no additional verification is performed to ensure cumulative uncertainties don't reach a significant quantity. Improvements in measurement precision will reduce the material uncertainty, but many of the more accurate classical techniques have a long wait time (3 – 4 months) due to transportation of samples off-site and the time it takes to do the analysis. For these reasons, many experts have called for improvements in safeguards methods to reconcile these and other difficulties and uncertainties [38-40].

Many of the improvements envisioned require drastic increases in manpower, technology or both. Shortening the material accountancy period or increasing the number of material balance areas requires many more person-hours to accomplish a timely analysis. Currently, the Rokkasho reprocessing plant requires between 1000 and 1200 Person Days of Inspection (PDI) out of the approximately 10,000 PDIs deployed by the IAEA's Safeguards Department worldwide [7]. A standard GCEP requires an annual inspection with 2-4 IAEA inspectors and monthly interim inspections is 1-2 inspectors. The annual inspection may take between 10 and 15 days, while the monthly interim inspection only requires at most 4 days. Additional process information would provide inspectors means for verifying the process and preventing diversion of nuclear material, without the need for drastic increases in manpower or technology [38].

2.2.1 Process Monitoring for Nuclear Safeguards

While a number of evaluation systems have been developed, process monitoring for safeguards is still in its infancy. Certain systems, through experiences on real facilities, are evolving algorithms useful for evaluating real-world scenarios rather than idealized simulations [41, 42]. The two types of facilities most examined are a gas centrifuge enrichment plant and reprocessing plants.

Process monitoring as an international safeguards tool has its root in the late 1970s. Prior to the 1970s, reprocessing plants were small and limited to the nuclear weapon states. By the late 1970s, the Tokai Reprocessing Plant in Japan was nearing completion and there was the possibility of a large-scale commercial reprocessing plant at Barnwell in the United States. Several countries with reprocessing technology, including France, Germany, Great Britain, the United States, and Japan, joined the IAEA in the Tokai Advanced Safeguards Exercise (TASTEX) to investigate advanced techniques to improve international safeguards capabilities for larger commercial plants. One of the recommendations was for the implementation of process monitoring. Research originated at the Idaho Chemical Processing Plant to detect valve setting errors and prevent inadvertent transfers between tanks [12]. The nature of the instrumentation of the Tokai made interfacing for digital data collection difficult. However, the Barnwell plant was built with more modern instrumentation, including a computer-based data collection system, so development of process monitoring continued at that facility. Further development continued at the Integrated Equipment Test (IET) facility at the Oak Ridge National Laboratory, but the reprocessing equipment developed there was never intended for commercial operation. The earliest efforts of process monitoring were referred to as solution monitoring because they concentrated on monitoring vessel volumes and the flow of solution between tanks. The development of process evaluation software lagged significantly behind. However, research at Barnwell and at the EIT tried to predict special nuclear material concentrations based on density measurements and some acid concentration measurements [54].

Until the early 1990's, the techniques developed under TASTEX remained the international standard. Applications expanded slightly as the scope of facility operations expanded, but measurement and data-collection was patterned after the Tokai system. All data was collected on-site and for the most part evaluation was based on the capabilities of the inspector to use packaged graphical interface routines. Without the development of process evaluation software, the inspectors were left with a wealth of data but very limited ways to efficiently analyze it. However, in the early 2000's an IAEA and Glasgow University collaboration made significant progress in evaluation software. The software was used to extract pertinent features of the data, observe specific plant operations, and to detect and isolate anomalies. It still required that knowledgeable inspectors interact with qualitative graphical representations, but it also represented the first real effort to provide inspectors with both quantitative and partially diagnosed information [54].

Currently, France, Japan, Russia, and the United Kingdom operate commercial reprocessing plants, but only the Japanese plants at Rokkasho and Tokai fall under IAEA safeguards. In existing reprocessing plants, operators have extensive and cutting edge systems for monitoring and control including cameras, tamper-indicating devices, radiation detectors, and the typical process control instrumentation. While this broad range of data could provide very detailed information for the purposes of verification, operators have traditionally been reluctant to allow the IAEA access to the data due to concerns about revealing proprietary information [12].

The first application of process monitoring for safeguards was in reprocessing facilities, but in the early 1970's GCEPs were also a major concern. Initially safeguards at GCEPs did not involve access to the cascade halls. Recognizing the minimal effort required for a commercial GCEP to be reconfigured to produce highly enriched uranium, the Hexapartite Safeguards Project (HSP) convened in the early 1980's with the objective of developing an effective safeguards approach for GCEPs. The HSP approach combined traditional nuclear accountancy measures outside the cascade hall with limited frequency, unannounced access (LFUA) to the cascade hall. The traditional measures outside the cascade hall mainly consisted of weight measurements of the

feed, tail, and product tanks in addition to the verification that the accountability scales are properly calibrated. Inside the cascade, measures included non-destructive assay measurements on piping within the cascade hall and visual observation to verify design information. In the early 1990's, gas centrifuge technology reached a stage where higher enrichment levels could be produced without visible changes to the cascade design, leading to the development of a continuous enrichment monitoring system comparable in purpose to the process monitoring systems developed for reprocessing facilities [45].

Safeguards at GCEPs focus on determining whether or not undeclared uranium is being enriched. The most prevalent method is through load cell monitoring to verify nuclear material accountancy. The proposed systems analyze cylinder weight data that is collected frequently from cylinders located in feed, tails, and product stations. Research exploring PM at gaseous centrifuge enrichment plants for the detection of undeclared production by monitoring the enrichment cascade and the load cell stations is described in [43-48]. A typical load cell profile is shown in Figure 2-7.

The profile shown in Figure 2-7 lacks any noise components that would be seen in an actual plant. An actual load cell profile would be contaminated with noise, most significantly seen as a period of instability when a tank is first loaded or removed. Material accountancy is then simply implemented using:

$$ID = \sum \Delta_{\text{feed stations}} - \left(\sum \Delta_{\text{tail stations}} + \sum \Delta_{\text{product stations}} \right) \quad 2.10$$

where ID is the inventory difference, $\sum \Delta_{\text{feed stations}}$ is the summation of the change of tank weights for all the feed stations, $\sum \Delta_{\text{product stations}}$ is the summation of the change of tanks weights for all product stations, and $\sum \Delta_{\text{tail stations}}$ the summation of the change of tanks weights for all tail stations. When the plant inventory is perfectly balanced, ID will equal zero. In reality, all plant experience varying amount of material holdup in pipes, valves, and other components, so that ID is never zero. However, by characterizing the




Figure 2-7. A hypothetical load cell profile as a feed tank is loaded, processed, and removed from a GCEP feed station. At an actual plant, the profile would be contaminated with noise spikes, most significantly when a tank is first loaded or removed.

expected facility holdup, equation 2.10 can be expanded to include a facility holdup term and a more meaningful value for ID is achieved.

Accurately determining $\sum \Delta_{stations}^{feed}$, $\sum \Delta_{stations}^{product}$, and $\sum \Delta_{stations}^{tail}$ first requires that the load cell data can be matched to a template to accurately identify the number of cylinders placed on a feed, tail, or product station. Then correctly identifying the initial and final tank weights is used to determine the material throughput. Currently, the IAEA does not make use of load cell data for safeguard purposes, instead relying solely on authenticated accountability scales. Incorporating load cell data with authenticated accountability measurements provides a measure of validation that no undeclared material was processed. However, operators have two key concerns: the unnecessary transmission of data out of a facility, and unnecessary inspector access to detailed process monitoring. In response to these concerns, algorithms have been proposed to automatically and locally count the number of cylinders at a GCEP station, thereby limiting the transmission of data offsite.

An outline using template matching to count the number of cylinders processed at a load cell station is summarized from John Howell's paper entitled "Algorithms to Count the Cylinder Throughput of a GCEP Feed Station" [49]. For this template matching procedure, it is assumed that the GCEP operator will allow local automated feed cylinder counting from their load cell data, but that the operator is under no obligation to operate a facility in a prescribed manner. For instance, the operator may: (1) place a part-filled cylinder into a station, (2) temporarily halt station output, (3) remove a cylinder before it is completely empty, or (4) vary the rate at which a cylinder is emptied. The load cell data would consist of load cell weights collected frequently and at regular intervals. This data would be saved locally in a fixed format. The counting algorithms would search through the data and extract the approximate times when cylinders were changed and their initial and final weight. The basic profile of a feed station was shown in Figure 2-7. The basic profile changes under different operations are shown in Figure 2-8.




Figure 2-8. Changes in a GCEP feed station profile under different operating conditions. The spikes when a tank is added or removed from the scale represent expected fluctuations.

The template outlined from Howell's paper [49] is shown in Figure 2-9. Section A marks the period when a cylinder is emptying, which is a function of the feed rate (possibly unspecified and varying), the initial cylinder weight, and the final cylinder weight. Section B marks the period when an empty cylinder still resides on the station load cells. Section C marks the delay before a new feed cylinder is placed on the station. Section D marks the period when a new cylinder is first placed on the station and the weight value is still fluctuating. Section E marks the period when a new cylinder is ready to be emptied. The weight value may fluctuate due to various processes associated with bringing a new cylinder online. With the exception of section A, all other sections would be of concern only to the operator, since section A contains all the information about the amount of material processed and the initial and final tank weight.

Each template would be specified using predefined parameters corresponding to a weight threshold, a change of weight threshold, a time threshold, or a combination of all three. For instance, section C is specified by setting a lower weight threshold corresponding to an empty scale weight (with some error allowance for measurement noise) and a length of time that the feed station must remain in section C. Section B could then be specifying that a lower threshold associated with the weight of an empty tank and by specifying that the data is not in section A, D, or E. Additional templates would be created for partially filled cylinders being placed on the station or for a cylinder being withdrawn prematurely. Once the templates have been matched, the number of cylinders processed can easily and automatically be counted by counting the number of times a specific template has been matched. With this automated system, the need for manual analysis would be limited to cases where the number of cylinders reported is not in agreement with the number of cylinders declared by the operator.

Cylinder counting could be used as part of a wider scheme that collects, analyzes and integrates data from the load cells, accountancy weigh scales, and other key measurements taken in the facility. Several key processes are already collected by the operator for use in process monitoring, but are not specifically available to the IAEA for safeguards. Research in safeguards at reprocessing facilities has focused on solution monitoring. Solution monitoring is a specific form of process monitoring for




Figure 2-9. A general feed template for a GCEP feed station. Section A marks the period when a cylinder is emptying. Section B marks the period when an empty cylinder still resides on the station load cells. Section C marks the delay before a new feed cylinder is placed on the station. Section D marks the period when a new cylinder is first placed on the station, and the weight value is still fluctuating. Section E marks the period when a new cylinder is ready to be emptied [49].

chemical facilities. As material is transferred from one process tank to the next, process variables are recorded. The standard process measurements are level, density and temperature of a solution tank. Several of these measurements are highly correlated with the flow of SNM through a facility and therefore could naturally lend themselves to drawing safeguards conclusions. For example, consider the three tank example outline in [50]. In this example, a generic feed, accountability, and buffer tank are connected in series, as shown in Figure 2-10. Suppose that flow rates F_{11} and F_{12} into the feed tank, F_{21} and F_{22} into the accountability tank, and F_{31} and F_{32} into the buffer tank and F_{41} out of the buffer tank are measured in real time and recorded every 5 minutes. The temperature (C°), density (ρ), and level (which is calibrated to the volume V), are also recorded every 5 minutes. Different assumptions regarding the availability of measurements of the nitric acid concentration H , the Plutonium concentration Pu , and the Uranium concentration U can be examined. Typically concentrations are not measured in-line, but rather by taking a stream sample and analyzing it.

Traditional nuclear materials accountability would compare the SNM mass entering the feed tank with the SNM mass exiting the buffer tank. The throughput in tank 2 would be ignored. The inventory difference for the material balance is defined as:

$$ID = T_{in} + I_{begin} - T_{out} - I_{end} \quad 2.11$$

Where T_{in} and T_{out} are transfers into and out of the tank, and I_{begin} and I_{end} are the beginning and ending material inventory (or level) in the tank. The inventory difference reported by the operator, $ID_{operator}$, is calculated using equation 2.11. However, it is not feasible for an inspector to re-measure every tank. Therefore the inspector will select a random subset of tanks from. The inspector will then independently measure each of those tanks and compare them to the operator's measurement. A difference statistic, D , is calculate by:

$$D = N \sum_{j=1}^n (o_j - i_j) / n \quad 2.12$$

MBA Boundary




Figure 2-10. A generic feed, accountability, and Buffer tank connected in series [50]. A material balance area surrounds the three tanks. The flow rates are denoted by F , the temperature by C° , the density by ρ , and the volume by V . Measurements of the nitric acid concentration, plutonium concentration, and uranium concentration are denoted by H , Pu , and U respectively [50].

Where N is the total number of tanks, n is the subsample size, o_j is the operator's measurement of sample j , and i_j is the inspector's measurement of sample j . The inventory difference as calculated by the inspector is then expressed as:

$$ID_{\text{inspector}} = ID_{\text{operator}} - D \quad 2.13$$

One limitation of this approach is that the uncertainty of ID_{operator} is generally smaller (due in part to the larger number of measurements available) than the uncertainty of D and therefore the uncertainty of $ID_{\text{inspector}}$. This lower uncertainty of ID_{operator} means that the operator is better able to detect diversions compared with an inspector. For example, at the Rokkasho plant 12 important vessels are equipped with instruments installed and controlled by the IAEA (in addition to instruments installed by the operator) to make independent and externally authenticated measurements (for use in calculating the D statistic). However, approximately 80 vessels only use operator's instruments and, therefore, are not authenticated. Process monitoring provides a means to incorporate unauthenticated operator data with authenticated inspector data to lower the uncertainty of $ID_{\text{inspector}}$ and therefore increase the detection probability of unauthorized activity. In the three tank example, errors from one measurement would be correlated with errors in subsequent measurements. For example, if the measured flow rate F_{21} into tank two is too high (meaning the actual flow rate into tank two is lower than measured), then the predicted volume in tank 1 would be too small, but the predicted value in tank 2 would be too high. Monitoring these errors using multivariate statistical methods could ensure that the 3 tanks are operated as declared, providing a measure of authentication of each of the operator's measurement. In this sense, process monitoring is self-authenticating. Several other examples of solution monitoring for safeguards are available in [51-53].

Applying process monitoring techniques to safeguards measurements provides a cost-effective way to reduce the burden of IAEA inspectors. Using the self-authenticating nature of process monitoring, the need for additionally independent and externally authenticated measurements is reduced and provides a possible way for remote monitoring of a facility. Additionally, more sophisticated monitoring techniques

would lower the uncertainty of $ID_{inspector}$ which could translate into extending the intervals between inspections, further reducing the IAEA manpower needed to monitor a single facility.

While initial safeguards approaches for reprocessing facilities and for GCEPs had different focuses and specific techniques, several of the techniques are equally applicable for both facilities. For example, the load cell algorithms developed for GCEP monitoring could be applied to tank monitoring at reprocessing facilities, and the solution monitoring techniques developed for reprocessing facilities could be adapted to monitor the gas flow inside a GCEP cascade hall. While on-line monitoring has been well established in other areas, its application to nuclear safeguards is still new. Current research implementing process monitoring for safeguards at both reprocessing facilities and GCEPs is being conducted by Tom Burr of the Los Alamos National Laboratory and John Howell of the University of Glasgow, UK. Several other laboratories and agencies are involved in this research including: the Oak Ridge National Laboratory [56 - 59], The U.S. Department of Energy, The International Atomic Energy Agency, and the Nuclear Nonproliferation and Technology Center, Japan. A complete review on PM for safeguards, including international collaborations, is available in [54, 55].

3 METHODOLOGY

Traditional safeguards techniques rely on material accountancy. Improvements in monitoring techniques can reduce the burden of IAEA inspectors by extending the interval between inspections [38]. Historically, process monitoring techniques have been applied specifically for condition monitoring for maintenance scheduling [9, 13]. More recently a three volume NUREG/CR-6895 series [14-16] was developed for the Nuclear Regulatory Commission (NRC) to provide background, technical guidance and explore implementation issues related to the use of PM for the extension of safety critical sensor calibration intervals. The application and expansion of PM techniques for safeguards could provide a framework for reducing the burden on IAEA inspectors without compromising their nuclear safeguards mission.

3.1 Safeguards Monitoring of a Simulated Uranium Blend-Down Facility

The second pillar of the Nuclear Non-Proliferation Treaty calls for a reduction in the number of nuclear warheads held by the nuclear states. To this end, the United States and Russia signed an HEU Purchase Agreement in 1993. Under this agreement, Russia would dismantle part of their nuclear stockpile, blend-down their highly enriched uranium (HEU) components to reactor grade uranium (approximately 4% ^{235}U enriched), and sell the reactor grade uranium to the United States as reactor fuel. The Blend-Down Monitoring System (BDMS) was designed to provide verification that the requirements of the HEU Purchase agreement were being met. To monitor the ^{235}U flow, the Oak Ridge National Laboratory (ORNL) developed the Fissile Mass Flow Meter (FMFM). The FMFM measures the delayed gammas from induced fissions of ^{235}U . Figure 3-1 shows a schematic of the FMFM used in the BDMS. THE FMFM is a two part system. The first part is a neutron source assembly. The FMFM neutron source is a Californium-252 used to induce fissions in a uranium hexafluoride stream. The source is incased in polyethylene to moderate the source neutrons. During operation a shutter is opened to allow moderated neutrons to enter uranium hexafluoride stream and induce fissions. The second part of the assembly is a shielded

FMFM Source Modulator Subassembly

FMFM Detector Assembly




Figure 3-1. A schematic of the Fissile Mass Flow Meter (FMFM) used in the blend-down monitoring system. When the shutter is opened a Californium-252 neutron source induces fissions in a uranium hexafluoride stream. Downstream of the neutron source is a shielded detector array used to detect delayed gammas from Uranium-235 fission. The number of delayed gammas detected is proportional to the amount of Uranium-235 in the stream. The time delay between the source shutter opening and the initial detection of fission gammas can be used to calculate the flow velocity.

detector located downstream from the neutron source. The delayed gammas from Uranium-235 fission are detected as the stream passes the detector. The number of delayed gammas detected is proportional to the amount of Uranium-235 in the stream. The signal from the FMFM is the number of counts per second [56]. Additionally, the time delay between the source shutter opening and the initial detection of fission gammas can be used to calculate the flow velocity. A complete explanation of the BDMS and the FMFM can be found in [56 - 59].

A MATLAB Simulink model of a uranium blend-down facility was used to create simulation data for modeling and testing. A full description of relevant equations used to develop the Simulink model is given in Appendix A. Conventional process monitoring methods were augmented through the development of algorithms that allow for the addition of the FMFM signal and through data reconciliation techniques. In this hypothetical uranium blend-down facility, uranium of two different enrichments is blended into a product with a desired target enrichment. The facility is set up to run in a batch mode, where tanks of high-enriched uranium-hexafluoride (HEUF_6) and low enriched uranium-hexafluoride (LEUF_6) are fed through instrumented lines, blended together, and then stored as product low-enriched uranium-hexafluoride (PLEUF_6). At standard normal atmospheric pressure and temperature, uranium-hexafluoride is nominally a solid, but for the purpose of these simulations it was assumed gaseous. The combination of sensors are used to monitor the process are shown in Figure 3-2. All three legs are instrumented with flow meters and weight sensors. Additionally the HEUF_6 and PLEUF_6 are instrumented with FMFMs.

The simulated measurements from the three flow meters, three weight sensors, and two FMFMs were used to create empirical auto-associative kernel regression (AAKR) models that encompass normal operating conditions. Velocities were varied to simulate actual operations, but enrichment was held constant. Abnormal conditions were simulated by allowing for the diversion of material from the HEUF_6 leg just before the blending tee. Because the removal happens after the HEUF_6 leg instrumentation, its effects should only be seen in the PLEUF_6 leg instrumentation. Table 3-1 summarizes the modeled instrumentation on the hypothetical BDMS system.




Figure 3-2. Schematic of the simulated uranium blend-down facility. All legs are instrumented with flow sensors and tank weight sensors. Additionally, the HEUF₆ and PLEUF₆ legs are instrumented with FMFMs. Diversions were simulated by removing material from the HEUF₆ just before the blending tee.

Table 3-1. Summary of the simulated instrumentation of the Blend Down Facility.

Sensor Number	Sensor Type
1	HEUF ₆ weight sensor (g)
2	LEUF ₆ weight sensor (g)
3	PLEUF ₆ weight sensor (g)
4	HEUF ₆ flow meter (cm/s)
5	LEUF ₆ flow meter (cm/s)
6	PLEUF ₆ flow meter (cm/s)
7	Fissile Mass Flow meter on HEUF ₆ leg (gU ²³⁵ /sec)
8	Fissile Mass Flow meter on PLEUF ₆ leg (gU ²³⁵ /sec)

To simulated actual measurements, 1.0 % process noise was added to both the HEUF₆ and LEUF₆ fluid velocities. Additionally, independent sensor noise was added to each sensor. The flow meters and weight sensors had 1.0% and 0.1% respectively of their maximum sensor reading added as Gaussian noise. Uncertainty for radiation detectors is not quantified the same way as process sensors, but for simulation simplicity the FMFM was assumed to have 1.0% Gaussian noise. Six scenarios were simulated in which the magnitudes of the HEUF₆ diversion were varied. In the first scenario there was no HEUF₆ diversion. The remaining five scenarios had HEUF₆ diversions of 0.1%, 0.5%, 1.0%, 2.0%, and 10% respectively.

3.1.1 Building an AAKR Model

Four different models were created. In each model, an AAKR architecture was used. The first model used only data from the flow meters and weights sensors. The second model used the same data as the first model, but data reconciliation techniques were used. The third model did not use data reconciliation techniques, but the AAKR architecture was expanded to include FMFM signals. Finally, the fourth model incorporated both the FMFM signals and used data reconciliation. The objective of these models was to: 1) quantify the benefits of incorporating radiation sensors into an AAKR model and, 2) quantify the benefits of incorporating data reconciliation with the AAKR model. The main metric will be the ability of the each model to detect a HEUF₆ diversion. A summary of the different models is given in Table 3-2. A block diagram of the models is given in Figure 3-3.

The first step to building an AAKR model is to select which measurements should be included as inputs. These variables were selected based on their correlations. Measurements that are highly correlated perform better as predictors. In the simulated facility, all 8 sensors are highly correlated (correlations coefficients are all above 0.7), so all will be included in the model. However, for a real world application with a more complicated system, there are likely to be measurements that are not well correlated and therefore should not be included in the model.

Table 3-2. Summary of monitoring models for the Simulated Blend Down Facility. While each model was based on an empirical AAKR model, there were a total of 9 different parameters available to each model – 8 different sensors and the possibility of preprocessing the measured signal with data reconciliation. An “X” indicates that the sensor or technique was included in the AAKR model.

	Model 1	Model 2	Model 3	Model 4
HEUF ₆ weight sensor	X	X	X	X
LEUF ₆ weight sensor	X	X	X	X
PLEUF ₆ weight sensor	X	X	X	X
HEUF ₆ flow meter	X	X	X	X
LEUF ₆ flow meter	X	X	X	X
PLEUF ₆ flow meter	X	X	X	X
HEUF ₆ FMFM			X	X
PLEUF ₆ FMFM			X	X
Data Reconciliation		X		X




Figure 3-3. Block diagram of AAKR model with data reconciliation. The data selection, model building and predictions, and computing the residuals are the exact same as an AAKR model without data reconciliation. The purpose of including data reconciliation is to better constrain the measurement values using known physical laws about the system.

Once the relevant sensors were identified a normal dataset was generated. This dataset should be fault free and cover the expected operating conditions. The normal dataset serves as the historical operation data and should be large enough to cover the extremes of normal operation. In the case of the uranium blend-down facility, the normal set should include data representing the minimum and maximum material flow rate as well as everything in between. If the normal data does not contain all the expected operating conditions, then the model will perform poorly whenever the inquiry vector is not within the bounds of the normal data. For the simulated facility, the normal dataset was built from the data generated with no diversion.

The normal dataset was then denoised and standardized. Denoising, commonly referred to as smoothing, attempts to remove the uncorrelated and independent measurement noise inherent in all instrumentation as well as noise spikes resulting from outside processes (such as a vibration spike due to somebody hitting a pipe). In the simulated facility each measurement was contaminated with 1% gaussian noise. The data was denoised with a median filter with a window of 5. Standardizing the data is the process of mean-centering the data and scaling the data so that it has a unit-variance. To mean center the data, the mean value for each sensor in the normal dataset is simply subtracted. To give the data a unit variance, each value is divided by the variance of the normal dataset. The denoising and standardizing, or scaling, parameters used to mean center and unit variance the normal data set are applied to each subsequent dataset. This means that all subsequent datasets were: denoised with a median filter with a window of 5, had the mean of the normal dataset subtracted for each sensor, and had each sensor divided by the variance of the normal dataset.

Next a test and validation dataset were generated. These datasets were denoised and scaled in the same way the normal dataset is. The test dataset is used to optimize the AAKR model and the validation dataset is used to determine the model's performance. These datasets do not have to cover every operating condition as the normal dataset did, but should cover a sample of each. Each dataset should be independent, for example the testing and validation dataset should not be the exact

same data. These datasets were also generated from the simulated scenario with no diversion.

When an inquiry vector is processed, the vectors in the normal dataset are selected based on their similarity to the inquiry vector. The similarity is determined by calculating the Euclidean distance between the two vectors, and the weights are calculated using a Gaussian Kernel. Initially, the kernel bandwidth is set to 1, but was then optimized using the test dataset. The bandwidth was optimized so that the model uncertainty is minimized. Once the similar vectors have been selected, the prediction is calculated by taking a weighted sum of the similar vectors. The effects of bandwidth optimization are shown in Figure 2-3 and Figure 2-4.

Once the kernel has been optimized using the test dataset, the validation dataset was used to quantify the performance metrics of a model. The validation dataset was denoised and scaled the same way as the normal dataset. The bandwidth for the Gaussian kernel was previously optimized with the testing dataset. The Euclidean distances and observation weights are determined, and a prediction is made using the weighted sum of the similar vectors. Once a prediction is made, the residual was calculated by finding the difference between the prediction and actual measurement. Model performance characteristics such as accuracy, uncertainty, cross sensitivity, auto sensitivity, error-uncertainty limit monitoring (EULM) detectability, and SPRT detectability were calculated using this dataset. A complete description of these metrics can be found in [60-62].

Once the AAKR model was optimized and its performance metrics characterized, data from the other 5 scenarios (the diversions) were fed through the model and the residuals were analyzed using the SPRT. A fault was declared when detected using the SPRT alarm. The SPRT alarm requires four pieces of information: the mean of the expected residuals, the variance of the expected residuals, a false alarm probability, and a missed alarm probability. The expected mean and variance of the residuals was captured when quantifying the performance metrics. The false alarm and missed alarm probability were set to 1% and 10% respectively. To compensate for the high false

alarm and missed alarm probability, several consecutive residuals were required to be detected using the SPRT alarm before the fault was considered detected. So while the probability of a single residual being incorrectly detected using the SPRT alarm may be high, the probability of several consecutive residuals being incorrectly detected is much lower. Therefore, the specified false alarm and missed alarm probabilities are not truly representative of the actual false alarm and missed alarm probabilities. 5 consecutive residuals had to be detected before the fault was considered identified. In applications where there are not a large number of predictions, and therefore not a large number of residuals, the false alarm and missed alarm rates should be lowered.

3.2 Load Cell Monitoring at a Mock Feed and Withdrawal Facility

The Oak Ridge National Laboratory (ORNL) mock feed and withdrawal facility is designed to mimic large plant feed and withdrawal facilities. A schematic of the facility is shown in Figure 3-4.

In this facility, water is pumped from a feed tank into a surge tank, which then drains by gravity into a product tank and a tail tank. The only instrumentation is the weight sensors that monitor the feed, product, and tail tanks. Initially, the height of water in the surge tank was maintained via manually adjusting valves on the surge outlet. This control setup led to inconsistent operation depending on the attentiveness of the operator. A control valve and pressure sensor was later installed to replace operator (described in section 4.2.3). Throttling valves are used to fine tune the ratio of product to tail flow. This setup is analogous to a batch operation where tanks are placed on a scale and material is processed ,resulting in product and tail material. The actual facility (shown in Figure 3-5) has three feed, three product, and two tail tanks to allow for continuous operation.

For example, in a gaseous centrifuge enrichment plant, a tank of feed material would be weighed and fed though a cascade of centrifuges, resulting in product (enriched) material and tail (depleted) material. The weight of the tank would be monitored by a load cell station, which includes a weight scale. The surge tank serves




Figure 3-4. Schematic of the ORNL mock feed and withdrawal facility. Tanks of water are placed on the feed station and the water is pumped into a surge tank. Water is then partially routed to a tail and production station. In this setup, the pump, surge tank, and stream splitter serve as a mock cascade area.




Figure 3-5. The actual ORNL mock feed and withdrawal facility. Three feed stations are pumped into a Surge tank, which then drains into the product and tail stations. Before and after a tank is processed at each station, it is weighed on an accountability scale.

as a black-box between the feed, product, and tail material. The inventory difference is defined as the material unaccounted for between the beginning and end of the process. In a gaseous centrifuge plant, the surge tank would be analogous to material in the centrifuge cascade, including material that leaks or plates out on piping.

As tanks are brought into the facility, they are weighed on an accountability scale and are then stored. When a tank is ready to be processed it is placed on the appropriate station process scale, the material is processed, and it is removed. A tank is then re-weighed on the accountability scale. The total amount of material processed is calculated from weight recorded by the accountability scales, and it is recorded on a material declaration sheet. In an actual GCEP, the accountability scales require more precision and calibration than the process scales since they are used to declare the amount of material. During an inspection, the IAEA will check the calibration of the accountability scales and check the weight of a sample of cylinders. A schematic of a tank life cycle and the types of data generated during operation is shown in Figure 3-6.

In the ORNL facility, the accountability scale data is a step function that represents the weight of a tank placed on the accountability scale. The process scale data is a continuous function that shows the instantaneous weight of a cylinder as material is being processed. Process scale data is collected from the weight scales at each of the feed, tail, and product stations. The accountability scale data and process scale data at a large GCEP is directly analogous the accountability scale data and process scale data at the ORNL mock feed and withdrawal facility. The facility control data is the data used to control the flow of material inside the mock cascade area. In the ORNL mock feed and withdrawal facility, this data is limited to the signals from the feed pump current, the control valve and pressure transducer. In an actual GCEP the facility control data would encompass all of the controls in a cascade area including: control valves between centrifuges; motor signatures for each centrifuge; flow, pressure, and temperature measurements within each centrifuge; etc. A sample of the different data types is shown below in Figure 3-7. The accountability scale data is used to make mock operator declaration sheets. The current inspection process only uses data from the accountability scales. A sample declaration sheet is shown in Table 3-3.




Figure 3-6. The life cycle of a tank at the ORNL mock feed and withdrawal facility and the type of data generated during each stage of the tank life cycle. The black boxes and arrows represent how the physical material is moved through the facility. The blue boxes and dashed arrows represent the data collected from each physical location.




Figure 3-7. A sample of the different data types collected at the ORNL mock feed and withdrawal facility. The accountability scale data (a) is a step function of a tank's weight before and after processing. The process scale data (b) is continuous and shows the instantaneous weight of a tank on a station as material is being processed. The feed pump power (c) data is a type of facility control data and may either be a step function or continuous depending upon the instrumentation.

Table 3-3. A material declaration sheet from the ORNL mock feed and withdrawal facility. The sheet list the initial and final tank weight, as read by the accountability scales, before and after processing.

Tank ID	Initial Weight (kg)	Date	Time	Final Weight (kg)	Date	Time
25G04	85.09	14-Oct-10	3:30 PM	35.22	15-Oct-10	11:05 AM
25G06	141.44	15-Oct-10	10:15 AM	35.56	15-Oct-10	12:05 PM
25G05	35.32	15-Oct-10	10:15 AM	141.01	15-Oct-10	12:05 PM
10L02	1.148	15-Oct-10	10:15 AM	10.399	15-Oct-10	1:05 PM
10L03	1.152	15-Oct-10	10:15 AM	1.148	15-Oct-10	1:05 PM

To find undeclared activity, the inspector would first confirm the calibration of the accountancy scales and then reweigh a sample of tanks and compare them to the declaration sheets. In a large facility, the sheer amount of material processed limits the ability of an inspector to efficiently search for undeclared activity. However, an automated system that tracks the tank weights as seen by the process scale would provide a measure of cross-validation to the declaration sheets. Additionally, implementing process monitoring techniques that are sensitive to subtle process changes would increase an inspector's sensitivity to undeclared activity and possible special nuclear material diversions.

At the ORNL facility, sensors are sampled at 1 hertz and stored in a computer database. A MATLAB graphical user interface (GUI) was developed to read the database output and generate a report mimicking a declaration sheet by:

- 1 Identifying the state of each station (Empty, Static, Filling, or Draining) based on the profile of the process scale data. The process scale data shows the instantaneous tank weight as material is being processed, rather than just the initial and final weight, and has significant features that give insight into the state of each station. The station states were the foundation in building a load cell monitoring system.
- 2 Providing a declaration sheet based on the weights read by the process scales via a cylinder counting algorithm. The counting algorithm extracted how many tanks were placed and removed from each process station and the initial and final weight of each tank. The GUI declaration sheet is similar in form to the facility declaration sheet and provided an efficient way to ensure that the facility declared all processing activities.
- 3 Tracking the cumulative inventory difference of the facility identifying atypical process operations that could be indicative of material diversion or other abnormal activity.

The state specific models mentioned in item 3 would be based on a process monitoring scheme similar to the one described for the BDMS. Initially the model would only use data from the process scales. However, the effects of adding other facility data into the model will also be described. A MATLAB Simulink model of the facility was also developed for the ORNL facility to evaluate the effect of additional sensors not originally present in the actual facility. The relevant facility equations for the Simulink model are listed in Appendix B.

4 APPLICATIONS AND RESULTS

This section presents the results of the Uranium Blend-Down Facility simulation and the development of a load cell monitoring system for the ORNL mock feed and withdrawal facility. First, the effects of incorporating radiation detector signals and combining data reconciliation with the AAKR model are investigated with simulated data from the Uranium blend-down facility. Second, the development of a load cell monitoring system for the ORNL facility is described.

4.1 The Uranium Blend-Down Facility Simulation

A MATLAB Simulink model was developed to simulate data from the Blend-Down Monitoring System. The model was used to simulate both normal plant operations and anomalous operations resulting from a diversion from the HEUF₆ leg. The simulated data was used to develop an AAKR model to monitor the blend-down process. The AAKR model requires correlated sensors to make accurate predictions. However, the instantaneous tank weight is not well correlated. Therefore, the derivative of the weight sensors was approximated by the difference between two consecutive observations to determine the instantaneous rate of change in weight, also referred to as the mass flow rate. The equation relating mass flow rate and gas velocity is:

$$\dot{m} = \rho \cdot Pipe_{axs} \cdot v \quad (4.1)$$

where $\dot{m}$ is the mass flow rate in g/sec, ρ is the gas density in g/cm³, $Pipe_{axs}$ is the cross sectional area of the pipe in cm², and v is the gas velocity in cm/sec. The gas density is assumed to be constant during operation, so the mass flow rate and gas flow rate are directly proportional to each other. Additionally, the FMFM is modeled as linearly proportional to the mass flow. The linear approximation is valid for this system since the k-effective of the gas stream is low (<.1) and therefore the subcritical multiplication of the system is also low. Therefore, all measurements essentially give the same information. But, because the data comes from independent sensors they

provide a level of redundancy. Figure 4-1 shows typical data generated by the Simulink model. While only data for the LEU leg is shown, each leg is instrumented identically.

Even though the instantaneous weight sensors are not well suited for the AAKR method, they still provide useful information. A bleed-off or diversion scenario will eventually be detected by tracking the tank weight residual, i.e., the initial total weight of the HEUF₆ and LEUF₆ less the current sum of the three weight measurements. This residual will steadily increase as more HEUF₆ is diverted.

The AAKR model is able to predict the 1% process noise because it is common in each sensor. However the independent measurement noise cannot be predicted because it is not correlated within the sensors. The accuracy is less than 1% because the model tends to average the sensor noise through the inherent redundancy in the physical arrangement.

Four different models were considered and were described in section 3.1. All four models have similar performance on the unfaulted data. The unfaulted data performance describes the ability of the model to predict, or confirm, the measured sensor values. Table 4-1 summarizes the average accuracy expressed as percent error and average uncertainty, expressed as a percent of the mean sensor value, for each of the four AAKR models built from unfaulted data. For each model, a diversion was declared detected when the SPRT alarm indicated a change in the residuals.

Neither model 1 nor model 2 incorporated the FMFM's and therefore relied only on the derivative of the instantaneous tank weight and gas flow velocities to build a model and detect a fault. Any diversion from the HEUF₆ would be seen as a change in the residuals between the measured and predicted values in any of the PLEUF₆ sensors, shown in Figure 4-2. While neither model 1 nor model 2 was able to detect the 10% diversion instantaneously, the diversion could still have been detected by tracking all three tanks weights (the HEUF₆ tank, the LEUF₆ tank, and the PLEUF₆ tank) and ensuring that total outflow from the HEUF₆ and LEUF₆ tank equal the inflow to the




Figure 4-1. Simulated data from the Uranium blend-down facility Simulink model. The simulated data represents unfaulted operation. Initially, the LEU process scale is vacant, then a tank is added and its material processed, and finally the tank is removed. The LEU Flow and FMFM show a step change corresponding to when a tank begins processing. The LEU mass flow is not a direct output of the model but is approximated by taking the derivative of the LEU tank weight.

Table 4-1. Summary of AAKR model performance for unfaulted data from the blend-down monitoring system simulation. The accuracy is expressed in terms of percent error, so a lower value is indicative of less error. The uncertainty is expressed as a percent of the mean sensor value.

	Model 1	Model 2	Model 3	Model 4
Accuracy (%)	0.45	0.27	0.58	0.43
Uncertainty (%)	1.58	1.94	2.12	1.73

P-LEU Velocity Sensor Residuals for 10% Bleed-off Case (Model 2)




Figure 4-2. Residuals and SPRT Alarm for the P-LEU velocity sensor (model 2) 10% diversion scenario. Because there is no statistical change between the predicted and measured values, a 10% diversion could not be detected with this model. Note that in the top graph the residuals show no change and the SPRT does not alarm.

PLEUF₆ tank. The cumulative weight difference between the three tanks is shown in Figure 4-3. Under normal conditions, the residual should fluctuate around zero because the mass flow rate out of a tank is balanced with the mass flow rate into another tank. Under a diversion, there is a growing residual proportional to the cumulative amount of material diverted. This diversion assumes that the diverter does not introduce foreign material into the PLEUF₆ tank at the same rate that he diverts from the HEUF₆ leg. Doing so would ensure the mass flow rate into and out of the tanks remains equal, meaning that the residuals would show no indication that a diversion was taking place.

Models 3 and 4 incorporated the FMFMs and therefore had a sensor that was very sensitive to only the ²³⁵U flow (instead of the total Uranium flow). The total mass flow rate from the LEUF₆ leg is approximately 32x larger than the mass flow rate from the HEUF₆, so the PLEUF₆ leg mass flow rate is dominated by the LEUF₆ leg flow. Therefore, diversions from the HEUF₆ leg have a minimal impact on the PLEUF₆ tank weight and flow measurements. However, when examining the ²³⁵U flow instead of the total Uranium flow, the HEUF₆ is much more significant since the flow enrichment is 90%, while the LEUF₆ leg enrichment is less than 1%. Therefore, diversions from the HEUF₆ leg have a significant impact on the total ²³⁵U flow seen by the PLEUF₆ FMFM. Figure 4-4 shows the measured and predicted values for the PLEUF₆ FMFM in the 10% diversion scenario. The diversion can immediately be identified by the step change in the residuals.

Comparing the residuals, Figure 4-2 and Figure 4-4 shows that with the 6 sensor model (models 1 and 2), a 10% diversion could not be detected, but the 8 sensor model (model 3 and 4) could detect the diversion. The next step was to determine the minimal diversion that could be detected with the 8 sensor model. The Sequential Probability Ratio Test (SPRT) is a statistical technique developed to detect subtle changes in process parameters [60]. The SPRT essentially measures if a sequence of numbers comes from a specified distribution or not. If the model were perfect, then the residuals would only reflect the measurement noise of each sensor. Since the measurement noise of each sensor comes from a Gaussian distribution, the residuals should also be normally distributed. During normal operation, the distribution of the residuals will




Figure 4-3. Cumulative weight difference between the three tanks during a diversion. Under normal operation, the cumulative weight difference should be zero, since everything that leaves the HEUF₆ and LEUF₆ tanks enter the PLEUF₆ tank. During a 10% diversion, there is a growing residual as the cumulative amount of material diverted increases.




Figure 4-4. Predicted and measured PLEUF₆ flow rates for the 10% diversion scenario for the 8 sensor model (models 3 and 4). In the top plot the measured values are in blue and the prediction values are in red. There is a large statistical change between the predicted and measured values. A 10% diversion could be detected with this model. Note that in the top graph, the predicted values are indistinguishable from the measured values before the diversion takes place.

remain normal and the residuals will pass the SPRT test. The SPRT assumes the residuals come are gaussian white noise. In reality the residuals do not always meet those assumptions. When the assumptions are not fully met the SPRT will have a higher false alarm rate. However, by initially specifying a high false alarm rate and requiring several consecutive residuals to fail before declared a fault identified, the true false alarm rate is much lower. For instance in Figure 4-2 there are no alarmed states, even though with over 7000 data points and a specified false alarm rate of 1% one would expected about 70 false alarms. While requiring consecutive residuals to fail reduces to true false alarm rate, it tends to lengthen the time to detect a fault. But, this is a small effect in detecting a diversion, only requiring the alarm to occur a few time steps after the diversion starts. During a diversion the distribution of residuals will change and if the change is large enough then the residuals will fail the SPRT test. For the 8 sensor model, the SPRT detectability was found to be 1.10%. Therefore the model can detect a 2.0% diversion, but it cannot consistently detect the 1.0%, 0.5%, or 0.1% diversion. Figure 4-5 shows the measured and predicted values for the $PLEUF_6$ FMFM in the 2.0% diversion scenario. A missed alarm would be seen as an unalarmed state in Figure 4-5 after the 200th time step, which is when the diversion is started. With a specified missed alarm rate of 10% and about 1000 residuals in the faulted state, the expected number of missed alarms would have been around 100. However, requiring consecutive residuals to fail the also lowers the true missed alarm rate just as it did for the false alarm rage. Figure 4-5 shows now unalarmed states between 200 and 1200, so the true missed alarm rate is much lower than 10%.

The models that incorporated data reconciliation had similar accuracy and uncertainty measurements as well as fault detection characteristics. The improvements caused by data reconciliation are not visible by simply inspecting the residual plots. Table 4-2 summarizes some additional performance metrics. With each metric a lower value is desired. The accuracy is expressed in terms of percent error, so a lower value indicates less error. The analytical uncertainty is also expressed as a percent, with a lower uncertainty corresponding to a higher precision. Auto-sensitivity is a measure of robustness, with a lower value indicated a higher robustness. The auto-sensitivity




Figure 4-5. Predicted and measured PLEUF₆ flow rates for the 2% diversion scenario for the 8 sensor model (models 3 and 4). In the top plot, the measured values are in blue and the prediction values are in red. The SPRT identifies statistical change in the residuals, indicating that a 2% diversion could be detected with this model.

measures how much a faulted sensor input affects the prediction of the same sensor. The cross-sensitivity is a measures the spillover effect. The spillover effect is the effect a faulted sensor input has to the prediction of a different sensor. The Error Uncertainty Limit Monitoring (EULM) detectability indicates the smallest fault that is detector using Error Uncertainty Limit Monitoring, which monitors the uncertainty of the prediction errors relative to some specified tolerance. The EULM metric defines how small of a drift can be detected with 95% certainty. Finally, the SPRT Detectability measures the smallest process parameter change that can be detected using the Sequential Probability Ratio Test. A more complete explanation of these metrics can be found in [60-62]. While most of the metrics are very similar between model 3 and model 4, the EULM detectability and SPRT detectability were significantly improved by the application of data reconciliation to model 4. Each metric is about 30% smaller in model 4, meaning that the model's fault detection capability is about 30% better. It is important to note that these models were created using simulated data and in a highly idealized case. In a real world application, the signals may not be as clean and the correlations between sensors would not be as high. However, the basic relationships between the sensors would not change. The effect of real data would degrade each metric in Table 4-2 and decrease overall model performance. However, the effects would be seen in each of the 4 models. While the performance metrics listed in Table 4-2 are better than would be expected in an actual model, they are still useful for comparing how each model would perform in comparison to the other models.

4.2 The ORNL Mock Feed and Withdrawal Facility

The previous section showed an application of process monitoring for safeguards on a simulated facility with simulated data. The ORNL Mock Feed and Withdrawal facility provides a more realistic test bed for the development and application of a safeguards monitoring system. Whereas with the simulated uranium blend-down facility detecting material diversion was the sole safeguards option, with the ORNL facility the safeguards goal is to detect undeclared activity and to alert an inspector to atypical facility operation – not necessarily limited to material diversions.

Table 4-2. Model performance metrics for model 3 (No Data Reconciliation) and Model 4 (With Data Reconciliation). With each metric, a lower value is desired. The accuracy is expressed in terms of percent error, so a lower value is indicative of less error.

	Model 3 (No Data Reconciliation)	Model 4 (With Data Reconciliation)
Accuracy (percent)	0.58	0.43
Uncertainty (percent)	2.12	1.73
Auto-sensitivity	0.65	0.53
Cross-sensitivity	0.15	0.13
EULM Detectability	6.07	4.30
SPRT Detectability	1.58	1.09

The ORNL facility has limited instrumentation. Initially, the only data available came from the process load cell stations. Using load cell data, an automated method for reconciling process scale data and accountability scale data to ensure that no undeclared material processing has taken place was developed similar to the template matching procedure outlined in [49]. In a real GCEP, undeclared production may indicate that the facility is stockpiling special nuclear material for possible clandestine purposes. The developed algorithms determine if the amount of material processed according to the accountancy scale data matches the amount of material processed according to the process scale data (Figure 4-6). Additionally, the cumulative inventory difference (CID) was analyzed to determine if any material diversion took place during operation.

Initially, the stream splitter shown in Figure 3-4 consisted of three hand valves: a master on/off valve and two small throttling valves. Maintaining a steady surge tank level required an operator to constantly adjust the throttling valves. In practice, it is difficult to precisely control the manual valves to keep the surge tank constant while routing the proper amount of flow to the tail and product legs. This type of control scheme led to large inconsistencies in the instantaneous inventory difference and the material inflow and outflow of the surge tank. A picture of the original control scheme is provided in Figure 4-7. A typical feed, product, and tail process scale profile from this control scheme is shown in Figure 4-8. Figure 4-8 shows several problems arising from the manual control scheme originally employed. Most notably, material is drained into the tail tank for an additional 50 minutes after the feed tank has been emptied. The two blue circles in Figure 4-8 represent changes to the tail tank flow, but they do not correspond to changes to either the feed flow or product flow. The manual control scheme led to a highly variable system where the relationship between the inflow and outflow did not remain stable. At an actual GCEP, the material flow into and out of the cascade area are tightly controlled, and the amount of holdup material in the facility low. In the ORNL facility, the surge tank provides an unrealistically large hold-up volume when compared to a GCEP. However, the larger holdup volume is required because the pressure of the water column in the surge tank is the driving force of the material




Figure 4-6. An essential function of the MATLAB GUI is to determine if the amount of material processed according to the accountancy scale data (and listed in the declaration sheet) agrees with the amount of material processed according to the process scale data.




Figure 4-7. Original Facility control setup. The yellow-handled valve behind the throttling valves acted as a master on/off valve for the surge tank. During operation, the master valve was opened and the two throttling valves were used to route the proper amount of material to the product and tail legs. Proper operation required that the operator be aware of the surge tank level and the ratio of material routed to the product and tail legs. In practice, balancing material flow into and out of the surge tank while simultaneously routing the proper amount of material to each leg was difficult to achieve by simply adjusting the throttling valves. The control system was very sensitive to operator attentiveness and to changes in operator personnel.




Figure 4-8. Facility data from March 27, 2009 when the facility was still under manual control. The black line at 125 minutes represents when operation should have finished. There is no more feed material entering the mock cascade area, and the product tank is completely filled. However, because the operator was unable to balance the material flow within the surge tank, approximately 30 kg of water accumulated in the surge tank which then had to be drained into a tail tank. Additionally the blue circles highlight times when the tail tank flows were adjusted, but without any corresponding adjustments to the feed or tail flow. This control scheme led to high variability in the system, where inflow to the mock cascade area was no longer an accurate predictor of outflow of the mock cascade area.

outflow. In an actual GCEP, pumps move the material into and out of the cascade area.

4.2.1 Types of Operation at the ORNL Facility

The ORNL facility's purpose is to provide test data for on-line monitoring techniques. Three different types of runs were performed to provide case studies for the load cell monitoring system: normal operation, undeclared activity, and product leg diversions. During normal operation, tanks are weighed on an accountability scale, placed on the appropriate process station, and material is processed. When processing is complete, the tanks removed from the process scales are re-weighed on the accountability scales. The weights from the accountability scales are recorded on the material declaration sheet. The material declaration sheet is the official record of what was processed, and the accountability scales are considered authenticated scales that would be used during an inspection. Data from the process scales are stored in a central database. The process scales are not authenticated, and the data from the process scales are used by the facility for process monitoring. The process scales would not be used by an inspector because they are not authenticated.

Most of the runs performed were normal operation. The feed rates and stations used varied in each run to create a realistic operational history. The data from normal operation was used to develop and train the monitoring system. The undeclared activity runs were essentially the same as the normal operation runs, except the tanks were not weighed on the accountability scales and therefore do not appear on the material declaration sheet. This scenario was meant to represent a facility trying to hide its material throughput from an inspector, for the purpose of obtaining product material. While the ORNL facility only processes water, at a GCEP the product material would be HEUF_6 , possibly for the purpose of developing nuclear weapons. In practice, normal operation runs doubled as undeclared activity runs by simply ignoring the material declaration sheet. The life cycle of a tank during an undeclared activity run is shown in Figure 4-9. Comparing a material declaration sheet to a process scale declaration sheet can identify undeclared activity. For undeclared activity, it was usually assumed




Figure 4-9. The life cycle of a tank during undeclared activity. This scenario looks like normal operation except no authenticated accountability data is generated. This scenario represents a facility processing undeclared material in an attempt to make HEUF₆ for their own, possible clandestine, purpose.

that a feed, tail, and product tank was not declared. However, any combination or number of the tanks could have been undeclared.

The third type of run was a product leg diversion. In this scenario some material was diverted from the product leg. There were two different diversion scenarios. In the first scenario no make-up material was added to the tanks to mask the diversion. In the second scenario make-up material was introduced. The ORNL facility has a small diversion leg installed on the product leg, as shown in Figure 4-10. In a diversion scenario the material declaration sheet will match the process scale declaration sheet, but looking at the CID of the facility will show a non-stable value as more and more material is diverted.

The final scenario was a masked enrichment operation. In this scenario an undeclared feed, tail, and product tank were processed in parallel with a declared feed, tail, and product tank. However, the undeclared tanks were not placed on a process scale, so there was no process scale data either. This scenario was the limiting case study. The CID shows some hint of atypical facility operation, but was not able to definitively identify that undeclared material was processed off-scale. Figure 4-11 describes the masked enrichment scenario.

4.2.2 Development of a Load Cell Monitoring System

While the relationships between the flows were highly variable, the general profile from each station follows a predictable pattern. When a station is empty the station weight readout fluctuates around zero. Then a tank is placed on the station resulting in a step change followed by a brief period of high variability as the weight settles. As material is processed then weight of a tank either increases or decreased until the tank is full/empty, and then a tank is removed. There is another brief period of high variability as the tank. Figure 4-12 shows the raw profile for a feed, tail, and product station. The raw profile is noise free except for the spikes when a tank is first loaded or unloaded from a station. Part c of Figure 4-12 shows data from all the process scales during operation. The first requirement of the load cell monitoring




Figure 4-10. Diversion path at the ORNL facility.




Figure 4-11. Description of an enrichment masking scenario at the ORNL facility.




Figure 4-12. Raw tank station profiles. (a) shows the profile from a Feed Station, while (b) shows the profile from a tail and product station. Both (a) and (b) show the large spikes when a tank is added or removed (circled in red). In (b) the product station (blue) spikes are not as significant because the product tanks are much smaller. (c) Shows the profiles from all 8 stations (3 feed, 3 product, and 2 tail) at the ORNL facility.

system was to take the data from Part c of Figure 4-12 and automatically generate a summary sheet for comparison to the material declaration sheet.

To calculate the number of cylinders and the amount of material processed at each station, a method similar to the template matching procedure outlined in [49] was employed. Rather than trying to match each profile to a specific template, the state of each process scale was generally categorized as Empty, Static, Filling, or Draining. States were identified by looking at the magnitude of the weight, and the derivative of the weight reading, approximating the mass flow rate. Because identifying the state required approximating derivatives, the large spikes associated with the loading and unloading of tanks were first removed or the derivative at those times would not have been meaningful.

While the process scale data was sampled at 1Hz, it was downsampled to 0.2 Hz during analysis. Downsampling was done to limit the amount of memory required process a run. For instance, a month of process scale data contained over 25 million data points. Usually only a single day's worth of data was analyzed at a time because operating procedures at ORNL did not allow the facility to run over night. Downsampling the data reduced the amount of data exported from the facility while still preserving all the transients associated with operation of the facility.

A median filter was used to remove the noise spikes. A median filter uses a window of values about a point, finds the median value within the window, and assigns the value to the original point. A median filter is useful for removing spikes without changing the original trend or mean of the data. For the ORNL facility a window size of 11 consecutive points was used. Since the data was first downsampled to 0.2 Hz, a window size of 11 actually corresponds to a window of 55 seconds. Figure 4-13 shows the smoothing effect of a median filter on a feed tank profile.

Once the noise spikes have been removed, the state of an individual station is categorized in a two step process. The first step determines when the scale is empty by looking for a weight reading of below 0.5 kg. 0.5 kg is above the noise level of an empty




Figure 4-13. The effect of using a median filter on a feed scale profile. The raw data shown in blue is noise free except for spikes when a tank is unloaded and loaded – at approximately 600 and 650 minutes respectively. The dotted red line shows the smoothed data. It follows the same profile as the raw data, only without the noise spikes.

scale, but still well below the weight of a scale occupied with an empty tank. There is no mass flow rate threshold to identify an Empty state. The other three states, Static, Draining, and Filling, do have mass flow rate thresholds. Static is used to identify the times when a scale is occupied but material from that specific station is not being processed, and Draining and Filling are used to identify when material is actually being processed at the station. The mass flow rate limits for a Static state are set to ± 0.2 kg/sec for feed and tail tanks and ± 0.02 kg/sec for product tanks. Product tanks have a smaller window because they are about $1/10^{\text{th}}$ the size of a feed and tail tank and because the product station scales have less noise. The mass flow rate limits for the Draining and Filling state are -3 to -0.2 kg/sec and 0.2 to 3.0 kg/sec for feed and tail tanks and -0.3 to -0.02 kg/sec and 0.02 to 0.3 kg/sec for product tanks respectively. The derivative thresholds were chosen based on known properties of the system. For instance, the maximum feed rate is limited by the pumps at 2.0 kg/min, while normal operation is 1.5 kg/min. A summary of the state identification parameters is given in Table 4-3. Figure 4-14 graphically shows a raw tank profile and its approximated mass flow rate, and it shows the application of each threshold for a feed tank profile. Figure 4-15 shows the raw data from a feed station in the top plot and the results of smoothing and categorizing the station state in the bottom plot. In the bottom plot of Figure 4-15, the spikes associated with the loading and unloading of the tank have been smoothed, and the times when the station is Empty, Static, and Draining have been properly identified by looking at the derivative and absolute value of the raw data.

Once the states of an individual process station have been identified, the life cycle of a tank can be viewed with respect to possible station states. Figure 4-16 outlines the possible state transitions for the process stations. A station started in either the Empty state or the Static state, depending on whether or not a tank was already on the scale. If initially empty, once a tank was placed on the appropriate station, the state changed to the Static state. As material was processed, the station transitioned to either a Draining or Filling state depending, depending on whether the tank was receiving material or was having material removed. Once processing was complete, the station returned to the Static state. If the tank was removed, the station state returned

Table 4-3. Summary of thresholds for identifying station states.

State	Magnitude Threshold	Derivative Threshold (Lower)	Derivative Threshold (Upper)
Empty (Feed or Tail)	< 0.5 kg	N/A	N/A
Static (Feed or Tail)	> 0.5 kg	-0.2 kg/sec	0.2 kg/sec
Draining (Feed or Tail)	> 0.5 kg	-3.0 kg/sec	-0.2 kg/sec
Filling (Feed or Tail)	> 0.5 kg	0.2 kg/sec	3.0 kg/sec
Empty (Product)	< 0.5 kg	N/A	N/A
Static (Product)	> 0.5 kg	-0.02 kg/sec	0.02 kg/sec
Draining (Product)	> 0.5 kg	-.30 kg/sec	-0.02 kg/sec
Filling (Product)	> 0.5 kg	0.02 kg/sec	.30 kg/sec




Figure 4-14. A feed station profile and its approximated mass flow rates. The top plot shows a feed station profile with the 0.5 kg limit plotted as a dotted line. When the weight reading is below the dotted line, the station is said to be in the Empty state. The bottom plot shows the approximated mass flow rate for the same profile. The thresholds for the Static, Filling, and Draining state are shown as dotted lines. In addition to the mass flow rate falling within the specified region, the magnitude of the weight must be above 0.5.




Figure 4-15. The top plot shows the raw data from a feed station. The bottom plot shows the same station only the data has been filtered and the states identified. The profile portions are correctly identified as Empty, Static, or Draining based on the derivative and the absolute value of the data.




Figure 4-16. The life cycle of a process station with respect to 4 possible states.

to Empty. Alternatively, if the tank was not completely empty or full, more material was processed and the station state would have returned to either Draining or Filling. This cycle repeated itself as more tanks were added or removed. To count the number of cylinders processed, one looked for when a station first enters an Empty state and then when the station returned to an Empty state. If between those two states material was processed, then increment the number of cylinders processed. Special cases were needed if station did not start in the Empty state (which happened if a tank was already on the station at the beginning of the analysis) and if the station did not end in the Empty state (which happened if the tank was not removed by the end of the analysis).

To calculate the amount of material processed, the difference between the beginning and ending weight values of each Filling or Draining state was calculated and summed together for each tank. By only looking at the Filling or Draining states instead of the weight before and after the tank is placed on the station, the influence of unfiltered noise spikes was minimized. Since the data was time tagged, more information was available than from the declaration report. Figure 4-17 shows an accountancy scale profile overlapped with a process scale profile. The declaration report is generated from authenticated process scale data, while the process summary report comes from unauthenticated process scales. The information in the process summary report can then be used to validate the declaration report to provide an inspector with greater assurance that all the material processed by the facility is properly declared for in the declaration report. The process summary report contains more information than the declaration report, but information barriers could be used to limit the output or change the format to exactly mimic a declaration report.

Once a process summary report has been generated, the final step is to compare it to a material declaration sheet to search for undeclared activity. Undeclared activity means that a tank is not weighed on the accountancy scales and recorded in the declaration report, but is still processed at the facility. Therefore the amount of cylinders recorded by a material declaration sheet and indicated by the process summary report would differ. However, a direct comparison between the material declaration sheet and the process summary sheet is hindered by two facts: tanks ID numbers are recorded




Figure 4-17. A sample accountancy scale profile and process scale profile. The dark blue bars represent data from the accountancy scales. These values are step functions that represent the authenticated tank weights. The accountancy scale values are then used to generate the declaration report. The light blue line represents a station profile as recorded by the process scales. This information is used to generate a process summary report which can be compared directly to the declaration report, providing the inspector with more assurance that all the material processed by the facility was recorded in the declaration report.

on the material declaration sheet, but are not available with the process scale data, and the process scales may not be as calibrated or precise as the accountability scales. In an actual facility, the difference between an accountability scale reading and a process scale reading may be as much as 30 kg. However, this difference usually represents a near constant offset. Since matching based on tank ID's is not available (because the tank ID is not included in the process scale data) and matching based on the initial and/or final weights may not be practical (due to the difference in scale calibration and precision), the next best method was to try and match tanks based on the delta between the initial and final tank weights. The comparison algorithm calculates the delta for each tank declared on the material declaration sheet and for each tank listed in the process summary report, and matches them most similar deltas. A report is then generated listing each matched and unmatched tank. Unmatched tanks from the process summary report would be indicative of undeclared activity. A sample case is shown in section 4.2.4.

4.2.3 Addition of a PI Controller

Because of the variability introduced with the manual control scheme, a rule-based monitoring system was necessary. A rule-based system monitoring system is a system where one identifies a set of rules specific to a facility. These rules require a thorough understanding of the facility and its operation, therefore a system developed for one facility is generally not applicable to another. For the load cell monitoring system, a set of rules were developed to identify the station states and to count the number of cylinders processed. While a rule-based system was sufficient for cylinder counting and verifying the material declaration sheet, it provided a poor foundation for identifying atypical activity. The MATLAB Simulink model of the facility was used to scope what improvements could be made if a different control scheme was employed. Removing the human element from the control scheme would remove most of the variability in the system, which in turn would make facility operation more consistent. The most intuitive method was to replace the manual control valves with an automated PI controller. Figure 4-18 shows the simulated data when a PI controller is employed rather than manual control. Comparing Figure 4-18 to actual facility data generated with




Figure 4-18. Simulated data from incorporating a PI controller instead of a manual control scheme. In the simulated data, the PI controller keeps a constant level in the surge tank, so the mass flow into and out of the surge tank are kept equal. When the mass flow into and out of the surge tank is equal, the correlations between each individual flow rate increase, and more sensitive monitoring techniques can be employed to search for atypical facility operation.

the manual control scheme shown in Figure 4-8, Figure 4-18 does not show any abrupt changes in the tail flow and all three tanks are filled equally. Adding a PI controller removes most of the irreproducibility from each run and allows for more sophisticated monitoring techniques to be effectively implemented.

A PI controller and pressure transducer were installed at the surge tank outlet. The pressure transducer measured the height of water in the surge tank, and the PI controller would adjust a ball valve to maintain a set water level. In this setup, the total outflow from the surge tank is completely controlled by the PI controller. The ratio between the product and tail legs is still controlled by the throttling valves, but with the total outflow controlled by the PI controller the throttling valves do not have to be constantly adjusted during operation. As a result, the ratio between the product and tail legs remains consistent and the mass balance of the system is more stable. Figure 4-19 shows the PI controller and pressure transducer system installed at the ORNL Facility.

Tracking the CID of the facility is the primary way to search for atypical activity. The CID is calculated by extracting the times when the station is in the Draining or Filling state, taking the derivative of the instantaneous station weight to approximate the mass flow rate, and then summing the mass flow rate for each station. When the stations are in the Static or Empty state, the mass flow rate should be zero. But, because of sensor noise the mass flow rate may be non-zero. Including these times in the CID calculation increases the variation. During normal operation, the CID should remain stable. However, with the manual control scheme, the CID grows to almost 40 kg before returning to approximately 3 kg. Figure 4-20 shows the CID plot for the March 27, 2009 experimental run (the process station profiles are shown in Figure 4-8). The top plot in Figure 4-20 shows the CID for each station type. The inventory difference was calculated with respect to material entering and leaving the surge tank, so feed stations are positive as their material was entering the surge tank, and product and tail stations are negative because their material was leaving the surge tank. The CID for each station shows how much material was processed at each station. The bottom plot of Figure 4-20 is just the sum of the CID for each station. This plot shows how the




Figure 4-19. The updated control scheme at the ORNL Facility. The installation of a pressure transducer with a PI control valve removed a human element from control of the facility, removing the largest source of variability in facility operation. The pressure transducer gave a measure of the water level in the surge tank, and the PI controller adjusted the valve to maintain a set level.




Figure 4-20. Plots of the cumulative inventory difference from the March 27, 2009 run, when the facility was still under a manual control scheme.

amount of material in the surge tank is changes. Ideally this value should fluctuate around zero, but with a manual control scheme the operator was not able to effectively control the surge tank level. The same plots are shown in Figure 4-21, but in this run the PI controller had been installed in the facility. Importantly, the system CID peaks at 1.3 kg and reaches a stable value of 0.6, whereas in Figure 4-20 the system CID peaks at 40 kg and never has a stable value. The bottom plot of Figure 4-21 shows a few features related to the facility operation. The large initial spike comes from the material transport delay of the system. As water is pumped from the feed stations, it takes a finite amount of time to reach the surge tank. Even more significant is that the PI controller is not turned on immediately, allowing the surge tank to accumulate water. After the PI controller is turned on, it takes more time for the surge tank level to return to baseline. After this startup time (usually about 15 minutes), the total CID reaches its steady state value. The larger spikes during this stable period are correlated with the switching of feed tanks. The large negative spike in Figure 4-21 at ~18:15:00 is caused because the tank at Feed Station 3 is taken offline. The system quickly returns to its steady state value as Feed Station 1 is brought online. The final downward spike at the end of processing is again due to transport delay. As the Feed tanks are fully emptied and no new tanks are brought online, the material takes a few seconds to be processed and enter the product and tail tanks. At the end of the run, -0.3 kg is the total inventory difference. Generally one would expect these values to be positive due to material holdup in the system (i.e. more material enters the surge tank or mock cascade area than leaves it), but at the ORNL facility the large volume of water in the surge tank allows more material to be removed from the mock cascade area than enters. In a real facility, this would be analogous to the removal of residual holdup from previous operation.

The Sequential Probability Ratio Test was used to search for abnormal facility operation. This test determines if a sequence of numbers comes from a specified distribution or not. Since during normal operation the total CID should remain stable (except for at the beginning and end of operation, as explained previously), the SPRT was used to determine if the total CID changed in a statistically significant way during




Figure 4-21. Cumulative inventory difference from July 19, 2010, after installation of the PI controller. Plot (a) is the CID for each station, and plot (b) is the total CID calculated by summing the data in plot (a).

operation. The input to the SPRT is the mean and variance of the known distribution, and the probability of a missed alarm and false alarm. The probability of a missed alarm and false alarm was set to 10% and 1% respectively. However, as the uranium blend-down facility analysis, requiring several consecutive points to fail the SPRT significantly reduces the false alarm and missed alarm rates. 5 consecutive points were required to fail the SPRT before abnormal operation was identified. Requiring consecutive points to fail the SPRT slowed down the speed with which abnormal operation was identified. Since the data was downsampled to 0.2 Hz, abnormal facility operation could be detected in a minimum of 25 seconds. Considering that a typical run took over 3 hours, delaying detection by 25 seconds was insignificant. The mean and variance was directly calculated by at a window of 50 points in the total cumulative inventory plot. The window was chosen about 15 minutes into the run, so that it was not contaminated by the facility startup transients. Additionally, the CID was smoothed using a median filter to remove the spikes associated with bringing tanks online and offline. The raw and smoothed CID and the results of the SPRT are shown in Figure 4-22. In the bottom plot, the red circles represent times when the data failed the SPRT and blue represents normal operation. Even during normal operation, the data was expected to temporarily fail the SPRT during startup and shutdown transients. Times when the data failed the SPRT will indicate when the facility was not operating at under normal steady state conditions. In Figure 4-22, the stable CID lasted from 17:45 to about 19:00, which was a total of 900 points at a downsampled rate of 0.2 Hz. There are no false alarms during this period, so even though the specified false alarm rate was 1%, the true false alarm rate is much smaller.

The described process state identification, cylinder counting algorithms, and SPRT alarms were written in a simple, intuitive MATLAB GUI. The MATLAB GUI provides interactive plots of all the station scales, automatically generates a summary report, and provides the option of automatically comparing a summary report to a material declaration sheet. The GUI incorporates all of the previously described analysis and is designed to be a software tool for a facility inspector.




Figure 4-22. The cumulative inventory difference and SPRT results for the July 19, 2009 run. The cumulative inventory difference was smoothed with a median filter, and the mean and variance was taken from a window of data points right after the startup transient. The bottom plot shows the result of the SPRT. The red circles represent when the data failed the SPRT, and blue represents normal operation.

4.2.4 GUI Development

LoadCellEvents is a MATLAB based Graphical User Interface (GUI) designed as a tool for load cell monitoring at the ORNL Mock Feed and Withdrawal facility. The GUI provides an interactive platform for plotting load cell data, calculating the total number of cylinders and the amount of material processed, and examining the cumulative inventory difference. There are three inputs to the GUI: a text file from the process scales, a material declaration sheet, and a process scale declaration sheet. The outputs are: plots of the raw data, station state plots, cumulative ID plots, ID analysis plots, SPRT alarm plots, a process scale summary report, a process scale declaration sheet, and a comparison sheet used to reconcile a material declaration sheet and process declaration sheet. The program has two main parts: PlotEvents, which loads and plots the process scale data, and AnalyzeEvents which performs all of the analysis. A block diagram of the program is provided in Figure 4-23.

4.2.4.1 Using PlotEvents

PlotEvents is the program section that loads the process scale data and provides plots of each station. The input to the GUI is a simple text file containing 10 columns: 1) the data in yyyy-mm-dd format, 2) the time in HH:MM:SS format, 3)-10) the weights, in kg, for the 3 feed, 3 products, and 2 tail stations respectively. A sample input file is provided in the Appendix.

Double clicking the LoadCellEvents executable opens the PlotEvents window shown in Figure 4-24. Selecting the “Load Data” button prompts the user to select an input file. Once the file is selected, the raw data from the file is plotted. The user can select which stations to plot by clicking the appropriate tank station button. The time period of interest can be entered in the time box. The plot will not update until the “Replot” button is selected. The “New Window” button opens the current plot in a separate window. The “Analyze Events” button is explained in the next section.




Figure 4-23. Block Diagram of the MATLAB GUI. The inputs to the GUI are shown as blue lines and the outputs are shown as dotted lines. The solid black lines represent different sections of the GUI.




Figure 4-24. The opening PlotEvents window after data has been loaded. Clicking the Load Data button prompts the user to load a data file. Initially every station is plotted, but the user can select difference stations and then use the Replot button to update the plot.

4.2.4.2 Analyzing Events

Once the operator has inspected the raw data, clicking the “Analyze Events” button will open a new window, AnalyzeEvents shown in Figure 4-25. Four different types of analysis are available and can be selected using the four toggle buttons in the top right corner: System states, Cumulative ID (Inventory Difference), ID Analysis, Prognostics. In addition to the four types of analysis, there are three buttons available to the user: Replot, New Window, and Write Summary. The Replot and New Window button are used for updating plots and plotting data in a new window, respectively. The Write Summary button will be described after the types of analysis are discussed.

The first analysis, System States, provides plots of the station state identification algorithms. The allowed states are Empty, Static, Draining, or Filling. The Empty state is for when there is nothing placed on the process station. The Static state is when there is a tank on the station but no material is being removed or added to the tank. The Draining and Filling state is when material is actively being removed or placed into a tank. Different stations can be selected using the Tank Selection buttons provided next to the plot; however, only an individual station can be plotted at a time. These plots were included as a visual check to ensure the state categorization algorithms were performed correctly. If there are any discrepancies between the summary report and a material declaration sheet, these plots provide assurance that the discrepancies were due to incorrect station analysis.

The second type of analysis is the Cumulative ID. As with the System States analysis, the primary purpose is to provide a visual check. The Cumulative ID analyses the cumulative amount of material processed at each station: positive values for when material enters the mock cascade area (feed material) and negative values if material leaves the mock cascade area (product and tail material). Additionally, the total CID is plotted. Individual stations can be selected via the buttons next to the plot. Any combination of stations can be selected, but the figure will not update until the Replot button is pressed. Figure 4-26 shows the AnalyzeEvents window when the Cumulative ID analysis is plotted.




Figure 4-25. The window opened when the Analyze Events button was pressed. Four different types of analysis are available and can be selected using the four toggle buttons in the top right corner. The initial analysis is the System States.




Figure 4-26. The AnalyzeEvents window when the Cumulative ID analysis is selected. Different stations can be selected and plotted together. Additionally, the total cumulative inventory difference can be plotted.

The third type of analysis is called ID Analysis. This analysis looks only at the total cumulative ID analysis. Figure 6 shows the AnalyzeEvents window when ID analysis is selected. This window provides two plots. The top plot can be toggled between the cumulative ID and a Gaussian fit to the cumulative ID. The first option plots the raw and smoothed CID (Figure 4-27). The second option plots the makes a histogram of the CID values and fits them to a normal curve (Figure 4-28). The fourth analysis, Prognostics, is currently non-functional. The button is included in anticipation of future research, being performed by David Hooper, to predict the amount of material (if any) that will be unaccounted for if the current operation trend continues.

The “Write Summary Report” button examines each station to calculate the number of tanks placed on the scale, the number of tanks processed, the start and end time of a tank processed, the initial and final weight of each tank processed, and the total inventory difference ($\text{Feed} - \text{Tail} + \text{Product}$). A sample report is shown at the end. Additionally, a process scale declaration sheet is also created. The process scale declaration sheet is compared to a material declaration sheet to ensure that all the material processed was declared on the material declaration sheet. A sample process summary report is provided in the Appendix. A sample process scale declaration sheet is shown the undeclared activity case study.

Once the “Write Summary Report” is clicked, a new button becomes available: “Compare to Declaration”. This button prompts the user to select a material declaration sheet and a process declaration sheet. These two declaration sheets are then compared to each other, matching tanks based on their initial and final weights. Ideally, tanks would be matched based on their unique tank ID, but the tank ID is not recorded in the process scale data. Unmatched tanks are highlight as undeclared activity. Figure 4-29 shows the Compare to Declaration button available after the Write Summary Report button has been clicked.




Figure 4-27. The AnalyzeEvents window when ID analysis is selected. The default top plot is the raw and smoothed cumulative ID difference. The bottom plot is the SPRT alarm. Red indicates an atypical state; blue represents a normal state. Startup and shutdown transients are recognized as atypical operation. The bottom plot is the Sequential Probability Ratio Test (SPRT) alarm. When the facility is under normal steady state operation, the CID should be stable. The SPRT alarm is a statistical test to determine if the CID is remaining stable. The red portion indicates that the CID is not stable, while the blue portion means it is stable. The SPRT alarm provides indication of changes to facility operation and is meant as an inspector tool to recognize atypical facility operations.




Figure 4-28. The AnalyzeEvents window when ID analysis is selected and the Gaussian ID plot is selected. Instead of the CID vs. time, the top plot is a histogram of the CID fitted to a normal curve. The bottom plot remains unchanged.




Figure 4-29. The AnalyzeEvents window after the Write Summary Report button has been executed. A new button becomes available: Compare to Declaration, which allows the user to select a material declaration sheet and process declaration sheet

4.2.4.3 Finding Undeclared Activity

Undeclared activity is when an operator processes material without first weight the tanks on an accountability scale. Therefore, the tank will appear on the process declaration sheet but not the material declaration sheet. For example, consider the material declaration sheet from October 14 and 15, 2010 (Table 4-4). This sheet provides a tank Identification number, an initial and final weight, and the date and time the tank was weighed.

Using the PlotEvents software to analyze the process scale data from that time period and selecting the Write Summary Report in the AnalyzeEvents window will generate a process summary report. In the Process Summary Report from October 14 and 15, 2010, (Table 4-5) the start and stop times are not recorded, only the date. Most importantly, the tank IDs are not recorded because the tank IDs are not recorded when a tank is placed on a process scale.

Once the “Write Summary Report” button has been selected, the material declaration sheet and process declaration sheet can be automatically compared using the newly available “Compare to Declaration” button. This button will prompt the user to load a material declaration sheet and process summary sheet for comparison.

Because the tank IDs are not available in the process summary report, reconciling the material declaration sheet with the process declaration sheet was done by comparing each tanks initial and final weight (Table 4-6). Reconciled tanks are displayed first, showing the tank ID and the process station they were at. Tanks appearing on the material declaration sheet but not the process declaration sheet are shown next. These represent tanks onsite but that have not been processed. Finally, the undeclared tanks appearing in the process declaration sheet are shown. These are the tanks indicative of undeclared activity.

Table 4-4. Material Declaration Sheet for October 14 and 15, 2010

Tank ID	Initial Weight (kg)	Date	Time	Final Weight (kg)	Date	Time
25G04	85.09	14-Oct-10	3:30 PM	35.22	15-Oct-10	11:05 AM
25G06	141.44	15-Oct-10	10:15 AM	35.56	15-Oct-10	12:05 PM
25G05	35.32	15-Oct-10	10:15 AM	141.01	15-Oct-10	12:05 PM
10L02	1.148	15-Oct-10	10:15 AM	10.399	15-Oct-10	1:05 PM
10L03	1.152	15-Oct-10	10:15 AM	1.148	15-Oct-10	1:05 PM

Table 4-5 Process Declaration Sheet for October 14 and 15, 2010

Station	Initial Weight (kg)	Date	Final Weight (kg)	Date	Delta (kg)
Feed1	141.74	15-10-2010	35.91	15-10-2010	105.83
Feed2	140.76	15-10-2010	35.49	15-10-2010	105.27
Feed3	85.43	15-10-2010	36.08	15-10-2010	49.35
Product1	1.15	15-10-2010	3.24	15-10-2010	-2.09
Product2	1.16	15-10-2010	10.4	15-10-2010	-9.24
Tail1	35.75	15-10-2010	141.38	15-10-2010	-105.63
Tail2	35.48	15-10-2010	81.88	15-10-2010	-46.4
Tail2	36.2	15-10-2010	132.37	15-10-2010	-96.17

Table 4-6. Results of comparing the Material Declaration Sheet and Process Declaration Sheet. Reconciled tanks are shown first, then un-matched tanks listed on the material declaration sheet, and finally undeclared tanks.

Reconciled Tanks						
Tank ID	Station	Date	Initial Weight (kg)	Date	Final Weight (kg)	Delta (kg)
25G04	Acc't	14-Oct-10	85.09	15-Oct-10	35.22	49.87
	Feed3	15-Oct-10	85.43	15-Oct-10	36.08	49.35
25G06	Acc't	15-Oct-10	141.44	15-Oct-10	35.56	105.88
	Feed1	15-Oct-10	141.74	15-Oct-10	35.91	105.83
10L02	Acc't	15-Oct-10	1.148	15-Oct-10	10.399	-9.251
	Product2	15-Oct-10	1.16	15-Oct-10	10.4	-9.24
25G05	Acc't	15-Oct-10	35.32	15-Oct-10	141.01	-105.69
	Tail1	15-Oct-10	35.75	15-Oct-10	141.38	-105.63

Un-Matched Tanks From Material Declaration Sheet							
Tank ID	Initial Weight (kg)	Date	Time	Final Weight (kg)	Date	Time	Delta (kg)
10L03	1.152	15-Oct-10	10:15 AM	1.148	15-Oct-10	1:05 PM	0.004

Undeclared Tanks at Process Stations					
Station	Initial Weight (kg)	Date	Final Weight (kg)	Date	Delta (kg)
Feed2	140.76	15-10-2010	35.49	15-10-2010	105.27
Product1	1.15	15-10-2010	3.24	15-10-2010	-2.09
Tail2	35.48	15-10-2010	81.88	15-10-2010	-46.4
Tail2	36.2	15-10-2010	132.37	15-10-2010	-96.17

4.2.4.4 Finding A Diversion

In the case of material diversion, the process declaration sheet and material declaration sheet will be in agreement. For small diversions the material holdup ,may fall within acceptable values. Even for large diversions, makeup material could be placed in the appropriate tank to satisfy the mass balance. However, the total CID plot provided by the ID analysis option in the AnalyzeEvents window provides a way to search for diversions. A plot from normal operation is shown in Figure 4-30. The CID remains stable with the exception of startup and shutdown transients. The SPRT alarm signals on the transients, but does not alarm during steady state operation. Figure 4-31 shows the same analysis during a diversion. The diversion data is from June 7, 2010 operation. During this operation material was diverted from the product leg starting around 17:45, the diversion continued until approximately 19:12, but material processing continued until 20:38. Around 0.6 kg of material was diverted. However, once material was finished processing, material was reintroduced into the product tank so that the final cumulative inventory was within normal bounds. The SPRT alarm triggers within a few minutes and would alert an inspector that there was an atypical operation.

4.2.4.5 Detecting Masked Enrichment

In the case of masked enrichment, undeclared material was processed in parallel with declared material and the undeclared tanks were not placed on a station scale. Therefore process scale data was not available to identify the undeclared material. Additionally, the total inflow and outflow of the mock cascade area was equal. The operator was trying to simultaneously divert undeclared material away from the process scale while ensuring that the correct mass to the declared tanks. Therefore, the CID was not as stable as during normal operation. Figure 4-32 shows the CID during the masked enrichment run and the associated SPRT alarms. While the SPRT did alarm, it was not as definitive as in Figure 4-31. With a more sophisticated facility, this scenario could be hidden from an inspector. The CID did not give much insight in this scenario because most of the effected systems were inside the mock cascade area. If facility




Figure 4-30. The ID analysis of normal operation. The CID remains stable with the exception of startup and shutdown transients. The SPRT alarm triggers during the transients, but not steady state operation.




Figure 4-31. The ID analysis of a diversion operation. The cumulative inventory difference does not remain stable. The SPRT alarm identifies the deviation from steady state operation.




Figure 4-32. SPRT alarms during a masked enrichment run.

data were used in addition to process scale data, this scenario could be identified by looking for changes in the facility data. To this end, the MATLAB Simulink model was used to generate data for a masked enrichment scenario. The simulated data included both the process scale data and the facility control data. An AAKR model was trained using the mass flow rate (approximated from the derivative of the process scale data), and PI controller voltage.

For this simulation, the undeclared feed material was 10% of the declared feed flow. The expected feed flow was lower than the actual feed flow because actual feed flow included undeclared material processed in parallel. Because the PI controller was sensitive to the total feed flow, the additional undeclared feed flow causes the PI controller valve to open more than expected. Therefore, this scenario was identified by looking at the residuals between the measured PI voltage and the predicted PI voltage. Figure 4-33 shows feed flow rates, the PI voltage residuals, and the SPRT alarm corresponding to the masked enrichment simulation. This scenario represented a limiting case for load cell analysis. However, by incorporating facility control data and developing an AAKR model, the atypical facility operation would have been identified.




Figure 4-33. Simulated masked enrichment detection. Because the feed flow was higher than expected, the PI voltage was larger than predicted. Using an SPRT alarm on the residuals clearly identified that atypical operation had taken place.

5 CONCLUSIONS

Developing automated routines for safeguards monitoring is a burgeoning field in the safeguards community. While process monitoring has been applied across several industries, its application and development for safeguards is still in its infancy. Much of the research in safeguards monitoring has focused specifically on two types of facilities: nuclear fuel reprocessing facilities, and GCEPs. Traditional safeguards relied on inspection and authenticated measurements to perform material balances. However, integrating the wealth of unauthenticated data typically collected by facilities for process monitoring with traditional safeguards measures provides a cost-effective way to reduce the inspector burden while providing a more complete safeguard monitoring system.

The first step of this research examined the application process monitoring techniques applied to a simulated uranium blend down facility. The goals of this research were to identify what safeguards conclusions could be drawn from the application of an on-line monitoring system and to see if the monitoring system could be expanded to include signals from radiation detectors specific to nuclear processing facilities. To this end, a MATLAB Simulink model was used to generate normal facility data and diversions ranging from 0.1% to 10% of the HEU flow. The Simulink model included three flow sensors, three weight scales, and two fissile mass flow meters. The fissile mass flow meters incorporated radiation sensors to measure the mass flow of ^{235}U (a complete description of the fissile mass flow meters can be found in [56]). Four different models were created based on the Auto-Associative Kernel Regression. In each model, simulated data from normal operation was used to train the model. The model was then tested on the diversion data to determine its sensitivity to a diversion. In the first two models, signals from the fissile mass flow meters were not included as part of the monitoring system. The first and second model differed in that the second model's data was preprocessed with data reconciliation. Models three and four were similar to models one and two, except the fissile mass flow meter signals were included in the model. Comparing the models provided a quantitative measure of the improvements gained by including radiation sensors and preprocessing the data with data reconciliation. All models had similar performance on the unfaulted data, but

neither of the first two models was able to detect a 10% diversion of the HEU leg. However, models 3 and 4 were both able to detect a 1% diversion from the HEU leg, showing that incorporating simulated radiation signals is expected to greatly increase the sensitivity of the monitoring system to a diversion. While the application of data reconciliation in model 4 did not greatly increase the accuracy compared to model 3, it did significantly improve the ELUM and SPRT detectability. This research showed the development of an auto associative kernel regression model for safeguards monitoring and the benefits of preprocessing with data reconciliation.

The second phase focused on developing a load cell monitoring system for the ORNL mock feed and withdrawal facility. The ORNL mock feed and withdrawal facility is a continuous batch feed water processing facility. This facility provides realistic data to develop and test automated monitoring techniques. The load cell monitoring system employed an automatic cylinder counting algorithm to extract the number of tanks and the amount of material processed at each process scale. This data was then automatically compared to a material declaration sheet to search for undeclared activity. Traditionally, verifying the information on a material declaration sheet required an inspector to tediously reweigh every cylinder on an authenticated scale. However, due to the sheer number of cylinders at a facility, only a small random sample is reweighed during a monthly inspection. Detecting undeclared activity would require an inspector to find an undeclared cylinder while reweighing a small random sample. Applying load cell monitoring to a facility, even though the data comes from unauthenticated scales, provided an efficient way to search for undeclared activity. Additionally, atypical facility operation could be identified by tracking the CID. For instance, material diversion could be detected by tracking the CID. Currently, facilities are not required to report the CID. The masked enrichment scenario represented a limiting case. This scenario could not be easily identified only using load cell monitoring. However, simulated data showed that incorporating load cell data with the PI voltage in an AAKR model would have been able to identify the atypical operation. The load cell monitoring system cumulated in a MATLAB-based-GUI that incorporated process scale state identification, automatic cylinder counting algorithms, material declaration sheet verification (through process

scale cylinder counting), and incorporated a SPRT alarm system for when the total CID did not remain stable. While a SPRT alarm did not specifically signify clandestine intentions, it did highlight times when the facility was not operated as expected. An inspector could use this data as part of an information driven inspection, rather than just examining small random samples. The GUI was designed to be a software tool for inspectors, possibly providing a means to remotely authenticate facility declarations and reduce the frequency and manpower required for a facility inspection.

5.1 Summary of Contributions

This research explored the development of a safeguards monitoring system for two specific facilities: a simulated uranium blend down facility and the ORNL mock feed and withdrawal facility. For the simulated facility, a safeguards model was developed based on the Auto-associated Kernel Regression architecture which was augmented to include data-reconciliation and radiation sensors. Incorporating radiation sensors and data-reconciliation should allow diversions as small as 2% of the HEU leg to be detected.

For the ORNL facility, a load cell monitoring system was developed. This system included an automated routine to calculate the amount of material processed, count the number of cylinders, and search for atypical facility operations using the CID and the SPRT alarm. The load cell monitoring system cumulated in the development of a MATLAB-based GUI which could be used as a tool for inspectors to quickly and possibly remotely use facility data to draw safeguards conclusions.

5.2 Future Work

The area of safeguards monitoring holds many opportunities for continuing research beyond the scope of this dissertation. Several such areas have been mentioned throughout this report; a few specific ones are outlined below.

This research focused on developing a monitoring system used specifically to draw safeguards conclusions. The main monitoring method for the uranium blend down

facility was based on the AAKR, which was augmented with data reconciliation. However several different types of models have historically been used for process monitoring and optimization, and future work could examine the strengths and weaknesses of using different monitoring models. Also, looking at a number of different facilities would provide a better foundation for determine the optimum safeguards monitoring model. Additionally, fault detection was determined using the SPRT test. The specified false and missed alarm rates were 1% and 10%, respectively. But, by requiring 5 consecutive points to fail before a fault is declared, the true false and missed alarm rates are much lower. Quantifying the true false and missed alarm rates was beyond the scope of this research, but is an important step in quantifying a real-world monitoring system. One major limitation of this research was that the data was generated with a highly idealized MATLAB Simulink model. Real world data would likely have many other features, such as increased noise or a more complicated operation, which would degrade the performance of the model. A future area of research could examine developing these models for an actual facility, finding the limits of detection, and identifying the true missed alarm and false alarm rate.

The second phase of research examined developing a load cell monitoring. This load cell monitoring system used unauthenticated process scales to draw safeguards conclusions about a facility. A general cylinder counting algorithm provided an efficient way to validate facility declarations and to search for undeclared activity. With a modern control scheme, atypical facility operation could be identified by employing a SPRT alarm. Future work could examine incorporating other facility data and its effects on safeguards monitoring. For instance, the masked enrichment scenario represented the limiting case when only the process scales are included in the model. However, using simulated data, it was shown that including other facility data, namely the PI controller voltage, into an AAKR model would easily detect this scenario. While the SPRT was useful for finding atypical operation, more sophisticated residual analysis could provide inspector with a way to identify different scenarios (such as a diversion, sensor drift, or new equipment installation), rather than just identifying atypical operation.

One major hurdle to implementing online safeguards monitoring is what data are made available to an inspector and the possibility of transferring data offsite. The effect of incorporating other facility data and its benefit to safeguards monitoring should be examined. Using simulated data, it was shown that incorporating facility control data with load cell data into an AAKR model provided a way to identify atypical operation that would not have been possible using only load cell data. While safeguards monitoring is usually viewed from an inspector perspective and seen as a burden to a facility, further research could identify the potential benefits of a facility incorporating a safeguards monitoring system, specifically in reducing the frequency of inspections and the length of each inspection.

LIST OF REFERENCES

1. IAEA Department of Safeguards, "IAEA Safeguards: Staying Ahead of the Game," Vienna, Austria: The International Atomic Energy Agency, July 2007.
2. International Atomic Energy Agency, "IAEA Safeguards Glossary, 2001 Edition," Vienna, Austria: The International Atomic Energy Agency, 2001.
3. International Atomic Energy Agency, "Safeguards Techniques and Equipment, 2003 Edition," Vienna, Austria: The International Atomic Energy Agency, 2003.
4. International Atomic Energy Agency, "IAEA Safeguards Glossary, 2001 Edition," Vienna, Austria: The International Atomic Energy Agency, 2001.
5. INFCIRC/153 (Corrected), The Structure and Content of Agreements between the Agency and States Required in Connection with the Treaty on the Non-Proliferation of Nuclear Weapons, 1972.
6. J. E. Doyle, Editor. Nuclear Safeguards, Security, and Nonproliferation. Butterworth-Heinemann (Elsevier), Burlington, MA. Chapter 5, "International Safeguards Inspection: An Inside Look at the Process" by B. Boyer and M. Schanfein, p. 91, 2008.
7. U.S. Congress, Office of Technology Assessment, "Nuclear Safeguards and the International Atomic Energy Agency, Appendix A," OTA-ISS-615. Washington, D.C.: U.S. Government Printing Office, June 1995.
8. O. Yamamura, R. Yamamoto, S. Nomura and Y. Fujii, "Development of safeguards and maintenance technology in Tokai Reprocessing Plant," Progress in Nuclear Energy, Vol. 50, p. 666-673, 2008.
9. Heo, G. Y. 2008. Condition Monitoring Using Empirical Models: Technical Review and Prospects for Nuclear Applications, Nuclear Engineering and Technology, Vol. 40, No. 1.
10. Hines, J. W., and E. Davis. 2005. Lessons Learned From the U.S. Nuclear Power Plant Online Monitoring Programs, Progress in Nuclear Energy, Vol. 46, No. 3-4.

11. Kano, Manabu, Nagao, Koji, Hasebe, Shinji, Hashimoto, Iori, Hiromu, Ohno, Strauss, Ramon, and Bakshi, Bhavik. 2000. Comparison of statistical process monitoring methods: application to the Eastman challenge problem, *Computers and Chemical Engineering*, Vol 24, No. 175-181.
12. Burr, T., M. Ehinger, and J. Howell. 2008. A Review of Process Monitoring for Safeguards, 8th International Conference of Facility Operations-Safeguards Interface.
13. Li, F., B. R. Upadhyaya, and L. Coffey. 2009. Model-based Monitoring and Fault Diagnosis of Fossil Power Plant Process Units using Group Method of Data Handling, *ISA Transactions, Journal of the International Society of Automation*.
14. Hines, J. W., and R. Seibert. 2006. Technical Review of Online Monitoring Techniques for Performance Assessment, Volume 1: State-of-the-Art, NUREG/CR-6895, Vol. 1, U.S. Nuclear Regulatory Commission, Washington, D.C.
15. Hines, J. W., D. Garvey, R. Seibert, and A. Usynin. 2007. Technical Review of Online Monitoring Techniques for Performance Assessment, Volume 2: Theoretical Issues, NUREG/CR-6895, Vol. 2, U.S. Nuclear Regulatory Commission, Washington, D.C.
16. Hines, J. W., J. Garvey, D. Garvey, and R. Seibert. 2008. Technical Review of Online Monitoring Techniques for Performance Assessment, Volume 3: Limiting Case Studies, NUREG/CR-6895, Vol. 3, U.S. Nuclear Regulatory Commission, Washington, D.C.
17. Xu, Haiyong, Schoenberg, Frederick. Kernel regression of direction data with application to wind and wildfire data in Los Angeles County, California.
18. Meadows, Donella H., Dennis L. Meadows, Jørgen Randers, and William W. Behrens III. (1972) *The Limits to Growth*. New York: University Books. ISBN 0-87663-165-0

19. Bickel, Peter, J., Doksum, K. A. (2001) Mathematical Statistics: Basic and Selected Topics, Volume 1, Pearson Prentice-Hall, 2001
20. Davidson, A.C. (2003). Statistical Models, Cambridge University Press, 2003
21. Freedman, D.A. (2009). Statistical Models: Theory and Practice, Second Edition, Cambridge University Press, 2009
22. Seber, G. A. F., Lee, A. J. (2003). Linear Regression Analysis, 2nd Edition, Wiley-Interscience, 2003.
23. Atkeson, Christopher G., Andrew W. Moore, and Stefan Schaal (1997a). Locally Weighted Learning, Artificial Intelligence Review, Vol. 11, pp. 11-73: 1997.
24. Garvey, D., and J. W. Hines. 2006. Robust Distance Measures for Online Monitoring: Why Use Euclidean? 7th International Conference on Fuzzy Logic and Intelligent Technologies in Nuclear Science (FLINS), Genova, Italy.
25. Fan, J., Gijbels, I. (1996a), Local Polynomial Modeling and Its Applications, Chapman & Hall/CRC, New York, NY: 1996.
26. Hines, J.W. (2008b). Class notes for NE 671, Empirical Modeling
27. Kuehn, D.R., and H. Davidson (1961), "Computer Control II . Mathematics of Control", Chem. Eng. Progress, 57, 44-47.
28. Stanley, G.M. and R.S.H. Mah (1977), "Estimation of Flows and Temperatures in Process Networks", AIChE Journal, 23, pp. 642-650.
29. Knepper, J.C., and J.W. Gorman (1980), "Statistical Analysis of Constrained Data Sets", AIChE Journal, 26, 260-264.
30. Liebman, M.J., T.F. Edgar, and L.S. Lasdon (1992), Efficient Data Reconciliation and Estimation for Dynamic Processes Using Non-Linear Programming Techniques", Computers Chem. Engng., 16, 963-986.

31. Rollins, D.K., S. deVenathan, M.V.B. Bascunana (2002), "Measurement Bias Detection in Linear Dynamic Systems", *Computers in Chemical Engineering*, 26, 1201-1211.
32. Langenstein, M., and J. Janskey. 2003. Process Data Reconciliation in Nuclear Power Plants, 17th Structural Mechanics in Reactor Technology Conference (SMIRT 17), Paper #D435, Prague, Czech Republic.
33. Zhao, K. and B. R. Upadhyaya. 2004. Data Reconciliation and Gross Error Detection for Iris Helical Coil Steam Generators, *Transactions of the American Nuclear Society*.
34. Wald, Abraham (June, 1945). "Sequential Tests of Statistical Hypotheses". *Annals of Mathematical Statistics*
35. Humenik, K.E. and K.C. Grocc (1990), "Sequential Probability Ratio Test for Reactor Signal Validation and Sensor Surveillance Applications", *Nuclear Science and Engineering*, Vol. 105, pp. 383-390
36. Dwight D. Eisenhower. "Atoms for Peace," An address given to the 470th Plenary Meeting of the United Nations General Assembly, December 8th, 1953, accessed at http://world-nuclear-university.org/html/atoms_for_peace/ (January 2009).
37. International Atomic Energy Agency, "IAEA Safeguards: Stemming the Spread of Nuclear Weapons," Vienna, Austria: The International Atomic Energy Agency, accessed electronically at http://www.iaea.org/Publications/Factsheets/English/S1_Safeguards.pdf, March 2009.
38. U.S. Congress, Office of Technology Assessment, "Nuclear Safeguards and the International Atomic Energy Agency," OTA-ISS-615. Washington, D.C.: U.S. Government Printing Office, June 1995.

39. P.C. Durst et al., "Advanced Safeguards Approaches for New Reprocessing Facilities," ASA-100 report, Pacific Northwest National Laboratory, PNNL-16674, June 2007.
40. International Atomic Energy Agency, "Report of the LASCAR Forum: Large Scale Reprocessing Plant Safeguards," STI/PUB/922, IAEA, Vienna, Austria, 1992.
41. Advisory Group paper on authentication techniques for in-plant NDA equipment applied to IAEA safeguards, Vienna, 10-13 Nov, AG-336, 1981,.
42. Avenhaus, A., Canty, M., Compliance Quantified, Cambridge Press, Cambridge, 1996.
43. Delbeke, J. F. A., J. Howell, G. Eklund, G. Janssens- Maenhout, P. Peerani, and W. Janssens. 2208. The Real- Time Mass Evaluation System as a Tool for the Detection of Undeclared Cascade Operation by GCEPs, 8th International Conference of Facility Operations-Safeguards Interface.
44. Delbeke, Jochen, J. Howell, and W. Janssens. 2007. The Detection of Undeclared LEU Production at GCEP by Real Time Mass Balancing, Proceedings of the 48th Institute of Nuclear Materials Management Annual Meeting.
45. Cooley, Jill N. 2007. Model Safeguards Approach and Innovative Techniques Implemented by the IAEA at Gas Centrifuge Enrichment Plants, Proceedings of the 48th Institute of Nuclear Materials Management Annual Meeting.
46. Laughter, N. D., A. Krichinsky, J. B. Hines, D. N. Kovacic, and J. R. Younkin. 2008. Simulated Process Test Bed for Integrated Safeguards Operations Monitoring," 8th International Conference of Facility Operations- Safeguards Interface.
47. Lenarduzzi, R., J. Begovich, J. B. Hines, D. N. Kovacic, M. Whitaker, and J. R. Younkin. 2007. Technologies for Real Time Monitoring of Load Cells for

International Safeguards Applications, Proceedings of the 48th Institute of Nuclear Materials Management Annual Meeting.

48. Richir, P., P. Peerani, J. Delbeke, L. Dechamp, W. Janssens, and P. Meylemans, 2008. Current and Future Developments for Improved Safeguards in Sensitive Nuclear Facilities, Proceedings of the 2008 INMM/ESARDA Tokyo Workshop.
49. Howell, J. 2010. Algorithms To Count The Cylinder Throughput Of A GCEP Feed Station
50. Burr, T., Ehinger, M., Howell, J., Pomeroy, G. 2008. Reducing Inspector Uncertainty via Solution Monitoring, Proceedings of the 49th Institute of Nuclear Materials Management Annual Meeting
51. Burr, T., Howell, J., Longo, C. Suzuki, M. 2009. Change Detection in Solution Monitoring for Safeguards, Proceedings of the 50th Institute of Nuclear Materials Management Annual Meeting
52. Burr, T., Hamada, M. 2009. Measurement Error Modeling and Simulation for Solution Monitoring, Proceedings of the 50th Institute of Nuclear Materials Management Annual Meeting
53. Howell, J. 2009. Differentiating Between Normal and Abnormal Situations With Solution Monitoring Systems, Proceedings of the 50th Institute of Nuclear Materials Management Annual Meeting
54. Ehinger, M., Pomeroy, George, and Budlong-Sylvester, Kory. 2009. Process Monitoring for Nuclear Safeguards, Proceedings of the 50th Institute of Nuclear Materials Management Annual Meeting
55. Howell, J. 2009. International Collaborations In Process Monitoring, Proceedings of the 50th Institute of Nuclear Materials Management Annual Meeting
56. Uckan, T. et al., "Measurement Methodology of the Fissile Mass Flow Monitor for the HEU Transparency Implementation Instrumentation in Russia," Forty-Second

Annual Meeting of the Institute of Nuclear Materials Management, Indian Wells, California, July 20, 2001.

57. Uckan, Tanner, March-Leuba, Jose, Powell, Danny, and Wright, Michael, 2005. "Blend Down Monitoring System Fissile Mass Flow Monitor Implementation at the ElectroChemical Plant, Zelenogorsk, Russia. ORNL TM-2005/193
58. Uckan, Tanner, March-Leuba, Jose, Brittain, Ray Carlton, Powell, Danny, and Glaser, Joe, 2005. "Blend Down Monitoring System Fissile Mass Flow Monitor Implementation at the Siberian Chemical Enterprise, Seversk, Russia.
59. Uckan, T. et al., "The Blend Down Monitoring System Demonstration at the Padijcah Gaseous Diffusion Plant," Fortieth Annual Meeting of the Institute of Nuclear Materials Management, Phoenix, Arizona, July 25, 1999.
60. Wald, A. 1947. Sequential Analysis, John Wiley & Sons, New York, NY.
61. Garvey, D., and Hines, J. W. 2006. Development and Application of Fault Detectability Performance Metrics for Instrument Calibration Verification and Anomaly Detection, Journal of Pattern Recognition Research, Vol. 1, pp. 2-15.
62. Garvey, D., and Hines, J. W., 2006. Sensor Fault Detectability Measures for Autoassociative Empirical Models, 16th Annual Joint ISA POWID/EPRI Control and Instrumentation Conference, 49th Annual ISA POWID Symposium, San Jose, CA.
63. J. Howell, P. Friend, D. Jones, C. Taylor. Load-Cell-Based Mass Evaluation Systems Re-assessed on the Basis of URENCO (Capenhurst) Load Cell Data. ESARDA, Lithuania, 2009.

APPENDIX A: Relevant Equations to the Uranium Blend-Down Model

The simulated uranium blend down facility processes single tanks of UF₆ at each station leg. A tank is loaded on each station, and then HEUF₆ and LEUF₆ are blended together to produce PLEUF₆. The three types of sensors are: weight (kg), gas velocity (m/sec), and the Fissile Mass Flow Meter (counts/sec). This section will give a brief overview of the relevant physics. Default values were taken from [56 - 59] or approximated based on system geometry.

For the HEU and LEU legs, a mass flow rate (kg/sec) is first specified. The instantaneous tank weight at LEUF₆ and HEUF₆ legs is given by:

$$m_t = m_i - \int_0^t \dot{m} dt$$

Where m_t is the mass in the tank (kg) at time t , m_i is the initial mass of material in the tank, and $\dot{m}$ is the mass flow rate (kg/sec). This default values for the mass flow rate are $8.6 * 10^{-3}$ kg/sec for the LEUF₆ leg and $2.6 * 10^{-3}$ kg/sec for the HEUF₆ leg. For the PLEUF₆ the tank weight is given by:

$$m_t = m_i + \int_0^t (\dot{m}_{leuf6} + \dot{m}_{heuf6} - \dot{m}_{diversion}) dt$$

Where $\dot{m}_{leuf6}$ is the LEUF₆ mass flow rate, $\dot{m}_{heuf6}$ is the HEUF₆ mass flow rate, and $\dot{m}_{diversion}$ is the mass flow rate of the diversion. The mass flow rate of the diversion is usually specified as a percent of the HEUF₆ leg. For normal operation it is zero.

The gas flow velocities are given by:

$$v = \frac{\dot{m}}{\rho \times A}$$

Where $\dot{m}$ is the mass flow rate (kg/sec), v is the drain rate (m/sec), ρ is the density (kg/m³), and A is the pipe cross sectional area (m²). The pipe cross sectional area is $9.2 * 10^{-3}$ m² for the HEUF₆ leg and $3.77 * 10^{-2}$ for the LEUF₆ and PLEUF₆ legs. The density of the UF₆ is 0.57 kg/m³.

The PLEUF₆ enrichment is calculated in the blending tee and is governed by the following equation:

$$E_p = \frac{\dot{m}_l \times E_l + (\dot{m}_h - \dot{m}_d) \times E_h}{\dot{m}_l + \dot{m}_h - \dot{m}_d}$$

Where E is the percent enrichment for each leg. The default values are 90% and 1.5% for the HEUF₆ and LEUF₆ leg, respectively. $\dot{m}$ refers to the mass flow rate in kg/sec. The subscripts p , l , h , and d , refer to the PLEUF₆ leg, LEUF₆ leg, HEUF₆ leg, and the diversion.

The final piece of instrumentation is the Fissile Mass Flow Meter. The FMFM measures the ²³⁵U flow by counting the number of delayed gammas that appear after a fission event. The ²³⁵U flow is given by

$$\dot{m}_{U235} = \dot{m}_{total} \times E \times MF$$

where $\dot{m}$ refers to the mass flow rate (kg/sec), E is the material enrichment, and MF is the mass fraction of the uranium in UF₆ :

$$MF = \frac{(235 \times E) + (238 \times (1 - E))}{(235 \times E) + (238 \times (1 - E)) + 6 \times 19}$$

235 is the atomic mass of ²³⁵U, 238 is the atomic mass of ²³⁸U, 19 is the atomic mass of Fluorine, and 6 is the number of Fluorine atoms per Uranium atom in UF₆. For this simulation, the Fissile Mass Flow Meter response was assumed to be linearly proportional to the ²³⁵U flow. The equation governing the ²³⁵U mass flow to Fissile Mass Flow Meter counts is:

$$FMFM = \dot{m}_{total} \times S \times \frac{Solid\ Angle}{4\pi} \times P_{Fission} \times P_{Gamma} \times E_{detector}$$

where S is the neutron source strength in neutrons per sec, *Solid Angle* is the solid angle seen by the source neutron (dimensionless), $P_{fission}$ is the probability that a neutron that has entered the material flow interacts with ²³⁵U atom and causes a fission(units of fission*sec/neutron), P_{gamma} is the probability of a delayed gamma from a

fission (dimensionless), and $E_{detector}$ is the detector efficiency (dimensionless). The simulated values are $S = 1.04 * 10^8$ neutrons/sec, Solid Angle = 0.2, $P_{fission} = 6.75 * 10^{-5}$, $P_{gamma} = 0.1$, and $E_{detector} = 0.22$. Multiplying Solid Angle, $P_{fission}$, P_{gamma} , and $E_{detector}$ gives the probability of a source neutron causing a fission in a ^{235}U atom to be $2.97 * 10^{-7}$. With the simulated neutron source, there are 30 counts produced per a kg of ^{235}U irradiated.

The actual physics of the FMFM meter are much more complicated than described. The source strength was chosen to mimic a 3-4 micrograms of Cf-252 spontaneous fission source. The solid angle was chosen assuming that only $1/5^{\text{th}}$ of the source neutron have a chance to interact with the UF_6 stream. This assumes that the pipe encompasses approximately $1/5^{\text{th}}$ of the area seen by the neutron source. $P_{fission}$ was calculated by multiplying the U^{235} fission cross section by the number density of the UF_6 gas.

APPENDIX B: Relevant Equations to ORNL Facility Model

The simulated ORNL mock feed and withdrawal system was used to identify what the effects of changing the control scheme would be and to simulate data for development of an AAKR model. In this facility, a cylinder is loaded on each station, and feed water is pumped into a surge tank and then drains into a product and tails tank. The process scales are the only sensors. This section will give a brief overview of the relevant physics.

For a given mass flow rate, the mass of water inside a cylinder is:

$$m_t = m_i - \int_0^t \dot{m} dt$$

where m_t is the mass in the cylinder (kg) at time t , m_i is the initial mass of material in the cylinder, and $\dot{m}$ is the mass flow rate (kg/sec). The max feed mass flow rate is approximately 2 kg/min.

The feed water flows into the surge tank, and the volume of the tank changes according to:

$$\Delta H = \frac{\int_0^t (\dot{m}_{in} - \dot{m}_{out}) dt}{A_{surge} \times \rho}$$

where $\dot{m}_{in}$ and $\dot{m}_{out}$ are the mass flow rates into and out of the surge tank (kg/sec), ρ is the density of water (kg/m³), and A_{surge} is the effective Surge Tank cross sectional area (m²). The mass flow rate out is proportional to the height of water in the tank and the valve positioning:

$$\dot{m}_{out} = ValvePos \times Outlet\ Axs \times \rho \times \sqrt{2gH}$$

where $ValvePos$ is the amount the outlet valve is opened (ranges from 0 to 1), $Outlet\ Axs$ is the outlet pipe cross sectional area (m²), ρ is the density of water (kg/m³), g is the gravitational constant 9.8 m/sec², and H is the water height in the tank (m).

ValvePos is determined by the PI controller based on the difference between the setpoint height of water in the surge tank and the measured height of water in the surge tank. The PI constants are adjusted based on the desired responses of the system.

The error function is:

$$Error = H_{measured} - H_{setpoint}$$

when the measured height is larger than the setpoint, the error is larger and the valve opens more.

The surge tank water level is measured by a pressure transducer; the equation relating the pressure to the height of water in a tank is:

$$P = \rho \times g \times H$$

where P is the pressure in Pascal, ρ is the density of water (kg/m^3), g is the gravitational constant 9.8 m/sec^2 , and H is the height of water in the surge tank (meters).

The flow to the product and tail legs is controlled by throttling valves. These valves are opened to a set point and are then left alone during operation. Therefore, the flow to the tail or product leg is just a percent of the flow out of the surge tank. Typically, about 3% goes to the product leg and 97% goes to the tail leg.

The weight of a tail or product cylinder is given by:

$$m_t = m_i + \int_0^t \dot{m} dt$$

where m_t is the mass in the cylinder (kg) at time t , m_i is the initial mass of material in the cylinder, and $\dot{m}$ is the mass flow rate (kg/sec).

APPENDIX C: Descriptions of Algorithms used by PlotEvents

This section will describe some of the algorithms used by the PlotEvents software. A combination of inherent MATLAB functions and user-written functions are called at various times throughout the program. When feasible, they are presented in the order in which they are called.

1) When the “Load Data” button is clicked in the PlotEvents reader and a data file is selected, the pathname and file name are passed to the ReadORNL.m function. This is a user-written function that simply loads the data file into two matrices: **data** - 7 columns for the 7 process stations, with the rows corresponding to the weights at each time step, and **time** - one column corresponding to the time stamps.

2) The data is then downsampled using the MATLAB inherent downsample.m function. The data is downsampled from 1 hertz to 0.2 hertz (one sample every 5 seconds instead of every second). This is done primarily to reduce the amount of noise associated with loading and unloading tanks and to reduce the amount of storage space required.

3) The data is plotted and the xticklabel.m function is called. This is a user-defined function which correctly fixes the x-axis ticks on the plots. This function is called whenever the plot is updated or a when the data is plotted in a new window.

4) When the “AnalyzeEvents” Button is pressed the raw data is passed to the AnalyzeEvents Window. A new matrix, **cddata**, is defined being the same size as the **data** matrix, but all zeros. This matrix will hold the state number for each station at each instant in time. The raw data is smoothed by the SmoothORNL.m function. This is a user-defined function that selectively smoothes the data based on its state, defined in the **cddata** matrix. The smoothing function is a median filter (inherent to MATLAB) with differing time windows. The smooth data is recorded in the **sdata** matrix.

5) Initially, all tank states are defined as zero (since **cddata** is initially all zeros). The stations are then initially categorized by the ClassifyWeightORNL.m function. This

is a user-defined function that defines a state number based on the current station weight. The only state defined in the `ClassifyWeightORNL.m` function is the Empty state. It uses a hard threshold to define when a station is empty. If station weight is below 0.5 kg then the station is said to be in the Empty state. The state number or 1 corresponds to the Empty state.

5) Once the empty states have been defined and the **cdata** matrix has been updated accordingly, the `SmoothORNL.m` function is used to selectively resmooth the stations during the Empty state. This helps removes some spikes associated with the loading and unloading of a tank.

6) Once the data has been resmoothed and the **sdata** matrix updated, the stations are further categorized by the `ClassifyDiffWeightORNL.m` function. The first categorization performed by the `ClassifyWeightORNL.m` function essentially allows the state to be either Empty (state number = 1 in **cdata**) or undefined (state number = 0 in **cdata**). The `ClassifyDiffWeightORNL.m` is a user-defined function that determines the station state based on the instantaneous derivative of the weight. The weight derivative during the times when the system state is undefined is calculated and compared to defined thresholds. If the derivative is between ± 0.2 kg/sec (for feed and tail tanks) or between ± 0.02 kg/sec (for product tanks) the station state is defined as Static and the state number assigned is 2. If the derivative is between -3 and -0.2 kg/sec (for feed and tail tanks) or between -0.3 and -0.02 kg/sec (for product tanks) the station state is defined as Draining and the state number assigned is 3, If the derivative is between +0.2 and +3 kg/sec (for feed and tail tanks) or between +0.02 and +0.3 kg/sec (for product tanks) the station state is defined as Filling and the state number assigned is 4. Derivatives that do not fall within these thresholds are left undefined (state number = 0)

7) The next step is to remove the undefined (state number = 0) states in **cdata**. This is accomplished by the `CombineZerosORNL.m` function. This is a user-defined function that sets all undefined states to the start occurring before the undefined period. For example if the station is in the Empty state and then briefly falls into the undefined state before entering the Static state, the period of time when the state is undefined

would be assigned the Empty state. This generally occurs when tanks are loaded and unloaded as large spikes may cause large derivatives that fall outside the thresholds defined in the ClassifyDiffWeightORNL.m function.

8) At this point, all elements of **cdata** should be non-zero. The amount of time a station remains in each state is checked by the CheckLengthORNL.m function. This is a user defined function that checks how long a station remains in a specified state, compares it to a threshold, and then removes the state if its length is below the threshold. For example, if a station is in the Static state and then for a single time step enters the Draining state (most likely as the result of measurement noise), followed by immediately re-entering the Static state, then the Draining state is reassigned as static. The static and Empty states must last at least 5 consecutive time steps, while the Draining and Filling states must last at least 30 consecutive time steps.

9) A final check is performed by the CheckStaticORNL.m function. This is a user defined function that checks the Static state for derivatives that would correspond to the addition or removal of a tank. For instance, a -35.0 kg/sec derivative or higher is associated with removal of an empty feed tank. These times are assigned the Empty state.

10) Now the station states are completely defined in the **cdata** matrix. The TanksProcessedORNL.m function extracts the number of tanks placed on the stations, processed on the station, the start and end time of all processed tanks, and the start and end weights of all processed tanks. This is done by examining the **cdata** matrix and the state transitions. This information is printed to a summary report when the "Write Report" button is clicked in the AnalyzeEvents window

11) The **cdata**, **sdata**, and **data** matrices are then passed to the InventoryDifference Window. This happens automatically once the TanksProcessedORNL.m function has completed. The CumIDORNL.m function is used to calculate the CID for each station. The CID starts as zero until the station state becomes either Draining or Filling. Then the total amount of material processed each

time step is cumulatively added. During times when the station is Static or Empty, the differential inventory difference is assumed zero.

APPENDIX D: Sample Input File For PlotEvents

Date	Time	Feed1	Feed2	Feed3	Product1	Product2	Product3	Tail1	Tail2
2010-05-25	16:00:00	141.1100	142.4434	-0.0200	1.1530	-0.0010	7.2870	35.5600	70.1300
2010-05-25	16:00:01	141.1100	142.4434	-0.0100	1.1530	-0.0010	7.2870	35.5500	70.1300
2010-05-25	16:00:02	141.1100	142.4434	-0.0200	1.1530	-0.0010	7.2870	35.5500	70.1300
2010-05-25	16:00:03	141.1100	142.4434	-0.0100	1.1530	-0.0010	7.2870	35.5500	70.1300
2010-05-25	16:00:04	141.1100	142.4434	-0.0200	1.1530	-0.0010	7.2870	35.5600	70.1300
2010-05-25	16:00:05	141.1100	142.4434	-0.0100	1.1530	-0.0010	7.2870	35.5600	70.1300
2010-05-25	16:00:06	141.1100	142.4434	-0.0200	1.1530	-0.0010	7.2870	35.5600	70.1300
2010-05-25	16:00:07	141.1100	142.4434	-0.0100	1.1530	-0.0010	7.2870	35.5500	70.1300
2010-05-25	16:00:08	141.1100	142.4434	-0.0200	1.1530	-0.0010	7.2870	35.5500	70.1300
2010-05-25	16:00:09	141.1100	142.4434	-0.0100	1.1530	-0.0010	7.2870	35.5500	70.1300
2010-05-25	16:00:10	141.1100	142.4434	-0.0200	1.1530	-0.0010	7.2870	35.5600	70.1300

APPENDIX E: Process Summary Report

Data from 19-Jul-2010 to 19-Jul-2010

Tank information for Feed1

Number of tanks processed 1

Tank 01

Start Time: 19-07-2010
18:17:04

Start Weight: 144.22

Stop Weight: 84.61

Stop Time: 19-07-2010
19:08:44

Delta: 59.61

The Station is empty 0.00% of the time

The Station is static 65.52% of the time

The Station is being utilized 34.48% of the time

Tank information for Feed2

Number of tanks processed 1

Tank 01

Start Time: 19-07-2010
17:37:27

Start Weight: 140.86

Stop Weight: 35.49

Stop Time: 19-07-2010
19:15:18

Delta: 105.37

The Station is empty 6.42% of the time

The Station is static 28.33% of the time

The Station is being utilized 65.25% of the time

Tank information for Feed3

Number of tanks processed 1

Tank 01

Start Time: 19-07-2010
17:34:53

Start Weight: 80.04

Stop Weight: 36.04

Stop Time: 19-07-2010
18:17:04

Delta: 44.00
The Station is empty 9.41% of the time
The Station is static 62.42% of the time
The Station is being utilized 28.17% of the time

Tank information for Product1

Number of tanks processed 1

Tank 01
Start Time: 19-07-2010
17:35:52

Start Weight: 1.15

Stop Weight: 9.78

Stop Time: 19-07-2010
19:14:13

Delta: 8.62
The Station is empty 2.66% of the time
The Station is static 31.77% of the time
The Station is being utilized 65.58% of the time

Tank information for Product2

Number of tanks processed 0

The Station is empty 0.00% of the time
The Station is static 100.00% of the time
The Station is being utilized 0.00% of the time

Tank information for Product3

Number of tanks processed 0

The Station is empty 0.00% of the time
The Station is static 100.00% of the time
The Station is being utilized 0.00% of the time

Tank information for Tail1

Number of tanks processed 1

Tank 01
Start Time: 19-07-2010
18:26:57

Start Weight: 35.37

Stop Weight: 130.81

Stop Time: 19-07-2010
19:15:38
Delta: 95.44
The Station is empty 3.49% of the time
The Station is static 64.03% of the time
The Station is being utilized 32.48% of the time

Tank information for Tail2
Number of tanks processed 1
Tank 01
Start Time: 19-07-2010
17:35:27
Start Weight: 35.77
Stop Weight: 141.04
Stop Time: 19-07-2010
18:26:52
Delta: 105.27
The Station is empty 35.42% of the time
The Station is static 30.27% of the time
The Station is being utilized 34.31% of the time

Number of FEED tanks processed in this time: 3
Number of PRODUCT tanks processed in this time: 1
Number of TAILS tanks processed in this time: 2
The total inventory difference for these events is -0.35 kg.
Positive inventory difference means more material entered the mock cascade area than left it.

VITA

James Joseph Henkel was born on March 4, 1985 in Honolulu, Hawaii. In 1997 he and his family moved to Knoxville, TN. There, he graduated from Karns High School with honors in May 2003. He then attended the University of Tennessee, Knoxville, where he graduated with a Bachelor of Science in Nuclear Engineering in May 2007 and with his Masters of Science in Nuclear Engineering in May 2008.

As an undergraduate, James interned at the Y-12 National Security Complex investigating the feasibility of implementing a microwave casting furnace. Then in 2004 he began part time work with the Oak Ridge National Laboratory's Nuclear Materials Detection and Characterization group under Dr. John Mihalczko. That work continued throughout his graduate studies. Since that time, he has worked on several projects for verifying or detecting special nuclear material using active and passive interrogation techniques. His work focuses primarily on the computer modeling of measurements. Since coming to the lab, he has authored or co-authored over a dozen papers for journals and conference proceedings.

In May, 2007, he began research with Dr. J. Wesley Hines investigating the application of auto-associative model architectures for nuclear safeguards monitoring. This work led to a partnership with the Oak Ridge National Laboratory's International Safeguards Group to develop a load cell monitoring system for their mock feed and withdrawal facility. His current research is in the area of empirical modeling methods for nuclear safeguards monitoring. He completed his Ph.D. in Nuclear Engineering in July, 2011.

As an undergraduate, he was elected president of the University of Tennessee's American Nuclear Society student chapter. As a graduate student, he was a founding member of the University of Tennessee's Institute of Nuclear Materials Management student chapter. He is also a nation member of both the American Nuclear Society and the Institute of Nuclear Materials Management. In 2007, he passed the Fundamentals of Engineering exam to be recognized as an Engineering Intern.