

Development of a Prognostic Method for the Production of Undeclared Enriched Uranium

A Dissertation Presented for
The Doctor of Philosophy Degree
The University of Tennessee, Knoxville

David Alan Hooper

August 2011

Copyright © 2011 by David A. Hooper

All rights reserved.

Acknowledgments

First, thank you to my dissertation committee. Dr. Hines, thank you for taking me under your advisement through the last few years of my education. I appreciate that more than I ever expressed. To the other members of the committee (Dr. Bozdogan, Dr. Upadhyaya, and Dr. Hayward): thank you for being mentors, friends, and critics as I best needed. To the rest of the Nuclear Engineering department: thank you for the instruction and guidance. At risk of omitting a name, thank you to my teachers: Drs. Upadhyaya, Ruggles, Townsend, Pevey, Gehin, Miller, and Maldonado.

To the Statistics department: thank you for the opportunity to learn via the Intercollegiate Graduate Statistics Program. That is truly one of the biggest gems of the university, and I only regret that I couldn't talk more graduate students into going through it. Drs. Bozdogan, Leitnaker, Petrie, Mee, Schmidhammer, Younger, and Leon, thank you for your instruction.

To the UTNE support staff: Thank you. You bailed me out on many occasions and acted like it was no big deal. It was a big deal to me.

Special thanks must be given to Dr. Vince Jodoin of Oak Ridge National Laboratory: thank you for the terrific experience while working for you between my MS and PhD.

Thank you also to the people at ORNL, in Washington, and in the IAEA who made this project possible. I've had a terrific time, and perhaps the greatest learning experiences had nothing to do with technical ability. In particular, thank you, Michael Whitaker, for managing me through this project.

Most importantly, thank you to my wife Teresa (who understands better than anybody when I say tl;dr), my parents, and my sister Karen. It must have seemed like forever waiting for me to finish my academic career.

Abstract

As global demand for nuclear energy and threats to nuclear security increase, the need for verification of the peaceful application of nuclear materials and technology also rises. In accordance with the Nuclear Nonproliferation Treaty, the International Atomic Energy Agency is tasked with verification of the declared enrichment activities of member states. Due to the increased cost of inspection and verification of a globally growing nuclear energy industry, remote process monitoring has been proposed as part of a next-generation, information-driven safeguards program. To further enhance this safeguards approach, it is proposed that process monitoring data may be used to not only verify the past but to anticipate the future via prognostic analysis. While prognostic methods exist for health monitoring of physical processes, the literature is absent of methods to predict the outcome of decision-based events, such as the production of undeclared enriched uranium.

This dissertation introduces a method to predict the time at which a significant quantity of unaccounted material is expected to be diverted during an enrichment process. This method utilizes a particle filter to model the data and provide a Type III (degradation-based) prognostic estimate of time to diversion of a significant quantity. Measurement noise for the particle filter is estimated using historical data and may be updated with Bayesian estimates from the analyzed data. Dynamic noise estimates are updated based on observed changes in process data. The reliability of the prognostic model for a given range of data is validated via information complexity scores and goodness of fit statistics. The developed prognostic method is tested using data produced from the Oak Ridge Mock Feed and Withdrawal Facility, a 1:100 scale test platform for developing gas centrifuge remote monitoring techniques. Four case studies are considered: no diversion, slow diversion, fast diversion, and intermittent diversion. All intervals of diversion and non-diversion were correctly identified and significant quantity diversion time was accurately estimated. A diversion of 0.8 kg over 85 minutes was detected after 10 minutes and predicted to be 84 minutes and 10 seconds after 46 minutes and 40 seconds with an uncertainty of 2 minutes and 52 seconds.

Executive Summary

As with many technologies, the peaceful benefits of nuclear energy such as clean power and production of medical isotopes must be balanced against potential harms. The appeal of the low lifetime costs, high energy output, and clean emissions has kept nuclear power as an attractive energy source for much of the world, even in the face of the potential risks for nuclear materials to be used in weapons such as traditional nuclear bombs or "dirty" bombs. As of this writing, well over 400 commercial nuclear power plants operate across the world in over 30 countries, accounting for almost 400 GWe of electrical production and roughly 14% of the world's total electrical energy supply. Additionally, nearly 250 research reactors and a little less than 200 naval vessels (e.g. aircraft carriers and submarines) rely on nuclear power in a total of over 50 nations. Global nuclear capacity is continually increasing, and the demand for nuclear resources continues to grow despite the negative publicity of incidents such as the event at Fukushima-Daiichi.

Unfortunately, the process of developing raw materials into fissionable sources of energy productions is nearly identical to the development of weapons-grade fission sources. For example, a typical commercial nuclear reactor might consume fuel with the Uranium 235 content enriched to about 5%, while a nuclear bomb using U-235 as the fissionable fuel would require far higher enrichment levels. (Precise estimates of uranium enrichment required for weapons-grade material are highly sensitive and typically classified; for sake of discussion, let us suppose that Uranium 235 enrichment greater than 50% is necessary.) While the difference in enrichment levels may seem significant, the equipment that can produce low levels of enrichment for power generation may be used for producing far higher levels of enrichment. The only difference is how long the enrichment process is allowed to run before extraction of the enriched material is finalized. Additionally, the process of enriching uranium 235 from its natural level of about 0.7% to about 5% requires as much effort as enrichment from 5% to levels closer to 50%.

To balance the peaceful benefits of nuclear energy against the potential dangers, the members of the United Nations have agreed to task the International Atomic Energy Agency (IAEA) with the verification of declared production and processing of nuclear materials, as outlined by the Nuclear Nonproliferation Treaty. In the context of uranium enrichment, the IAEA monitors the amount of material processed by enrichment facilities and reports whether their observations support the declarations that the facility makes of its production. While the IAEA does not have enforcement capacity, their verifications play a major role in international relationships: countries that demonstrate effective safeguarding and

application of nuclear materials benefit from the use of nuclear energy without the stigma of being perceived as a significant nuclear risk while noncompliant nations are at risk of economic sanctions or other negative effects of impaired international relationships.

Historically and presently, the IAEA monitors uranium enrichment by periodically sending inspectors to declared enrichment facilities and comparing the operator's declarations to their own accountancy of the material processing (i.e. the amount of feed, product, and tails material moving in and out of the facility). This process relies entirely on visual inspection and measurement of the facility's activities and typically involves about a dozen trips by inspectors to enrichment facilities each year. As global nuclear power demand increases, more enrichment facilities are built and the need for more inspectors continues to rise. In response, the IAEA is seeking to reduce their operating cost by incorporating remote monitoring practices into their inspection regime. Conceptually, the IAEA would collect process data in their own on-site database and analyze the data to generate conclusions of facility compliance with the NPT. These conclusions would not replace the accountancy system in place; rather, the hope is that remote monitoring may provide greater assurance of compliance between inspections and therefore allow the IAEA to conduct fewer live inspections on an "information-driven" schedule rather than the current periodic schedule that has been employed for several decades. In return, enrichment facilities stand to benefit from the reduced live inspection periodicity; with fewer interruptions in the enrichment process, more material may be processed in a facility and the holding time of material cylinders may be reduced, resulting in significant economic savings.

Should remote monitoring be incorporated into the IAEA's safeguards program for uranium enrichment, the information collected from the monitoring system may be useful for more than simply explaining the past (i.e. what has already occurred at a facility). While traditional inspection practices are backward looking in nature and verify the activities that have already occurred, remote monitoring data may be coupled with prognostic techniques developed for health monitoring applications to anticipate how current trends at a facility may play out into the future. For example, if the process data suggests that a protracted diversion event is occurring (i.e. that material is slowly being removed from the process stream and accumulated without being declared), the traditional approach to the data would only allow the IAEA to estimate the amount of material that has already been lost. With prognostics, the IAEA may be able to forecast the time until a specific quantity of material is unaccounted for in the process; such information may be used to decide how urgently the Agency needs to consider a

response. If the apparent loss of material is extremely slow, perhaps the next scheduled inspection is sufficient to identify and resolve the issue. If the material loss is quick, the Agency may need to consider an unannounced inspection prior to the next scheduled inspection. Such knowledge would be well-suited to the concept of information-driven safeguards and inspections, and would help maximize inspection efficiency.

This research focuses on the development of a prognostic method to forecast the time until a predetermined quantity of material has been diverted from an enrichment facility. Unlike traditional health monitoring applications such as crack propagation in metal components or degradation of valves, inspection monitoring is not focused on modeling mechanical processes but identifying human decisions and changes in operating conditions as a result of these decisions. This distinction introduces several problems to the process monitoring approach. First, diversion may not actually occur and the material balance may never truly "degrade" due to the loss of material. Second, there is no guarantee that a diversion, if present, will follow a predetermined pattern. Diversion may be performed at a constant rate, intermittently, at fast or slow rates, or a myriad of other patterns. Third, diversion is not the only means by which the material balance may be affected; measurement biases in weight sensors, holdup in the enrichment cascade, and normal processes such as degassing of feed cylinders may all affect the material balance.

To accommodate the variety of possible diversion and non-diversion scenarios that may be encountered, this research employs a prognostic method to predict the time at which a predefined "significant quantity" of material is lost from the material balance of the enrichment process. The particle filter uses a Monte Carlo procedure to estimate trends in the materials balance and generate predictions and uncertainty estimates of the time to significant quantity diversion. The particles are constrained to linear predictions, so that any estimated time to diversion of a significant quantity assumes that the diversion rate remains constant. (This is in alignment with typical enrichment processes, where maximal efficiency and minimal equipment wear are achieved by maintaining a constant rate of processing.) The particle filter process, however, affords nonlinear flexibility by updating the particle weights and trajectories as more data are received.

The prognostic method is applied to data collected from the Oak Ridge Mock Feed and Withdrawal Facility, an analog of the gas centrifuge enrichment process that uses water as the simulated enrichment medium and is scaled roughly 1:100 in both time and mass to a gas centrifuge enrichment plant. The

particle filter is applied to four data sets: a non-diversion process, a slow diversion, a fast diversion, and an intermittent diversion. The significant quantities of concern for the non-diversion and slow diversions are 0.8 kg of water, or analogously 80 kg of low-enriched uranium; for the fast and intermittent diversions, the significant quantity limits are 2 kg, which scales to approximately 200 kg. (The varying thresholds for diversion were designed to accommodate laboratory restrictions for operating times. However, the IAEA definition of a significant quantity for low enriched uranium is the amount that contains 75 kg of Uranium 235, or roughly 2000 kg of 5% enriched uranium hexafluoride, which scales to roughly 20 kg of diverted water at the mock facility. The lower limits thus represent a conservative threshold for diversion rather than an absolute time to undeclared significant quantity production.)

For the non-diversion case, the particle filter correctly predicts that diversion does not produce 0.8 kg of material within the prognostic window of roughly 14 hours. For slow diversion, the particle filter first detects the diversion about 10 minutes after initiation; by 46:40 minutes, the actual time to diversion of 85 minutes lies within the 95% confidence interval of the particles, and the final prediction at 6 minutes and 40 seconds accurately predicts 0.8 kg of diversion between 81:58 minutes and 92:46 minutes with 95% confidence. For fast diversion, the diversion is again detected within 10 minutes (the earliest allowable detection time due to particle learning time) with the 120 minutes to 2 kg of diversion predicted with 95% confidence immediately. With intermittent diversion, the diversion was initiated after approximately 60 minutes and paused from 96:40 minutes to 103:20 minutes. The particle filter accurately detected the diversion immediately and recognized the pause in diversion, giving a temporary prediction of no significant quantity production during this time. Once diversion resumed, the particle filter predicted the correct time to significant quantity production of 135 minutes with 95% confidence, with a 95% confidence interval prediction at 103:20 minutes of approximately [283:20 66:40] minutes and a 95% confidence interval at 126:40 minutes of [167:40 135:50] minutes. The very large confidence intervals at 103:20 minutes were reflective of the very high uncertainty associated with a prediction immediately after a change in regime.

The precision of predictions of time to significant quantity diversion were heavily dependent on the constancy of the diversion rate itself. Not knowing the intent of the operator, the particle filter automatically adapted its confidence intervals based on the dynamics of the data, increasing the variance estimates with dynamic diversion rates and converging on the time to significant quantity

diversion when the diversion rate was stable. In all cases, the presence of diversion was correctly identified and the estimates converged toward the correct time to diversion as more data were collected.

Table of Contents

1	INTRODUCTION	1
1.1	Problem Statement	3
1.2	Original Contributions	4
1.3	Organization of the Document	5
2	LITERATURE SURVEY	6
2.1	IAEA Responsibility for Enrichment Verification	6
2.2	Current Process for GCEP Inspection	7
2.3	Next-Generation Techniques for GCEP Inspection and Monitoring	8
2.3.1	Load Cell Monitoring and Automated Analysis	9
2.3.2	Non-Destructive Assay	15
2.3.3	Facility Surveillance	16
2.3.4	Cylinder Identification	16
2.3.5	Inspection Periodicity	16
2.3.6	Mailbox Declarations	17
2.3.7	Automated Analysis at the Mock Feed and Withdrawal Facility	17
2.4	Process Monitoring and Prognostics	23
2.4.1	Type I, II, and III Prognostics	24
2.4.2	Methods for Type III Prognostics	25
2.4.3	Prognostics within Verification and Safeguards	39
2.5	Model Selection Criteria	39
3	METHODOLOGY	43
3.1	Mock Feed and Withdrawal Facility	43
3.2	Prognostics and Uncertainty Estimates	44
3.2.1	Identification of Facility Operation Features	45
3.2.2	Estimation of Diversion (TMUF)	45
3.2.3	Automated Identification of Undeclared Activity	46
3.2.4	General Path Method	48
3.2.5	Particle Filter Method	53
3.3	Software Development	66
4	APPLICATION AND RESULTS	67

4.1	Mock Feed and Withdrawal Facility Improvements	67
4.1.1	Automated Control of the Surge Tank Control Valve	67
4.1.2	Reduction of Pressure Transducer Noise	70
4.2	Mock Feed and Withdrawal Facility Operation	77
4.3	Mock Feed and Withdrawal Observability.....	79
4.3.1	Feed and Withdrawal Station State Variables	79
4.3.2	PI Control State Variable	81
4.3.3	PI Controller State-Space Representation	82
4.3.4	State-Space Model of the Mock Facility	83
4.3.5	Observability of the Mock Facility.....	86
4.3.6	Observability of a Cascade System	87
4.3.7	Observability during Diversion.....	89
4.4	Predicting Time to Diversion of a Threshold Quantity.....	92
4.4.1	Definition of the Prognostic Parameter	93
4.4.2	Markov Chain Model.....	95
4.4.3	General Path Model	100
4.4.4	Particle Filter Model.....	105
4.4.5	Prognostic Model Selection	112
4.5	Case Studies	113
4.5.1	Legitimate Operation (No Diversion).....	114
4.5.2	Slow Diversion.....	120
4.5.3	Fast Diversion.....	125
4.5.4	Intermittent Diversion	130
4.5.5	Case Studies: Concluding Remarks	139
4.5.6	Limitations on Particle Filter Reliability	140
5	Conclusions	142
5.1	Summary of Contributions.....	145
5.2	Future Work	146
	REFERENCES.....	148
	APPENDICES	159
	Appendix A: Description of the Mock Feed and Withdrawal Facility	160

F&W Facility Components.....	160
Basic Facility Operation.....	172
Appendix B: Measurement and Dynamic Noise Estimates.....	174
Estimation of the Measurement Noise	174
Influence of Measurement Noise Illustrated	175
Estimation of the Dynamic Noise.....	179
Updating the Dynamic Noise Estimate Based on Data Behavior.....	183
Appendix C: Case Studies	188
Legitimate Operation	188
Slow Diversion.....	188
Fast Diversion.....	188
Intermittent Diversion	195
VITA.....	198

Table of Figures

Figure 1 - MUF Calculation from Capenhurst Load Cell Evaluation [Howell, et al., 2009]	10
Figure 2 - MUF Calculation from the Mock F&W Facility.....	11
Figure 3 - <i>PlotEvents</i> Main Screen: Process Station Weight History over Time	19
Figure 4 - <i>PlotEvents</i> Analysis of Process Station Usage over Time.....	20
Figure 5 - <i>PlotEvents</i> Depiction of Process Scale Inventory Differences	21
Figure 6 - <i>PlotEvents</i> Net Inventory Difference for a Run of the Mock F&W Facility	22
Figure 7 - Chinnam's Historical Drill Bit Failure Data [Chinnam, 1999]	28
Figure 8 - Fatigue Crack Propagation in a Metal Plate [Lu, et al., 1993].....	31
Figure 9 - GPM Failure Prediction of an Elevator Door [Yan, et al., 2004]	32
Figure 10 - Sample Estimates of the System State Prior to Weighting.....	55
Figure 11 - Sample Estimates of the System State after Weighting	56
Figure 12 - PI Control Stabilization and MUF Calculation	61
Figure 13 - Fitting Normal Curves to Normal and Beta Distributions	65
Figure 14 - Original Surge Tank Control Valve Setup	68
Figure 15 - Pressure Transducer Voltage with no Surge Tank Flow.....	71
Figure 16 - Unfiltered Pressure Transducer Signal - Fast Fourier Transform.....	72
Figure 17 - Filtered and Unfiltered Pressure Transducer Signals.....	74
Figure 18 - Filtered Pressure Transducer Signal Frequency Spectrum	75
Figure 19 - Voltage Standard Deviation as a Function of Sample Rate	76
Figure 20 - Simplified schematic of the mock F&W facility	80
Figure 21 - Mock Facility with a Cascade of Surge Tanks.....	88
Figure 22 - Schematic of the Mock Facility with Diversion	91
Figure 23 - Inventory Difference for a Legitimate Run	94
Figure 24 - Modified Inventory Difference for a Legitimate Run.....	96
Figure 25 - Markov Chains for Fast Diversion after 800 Time Steps.....	99
Figure 26 - General Path Model Prior Function Development	101
Figure 27 - General Path Model Predictions for Diversion	102
Figure 28 - General Path Model Predictions for a Legitimate Run	104
Figure 29 - Particle Filter Prediction at 16 Minutes, 40 Seconds.....	107
Figure 30 - Particle Filter Prediction at 33 Minutes, 20 Seconds.....	108
Figure 31 - Particle Filter Prediction at 50 Minutes	109
Figure 32 - Particle Filter Prediction at 66 Minutes, 40 Seconds.....	110
Figure 33 - Particle Paths for the Diversion Trial Run	111
Figure 34 - Weighted Mean Particle Trace for Non-Diversion.....	115
Figure 35 - Particle Filter Statistics for Non-Diversion after 66 Minutes, 40 Seconds.....	116
Figure 36 - Prognostic Particle Traces for Non-Diversion after 66 Minutes, 40 Seconds	117
Figure 37 - Predictions for Non-Diversion up to 66 Minutes, 40 Seconds.....	119
Figure 38 - Particle Statistics for Slow Diversion.....	121
Figure 39 - Particle Traces for Slow Diversion	122
Figure 40 - Predictions of SQ Production during Slow Diversion.....	124
Figure 41 - Particle Statistics for the Fast Diversion	126
Figure 42 - Particle Traces for the Fast Diversion	127
Figure 43 - Predictions of Significant Quantity Prediction for Fast Diversion.....	129
Figure 44 - Particle Statistics for Intermittent Diversion at Time = 60 minutes.	131

Figure 45 - Particle Traces for Intermittent Diversion at Time = 60 minutes.	132
Figure 46 - Particle Statistics for the Intermittent Diversion Run at Time = 125 minutes.	134
Figure 47 - Particle Traces for Intermittent Diversion at Time = 125 minutes.	135
Figure 48 - Predictions of Significant Quantity Production during Intermittent Diversion	138
Figure A1 - Mock F&W Facility - A Snapshot.....	161
Figure A2 - A Tank on Scale at a Tails Station	163
Figure A3 - A Product Tank on Station	164
Figure A4 - Control Panel for a Feed Pump.....	165
Figure A5 - Product and Tails Needle Valves, and Surge Tank Cutoff Valve	167
Figure A6 - OHAUS Scale Monitor	168
Figure A7 - Diversion Line with Collection Tank in Place	169
Figure A8 - Diversion Line Control Wheel	170
Figure A9 - PI Control Computer with LABView™ Screenshot	171
Figure A10 - PI Controller Screen Capture	173
Figure A11 - Low Measurement Noise Estimate and Near-Constant Particle Weights.....	176
Figure A12 - High Measurement Noise Estimate and Uncontrolled Particle Distribution	177
Figure A13 - Accurate Dynamic Noise and Balanced Particle Distribution	178
Figure A14 - Low Dynamic Noise Estimate and Slow Particle Response	181
Figure A15 - High Dynamic Noise Estimate and Excessive Particle Response	182
Figure A16 - Weighted Particle Distribution when System Measurements are out of Range	184
Figure A17 - Weighted Particle Distribution when System Dynamics are Low	186
Figure A18 - Dynamic Noise Estimate Updating and Controlled Particle Variance	187
Figure A19 - Feed and Withdrawal Profile for Legitimate Operation	189
Figure A20 - Cumulative Inventory Difference for Legitimate Operation	190
Figure A21 - Load Cell Profiles for Slow Diversion	191
Figure A22 - Cumulative Inventory Difference for Slow Diversion	192
Figure A23 - Load Cell Profile for Fast Diversion.....	193
Figure A24 - Cumulative Inventory Difference for Fast Diversion	194
Figure A25 - Load Cell Profiles for Intermittent Diversion.....	196
Figure A26 - Cumulative Inventory Difference for Intermittent Diversion	197

Table of Equations

Equation 1 - Inventory Difference for Special Nuclear Material.....	12
Equation 2 - Historic Formulation of Material Unaccounted For (MUF)	12
Equation 3 - Likelihood Ratio Test for Cumulative MUF.....	12
Equation 4 - Likelihood Ratio Test for Page's Test.....	13
Equation 5 - Functional Form for the General Path Model	26
Equation 6 - Minimum Time to Failure for Multiple Failure Modes.....	26
Equation 7 - Weighted Least Squares Algorithm	29
Equation 8 - Exponential Smoothing Scheme for Weighted Least Squares Coefficients	29
Equation 9 - Kalman Filter Target Location.....	33
Equation 10 - Kalman Filter Measurement Relationship	34
Equation 11 - Chapman-Kolmogorov Equation	34
Equation 12 - Bayes Rule.....	34
Equation 13 - Target Vector Prediction under the Gaussian Assumption	34
Equation 14 - Degeneracy Algorithm for Sequential Importance Resampling.....	36
Equation 15 - RUL Prediction from Orchard's First PF Prognostic Method	37
Equation 16 - Regularization of the PF Prediction with the Quadratic Kernel	38
Equation 17 - AIC Measure for Model Selection.....	40
Equation 18 - Bias-Corrected Information Complexity.....	40
Equation 19 - Takeuchi's Information Criterion.....	40
Equation 20 - Information Complexity and Inverse Fisher Information, General Form.....	41
Equation 21 - C1 Covariance Complexity for ICOMP(IFIM).....	41
Equation 22 - ICOMP(IFIM) for Protection against Misspecification	41
Equation 23 - C1 Covariance Complexity for ICOMP(IFIM)Misspec	42
Equation 24 - Conceptual Definition of TMUF.....	46
Equation 25 - Unit Degradation Relationship for the General Path Model.....	48
Equation 26 - Fitted General Path Model	49
Equation 27 - Linear Regression Equation in Matrix Form	50
Equation 28 - Pseudoinversion to Solve for the Parameter Vector.....	51
Equation 29 - General Path Historical Model for Safeguards	51
Equation 30 - Bayes' Theorem	53
Equation 31 - Importance Function for Sequential Importance Sampling	54
Equation 32 - Definition of Sampling Weights.....	54
Equation 33 - Approximation for Sampling Weights	54
Equation 34 - Particle Filter State Estimation	57
Equation 35 - Particle Filter Future Failure Probability Estimation	57
Equation 36 - Particle Filter Future Failure Probability Distributions.....	57
Equation 37 - Weight Update Equation	58
Equation 38 - Linear model for MUF Estimation	59
Equation 39 - Model Coefficient Update Equation.....	60
Equation 40 - PI Control Equation.....	69
Equation 41 - Relationship Between Water Level and Pressure	69
Equation 42 - Observability Relationship.....	79
Equation 43 - Definition of Product Flow parameter p_e	81
Equation 44 - Relationship between Product and Tail Flow	81
Equation 45 - Mass Flow into the Surge Tank.....	81
Equation 46 - Exit flow Relationship from the Surge Tank	81

Equation 47 - Linearization of the Water Level / Flow Relationship	82
Equation 48 - PI Controller governing equation	82
Equation 49 - Variable Substitution for the PI Integral Term	83
Equation 50 - PI Controller Equation in Linear Coupled Form.....	83
Equation 51 - Surge Tank Mass Relationship	83
Equation 52 - Matrix Form of State-Space Representation	85
Equation 53 - Definition of c_1 constant in State-Space Equation.....	85
Equation 54 – Column Definitions for C Matrix for Measurement of Product and Tail Mass.....	85
Equation 55 - C Matrix with Feed, Product, and Tails Observed	86
Equation 56 - Observability Matrix (Numeric).....	86
Equation 57 - C Matrix with Feed and Tails Observed	87
Equation 58 - State Space Model for the Surge Tank Cascade	89
Equation 59 - State Equation for Diversion Flow	90
Equation 60 - State Equation for Product Flow with Diversion Present.....	90
Equation 61 - State Space Representation with Diversion	90
Equation 62 - C Matrix for the Diversion Scenario	92
Equation 63 - Probability Matrix for Fast Diversion after 800 Time Steps	97
Equation 64 - Prior Function for the General Path Model.....	100
Equation A1 - Sequential Importance Sampling Update Equation	174

List of Abbreviations

AIC	Akaike's Information Criterion
CIC	Bias Corrected Information Criterion
BOL	Beginning of Life
CID	Cumulative Inventory Difference
CUMUF	Cumulative Material Unaccounted For
EOL	End of Life
F&W	Feed and Withdrawal
FFT	Fast Fourier Transform
FS	Feed Station
GCEP	Gas Centrifuge Enrichment Facility
GP/GPM	General Path (Model)
HEU	Highly Enriched Uranium
IAEA	International Atomic Energy Agency
IC	Information Criterion
ICOMP	Information Complexity
IFIM	Inverse Fisher Information Matrix
LEU	Low Enriched Uranium
MC	Monte Carlo
MLE	Maximum Likelihood Estimator
MSE	Mean Square Error
MUF	Material Unaccounted For
NPT	Nuclear Nonproliferation Treaty
ORNL	Oak Ridge National Laboratory
PF/PFM	Particle Filter (Model)
PI	Proportional-Integral
PS	Product Station
RTMES	Real Time Mass Evaluation System
RUL	Remaining Useful Life
SIR	Sequential Importance Resampling
SIS	Sequential Importance Sampling
SQ	Significant Quantity
SSE	Sum of Squared Error
TIC	Takeushi's Information Criterion
TPM	Transition Probability Matrix
TS	Tails Station
TTF	Time to Failure

1 INTRODUCTION

One of the central missions of the International Atomic Energy Agency (IAEA) is to monitor the enrichment of uranium at declared facilities such as gas centrifuge enrichment plants (GCEPs) [IAEA, 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]. The two primary goals of the inspection process are to verify that the facility is enriching only as much material as is declared, and to verify that "Significant Quantities" (SQs) [IAEA, 2002] of undeclared special nuclear material such as low-enriched uranium (LEU) or high-enriched uranium (HEU) are not produced [IAEA, 2002]. (For LEU, a Significant Quantity is the amount of material that contains 75 kg of U-235. For example, approximately 2,000 kg of 5% enriched UF_6 meets this standard.) The current method of verification of GCEP activity is to periodically send inspectors to the facility to ensure that only declared feed cylinders are brought on site, while only declared tails and product cylinders are removed from the facility. Additionally, the material balance is reviewed alongside assay of the feed, product, and tails cylinders to ensure that the material feed into the enrichment process has been extracted and accounted for, and that the balance of U-235 is maintained. The process is rather cumbersome to the IAEA, who must maintain a sufficient staff and travel budget of inspectors, as well as the enrichment facility operator, who must often wait on the inspector to verify cylinders before using them, or who must occasionally stop the enrichment process so that the cylinders at feed or withdrawal stations may be measured and verified.

Currently, the global increase in nuclear power is causing a continual and sharp increase in the number of nuclear power plants (and consequently, the number of enrichment facilities) [Sokolov, September 2006]. As the number and size of enrichment facilities continues to grow, the financial and manpower burdens on the IAEA increase correspondingly. However, the operating budget of the IAEA is not currently growing at the same pace as worldwide uranium processing, a discrepancy that may eventually preclude the IAEA from effectively monitoring GCEPs and other enrichment facilities and fulfilling their verification mission. Knowing this, the IAEA is interested in developing alternative monitoring methods that may save time and manpower without sacrificing inspection effectiveness (and possibly even increasing the effectiveness of IAEA monitoring).

One solution that lends itself naturally to the problem is to adapt remote process monitoring techniques for verification of GCEP operation [Dixon, et al., 2007; Laughter, 2009; Laughter, et al., 2010]. Use of

process monitoring as a verification tool is neither a new idea [Speed, et al., 1986] nor is limited to GCEPs, and process monitoring has also been proposed to assist verification for other nuclear processes, particularly for reprocessing facilities where transuranic elements (especially plutonium) are involved [Burr, et al., 2008]. In addition to the continuation of research at Los Alamos National Laboratory into statistical and methodological processes for monitoring, coordinated research is currently being conducted at Oak Ridge National Laboratory (ORNL), Savannah River National Laboratory (SRNL), and Sandia National Laboratory (SNL) to develop a process monitoring system that could meet the verification requirements of the IAEA as well as the data protection requirements of the GCEP operators. The research of these national laboratories focuses largely on load cell monitoring of the feed, tails, and product stations to maintain a constant account of the material entering and leaving the "cascade" area of a GCEP – the network of centrifuges that perform the actual task of enrichment. Many challenges must be solved to develop a reliable monitoring system, such as the development of a data communication network that cannot transmit proprietary operator information off-site, the certification of accountancy scales that can verify cylinder weights without necessarily having an inspector on-site, and adequate protection of the system against spoofing and other forms of data manipulation. The research of this dissertation, however, will focus on analyzing the load cell and accountancy scale data and developing automated methods of drawing safeguards conclusions that serve the IAEA's mission of verification.

As with the current inspection process, any new techniques must protect the operator's proprietary and confidential information [IAEA, 2002]; for example, inspection cannot be so invasive as to learn the operator's methods for maximizing efficiency of the enrichment process. To meet this need, the amount of data that is monitored must be minimized to that which is necessary for verification, and only the necessary conclusions and information may be passed beyond the physical boundaries of the facility. For this reason, any process monitoring-based method of inspection and verification must be robust enough to be applied on-site without requiring constant inspector presence while still communicating the necessary information to IAEA headquarters in Vienna, Austria.

At ORNL, a mock feed and withdrawal (F&W) facility has been developed so that process monitoring techniques may be tested and refined without disturbing the operations of a GCEP [Krichinsky, et al., 2009]. The facility uses water as a substitute for UF_6 , greatly reducing cost and eliminating the need for radiation and special materials controls. With three feed stations, two tail stations, and three product

stations, the facility may be run continuously and has provided data that can be analyzed in a similar manner to that of GCEP operation [White, et al., 2010]. Though some features of the enrichment process are lost (e.g. the ability to perform destructive analysis (DA) and nondestructive analysis (NDA) of the product and tails materials to determine enrichment levels), the system as a whole serves as a suitable analog to GCEP load cell activity, allowing for the development of load cell monitoring techniques at the "proof of concept" level. Further discussion of the mock F&W facility is provided in Appendix A.

For this research, the mock F&W facility will be used in place of real GCEP data. All conclusions are intended to illustrate the techniques at a scale level so that the IAEA and GCEP operators may understand how an automated process monitoring and prognostic system may work. With the potential for upscaling to real GCEP monitoring in mind, any techniques developed in this research will be partially judged on their apparent portability to full-scale GCEP monitoring as well as the confidence and accuracy at which any safeguards conclusions are made.

1.1 Problem Statement

The shift toward information-driven safeguards of GCEP activity represents a fundamental change in the way the IAEA fulfills its verification mission. Where the traditional method of enrichment activity verification relied purely on agreement between inspector observation and facility operator declaration, the information-driven approach adds analysis of process monitoring data as a complement to inspector observations. By adding plant performance data to the inspection metrics, the IAEA hopes to reduce the gaps in material accountability (e.g. the inability to verify material processed to and from cylinders that are still in-process), increase the level of sophistication required to defeat IAEA inspection, and to establish a smarter, more efficient inspection routine. Improved inspection efficiency holds particular value to both the IAEA and the GCEP operator; by utilizing "on-demand" inspection periodicities, the overall number of inspections (as well as the time required to perform the inspections) may be reduced for GCEPs where process data indicates a high likelihood of agreement between declarations and actual enrichment activity. Reduction in the frequency and length of time of inspections is expected to provide significant savings to the IAEA through reduced inspector hours and travel, and to provide reduce cost to the operator, whose enrichment operations would experience fewer interruptions in the process.

The heart of the information-driven safeguards system the IAEA hopes to implement lies in the ability to identify behaviors in GCEP operation that may indicate a need for inspections, with a particular concern

toward identifying potential diversion or undeclared activity. Extending this goal, process monitoring information may also provide the ability to anticipate the severity of an event. (For example: if diversion is occurring, how long until a SQ of material is diverted?) Prognosis of safeguards-related faults in GCEP monitoring may further refine the ability of the IAEA to decide upon appropriate action. This research focuses on the development of prognostic methods using on-line weight measurements of feed, product, and tails cylinders at process stations.

The mock F&W facility at Oak Ridge will be used as a test platform for the development of prognostic techniques to predict the time to SQ production through either diversion or undeclared activity. Predictions of SQ production must also discriminate between illicit activities and normal features of operation that affect the MUF calculations, such as cold trap operation, cascade holdup, and instrumentation error. Additionally, uncertainty estimates must be developed so that the results of the prognosis do not merely identify a prediction of time to SQ production, but the likelihood that the data is in fact suggesting SQ production as well as the perceived variance in the prediction. Competing prognostic models will be developed using General Path and Particle Filter techniques, and the models will be compared based on the accuracy and uncertainty of their predictions.

1.2 Original Contributions

The research within this document contains distinct, original contributions to the fields of prognostics and of international safeguards. These contributions center on the development of prognostic methods for uranium enrichment verification with a specific focus of predicting the time at which a significant quantity of diverted material may be produced. In contrast to traditional health monitoring prognostics, this research focuses on diagnosing and prognosing the effects of human decisions instead of component wear and failure. Such prognostics methods are also novel to safeguards application, which traditionally focus on diagnosing the past rather than anticipating the future. The primary original contributions are described here:

1. Adaptation of prognostic methods for safeguards applications by the development of a prognostic method to predict the estimated time until diversion of a "significant quantity" of material from the mock feed and withdrawal facility.
2. Development of uncertainty estimates for the time until significant quantity diversion prediction and metrics to quantify the validity of the prognostic method.

3. Development of methods to update measurement and dynamic noise estimates to account for potentially sudden and significant decision-based changes in system behavior, such as the initiation or cessation of diversion.
4. Development of methods to discriminate anomalous behavior (e.g. diversion) from normal perturbations of the mass balance data from the mock F&W facility, such as the loss of material to cold trap emulation.

1.3 Organization of the Document

Chapter 2 reviews the literature in international safeguards and prognostics as pertinent to this research. The state of the art for prognostics methods is examined, as well as research in remote monitoring for gas centrifuge enrichment facilities. Chapter 3 discusses the equipment and methods used in this research. The Oak Ridge Mock Feed and Withdrawal Facility is discussed, and detailed information of the facility may be found in Appendix A: Description of the Mock Feed and Withdrawal Facility. The general path model and the particle filter method are described in detail in relation to prognostics applications. Chapter 4 provides the results of the research. The changes made to the Mock Feed and Withdrawal Facility are described in detail, including procedural and equipment modifications to improve data consistency. A preliminary comparison of the various prognostic models is made and the particle filter is justified as the prognostic model of choice for this research. The particle filter method is then applied to the four case studies described in Appendix C: Case Studies. Finally, Chapter 5 provides conclusions drawn from this research, as well as possibilities for future work beyond the scope of this research.

2 LITERATURE SURVEY

2.1 IAEA Responsibility for Enrichment Verification

The authority for IAEA verification of enrichment processes stems from treaty between nation states and the United Nations (UN) to promote the non-proliferation of nuclear weapons [IAEA, 1972]. Through independent verification of the enrichment facility operator's activities, the ability for a declared enrichment plant to produce weapons-grade enriched uranium or to produce LEU that is earmarked for future weapons production without international awareness is reduced. While the IAEA has no enforcement authority (that is, the IAEA cannot prevent an enrichment facility from producing HEU or LEU for weapons production), their independent verification that a facility (and, by extension, the nation) is operating only for energy production purposes is a key contributor in determining whether a nuclear nation is a significant proliferation threat [IAEA, 2009].

The primary objectives of IAEA inspections are to ensure that the enrichment facility is not diverting declared nuclear material and that the facility is not producing undeclared nuclear material [IAEA, 2002]. In both cases, the goal is to assure that an enrichment facility is not generating a "Significant Quantity" (SQ) of nuclear material, where an SQ is defined as the minimum amount of nuclear material needed to produce a pure fission weapon. Currently, the IAEA defines a SQ of LEU as the amount of LEU that contains 75 kg of U-235, and of HEU as the amount of HEU that contains 25 kg of U-235 [IAEA, 2002]. (Note: [Cochran, 1995] suggested that far lower levels are needed based on more conservative estimates of the material needed for weapons production, but the IAEA convention will be followed in this research.) This verification is considered the fulfillment of the IAEA's inspection purpose under the IAEA Statute, Article III.A.5 [IAEA, 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]:

To establish and administer safeguards designed to ensure that special fissionable and other materials, services, equipment, facilities and information made available by the Agency or at its request or under its supervision or control are not used in such a way as to further any military purpose; and to apply safeguards, at the request of the parties, to any bilateral or multilateral arrangement, or at the request of a State, to any of that State's activities in the field of atomic energy.

By verifying that no material is removed from the balance of declared nuclear material and that no undeclared material is processed within the enrichment facility, the facility is deemed to be operating for the peaceful purpose of producing enriched uranium for energy production. The authority of the

IAEA to inspect enrichment plants is provided by the IAEA Statute in fulfillment of this statement as agreed to by the nations who have signed to the Nuclear Nonproliferation Treaty (formally the Treaty on the Non-Proliferation of Nuclear Weapons, or the NPT).

2.2 Current Process for GCEP Inspection

Currently, inspections of GCEPs are conducted by IAEA inspectors who travel on-site to verify compliance with the NPT [IAEA, 2009]. The inspectors verify the accountancy weights of all feed, product, and tails cylinders that are either arriving or departing the facility, ensure that all declared cylinders are properly accounted for, and ensure that there is no undeclared material present at the facility. Every cylinder must be weighed prior to being placed in the enrichment process, and must again be weighed after processing prior to the cylinder's removal from the facility; both of these weights (called accountancy weights) are taken in the presence of the IAEA inspector. Assay of the feed, product, and tails materials are also taken to measure their respective enrichment levels. These accountancy weights are the heart of the verification process; if the processed feed stock weights match that of the of the withdrawal weights, and if the enrichment balances of the assay agree with the accountancy weights, then the declared material is considered to have been properly accounted for. The inspection process may also look for evidence of undeclared activity, such as the presence of feed and withdrawal stations that are not listed in the facility declarations.

In conjunction with the inspector's actions, the facility operator is responsible for maintaining a current and accurate inventory of all nuclear material within the facility, recording all transport of nuclear material to and from the facility, and submitting routine inventory reports to the IAEA [IAEA, 1993]. The operator's declarations are compared to the inspector's findings to determine if sufficient agreement exists between the two observations. If any discrepancies (referred to as 'material unaccounted for', or MUF [IAEA, 2002]) exist, they are resolved by searching for reasonable technical origins, such as measurement uncertainties. If MUF cannot be satisfactorily explained, then the inspector reports an inability to verify that no diversion has occurred. With the exception of measured added under the additional protocol of 1998 [IAEA, 2009; IAEA, Model Protocol Additional to the Agreement(s) between State(s) and the International Atomic Energy, 1998], the inspection process has remained largely unchanged since its inception.

2.3 Next-Generation Techniques for GCEP Inspection and Monitoring

A shift toward the implementation of an "information-driven" approach to safeguards has been documented since the discovery of the Iraq covert nuclear program in 1991 [Frazar, et al., 2010; IAEA, 2005]. Applying information-driven concepts to GCEP inspections is a topic that has been discussed as early as the 1970s [Speed, et al., 1986] with a revival of interest in the last fifteen years [Howell J. W., 1995], though application has yet to materialize largely due to the extensive amount of negotiation between the IAEA and enrichment facility operators [Friend, 2008; Friend, 2010]. Much of the interest in enhanced monitoring techniques is driven by the desire to improve the timeliness of detection [Tsvetkov, 2007] with regards to diversion, undeclared activity, and higher than declared enrichment [Pickett, et al., 2008] and to close existing gaps in continuity of knowledge of verified activity at GCEPs [Curtis, 2009]. Another reason for interest is the desire to increase inspection efficiency so that inspection costs may be reduced for both the IAEA and the facility operator [Gyane, 2010].

The concept of information-driven safeguards is hardly new [Cobb, 1981; Shipley, 1983] and is not isolated to monitoring of GCEPs. Over the past 15 years, a significant amount of research has focused on remote monitoring techniques for spent fuel reprocessing facilities [Burr T. , et al., 1995]. The presence of transuranic elements in spent fuel (especially plutonium) and their isolation during reprocessing has resulted in a greater focus on reprocessing facilities than enrichment facilities during this time [Burr, et al., 1999; Longmire, et al., 2002; Burr T. , et al., 2003; Burr, et al., 2006] in part due to the greater sensitivity toward undeclared plutonium production within the IAEA's definitions of significant quantities. Where uranium enrichment SQ limits are focused on U-235, the presence of plutonium in reprocessing invokes the IAEA SQ limit of material containing 8 kg of plutonium [IAEA, 2002]. (The far more restrictive limit for plutonium and the specific elemental separation of plutonium that occurs in reprocessing results in a greater interest in verifying that undeclared plutonium is not diverted. This attention is largely responsible to the greater research interest that was shown in reprocessing monitoring over enrichment monitoring from approximately 1995 to 2005.) The research efforts in the development of online monitoring techniques for spent fuel reprocessing may be seen to mirror efforts applied to GCEPs. In both cases, the primary focus is placed on closing "knowledge gaps" in material inventory balances and in isotopic assays through remotely monitored sensing equipment so that both the between-inspection uncertainties and the inspection periodicities may be reduced (compare [Burr, et al., 2008] and [Curtis, 2009]).

2.3.1 Load Cell Monitoring and Automated Analysis

Research efforts are currently being conducted within the United States [Krichinsky, et al., 2008; Krichinsky A. M., 2009; Garcia, 2010; Boyer, 2010; Durst, 2008; Lockwood, 2010; Laughter, et al., 2008] and internationally [Howell, et al., 2009; Delbeke, et al., 2008; Delbeke, et al., 2007; Dixon, et al., 2006; Howell J. , 2008; Howell, et al., 2007] to develop new monitoring techniques that help the IAEA meet their enrichment monitoring mission within the constraints of their budget. Load cell monitoring has been central to most research efforts for next-generation verification of GCEP and other enrichment facility activity [Howell J. , 2008; Krichinsky A. , 2008; Whitaker, et al., 2009; Laughter, 2009]. The central focus of load cell monitoring has been to maintain a running calculation of the MUF within a facility. The MUF value may include losses from cold traps, the initial charging of the cascade, and natural variations in holdup within the cascade due to normal operations. A typical example of a MUF calculation over time is presented by Howell in Figure 1 from an experiment in load cell calculations at Capenhurst [Howell, et al., 2009]. Natural fluctuation in the MUF value may be seen, and the positive upward trend was attributed to a scale bias in the tails station load cell.

Similar calculations have been performed on operational data from the mock F&W facility, as shown in Figure 2. Unlike Figure 1, which shows MUF as a positive value, the mock F&W facility data represents MUF as a negative value. Here, the effects of holdup can be seen in facility startup and shutdown (near 400 and 1600 units of time, respectively), but the MUF remains relatively constant over time during operation. Normal fluctuation due to facility operation may be seen, and perturbations near the 1000 units of time mark are attributable to switchovers between feed and tails tanks. (Further discussion of the mock F&W facility may be found in Appendix A.)

As with Howell's work at Capenhurst, the remaining literature on analytical methods for verification of enrichment activities focuses on calculation of the MUF and application of statistical measures to evaluate and interpret the data. In virtually all instances, the central statistic for monitoring the material balance is the inventory difference (ID), which is defined by [Burr, 2008] as Equation 1, where ID_t is the inventory difference at time t , BI_t is the beginning inventory (typically calculated by an audit), R_t is the receipt of new material, EI_t is the ending physical inventory (also by audit), and S_t is the shipment of material out of the control area. By traditional inspection methods, where the inspectors calculate the various terms of Equation 1 by manual audit of the inventory at a facility, the measurements are taken with a periodicity of the inspection visits themselves, which is typically performed monthly.




Figure 1 - MUF Calculation from Capenhurst Load Cell Evaluation [Howell, et al., 2009]




Figure 2 - MUF Calculation from the Mock F&W Facility

$$ID_t = BI_t + R_t - EI_t - S_t$$

Equation 1 - Inventory Difference for Special Nuclear Material

Indeed, Equation 1 represents the core of all inspector-based verification performed by the IAEA and is identical to Equation 2 [Speed, et al., 1986], where $I(n-1)$ and $I(n)$ are the respective inventories at the beginning and ending times of the evaluation interval and $T(n)$ is the net transfer of material in and out of the material balance boundary. Material Unaccounted For (MUF), is equivalent to ID and is the common term within the IAEA for the material balance discrepancy. The annual facility-wide audits provided the estimates for the beginning inventory and ending inventory for a given interval, and the monthly inspections provided calculations for the receipts of new feed material and shipments of the enriched product and depleted tails materials. If the calculated inventory differences fell outside previously determined thresholds, further inquiry could be performed to resolve the apparent loss (or gain) of material from the facility.

$$MUF(n) = I(n - 1) - I(n) + T(n)$$

Equation 2 - Historic Formulation of Material Unaccounted For (MUF)

Both [Burr, 2008] and [Speed, et al., 1986] point to statistical tests derived from the calculation of the inventory difference. One such test is the CUMUF (Cumulative Material Unaccounted For) test, where the MUF is calculated at each inspection interval as an accumulation since the previous facility audit. In other words, the material balance at the previous audit resets the MUF to zero; each inspection thereafter adds or subtracts the material balance discrepancy to the current MUF tally, creating an integrated calculation of the MUF between audits. Using this approach, a likelihood ratio test was formulated in [Cobb, 1981] where the CUMUF fails the test if Equation 3 is true. This particular test is suspect due to its assumption that each measurement X_i is independent, which cannot be guaranteed.

$$\sum_{i=1}^n X_i > c \sqrt{(n\sigma_T^2 + 2\sigma_I^2)}$$

Equation 3 - Likelihood Ratio Test for Cumulative MUF

Page's test was commonly proposed as an alternative procedure around the same time [Woods, et al., 1980; Pike, et al., 1982; Jones, 1984]. The test is applicable with ordered observations (e.g. time series) and allows the testing of the hypothesis of a trend against its null. In this case, the suggested hypothesis is that the MUF remains constant, which is tested against the null of a nonconstant MUF. Accordingly, the likelihood ratio for Page's test reduces to Equation 4, where m_b and m_e are the beginning and end of the evaluation period, respectively, and θ_0 is positive. The literature suggests several variations on this test, such as applying the test to standardized residuals or adding robustness measures.

$$1 \leq m_b \leq m_e \leq N \left[\sum_{i=m_b}^{m_e} \left(X_i - \frac{1}{2} \theta_0 \right) \right]^+ > c$$

Equation 4 - Likelihood Ratio Test for Page's Test

Even in the early stages of IAEA inspection, a significant amount of research was directed toward the development of monitoring methods for continuous data streams rather than monthly inspection reports. As James Shipley noted [Shipley, 1983]:

Currently, one technique being developed for generating timely materials accounting data is near-real-time accounting, which is based on obtaining inventory information without interrupting the processing of nuclear material. The purpose of near-real-time accounting systems is to detect anomalies, possibly resulting from diversion, in a timely fashion and then to localize them for investigation.

Shipley then proposes several possible methods to realize near real-time accounting, focusing especially on various formulations of the sequential ratio probability test (SPRT) such as Wald's SPRT, a uniform diversion test, a sequential variance test, and SPRT derived from the CUSUM. Such tests have obviously yet to be realized in practice, as the IAEA is still relying on monthly inspection reports to satisfy their mandate for verification of enrichment declarations, but the delay in transitioning to a near real-time monitoring system certainly has not been for lack of desire on the part of statisticians and researchers.

Shortly after [Speed, et al., 1986], interest (or, more likely, funding) in process monitoring of enrichment facilities stagnated. Throughout the late 1980s through the 1990s, most of the research in near real-time accountancy was performed for reprocessing scenarios rather than enrichment, due largely to the political interest in plutonium management in spent nuclear fuel, as in [Burr, et al., 1995]. Although not directly related to this research, the reprocessing monitoring efforts of this time are noteworthy in that

the period represented a "changing of the guard" of the researchers involved, as the next generation of process statisticians were focused on the reprocessing problem at the time. Like enrichment monitoring, near real-time techniques were proposed for reprocessing, usually under the term *solution monitoring* [Burr, et al., October 1997] (also printed as [Burr T. , et al., 2003]).

In Europe, current work in process monitoring for enrichment verification is centered on the development of the Real Time Mass Evaluation System (RTMES), discussed in [Delbeke, et al., 2008; Delbeke, et al., The Detection of Undeclared LEU Production at a GCEP by Real-Time Mass-Balancing, 2007]. In this report, the typical arrangement of load cells for product, feed, and tails stations is anticipated, and the data is used to generate continuous estimations of the MUF, along with cumulative (integrated) MUF calculations. In this report, preliminary demonstrations of results that might be expected from the RTMES assume that a protracted diversion is accompanied by an increase in product flow rate such that the flow rate to declared product stations remains constant (and the excess is diverted as undeclared product). While not a necessary step for diversion, such a scheme is intended to disallow use of the product load cell data only as an indicator of diversion, since the declared product flow rate would appear constant. The results of such a system could then be compared to the operator's declarations [Howell, et al., Data Consistency Evaluation in GCEPs, 2007] such that the total material transfer is checked against the declarations and the internal processes are checked via CUMUF, likelihood tests, and assay measurements to generate a holistic picture of enrichment activity. Early case study literature cites data taken from the URENCO (Capenhurst) Limited (now URENCO UK Limited, or informally Capenhurst) for proof of concept of load cell monitoring [Howell, et al., 2009]. While care would obviously be necessary to differentiate safeguards-relevant and proprietary information so that IAEA inspection routines do not become a potential leak of trade secrets, this is a topic that has yet to be resolved between the IAEA and enrichment facility operators regarding remote process monitoring systems.

It should be noted, however, that while the term RTMES implies a real-time system, where results are provided instantly, the actual concept envisioned in [Delbeke J. , et al., 2008] is in fact a near real-time system, with the "near" dropped due to the rapid analytical response of the system relative to the IAEA's reaction time (i.e. that the system may report results daily though the IAEA may not be able to send inspectors on a triggered inspection for several days due to travel logistics). In this sense, the term "real time" is a simplification for the sake of an easier name (RTMES vice NRTMES). This convention is carried

over from [Dixon, et al., 2006]. There are currently no true "real time" processes in the existing literature, likely due to the significant logistical overhead presented by real-time data handling and processing.

In the United States, the "renaissance" (for lack of a better term) in process monitoring research has been more holistic. The proof-of-concept of load cell monitoring was performed at the Portsmouth gaseous diffusion facility [Krichinsky, 2008]. Preliminary development and testing of monitoring algorithms has been conducted using the Mock Feed and Withdrawal Facility at Oak Ridge National Laboratory [Krichinsky, et al., 2009], including the work in this research. Such efforts represent the most thorough attempts to date to model diversion, off-scale activity, altered declarations, and other scenarios on the basis of load cell data. Use of data generated from a physical facility – even if just an analog – represents the first "step up" from computer simulated data. The mock facility is also used to develop and test data management solutions to find and demonstrate acceptable means to record, transmit, analyze, and protect load cell monitoring data [Garner, et al., 2011].

Additionally, research for next-generation safeguards verification is not only being conducted on load cell monitoring techniques [Krichinsky, et al., 2008; Krichinsky, 2008; Lenarduzzi, et al., 2007] but is also focused on all other facets of enrichment facility inspection (e.g. non-destructive assay (NDA), surveillance, and cylinder identification) [Laughter, et al., 2010]. The goal of these research efforts is to produce and demonstrate a comprehensive set of tools that the IAEA may use to increase the efficiency and reliability of their inspection processes while reducing the financial and manpower burdens on the agency [Whitaker, et al., 2009]. These additional fields are mentioned briefly, but are not expounded upon in this report as their development is independent of the development of load cell monitoring methods.

2.3.2 Non-Destructive Assay

One facet of enrichment process monitoring is the verification that feed, tails, and product cylinders contain uranium at the levels of enrichment that the facility operator has declared. The difficulty with spectroscopy is the thickness of the cylinder and material, which self-shields the signature from the innermost contents of the cylinder. Research is being performed to identify and develop NDA techniques that may provide information about the contents of the cylinder at a greater depth and with greater reliability.

2.3.3 Facility Surveillance

The surveillance of cylinder storage and processing areas remains a relatively undeveloped area of improvement in enrichment monitoring due to the delicate balance between comprehensive materials accountability and the protection of facility intellectual property. While constant video surveillance by the IAEA would allow for aggressive inspection and verification of facility activities, such data could also be used to provide insight into a particular facility's operational tendencies. Enrichment facility operators are rightly concerned that such information may allow others to learn their trade secrets, nullifying any economic advantage to maintaining the secrets. Additionally, public proliferation of videos of GCEP operations may conceivably aid other non-nuclear states in developing their own nuclear programs, resulting in an even greater proliferation risk. The (very legitimate) facility-side concerns have made the video surveillance problem a mostly logistical issue, and the development of surveillance protocols that aid the IAEA inspectors while protecting the facility's proprietary knowledge is a matter of cooperation rather than innovation.

2.3.4 Cylinder Identification

A considerable amount of inspector's effort is spent maintaining inventory of feed, tails, and product cylinders. Inspectors must visually identify each cylinder in order to maintain continuity of process knowledge from the time the cylinder is brought into the facility until it leaves. If a cylinder is not in its expected location, the facility must be visually searched to find the location of the cylinder. Various cylinder identification techniques (e.g. radio frequency identification (RFID) or global serial numbers) [Hori, et al., 2008; Pickett, et al., 2008; Pickett, et al., Results from a "Proof-of-Concept" Demonstration of RF-Based Tracking of UF₆ Cylinders During a Processing Operation at a Uranium Enrichment Plant, 2008] are currently being researched to find ways to automate the inventorying process and to provide more comprehensive and reliable data on the life cycles of the cylinders. Such efforts are intended to enable automated cylinder identification and tracking so that the location and disposition of a cylinder may be determined without the need for the inspector to arrive at the facility and visually identify the cylinder.

2.3.5 Inspection Periodicity

Currently, inspections are performed on a regular schedule to meet an IAEA requirement that no more than one month elapses between inspections. There is ongoing discussion about shifting from a periodic inspection routine to a random and/or event-triggered routine; by moving away from periodic inspections to a routine that does not allow the facility to know the date of the next inspection, it is

hoped that the number of inspections per annum may be reduced. By incorporating improved monitoring techniques, current estimates suggest that as few as five or six inspections per annum may be sufficient to provide the same verification of declared facility operation as the current twelve-per-annum routine, and that inspections may be smartly timed to coincide with facility operations such as cylinder loading or unloading [Laughter, et al., 2010]. This reduction in inspection frequency would reduce the burden on the IAEA as well as reduce the number of work stoppages required by the facility; both parties would benefit financially from a reduced inspection frequency, making this a particularly desirable improvement in the inspection process.

2.3.6 Mailbox Declarations

An additional safeguards technique of interest to the IAEA is the concept of mailbox declarations. In a mailbox declaration, the facility operator would provide a timely electronic message to the IAEA regarding the disposition of the facility. For example, at the time the mailbox declaration is made, the mailbox declaration may state which feed and withdrawal stations have cylinders present, and the declaration may state the ID of the cylinder on the station. By providing basic information of the condition of the cylinders, the IAEA has more information between inspections that may be used for verification of activity when the next inspection occurs [Korbmacher, 2008]. For example, if the most recent mailbox declaration stated that a certain cylinder had just been loaded onto a certain feed station and the inspection reveals that there is no cylinder at that feed station, then the discrepancy could raise a red flag that would warrant further investigation into the activities of the facility. Likewise, if the activity indicated in the mailbox declarations match the activity observed by inspection, then the IAEA may have increased confidence that the facility is operating as declared. The declarations are intended to help increase the continuity of knowledge and help the IAEA maintain a faster response time to potential abnormalities.

2.3.7 Automated Analysis at the Mock Feed and Withdrawal Facility

Part of the research effort with the Oak Ridge mock feed and withdrawal facility has focused on demonstrating automated analysis tools for process verification in a practical setting. The facility itself is designed to allow the demonstration of nearly all the aforementioned analytical techniques (excluding isotopic assay), including inventory monitoring, cylinder identification, on-line inspections, and database architecture [Garner, et al., 2011]. The process data recorded from the mock F&W facility may be analyzed using a MATLAB™ based software tool titled *PlotEvents*, written originally by James Henkel

[Henkel, 2010] and updated by both Henkel and this author [Hooper, et al., 2011; Henkel, et al., 2011]. Based on process station load cell data, the *PlotEvents* tool counts cylinders, identifies station events (e.g. loadings and unloadings), calculates the amount of material processed as functions of time, calculates inventory differences, and provides automatically generated reports on the operation of the facility for time periods of interest.

Sample screenshots of the *PlotEvents* tool are given in Figure 3 through Figure 6. Figure 3 displays the main *PlotEvents* screen, where the process data is illustrated for visual review. After clicking the "Analyze Events" button, Figure 4 is presented, where individual process stations may be reviewed for their usage. The Cumulative ID may also be analyzed in this window, as shown in Figure 5 and Figure 6. The "Write Summary Report" button provides a detailed overview of the facility operation for the time in question, including the number of tanks processed, the amount of material processed, and the change in MUF over the process. In Figure 4 through Figure 6, the "ID Analysis" and "Prognostics" radio buttons in the "Select Analysis" window are reserved for implementation of work in this research and that of James Henkel in his concurrent research.

With the *PlotEvents* software, the mock F&W facility data may be rapidly analyzed for operational features of interest to an inspector. By using a simple, graphics-heavy interface, the inspector would not require intimate understanding of the algorithms. Instead, inspector time and effort would be focused on interpreting the data more intuitively and using the findings to enhance and target the upcoming inspection efforts. Automated report generation of visual representations of the facility operation and of text-based interpretations of the events during the period of interest also serve to reduce the inspector's workload by providing the data in pre-formatted and easily understood terms that the IAEA and facility operators can review. Use of automated tools like *PlotEvents* to analyze the data and to generate results allows the inspection process to focus on known issues rather than blindly reviewing the facility's operational history.




Figure 3 - PlotEvents Main Screen: Process Station Weight History over Time




Figure 4 - PlotEvents Analysis of Process Station Usage over Time




Figure 5 - PlotEvents Depiction of Process Scale Inventory Differences




Figure 6 - PlotEvents Net Inventory Difference for a Run of the Mock F&W Facility

2.4 Process Monitoring and Prognostics

The field of process monitoring includes many disciplines; system monitoring, diagnostics, fault detection, prognostics, and health monitoring are a few such topics within the process monitoring umbrella. Process monitoring may be used to understand the state of a system, help plan and execute maintenance of processes (and process components), search for inefficiencies, and anticipate future changes to the system such as remaining useful life (RUL) or time to failure (TTF) estimates. Such an approach is typically used for mechanistic systems where wear of components is a natural consequence of their use. While this is a departure from the verification process (where undeclared activity would be considered a departure from normal use and therefore not a natural consequence), the methodology may be applied by analogy to the safeguards realm. Here, a discussion of prognostics is presented to provide context and an understanding of the discipline. Application of prognostics to safeguards is then presented in Sections 3 and 4 of this report.

Prognostics is the discipline of process monitoring that is of primary concern in this research. Definitions of prognostics with regards to process monitoring vary widely; this research will focus on the RUL approach to prognostics, where the prognosis is the estimation of the RUL of a system or component. As stated by Lu and Meeker [Lu, et al., 1993], such prognostics are concerned with developing a measure of the degradation of a component (or system), determining correlations between the degradation parameter and factors within the environment and the component, and using those correlations to make a prediction of the time remaining until failure of the component.

Traditional RUL prognostics focus on the time at which a system component reaches a predefined end of life criteria, either a hard failure or a soft failure [Orsagh, et al., 2006]. RUL estimates of actual failure are often termed hard failure estimates. Such failures typically center on component breakdown, like an incandescent light bulb with a burned-out filament. Soft failures, on the other hand, define a level of wear or degradation where a component or a system can no longer be trusted to meet its design specifications, like a light bulb that no longer emits its rated luminosity. The maintenance of automotive motor oil is an example of soft failure; the recommended oil change periodicity (typically 3,000 miles) is not based on an actual failure of the oil, but rather on the time when the engine-cleaning additives within the oil can no longer be trusted to effectively remove carbon deposits within the engine. With hard failures, the actual time to failure is seen as a distribution where both the mean and the variance of the RUL are desired. Soft failures, conversely, are defined limits (such as 3,000 miles for oil changes)

that do not have associated distributions; once the soft limit is reached, the component is declared to be failed, even if it appears to be working normally.

For the purposes of adapting prognostic techniques to GCEP verification, the SQs as defined by the IAEA [IAEA, 2002] may serve as soft failure limits. For diversion of LEU or undeclared production of LEU, a soft limit of 75 kg is in keeping with the IAEA's objectives for verification; for production of HEU or LEU at higher enrichment level than declared, the soft limit of 25 kg is appropriate. As was mentioned previously [Cochran, 1995], these SQ limits may be debatable as effective levels for safeguards, but the actual value of the soft limit generally does not affect the utility of a prognostic technique and is largely a moot point with regards to this research.

2.4.1 Type I, II, and III Prognostics

When developing predictions about the RUL or TTF of a component or system, the prognostic model must account for the characteristics of the component and the environment in which the component is used. By considering any potential influences on RUL within these two umbrella categories, three different types of prognostic approaches may be considered [Hines, et al., 2008].

Type I prognostics are the oldest methods, providing TTF estimates based on average environmental conditions and average lifetimes of components or systems of interest. Light bulb life expectancies are an example of Type I prognostics; if a light bulb is rated for 10,000 hours of use, that rating is based on an average light bulb used in a region of some standard temperature and humidity. (Note: the limit for light bulbs is typically a soft limit set to assure a minimum life for most bulbs; however, the classification of the limit as Type I is not based on the type of limit but rather the assumptions made in setting the limit.) Type I prognostics include distribution analyses, such as Weibull or exponential analysis.

In Type II prognostics, components are still evaluated as the historical "average" component in terms of durability and lifespan, but variations in environmental factors are also considered. If environmental conditions can be monitored and if correlations can be drawn between these conditions and the life of the component, then this additional information can be factored into the prognostic model to provide a more knowledgeable (and nominally, more accurate) estimate of RUL. Type II methods include shock models, Markov Chains in some cases, and proportional hazards models. In the case of proportional hazards models, Type II prognostics are often used to develop accelerated life testing models, where environmental factors are driven to harsher extremes (e.g. elevated temperatures or electrical currents)

to hasten wear. If the relationship between the wear rate and the environmental stressors are understood *a priori*, accelerated life testing can provide EOL data for components that are expected to last for a long time. In such cases, Type I EOL testing may not provide results in a timely manner. Type I models, however, may sometimes be updated using Bayesian knowledge of the environment to produce type II models. Such is the case with Markov Chains, where monitoring the state transitions of a component can allow for Bayesian updating of the transition matrix (TM).

Type III prognostics incorporate both environmental conditions and component condition (estimated by sensor data) into the RUL model. By accounting for environment and specific component characteristics, Type III models have the potential to provide the most reliable predictions of RUL. The General Path Model (GPM) is the most common example of a Type III model, but all models of this class feature the ability to monitor individual components within individual environmental conditions. This class of prognostic models is particularly useful for "critical" applications such as safety systems. Type III prognostics are often identified by a Bayesian process that updates both the environmental and component-specific influences on RUL.

2.4.2 Methods for Type III Prognostics

Two methods appear to have significant favor for Type III prognostics: the general path model and particle filtering.

2.4.2.1 General Path Prognostics

Originally developed by Lu and Meeker [Lu, et al., 1993], the general path model was among the first Type III methods to gain widespread favor and is based on establishing measures of degradation rather than simply tracking failure times. General Path models develop a degradation path based on historical data to provide information on the degradation of a component from beginning of life (BOL) to a defined end of life (EOL) as well as correlations to environmental and operational conditions so that the degradation of an object may be estimated over time by measuring the stressors as well as the component performance.

The General Path Model is a method to estimate the degradation path of a unit component for a single failure mode. GPM first begins with an assumption of a functional form that can be modeled using the historical data. This form is typically monotonic and may be considered by Equation 5.

$$x_{ij} = \eta(t_j, \phi, \theta_i) + \varepsilon_{ij}$$

Equation 5 - Functional Form for the General Path Model

Here, x_{ij} represents the degradation path of the i^{th} unit at a time t_j . The function η may be of a simple form (e.g. linear, exponential, or quadratic), but can be any valid function that provides a valid closed-form cdf from beginning of life to failure [Lu, et al., 1993]. The historical data is used to develop estimates of the model parameters ϕ and θ_i , where ϕ represents fixed effects that are constant among the population and θ_i represents variance from unit to unit. (The term ε_{ij} represents the customary measurement error.) Depending on the form of the function, the degradation paths of the historical data may be regressed onto Equation 5 to estimate the parameters, or some alternative estimation procedure (e.g. bootstrapping) may be employed. The estimation of the functional form of degradation is considered the first stage in GPM prognostics.

If the appropriate form for the GPM function is not known a priori, several competing models may be estimated using the historical data. A metric of comparison may then be observed to determine the appropriate model for the problem at hand; example comparisons may include estimation of point parameters such as the average degradation at beginning of life (BOL), or continuous parameters such as the shape of the path (e.g. linear vice exponential). Some models, such as polynomial models, may even be simplified based on the likelihood that individual parameters are significant. For example, a quadratic model may be simplified to a linear model if the coefficient of the second-order term is not significantly different from zero.

If multiple failure modes are available, the GP method may be adapted to include the known failure modes. One approach to multiple failure mode inclusion is to determine the minimum time to failure based on the time to failure of all known modes, as in Equation 6, where the various failure modes are encapsulated from mode 0 through mode S [Haghighi, et al., 2010].

$$\tau = \min (T^0, T^1, \dots, T^S)$$

Equation 6 - Minimum Time to Failure for Multiple Failure Modes

Extrapolation of the component's degradation path may then be performed with Bayesian updating of the degradation estimates based on both the environmental stressors and the component's performance. By using a degradation measure (sometimes termed a prognostic parameter [Coble, 2010]), the GPM can be used to extrapolate current degradation information of a component and can include censored data when building a failure model. GPM prognostication has been since expanded to include a large variety of techniques, including regression models [Upadhyaya, et al., 1994], neural networks [Upadhyaya, et al., 1994; Chinnam, 1999], and first-principles relationships [Luo, et al., 2003]. Uncertainty estimates of GPM techniques are relatively undeveloped in comparison to the GPM methodology itself, though significant work has been accomplished [Engel, et al., 2000; Byington, et al., 2004; Hines, et al., 2006].

The GPM approach is limited by the need to provide a closed-form function for the degradation path and the need to define a failure threshold [Garvey, et al., 2007]. If a convenient functional form is not readily apparent, the closed-form function may be approximated with more sophisticated methods, such as piecewise regression, fuzzy logic, or neural networking. The failure threshold may be defined by actual failure (hard limits), or by an arbitrary amount of degradation prior to failure (soft limits); in either case, these limits must be satisfactorily addressed for general path modeling.

2.4.2.2 General Path Model Prognostics – Case Studies in Literature

The method proposed in [Chinnam, 1999] applies the general path method to two case studies: the life expectancy of a drill bit that drills holes into stainless steel plates, and time to component failure due to fatigue crack growth. For the drill bit case, the force required to cut a hole into the steel plate was recorded and used to generate the historical failure path model by correlating drilling force to failure time (in number of operational cycles). Based on the historical data of eight drill bits employed to failure, a polynomial degradation model was fitted and used as the prior model. A failure threshold of a critical thrust force was defined based on the historical failures. The failure paths are shown in Figure 7 [Chinnam, 1999].




Figure 7 - Chinnam's Historical Drill Bit Failure Data [Chinnam, 1999]

Rather than using a Bayesian approach for GP model predictions of RUL for individual drill bits, Chinnam instead uses the historical data to develop the model type (e.g. polynomial vs. exponential vs. linear). As the drill bit of concern is used and force measurements are taken, the data is used to estimate the model parameters using a weighted least squares approach, given in Equation 7.

$$Q_w = \sum_{j=1}^m w_j \left(y_j - (\theta_0 + \theta_1 t_j + \theta_2 t_j^2 + \cdots + \theta_{p-1} t_j^{p-1}) \right)^2$$

Equation 7 - Weighted Least Squares Algorithm

Here, Q_w is minimized by adjusting the model parameters (θ). The weighting w_j may be adjusted to place greater emphasis on more recent observations by employing the exponential smoothing scheme given in Equation 8.

$$w_j = \beta(1 - \beta)^{m-j}, \quad 0 < \beta < 1$$

Equation 8 - Exponential Smoothing Scheme for Weighted Least Squares Coefficients

The coefficient β determines the relative weights given to the data points, with a recommended value between 0.001 and 0.3 [Johnson, et al., 1974]. Chinnam used a coefficient value of $\beta = 0.15$ for his model.

After the model was created, the data was bootstrapped to the model to derive confidence intervals for the model degradation path. An alpha level of $\alpha = 0.01$ was used, which resulted in wider confidence intervals than the customary $\alpha = 0.05$ used in most statistical applications. As more data is collected, the process of estimating model parameters and bootstrapping confidence intervals is repeated as frequently as desired. Despite the lack of a Bayesian framework to incorporate the historical data in model parameterization, the model works well in part because the measurement noise is well below the measurement values, so the initial data of the drill bit correlates well to its failure path.

Chinnam then applies the same method to fatigue crack growth in plates [Chinnam, 1999] with data taken directly from Lu and Meeker [Lu, et al., 1993] and may be seen graphically in Figure 8. (The historical data was generated by Lu and Meeker by visual inspection of fatigue crack growth plots in [Bogdanoff, et al., 1985].) The historical data of Lu and Meeker is used to verify a third-order polynomial model for estimation of crack growth. The model parameters are then estimated via bootstrapping for one plate using partial degradation data. The progression of the RUL prediction is tracked by estimating RUL at various times in the life of the plate, along with confidence intervals. This approach again provides usable estimates of time to failure because the noise in the data is low and the growth of the fatigue crack is very strongly correlated to the polynomial model. In both cases, however, the lack of a prior model leaves the predictions susceptible to variance in early predictions.

Another example of the General Path method may be seen in Figure 9 [Yan, et al., 2004]. Here, the probability failure was modeled using logistic regression to determine the likelihood that the door would fail to open or shut correctly. The "Windows" represent the time until an estimated 95% probability of failure, with the estimate for Window 2 being taken at some time after the estimate of Window 1. The updated prediction can be seen from the Bayesian updating of the parameter estimates based on the information collected between the two predictions. At 200 cycles, the routine maintenance was performed, at which time the general path model for the door would need to be reset to account for the change in component condition.

The Bayesian updating procedure allowed Yan to use the historical data in early failure predictions to stabilize the variance of the predictions. As further degradation data was collected, the individual data took increasing precedence over the historical data, allowing the GP model to more accurately reflect the degradation within the individual door. Since the signal-to-noise ratio was greater in the door data than in Chinnam's case studies, the Bayesian approach protected Yan from instability in the early prognostic estimates due to the variance of the model coefficient estimates.




Figure 8 - Fatigue Crack Propagation in a Metal Plate [Lu, et al., 1993]




Figure 9 - GPM Failure Prediction of an Elevator Door [Yan, et al., 2004]

2.4.2.3 Particle Filter Prognostics

Particle Filter (PF) methods are relatively new techniques in prognostics [Orchard, et al., 2005; Cadini, Particle Filtering for Diagnosis, Prognosis and On Condition Maintenance, 2009] and are an adaptation of particle filtering methods for statistical tracking of an object's position and motion [Ristic, et al., 2004; Marseguerra, et al., 2009]. In particle filtering, a series of Monte Carlo (MC) simulations of an object's state (position or vector motion traditionally, but may be defined as a degradation parameter in failure analysis) are performed based on the last known state of the object and a distribution of the expected probabilities for changes in the object's state. These Monte Carlo simulations (particles) are then compared to measurements of either the actual object's state or to metrics that correlate to the object's state, and then the particles are weighted based on their probability of representing the correct state. To ensure that a sufficient number of particles maintain enough weight to create a statistical sample, particles with weights below some predetermined threshold are typically eliminated, and then the particles are redistributed based on the distribution of the remaining particles and their weights [Doucet, et al., Sequential Monte Carlo Methods in Practice, 2001]. PF methods are of particular interest due to their ability to model non-linear processes [Cadini, et al., Monte Carlo-Based Filtering for Fatigue Crack Growth Estimation, 2009] as well as the elimination of the need to assume a particular probability density function (pdf) distribution of the process itself [Gordon, et al., 1993]. PF also provides a built-in estimation of uncertainty by evaluation of the particle distributions [Orchard, et al., 2008], with ongoing research in reduction of uncertainty over long-term predictions [Orchard, et al., Outer Feedback Correction Loops in Particle Filtering-based Prognostic Algorithms: Statistical Performance Comparison, 2009].

The particle filter is an evolution of the linear filtering method proposed by Kalman [Kalman R. E., 1960], commonly known as the Kalman filter. Kalman's approach assumed a linear state-space relationship with a target vector $\mathbf{x}$ and measurement vector $\mathbf{z}$. Under these assumptions the relationship between the measurements at time $k-1$ and the target location at time k could be expressed by Equation 9. The relationship between the target location and the measurements at time k could be expressed by Equation 10 [Arulampalam, et al., 2002]. In these equations, F and H are state matrices relating $\mathbf{x}_{k-1}$ to $\mathbf{x}_k$ and $\mathbf{x}_k$ to $\mathbf{z}_k$, and $\mathbf{v}_{k-1}$ and $\mathbf{n}_k$ are sampled from Gaussian distributions.

$$\mathbf{x}_k = F_k \mathbf{x}_{k-1} + \mathbf{v}_{k-1}$$

Equation 9 - Kalman Filter Target Location

$$\mathbf{z}_k = H_k \mathbf{x}_k + \mathbf{n}_k$$

Equation 10 - Kalman Filter Measurement Relationship

With Kalman filtering, target $\mathbf{x}_k$ is estimated based on the knowledge of $\mathbf{x}_{k-1}$, and the prediction is then compared to the measurement vector $\mathbf{z}_k$. If the pdf $p(\mathbf{x}_{k-1} | \mathbf{z}_{1:k-1})$ is available at time $k-1$, then the pdf for $\mathbf{x}_k$ is estimated using the Chapman-Kolmogorov equation, given in Equation 11. The resulting solution provides the estimate for the prior (also known as the normalizing constant) in Bayes Rule given in Equation 12.

$$p(\mathbf{x}_k | \mathbf{z}_{0:k-1}) = \int p(\mathbf{x}_k | \mathbf{x}_{k-1}) p(\mathbf{x}_{k-1} | \mathbf{z}_{1:k-1}) d\mathbf{x}_{k-1}$$

Equation 11 - Chapman-Kolmogorov Equation

$$p(\mathbf{x}_k | \mathbf{z}_{0:k}) = \frac{p(\mathbf{z}_k | \mathbf{x}_k) * p(\mathbf{x}_k | \mathbf{z}_{0:k-1})}{p(\mathbf{z}_k | \mathbf{z}_{0:k-1})}$$

Equation 12 - Bayes Rule

At its heart, the Kalman filter is a method to estimate the prior in Equation 12, where the term $p(\mathbf{x}_k | \mathbf{x}_{k-1})$ in Equation 11 is estimated by direct application of Equation 9 and Equation 10. The Gaussian assumption was enforced by Kalman to ensure a symmetric, convex distribution of the MC routine so that the prediction estimate of the target is simply the mean of the process, as in Equation 13 [Kalman R. E., 1960; Arulampalam, et al., 2002].

$$\mathbf{x}_{k|k-1}^* = E[\mathbf{x}_k | \mathbf{z}_{0:k-1}]$$

Equation 13 - Target Vector Prediction under the Gaussian Assumption

The assumption of a linear relationship between the measurement and target vectors allowed for simplifications to the recursive solution routine [Kalman, et al., 1961]. With linearity, the F and H

relationships in Equation 9 and Equation 10 are directly invertible, allowing for rapid solution of the Bayesian process. With the limitations in computational power during the time of Kalman's work, this assumption was somewhat mandatory for the filtering process to be timely and effective.

The particle filter method is an evolutionary extension of the Kalman filter [Kalman R. E., 1960] where the two fundamental restrictions of the Kalman filter are eliminated. Rather than assume a Gaussian distribution to the prior pdf, the distribution is estimated using a Monte Carlo process [Ristic, et al., 2004]. If the target vector $\mathbf{x}_{k-1}$ and measurement vector $\mathbf{z}_{k-1}$ are known, then the prediction of the future target vector $\mathbf{x}_k$ may be estimated using a MC process. The MC particles are then compared to the measurement vector $\mathbf{z}_k$ to determine the likelihood that each particle could represent the true target vector $\mathbf{x}_k$ given $\mathbf{z}_k$. The particles are then weighted according to their likelihoods in a process termed Sequential Importance Sampling (SIS) [Doucet, et al., 2001; Doucet, et al., On sequential Monte Carlo sampling methods for Bayesian filtering, 2000]. SIS has also been referred to as bootstrap filtering [Helferty, et al., 1993], condensation [MacCormick, et al., 2000], and survival of the fittest [Kanazawa, et al., 1995].

The SIS step replaces the Kalman filter's direct estimation of $p(\mathbf{z}_k | \mathbf{z}_{0:k-1})$ in Equation 12 with a stochastic estimation that is free of both the linear and the Gaussian assumptions. Such approaches were proposed as early as 1954 [Hammersley, et al., 1954] though limitations in computational power and the problem of particle degeneracy suppressed the use of SIS until the 1990s [Ristic, et al., 2004] when Sequential Importance Resampling (SIR) was introduced [Helferty, et al., 1993].

With each successive application of SIS, particles with low likelihood of representing the true target vector would continually have their weights reduced to the lower bound of zero weight, while the particle with the highest likelihood of representing the target vector would eventually receive increased weight to the normalized upper bound of one. This degeneracy of the particles guaranteed that the MC distributions would eventually approach a zero-variance limit. To overcome this limitation, a resampling of the particles was introduced. After the SIS step, a distribution of the prior measurement vector could be obtained from the weighted particles. A new set of equally-weighted particles could then be producing by randomly sampling from the estimated prior distribution; with a sufficient number of particles, the distribution of the newly-distributed particles converges to the distribution of the old, weighted particles. At this point, the previous particles could then be discarded in favor of the new particles.

The resampling process can be performed at a fixed number of steps (e.g. after every SIS), but is usually performed when degeneracy is deemed significant. The most common degeneracy metric, N_{eff} , is given by Equation 14 [Ristic, et al., 2004]. If the weights of the particles are well-distributed, then N_{eff} will retain a low value. If the weights degenerate, then one particle's weight will grow to the upper limit of one and N_{eff} will likewise approach a value of one. By comparing N_{eff} to some threshold limit N_{eff} , a decision can be made whether resampling is necessary. Therefore, with SIR, degeneracy may be avoided and the particle filter process could be repeated indefinitely.

$$N_{eff} = \frac{1}{\sum_{i=1}^N (w_k^i)^2}$$

Equation 14 - Degeneracy Algorithm for Sequential Importance Resampling

To apply particle filter methods to prognostic applications, some assumptions must be made of the target vector and its changes over time. For tracking (i.e. non-prognostic) applications, the MC particles need not follow an optimal path from the current time step to the next, so long as their distribution encompasses that of the range of possibilities of the actual target vector, thereby ensuring that the target vector pdf is entirely within the sampling. The SIS procedure provides an automatic correction to the particle distribution via the weight changes. (While the distribution need not be ideal, distributions of the particles that more closely represent the true distribution of measured data will naturally perform more efficiently.) For prognostics, however, the particles will be used to sample a distribution many time steps away from the last known measurement to derive conclusions about the target vector well into the future. Without some control over the particle trajectories, the resulting distribution of particles may provide predictions with unreasonably high variances or with inaccurate understanding of the future path. Such control over the particle paths may be implemented with a priori knowledge of the degradation path, or by shape estimation of the path either by weight sampling or particle resampling [Orchard, et al., A particle-filtering approach for on-line fault diagnosis and failure prognosis, 2009]. Such applications have been developed in independent efforts for fatigue crack growth prognostics [Patrick, et al., 2007; Cadini, et al., Monte Carlo-Based Filtering for Fatigue Crack Growth Estimation, 2009] as well as for battery health management [Goebel, et al., 2008].

2.4.2.4 Particle Filter Prognostics – Case Studies in Literature

A notable example of particle filtering as a prognostics tool is given by Orchard [Orchard, et al., 2008].

In this approach, a particle filter algorithm is used as a diagnostic tool to initially detect the presence of a fault. Once the fault is detected, the current state-space pdf estimated by the PF model is used as the initial condition for the prediction of RUL. Continued monitoring of the system provides Bayesian updating to the RUL estimate by updating the particle densities. A case study is presented where the growth of a seeded crack in a planetary gear is detected and the RUL of the gear is predicted.

Orchard compared two approaches to updating the long-term predictions. The first approach, termed a "p-step ahead" approach, effectively freezes the particle weights and distribution at the last measurement, disallowing the redistribution of particles via SIR, and propagates the particles to the predefined end of life condition. The RUL estimate is simply the expectation generated from the continued particle paths, as in Equation 15. Here, the particle position estimates are given by $\hat{x}$, and the particle weights are given by ω . This approach is nothing more than simply carrying the weighted particles to their natural EOL conclusion.

$$\hat{x}_{t+p}^{(i)} = E \left[f_{t+p} \left(\tilde{x}_{t+p-1}^{(i)}, \omega_{t+p} \right) \right]; \hat{x}_t^{(i)} = \tilde{x}_t^{(i)}$$

Equation 15 - RUL Prediction from Orchard's First PF Prognostic Method

The second approach by Orchard attempts to refine the prediction by resampling the particles and equalizing their weights rather than simply leaving the particles with their current weights. This is a final re-application of the Sequential Importance Resampling step common to PF methods. To guard against possible irregularities from the resampling step (e.g. "holes" in the particle distribution where the random sampling process failed to provide new particles), the resampling process is regularized with a quadratic kernel based on an available regularization algorithm [Musso, et al., 2001]. This method updates the particle position with the quadratic kernel applied to the uncertainty, as in Equation 16. The random variable ε is randomly sampled from the kernel. $\hat{D}$ is computed from the covariance matrix of the former and current particle positions, and h is a normalizing constant based on the number of particles.

$$\hat{x}_{t+k}^{(i)*} = \hat{x}_{t+k}^{(i)} + h_{t+k}^{opt} \hat{D}_{t+k} \varepsilon^i$$

Equation 16 - Regularization of the PF Prediction with the Quadratic Kernel

When applied to the case study, the regularization appeared to reduce the variance of the predictions without noticeably affecting their accuracy.

Another example of PF methods for prognostics is given by Matthew Daigle [Daigle, et al., 2009]. Here, Daigle utilized a *fixed-lag* filter, where the state estimates at time t are not calculated until after some future measurements $t+k$ have been observed. This delay in state estimation is in itself undesirable, but if the delay is small in comparison to the length of time being predicted (e.g. if a 10-time step delay is used but predictions are made for 100 or more time steps in advance), then the detrimental effect is minimized. The benefit to fixed-lag is that the state estimation at time t can be improved by using measurements before, during, and after time t . In short, using a fixed-lag filter reduces the measurement error inherent in the state observation. Daigle presented a case study of a pneumatic valve where the wear and eventual failure of the valve is predicted by particle filtering. The fixed-lag approach provided some benefit dependent on the actual lag of the filter. In this case study, the maximum lag ($L=3$ steps) provided the best benefit, but at cost of the greatest time delay for state estimation. If the RUL prediction is expected to be many steps greater than the lag, then the fixed-lag approach to variance reduction appears to have merit.

Other prognostic applications of particle filters include Saha [Saha, et al., 2009] and Abbas [Abbas, et al., 2007]. The first article applies particle filters to the problem of EOL estimates for lithium-ion batteries, and the latter article applies PF to prediction of the corrosion of a battery grid. In both cases, the PF approach provided useful predictions to time to failure. Both case studies, however, illustrated the common thread to all PF methodologies – the need for a reasonable understanding of the failure mechanism and the degradation path. This may be provided by a first-principles approach or by collecting historical data, but in all cases the predictions (RUL or EOL) required a priori knowledge of the failure path. In the case of prognosing a human decision, such as the production of undeclared LEU or HEU, such a priori knowledge is not readily available and the shape of the failure path must be defined as an assumption rather than a common trend. Currently, the literature is devoid of PF or GP prognostic models where the failure path itself is unknown.

2.4.3 Prognostics within Verification and Safeguards

As prognostics are not yet formally applied to verification of activity within a safeguards context, existing literature is deficient in the formal discussion of prognostic types for verification and safeguards purposes. For this research, it is assumed that Type I prognostics are not effective models for diversion or undeclared enrichment because both activities represent a decision made by the operator (or possibly by a malicious third party in the case of diversion) and therefore represent a change in operating condition. If an "average" diversion or undeclared enrichment condition is assumed, then *all* prognostics of a plant would either predict the eventual production of an undeclared SQ or would never predict undeclared SQ production at all. In either case, no new safeguards information is gained.

Likewise, Type II prognostics are of limited use; though environmental changes may be accounted for in these models, they do not account for system variations such as scale biases or variations in the enrichment process itself. Such system changes can have a profound influence on MUF calculations, which are essential in estimating diversion and undeclared production. This research will therefore focus on Type III prognostics to account for both environmental factors (e.g. cold traps, MUF calculations, and plant operation) as well as system factors (e.g. scale biases). Further discussion of the prognostic modeling in this research is discussed under Section 3.

2.5 Model Selection Criteria

One of the cornerstone fields of statistical research is the selection of an appropriate model for fitting to data. Statistical models traditionally rely on measures of likelihood that estimate whether the suggested model is sufficiently correlated to the underlying data to provide a meaningful relationship between the measurement and target vectors. A variety of statistical tools, such as t-tests, chi-square test, likelihood ratio tests, maximum likelihood estimators (MLE), mean square error (MSE) or sum of square error (SSE) measurements, and other goodness-of-fit tests have been developed to provide measures for model selection. A thorough review of such techniques is a research topic to itself, and this review will merely acknowledge the vast array of techniques available for model selection criteria.

Traditional statistical methods, such as MLE and MSE tests, have two defined limitations: their reliance on arbitrary thresholds of "goodness" such as alpha values, and the assumption that the nature of variance of the data is known (e.g. normal, Cauchy, or lognormal). Attempts to eliminate these limitations from the model selection process have led to the Information Criteria (IC) class of statistical measures [Bozdogan, 1987]. The seminal work in the field of IC is Akaike's entropic information criterion

[Akaike, 1973], known as Akaike's Information Criteria (AIC). In summary, Akaike extended the maximum likelihood estimator test of model goodness of fit by providing a penalty term for the number of variables. $\hat{\theta}$ represents the parameter vector. The penalty term k (the number of parameters in the model) in Equation 17 accounts for the lack of parsimony of overfitted models by the addition of new explanatory variables to the model.

$$AIC = -2 \log L(\hat{\theta}) + 2k$$

Equation 17 - AIC Measure for Model Selection

The penalization of model parsimony provided a measure of protection against overfitting. Compared to traditional statistical measures of model validity, such as alpha tests, the AIC score provided a direct comparison between models rather than a comparison between a model and a predefined arbitrary standard. The model with the lowest AIC score was regarded as the "best" fit in terms of maximal likelihood and in terms of model complexity [Akaike, 1974].

Since Akaike introduced his penalization of model complexity as an integrated part of model selection, several penalty terms have been suggested. A bias-corrected form of the penalty term has been suggested [Bozdogan, 2000] where the number of parameters k is replaced by a bias-corrected penalty in Equation 18.

$$BCIC = -2 \log L(\hat{\theta}) + 2nb$$

Equation 18 - Bias-Corrected Information Complexity

Alternate formulations of the Information Criterion penalty concept replace the number of terms k in Equation 17 with a generalized expression for model complexity. Takeuchi [Takeuchi, 1976] proposed a complexity penalty based on the trace of the inverse Fisher Information Matrix $\mathcal{F}^{-1}$ and the outer product form $\mathcal{R}$, shown in Equation 19.

$$TIC = -2 \log L(\hat{\theta}) + 2tr(\hat{F}^{-1}\hat{R})$$

Equation 19 - Takeuchi's Information Criterion

The addition of Fisher information to the IC class of model selection criteria is fundamental to the Information Complexity (ICOMP) measures for model selection [Bozdogan, 1988], in which the IC concept of model selection has been extended to penalize the profusion of complexity in a model in addition to the penalties for lack of fit and lack of parsimony. The general form of ICOMP utilizing the inverse Fisher information matrix (IFIM) is given in Equation 20. Theoretical development and justification for the Fisher information matrix as a measure of model complexity may be found in [Kullback, et al., 1951; Bozdogan, 1988; Li, et al., 1996; Stigler, 1999].

$$ICOMP(IFIM) = -2 \log L(\hat{\theta}) + 2C_1(\hat{F}^{-1}(\hat{\theta}))$$

Equation 20 - Information Complexity and Inverse Fisher Information, General Form

The C_1 function is a measure of the maximal covariance complexity and accounts for lack of parsimony and profusion of complexity. As with Takeushi's information criterion, C_1 incorporates the trace of the contained matrix as a measure of parsimony; C_1 then utilizes the determinant of the contained matrix to add a penalty for complexity. For ICOMP(IFIM), the C_1 estimate of covariance complexity is shown in Equation 21. Here, s is the rank of $\mathcal{F}^{-1}$. Further discussion of C_1 may be found in [Bozdogan, 1990].

$$C_1(\hat{F}^{-1}) = \frac{s}{2} \log \left[\frac{tr \hat{F}^{-1}}{s} \right] - \frac{1}{2} \log |\hat{F}^{-1}|$$

Equation 21 - C1 Covariance Complexity for ICOMP(IFIM)

Extending the ICOMP theoretical framework for model selection, the assumption of proper model specification (i.e., that the fitted model is correct) may also be eliminated as in Equation 22, where the ICOMP score includes a penalty for misspecification [Bozdogan, et al., 2009]. In this form, the complexity measure penalizes against lack of parsimony and profusion of complexity using the robust covariance matrix $\mathcal{F}^{-1} \mathcal{R} \mathcal{F}^{-1}$, which is valid if the model is correctly or incorrectly specified. If the fitted model is indeed correct, then $\mathcal{F} = \mathcal{R}$, and the covariance matrix reduces to $\mathcal{F}^{-1}$, as in Equation 20.

$$ICOMP(IFIM)_{Misspec} = -2 \log L(\hat{\theta}) + 2C_1(\hat{F}^{-1} \hat{R} \hat{F}^{-1})$$

Equation 22 - ICOMP(IFIM) for Protection against Misspecification

Expanding the C_1 covariance complexity for $\text{ICOMP(IFIM)}_{\text{Misspec}}$ provides the result in Equation 23. If the model is properly specified, Equation 23 reduces to Equation 21 and the information criterion in Equation 22 simplifies to the standard ICOMP(IFIM) measure in Equation 20. This represents the optimal case, where the model is properly specified to the data and the information measure is minimally penalized. Should the model be improper for the data, the covariance complexity would depart from the optimal case and penalize the score of the model accordingly.

$$C_1(\hat{F}^{-1}) = \frac{s}{2} \log \left[\frac{\text{tr}(\hat{F}^{-1} \hat{R} \hat{F}^{-1})}{s} \right] - \frac{1}{2} \log |\hat{F}^{-1} \hat{R} \hat{F}^{-1}|$$

Equation 23 - C1 Covariance Complexity for $\text{ICOMP(IFIM)}_{\text{Misspec}}$

The penalization of a misspecified model makes the $\text{ICOMP(IFIM)}_{\text{Misspec}}$ criteria particularly well suited for comparing the goodness of fit of multiple models. As with any IC class of model selection, the model with the lowest score is generally considered the "best" model in terms of fit, parsimony, and complexity, and unitary differences in model scores are generally considered negligible. The misspecified form of ICOMP may then be used to gauge whether a model fitted to data is sufficiently valid for use as a model of the data by comparing the score of the hypothesized model to the score of a constant term or random noise.

This chapter presented a survey of the existing literature relevant to this research. The role of the IAEA as the monitoring agency was considered. Current and proposed next-generation inspection techniques were reviewed along with the driving factors behind the IAEA's desire for remote monitoring of enrichment facilities. Existing prognostics techniques were reviewed, with a particular focus on general path and particle filter models. Finally, model selection as a means for discriminating system behavior was discussed. In the next chapter, the methods of this research will be introduced. An overview of the mock feed and withdrawal facility at Oak Ridge will be presented. Prognostic techniques for predicting the time to diversion of significant quantities and their uncertainty estimates will be developed. The *PlotEvents* software tool will be discussed as a means for automated analysis of the process monitoring data from the mock feed and withdrawal facility.

3 METHODOLOGY

This research required the development of a prognostic method for safeguards and the validation of the prognostic method through case studies that tested the method under a variety of operating conditions. The mock F&W facility provided a platform to generate process data for prognostic model selection, development, and validation. Load cell data from the feed, tails, and product stations were recorded at 1 second intervals in a central database. The data was then analyzed for operational features such as tank loading/unloading or cold trap operation, and the MUF estimate was calculated as a function of time. The operational features and the MUF estimate were then analyzed to search for diversion or other undeclared activity.

Prior to analyzing the load cell data for MUF, the data is smoothed to eliminate the effects of loading/unloading and other perturbations like tube attachment, and then downsampled to produce a smaller, more manageable data set. Downsampling is typically performed by selecting every fifth data point so the analysis is performed on one data point for every five seconds of operation, or 0.2 Hz. The smoothing algorithms were developed primarily by James Henkel [Henkel, 2010]. With the smoothed data, the inventory differences could be calculated by summing the changes in load cell weights and integrating over time. This produced the MUF estimates from the load cell data.

After developing the MUF, monitoring and prognostic analytics are applied to the data to diagnose the likelihood of diversion, its severity, and a prediction of the time to SQ production. In this section, the theory behind the prognostic models (i.e. GP and PF models) will be developed. Because the theory has been developed in process monitoring applications such as health and condition monitoring, the terminology used to develop the theory here will be consistent with this canon of literature. As the methods are applied to the safeguards application of this research, the analogy between the process monitoring methods and the safeguards verification application will be explained and justified.

3.1 Mock Feed and Withdrawal Facility

A preliminary discussion of the mock F&W facility may be found in Appendix A.

For the mock F&W facility to serve as a sufficient analog of a GCEP, the load cell data produced must bear adequate resemblance to GCEP load cells. To this end, the mock F&W facility required some modifications to produce consistent data that had the same features as GCEPs. The capacity to bleed water from the feed tubing (between the feed pumps and the surge tanks) allows for cold trap

emulation. The ability to mask diversion was added by installing an independent tube with a funnel on one end and a throttle valve on the other. The masking assembly can easily be fixed so that the water in the funnel pours into a product or tail tank, and the flow rate can be regulated with the throttle valve to match the flow rate of the diversion.

Procedural changes were added to the mock F&W facility operation. The product and tail lines used to contain a significant amount of air as a consequence of valve operation for the individual product and tail lines. The air entrainment was found to cause inconsistent product and tail flow to a degree that was not in keeping with the GCEP analogy. As a result, valve operation procedures were changed to ensure that the product and tail tubes remained watertight throughout facility operation. The change in procedures have been observed to produce more consistent water flow, eliminating the need to adjust the product and tail throttle valves, particularly during facility startup and shutdown. Additionally, time lags have been added in between tank loading/unloading and tube connection/disconnection. The time lags generate data that is more similar in structure to GCEP operation and allow for the development of automated techniques to observe and monitor the handling of tanks on load cells.

3.2 Prognostics and Uncertainty Estimates

This research develops and characterizes methods that provide a prediction of the time at which a SQ of LEU is diverted from the process stream. This quantity is not directly measureable by the IAEA, as a facility that is illicitly diverting LEU is hardly going to allow such a measurement to be recorded by IAEA sensors. For this reason, the diversion of material is considered the target state – the unmeasured quantity of interest that is estimated by measurement of correlated parameters. The calculation of MUF using the load cell measurements provides a correlated estimate of the diverted material, but MUF is not only attributable to diversion or undeclared production. Normal process influences like cold trap operation and holdup also affect MUF calculations. A better measure of unaccounted material attributed to diversion will be therefore sought.

With a target state in mind, general path and particle filter models were developed to estimate the amount of diverted material as well as to predict the time until such a diversion would be expected to produce a significant quantity. Development of two competing models allowed for the model performances to be compared so that a preferential model can be determined. The models were designed to not only provide estimates of missing material and time to SQ production, but also to

provide estimates of the uncertainty of their predictions and, if possible, the likelihood that missing material may be attributable to diversion rather than measurement errors, such as scale biases.

3.2.1 Identification of Facility Operation Features

During normal operation of the mock F&W withdrawal facility, certain data features are always present: loading/unloading of cylinders, attachment of process tubes, feed and withdrawal activity, and cold trap operation. These features are easily detectable by visual inspection of the raw data from the load cells, but this data cannot be transmitted off-site at a GCEP because of the risk that the load cell data may reveal proprietary information about GCEP operation. The identification of these normal features of operation must therefore be performed by automated algorithms so that the normal operation of the plant may be verified.

Abnormal data features may occur, such as spikes in the load cell weight readings due to a person stepping on the scale. These abnormalities do not necessarily indicate illicit activity and should not be instantly identified as "red flags" that trigger inspections. The analysis of such data features should therefore place emphasis on identification of the necessary normal operational characteristics and reduce the priority of abnormalities unless a particular abnormality can be correlated to illicit activity. The capacity to identify data features and provide a summary analysis of their significance to facility operation has yet to be developed for the mock F&W facility and is a continuing part of this research.

3.2.2 Estimation of Diversion (TMUF)

In process monitoring, the estimated state of the system is typically defined as either the measure of degradation or the health of the system (or individual components). The measurement vector is a selection of system performance metrics (e.g. flow rates, temperatures, or voltages) that are correlated to the degradation or health state. For this research, the state will be defined as the MUF of the enrichment process. It should be noted at this point, however, that the MUF is generally inclusive of all missing material, including holdup and mass lost due to normal processes, such as holdup. The real quantity of interest is not the total MUF, but rather the unaccounted material after considering normal deviations from the mass balance. One may consider a refined MUF estimate, termed here as the TMUF, or T r u e M a t e r i a l U n a c c o u n t e d F o r m e a s u r e m e n t, which may be considered similar to Equation 24, where L_{CT} is the material loss from the cold trap, H is holdup, and $L_{O/P}$ is any other identified process loss.

$$TMUF = MUF - L_{CT} - H - L_{O/P}$$

Equation 24 - Conceptual Definition of TMUF

The TMUF is the optimal state estimation for prediction of time to SQ production rather than the MUF. If this quantity can be estimated through modeling of the cold trap, holdup, and other process factors (such as tank switchover), the TMUF will be used as the target state. Otherwise, the MUF will be estimated and the effects of normal processes will be accounted for afterward.

At a minimum, the measurement vector will be defined as the load cell measurements (which are used to calculate the MUF) along with any other measurements that can provide information on the normal process deviations. Examples of other possible measurements include the loading/unloading of cylinders, feed pump power, surge tank control valve position, and cold trap operation. In normal process monitoring, these variables may be used based on optimization of the monitoring process; in safeguards, such variables are often constrained by negotiation between the IAEA and the GCEP operator. For this reason, a minimal measurement set is sought that will include, at a minimum, the load cell measurements. It is anticipated that a reliable estimate of TMUF may be provided by load cell measurements alone; if not then justification for other measurements should have a minimum likelihood of revealing proprietary operational information so that their inclusion in the prognostic model may be acceptable to facility operators.

3.2.3 Automated Identification of Undeclared Activity

In this context, undeclared activity refers to the act of feeding material into the cascade from undeclared cylinders in order to generate undeclared LEU or HEU. With the traditional inspection method, there is a risk that undeclared activity might not be detected if a GCEP can complete the entire undeclared feed and withdrawal evolution in-between inspections. (Note: due to the comprehensiveness of the inspections, undeclared activity is not quite as simple to conceal as the preceding sentence indicates, but the concept is sufficiently explained in that sentence without attempting to write a "how-to" for undeclared activity under current inspection procedures.) With load cell data, undeclared activity can be detected by identification of feed and withdrawal that does not coincide with declared activity. If, for example, the daily mailbox declarations indicate that a cascade is not being used but the load cell data indicates activity, an inspection might be triggered to attempt to resolve the discrepancy. For the mock F&W facility, proper identification of facility usage can enable the

automated generation of summary reports that allow the IAEA to compare the load cell data to the daily mailbox declarations without transmitting the load cell data itself. This feature would be of particular value in detecting undeclared activity and has to be developed for the mock F&W facility. Reporting load cell activity will be performed by creating software algorithms that identify the features of facility operation and then automatically write reports to describe the operation.

3.2.3.1 Information Complexity and Model Selection

It is postulated that an undeclared removal of material from the mock F&W facility would represent a change in the operating regime of the process. Specifically, the MUF calculations would have different temporal relationship than when the facility is operated without any undeclared production (i.e. diversion). This difference in regime would change the optimal statistical model for the data by affecting cumulative inventory difference of the facility during operation. Additionally, the ability to fit a model to the data does not necessarily indicate that the model is an appropriate description of the information contained within the data.

To determine whether a trend that indicates potential diversion is a valid model for the data, the information complexity metric $ICOMP(IFIM)_{Misspec}$ is considered. From Equation 22, the proposed model for the data may be compared against the data using both the maximum log likelihood criterion and the maximal information complexity of the misspecified covariance matrix. This relationship imposes penalties for lack of fit (via maximum log likelihood) as well as lack of parsimony, complexity, and misspecification (via the misspecified covariance matrix). The fitted model for the data can then be tested against an alternative hypothesis by recording the score of the misspecified ICOMP test and comparing to the score of the alternative. In this case, two alternative hypotheses can be tested. The first alternative is that the data indicates a constant MUF. If this is true, then a constant model would provide the best fit, greater parsimony, and lower complexity and would then receive a lower ICOMP score than the fitted model. The second model is that of random data. If the MUF data in question is not well suited for a constant fit or for the fitted model, then neither would perform as well as a model of random data, and the random ICOMP model, with the lowest parsimony, would score equivalently or lower than the other models. Where the constant model would best represent a non-diversion regime with constant MUF, the random model would indicate a lack of coherence between the MUF data and any attempt to fit a model. Such a region would be seen as a transitional region where no information is usable for model selection.

The ICOMP score, therefore, may be used as a validation of the fitted linear model. If a model renders a prediction of SQ production yet the ICOMP score does not justify the use of the model over a constant or random term model, then the prediction cannot be considered reliable. Likewise, if the random model outscores the fitted or the constant models, then the data would suggest a lack of meaningful information within the inventory difference calculations. In this case, the data may suggest alternative remedy. For example, if model fitting procedures were applied to GCEP load cell data and the models suddenly lost relevance with respect to a random model, then the change in load cell relationships might indicate anomalous behavior that required further inspection.

3.2.4 General Path Method

The General Path Model was traditionally (and still remains) the most common method for Type III prognostics in process monitoring and condition-based maintenance applications. First proposed by Lu and Meeker [Lu, et al., 1993], the basic approach to GPM prognostics is to first build a prior model of the degradation path based on unit-by-unit variance and the influence of environmental factors. The historical data used to build the prior model may also include censored data, as information on the variance of unit performance is as important to the GPM as the environmental stressors. The historical data is then used to build a functional model of the degradation path of the component. For each unit in the historical data, the value of the degradation parameter may be defined as a variable x_{ij} , which is read as the degradation of the i^{th} component at time j . The degradation parameter may then be related to the time of life, the environment, and the individual unit characteristics using Equation 25.

$$x_{ij} = \eta(t_j, \phi, \theta_i) + \varepsilon_{ij}$$

Equation 25 - Unit Degradation Relationship for the General Path Model

In Equation 25, ϕ is a vector of fixed effects that are characteristic of the entire population, θ_i is a vector of the individual (usually random) effects that are unique to each unit, and ε_{ij} is included to account for measurement error. The term η is a function that describes the relationship between the causal (i.e. environmental and component) data and the degradation parameter. This function may be determined by first-principles models or by correlating the degradation parameter to the measured data. Unlike failure-time prognostic models, the prior function in Equation 25 represents an estimate of the *condition* of a component based on the process as opposed to the likelihood of failure. For this reason, even

censored historical data may be included in the historical model as the censored data contains pertinent information of unit condition through a process. (In general, a few units should be operated to failure when building the historical model so that the model contains an estimate of failure itself.)

For each degradation case in the historical library, the degradation history is regressed to the chosen degradation function described in Equation 25 and the parameters t_j , ϕ , and θ are recorded. Mean parameters for the degradation function are then calculated using the parameters for the individual regression cases to develop a mean function. This is the general path prior model, as given in Equation 26.

$$\bar{x}_j = \eta(t_j, \bar{\phi}, \bar{\theta}) + \bar{\varepsilon}_j$$

Equation 26 - Fitted General Path Model

The fitted model in Equation 26 provides a relationship for the mean degradation of the component over time, with information of the degradation path available at every step j . The variance of each parameter is also calculated at this time. If no further effort were made into development of this GP model, the model would be applied to a new degradation case by solving for the initial time of life of the component based on the initial degradation value, then predicting the degradation path of the component using the mean model in Equation 26. Uncertainty of the model may then be calculated using the uncertainties of the model parameters. A refined approach to extrapolation of the degradation path of an individual component is provided by Uphadyaya et al. in [Upadhyaya, et al., 1994].

When developing the historical model, runs should be included which represent the expected range of operation for the component. If a normal range of operation is not adequately represented in the historical data, extrapolation of the model may produce errant predictions of degradation, particularly if the degradation / process relationships contain nonlinearities.

If the model contains sufficient degradation data for and is sufficiently representative of the expected range of process data for future components in question, then a failure limit must be established. In the case of hard failure, this may be the level of degradation at which a component is expected to break, and typically includes both an expectation and variance for the degradation level at failure. If a soft

failure criterion is employed, then the relationship between the degradation level and the soft failure limit must be established. In either case, this critical degradation limit may be termed D .

With the degradation model, the process variables may be measured for a non-historical component to provide an estimate of the degradation imputed upon the component from its use. An initial degradation level must first be established for the component; estimation of the initial degradation may be as simple as assuming an average initial degradation for every component (e.g., zero degradation) or establishing a method to provide a unit-specific estimate of degradation. Such estimates will not be discussed here; in the problem of interest to this research, the 'degradation', or production a SQ of material will generally be assumed to start with zero SQ produced. As the unit is subjected to the process, the process variables of the GPM are measured and used to calculate the expected level of degradation of the component.

At this point, the GP model is only providing a measure of the current level of degradation of the component (i.e. monitoring the health of the component). If the model is to provide a prediction of RUL or TTF, a systematic estimation of the relationship between the unit's current state of degradation and its degradation level at failure must be defined. The simplest approach is to calculate the historical average TTF between the current and failure degradation levels and declare such time as the estimation of TTF or RUL. Unfortunately, this Type I approach to prognostics ignores the particular environmental stressors and characteristic response of the component itself. For true condition-based Type III prognostics, an estimate should account for these individual effects. Bayesian updating of the failure path is the most common approach to providing this Type III data.

The model may be further refined using Bayesian updating so that the prior information of the individual component in question may be used to provide insight into its likely failure path. Application of Bayesian updating is explained in Robinson and Crowder [Robinson, et al., 2000]; a brief summary of the Bayesian method is provided here. The method utilizes pseudo-inversion of the data matrix in Equation 27 to estimate the parameter vector, where in the mock F&W facility, X is the MUF estimate derived from the load cell measurements, Y is the material diverted, and b is the relationship between the two.

$$Y = bX$$

Equation 27 - Linear Regression Equation in Matrix Form

$$b = (X^T \Sigma_y^{-1} X)^{-1} X^T \Sigma_y^{-1} Y$$

Equation 28 - Pseudoinversion to Solve for the Parameter Vector

Prior information about the model parameters may be appended by adding additional rows to X and Y where the added row in X has a value of one at the j^{th} column (and zero elsewhere), and the prior value of the new parameter is added to Y in the j^{th} position. The square matrix of variances is also expanded to incorporate the variance and covariances related to the prior information at the end of the matrix. This may be repeated multiple times to include prior information about any regression parameter. While this approach is a very simple, quick, and straightforward application of Bayesian updating, it has been shown to provide reliable predictions of RUL that account for both the historical data and the prior knowledge of the individual component. The posterior estimate provided by appending the matrices X and Y, then solving for the parameters b can then be used to extrapolate a new failure path for the individual component.

At this time, the GP model uses historical knowledge yet calculates a degradation path specific to the component in question. This is the model that will be used in this research for the prediction of time to SQ production.

3.2.4.1 Application of GPM to Diversion Detection and Prognostics

The general path method was designed for application to process monitoring of plant systems and components where the observer is interested in failure; this author believes that the method can be applied by analogy to safeguards. To complete this analogy, the variables in the general path method must be redefined to verification variables. Revisiting Equation 26, the verification analog is identical and given in Equation 29.

$$\bar{x}_j = \eta(t_j, \bar{\phi}, \bar{\theta}) + \bar{\epsilon}_j$$

Equation 29 - General Path Historical Model for Safeguards

The analogy is straightforward; the "degradation" term, x_{ij} , is the undeclared production of LEU or HEU where the soft limit of failure is the production of one SQ of material. For LEU, this limit is 75 kg and for HEU, the limit is 25 kg [IAEA, 2002]. In most cases, the initial production of undeclared material will be

assumed to be 0 kg so that the initial "degradation" is zero. (If a nonzero initial production of undeclared material is believed to exist, the soft limit may be adjusted accordingly to monitor for a cumulative production of undeclared material of one SQ.) The fixed effects vector, ϕ , includes measurements common to all GCEP monitoring applications, such as the load cell data and cold trap operation. The random effects vector, θ_i , represents any particular characteristics of an individual GCEP that may be relevant to the general path model. For example, if a GCEP has a standard procedure of sampling of the feed material prior to processing into the cascade, that sampling may be accounted for within the general path model for that plant. The measurement error, ε_{ij} , is known based on the precision of the load cells, variance of mass lost due to cold trap operation, etc.

The historical data was created using runs previously performed on the mock F&W facility. Since a diversion event is most likely linear in terms of the rate of undeclared production, the model was assumed linear. Also, because a diversion event is triggered by an external decision (made either by the operator or a third party), the model assumed that diversion is not necessarily constant at all times. For example, if a GCEP runs without any diversion for three weeks, any monitoring approach must remain sensitive to the possibility that diversion may be initiated in the fourth week. To maintain the flexibility of detection of a changing enrichment regime with a linear model, a piecewise approach was used where the GP was applied to a moving window of the facility data.

Bayesian updating was performed using the measurement data collected during the run being evaluated as the prior knowledge. The parameter estimates from this data were appended into the X and Y matrices to provide new estimates of the time to SQ production, as well as the associated uncertainty estimates.

Since GCEP monitoring requires that all observed data must first be agreed upon by the IAEA and the facility operator, determining the minimum amount for effective verification is desirable; monitoring unnecessary variables might not provide sufficient benefit to justify the risk of revealing proprietary operational information. Therefore, only the minimal number of necessary variables was measured. Once this minimum level is established, the uncertainty estimates of the current level of undeclared production and of the prediction of time to SQ production were generated. These estimates were used to measure the ability of the model to reliably detect diversion and predict time to SQ production.

The model was tested against a variety of legitimate and diversion runs to test for false positive errors (i.e. declaring diversion when none exists) and false negative errors (not declaring diversion when present). The moving average window of time for evaluation was tested to find a balance between stability of the model, where longer time windows minimize the effect of noise, and responsiveness, where shorter time windows allowed for quicker detection of diversion.

3.2.5 Particle Filter Method

Particle filtering was originally developed to provide an estimation of the marginal probability in Bayes' Theorem that would allow for modeling of nonlinear systems and potentially non-Gaussian noise [Cadini F. , 2009]. The PF method utilizes Monte Carlo simulation to provide an approximate solution to the marginal distribution by generating artificial random samples and comparing their distribution to that of the measurements. The particle filter method first starts with Bayes' Theorem, given in Equation 30.

$$p(\mathbf{x}_k | \mathbf{z}_{0:k}) = \frac{p(\mathbf{z}_k | \mathbf{x}_k) * p(\mathbf{x}_k | \mathbf{z}_{0:k-1})}{p(\mathbf{z}_k | \mathbf{z}_{0:k-1})}$$

Equation 30 - Bayes' Theorem

Here, $\mathbf{x}$ represents the state-space vector of the system, which is not directly measured. In condition monitoring, this is typically the prognostic parameter or other measure of system health. The term $\mathbf{z}$ represents the vector of measurements. Measurements are first taken at time = 0; the "current" time, or the time of interest is represented by the subscript k , and the previous time step is $k-1$. The term $p(\mathbf{x}_k | \mathbf{z}_{0:k})$ is defined as the posterior distribution and represents the distribution of the likelihood of a system state $\mathbf{x}_k$ existing given the measurements $\mathbf{z}_{0:k}$ (i.e. all measurements, including the current measurement). The term $p(\mathbf{z}_k | \mathbf{x}_k)$ is the conditional probability and represents the likelihood that a given state would yield the current measurements. The term $p(\mathbf{x}_k | \mathbf{z}_{0:k-1})$ is the prior distribution and is the likelihood that a given state would exist based on $\mathbf{z}_{0:k-1}$, where $\mathbf{z}_{0:k-1}$ represents all measurements prior to the current measurement. Finally, $p(\mathbf{z}_k | \mathbf{z}_{0:k-1})$ is the marginal probability and represents the likelihood that the current measurements would occur given all previous measurements.

The long-standing difficulty with Bayes' Theorem is determining the marginal probability, and this is the purpose of particle filtering. The first step in particle filtering is Sequential Importance Sampling (SIS), where a known distribution is used to generate random samples for $\mathbf{x}_{0:k}$. (In contrast to GPM, SIS

updates all particles simultaneously at a given time step rather than updating a single particle all the way to failure.) The distribution only needs to ensure that the range of possibilities is covered, though a distribution that closely resembles the true distribution of probabilities should provide faster convergence and more reliable results. This distribution is defined as an *importance function*, as shown in Equation 31.

$$q(\mathbf{x}_k | \mathbf{z}_{0:k})$$

Equation 31 - Importance Function for Sequential Importance Sampling

Next, weights for the sample particles are defined by relating the importance function to the posterior distribution, as in Equation 32.

$$w_k^i \propto \frac{p(\mathbf{x}_k^i | \mathbf{z}_{0:k})}{q(\mathbf{x}_k^i | \mathbf{z}_{0:k})}$$

Equation 32 - Definition of Sampling Weights

In Equation 32, both terms in the ratio are unknown. However, the weights from the previous time step are known and can be used to approximate the weights using Equation 33. The approximation may require normalization to form a true pdf; this can be readily performed after all weights are estimated.

$$w_k^i \propto w_{k-1}^i * \frac{p(\mathbf{z}_k | \mathbf{x}_k^i) p(\mathbf{x}_k^i | \mathbf{x}_{k-1}^i)}{q(\mathbf{x}_k^i | \mathbf{x}_{0:k-1}^i, \mathbf{z}_{0:k})}$$

Equation 33 - Approximation for Sampling Weights

A generic visual example of the weighting process may be seen in Figure 10 and Figure 11. Prior to weighting, all samples have equal weight. At time step 20, those particles whose states have a higher likelihood of representing the true state of the system (as estimated through measurements) receive greater weight; the particles with less likelihood of representing the system receive less weight. Having re-weighted the particles, they now represent an updated posterior distribution of the marginal probability and the new posterior distribution may be estimated.




Figure 10 - Sample Estimates of the System State Prior to Weighting




Figure 11 - Sample Estimates of the System State after Weighting

As this process is repeated, the particle weights are continually updated every time Bayes' Rule is applied. Since SIS will add weight to particles with the highest likelihood at the expense of lower-likelihood particles, the process will eventually drive the weights of all particles to zero except for the highest-likelihood particle, a problem known as degeneracy. To avoid degeneracy, the particles are occasionally redistributed by Sequential Importance Resampling (SIR).

SIR may be conducted by a variety of methods, but the general approach is to replace the existing weighted particles with new unweighted particles chosen by the posterior distribution provided by SIS. In review, SIS may be seen as a filter of particle weights where the weights are updated to fit the estimated posterior distribution. SIR may therefore be seen as a filter of the particles themselves, where new particles are chosen based on the weighted particles after SIS. This two-stage process is the heart of particle filtering.

Because particle filtering employs a Monte Carlo process, uncertainty estimates may be readily provided by the existing particle distributions. The state estimate at the present time is given by Equation 34. The failure probability estimate at a future time $k+i$ is given by Equation 35. The failure time distribution at a future time $k+i$ is given by Equation 36.

$$p(\mathbf{x}_k | \mathbf{z}_{0:k}) \approx \sum_{i=1}^N w_k^i \delta(\mathbf{x}_{0:k} - \mathbf{x}_{0:k}^i)$$

Equation 34 - Particle Filter State Estimation

$$\hat{p}(k+i) \approx \frac{\sum_{m, x_k^m > d^*} w_k^m}{\sum_{n, x_k^n > d^*} w_k^n}$$

Equation 35 - Particle Filter Future Failure Probability Estimation

$$p(\tau_k | \mathbf{z}_{0:k}) \approx \sum_{i=1}^N w_k^i \delta(\tau_k - t^i)$$

Equation 36 - Particle Filter Future Failure Probability Distributions

Additional statistical inferences may be made via the particle distributions, such as 95% confidence intervals, hypothesis tests, etc.

3.2.5.1 Estimation of the Measurement and Dynamic Noise Terms

Particle filters require an estimate of two types of noise within the system: measurement noise and dynamic noise. Measurement noise is the error associated with the measurements of the state variables. The greater the certainty of the measurement, the more significant the difference between the measured system state and the particle state becomes. The weight updates during SIS are greatly influenced by measurement noise, as high uncertainty of measurement noise reduces the effect of SIS on particle weights. (See Equation 37, where the term σ_m^2 represents the measurement noise estimate.) Dynamic noise is the rate at which the system dynamics can change (i.e. how fast the system can accelerate between states). If a system can rapidly change its behavior (e.g. quickly transition from a constant MUF to a rapidly increasing MUF), then the particles must be distributed widely enough to ensure that the system dynamics are captured.

$$w_k^i \propto w_k^i * e^{\frac{-(Y-\hat{Y})^2}{2*\sigma_m^2}}$$

Equation 37 - Weight Update Equation

High measurement noise and dynamic noise both lead to increased uncertainty in the predictions of future system states, albeit by different mechanisms. High measurement noise reduces the effect of particle re-weighting, allowing outlier particles to maintain more weight and thereby increasing the weighted variance of the particles. High dynamic noise demands a greater spread of particles to effectively cover the possible range of future system states, which also increases particle variance. The noise estimates must be properly matched to system behavior to provide realistic uncertainties as well as particle behavior. A detailed discussion of the methods used to estimate and control measurement and dynamic noise terms is provided in Appendix B.

3.2.5.2 Application of Particle Filters to Linear Equations

In a typical particle filter tracking application (not a prognostic application), the particles are distributed around the estimated position of the target at some initial time, then given random velocity vectors based on the dynamic noise estimate of the system. As the particles move, and as new measured estimates of the target position are taken, the particles are weighted according to their likelihood of

representing the true position of the target, using Equation 33. Once degeneracy creates a sufficient imbalance of weights among the particles, they are redistributed according to their weighted distribution and the process continues.

To apply the particle filter to prognostics, an expected path shape must be provided for the particles to follow or their paths will simply define an expanding region around the last known target measurement without any regards to the particle's trending motion. In [Orchard, et al., 2009], the path was defined as an exponential growth of a crack, and the critical information was the initiation of the crack growth and its severity. For the problem of tracking and prognosing the mock F&W facility, there can be no genuine assumption about the future shape of the MUF trajectory as it is a function both of the system and of human decision. In a well-operated condition with no diversion or other impedance on the MUF, the MUF tends to remain constant over time with only oscillatory deviations due to the natural operation of the facility. If other influences on the facility operation are introduced (such as diversion, a biased load cell, or the use of an off-scale cylinder), the MUF may take any variety of shapes.

Because of the lack of any foreknowable MUF trajectory, no one model may be considered more viable than another based on any mathematical rigor. Instead, a linear model (Equation 38) is adopted with the implicit understanding that any prognosis merely answers the question: "if the system were to continue as operated, what is the expected outcome?" Yet, the linear model provides a simple elegance: if the mock F&W system is not experiencing diversion, bias, or other impingement on the MUF, then the MUF level should remain constant and the linear model should merely provide an average of zero for the coefficient vector **A**.

$$MUF_{i,t} = A_i * t + b_i$$

Equation 38 - Linear model for MUF Estimation

If nonlinear events occur, such as time-variant diversion or intermittent diversion, the piecewise nature of the particle filter should track the MUF with the best linear approximation within each operating region. If a prediction is requested at any time, that prediction would therefore be the most recent linear trend of the data. The linear assumption is therefore not grounded in a verifiable physical attribute of the mock F&W facility, but rather based on its versatility within the particle filter framework.

3.2.5.2.1 Initial Linear Model Coefficients

When the particle filter is first initiated, the particles are initially distributed about the initial measured MUF value. This value is taken as a zero reference, which means that all MUF calculations are taken relative to this point. The initial starting point for the particle filter is based on the initiation of the mock F&W facility pumps. Observation of the system has shown that the PI controller achieves dynamic equilibrium around 1,000 seconds after the initiation of the first pump, whether the system is in single-pump or dual-pump operation. This may be seen in Figure 12, where the first pump is initiated at roughly 2000 seconds and the PI control stabilizes by about 3000 seconds. (The effects at shutdown around 8000 seconds are an artifact of the mock F&W facility and are not considered in this study as the final MUF value is influenced by the manual closing of shutoff valves and is not an analog to GCEP operation.)

From this starting point, the particles are evenly distributed from the calculation of a severely increasing MUF to a severely decreasing MUF by varying the first-order coefficient of Equation 38. As the particles travel and load cell measurements are taken, the particles that best follow the path of the MUF receive the most weight. As SIR commences, the particles are redistributed according to the distribution of weights among the existing particles, providing the Bayesian framework for tracking MUF using a particle filter.

When the particles are resampled in SIR, the new first-order coefficients are chosen based on the linear path between the previous and current SIR position. (For example, if the mean position at the first SIR point calculated a MUF of 0.1 kg and the mean position at the second SIR was a MUF of 0.15 kg, then the mean first-order coefficient would equal the gain of 0.05 kg of MUF divided by the time between SIR points.) Given the mean coefficient, each particle's new coefficient is chosen by adding a random number to the mean coefficient, where the random number is normally distributed with a standard deviation equal to the estimated dynamic noise of the system. The basic form is given in Equation 39.

$$A_{i,t_2,new} = \bar{A}_{t_2-t_1} + \sigma_{dynamic} * randn(mean = 0, std = 1)$$

Equation 39 - Model Coefficient Update Equation




Figure 12 - PI Control Stabilization and MUF Calculation

3.2.5.3 Application of Particle Filters to Diversion Detection and Prognostics

The PF model was tested in the same manner as the GP model. Similar to the optimization of the GP time window, the PF model was tested to determine optimal frequencies of SIS and SIR updating to provide a balance between model stability and model responsiveness. The particle filter model was also constrained to piecewise linear particle trajectories to provide sensitivity to the decision of the operator to either divert or not divert material at any time. The particle filter model was built using an existing set of legitimate and diversion facility runs, and then tested against new runs to validate the model. As described above, undeclared production, prognostic, and uncertainty estimates were all automatically generated by the particle distributions.

3.2.5.4 Validation of the Fitted Particle Filter Model

Particle filters are particularly appealing for modeling nonlinear and potentially nonmonotonic data. The Monte Carlo process allows the particle filter to adaptively track the data trends without knowing the underlying structure, as the SIS reweighting provides the model information regarding the most likely data trends and the SIR resampling avoids degeneracy and allows the particles to adapt to data trends. The risk associated with such model flexibility is that the particle filter may adapt to any data, regardless of whether there is any underlying information associated with the data. For example, if a particle filter is applied to purely random data, the filter process will provide a model to the data, but no indication that the data is random and has no structure. To justify the particle filter model, its coherence to the data was tested using the ICOMP scoring system discussed in §3.2.3.1.

For a given region between SIR updates (i.e. when particles are only re-weighted, not resampled), the weighted mean of the particle paths was considered a model to the data. This model, referred to as the "linear" model due to the linear restrictions within the particle filter, will be scored based on its fit to its region of data using the misspecified ICOMP criteria. (Again, the misspecification feature of this metric adds additional penalty if the particle filter is indeed the "wrong" model for the data.) The data were also scored against a constant term model, which was the mean of the data in this region, as well as a random model, in which no coefficients were specified and the predictor data were a vector of normal random data with zero mean and unitary covariance. An example scoring is shown in Table 1. In this example, the particle filter scored the best in regions one and two, but then the random model suddenly provided the best score for the data from 2000 to 3000 seconds. In this region, the particle filter model would be considered invalid and any results from the particle filter in this region would be questioned.

Table 1 - Sample ICOMP Scores for Theoretical Mock F&W Facility Data

		Model Type		
Start Time	Stop Time	Linear	Constant	Random
0	1000	-100	-50	200
1000	2000	-250	-100	150
2000	3000	-50	-75	-200
3000	4000	-200	-250	120

Meanwhile, the data from 3000 to 4000 seconds is best represented by the constant term model. Again, the particle filter's results would be considered less valid than treating the MUF as a constant through this region. The difference in the interpretation between regions three and four is that region four would be considered a coherent operating range of legitimate operation, while region three was poorly modeled by the particle filter and by the constant MUF assumption. This would be a flag for an event requiring further inquiry, such as data corruption. By using ICOMP rather than traditional statistical measures like p-value tests, the model was guarded against overfitting (e.g. fitting a linear model to constant data), and model selection was no longer restricted by an arbitrary limit such as an alpha value.

In addition to ICOMP scores, the histogram of the particle filter prediction will be evaluated. If the MUF data is well-defined by the particle filter, then the weighted histogram of particle filter predictions should approximate a normal distribution. Even if the measurement noise of the data were not normal, the accumulation of weight updates would render the effect of the noise as a normal distribution upon the particles based on the Central Limit Theorem. If the particles did not effectively capture the underlying trends to the data, then abnormal deviations to the histogram would appear, such as skewness or bimodality. Such features would indicate that the data is in transition relative to the particles, which does not necessarily indicate that the model is insufficient, but that any predictions based on the particle filter in this region are suspect and may not be reliable.

This effect is visible in Figure 13, where a normal curve is fit to a normal distribution and a right-skewed beta distribution. Applying a Lillieford test for normality to each distribution where the null hypothesis states that a normal distribution fits the data, the normally distributed data does not reject the null with a p-value of 0.276 while the beta distributed data rejects the null with a p-value of approximately zero. The first case suggests that the particles are well centered on the data in the region of inquiry, while the second indicates that the particles are distributed unevenly with regard to the data. In this second case, any predictions of SQ production will be considered less reliable.




Figure 13 - Fitting Normal Curves to Normal and Beta Distributions

3.3 Software Development

The monitoring and prognostic methods developed in this research will be developed into automated software using MATLAB™. This software will be designed for observation of the mock F&W facility and will provide a simulation of the manner in which the monitoring of a GCEP may be automated and performed by on-site IAEA computers. Basic automated analysis of the mock F&W load cell data has been provided by an already-developed toolkit called *PlotEvents* [Henkel, 2010]. *PlotEvents* will be expanded to include accountancy scale data and any other measurements deemed relevant to the verification process. The software will be developed in two categories: summary reports that can be sent on to IAEA headquarters, and analytical tools that can be used by an inspector during an on-site visit. The summary reports represent the software's declaration of the state of the facility and will only include data that is deemed acceptable for transmission beyond the facility. The analytical tools will provide an inspector with the ability to diagnose the software's declarations so that the inspector can determine the source of any warnings as well as search for reasons why a warning might not have been transmitted as expected.

This chapter presented an overview of the mock feed and withdrawal facility as an analog test platform for developing remote monitoring techniques for gas centrifuge enrichment plants. General path and particle filter prognostic models were developed to predict the time until diversion of significant quantities of material from the mock feed and withdrawal facility. Finally, the *PlotEvents* software tool for automated analysis of the load cell monitoring data was introduced. The next chapter presents improvements made to the mock feed and withdrawal facility over the course of this research and the results of an observability analysis of the mock feed and withdrawal facility and its relevance to the monitoring process. Analyses of the prognostic methods are presented to justify the choice of the particle filter as the prognostic model for this research. Finally, four case studies are analyzed with the particle filter to estimate the time to diversion of a predefined significant quantity of material. The results of the case studies are evaluated based on accuracy, precision, flexibility, and uncertainty.

4 APPLICATION AND RESULTS

The results of this research may be divided into two distinct categories: the work related to the improvement and operation of the mock F&W facility, and the analytical work performed to analyze the data collected from the mock facility. Subsections 4.1 through 4.3 discuss applied work performed to improve the mock F&W facility and to parameterize its performance. This includes a discussion of the changes to the mock F&W facility and its operation, parameterization of the facility, an observability analysis of the facility, and efforts to reduce noise when measuring surge tank water level. Subsection 4.4 discusses preliminary evaluation of prognostic models that builds towards the ultimate objectives of this research.

4.1 Mock Feed and Withdrawal Facility Improvements

Several improvements were made to the mock F&W facility to improve the consistency of the data and to increase the flexibility of the facility to mimic potential events at an enrichment facility. Data consistency was improved by automating control of the surge tank water level using a proportional-integral (PI) controlled valve instead of the operator's visual estimation of water level. PI control itself was subsequently optimized via low-pass filtering and trimmed mean averaging of the water pressure transducer signal to provide a clean, low-variance reading of the water level, and by eliminating physical sources of signal noise within the facility.

4.1.1 Automated Control of the Surge Tank Control Valve

Originally, flow from the surge tank was controlled by two manual throttle valves – one valve on the tubing leading to the product stations and one valve on the tubing leading to the tails stations. The facility operator would visually observe the water level inside the surge tank and adjust the throttle valves as necessary to maintain the desired water level within the surge tank. This configuration led to problems of consistency: the surge tank level was controllable only within the operator's ability (and interest) to maintain a constant level, and the "enrichment" (i.e. the product / tails ratio) was consistent only as far as the operator's ability to guess the appropriate throttle valve adjustments. Figure 14 shows the two throttle valves (left) and the cutoff valve (center) on the outlet of the surge tank.



Figure 14 - Original Surge Tank Control Valve Setup

Without consistent flow from the surge tank, the data collected from the mock F&W facility could not effectively be used as an analog to GCEP operation. The system of manual throttle valve control of surge flow was replaced by a proportional-integral (PI) controlled ball valve (the "surge tank control valve") to regulate the quantity of flow. The two throttle valves were then used to regulate the flow from the surge tank control valve and maintain the desired product / tails ratio. The PI controller was created using LABView™ on a Dell™ D620 laptop. The governing equation for the PI controller is given in Equation 40.

$$u(t) = K_p e(t) + K_I \int_0^t e(\tau) dt$$

Equation 40 - PI Control Equation

In Equation 40, the term $u(t)$ is defined as the change in actuator voltage and the error term $e(t)$ is defined as the difference between the surge tank level setpoint and the actual height of water in the surge tank (actual height – setpoint = error). If the surge tank level is too high, a positive error term would increase the voltage to the valve actuator, opening the valve. K_p is defined as the proportionality constant, and K_I is the integral constant.

The water level was measured using a pressure transducer placed at the outlet of the surge tank. The transducer signal was converted from a voltage to a water pressure, which could then be related to the height of the water above the transducer through Equation 41. By accounting for the difference in height of the pressure transducer and the surge tank outlet, the water level relative to the surge tank control valve could be calculated.

$$P_{water} = \rho_{water} * g * h_{water}$$

Equation 41 - Relationship Between Water Level and Pressure

Through experimentation, effective values for the constants were determined to be $K_p = 12 \text{ V}/(\text{kg}/\text{min})$ and $K_I = 0.5 \text{ V}/(\text{kg}\cdot\text{sec}/\text{min})$. The steady-state drift of the surge tank level was observed to be about 0.02 inches, and the valve actuator typically changed position only once about every four to five seconds, with incremental changes of less than 0.1 V.

4.1.2 Reduction of Pressure Transducer Noise

When first applied, the PI control system was highly susceptible to noise from the pressure transducer, resulting in excessive operation of the valve actuator. The 1 Hz updates in valve position caused constant actuator motion with changes in the actuator control signal (see Equation 40) commonly as high as 0.5 V (out of a 10 V range of operation). With mock F&W facility operational runs routinely lasting two hours and occasionally five to eight hours, actuator wear was a concern.

According to Equation 40, reducing the changes in the control voltage to the actuator could either be accomplished by reducing the constants or by reducing the changes in the transducer voltage. Reducing the constants K_p and K_i would reduce the effectiveness of the PI control scheme and was therefore undesirable. Reducing the changes in the transducer voltage was achieved through: signal analysis to isolate signal noise, signal sampling and smoothing optimization, and identification and correction of physical sources of noise within the mock F&W facility.

4.1.2.1 Signal Analysis

The pressure transducer signal was recorded at 256 Hz for several minutes while the system was static (i.e. no water flowing into or out of the surge tank). A subset of the raw signal is shown in Figure 15. The transducer signal had a standard deviation of 0.00955 V, which, when converted to water height, corresponds to a standard deviation of water height of 0.47 inches. In comparison, the facility was expected to operate with extreme surge tank levels of 19.5 to 20.5 inches, with normal operation within 19.95 to 20.05 inches, indicating that signal noise dominated the underlying signal of changes in the water level. To examine the noise of the pressure signal, a Fast Fourier Transform (FFT) algorithm was applied by importing the signal into MATLAB™. The transformed signal is shown in Figure 16.




Figure 15 - Pressure Transducer Voltage with no Surge Tank Flow




Figure 16 - Unfiltered Pressure Transducer Signal - Fast Fourier Transform

A very strong 60 Hz 'hum' can be seen in the transducer signal, as well as a 120 Hz signal. The 60 Hz noise was an artifact of the AC power supply used to energize the transducer, and the 120 Hz noise was most likely harmonic noise from the same AC source. A third-order Butterworth low-pass filter was applied with a break frequency of 5 Hz. This filter was applied to the exact same signal as in Figure 15; the results are shown in Figure 17 and Figure 18.

The low-pass filter eliminated the severe 60 Hz and 120 Hz noise, and the frequency with the largest noise signal was 4 Hz, which was more than a full order of magnitude lower in strength than the 60 Hz hum. The standard deviation of the filtered signal was 0.001 V (~ 0.05 in), which was an order of magnitude lower than the standard deviation of the unfiltered signal.

4.1.2.2 Signal Sampling Rate and Smoothing Optimization

Having eliminated the electrical noise from the pressure transducer signal, the sampling rate was varied to quantify its relationship to signal variance. A plot of the sample rate dependence of transducer standard deviation is shown in Figure 19. Because of the desire to use fast Fourier transforms (FFTs) to filter the signal, sampling rates were restricted to powers of two (e.g. 256 Hz, 512 Hz, 1024 Hz, etc.).

As the sampling rate increased, the signal standard deviation decreased. The relationship was linear in log-log form up to a sampling rate of 8096 Hz, after which no significant reduction in noise was observed. Since operation of the system with high sampling rates did not inhibit the LABView™ PI controller or generate significant heat within the laptop, a sampling rate of 8096 Hz was used to minimize signal noise.

With the sampling rate and low-pass filter in place in the PI controller, a trimmed mean filter was used to average all of the data points for each second. The highest ten percent and lowest ten percent of the data values were eliminated, with the mean of the remaining eighty percent of the signal serving as the transducer voltage reading for that second. This value was then converted to a pressure value, and ultimately into a water level measurement so that the water level error $e(t)$ could be calculated for Equation 40.




Figure 17 - Filtered and Unfiltered Pressure Transducer Signals




Figure 18 - Filtered Pressure Transducer Signal Frequency Spectrum




Figure 19 - Voltage Standard Deviation as a Function of Sample Rate

4.1.2.3 Reduction of Physical Sources of Pressure Transducer Noise

Two sources of pressure variation were determined to exist in the product and tails flow tubes. The first was the presence of air in the tubes. Initially, no attempts were made to maintain an air-free environment within the tubes, but inconsistent flow out of the surge tank would often cause control problems, particularly during low-flow conditions during startup and shutdown. During such low-flow conditions, flow was often observed to suddenly cease or surge, requiring significant action of the throttle valves to regulate the flow. Once flow rates were of sufficient magnitude, these inconsistencies would generally disappear. To eliminate the flow inconsistencies, air was eliminated from the tubes by changing operating procedures. The air-free product and flow tubes provided more consistent flow during startup and shutdown, eliminating the need for throttle valve adjustment and providing more consistent flow throughout operation of the facility.

The second source of noise was a loose tube seal at the shutoff valve for tail station 2 (TS-2). When TS-2 was being filled, air would periodically enter the flow stream at the shutoff valve. The air introduction affected the static pressure inside the tube, causing a backpressure variation that altered the pressure transducer voltage. This effect caused a fluctuation in the transducer voltage with a frequency of about 4 Hz, which in turn caused the PI controller to move the actuator with a 4 Hz oscillation about the nominal actuator voltage for the current operating regime. The TS-2 valve seal was fixed by replacing the tube section exiting the shutoff valve, eliminating the air entrainment and the subsequent pressure variation.

4.2 Mock Feed and Withdrawal Facility Operation

The quality of the mock F&W analog is only as good as the manner in which the facility is operated. By converting surge tank level control from an operator's visual estimation of level to a PI controller, the response of the facility to changes in feed rates (e.g. startup, shutdown, and changes in feed tanks) was made more consistent. Still, the mock F&W facility is largely under manual operation, with the feed pumps and all tank loadings/unloadings performed by human interaction. To produce consistent data, the operational procedures of the mock F&W facility were modified and formalized.

4.2.1.1 Separation of Tank Loading and Connecting Features

Originally, as a tank was loaded onto a feed, tails, or product station, the feed or withdrawal tube was connected immediately after removing the lift truck from the tank. However, discussions with ORNL, SRNL, and the IAEA revealed that there is usually a notable time lag between a cylinder loading and its

connection to the cascade. When GCEP load cell data is analyzed, these two events can be distinctly observed as separate events – a feature that was heretofore missing from the mock F&W load cell data. The operation of the mock F&W facility was therefore modified so that a minimum wait time of one full minute was enforced between the loading of a tank and its subsequent connection to the process. A similar time lag was also observed to occur during offloading, where the cylinder would be offloaded at some measureable time after the cascade connection was removed. The offloading procedure of the mock F&W facility was similarly modified to produce these distinct data features.

4.2.1.2 Cold Trap Emulation

In a typical UF_6 feed cylinder, a small amount of light gases are present in the feed stock; these light gases are an unavoidable side effect of the refinement and conversion processes that produce the UF_6 from uranium ore, but are highly undesirable in the enrichment processes, as the light gases would immediately travel to the interior of a centrifuge and inhibit the separation of U-235 from U-238. To relieve the cylinders of any light gases, the feed cylinder is first heated to release the light gases from the UF_6 . The light gases are then allowed to pass out of the cylinder through the cascade connection and diverted to a vent stream. Since a small amount of UF_6 gas inevitably accompanies the light gases during venting, the vented gas is passed through a cold trap where the UF_6 is precipitated out of gaseous phase and collected. The UF_6 , while a very small amount, may be recycled if the facility desires.

The cold trap procedure affects the MUF calculation due to the loss of mass in the feed cylinder represented by the vented gases. In the mock F&W facility, no emulation of the cold trap had been possible prior to work on this research, but discussions with the IAEA indicated interest in demonstrating the effect of cold trap operation on MUF calculations and subsequent diversion estimates. To meet this interest, a bleed valve on the feed tubing of the mock F&W facility was reserved for removal of a small quantity of the feed water. This removal was performed by pumping a small amount of water from a new feed tank and relieving the pumped water through the bleed line rather than into the surge tank. The removal of water can be performed for the first cylinder prior to actual startup of the mock enrichment process, and can be performed in-line for any other cylinders without interrupting the process.

4.2.1.3 Sampling

Sampling of the product tanks was performed by siphoning after the product tanks were filled but prior to its removal from the filling station. The sampling process imitates sampling of product LEU that is

performed to both verify the enrichment levels of the material and to demonstrate the quality of the product to purchasers. The sampling may be performed while the mock enrichment process is still running (i.e. after the product tank in question has been filled and the mock process is filling a different product tank) or at the end of the process (to sample the final product tank). Each sampling process was restricted to the removal of about 0.05 kg of water – an amount small enough to be discrete, yet large enough to measurably affect the MUF calculations.

4.3 Mock Feed and Withdrawal Observability

Prior to analysis of the mock feed and withdrawal facility for detection of diversion, the system must first be determined to be "observable". An observable system allows for the prediction of one state variable with the knowledge of a defined subset of the remaining state variables. Without observability, knowledge of a subset of state variables may lead to multiple valid predictions of the appropriate value of the final state variable, rendering impossible any verification of state variables.

The classical approach to determining observability is to determine the rank of the observability matrix as defined by Equation 42 . If the rank is of the same order as the total number of state variables, the system is determined to be observable; that is, the variables do contain a strict, one-to-one relationship with each other that guarantees that the knowledge of one state variable may be gained through measurement of the remaining state variables.

$$O_b = \begin{bmatrix} C \\ CA \\ \dots \\ CA^{n-1} \end{bmatrix}$$

Equation 42 - Observability Relationship

4.3.1 Feed and Withdrawal Station State Variables

The mock F&W facility was defined as having three feed stations, two tail stations, and three product stations. Any decisions as to which stations were in use at any time were made by the operator and were not bounded by any logical or mathematical relationship. For example, feed could arbitrarily be introduced through feed station one, two, or three; the feed station used also has no bearing on which product and tails stations are used. Rather than focus on the individual stations, the feed, product, and tails flows were defined as aggregate from all stations, resulting in three station state variables: feed flow, product flow, and tails flow. The simplified schematic of the mock F&W system shown in Figure 20 illustrates this simplification with cumulative feed, product, and tail flow rates.




Figure 20 - Simplified schematic of the mock F&W facility

In a GCEP, the relationship between the product and tail flow rates is determined by the level of enrichment desired by the plant operator as well as the enrichment efficiency of the plant. In the mock F&W facility, the relationship was governed by two throttle valves – one for each flow – that maintained a constant relationship between product and tail flow. The product and tail flow rates could then be defined by their relationship to the flow rate exiting the surge tank by defining a product flow parameter, p_e , as in Equation 43.

$$p_e = \frac{\text{Product Flow Rate}}{\text{Surge Tank Exit Flow Rate}}$$

Equation 43 - Definition of Product Flow parameter p_e

The tail flow could then be defined with respect to the product flow as in Equation 44.

$$\text{Tail Flow} = \frac{1 - p_e}{p_e} \text{Product Flow}$$

Equation 44 - Relationship between Product and Tail Flow

4.3.2 PI Control State Variable

The feed into the surge tank was seen as a forcing function dictated by the feed flow rate, but the flow out of the surge tank was controlled by the PI controller, which actuated the surge valve at the exit of the surge tank. The flow into the surge tank is given in Equation 45, and the PI controller was governed by Equation 46:

$$Q_{in} = \frac{dF_{mass}}{dt}$$

Equation 45 - Mass Flow into the Surge Tank

$$Q_{out} = \rho a \sqrt{gh}$$

Equation 46 - Exit flow Relationship from the Surge Tank

Here, ρ is defined as the density of the fluid (water), a is the cross-sectional area of the surge tank exit pipe, g is the acceleration due to gravity, and h is the height of the water in the surge tank, with respect

to the exit pipe. The $\sqrt{h}$ term presents nonlinearity to the system. However, the nominal height of water in the tank (relative to the exit pipe) is 20 inches, and normal variation of water height does not exceed 19.9 to 20.2 inches during operation. (During operation, the water level varies between 19.98 and 20.02 inches. However, the level was observed to vary from 19.9 to 20.2 inches during startup.) Assuming that water is incompressible and has constant density, water levels of 19.9 and 20.2 inches experience only about a 0.5% difference in flow rate compared to a water level of 20 inches. The height-flow rate relationship may thus be linearized with acceptable error as in Equation 47:

$$\frac{dQ_{out}}{dh} \cong \frac{\rho a}{2} \sqrt{\frac{g}{h_o}} \bigg|_{h_o=20 \text{ inches}} * h$$

Equation 47 - Linearization of the Water Level / Flow Relationship

4.3.3 PI Controller State-Space Representation

The flow through the surge valve was regulated by a PI controller, which adjusted the position of the throttle valve by changing the control voltage of the valve actuator. The actuator would shut the valve entirely with a control voltage of 0 V, and would open the valve entirely with a control voltage of 10 V. Typically, the control voltage of the actuator would be about 3 to 5 V during operation. The PI controller was governed by the relationship in Equation 48:

$$\frac{dV_{actuator}}{dt} = K_p \delta h + K_I \int_0^t \delta h dt$$

Equation 48 - PI Controller governing equation

K_p represented the proportionality constant, and K_I represented the integral constant for the PI controller. The variable δh was defined as the difference between the surge tank water level setpoint and the actual height of water in the surge tank. If the water level was lower than the setpoint, δh would be a negative value and would decrease the control voltage to the actuator, thereby restricting the surge tank control valve and restricting flow out of the surge tank. The control scheme was dominated by proportional control, and the integral term eliminated the steady-state error of the controller.

The integral term presented a nonlinearity in the PI control equation. However, Equation 48 could be linearized by defining a variable substitution for the integral term [Chau, 2002], as in Equation 49:

$$K_I \int_0^t \delta h dt = x_1$$

$$\frac{dx_1}{dt} = K_I \delta h$$

Equation 49 - Variable Substitution for the PI Integral Term

This substitution generates the coupled system of linear equations for the PI controller in Equation 50:

$$\frac{dV_{actuator}}{dt} = K_p \delta h + x_1$$

$$\frac{dx_1}{dt} = K_I \delta h$$

Equation 50 - PI Controller Equation in Linear Coupled Form

4.3.4 State-Space Model of the Mock Facility

With the surge tank water height / outlet flow and the PI controller equations linearized, the entire facility may be written in state-space form. The forcing function is the feed flow, which is controlled by manual operation of the feed pumps. The state variables are: the perturbation in surge tank mass from its setpoint, the actuator voltage of the control valve, the integral term of the PI controller (as in Equation 50), the product stream mass, and the tail stream mass. The mathematical variables used for each state variable are given in Table 2.

The surge tank mass is directly related to the height of the water in the surge tank by the relationship in Equation 51:

$$S = \rho_{water} A_{xs} h$$

Equation 51 - Surge Tank Mass Relationship

Table 2 - State Variables of the Mock Facility

Symbol	Definition
δS	Perturbation of surge tank mass
V	Actuator Voltage
x_1	Integral component of PI controller
P	Product Mass
T	Tails Mass

Since the surge tank is a cylinder with a smaller internal cylinder occupying some of the volume, the cross-sectional area of the tank is the difference in the tank cross-sectional area and the cross-sectional area of the internal cylinder.

In matrix form, the state-space representation of the system is given in Equation 52:

$$\frac{d}{dt} \begin{bmatrix} F \\ \delta S \\ V \\ x1 \\ P \\ T \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & -c_1 & 0 & 0 & 0 \\ 0 & Kp & 0 & Ki & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & p_e/(1-p_e) & 0 \\ 0 & 0 & c_1 * (1-p_e) & 0 & 0 & 0 \end{bmatrix} * \begin{bmatrix} F \\ \delta S \\ V \\ x1 \\ P \\ T \end{bmatrix} + k_{pump} * \begin{bmatrix} Pump\ Power \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

Equation 52 - Matrix Form of State-Space Representation

The matrix of coefficients is typically defined as the A matrix and contains constants that linearly relate the state variables to their time derivatives. The term c_1 represents the relationship between the actuator voltage (which relates to the exit flow aperture) and the flow rate of the water from the surge tank and is given in Equation 53.

$$c_1 = \frac{a}{A} \sqrt{\frac{g\rho A}{2S_0}}$$

Equation 53 - Definition of c_1 constant in State-Space Equation

4.3.4.1 Definition of the C Matrix

In observability theory, the C matrix defines which variables are observed (i.e. measured) and which variables are not observed. For example, if only the product and tail flow rates are measured for the system defined in Equation 52, then only the two final variables of the five state variables are measured. These measured variables are represented in the C matrix with a value of 1, and the unobserved variables are given a value of 0. The column definitions of the C matrix are shown in Equation 54.

$$C = [Feed \quad Surge \quad PI\ Voltage \quad PI\ Integral \quad Product \quad Tails]$$

Equation 54 – Column Definitions for C Matrix for Measurement of Product and Tail Mass

The *Feed*, *Product*, and *Tails* variables represent the measured weights of the respective feeding and filling stations. The *Surge* variable represents the holdup volume within the surge tank. The *PI Voltage* and *PI Integral* variables are the two terms necessary to linearize the PI controller as in Equation 50. If the observed variables are the feed, product, and tails weights (i.e. a load cell system), then the C matrix would be written as in Equation 55, where the number of rows equals the number of outputs.

$$C = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Equation 55 - C Matrix with Feed, Product, and Tails Observed

Equation 55 represents the normal observation status of a load-cell based monitoring system, where the actual controls of the surge volume are not observed by the monitoring party. The "1" values indicate that the variable is observed, while the "0" values indicate that the variable is not observed.

4.3.5 Observability of the Mock Facility

Using Equation 42, Equation 52, and Equation 54, the observability of the system may be tested knowing only the product and tail flow rates. In this case, the observability matrix numerically reduces to Equation 56:

$$O_b(A, C) = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.0526 \\ 0 & 0 & 1.27e-3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 6.68e-5 & 0 & 0 & 0 \\ 0 & 2.05e-3 & 0 & 6.853-6 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1.08e-4 & 0 & 3.60e-7 & 0 & 0 \\ 2.05e-3 & 6.85e-6 & -2.75e-6 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 1.08e-4 & 3.60e-7 & -1.45e-7 & 0 & 0 & 0 \\ 6.85e-6 & -4.44e-6 & -9.15e-9 & -1.48e-8 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 3.60e-7 & -2.34e-7 & -4.82e-10 & -7.80e-10 & 0 & 0 \\ -4.44e-6 & -2.96e-6 & 5.94e-9 & -4.94e-11 & 0 & 0 \end{bmatrix}$$

Equation 56 - Observability Matrix (Numeric)

The observability matrix of the facility was rank 6, which matched the number of variables and indicated that the system was completely observable. By knowing the product and tail flows, there is a one-to-one correspondence to the feed flow (within the limits of the linearization of the surge flow / water level relationship and the PI controller governing equation). By monitoring the product and tail flow rates over time, the feed flow rate could also be observed over time. If a load cell variable is omitted, then the system can no longer be considered observable. For example, if the product load cell was not observed, then the C matrix becomes Equation 57.

$$C = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Equation 57 - C Matrix with Feed and Tails Observed

The resulting observability matrix (which is not shown for brevity) is only of rank 5. The lack of product load cell data renders the system unobserved, which indicates that the feed and tails observations are not sufficient to monitor the facility. If another relationship could be found to accurately monitor the ratio of product and tails flow, such as measurement of the product and tails flow rates through their respective flow tubes, then observability could be restored. This could be a redundancy that could verify the load cell measurements or provide back information in the event of the loss of load cell data, but if the discussion is confined to only load cell observation, then all load cells must be observed.

4.3.6 Observability of a Cascade System

While the mock F&W facility may be observable, any analysis performed on the mock facility is of value only if a similar analysis may be performed on a real GCEP. A rigorous observability analysis of a GCEP is not possible due to the proprietary nature of the enrichment process. Therefore, the utility of observability may only be estimated by determining if observability is maintained when the mock F&W model is increased in complexity to more realistically mimic a GCEP.

The fundamental difference between the mock facility and a real enrichment plant is the level of complexity; while the mock F&W facility has only a single stage between the feed and the product/tail flows (i.e. the surge tank), a centrifuge enrichment facility is composed of a cascade of centrifuges. To determine whether analysis based on the observability of the mock facility may be extrapolated to a cascade system, the analysis was performed again on a theoretical facility with two surge tanks. The theoretical facility is shown in Figure 21.




Figure 21 - Mock Facility with a Cascade of Surge Tanks

With the added tank, the state space model and C matrix are shown in Equation 58:

$$\frac{d}{dt} \begin{bmatrix} F \\ \delta S_1 \\ \delta S_2 \\ V \\ x1 \\ P \\ T \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & -c_0 & 0 & 0 & 0 & 0 & 0 \\ 0 & c_0 & 0 & -c_1 & 0 & 0 & 0 \\ 0 & 0 & Kp & 0 & Ki & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & p_e/(1-p_e) & 0 \\ 0 & 0 & 0 & c_1 * (1-p_e) & 0 & 0 & 0 \end{bmatrix} * \begin{bmatrix} F \\ \delta S_1 \\ \delta S_2 \\ V \\ x1 \\ P \\ T \end{bmatrix} + k_{pump} \begin{bmatrix} Pump \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

$$C = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Equation 58 - State Space Model for the Surge Tank Cascade

With the second surge tank, the observability matrix is of rank seven, which is indicative of all seven state variables and indicates that the product and tail measurements are sufficient to observe the entire system. In both the single-tank and double-tank analyses, the surge level control data were not measured and remained unknown throughout the entire process. The feed flow history was observable despite not knowing the surge tank water level or the PI control.

4.3.7 Observability during Diversion

Since diversion is the undeclared removal of material from the enrichment process, it is assumed that any attempted diversion would not be monitored (i.e. measured by instruments such as load cells) by the IAEA. Therefore, the diversion would not be an observed variable in the enrichment process. The question is then whether the mock facility is observable during a diversion event without measurement of the diversion variable.

In the protracted diversion scenario utilized at the mock facility, the material is diverted from the product stream since the enriched product is the more valuable commodity at a GCEP. Figure 22 illustrates the location of the diversion line within the simplified schematic of the mock facility.

With the new variable in the system, the diversion may be viewed as a fraction of the normal product flow, where p_{nd} is the proportion of flow that is *not* diverted (i.e. the fraction of product flow that actually goes to the product tank). The quantity $(1-p_{nd})$ is the fraction of product flow that is actually

diverted. The state equation for diversion flow is therefore a linear function of the product flow, as in Equation 59:

$$\frac{dD}{dt} = \frac{(1 - p_{nd})}{p_{nd}} P$$

Equation 59 - State Equation for Diversion Flow

The product flow is now given in Equation 60:

$$\frac{dP}{dt} = \frac{p_e * p_{nd}}{(1 - p_e)} T$$

Equation 60 - State Equation for Product Flow with Diversion Present

With Equation 59 and Equation 60, the state space representation of the system may be written as Equation 61. This is still a linear state space system of equations and is valid for observability analysis.

$$\frac{d}{dt} \begin{bmatrix} F \\ \delta S \\ V \\ x1 \\ P \\ T \\ D \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & -c_1 & 0 & 0 & 0 & 0 \\ 0 & Kp & 0 & Ki & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & p_e * p_{nd} / (1 - p_e) & 0 & 0 \\ 0 & 0 & c_1 * (1 - p_e) & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & (1 - p_{nd}) / p_{nd} & 0 & 0 & 0 \end{bmatrix} * \begin{bmatrix} F \\ \delta S \\ V \\ x1 \\ P \\ T \\ D \end{bmatrix} + k_{pump} * \begin{bmatrix} Pump\ Power \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ D \end{bmatrix}$$

Equation 61 - State Space Representation with Diversion




Figure 22 - Schematic of the Mock Facility with Diversion

If the IAEA has normal monitoring of the load cells available, then the Feed, Product, and Tails load cells would be the observed variables and the Surge Height, PI control variables, and the Diversion would be unobserved. The resulting C matrix is given in Equation 62:

$$C = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

Equation 62 - C Matrix for the Diversion Scenario

In this case, the observability matrix is only of rank 6 when the Feed, Product, and Tails variables are observed. For a seven variable system, this is insufficient to correlate to output variables as defined in the C matrix.

The observability assessment of the mock facility offers two insights. First, the load cells are sufficient for a monitoring scheme as the behavior of the system can be satisfactorily tracked and the results are known to be consistent. If a monitoring system is utilized that accurately calculates the MUF from the load cell data when there is no diversion, then it should be consistently accurate during facility operation and any particular pattern of feed and withdrawal behavior should have a consistent interpretation. Second, (and by extension), if diversion is present, the load cell data is no longer able to provide reliable information about the process. Deviations between the load cell observations and their conclusions should therefore be readily apparent.

4.4 Predicting Time to Diversion of a Threshold Quantity

One question of interest for safeguards monitoring is: how much material is being diverted if a diversion of LEU is indeed occurring? Another way the question may be phrased is: how long until a "significant quantity" of material is diverted from the process? By viewing diversion as a 'degradation' of the enrichment process, the question may be cast into terms of prognostics of a process, where a threshold level of degradation is the typical variable of concern.

For LEU, the IAEA definition of a "significant quantity" is the amount of LEU necessary to contain 75 kg of U-235. For more highly enriched LEU (e.g. 10% enrichment), this amount is lower than LEU of lower enrichments, such as 5% enriched LEU. Because the mock F&W facility does not actually "enrich" water (or provide any analogous differentiation between the product and tails), establishing a significant quantity of diverted water in the mock facility requires some assumptions. For 4% enriched LEU in the

form of UF_6 , roughly 2000 kg of material is necessary to meet the IAEA's SQ threshold. In the mock F&W facility, if the product is assumed to be of similar "enrichment", the 1:100 scale of the facility would mean that 20 kg of product would have to be diverted to reach the same threshold. Diversion of 20 kg of water in the mock facility would require either unrealistically fast diversion rates or operation of the facility over multiple days. The former is impractical as analogous data, and the latter is impossible due to lab restrictions on working hours. Therefore, a lower threshold is used for the mock facility. In essence, this becomes a conservative estimation; by predicting the diversion of a smaller quantity than the significant quantity, a successful prognostic model would be more sensitive and would suggest a need for action well in advance of the actual generation of a mock significant quantity.

Three candidate models were tested to determine their effectiveness in predicting the time to significant diversion of material: Markov Chain, General Path, and Particle Filtering. In each case, the MUF was calculated as the prognostic parameter with the assumption that an increasing magnitude of MUF was the direct result of diversion. A threshold MUF quantity of 0.75 kg was set for preliminary model testing. (This threshold allowed for operation of the facility with slow and fast diversion rates while still keeping the hours of consecutive operation of the facility within acceptable lab limits.) The models were tested for progressively increasing amounts of process data to monitor the evolution of the predictions over time. The prognostic models were evaluated on how quickly they effectively determined whether a diversion was occurring, how quickly the model converged on a prediction of the correct time until a significant quantity was diverted, and how confident the model was in its prediction (i.e. the tightness of the confidence intervals of the prediction).

4.4.1 Definition of the Prognostic Parameter

The cumulative inventory difference calculation provided by *PlotEvents* was a direct calculation of the MUF of the system, as shown in Figure 23.




Figure 23 - Inventory Difference for a Legitimate Run

When the mock enrichment process started, the MUF value increased (seen in Figure 23 as a decrease in the cumulative inventory difference value), but remained steady throughout the process. At the end of the process, the MUF value decreases. The initial increase in MUF was a consequence of the manual operation of the facility and the lag time of the PI controller in responding to the system startup. Likewise, the final decrease in MUF due to system shutdown was a function of the PI controller lag in responding to system shutdown and the manual closing of the cutoff valves of the system. In both cases, the change in MUF was a perfectly legitimate part of the operation of the facility and therefore needed to be removed from the inventory difference trace. This was accomplished by removing the trace prior to the stabilization of the PI controller. Additionally, the trace was re-centered at an inventory difference value of 0 kg at the origin of the truncated trace, and the sign was reversed such that an increasing inventory difference trace corresponded to an increasing MUF value rather than a decreasing MUF value. The new trace for the inventory difference data in Figure 23 is shown in Figure 24. Here, the time scale is referred to as "time steps", where each time step equals one second.

By conditioning the data as in Figure 24, the inventory difference trace would always begin at a value of zero kg and not be influenced by the initial startup of the mock F&W facility.

4.4.2 Markov Chain Model

A Type II Markov Chain model was developed by defining a transition probability matrix (TPM) using data from the start of the prognostic trace up to the time of the prognostic observation. For example, if a prognostic analysis was performed at the 400th time step, the transitions from the 0th through the 400th time step were used to create the TPM. The TPM then provided a sense of the behavior of the model for the forecasting of the trace into the future. The TPM was defined with four possible states, as given in Table 3.




Figure 24 - Modified Inventory Difference for a Legitimate Run

Once the TPM was defined, Monte Carlo simulation of the process was performed for 100 predictive traces, and the results of the Monte Carlo simulation were used to build statistics on the mean and standard deviation of the time to diversion of a significant quantity. A threshold of 10,000 time steps was set to discriminate between a prediction of diversion or legitimate operation; if the significant quantity of 0.75 kg was not predicted to have been diverted by 10,000 time steps, then the model assumed that diversion was not occurring.

A representative result of the Markov Chain analysis is shown in Equation 63 and Figure 25 for a diversion run after 800 time steps have passed. Here, the significant quantity was reached at 950 time steps, so the prognosis was queried late in the diversion scenario when most of the significant quantity had already been diverted.

$$TPM = \begin{bmatrix} 0.1280 & 0.1220 & 0.1890 & 0.5610 \\ 0.1779 & 0.1718 & 0.2270 & 0.4233 \\ 0.2363 & 0.2143 & 0.2582 & 0.2912 \\ 0.2457 & 0.2595 & 0.2318 & 0.2630 \end{bmatrix}$$

Equation 63 - Probability Matrix for Fast Diversion after 800 Time Steps

In Figure 25, the data to the 800th time step is shown in red, the true data beyond the 800th time step is in black, and the Monte Carlo simulations are in blue. Most of the Monte Carlo simulations reach the 0.75 kg threshold very quickly and cannot be individually seen. A few, however, exhibit strong negative trends initially, and then take many time steps to finally provide a prediction of significant diversion. While the Markov Chain model did correctly predict diversion and provide a reasonable estimate of the mean time to diversion (as seen by the distribution estimate – the red line in Figure 25), the variance was so great that the prediction could not be considered more than an assertion that diversion was in fact occurring. That is, the Markov Chain model's time to significant diversion was not reliable due to the magnitude of normal process noise and variability, even with late-term diversion predictions and Bayesian updating of the TPM.

Table 3 - State Definitions for the Transition Probability Matrix

State	Definition (kg/min)
1	<-0.01
2	-0.01 to 0
3	0 to 0.01
4	>0.01




Figure 25 - Markov Chains for Fast Diversion after 800 Time Steps

4.4.3 General Path Model

The General Path model began with a preconceived prediction of time to significant diversion by developing a linear trend of diversion from eight training runs known as a prior model. The prior model was created by linear regression fitting to the eight data sets. The eight training runs and the resulting General Path model are shown in Figure 26.

Fitting the linear function $Degradation = a * Time\ Cycle + b$, the model coefficients were estimated to be $a = 0.00211 \pm 0.000341$, and $b = 0.0690 \pm 0.173$. Because the value $b = 0$ was well within a standard deviation of the parameter mean, the parameter b could be neglected as insignificant to the model. The resultant prior model is given in Equation 64.

$$Material\ Diverted\ (kg) = 0.00211 * (Time\ Steps)$$

Equation 64 - Prior Function for the General Path Model

The prior equation provided a starting estimate for the time to diversion. With the General Path model, the prior was updated by the data from a run in progress according to Bayes Theorem. The Bayesian updated model was then used to estimate the time until significant diversion (or to predict that diversion was not occurring if the time to significant diversion was greater than 10,000 time steps). For the diversion run analyzed with the Markov Chain model, the prognosis estimates are shown for predictions after 200, 400, 600, and 800 time steps in Figure 27 and Table 4.

The prediction of time to significant diversion from the General Path model converged toward the correct solution of 950 time steps, though the prediction never came within one full standard deviation, even after 800 time steps had passed. Still, the General Path model clearly outperformed the Markov Chain model, yielding a far more reliable prediction of time to significant diversion.

A case with no diversion is shown in Figure 28. In no case did the predicted time to significant diversion occur prior to 10,000 time steps, a value that suggests that diversion is not present in this run. The tendency of the prediction to "wander" between positive and negative rates of diversion was a function of the noise in the data




Figure 26 - General Path Model Prior Function Development




Figure 27 - General Path Model Predictions for Diversion

Table 4 - General Path Estimates of Time to Significant Diversion

Start of Prognostication (Time Steps)	Time to Significant Diversion From Time = 0	
	Expectation (Time Steps)	Standard Deviation (Time Steps)
(Prior Model)	3555	2196
200	1401	42
400	1090	20
600	974	15
800	966	13




Figure 28 - General Path Model Predictions for a Legitimate Run

4.4.4 Particle Filter Model

The particle filter model was applied using the techniques described in §3.2.5. Similar to the GP model, the particle filter was applied to diversion and non-diversion case studies and the predictions were compared to the known diversion rates. For the diversion case used in the GP model, the PF model provided the predictions and uncertainties given in Table 5. The critical MUF value was again 0.75 kg, and the actual time to SQ production was 79:10 minutes.

400 total particles were used in the PF model, with initial paths varying between -0.06 and +0.06 kg/min. Updates were allowed after 125 seconds from the initial distribution or after any SIR redistribution. (This equates to every 25 time steps, since measurements were available once every 5 seconds.) The 125 second delay in SIS and SIR updating allowed the particles to travel and spread, reducing the chance that high-likelihood particles would be demoted early in the process when all particles are close together.

As with the GP model, a threshold of 13 hours and 53:20 minutes was set as the outer limit for a prediction of diversion. If the PF model did not predict 0.75 kg of diversion after 13 hours and 53:20 minutes (2 hours and 46:40 minutes total time steps beyond the known data), then no prediction of diversion would be made. This limit censored predictions of critical MUF production due to negligibly low model coefficients.

The particle filter predictions are shown in Figure 29 through Figure 32. In Figure 29 and Figure 30, the particle filter is inhibited by the early "settling in" phase where the initial particle paths are uncorrelated to the actual MUF observations. The prediction variances are sufficiently high to warrant suspicion about the predicted time to SQ diversion, but the trend is apparent. By 50 minutes in Figure 31, the particle variance has reached relative equilibrium with the measured data. The prediction of 80:07 minutes to SQ production is within the margin of error of the actual time to SQ production of 79:10 minutes. The result is similar at 66:40 minute in Figure 32; the prediction of 4802 seconds with a standard deviation of 162 seconds is very reasonable, given the noise and the relatively constant increase in MUF. Figure 33 shows the particle paths for the 66:40 minute prediction case (other cases are not shown as the particle paths are subsets of the particles in Figure 33). The effect of particle weights is not readily apparent, as the particle paths are all illustrated with equal line weight in Figure 33.

Table 5 - Particle Filter Estimate of Diversion		
Start of Prognostication (Seconds)	Time to Significant Diversion From Time = 0	
	Expectation (Seconds)	Standard Deviation (Seconds)
1000	18,600	2907
2000	7388	464
3000	4807	163
4000	4802	163




Figure 29 - Particle Filter Prediction at 16 Minutes, 40 Seconds




Figure 30 - Particle Filter Prediction at 33 Minutes, 20 Seconds




Figure 31 - Particle Filter Prediction at 50 Minutes




Figure 32 - Particle Filter Prediction at 66 Minutes, 40 Seconds




Figure 33 - Particle Paths for the Diversion Trial Run

4.4.5 Prognostic Model Selection

Of the three prognostic models tested, the particle filter model was chosen for inclusion in the *PlotEvents* software for monitoring the mock F&W facility. The Markov Chain model was clearly unreliable as a prognostic model due to the high noise of the system. Also, the inability of the Markov Chain model to recognize the difference between a diversion and a nondiversion (as well as the time that a diversion may begin) rendered it useless as a prognostic model for determining whether a diversion event may be taking place.

The general path model performed well, but suffered from a lack of an internal mechanism for piecewise regression. Additionally, the error estimates from the general path model tended to be too small, resulting in overconfident predictions of diversion that were not within the margin of error from the actual time to SQ production. (This tendency also caused the GP model to appear to predict long-term diversions when no diversion was present, as the low variance of the prediction was easily interpreted to be a high confidence in the prediction rather than simple agreement among the multiple paths of the model.)

The particle filter model, on the other hand, was a piecewise model by design (where the "pieces" may be viewed as the regions between SIR updates) and naturally changed trends based on changes in the data. Additionally, the particle variance was directly influenced by the measurement and dynamic noise of the system through the SIS and SIR updates, respectively (see Appendix B: Measurement and Dynamic Noise Estimates).

The particle filter model did express two limiting attributes. First, the particle variances did not tend to settle out for about 33:20 minutes. This was partly a function of the high data noise and the inability of a model to make reliable prognostics without sufficient data to overcome the noise, but the convergence time is a necessary consequence of the initial particle spread. (The initial particle spread is in turn a consequence of the lack of foreknowledge of the MUF path at the time the particle filter is initiated.) Second, the particle filter exhibited a limiting factor of noise reduction. With both the 50 minute and 66:40 minute cases, the prediction had a standard deviation of 164 seconds. This is a consequence of the particle filter's reliance on noise to estimate the filter parameters.

The inherent flexibility and the self-regulated variance estimates were the primary factors for choosing the particle filter model. The higher variance of the particle filter (relative to the GP model) was an

admission that the noise of the system (both measurement and dynamic) limited the reliability of any prediction made using the data.

4.5 Case Studies

The particle filter prognostic model was applied to four unique case studies: legitimate operation, slow diversion, fast diversion, and an interrupted diversion. The legitimate operation tested the model's performance when no prediction of diversion should be made. The two diversion cases tested the model's ability to adapt to different rates of diversion. The interrupted diversion case featured a pause in the diversion process, then a continuation of the diversion. All four cases are explained in more detail in Appendix C: Case Studies.

For each case study, the particle filter was applied at 16:40 minutes (1000 seconds) after the initiation of the feed flow. This delay in filtering allowed the PI controller time to stabilize after the initial perturbation of feed flow. The cumulative inventory difference at the beginning of particle filtering was defined as the "zero point" for the inventory difference. For the non-diversion and the slow diversion cases, the significant quantity threshold was defined as 0.8 kg. For the fast diversion and the intermittent diversion cases, the significant quantity threshold was defined as 2.0 kg. The difference in threshold allowed for an accurate, known time of SQ production for the slow diversion case while also providing more run time for the particle filter in the face of faster diversion rates. (If scaled properly from the IAEA standard of 75 kg of U-235, the SQ threshold would normally be about 20 kg of material; however, this size of a diversion would require running the facility longer than laboratory regulations allowed. The threshold was therefore set to lower values. Fortunately, the threshold is a matter of scale, not of logistics, with particle filtering.)

4.5.1 Legitimate Operation (No Diversion)

When applied to the non-diversion scenario, the particle filter never predicted the diversion of a significant quantity throughout the available range of data. The weighted mean and weighted standard deviation can be seen during the range of data in Figure 34. The instantaneous jumps in the weighted particle mean occur when the particles are resampled and are a function of the programmed delays in weight updating after resampling. When the particles are used for prediction of SQ production, no such prediction is made by the time limit of 13 hours and 53:20 minutes (50,000 seconds), as is seen in Figure 35. Only a fraction of the particles have exceeded the 0.8 kg SQ threshold, resulting in the skewed distribution of times at which particles crossing the SQ threshold. The particle traces can be seen in Figure 36, again revealing a prediction of no SQ production. Also in Figure 36, the ICOMP measure of model validity indicated that the linear model became less valid than a constant-MUF model over time. This is indicated by the yellow shading over the SQ prediction alarm indicator, which shows a preference for the constant model.

Finally, a look at the particle filter predictions at every five second time interval indicates that there was never a prediction of diversion at any time during the legitimate run, as shown in Figure 37. An occasional prediction of SQ production could have been acceptable as false positives due to the noise inherent in the data, but in this case, all signs indicated that facility did not have an unaccounted loss of mass throughout this run.




Figure 34 - Weighted Mean Particle Trace for Non-Diversion.





Figure 35 - Particle Filter Statistics for Non-Diversion after 66 Minutes, 40 Seconds.



Green areas indicate that the linear model is best.
 Yellow areas indicate that a constant MUF assumption is best.
 Red areas indicate that linear models are no better than random models.

Figure 36 - Prognostic Particle Traces for Non-Diversion after 66 Minutes, 40 Seconds

Table 6 - ICOMP Scores for Non-Diversion Run

		Model Type		
Start Time	Stop Time	Linear	Constant	Random
0	600	-282.61	-268.88	-107.81
600	1200	-297.99	-295.48	-110.69
1200	2800	-582.20	-556.82	-168.15
>2800		-692.50	-701.35	-281.81




Figure 37 - Predictions for Non-Diversion up to 66 Minutes, 40 Seconds

4.5.2 Slow Diversion

The threshold for SQ production for the slow diversion case was set to 0.8 kg and was reached at 85 minutes after the initiation of the particle filter. When the particle filter is initiated, the particles require about 10 minutes to learn the data and to begin to track the MUF calculations. At 10 minutes, the first SIR process occurs and the particles begin to track the MUF based on Bayesian updating of particle trajectories rather than the initial random motion of the particles. The weighted mean and confidence intervals of the particles may be seen in Figure 38, along with the statistics for the prediction of time to significant quantity production after 66:40 minutes of particle filter operation. In this case, the particles predict SQ production of 0.8 kg at 87:02 minutes after the initiation of the particle filter, with a standard deviation of +/- 172 seconds, or 95% confidence intervals of [81:18 92:46] minutes. Viewing the data transversely, the second histogram in Figure 38 indicates a MUF at 5 seconds of 0.83 kg with a standard deviation of 0.03, or 95% confidence intervals [0.77 0.89] kg. The time to SQ production statistical measures passes the Lillieford test for normality while the MUF at SQ time fails to reach the $p = 0.05$ threshold, indicating that the particles are tracking well along the mean of the data and have stabilized with only minor non-normality, but that there may be outliers within the data.

The particle traces are shown in Figure 39. Only three total resampling instances were required to overcome degeneracy of the particles. It can be seen, however, that the particles required a considerable amount of time (the second resampling at 30 minutes) to converge to a uniform estimate of time to diversion. In the presence of high measurement noise and a slow diversion rate, such a slow refinement of the particle traces is hardly surprising. Also in Figure 38, the alarm indicator showed predictions of impending SQ production for all times after 10 minutes (i.e. after the first SIR sequence updated the initial random particle trajectories). For all particle sampling periods, the ICOMP measure indicated that the linear model was the best performer for the data. This is also seen in Table 7 where the ICOMP scores overwhelmingly favor the linear model except for the first region of interest. Here, the linear and constant model scores are too close to distinguish. (This is also the region where no prediction of SQ production is made.) Figure 40 shows the predictions at every 5 second measurement interval, with the actual time to SQ production falling within the confidence limits after 46:40 minutes have elapsed.




Figure 38 - Particle Statistics for Slow Diversion




Figure 39 - Particle Traces for Slow Diversion

Table 7 - ICOMP scores for Slow Diversion

		Model Type		
Start Time	Stop Time	Linear	Constant	Random
0	600	-443.71	-443.12	-424.13
600	1800	-805.67	-597.74	-224.78
1800	2800	-438.67	-330.56	131.24
>2800		-837.90	-681.55	370.63




Figure 40 - Predictions of SQ Production during Slow Diversion

4.5.3 Fast Diversion

The fast diversion scenario featured a diversion rate nearly twice that of the slow diversion. The diversion of 2 kg of material took about two hours, or 120 minutes. The diversion rate was constant and, other than noise from switching feed and tails tanks, the data was largely devoid of features.

Unlike the slow diversion case, the particle filter converged onto the measured data quickly, with the weighted 95% confidence intervals of the particles effectively converging after approximately 20 minutes, as seen in Figure 41. The particle tracking stays well-centered on the measured data, and the prediction of SQ production that is taken at 83:20 minutes yields a prediction of 149:25 minutes to SQ production with a standard deviation of 214 seconds and 95% confidence intervals of [142:18 156:32] seconds. The particle traces in Figure 42 indicate that resampling was only necessary twice in the first 83:20 minutes, once at 10 minutes (due to the particles learning the initial data) and again at 43:20 minutes. This low resampling rate suggests that the particles are not rapidly suffering degeneracy (see Equation 37), which means that most particles are representing the measured data well and the particle weights are not being strongly favored to a few particles.

In Figure 42, the weighted mean of the particles appears to be slightly skewed in favor of the slower SQ production rates. Still, the ICOMP scores suggest that the linear model is best for all regions of the data, and once the particles resample in favor of the data instead of the initial random trajectories, the SQ production alarm instantly signals the sensed loss of material from diversion. Table 8 shows the strong preference for the linear model, which is not surprising considering the linear relationship in the data.

In the fast diversion run, production of 2 kg occurred about 120 minutes after initiation of the particle filter. In Figure 43, the particle prognostics converge on the correct time to SQ production quickly with a prediction of approximately 120:50 minutes with 95% confidence intervals of [106:40 135] seconds, but then deviate in the presence of the noise from the tank switchovers. As the particle estimates converge, however, the effect of measurement noise can be seen as the prediction diverges after the second SIR update at 43:20 minutes. While still predicting that diversion will produce 2 kg within 13 hours and 53:20 minutes, the time itself is errant at roughly 8 hours and 53:20 minutes and slowly converges on the correct time to significant quantity diversion over time. The perturbation is not surprising; in a high-noise environment, small deviations can produce large changes in predictions. This effect is most apparent when the particles have converged to within measurement noise but still have a long time until the significant quantity is reached, as is seen at 43:20 minutes in Figure 42.




Figure 41 - Particle Statistics for the Fast Diversion




Figure 42 - Particle Tracks for the Fast Diversion

Table 8 - ICOMP Scores for Fast Diversion

		Model Type		
Start Time	Stop Time	Linear	Constant	Random
0	600	-412.39	-380.35	-179.01
600	2600	-1440.57	-229.93	564.87
>2600		-1250.84	-537.97	1304.62




Figure 43 - Predictions of Significant Quantity Prediction for Fast Diversion

4.5.4 Intermittent Diversion

The final case featured a diversion event that began approximately one hour after the start of the run, and included a small cessation in the diversion flow about 45 minutes later. The entire run lasted approximately two hours. During the diversion period, a total of 1.891 kg of water was diverted from the system; with the flow rate maintained during the active diversion, continuation of the run would have resulted in 2 kg of water diverted at approximately seven minutes later, or about 126:40 minutes from the start of the run.

4.5.4.1 Results Prior to the Pause in Diversion

For this run, two separate evaluation periods were observed, once at 10 minutes and once at 120 minutes, allowing a more detailed observation of the particle filter shortly before and after the pause in diversion. The particle statistics are shown at time equals 10 minutes in Figure 44, and the particle traces are in Figure 45. The particle statistics indicate some discrepancy between the particles and the measurements, as the predicted time to 2 kg of MUF shows some right skewness in the weighted histogram in Figure 45. The confidence intervals on the time to SQ production and on the amount of SQ production at 142:26 minutes are relatively large, in part due to the high dynamic noise of the system which has necessitated an increase in the dynamic noise estimate. In other words, the rapid shift from a non-diversion regime to diversion introduced a large change in the system dynamics, which was reflected in the particle filter by an increased variance in the first-order coefficients of the linear model. This adjustment allowed the particles to track the more rapid changes in the data at a cost of long-term precision in the prediction.

In the particle traces and SQ production alarm indicator in Figure 45, a brief alarm was raised at around time equals 16:40 minutes but disappeared throughout the remainder of the non-diversion regime. This temporary alarm was the result of the particles adjusting to the dynamics of the data. But even with this alarm, the ICOMP scores indicated that the constant model was better suited to the data than the linear model throughout most of the non-diversion period. This combination of events leads to a conclusion that diversion is not likely happening during the first half of the run.




Figure 44 - Particle Statistics for Intermittent Diversion at Time = 60 minutes.




Figure 45 - Particle Tracks for Intermittent Diversion at Time = 60 minutes.

4.5.4.2 Results after the Pause in Diversion

In Figure 46, the effect of the pause in diversion can be clearly seen on the particle filter response. The particles overshoot the measured data, and the subsequent SIR process causes the particles to track in a negative direction for a short time. The large change in particle response again increases the dynamic noise estimate by roughly a factor of two (1.15^5 , or 2.01 due to five total dynamic noise estimate increases). Again, the statistics belie a large uncertainty in the consistency of the data trend. The estimated time to SQ production is 151:45 minutes with a 95% confidence interval of 7:56 minutes, but the weighted particle estimate is largely skewed and the lack of normality indicates that the particles are not well centered on the data.

The particle traces in Figure 47 illustrates the skewness of the particles. A clear majority of the particles predict a slower rate of MUF production than the weighted mean, but since the measured data favors a faster MUF production rate, the slower ascending particles hold very little weight, creating the long tail seen in the histogram in Figure 46. During the pause in diversion, the significant quantity alarm indicator did not suggest a diversion within the allotted time frame as seen by the zero value in Figure 46 from about 96:40 minutes to 103:20 minutes. During this time, the ICOMP scores favored the constant model over the linear model as well, as shown by the yellow highlighting for the same time frame. The ICOMP scores for each region are shown in Table 9. Even though the high dynamic noise of the data caused large uncertainties in the model, the random model always received much higher ICOMP scores than the linear and constant models, indicating that the data had a coherent structure and that the noise was not purely random. Also notable in Figure 47 and Table 9 is the large number of SIR updates. With high dynamic noise in the system, the particle trajectories are spread more widely than in the slow and fast diversion cases (which were constant rate diversions with low dynamic noise); more particles inevitably deviate from the measured values, increasing the degeneracy rate and requiring more resampling over time. The high dynamic noise also increases the uncertainty of the predictions, reflecting the reduced expectation that any diversion would remain constant over time.




Figure 46 - Particle Statistics for the Intermittent Diversion Run at Time = 125 minutes.




Figure 47 - Particle Traces for Intermittent Diversion at Time = 125 minutes.

Table 9 - ICOMP Scores for Intermittent Diversions

		Model Type		
Start Time	Stop Time	Linear	Constant	Random
0	600	-308.53	-303.88	-199.84
600	800	-67.22	-51.95	-50.78
800	1200	-251.05	-264.44	-49.36
1200	1800	-318.63	-325.83	-52.92
1800	3000	-620.05	-621.17	-112.01
3000	3600	-315.95	-182.01	-29.62
3600	4200	-332.35	-174.37	-167.15
4200	5200	-560.72	-251.38	229.57
5200	5800	-262.66	-199.59	270.27
5800	6200	-226.18	-238.03	205.06
6200	6800	-365.76	-212.91	358.28
>6800		-414.43	-171.40	482.09

Figure 48 shows the predictions of SQ production for all measurement times. Again, an early prediction of SQ production was made when the particle filter was still "settling in" to the data, but the lack of an actual diversion event precluded the signal from sustaining over a significant period of time. The rapid change in regime from non-diversion to diversion can be seen around 60 minutes, where a long term prediction first begins. The unsteadiness of the prediction mirrors the dynamics of the data, though much of the diversion predictions do include the actual 2.0 kg diversion time within their estimated confidence intervals. The large increases in the error ranges are the result of the increased dynamic error estimates

In qualitative terms, the intermittent diversion agrees with the intuitive inference that, in the face of a dynamically unstable process, any prediction of a time to significant quantity diversion should be met with skepticism. With MUF trends like Figure A26, it is clear that the facility operation cannot be considered steady. Whether by natural dynamic variance of the enrichment process or (more likely) by willful changes in the enrichment operating regime by the operator or a malevolent third party, the assumption that trends will continue "as they are" into the future is suspect. In this condition, uncertainty estimates to the predictions *should* be large. This is especially true after 103:20 minutes, when the uncertainty of the particles has grown considerably due to the dynamics of the pause in diversion. Convergence of the particles (and, by extension their relative precision) requires stability in the data. The particle filter process automatically increased uncertainty in the face of increasing system dynamics, a feature that stands in contrast to initial testing of the general path model, which tended to be overly precise (see Section 4.4.3).




Figure 48 - Predictions of Significant Quantity Production during Intermittent Diversion

4.5.5 Case Studies: Concluding Remarks

In the broadest view, the particle filter successfully identified periods of increasing MUF related to diversion and periods where the MUF of the system was relatively stable. During legitimate operation (seen in both §4.5.1 and the first half of the data of §4.5.4), no prediction of imminent SQ production was made, and the ICOMP measure of model validity suggested that a linear model with increasing MUF did not explain the data as well as a constant model. When diversion was present, the particle filter successfully recognized the event and the linear model provided a projected time to the production of a defined SQ of material with error bounds regulated by the measurement noise of the data, the dynamics of the data, and the length of time until the estimated production of the significant quantity.

Though the particle filter was constrained to a linear model of potentially very nonlinear data, the piecewise solutions generated as a result of the sequential importance resampling allowed the prognostic model to successfully track changes in trend of the MUF. By allowing dynamic noise updating and recalculation of the resampling procedure when the particles failed to adequately fit the measured data (see Appendix B: Measurement and Dynamic Noise Estimates), the particle filter could effectively track sudden changes in MUF trends without requiring unnecessarily high dynamic noise estimations during relatively stable periods.

The results of the particle filter predictions cannot be interpreted beyond their assumptions. With a linear particle filter model, the assumption behind any prediction is that the data will continue in the same linear trend as the particles have identified. For example, Figure 48 shows several vastly different SQ production times for the intermittent diversion case, even when the diversion rate is relatively stable. The effect is particularly noticeable during the beginning of diversion because the SIS and SIR particle updates are retraining the particles to the new operating regime, and the distance between the current MUF value and the critical MUF value allow small changes in trend to cause large changes in predictions. Simply stated, the particle filter predictions may be viewed as: if things continue as they trend right now, how much time until an SQ is produced? Even then, further details, like the validity of the linear model over the constant model (i.e. the ICOMP scores) and how well-centered the particles are upon the measured data (the histograms seen in Figure 33, for example) are indicators of the viability of the prediction.

In summary, the particle filter provided reliable prognostic indication of the potential production of a significant quantity of undeclared material from the mock F&W facility. Unlike traditional health

monitoring prognostic applications, no prediction of time to SQ production can ever be considered wholly valid; for example, if diversion is occurring, the simple decision to cease undeclared production of material would render any previous predictions invalid despite the model's inability to account for such a decision. In this light, the prediction itself is best viewed as an estimate of the results of the trends of the current operating regime.

4.5.6 Limitations on Particle Filter Reliability

One core assumption with particle filters is that the target moves through its system space in a continuous manner. In the case of this research, it is assumed that the MUF approximation does not change discretely (i.e. a step change in value). A significant discrete (or near-discrete) change to the measured MUF approximation could potentially move the MUF measurements outside the probability distribution defined by the particle positions and weights. In this case, the particle filter would respond in one of two ways. First, if the measurement values lie outside the particle range but still close enough to preferentially weight closer particles, the particle filter would recognize the skewed probability distribution and adjust the dynamic noise estimates and bias the particle trajectories toward the measurements. However, the increase in the dynamic noise estimate and the bias would be so great that the particles would spread over an excessively wide area, resulting in excessively high uncertainties for any prognostic estimates. The second possible response to a discrete change in MUF occurs if the new measurements are not close to the particle distribution. In this case, *all* particles would have a near-zero likelihood of representing the measured values and the SIS step would not be able to meaningfully alter the particle weights. The weights might remain unaffected, or they may vary randomly (and chaotically). In any case, the particles would be "lost" in the state space.

If the prognostic technique proposed here is kept in context with the overall process monitoring regime, such a discrete change in MUF value would likely be cause for alarm from the monitoring algorithms. Such large and rapid changes in unaccounted material would be a severe departure from normal operating conditions and would likely create a signal that further inspection may be necessary to understand the issue. If so, then the particle filter's results need not be considered; the monitoring system will have triggered the necessary concern from the inspectors. If the monitoring system does not provide an alarm to a discrete change in MUF, then the particle filter would have to be "reset" by recognizing the discrete MUF change, re-initializing the particles on the new data center, and starting the tracking process over again. It should be noted as well that a discrete change in MUF would impair

any prognostic method, as MUF is typically a continuous variable and no prognostic method currently exists to anticipate discrete changes in normally continuous target locations.

Other methods of defeating a particle filter system would typically attempt to increase the measurement or dynamic noise to excessive values. For example, establishing a significant sinusoidal fluctuation in the MUF would require a high dynamic noise estimate to keep the particles trained on the measurements. In this case, a subsequent attempt to divert – even utilizing a constant-rate diversion scheme – would only produce predictions of SQ diversion with excessively high uncertainties. By the time the dynamic noise was trained down to suitable estimates, a significant quantity may have already been diverted. As with discrete changes in MUF, excessive variation in MUF may be noted by tracking algorithms and flagged as an indication that the process data requires attention. Alternatively, an alarm might be triggered if the dynamic noise estimate within the particle filter exceeds some factor of its current estimate. If, for example, the dynamic noise estimate that is five times the normal estimate may serve as an alarm.

5 Conclusions

Traditionally, the role of the inspector in international safeguards has been to look in the rear view mirror and make sense of the past. *Has undeclared material been produced? If so, how much? When did it happen?* If the next-generation inspection practices include remote monitoring techniques, however, the opportunity exists to provide forward-looking conclusions from the data: *If undeclared material is being diverted, when will a specified quantity be produced?* Answers to questions like this could provide useful to a fully information-driven inspection scheme where on-site visits by inspectors are driven by data rather than by predefined periodicity. If reliable forecasts of enrichment processes can be made, inspection frequencies could potentially be reduced even further, resulting in increased cost savings for the inspecting agency and for the facility operator.

This research explored the use of prognostics methods to predict the time that a significant quantity of material may be diverted from an enrichment process. After preliminary testing of various predictive methods, a particle filter technique was chosen for application to four case studies: a non-diversion process, slow diversion, fast diversion, and an intermittent diversion scenario. The four cases tested the particle filter for false positives, false negatives, adaptation to low or high dynamic noise, and nonlinear trends within the MUF calculations.

The particle filter utilized Bayesian updating of measurement and dynamic noise estimates to best approximate the characteristics of the data itself. Additionally, checks were built into the resampling procedure to verify that the particles' trajectories remained close to the measured data and to provide error correction if the particles failed to accurately represent the data. These checks increased or decreased the dynamic noise estimates based on the variability of the data trends (i.e. how quickly the trends in the data were observed to change) and prevented sudden, drastic changes in operation (e.g. the intermittent diversion) from defeating the particle filter's data tracking.

The data for this research was generated on the mock feed and withdrawal facility, a scaled-down analog of a gas centrifuge enrichment plant that uses water as the process medium. The mock F&W facility is scaled roughly 1:100 in both time and space to a GCEP so that 1 kg of material in the mock F&W facility is comparable to 100 kg of uranium hexafluoride and a similar scale in time. Using the mock facility, load cell measurements of feed, tails, and product cylinders can be recorded as functions of time.

The first case study featured fully declared operation of the mock feed and withdrawal facility (i.e. without diversion). At no time over the 66:40 minutes of available data did the particle filter predict the diversion of a significant quantity of material within a window of 13 hours and 53:20 minutes after the final measurement. In the second case study, a diversion of 0.8 kg over the course of 85 minutes was introduced. After 10 minutes (the time required for resampling of the particle from their initial random trajectories), the particle filter successfully identified the continual increase of unaccounted material as a possible diversion. After 46:40 minutes of measured data, the particle filter predicted the diversion of 0.8 kg by 84:10 minutes with a standard deviation of 172 seconds. From 46:40 minutes to the end of measured data at 66:40 minutes, the particle filter predicted the correct time to diversion of a significant quantity within 95% confidence intervals. Because the diversion was conducted at a constant rate, the convergence of the particle filter was expected.

The third case study featured the diversion of 2 kg over 120 minutes. The particle filter identified the diversion at the first resampling at 10 minutes, predicting that 2 kg would be diverted by 121:40 minutes with 95% confidence intervals of approximately [110:50 132:30] minutes. The prediction remained accurate with 95% or better confidence until 2800 seconds, when the resampled particles predicted diversion of 2 kg by 8 hours and 53:20 minutes. The prediction converged toward the correct time to significant quantity diversion, and by 83:20 minutes the particle filter predicted 2 kg of diversion by 149:25 minutes with 95% confidence intervals of [142:18 156:32] minutes. The loss of accuracy in this prediction revealed the sensitivity of the particle filter to the measurement noise.

Finally, the fourth case study introduced nonlinearity to the problem. The diversion was initiated 60 minutes after initiation of the particle filter, paused for approximately five to ten minutes at about 100 minutes, then continued for the rest of the run. A total of 1.89 kg had been diverted when the run was terminated, but 2 kg of material was expected to be diverted at approximately 135 minutes after the initiation of the particle filter. At 100 minutes, the particle filter predicted a time to 2 kg of diversion of 142:26 minutes with 95% confidence intervals of [125:02 159:50] minutes. If the actual time to diversion of 2 kg were adjusted by 5 minutes to 130 minutes to account for the pause in diversion, the prediction at 10 minutes fell well within the 95% confidence interval. After the pause in diversion, the prediction was again evaluated at 125 minutes and the particle filter predicted a time to 2 kg of diversion by 151:45 minutes with 95% confidence intervals of [135:52 167:37] minutes, which does not quite capture the time to diversion of 135 minutes. In Figure 47, the particle filter estimate appears to

lag the measured data, a consequence of the reduced sensitivity of the Bayesian updating of particle weights due to the measurement noise. Additionally, the high dynamic noise of the system due to the pause in the diversion caused large error estimates in the prediction. Several dynamic noise estimate updates were observed at the start of diversion and when the pause in diversion occurred, indicating that the dynamic noise of the system was much greater at these times than during the first 60 minutes when no diversion occurred. Such changes in dynamic noise in the system could potentially be used as indicators of changes in the operation of the facility, but any correlations between the dynamic noise and operational behavior were not pursued in this research.

It is particularly important to note that the utility of a prediction of the time to diversion of a specified amount of material is only valid within the context of the assumptions of the model itself. In this research, a linear particle filter was used for two reasons. First, the most likely scenario for protracted diversion is that of a constant removal rate of material. Since enrichment facilities are typically run at a constant processing rate to maximize efficiency and minimize wear of the centrifuges due to angular acceleration stresses, a protracted diversion would most likely be of a constant rate as well. Second, there is no a priori justification for a more sophisticated model for diversion rates (e.g. polynomial, exponential, or hyperbolic), so a linear diversion rate assumption is the most reasonable "first guess" for predicting time to diversion of a significant quantity. The consequence of the linear model is that any prediction assumes that operation of the enrichment process will continue on in the same manner as the last series of measurements. This leaves a qualifier to any prediction: *if trends continue as they are now, when would a significant quantity be diverted?* Therefore, it is unwise to treat the predictions as an absolute; instead, they should be viewed as a guideline that may influence decisions for inspection monitoring.

For any prognostic model for predicting time to diversion, understanding the interpretation of the uncertainty estimates is at least as important (and perhaps even more important) than the estimate itself. Diversion is, again, a process initiated and conducted by decision rather than by mechanical wear or degradation. There is no guarantee of monotonicity to a diversion trend, as temporary stops in diversion coupled with high noise and changes in holdup could indeed produce negative MUF trends. Additionally, there is no guaranteed pattern for diversion; while it is most efficiently performed at a constant rate due to the mechanics of gas centrifuge enrichment, diversion could be performed intermittently, or with time-variant rate changes. While the estimate is only as valid as its assumptions

(e.g. linear trends in linear vs. intermittent environments), the uncertainty estimates are more directly reflective of the data. Unexpected increases in uncertainty can be indicative of changes in operation, which has a great influence on the reliability of the prognostic model.

This research studied the application of health monitoring prognostic techniques to remote monitoring and verification of enrichment activities. A particle filter model was developed and tested against data generated from the mock feed and withdrawal facility at Oak Ridge National Laboratory. The accuracy of the predictions was measured at different time intervals, and the behavior of the model was analyzed based on these results. Bayesian updating of the measurement and dynamic noise estimates was automatically performed by the particle filter model, resulting in uncertainty estimates based on the inherent uncertainty in the data measurements and the relative stability of the trends in the MUF calculations.

5.1 Summary of Contributions

This research presented a prognostic method to predict the time until a significant quantity of material was diverted from the mock F&W facility at Oak Ridge National Laboratory using only process data from the feed and withdrawal stations. A Particle Filter method was chosen over Markov Chain and General Path models due to the Particle Filter's inherent flexibility to nonlinear data, lack of a Gaussian distribution constraint, inherent statistical dependence on measurement and dynamic noise, and the natural distribution and variance estimates of the particle predictions. While prognostic methods have been developed for health monitoring applications to predict the degradation of a component or system, this research represents the first effort to apply these prognostic methods to decision-based events.

In addition to the prediction of a time to future significant quantity diversion of material, the particle filter prognostic method in this research provided a thorough set of distribution and uncertainty measures to aid in understanding the reliability and significance of the prediction. Weighted particle distributions provided a measure of the certainty of the time to failure, and these distributions varied according to the measurement and dynamic noise estimates of the system. As noise sources increased in magnitude, the particle filter reflected the decreased predictability of the system with greater confidence intervals. Further, information complexity scores provided a measure of the suitability of the particle filter itself and demonstrated in clear quantifiable terms whether the model was a better fit to

the data than a constant-MUF assumption and whether the data held sufficient coherence to justify the model.

Measurement and dynamic noise estimates greatly affect the performance of particle filters. If these estimates are low, the particle distributions may not adequately describe the state space represented by the measured data, rendering any interpretations of the particle positions and trends inaccurate and meaningless. Likewise, excessively high noise estimates cause excessive variation in particle position and trajectory and generate high-uncertainty predictions that are likewise meaningless. It was seen in this research, however, that the dynamic noise of the MUF may be relatively low for long periods of time and high for brief periods when the operation of the mock facility is altered, such as when diversion is initiated or stopped. As discussed in Appendix B: Measurement and Dynamic Noise Estimates, data-driven methods were added to the particle filter to estimate the noise in the data and update the filter's noise estimates.

The mock feed and withdrawal facility was designed to be a fast, flexible, and reliable system for generating process station data analogous to that expected from a gas centrifuge enrichment plant. As a part of this research, several improvements were made to the mock F&W facility that allowed the simulation of degassing and sampling. These modifications emulated real and normal losses of material expected in a GCEP that could affect the material balance of the cascade area. In this research, these effects were identified in the data and accounted in the MUF calculations to remove their influence and provide a more accurate MUF-to-diversion estimate. A rules-based architecture was envisioned that would anticipate degassing and sampling, indicate an alarm if such events were missing, and account for their effects on MUF if these events were present.

5.2 Future Work

Application of prognostics to human intervention of mechanical systems is at best difficult. Unlike classical health monitoring applications, such as crack growth in plates, or bearing failure, events such as material diversion are the result of a decision to change the operating regime rather than a physical characteristic of a component. Better prognostic application in this field can only occur if the range of possibilities for human decisions can be better understood.

The prognostic model in this research was based on measurements taken every five seconds from the mock F&W facility. (The data was actually stored every second, but was then downsampled prior to use

in this research.) In real-world terms, a five second sample rate equates roughly to one sample every eight to ten minutes, based on a very general 1:100 scale in time. If the IAEA and enrichment facilities agree to remote monitoring, the data collection rate may be limited by agreement rather than the ideals of the inspectors and monitoring system designers. In this light, data collection at varying rates should be examined to determine their effect on the prognostics – or even if prognostics are feasible at the current state of technology. (This is particularly true for much slower collection rates, such as 1 data point per day in a GCEP, or about one point per 15 minutes in the mock F&W facility).

One limitation of the mock F&W facility is the lack of assay data that is present in enrichment processes. Since the water separated in the mock facility is not actually isotopically enriched, no correlation between product enrichment and tails depletion may be established. Without this knowledge, it is impossible to differentiate between legitimate non-diversion operations of the facility versus continuous replacement of diverted mass based on load cell data alone. If prognostics are applied to enrichment processes, research should be performed into incorporating assay information into the prognostic model to identify diversion masking events.

Most importantly, application of prognostic methods must be tested against full-scale data if forecasting the time to critical events such as the production of significant quantities of material is of interest to the IAEA (or other oversight agencies in similar scenarios). In the specific problem studied in this research, the particle filter method was adapted as a prognostic tool to estimate the effects of human intervention on a process, and the performance of the particle filter was tuned to match the characteristics of the data. For GCEP monitoring, a particle filter approach would have to include sufficiently accurate understanding of the measurement and dynamic noise terms in the data, as well as any relevant data features that might not be reproducible in the mock F&W facility.

REFERENCES

Abbas, M., Ferri, A., Orchard, M. E., & Vachtsevanos, G. J. (2007). An intelligent diagnostic/prognostic framework for automotive electrical systems. *2007 IEEE Intelligent Vehicles Symposium* , 352-357.

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control* , 716-723.

Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. *Second International Symposium on Information Theory* (pp. 267-281). Budapest: Akademiai Kiado.

Arulampalam, M. S., Maskell, S., Gordon, N., & Clapp, T. (2002). A Tutorial on Particle Filters for Online Nonlinear/Non-Gaussian Bayesian Tracking. *IEEE TRANSACTIONS ON SIGNAL PROCESSING* , 50 (2), 174-188.

Bogdanoff, J. L., & Kozin, F. (1985). *Probabilistic Models of Cumulative Damage*. New York: John Wiley & Sons.

Boyer, B. E. (2010). Analysis of the Effectiveness of Gas Centrifuge Enrichment Plants Advanced Safeguards. *51st Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Baltimore, Maryland: INMM.

Bozdogan, H. (2000). Akaike's Information Criterion and Recent Developments in Information Complexity. *Journal of Mathematical Psychology* , 62-91.

Bozdogan, H. (1988). ICOMP: A new model-selection criterion. In H. H. (Ed.), *Classification and Related Methods of Data Analysis* (pp. 599-608). Amsterdam: Elsevier-Science (North Holland).

Bozdogan, H. (1987, September). Model Selection and Akaike's Information Criterion (AIC): The General Theory and its Analytical Extensions. *Psychometrika* , pp. 345-370.

Bozdogan, H. (1990). On the Information-Based Measure of Covariance Complexity and its Application to the Evaluation of Multivariate Linear Models. *Communications in Statistics Theory and Methods* , 221-278.

Bozdogan, H., Howe, J., Katragada, S., & Liberati, C. (2009). *Misspecification Resistant Model Selection Using Information Complexity with Applications*. Catania, Italy: Italian Classification and Data Analysis Group (CLADAG).

Burr, T. (2008). Statistical Methods in Nuclear Nonproliferation Activities at Declared Facilities. In J. E. Doyle, *Nuclear Safeguards, Security, and Nonproliferation* (pp. 151 - 164). Oxford, UK: Butterworth-Heinemann.

Burr, T., Coulter, A., Hakkila, A., Kadokura, H., & Fujimaki, K. (1995). Statistical Methods for Detecting Loss of Materials Using NRTA Data. *Annual Meeting of the Institute for Nuclear Materials Management* (pp. 1032 - 1037). Palm Desert, Ca.: INMM.

Burr, T., Coulter, C., & Wangen, L. (October 1997). *Solution Monitoring: Qualitative Benefits to Safeguards*. Vienna, Austria: IAEA Symposium on International Safeguards.

Burr, T., Coulter, C., Howell, J., & Wangen, L. (2003). Solution Monitoring: Quantitative and Qualitative Benefits to Nuclear Safeguards. *Journal of Nuclear Science and Technology*, 40 (4), 256 - 263.

Byington, C., Watson, M., Edwards, D., & Stoelting, P. (2004). A Model-Based Approach to Prognostics and Health Management for Flight Control Actuators. *Aerospace Conference, 2004*. (pp. 3551-3562). Big Sky, MT: Proceedings of the IEEE Aerospace Conference.

Cadini, F. (2009). Particle Filtering for Diagnosis, Prognosis and On Condition Maintenance. Milan, Italy.

Cadini, F., Zio, F., & Avram, D. (2009). Monte Carlo-Based Filtering for Fatigue Crack Growth Estimation. *Probabilistic Engineering Mechanics*, 24 (3), 367-373.

Chau, P. (2002). *Process Control, a First Course with MATLAB*. New York: Cambridge University Press.

Chinnam, R. (1999). On-line Reliability Estimation of Individual Components, Using Degradation Signals. *IEEE Transactions on Reliability*, 48 (4), 403 - 412.

Cobb, D. D. (1981). Sequential Tests for Near Real-Time Accounting. *Journal of the Institute of Nuclear Materials Management*, 8 (2), 81-92.

Coble, J. (2010). *Merging Data Sources to Predict Remaining Useful Life – An Automated Method to Identify Prognostic Parameters*. Department of Nuclear Engineering. Knoxville: University of Tennessee.

Cochran, T. B. (1995). *The Amount of Plutonium and Highly-Enriched Uranium Needed for Pure Fission Nuclear Weapons*. Washington, D.C.: National Resources Defense Council. Inc.

Curtis, M. (2009). The Problem with Continuity of Knowledge in Enrichment Plant Process Monitoring. *50th Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Tucson, Arizona: INMM.

Daigle, M., & Goebel, K. (2009). Model-based Prognostics with Fixed-lag Particle Filters. *Annual Conference of the Prognostics and Health Management Society* (p. np). San Diego: PHM Society.

Delbeke, J., Howell, J., & Janssens, W. (2007). The Detection of Undeclared LEU Production at a GCEP by Real-Time Mass-Balancing. *48th Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Tucson: INMM.

Delbeke, J., Howell, J., Eklund, G., Janssens-Maenhout, G., Peerani, P., & Janssens, W. (2008). The Real Time Mass Evaluation System as a Tool for the Detection of Undeclared Cascade Operation at GCEPs. *8th International Conference on Facilities Operations - Safeguards Interface* (pp. on CD-Rom). Portland, Oregon: American Nuclear Society.

Dixon, E., Howell, J., Boyer, B., DeMuth, S., & Beddingfield, D. (2006). Evaluating New MC&A Concepts for Gas Centrifuge Enrichment Plants. *IAEA Symposium on International Safeguards* (pp. 625 - 635). Vienna: IAEA.

Doucet, A., De Freitas, N., & Gordon, N. (2001). *Sequential Monte Carlo Methods in Practice*. New York: Springer Science.

Doucet, A., Godsill, S., & Andrieu, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, 10 (3), 197-208.

Durst, P. (2008). An Analysis of Advanced Safeguards Approaches for New Large-Scale Uranium Enrichment Plants. *49th Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Nashville, Tennessee: INMM.

Engel, S., Gilmartin, B., Bongort, K., & Hess, A. (2000). Prognostics, the Real Issues Involved with Predicting Life Remaining. *Proceedings of the IEEE Aerospace Conference* (pp. 457 - 469). Big Sky, MT: IEEE.

Frazar, .., & Mladineo, S. (2010, June). Safeguards Culture: Lessons Learned. *ESARDA Bulletin* , 44, pp. 11 - 16.

Friend, P. (2010). URENCO Conference on GCEP Safeguards, held in December 2009. *51st Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Baltimore, Maryland: INMM.

Friend, P. (2008). Urenco's Views on International Safeguards Inspection. *8th International Conference on Facilities Operations - Safeguards Interface* (p. np). Portland, Oregon: ANS.

Garcia, H. L. (2010). Process Monitoring for Safeguards via Event Generation, Integration, and Interpretation. *51st Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Baltimore: INMM.

Garner, J., Gilligan, K., Henkel, J., Hooper, D., Krichinsky, A., Lockwood, D., et al. (2011). New Measures to Safeguard Gas Centrifuge Enrichment Plants. *33rd ESARDA Annual Meeting* (p. np). Budapest: ESARDA.

Garvey, D., & Hines, J. W. (2007). Dynamic Prognoser Architecture via the Path Classification and Estimation (PACE) Model. *AAAI Fall Symposium on Artificial Intelligence for Prognostics* (pp. 44 - 49). Arlington, VA: AAAI.

Goebel, K., Saha, B., Saxena, A., Selaya, J. R., & Christopherson, J. P. (2008, August). Prognostics in Battery Health Management. *IEEE Instrumentation and Measurement Magazine* , pp. 33 - 40.

Gordon, N., Salmond, D., & Smith, A. (1993, April). Novel Approach to Nonlinear/Non-Gaussian Bayesian State Estimation. *Radar and Signal Processing, IEE Proceedings F* , 140 (2), pp. 107-113.

Gyane, E. (2010). Information-driven safeguards: A country officer's perspective. *Symposium on International Safeguards* (p. np). Vienna: IAEA.

Haghighi, F., Nooraee, N., & Rad, N. (2010). On the General Degradation Path Model: Review and Simulation. In M. Nikulis, & e. al., *Advances in Degradation Modeling* (pp. 147-155). Boston: Birkhauser.

Hammersley, J. M., & Morton, K. W. (1954). Poor Man's Monte Carlo. *Journal of the Royal Statistical Society* , 23-38.

- Helferty, J. P., & Mugett, D. R. (1993). Optimal Observer Trajectories for Bearings only Tracking by Minimizing the Trace of the Cramer-Rao Lower Bound. *Proceedings of the 32nd IEEE Conference on Decision and Control* (pp. 939-939). San Antonio: IEEE.
- Henkel, J. H. (2010). *Development and Application of Process Monitoring Techniques for Transfer Facility Safeguards*. Oak Ridge: Oak Ridge National Laboratory.
- Hines, J. W., & Usnyin, A. (2008). Current Computational Trends in Equipment Prognostics. *International Journal of Computational Systems* , 1 (1), 95 - 109.
- Hines, J., Usnyin, A., & Urmanov, A. (2006). Prognosis of Remaining Useful Life for Complex Engineering Systems. *5th International Topical Meeting on Nuclear Plant Instrumentation, Control, and Human-Machine Interface Technology (NPIC&HMIT)*. Albuquerque: American Nuclear Society.
- Hori, M., & Mixui, M. (2008). Feasibility of Using RFID in the Material Accountancy and Safeguards Verification in the Nuclear Fuel Cycle. *49th Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Nashville, Tennessee: INMM.
- Howell, J. (2008). *A Review of Process Monitoring for Safeguards*. Richland: Pacific Northwest National Laboratory.
- Howell, J. W. (1995). *Advanced Integrated Safeguards Using Front-End-Triggering Devices*. Los Alamos: Los Alamos National Laboratory.
- Howell, J., & Bedell, J. (2007). Data Consistency Evaluation in GCEPs. *48th Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Tucson: INMM.
- Howell, J., Friend, P., Jones, D., & Taylor, C. (2009). Load-Cell-Based Mass Evaluation Systems Re-assessed on the Basis of Urenco(UK) Load-Cell Data. *31st ESARDA Annual Meeting* (p. np). Vilnius, Italy: ESARDA.
- IAEA. (1993). *Against the Spread of Nuclear Weapons: IAEA Safeguards in the 1990s*. Vienna: IAEA.
- IAEA. (2005). *Final Report on the IAEA Technical Meeting on Techniques for IAEA Verification of Enrichment Activities*. IAEA, Department of Safeguards, Division of Technical Support (SGTS), Division of Concepts and Planning (SGCP). Vienna, Austria: IAEA.

- IAEA. (2002). *IAEA Safeguards Glossary, 2001 Edition*. Vienna: IAEA.
- IAEA. (1998). *Model Protocol Additional to the Agreement(s) between State(s) and the International Atomic Energy*. Vienna: International Atomic Energy Agency.
- IAEA. (2002). *The Agency's Regime for the Protection of Safeguards Confidential Information*. Vienna: IAEA.
- IAEA. (2009). *The Safeguards System of the International Atomic Energy Agency*. Vienna: IAEA.
- IAEA. (1972). *The Structure and Content of Agreements Between the Agency and States Required in Connection with the Treaty on the Non-Proliferation of Nuclear Weapons*. Vienna: IAEA 1972.
- Johnson, L. A., & Montgomery, D. C. (1974). *Operations Research in Production Planning, Scheduling, and Inventory Control*. New York: John Wiley & Sons.
- Jones, B. (1984). Near Real Time Material Accountancy. *ESARDA Bulletin* , 7, 19-22.
- Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering* , 35-45.
- Kalman, R. E., & Bucy, R. S. (1961). New Results in Linear Prediction and Prediction Theory. *Journal of Basic Engineering* , 83 (D), 95-108.
- Kanazawa, K., Koller, D., & Russell, S. J. (1995). Stochastic Simulation Algorithms for Dynamic Probabilistic Networks. *Proceedings of the Eleventh Annual Conference on Uncertainty in Artificial Intelligence* (pp. 346 - 351). San Francisco: Morgan Kaufmann.
- Korbmacher, T. F. (2008). Field Trial of Mailbox Concept at a Urenco Enrichment Plant. *49th Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Nashville, Tennessee: INMM.
- Krichinsky, A. (2008). *Continuous Load Cell Monitoring: Verifying Nuclear Material Flows at Enrichment Plants*. Oak Ridge, Tennessee: Oak Ridge National Laboratory.
- Krichinsky, A. (2008). Evaluating Continuous Load Cell Monitoring as an Effective Safeguards Strategy for Feed and Withdrawal Systems. *49th Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Nashville: INMM.

Krichinsky, A. H. (2008). Evaluating Continuous Load-Cell Monitoring as an Effective Safeguards Strategy for Feed and Withdrawal Systems. *49th Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Nashville, Tennessee: INMM.

Krichinsky, A. M. (2009). Integrating UF6 Cylinder RF Tracking with Continuous Load Cell Monitoring for Verifying Declared UF6 Feed and Withdrawal Operations. *50th Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Tucson: INMM.

Krichinsky, A., Bates, B., Chesser, J., Koo, S., & Whitaker, J. M. (2009). *A Mock UF6 Feed and Withdrawal System for Testing Safeguards Monitoring Systems and Intended for Nuclear Fuel Enrichment and Processing Plants*. Oak Ridge: Oak Ridge National Laboratory.

Kullback, S., & Leibler, R. (1951). On Information and Sufficiency. *Annals of Mathematical Statistics*, 22 (1), 79-86.

Laughter, M. (2009). Optimizing and Joining Future Safeguards Efforts by "Remote Inspection". *Proceedings 2nd JAEA-IAEA Workshop on Advanced Safeguards Technology for the Future Nuclear Fuel Cycle* (p. np). Tokai-mura Ibaraki, Japan: IAEA.

Laughter, M. (2009). Safeguards for Alternative UF6 Enrichment Plants. *50th Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Tucson, Arizona: INMM.

Laughter, M., Gilligan, K., McGinnis, B., Morgan, J., & Whitaker, M. (2010). Utilizing Advanced Operator Instrumentation for Safeguards. *51st Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Baltimore, Maryland: INMM.

Laughter, M., Krichinsky, A., Hines, J. B., Kovacic, D. N., & Younkin, J. R. (2008). Simulated Process Test Bed for Integrated Safeguards Operations Monitoring. *8th International Conference of Facility Operations-Safeguards Interface* (p. np). Portland, Oregon: ANS.

Laughter, M., Shephard, A., Rowe, N., Whitaker, M., Martyn, R., & Fitzgerald, P. (2010). Interpretation of Field Test Data from UF6 Cylinder Accountancy Scales. *51st Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Baltimore: INMM.

Lenarduzzi, R., Begovich, J., Hines, J., Kivacic, D., Whitaker, M., & Younkin, J. (2007). Technologies for Real Time Monitoring of Load Cells for International Safeguards Applications. *48th Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Tucson: INMM.

Li, S. C., Lewandowsky, S., & DeBrunner, V. E. (1996). Using Parameter Sensitivity and Interdependence to Predict Model Scope and Falsifiability. *Journal of Experimental Psychology* , 125 (4), 360-369.

Lockwood, D. L. (2010). DOE Investigation of Advanced Safeguards Concepts for GCEPs. *51st Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Baltimore, Maryland: INMM.

Lu, C., & Meeker, W. (1993, May). Using Degradation Measures to Estimate a Time-to-Failure Distribution. *Technometrics* , 35 (2), pp. 161 - 174.

Luo, J., Namburu, M., Pattipati, K., Qiao, L., Kawamoto, M., & Chigusa, S. (2003). Model-based Prognostic Techniques [maintenance applications]. *IEEE Systems Readiness Technology Conference*, (pp. 330 - 340). Anaheim.

MacCormick, J., & Blake, A. (2000). A Probabilistic Exclusion Principle for Tracking Multiple Objects. *International Journal of Computer Vision* , 39 (1), 57 - 71.

Marseguerra, M., & Zio, E. (2009). Monte Carlo simulation for model-based fault diagnosis in dynamic systems. *Reliability Engineering and System Safety* , 94 (2), 180 - 186.

Musso, C., Oudjane, N., & Le Gland, F. (2001). Improving Regularized Particle Filters. In A. Doucet, N. de Freitas, & N. Gordon, *Sequential Monte Carlo Methods in Practice* (pp. 247 - 272). New York: Springer-Verlag.

Orchard, M., & Vachtsevanos, G. (2009). A particle-filtering approach for on-line fault diagnosis and failure prognosis. *Transactions of the Institute of Measurement and Control* , 31 (3-4), 221-246.

Orchard, M., Kacprzynski, G., Goebel, K., Saha, B., & Vachtsevanos, G. (2008). A Particle Filtering Approach for On-Line Fault Diagnosis and Failure Prognosis. *1st International Conference on Prognostics and Health Management*. Denver: PHM.

- Orchard, M., Tobar, F., & Vachtsevanos, G. (2009, December). Outer Feedback Correction Loops in Particle Filtering-based Prognostic Algorithms: Statistical Performance Comparison. *Studies in Informatics and Control* , 18 (4), pp. 295 - 304.
- Orchard, M., Wu, B., & Vachtsevanos, G. (2005). A Particle Filtering Framework for Failure Prognosis. *Proceedings of the World Tribology Congress*. Washington, D.C.
- Orsagh, R., Brown, D., Kalgren, P., Byington, C., Hess, A., & Dabney, T. (2006). Prognostic Health Management for Avionic Systems. Big Sky, MT: IEEE Aeroscape Conference.
- Patrick, R., Orchard, M., Zhang, B., Koelemay, M., Kacprzynski, G., Ferri, A., et al. (2007). An Integrated Approach to Helicopter Planetary Gear Fault Diagnosis and Failure Prognosis. *42nd Annual Systems Readiness Technology Conference* (pp. 547 - 552). Baltimore: AUTOTESTCON.
- Pickett, C. K. (2008). Results from a "Proof-of-Concept" Demonstration of RF-Based Tracking of UF6 Cylinders During a Processing Operation at a Uranium Enrichment Plant. *49th Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Nashville: INMM.
- Pickett, C. Y. (2008). Evaluation of an RF-Based Approach for Tracking UF6 Cylinders at a Uranium Enrichment Plant. *International Conference on Facility Operations -- Safeguards Interface*. Portland, Oregon: ANS.
- Pike, D., & Woods, A. (1982). Statistical Methods in Nuclear Materials Accountancy. *Nuclear Safeguards Technology 1982* , IAEA-SM-260/13, 359-372.
- Ristic, B., Arulampalam, S., & Gordon, N. (2004). *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Boston: Artech House.
- Robinson, M., & Crowder, M. (2000). Bayesian Methods for a Growth-Curve Degradation Model with Repeated Measures. *Lifetime Data Analysis* , 6, pp. 357 - 374.
- Saha, B., & Goebel, K. (2009). Modeling Li-ion battery capacity depletion in a particle filtering framework. *Proceedings of the Annual Conference of the Prognostics and Health Management*. San Diego: PHM Society.

- Shipley, J. (1983). Sequential Likelihood-Ratio Tests Applied to a Series of Material Balances. In F. Argentesi, R. Avenhaus, M. Franklin, & J. Shipley, *Mathematical and Statistical Methods in Nuclear Safeguards* (pp. 183 - 212). Switzerland: Harwood Academic Publishers.
- Sokolov, Y. M. (September 2006). Nuclear Power - Global Status and Trends. *Nuclear Energy 2006* , pp. 15 - 19.
- Speed, T., & Culpin, D. (1986). The Role of Statistics in Nuclear Materials Accounting: Issues and Problems. *Journal of the Royal Statistical Society , A 149* (Part 4), 281-313.
- Stigler, S. (1999). *Statistics on the Table: The History of Statistical Concepts and Methods*. Cambridge: Harvard University Press.
- Takeuchi, K. (1976). Distribution of information statistics and a criterion of model fitting. *Suri-Kakagu (Mathematical Sciences)* , 12-18.
- Tsvetkov, I. B. (May 22-24, 2007). Implementation of the IAEA's Model Safeguards Approach for Gas Centrifuge Enrichment Plants. *29th ESARDA Annual Meeting*. Aix en Provence, France.
- Upadhyaya, B., Naghedolfeizi, M., & Raychaudhuri, B. (1994, June). Residual Life Estimation of Plant Components. *P/PM Technology* , 7 (3), pp. 22-29.
- Whitaker, J. M., Lockwood, D., Lousteau, A., & Zhernosek, A. (2009). Demonstration Workshop on Safeguards Technologies for Uranium Enrichment Plants. *50th Annual Meeting of the Institute for Nuclear Materials Management* (p. np). Tucson, Arizona: INMM.
- White, J., Laughter, M., & Krichinsky, A. (2010). Results of Inspections of Operation of the ORNL Mock Feed/Withdrawal System. *50th Annual Meeting of the Institute of Nuclear Materials Management* (p. np). Baltimore, MD: INMM.
- Woods, A., Pike, D., & Rose, D. (1980). *Improved Statistical Procedures for Detecting a Loss in Nuclear Materials Accountancy*. Tech. Rep. 2/80/07, University of Reading, Department of Applied Statistics.
- Yan, J., Koc, M., & Lee, J. (2004). A prognostic algorithm for machine performance assessment and its application. *Production Planning and Control* , 15 (8), 796-801.

APPENDICES

Appendix A: Description of the Mock Feed and Withdrawal Facility

A full description of the development and construction of the mock F&W facility may be found in Krichinsky's report [Krichinsky, et al., 2009]. For convenience, a brief description is provided in this appendix.

The mock F&W facility simulates the enrichment process using water in place of uranium hexafluoride. Instead of a cascade of centrifuges, the enrichment "process" is mocked using a surge tank with an outlet flow that is split between product and tails stations. The product to tails flow rate ratio may be controlled using two needle valves – one for the product flow and one for the tails flow. The water is fed into a surge tank through any of three feed stations. An image of the facility is provided in Figure A1 from the Krichinsky report [Krichinsky, et al., 2009]. The three feed stations are in the back right of the figure, the two tails stations are in front, and the three product stations are to the left.

F&W Facility Components

Surge Tank. The surge tank is a 32 inch tall, 28 inch wide cylindrical polyethylene tank. A hinged lid on the top allows easy access if needed, but prevents evaporation when shut. A feed line has been added to the top of the tank, and discharge is gravity fed through an outlet at the bottom of the tank. Typically, the water level in the tank is maintained at 20 inches (measured from the outlet) so that sufficient pressure is always available to provide suitable withdrawal flow rates. A "volume-eating" cylinder is present within the surge tank to reduce the amount of water held within the tank and increase the sensitivity of the process to changes in feed or withdrawal flow rates.

It should be noted that in a GCEP, the amount of holdup in the cascade is typically only a few kg of UF_6 . The 20 inches of water in the surge tank, even after accounting for the empty volume within the surge tank, represents far more material than GCEP holdup when scaled up. However, the actual variation of the water level in the surge tank only ranges from 19.9 inches to 20.2 inches at the absolute maximum. The remaining 19.9 inches of water height serves only to provide static pressure for the withdrawal flow and is therefore not considered a part of the holdup of the process.



Figure A1 - Mock F&W Facility - A Snapshot

Feed and Tails Tanks. The feed and tails tanks are 25-gallon polyethylene tanks with an engineered stainless steel cradle for transport by lift truck. A square-shaped recess is at the bottom of each tank; by placing the end of the feed tube into the recess, virtually all of the water in a feed tank may be processed into the system. When empty, the tank/cradle assembly weighs approximately 35 kg; when full, the assembly weighs approximately 140 kg. A feed/tails tank is shown in Figure A2.

A hole has been drilled into the cap of each of the tanks so that the feed and tails tubes may be securely fastened to the tank. When the tube is inserted, a conical fitting on the tube hold the tube firmly in place. When the tank does not have a tube in place, a rubber stopper is inserted into the hole in the cape to prevent evaporation as well as possible spillage during transport.

Product Tanks. The product tanks are 10-gallon polyethylene tanks with no special cradle assembly. With the lid, each tank weighs about 1.14 kg empty and about 12 kg full. Due to the geometry of the product stations, the product tubes need not be secured as tightly as the feed and tails tubes, so no hole is drilled in the product tank lids. It is assumed that evaporation is not significant for the time required to fill a product tank (approximately two hours is typical), and the lid is fastened when the tank is not actively being filled. A product tank is shown in Figure A3.


Positive displacement pumps. The water is fed from the feed tanks to the surge tank by positive displacement pumps – one pump for each feed station. The pumps are rated for up to 80 L/hr flow rates with a rated head of up to 20 psig. The flow rate is varied by adjusting the stroke and the frequency of the piston cycle from zero to one hundred percent. The stroke and frequency adjustments are independent. During operation the stroke is typically maintained below 100% (e.g. 80% is typical) in order to avoid wear that may occur when the piston travels its entire stroke length. A feed pump control panel is shown in Figure A4.



Figure A2 - A Tank on Scale at a Tails Station



Figure A3 - A Product Tank on Station



Valves. For each feed and withdrawal station, there is a manually operated cutoff valve. These valves are actuated as either open or shut and are only used to allow or prevent flow; they are not used for throttling. There is a similar cutoff valve at the outlet of the surge tank.

There are two needle valves for the product and flow lines; one valve for each line. These valves provide fine control for the ratio of product and tails flow rates and are used to simulate the "enrichment" of the water. Typically, about 5% of the surge tank discharge is directed to the product stations and the remainder is directed to the tails stations. Figure A5 shows the product and tails needle valves (PV-H and TV-H, respectively) as well as the surge tank control valve, which is identifiable by the manual yellow handle.

Load Cells. The feed and tails stations consist of centering pans placed on 250-kg capacity scales with precision of 0.01 kg. The product stations are 25-kg capacity scales with precision of 0.001 kg. Two additional scales – one of each type – serve as the accountancy scales and have been calibrated with accuracy of 0.01% over the lower half of their range and 0.03% over the upper half of their range. The weight of each scale is read by an OHAUS scale monitor (Figure A6) that includes basic functions for re-zeroing and re-calibrating the scale readings if necessary.

Diversion Line. Diversion may be simulated using a bleed off line on the product flow line. The bleed off line is located between the product throttle valve and the product stations. A medical IV drip system is employed to allow fine control of the typically slow diversion flow rate. The diversion assembly may be seen in Figure A7, with a close up of the flow rate control wheel in Figure A8.

PI Controller. The surge tank water level is controlled by proportional-integral control (see Section 4.1.1) that actuates a globe valve between the surge tank outlet and the surge tank cutoff valve. The controller senses the height of water in the surge tank by a pressure transducer located at the outlet of the surge tank. The set point of the PI controller is typically 20 inches of water and the PI controller itself is written in LABView™, which resides on a Dell™ D620 laptop, as seen in Figure A9. A National Instruments™ data acquisition card simultaneously reads the pressure transducer signal and transmits the control voltage for the PI control valve actuator.




Figure A5 - Product and Tails Needle Valves, and Surge Tank Cutoff Valve



Figure A6 - OHAUS Scale Monitor



Figure A7 - Diversion Line with Collection Tank in Place



Figure A8 - Diversion Line Control Wheel





Figure A9 - PI Control Computer with LABVIEW™ Screenshot

Operation of the PI control is relatively simple and can be seen in Figure A10. The control program is defaulted to start with the control valve in the off position to avoid accidental discharge. The control scheme may be switched from manual control (where the user selects the actuator control voltage by an on-screen slider) to automatic (PI) control by a single mouse click. PI parameters are adjustable by on-screen entry and may be changed during operation, though such changes are never made in practice. The tank representation and the chart allow the user to observe the controller's interpretation of the surge tank water level as well as the surge tank set point and the control valve position.

Basic Facility Operation

The actual operation of the facility will depend on the type of simulation that is desired, but a few steps are common to all runs. These steps are briefly outlined here to give the reader an understanding of the mock feed and withdrawal process. During operation, a research logbook is maintained that describes all actions taken and records all relevant scale readings.

1. The tanks are weighted on accountancy scales and placed on the appropriate stations.
2. With the surge tank cutoff valve closed, the PI valve is manually perturbed to ensure proper operation. The PI valve is then shut and may be ready for automatic mode.
3. All facility cutoff valves are checked; the appropriate valves are opened in anticipation of the run.
4. Feed flow is initiated by activating the appropriate feed pump.
5. The PI valve is set to automatic mode (the cutoff valve was opened in step 3).
6. At the end of the run, the PI valve is set to manual mode and the control voltage set to zero.
7. All feed pumps are de-energized.
8. All cutoff valves are shut.
9. Tanks are re-weighed on accountancy scales as needed.



Appendix B: Measurement and Dynamic Noise Estimates

This appendix continues the discussion from §3.2.5.1, where the importance of measurement noise estimates and dynamic noise estimates in particles filters were discussed. Here, the noise estimate terms are developed for the inclusion in the particle filter prognostic model.

Estimation of the Measurement Noise

The measurement noise of the system is an estimate of the uncertainty of the state measurements themselves. (For example, a measured length variable may have a precision of +/- 0.1 m, or +/- 0.01 m.) As measurement noise increases, the effect of measurements on the particle weights during the SIS phase decreases. In Equation 33, increased measurement noise reduces the effect of the importance function by reducing the effect that distance from the measurement has on the updated weight of the particle in question. This is seen in the weight updating equation in Equation A1, where Y is the measured estimate, $\hat{Y}$ is the particle position, and σ_m^2 is the measurement noise. An artificially high estimate of measurement noise prevents the particle weights from updating efficiently, allowing outlier particles to maintain unnecessarily high weight. Likewise, an artificially low estimate of measurement noise causes the particle weights to degenerate too rapidly, reducing the importance of particles that may carry accurate information about the system and requiring more resampling steps than is necessary.

$$w_k^i \propto w_k^i * e^{\frac{-(Y-\hat{Y})^2}{2*\sigma_m^2}}$$

Equation A1 - Sequential Importance Sampling Update Equation

Equation A1 highlights the importance of both the particle error (the distance between the particle and the measurement) and the measurement uncertainty. High particle errors (the numerator in the exponential term) result in a rapid loss of weight for a particle, while high measurement uncertainties reduce the rate at which particles lose weight for a given error. While the particle error affects only the individual particle, the measurement error affects the weight updates of every particle.

The measurement noise in the system was estimated by utilizing a known sample of data from the system and estimating the noise with a median filter with a median window of size 8. By analyzing the error term between the known data sample and the smoothed data, an estimate of the measurement

noise was obtained. For the mock F&W facility, the measurement noise was estimated to be roughly $1.0\text{e-}3$ kg.

Influence of Measurement Noise Illustrated

The effect of the measurement noise estimate may be seen by the change in behavior of the particles for a given scenario. In Figure A11 through Figure A13, the particle filter prognostic method is applied to the same data three times, with only the measurement noise estimate varied between the three cases. In Figure A11, the measurement noise was estimated to be one tenth its nominal value. In Figure A12, the measurement noise was estimated to be ten times its nominal value. Figure A13 utilized the most accurate available estimate of the measurement noise.

In Figure A11, the low noise estimate quickly reduced the weights of all particles but a select few (those particles which were "lucky" enough to coincide with the early measurements). As the particles were redistributed in SIR, they were exclusively redistributed according to the positions of the weighted particles, resulting in gaps between the particle clusters and a near-constant variance of the particles throughout the ensuing series of measurements. Such a phenomenon reduced the effectiveness of the particle filter by ignoring particles that may contain relevant information about the system state.

Figure A12 illustrates the effect of an excessively high noise measurement estimate. Without the ability to efficiently de-weight outlier particles, the system variance is maintained higher than necessary. As a result, the precision of the prediction of time to SQ production is significantly reduced and the prediction itself is of less value.

When the measurement noise is balanced, the particle weights are changed greatly enough to favor the particles that better represent the system, yet do not change so rapidly as to automatically eliminate the importance of particles that do not coincide with the earliest of measurements (yet might well represent future measurements). This is seen in Figure A13, where the resampled particles are well distributed yet the variance of the prediction of the time to critical MUF is not excessive.




Figure A11 - Low Measurement Noise Estimate and Near-Constant Particle Weights




Figure A12 - High Measurement Noise Estimate and Uncontrolled Particle Distribution




Figure A13 - Accurate Dynamic Noise and Balanced Particle Distribution

Estimation of the Dynamic Noise

The dynamic noise is an estimate of how rapidly the system may change. (In the mock F&W facility, this is an estimate of how rapidly the *rate* of MUF accumulation changes, e.g. whether the MUF accumulation changes from 0 kg/hour to 0.1 kg/hour in the span of 5 minutes versus 10 minutes.) In tracking applications, the dynamic noise dictates how quickly the particles deviate from their "birth" positions after SIR. Higher dynamic noise estimates induce faster particle motion in order to ensure that the particles are always present in the possible range of future measurements of the target position. If the target has a low dynamic noise, then the particles need not move as quickly and the number of SIR updates may be reduced, resulting in a more efficient computational process.

In the extremes, a grossly high overestimate of dynamic noise may cause the particles to move so quickly that their SIS updates are rendered meaningless. In this condition, *all* particles end up with relatively little likelihood of representing the true system state and the SIS process cannot reliably weight the particles with respect to the actual measurements. Conversely, a grossly low underestimate of dynamic noise may cause the particle to lag behind the dynamics of the system. This condition also reduces the effectiveness of the SIS step by causing all particles to have low likelihoods of accurate state representation.

In a classical prognostic problem such as [Orchard, et al., 2008], the dynamic noise may be estimated by measuring the rates of system changes throughout the life of test cases. The maximal expected dynamic noise may then be used as an estimate of the system's total dynamic noise, ensuring that the particles may transition with sufficient speed to always cover the possible future ranges of the system state. For example, if a particle filter were applied to the data in Figure 9, the second derivative of the mean exponential growth of the crack length would provide a sufficient estimate for dynamic noise. If the exponential growth in Figure 9 were much more rapid, a higher dynamic noise estimate would be necessary.

Dynamic noise also affects the precision of the prognostic estimate of the particle filter. Low dynamic noise systems may have much higher prognostic precision because the particles "spread" apart more slowly. If the dynamic noise is great, then the future prediction has more uncertainty associated with the increased variance of the future system states, resulting in a less precise prediction. In this regard, particle filters naturally restrict their prognostic precision based on how rapidly the system state may

change. Particle filter confidence intervals are therefore a function of both the measurement error and the natural system dynamics of the system. This relationship is built into the particle filter by design.

Unlike classical health monitoring and prognostics problems, the dynamic noise of the mock F&W facility (and by extension, a GCEP) is not knowable a priori, as changes in the accumulation rate of MUF is a function of both the normal operation of the facility and the decisions of the operator. If the operator were to rapidly initiate a large diversion rate, the dynamic noise associated with the diversion would be much greater than a diversion that was slowly initiated and brought to full flow, even if the more slowly initiated diversion had the greater maximum diversion rate.

Worse, sudden changes in the system state might only be singular events. For example, if a rapid diversion were initiated, then the MUF accumulation rate might suddenly increase during initiation and cessation of the diversion, but otherwise be very low in magnitude. If the dynamic noise estimate is large enough to account for the singular changes, then the predictions provided by the particles would have high variances and low utility. If the dynamic noise estimate is instead maintained low enough to provide estimates with usable precision, then the particles would not be able to track the rapid change in MUF accumulation and would lose their validity. (SIR updating can only resample within the space occupied by the particles; if the actual system value is well outside this space, then resampling provides no benefit. Likewise the reweighting during SIS can no longer provide a meaningful distinction between particles and the particle filter model is rendered invalid.)

Figure A14 illustrates a scenario where the dynamic noise estimate is too low (in this case, by a factor of 5). The system is unable to respond adequately to a systematic change in state, and all particles become far removed from the actual measurements. In this case, *all* particles are estimated to have almost no weight during SIS and the update process loses physical relevance.

Meanwhile, Figure A15 illustrates a case where the dynamic noise estimate is too high by a factor of 5. The particles spread out quickly to cover the range of system state positions that the dynamic noise estimate suggests, but the state space coverage of the particles is excessive, and their great distances artificially inflate the variance of the state estimate.

Like the measurement noise estimate, the dynamic noise estimate must be balanced to provide a sufficiently agile response to changes in system behavior while still maintaining a precise variance of the particle position and prediction estimates.




Figure A14 - Low Dynamic Noise Estimate and Slow Particle Response




Figure A15 - High Dynamic Noise Estimate and Excessive Particle Response

Updating the Dynamic Noise Estimate Based on Data Behavior

To atone for the unknowable nature of the dynamic noise, an updating system was implemented. During normal operation with no diversion or other aberrant operator actions, the system dynamic noise was found to be roughly the same as the measurement noise (specifically, the dynamic noise varied between $1e-4$ and $1e-2$ kg/hour). Because of the similarity in values, the dynamic noise was initially set as equal to the measurement noise. (Note: the dynamic noise need not be as precise as the measurement noise, so long as it is sufficient and not grossly overestimated.)

Anytime a SIR procedure was implemented, the weighted distribution of the particles was estimated. Because the particles can only occupy a finite space based on their spread, their weighted distribution may be estimated as a beta function normalized to the maximum and minimum particle ranges. If the weighted distribution indicated that the actual system state was beyond the range of the particles (i.e. if the distribution were monotonically ascending or descending), then the dynamic noise was considered as insufficient. This monotonic behavior can be seen in Figure A16. The particle filter was returned to the previous SIR step (e.g. the particle filter was "rewound"). At this previous position, the dynamic noise estimate was then increased by a multiplicative factor and the coefficients of the linear model were biased in the direction of the system change. (If the weighted distribution indicated that the particles were too low in their MUF estimates, the coefficients were biased with greater values than the previous SIR step, and vice versa if the particles were too great in their MUF estimates.) The dynamic noise estimate update accounted for the dynamics of the system, while the biasing of the model coefficients allowed the particles to maintain presence in the system space occupied by the measured data. In this regard, the updating process allowed the particles to properly estimate the measured data in both high-dynamic and low-dynamic situations without requiring an excessively high dynamic noise estimate at all times. A multiplier of 1.15 was found to be effective for the data from the mock F&W facility.




Figure A16 - Weighted Particle Distribution when System Measurements are out of Range

Conversely, if the system dynamics were found to be lower than the estimate, the noise estimate was systematically reduced. When the system dynamic noise was less than the dynamic noise estimate, the weighted beta distribution of the particles was found to have a high peak value very near the midpoint of the distribution, as in Figure A17. When this condition was observed, the dynamic noise estimate was reduced by a specified factor. A lower limit was enforced such that the dynamic noise estimate was at least equal to the measurement noise to prevent excessive reduction of the dynamic noise estimate. In the mock F&W facility, a reduction factor of 1.1 was found to be effective (i.e. if the dynamic noise estimate needed to be reduced, the estimate was divided by 1.1).

The utilization of dynamic noise estimate updates allowed the particle filter to accommodate sudden, singular changes to the system state without requiring an unnecessarily high uncertainty of the prognostic estimates. This is seen in Figure A18, where the particles are able to track the sudden change from a non-diversion to a diversion even at around time = 4000 seconds. The high dynamic change of the system is successfully tracked, yet the particles do not have so much variance as to render a prediction meaningless. (In this case, a prediction of about 8000 seconds with a standard deviation of about 400 seconds was made.)




Figure A17 - Weighted Particle Distribution when System Dynamics are Low




Figure A18 - Dynamic Noise Estimate Updating and Controlled Particle Variance

Appendix C: Case Studies

Four separate runs of the mock F&W facility were used as test cases for the prognostic model. The first was a legitimate run where no diversion occurred; by extension, there was not a time at which a significant quantity of material could have been removed from the facility and the prognostic model should not provide a prediction.

Legitimate Operation

The entire feed and withdrawal profile for the legitimate run may be seen in Figure A19. This run was performed on July 7, 2010 from 17:30 to 19:15 GMT. A total of 209 kg of water was processed over a time of one hour and 42 minutes. The only significant perturbation came at the very end of the run when one pump was shut off and the flow rate of the system was significantly reduced. By calculating the cumulative inventory difference as in Figure A20, the stability of the system mass balance over time is evident.

Slow Diversion

The second case study was performed on June 28, 2010 over a total of two hours, 28 minutes from about 17:30 to 20:00 GMT. Over the course of the entire run, a total of 1.2 kg of material was diverted from the facility. With the slower diversion rate, the noise of the system was a factor in the stability of the predictions over time. The load cell profiles may be seen in Figure A21, and the cumulative inventory difference is in Figure A22.

Fast Diversion

The third case study featured a very fast diversion rate and was performed on October 14, 2010. The purpose of this run is to identify how quickly the prognostic model provides a sense of the strong diversion rate and whether the model can even find the diversion during the early seek time of the initial particle distribution. The load cell profiles are in Figure A23, and the cumulative inventory difference is in Figure A24. The run lasted for two hours, 56 minutes. A total of 1.719 kg of material was diverted over the course of one hour, 34 minutes. After the diversion, a batch replacement of material was performed, as can be seen in Figure A24.




Figure A19 - Feed and Withdrawal Profile for Legitimate Operation




Figure A20 - Cumulative Inventory Difference for Legitimate Operation




Figure A21 - Load Cell Profiles for Slow Diversion




Figure A22 - Cumulative Inventory Difference for Slow Diversion




Figure A23 - Load Cell Profile for Fast Diversion




Figure A24 - Cumulative Inventory Difference for Fast Diversion

Intermittent Diversion

The final test run was performed on September 21, 2010 over the course of one hour, 57 minutes. During the run, 1.89 kg of material was diverted. Diversion was halted midway through the process, however, when a tour of the facility was conducted. (The tour was known to be happening a priori, and the presence of the guests was seen as an opportunity to stop the run and attempt to hide the diversion. As an aside, the diversion was not visually identified by the guests, who did have some knowledge of the facility.) The load cell profiles are seen in Figure A25, and the cumulative inventory difference is in Figure A26. The pause in diversion occurred around the 15:50 GMT mark in the run and is barely visible in Figure A26. The pause lasted for about 5 minutes.

The run also feature the beginning of a diversion event well after the beginning of the run itself. With this feature, the particle filter will have to recognize the sudden change in regime introduced by the diversion even after the filter has learned the typical dynamic noise of the system.




Figure A25 - Load Cell Profiles for Intermittent Diversion




Figure A26 - Cumulative Inventory Difference for Intermittent Diversion

VITA

David Alan Hooper was born on October 8, 1976 in Billings, Montana and was raised in Casper, Wyoming. He graduated from Natrona County High School in 1995, and then attended the University of Wyoming. After two years as a Music Education major, he switched to Mechanical Engineering and graduated in the spring of 2001 as the Mechanical Engineering Student of the Year for his graduating class. After graduation, he married his wife Teresa and headed to Charleston, South Carolina where he served a commission in the United States Navy to instruct at the Naval Nuclear Power Training Command from 2001 to 2005. He finished his commission as a Lieutenant and was honorably discharged, having earned the honor of Master Training Specialist and having served as an officer class director as a Lieutenant Junior Grade (a role normally reserved for Lieutenants). He then drafted mechanical designs for buildings while waiting to start graduate school the following year.

In 2006, David enrolled at the University of Tennessee where he earned a Master's of Science in Nuclear Engineering in the fall of 2007. After spending time interning at Oak Ridge National Laboratory, taking extra classes at the University of Tennessee, volunteering at the Howard H. Baker, Jr. Center for Public Policy, and researching career options in law, policy, and engineering, David returned to full-time doctoral studies at the University of Tennessee in the fall of 2009. In the summer of 2011, David finished his PhD in Nuclear Engineering with a concurrent Master's of Science in Statistics (offered through the Intercollegiate Graduate Statistics Program of the College of Business at the University of Tennessee).