INL/EXT-10-20271

Safeguards Envelope Progress FY 10

Richard Metcalf

October 2010

INL/EXT-10-20271

Safeguards Envelope Progress FY10

Richard Metcalf

October 2010

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Office of Nuclear Energy
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

Safeguards Envelope Progress FY10

INL/EXT-10-20271

October 2010

Approved by

Richard Metcalf   Date

ABSTRACT

The Safeguards Envelope is a strategy to determine a set of specific operating parameters within which nuclear facilities may operate to maximize safeguards effectiveness without sacrificing safety or plant efficiency. This paper details the additions to the advanced operating techniques that will be applied to real plant process monitoring (PM) data from the Idaho Chemical Processing Plant (ICPP). Research this year focused on combining disparate pieces of data together to maximize operating time with minimal downtime due to safeguards. A Chi-Square and Croiser's cumulative sum were both included as part of the new analysis. Because of a major issue with the original data, the implementation of the two new tests did not add to the existing set of tests, though limited one-variable optimization made a small increase in detection probability. Additional analysis was performed to determine if prior analysis would have caused a major security or safety operating envelope issue. It was determined that a safety issue would have resulted from the prior research, but that the security may have been increased under certain conditions.

FOREWORD

The work described herein supported by the U.S. Department of Energy under DOE Idaho Operations Office Contract DE-AC07-05ID14517.

ACKNOWLEDGEMENTS

Mr. Duc Cao of Wisconsin University has been invaluable in helping code many of the methods to be demonstrated. Also to my reviewers, and Robert Bean of Idaho National Laboratory, I offer my thanks.

CONTENTS

ACRONYMS

AAKR	Auto associative kernel regression
DOE	Department of Energy
FAR	False Alarm Rate
IAEA	International Atomic Energy Agency
INL	Idaho National Laboratory
ICPP	Idaho Chemical Processing Plant
MAUA	Multi-Attribute Utility Analysis
MBP	material balance period
MC&A	material control and accountability
MUF	material unaccounted for
PM	process monitoring
PND	probability of nondetection
PP	physical protection
SNM	special nuclear material
SQ	significant quantity

Safeguards Envelope Progress FY10

INTRODUCTION

Modern nuclear systems will include process monitoring as part of their design to enhance their safeguards for both domestic and international safeguards approaches. These new requirements for enhanced safeguards must not incur significant costs, however, or nuclear facilities will be prohibitively expensive.

Though it is an obvious facet of design to be considered, safeguards have not historically been integrated into the design process, or even integrated fully into the operation of facilities. This has resulted in external, regulatory requirements adding synthetic costs to nuclear facilities because the industry has not embraced a systems engineering approach to safeguards. While the design phase systems engineering approach to safeguards would be Safeguards-by-Design, the operating-phase systems engineering approach is the creation of a Safeguards Envelope. The Safeguards Envelope program is currently working on a project to increase the security within nuclear facilities, using the Idaho Chemical Processing Plant (ICPP) as an example case.

This research has focused on maximizing an example material balance period (MBP) for special nuclear material (SNM). Two factors determine the optimum MBP. The false alarm rate (FAR) is the rate or percent of alarms which falsely declare a diversion scenario is taking place. The probability of nondetection (PND) is the probability of failing to detect a diversion. For a given set of parameters, decreasing the FAR usually requires relaxing the operational constraints and thresholds, but at the same time can increase the PND as it raises the possibility of hiding an abnormality. Thus, optimizing the MBP is also a problem of optimizing the FAR and PND.

Different statistical tests, however, provide different optimal FAR and PND. In this study, kernel regression analysis is applied to a declared ‘event’ from ICPP PM data to create a best fit curve. A trial data set is simulated from the ICPP data consisting of a ‘normal set’ and a ‘diversion set.’ Residual analysis and cumulative sum techniques are applied to determine optimum bounds for acceptable operating conditions based upon resultant FAR, PND, and MBP.

In prior work, several authors have found ways of increasing the probability of detection under different assumptions, with secondary papers refuting these claims due to the realism effects of safety, security, and plant requirements. [1] [2] [3] [4] [5] [6] [7] [8] Significant increases in detection had been accomplished in FY10 by the inclusion of new tests and methods of handling the data in the data sparse environment. However, the increases were ultimately flawed due to assumptions about the underlying data, and so only a small increase in the ability to detect material was accomplished. The underlying additional tests are available, but they remain untested.

Additionally, this year demonstrates the magnitude of the optimization problem. Similarly to the difficulty in analyzing nuclear fuel, optimization for reprocessing facilities through process monitoring and accountancy activities is extremely complex and large. Several methods for optimization were investigated, but all approaches remained too computationally expensive.

In addition, significant research this year was dedicated to evaluating the prior safeguards envelope operation compromised the safety or security envelopes on the example ICPP facility. Non-PRA safety analysis has shown an unfortunate result: because of issues in the original plant design, one of the requested changes to operation could have fatigued a major pipe and caused a significant leak. Other safety concerns, including criticality and increased dose, were found to be insignificant. Security concerns were also found to be insignificant, but the authors have not found an adequate method of assessing the vulnerability with existing common techniques.

THEORY OF SAFEGUARDS ENVELOPE

Definition of Safeguards

Safeguards are put into place to protect nuclear material from proliferation or other non-declared purposes, and are vital for securing the future of nuclear energy domestically and globally. This principle defines the need for fields such as nuclear nonproliferation, which guards against the theft or diversion of SNM. SNM is tracked through a nuclear facility, and that facility is responsible for ensuring that no more than one significant quantity (SQ) is unaccounted for in a given timeframe, ranging form one year to as little as 30 days. Depending upon the size of the plant, this can be an enormous and seemingly impossible task.

Definition of the Safeguards Envelope

Safeguards envelopes are operational spaces designed similarly to the idea of a safety envelope[1]. For years, the concept of an area of operation that does not needlessly endanger the public, personnel, or equipment of a nuclear facility has been a major component of nuclear facility design. This safety envelope methodology can just as easily be applied to safeguards, such that an operating space can be constructed that does not needlessly risk proliferation activities, while also not overburdening the operator with regulatory costs. The goal is simply to define a set of operational parameters which increase the probability of detecting a diversion of nuclear material[1] and apply them to operating and new nuclear facilities to make safeguards a point of optimization for operations instead of a fixed, ad hoc procedure. The most effective way to develop these parameters is to use real nuclear plant process monitoring data and perform statistical analyses and modeling methods.

PROBLEM STATEMENT

ICPP Facility

During its 43 years of operation, process monitoring was of great importance in the day to day operations of the Idaho Chemical Processing Plant. When the state-of-the art level/density scanner was introduced as part of the process monitoring system, the accuracy of the data improved drastically, allowing a greater confidence in nuclear materials control and accountability (MC&A). With this more reliable data, statistical analysis methods can be more effectively utilized to detect diversions of special nuclear material (SNM), and determine optimum operating parameters for both materials accountability and operator impact.

Diversion Detection

Process monitoring techniques and analysis methods are a primary focus in attempts to increase the ability to detect a diversion. The goal is to be able to detect as small of a diversion as possible without raising the false alarm rate (FAR) or the probability of nondetection (PND) above a reasonable level. If a FAR is too high it is not cost effective, for every alarm must be investigated which would be intrusive on the plant operator's other duties. With a high PND the issues are obvious; it is unacceptable to rely on a system for nuclear security when it fails to detect diversions. By utilizing advanced statistical analysis techniques, one can determine a balance of optimum working parameters and also obtain a better material balance period (MBP). In this study, MBP is used as the figure of merit because it can easily accept the FAR and PND into a single metric. This single figure of merit allows for a single optimization point rather than two, but these are both available for more detailed or plant specific studies.

STATISTICAL ANALYSIS

In this study, kernel regression was used to create a best fit function to the data received from ICPP. Kernel regression is a state estimation technique which is considered a nonparametric technique, for unlike linear regression, it does not assume a fundamental distribution in the data. The proposed algorithm compares historical and trial data sets and tests the ability to detect a diversion by looking at two items: degree of residual randomness and deviation from the mean. To determine the effectiveness of the statistical tests, we perform a Markov Monte Carlo simulation and run 500,000+ trials as a simple method for finding out the resultant FAR and PND values. A discussion of how this analysis is performed in detail is given in APPENDIX B. A more thorough description is provided in the prior reports in this series. The updates to this analysis are presented below.

Chi-Square Test

A Chi-Square test has been used to replace the original Z testing of the cumulative sums to identify if the deviation from the "true" values has the appropriate variance. An unusually high variance could represent diversion, mechanical fatigue, or sensor failure. While the typical Chi-Square test is used to evaluate the performance of a system, the application of this test to the residuals can provide a second measure to determine if the residuals are away from normal.

Equation 1 shows a standard Chi-Square test. In this equation, X^2 is our test statistic, Oi are the individual observations (residuals), and Ei are the expected values.

$$\chi^2 = \sum_{i=1}^{r} \left[\frac{(O_i - E_i)^2}{E_i} \right] \tag{1}$$

The numerator term appears to be a sum of the residuals, and the denominator term becomes the variance of the historical set. Similar to a student's-t test, very little is assumed about the data that is available and as a result, the Chi-Square test has differing thresholds for evidence based on the number of degrees of freedom. The degrees of freedom are one less than the number of observations in the set. Safeguards Envelope operation, therefore, not only increases the evidence set that is available for a statistical determination, but allows for higher thresholds to minimize the FAR.

This test has been integrated with the student's-t as part of the standard suite for detecting diversions. Specifically, this test provides a mechanism for combining all residuals positively to address the diversion scenario of removal of material during a statistically high event.[a]

The issue with adding a Chi-Square test is the increased FAR that is to be expected from adding additional tests. As discussed in the FY08 report, a union or intersection model can be created with the Chi-Square and cumulative sum test. Some diversion types would not typically be detected with the cumulative sum test, and so only a union model can be applied. This has an unfortunate disadvantage: the FAR must increase with the linearly with the FAR for each test, but the detection probability for some diversions is only derived from one test.

Croisier's CUSUM

Croisier's CUSUM is a cumulative sum method which updates the prior sum before moving to the next iteration. The update to the prior sum determines if the new sum will be moved towards zero (as given in Eqs. 2-7), or if the system will be reset to zero. This resetting to zero is expected to increase the PND but decrease the FAR and so may be preferred in applications where many measurements are taken in multiple locations. The reduction to FAR, which increases linearly to the number of measurements in the union model, is a crucial requirement for MBP and acceptance by operators.

$$C_t = \{(\mathbf{S}_{t-1} + \mathbf{e}_t)^T \mathbf{\Sigma}^{-1} (\mathbf{S}_{t-1} + \mathbf{e}_t)\}^{(1/2)} \tag{2}$$

$$\mathbf{S}_t = \mathbf{S}_{t-1} + \mathbf{e}_t - \mathbf{k} \text{ if } C_t > \lambda \tag{3}$$

[a] This diversion is outlined in APPENDIX A .

$$\mathbf{S}_t = 0 \text{ if } C_t \leq \lambda \tag{4}$$

$$\mathbf{k} = (\mathbf{S}_{t-1} + \mathbf{e}_t)\frac{\lambda}{C_t} \tag{5}$$

$$\mathbf{S}_t = (\mathbf{S}_{t-1} + \mathbf{e}_t)\frac{1-\lambda}{C_t} \tag{6}$$

$$Y_t = (S_t^T \Sigma^{-1} S_t)^{(1/2)} \tag{7}$$

In Eqs. 2-7, C_t is the existing and updating cumulative sum, S_{t-1} is the prior sum, the new S_t is the new sum added onto this group, k is a scalar (in the direction of S for the multivariate case), λ is a scaling parameter, and Y_t is the new test statistic.

The procedure for this analysis is very similar to other cumulative sum tests. The updated cumulative sum is used as a test statistic to determine if the root mean error is beyond a certain threshold with a given probability. Unique to this test is the parameter λ. λ is a scaling parameter for the impact of the most recent sum. In a students-t test, this factor is zero. However, if this parameter is nonzero, λ reduces the FAR but increases the PND because it adds an additional threshold for divergence on a given measurement before it is added to the cumulative sum, as expected by a system which has thousands of measurements.

One of the issues associated with Croisier's cumulative sum is that a control parameter, λ, is required as well as the standard threshold. Similarly to the Chi-Square test, this test has the potential to increase the optimization, but also synthetically increases the parameter space. In the event that Croiser's CUSUM's λ variable is highly sensitive, this test must be discarded. Unfortunately, this year an appropriate value for k was unable to be determined since implementation, and a basic student's-t statistic has been uniformly superior over the variables tested.

SIGNIFICANT ISSUES WITH BASE DATA

Original results for this year exceeded expectations of sufficient magnitude to require a second analysis. The FAR and PND were both significantly lower than 0.1%, allowing for several months MBP. This error took significant time to track down because the code associated with these tests had not been changed before running the analysis. It was believed that a heretofore unknown error within the code had produced a major bug in the reporting statistics.

Several techniques for debugging were applied, as well as external review by researchers not associated with this project. Finally as the code was determined to be accurate, the new test variables, a combination of level, density, and temperature, were checked to ensure no major data flaws.

An analysis for skewness on the density and level measurements provided no valuable results, even at the p=0.01 level. While at a glance the data is correct, a very careful examination shows that the density measurements begin later and end sooner than the level measurements. As a result, the error (which is most significant at the beginning and ending of the transient) was drastically misestimated, suggesting that process monitoring could perform much better than would be achievable in actual operations. This overlap is demonstrated in figure 1..

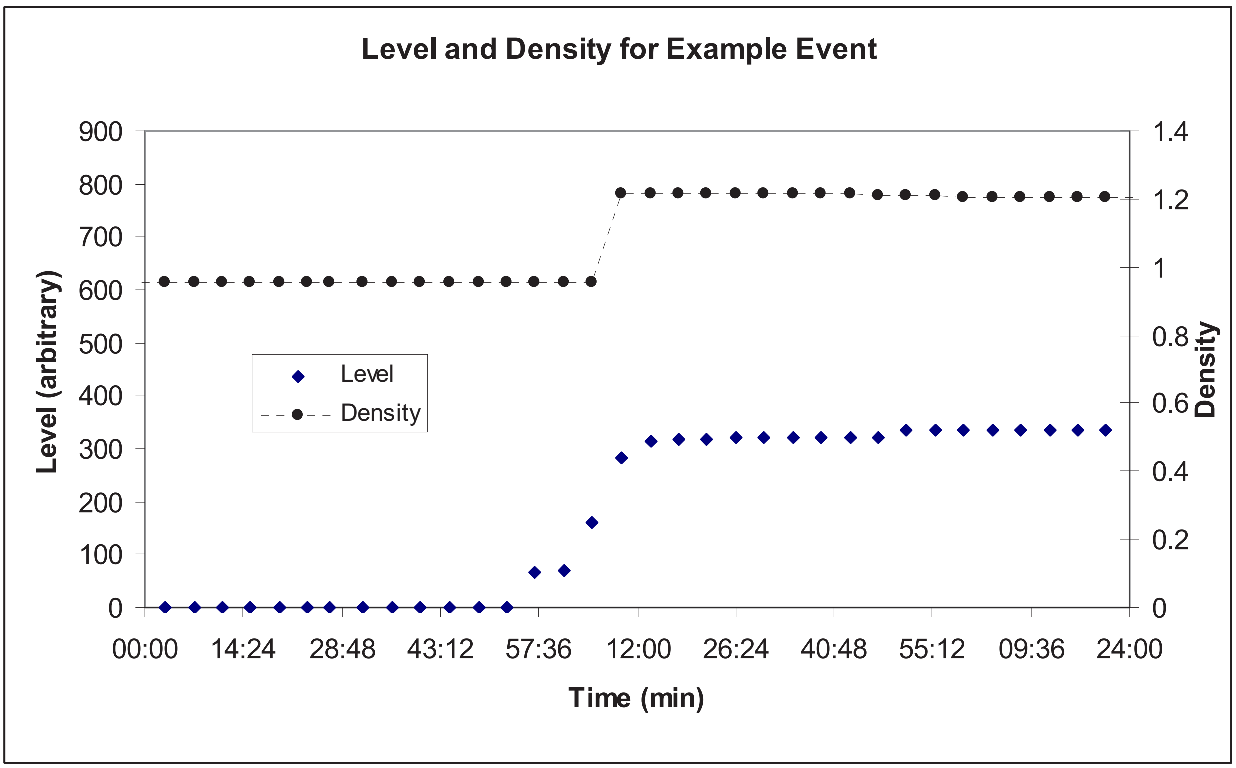

Figure 1: The data does not match in timestamp correctly.

The reason for the drastic increase in material balance period with the safeguards envelope operation was because of the synthetic increase in the number of points created points within the area in which the error was driven to zero. This compounded the underestimation of the FAR and PND and led to the drastic overestimation of the MBP. When these points were removed, the PND and FAR approached values more in line with that would be expected, given the assumptions presented in the analysis.

Results of the Limited Optimization

Initial results showed extreme improvement over FY08 and FY09 activities, but these results had to be ignored because of the data issues given in the prior paragraph. Unfortunately, a significant amount of time was allocated to combinational multivariate approaches in FY10, and so only the optimization approaches of early FY10 are still valid. Future research in this area will allow for an analysis of the time-lag and an estimate of its error for the remaining FY10 research. The optimization also was only performed on one variable (since this was intended to be applied to a single combined variable. The limited improvements did reduce the base PND while the FAR remained close to 1%.

TABLE 1

Tabulated Results Including Previous Work.

Data Type	Z-test	Kernel Reg.	CUMSUM	FY	FAR	PND
One-Var.	YES	YES	YES	FY09	1.15%	12.39%
One-Var.	YES	YES	YES	FY10	1.03%	10.10%

EXPLORATION OF THE SAFEGUARDS ENVELOPE PARAMETER SPACE

Exploration of the Parameter Space for the Safeguards Envelope:

Any safeguards envelope that is created should maximize the benefits to operations, safeguards, and security. Complexity in nuclear systems prevents an easy optimization, however. In a similar manner to the adjustments to load-cells required for relative humidity and air density, precise determination of the exact optimal parameters will require extensive start-up testing.

When simulations are available and these parameters are being estimated through models, it becomes much easier to explore the parameter space and determine potential optimal operations. Unfortunately, the size and versatility of bulk processing nuclear facilities and the flexible requirements of detection and false alarm for each subsystem create an exceptionally large number of variables. An example set of variables this analysis uses is provided:

L-Norm Level
Threshold for Student's-t test
Threshold for Chi-Square test
Operational changes-slowing down
Operational changes-location of slowdown
Kernel bandwidth
Weight per Kernel Residual
Acceptable Confidence Intervals
Number of intervals to test for Chi-Square
Number of tests performed per time for the cumulative sum
Amount of "rebaselining" per test

It must be assumed that:
1) Each variable contributes in at least a linear fashion,
2) Some variables contribute in nonlinear fashions, and
3) Independence is not expected from any variables.

Because the nonlinear nature of the interaction is unknown, this can be approximated by a series of exponentials or polynomials. For the sake of simplicity, this analysis will use polynomials. In this case, the final effect on the FAR, PND, and MBP are specific expressions of the generalized equation:

$$MBP(x_1, x_2, \ldots, x_N) = \sum_{i_1=0}^{\infty} \sum_{i_2=0}^{\infty} \cdots \sum_{i_N=0}^{\infty} \alpha_{i_1 i_2 \cdots i_N} x_1^{i_1} x_2^{i_2} \cdots x_N^{i_N} \tag{8}$$

In Eq. 8, the each alpha is a determined coefficient of the factors, which are given by the various x.Determining the coefficients of the factors of the infinite sum is not feasible with modern computational methods as the amount of factors limit to infinity. The nonlinear character of the interaction makes standard regression analysis useless unless enough cases were run to isolate the interaction of each variable.

Previous analysis using Markov Monte-Carlo evaluated the PND and FAR with a running time of roughly an hour per analysis. Each adjustment to each variable required an entirely new run. Under a broad assumption that each variable ten settings to explore the entire parameter space, the time required for a complete exploration is 10^N hours.

This can be accelerated by deconstructing the analysis into group families but the this can ultimately only reduce the amount of calculations to a loss of two degrees of freedom. While a reduction by two orders of magnitude may seem significant, this analytical method still cannot be pursued. The required time has led to the requirement of seeking secondary methods of evaluating the most relevant factors in the parameter space.

2^k Factorial Method

A method developed for exploring extremely large parameter spaces in operations analysis is the 2^k factorial analysis. In this analysis, an arbitrarily high and an arbitrarily low value is assigned for each one of the potential quantitative variables. Notes that some variables, such as diversion-type, are qualitative, and so a subset of qualitative variables must be chosen.

This can reduce the number of independent tests to factors of two. There are several assumptions in this model, however:

1) Each variable has only first order interaction with each other variable,
2) Each variable has only first order impact into the final function, and
3) Choice of "high" and "low" values are not the absolute limit and represent "appropriate" values for the quantitative measure.

The third assumption is not difficult to overcome for most of the conditions listed. For example, p values for the statistical tests of the Chi-Square and the student's-t test are unlikely to 0.50, and much more likely to be appropriate in the standard tests of p=0.05 or p=0.01. The first and second assumptions, however, are not appropriate for this model. It is reasonable to expect that a student's-t test on the residuals and a Chi-Square test on the absolute variance are poorly estimated by a “high” and a “low” variable test set. These functions should have multi-order impact, and not even be linearly independent. As a result, the proposed approach of the 2^k factorial method was not applied.

QUALITATIVE ANALYSIS OF SAFETY AND SECURITY FOR THE SAFEGUARDS ENVELOPE OPERATION

SAFETY

It was hypothesized that no significant safety effects would be found as a result of retaining part of the slurry mix in the G-105 tank for a long period of time. The G-105 accountancy tank is one of the first tanks that all of the head-ends of the ICPP feed into. Because of the variability in the head-ends for the ICPP, different slurry mixes are to be expected to be fed into G-105. One of these head-ends is for handling zirconium clad fuels, and at the time this was done with hydroflouric acid, this is shown graphically below in Figure 2.

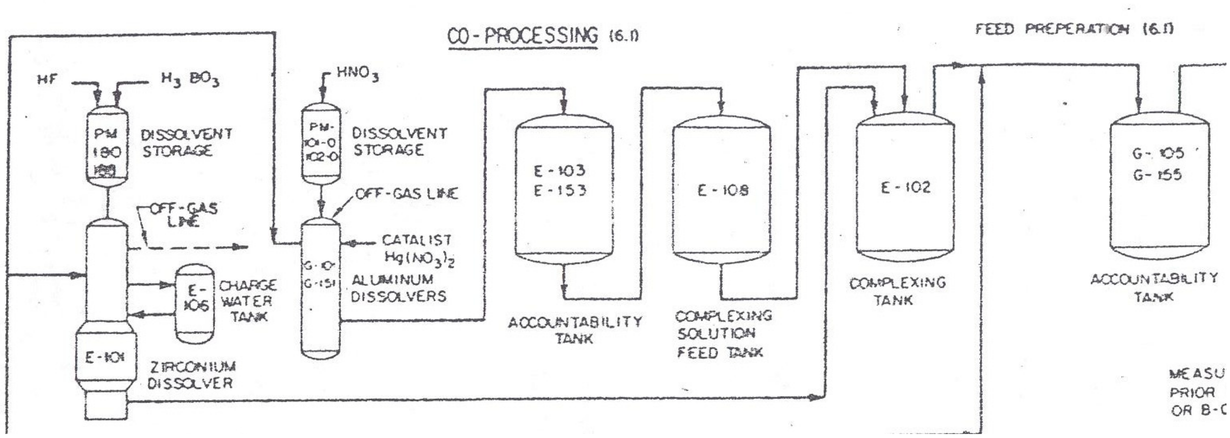

Figure 2: HF is used as part of the Zirconium dissolution process.

The second line appearing from the bottom of the figure is a second potential HF line from a different zirconium dissolver not using a coprocessing approach. In the event of a failure of the operator to clear the HF through the complexing tank in the coprocessing line, or a failure in the nitric acid complexing from the secondary line, G-105 would be flushed to remove the HF from the stainless steel in that tank. This flushdown would share very similar characteristics to the transient on which this analysis has focused.

Slowing the speed of the flushdown from HF in the stainless steel tank would have significantly increased the fatigue on the stainless steel because of the additional corrosion. This optimization, therefore, may not address the most dangerous diversion pathway: claims of mistake/valve failure followed by flushdown in the accountancy tank.

While it could be expected that this HF could be isolated in a nearby tank, the only tank available in this facility for storage from the G-105 tank is through the G-106 feed tank and then the G-108 rework tank. This is presented in Figure 3.

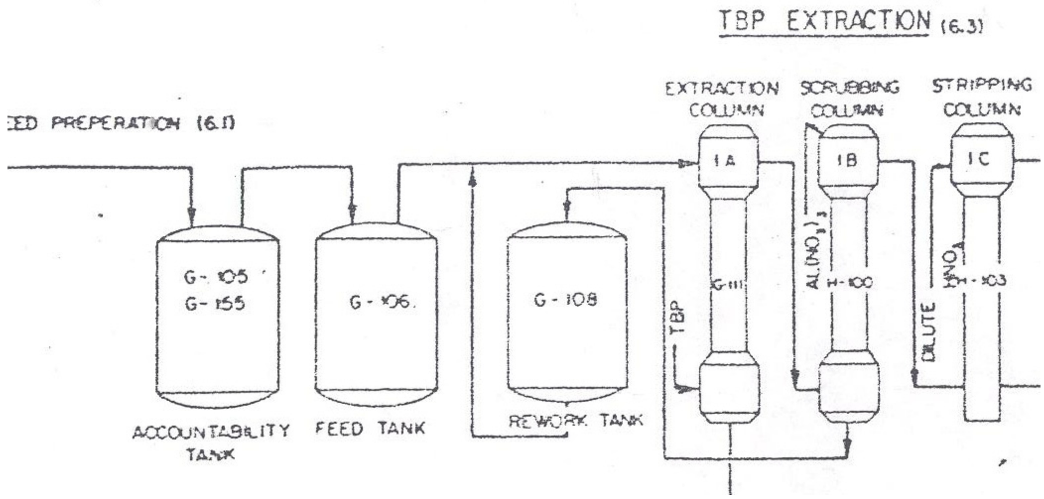

Figure 3: G-105 moves material through G-106, and then can be moved to the G-108 tank, or optionally through the extraction column.

A diversion over the transient of the fill for G-105 may not be able to even be tested with the G-106 tank if relief material (for neutralizing the HF) is added to G-106. Even if this is the case, the G-108 work tank is primarily for nitric acid, meaning the undue fatigue is to be expected. A rational actor, however, will likely maintain this error in the G-108 tank instead of contaminating and potentially destroying the extraction columns (tanks are less expensive than columns).

Therefore, the safeguards envelope operation suggested may have induced undue fatigue in not just the tanks, but the connected pipes. This reinforces the issue with analysis of all the variables quantitatively in the previous sections with off-normal conditions of the facility. With a full scale simulator of a reprocessing facility, it may be possible to analyze these conditions.

SECURITY

Changes to location of material in a nuclear facility can have significant impact on the security of that material to non-state threats. Knowledgeable insiders and external attackers both benefit from material remaining in a vulnerable location for longer periods of time. The safeguards envelope that has been proposed for the ICPP in the head-end process around G-105 increases the overall wait time during the transient.

Quantitative analysis of facility vulnerability that includes the length of time for vulnerability is rare and typically not applied in the DOE. Instead, since it is typically assumed that an attack will occur during the highest vulnerability (and therefore the window this remains open does not matter), the safeguards envelope that has been proposed would not indicate a higher level of vulnerability.

Furthermore, the vulnerability must be considered in context to other vulnerabilities in the facility. The operational changes on the example G-105 tank increase the transient

time by thirty minutes, an inconsequential time compared to the historically long wait times of the material in other normal operations (these often exceed several hours based on the available data set). The amount of total increased vulnerability, from the qualitative argument presented here, is insignificant as well.
Safeguards envelope operation combined with integration of process monitoring with safeguards systems could actually significantly increase the physical security of most locations in the facility. Assume a process monitoring alert (immediate) results in immediate heightened awareness of physical security systems. Ongoing research by Cipiti and Duran suggests that this would significantly increase adversary neutralization. As an example, assume the probability of detection by the immediate process monitoring methods triggers such a high alert. Process monitoring becomes a tri-use system because it supports operations, safeguards, and security collectively.

If process monitoring and other real-time sensors can be integrated into the safety, security, and operational envelope, a set of new metrics could be created for balancing the effects of reliance on any type of system. Accountancy methods do not contribute to physical security, but if process monitoring can, a suboptimal combination of process monitoring and accountancy for safeguards may be preferred because it provides benefit to physical security.

REMAINING CHALLENGES

Time-correlation Correction on Existing Data

An analysis of event beginning and ending in the existing data is crucial. Research by Burr had suggested that the start and stop of these events can have a significant impact on the uncertainty in a process monitoring system and this analysis has demonstrated that experimentally. Future analysis in this area must be able to correlate the start and stop of events accurately for the use of the optimization developed in FY10. Furthermore if this start and stop correlation can be combined with multi-tank analysis, it is possible that research associated with time-stamp events and propagation through complex systems could be applied, such work by Humberto, *et al.* [9] [10]

Determining Equipment Failure Versus Diversion

An alarm is raised because the plant begins operating outside the accepted parameters. However, the abnormality can be due to something as devious as material diversion, or as innocent as equipment failure. Since equipment does wear out and eventually fail, it could affect the process monitoring system and its ability to detect diversions effectively, raising the FAR. A code was developed in 1997 to address this specific problem, called IGENPRO, which was designed based on fuzzy logic and PM techniques. IGENPRO attempted to effectively estimate when a component might fail within the plant, and could be used to develop a more proactive maintenance schedule, rather than waiting until things failed completely. This system or possibly a more advanced code could alleviate the issue of increased FAR due to equipment failure. [11]

APPENDIX A

Diversions of Significant Concern

Statistically High Residual Diversion

Problem

Consider the case in which through random statistical fluctuations that the cumulative sums are unusually high. It is expected that during the lifetime of the facility that a deviation of four standard deviations. This represents an event that happens roughly one per forty-thousand measurements. In this case, the highly deviant event renders analysis of the cumulative residuals intrinsically flawed. While it may be possible to exclude highly anomalous events, if these events are not screened, the primary diversion-detection methodology will fail.

Different from a safety envelope, in which nature obeys statistical independence between events, diversion detection must include dependent response actions taken by a potential diverter. In this hypothetical case, the diverter knows that the cumulative sums are unusually high and chooses to divert during this example transient. Diversion of 0.5% would be undetected because of the anomalous event.

Solution

Integration of the Chi-Square test, which adds all residuals positively, will trigger an event-of-interest if larger than the defined threshold value. Chi-Square thresholds are determined by the FAR (p-value) and the degrees of freedom. Unfortunately this means that a new optimization threshold must be determined, but the diversion can be detected through this method.

Last Transient Point Diversion

Problem

Diversions from the last point in a transient are impossible to detect with a single tank analysis. The path between any given two points can be approximated with kernel regression or simply a line, but ultimately this estimation is only useful with regard to future points. In the data poor environment of a single point followed by zeros (such as at the end of a transient), there is no conclusion that can be drawn through level, temperature or density because all of these values are expected to go to zero.

This diversion type could remove 0.5% of material in the example transient that is used in the Safeguards Envelope analysis. Inclusion of additional points reduces the differential amount of material removed at each step, which further reduces the amount of material that could be stolen from this example diversion.

This diversion is made even more complicated with replacement of the material with nitric acid at temperature. In this case, the diverter may hide the diversion in the systematic error between the tanks.

Solution

This diversion type is significantly mitigated by increasing the number of points between the penultimate point and zero. As a result, operations may be recommended to maintain a small amount of material in the tank rather than letting it drain immediately. From a statistical standpoint, however, the only way to detect this type of diversion is tank to tank correlations. This adds significant error because each residual of these correlations will now include systematic tank error. This can be compounded in the event of draining a tank into multiple secondary tanks but is ultimately a solvable problem. For the replacement scenario, multi-tank, multi-variable analysis is the only detection that can be used.

APPENDIX B

Analysis of Safeguards Envelope

Theory

In reality, data always has noise, and due to this noise detecting small diversions is often difficult. To an approximation, we can assume that all measurements take the following form:

$$y_{measured}(t) = y_{true}(t) + \varepsilon_{calibration} + \varepsilon_{measurement} \qquad (1a)$$

where $\varepsilon_{calibration}$ is the calibration error and $\varepsilon_{measurement}$ is the measurement error. Calibration error is due to the non-perfect tuning of the measurement device and is usually a static additive error. The error however is randomly distributed from one device to another. The more familiar measurement error is that which arises from small fluctuations within the control volume (e.g. miniscule temperature fluctuations, or small movement) and is known to be normally distributed. As Eq. (1a) shows, both errors mask what the true value actually is and can hamper any type of verification process. Indeed, both can also be averaged assuming enough data exists to do so. Unfortunately, that is not the case in most scenarios, including our ICPP data. This is the realm in which statistical tests find their application as they look to the overall data trends to discover any abnormalities. Before tests are created, diversion behavior must first be understood.

Material diversions affect two components of measurement data: residual randomness and deviation from the mean or “expected” value. A residual is defined as the difference between the measured value and the true value where $y_{true}(t)$ would be an exact analytical value.

$$y_{residual}(t) = y_{measured}(t) - y_{true}(t) = \varepsilon_{calibration} + \varepsilon_{measurement} \qquad (2a)$$

As Eq. (2a) shows, a measurement residual should be nothing more than a time series of errors with a random distribution and mean of zero. In a diversion case however, the residual would take on an entirely different behavior. First, it is important to understand that abnormal data can be seen as normal data with an added deviation where *diverted(t)* is the nuclear quantity taken as a function of time as shown in Eq. (3a) below.

$$y_{abnormal}(t) = y_{true}(t) + \varepsilon_{calibration} + \varepsilon_{measurement} - \text{diverted}(t) \qquad (3a)$$

If the residual of this curve was computed with respect to the true values of a normal curve, illustrated in Eq. (4a) below, then it becomes obvious that the residual of an abnormal data curve is just a normal residual, such as Eq. (2a), but with an added non-random and/or non-zero mean function.

$$y_{abnormal}(t) - y_{true}(t) = \varepsilon_{calibration} + \varepsilon_{measurement} - \text{diverted}(t) \qquad (4a)$$

In other words, to determine whether or not a tank has been tapped, one simply needs to look at the residual of its data; if the residual has neither a purely random distribution nor a zero mean, then assume that a diversion has taken place. These tests can be performed with hypothesis Z-testing, standard deviation calculation, or cumulative sum examinations.

Unfortunately, detection with the above methodology is difficult for two reasons: not knowing $y_{true}(t)$, and having sparse data. Computing the most accurate residuals requires knowing before-hand what $y_{true}(t)$ is, which is technically impossible. In fact, knowing it would imply perfect measurements and make this entire statistical process pointless. However, what is known is the historical data, which tells what the measurement "ought" to be. With that, it becomes feasible to make good *approximations* of $y_{true}(t)$, especially with good fitting techniques. One must take caution, for approximations can be too uncertain if the base data is too sparse. Even the tests themselves can be misleading if not enough information is present. Again, advanced statistics become useful. Numerous techniques have evolved which take advantage of sparse data and create reliable models to work with (e.g. Principle Component Analysis, Least-Squares Fit, Student's T-testing). With both reliable historical data and advanced statistics, it becomes very possible to distinguish abnormal behavior from normal operating conditions.

Numerical theory aside, some important notes must be mentioned about the testing scenario. The setting involves a reprocessing facility tank filling and flushing a (assumed) homogenous nuclear material solution. Measurements of the solution's level, density, and temperature (LDT) are taken every four minutes and assumed to have a form similar to that of Eq. (2), but with an assumed zero calibration error. An artificial diversion is introduced in the same way as Eq. (4) and involves gradually taking 0.5% of the tank (in terms of level) in a linear fashion until flushing is complete. The exact start and stop times of the tank fill and flush is assumed known at all times.

Methodology

The basic premise of the algorithm is to take a trial data set and statistically compare it with a historical set by residual analysis. Before beginning residual construction, the data must first be collected and processed in the appropriate manner. A fully transient state was first isolated within the data logs and its LDT information was extracted into a data array with MATLAB. Because the state-of-the-art measurement systems ICPP had at the time, there was very little error within the data itself. Therefore, it was assumed that this information represented the "true" values, henceforth called the true curve, with which to build our simulated, noisier measurements.

In order to fully test the algorithm capabilities, a total of three curves were created: a historical curve, a normal trial curve, and a diversion trial curve. The latter two were meant to test the FAR and PND respectively. To build the "historical" curve, Gaussian noise with a standard deviation of 0.2% was added to the true set to simulate measurement error and labeled accordingly. To create a normal trial curve, henceforth called the normal curve, true curve values were again taken and similar noise was added. Creating the diversion trial curve, henceforth called the diversion curve, followed a similar process, but this time with a linearly increasing diversion function that peaked with a value of 0.5%.

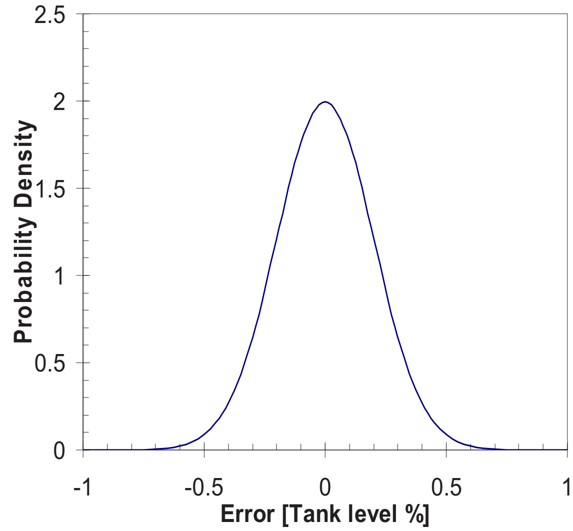

Figure 2. Probability density of measurement error. In this case, it is for the tank level.

Once the three simulated curves were created, Kernel regression was performed on the historical set to later approximate residuals. Kernel regression is a powerful state estimation technique designed to fit an approximate curve to noisy data. Unlike most familiar regression techniques, kernel regression is non-parametric and does not actually make any initial assumptions about the shape of the curve. Instead, it applies a Gaussian weight function centered at each data point and gives each neighboring point a contribution that is proportional to their distance.

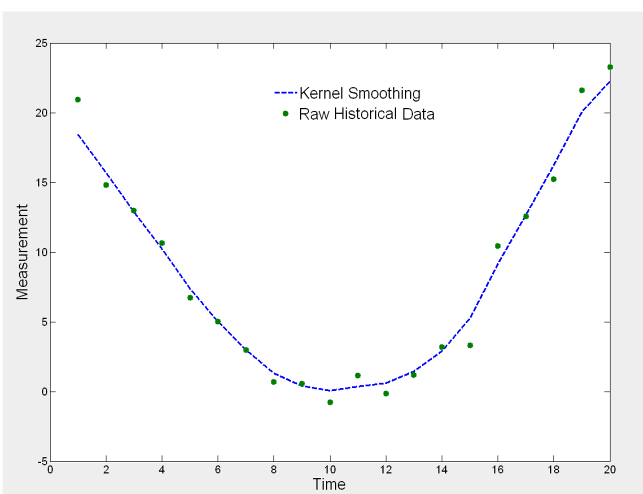

Figure 3. Kernel Smoothing on fake historical data. The example above is a noisy quadratic.

The degree of fitting is also a user-set parameter, called the kernel bandwidth. Too low of a value connects the dots poorly, while one too high will "over fit" and produce large errors. This is one of the parameters that can be optimized in the algorithm for best performance.

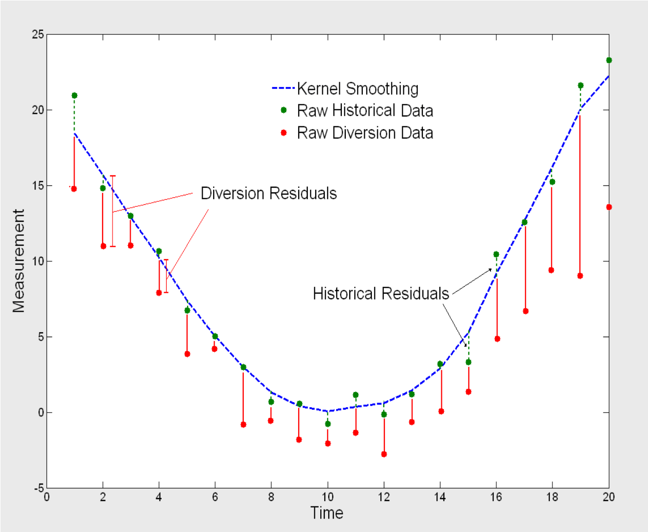

Figure 4. Diagram showing residual analysis with historical data and diversion data.

Once the kernel smoothed historical curve is obtained, the difference between that curve and the two trial curves (normal and diversion curve) give each trial curve their respective residual approximations. This is done by simply subtracting the raw data from the kernel smoothed curve for both the historical and trial case.

With the computed residuals, hypothesis Z-testing is then used to test for the criteria of randomness. To be considered normal, the residuals must have a mean of zero and a standard deviation similar to that of the historical residual (historical data minus kernel smoothed data). The resultant probability reveals how well the trial residual follows the stated constraints and can be compared to a threshold for judgment.

In addition to the Z-test, a cumulative sum threshold test was also implemented in order to measure mean and deviation behaviors in a way that the Z-test does not.

To this end, the residual vector components were each taken to the 1.5 power in order to better distinguish outliers (an L 1.5 norm). Then, a summation of the residual was taken and compared to a threshold to determine abnormal trends. Both these tests were used in OR fashion (if *either* test dismissed a case as a diversion, that result must be recorded).

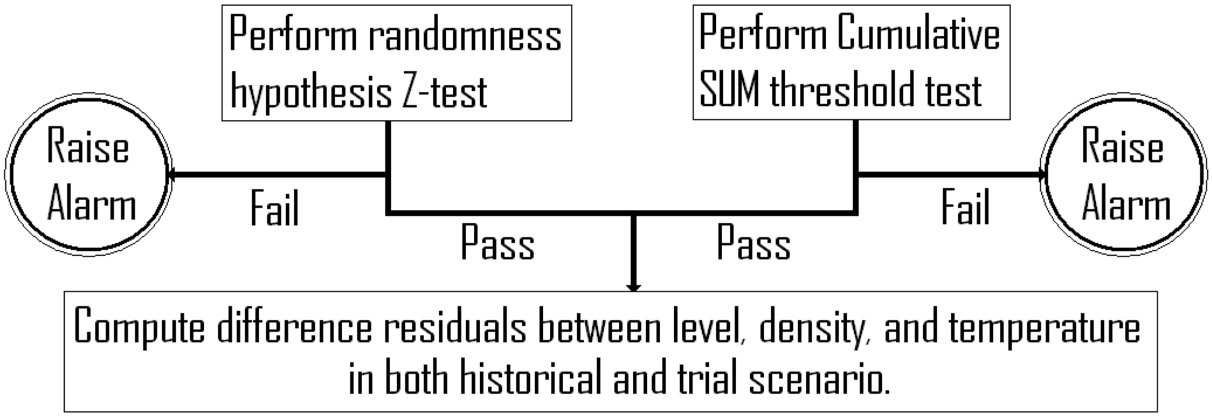

Figure 5. Flow-chart of the test priorities and the order in which they are executed, as well as how their results are interpreted.

Moving into the multi-variate tests, a difference residual comparison between, say, the level curve was compared against that of the density and temperature curves. This was done by computing the residuals for each individual curve with the same methodology as the one-variable scenario, and then taking the residual difference between curves. This was done to compute an “effective” difference between the two curves in terms of a normal distribution. Having these residuals, the same Z-test and cumulative sum threshold tests from before were used, but this time a stricter criterion was used to decide final judgment on a trial case. Because multi-variate calculations add more noise than a single, both the comparison between, say, a trial level curve and a density curve as well as a temperature must agree in result (there must be a majority rule). This result then was also used in areas where the one-variable tests were unable to detect anything abnormal. This way, the baseline PND could only be decreased at a cost of a small sacrifice of the FAR. A Chi-Square test was expected to replace the Z test for randomness, but the additional parameter could not be established.

Finally a *transient* operated at half speed is simulated in order to grab more data points. This was primarily to show that with more data points, and assuming easy diversion detection at steady-states, the FAR and PND values can be markedly decreased and the MBP should increase. To simulate half-speed at the transients, linear interpolation was done with the true curve between data points and applying the algorithm to the ‘new’ data to obtain new FAR, PND, and MBP values.

REFERENCES

[1] Metcalf, Richard; Bevill, Aaron; Bean, Robert. "Safeguards Envelope Progress FY08" INL/EXT-08-14915; September 2008.
[2] Cao, Duc, Metcalf, R. M. ; "Safeguards Envelope Progress FY09" September 2009.
[3] Burr, Tom L.; Wangen, Lawrence E.; Mullen, Mark F.; "Authentication of Reprocessing Plant Safeguards Data through Correlation Analysis"; LA-12923-MS (ISPO-374); (1995)
[4] Burr, Tom; Wangen, Larry; "Development and Evaluation of Methods for Safeguards Use of Solution Monitoring Data" LA-13185-MS, (LSPO-389), 1996
[5] Burr, Tom; Wangen, Larry; "Enhanced Safeguards Via Solution Monitoring" LA-13186-MS, (ISPO-390), 1996
[6] Burr Tom L.; Coulter C. Alton; Wangen, Larry E.; "Benchmark Data for a Large Reprocessing Plant for Evaluation of Advanced Data Analysis Algorithms and Safeguards System Design"; LA-13414-MS (ISPO-397); 1998
[7] Burr, Tom; Jones, James; Wangen, Larry. "Multivariate Diagnostics and Anomaly Detection for Nuclear Safeguards" LA-UR 94-2292
[8] Boland, T.J.; Deveraux, H.R. "Precision Level/ Density Scanner: Advanced Safeguards Systems Development for Chemical Processing Plants"; Report: ENICO-1080; January 1982
[9] Garcia, Humberto E.. "Integrating physical modeling, neural computing, and statistical analysis for on-line process monitoring." Nuclear technology 141.1 (2003):69-77
[10] Garcia, Humberto E.. "Operational analysis and improvement of a spent nuclear fuel handling and treatment facility using discrete event simulation." Computers & industrial engineering 38.2 (2000):235-249.
[11] Morman, J.A.; Reifman, J.; Vitela, J.E.; Wei, T.Y.C.; Applequist, C.A.; Hippely, P.; Kuk, W.; Tsoukalas, L.H.; "IGENPRO Knowledge Based Digital System for Process Transient Diagnostics and Management"; CONF-9709106; 1997