Application of Process Monitoring to Anomaly Detection in Nuclear Material
Processing Systems via System-centric Event Interpretation of Data from Multiple
Sensors of Varying Reliability

Humberto E. Garcia¹, Michael F. Simpson², Wen-Chiao Lin¹, Reed B. Carlson¹, Tae-Sic
Yoo¹

¹Idaho National Laboratory, Idaho Falls, Idaho 83415

²University of Utah, Department of Metallurgical Engineering, Salt Lake City, Utah
84112

Abstract

In this paper, we apply an advanced safeguards approach and associated methods for process monitoring to a hypothetical nuclear material processing system. The assessment regarding the state of the processing facility is conducted at a system-centric level formulated in a hybrid framework. This utilizes architecture for integrating both time- and event-driven data and analysis for decision making. While the time-driven layers of the proposed architecture encompass more traditional process monitoring methods based on time series data and analysis, the event-driven layers encompass operation monitoring methods based on discrete event data and analysis. By integrating process- and operation-related information and methodologies within a unified framework, the task of anomaly detection is greatly improved. This is because decision-making can benefit from not only known time-series relationships among measured signals but also from known event sequence relationships among generated events. This available knowledge at both time series and discrete event layers can then be effectively used to synthesize observation solutions that optimally balance sensor and data processing requirements. The application of the proposed approach is then implemented on an illustrative monitored system based on pyroprocessing and results are discussed.

Keywords: advanced safeguards; process monitoring; system-centric anomaly detection; integrated time- and event-driven analysis; nuclear fuel cycles;

1. Introduction and Motivation

It has been widely accepted that process monitoring (PM) of nuclear facility operations can significantly improve the detectability of facility misuse and/or material diversion activities. In order to monitor and assess the safeguards health of complex nuclear systems, it is often the interest to promptly detect the occurrence of abnormal (e.g., undeclared or unauthorized) operations, even before they may be discovered by utilizing the traditional nuclear material accountancy (NMA) approach. For example, one may be interested in how many times a particular material is routed to a certain workstation. It may also be useful to detect and record when a given unit operation exhibits abnormal behavior due to excessively high concentration of a certain chemical compound, for example, that may be a precursor of a nuclear proliferation-related activity—material diversion.

Pyroprocessing is an emerging technology that is currently being developed for potentially closing the back end of the nuclear fuel cycle. It features electrochemical separation processes at high temperatures with molten salt electrolytes. Its compact facility footprint makes it very appealing, especially for integrated or on-site reprocessing of used nuclear fuel from a fast reactor. The Republic of Korea (ROK) is currently researching pyroprocessing as its leading candidate technology to reduce the volume of its spent pressurized water reactor fuel and recover U/TRU for future fabrication of fast reactor fuel. From a policy perspective, the problem is that ROK is not a nuclear weapons state and has an active nuclear cooperation agreement with the United States that forbids pyroprocessing or any other reprocessing technology be applied to spent nuclear fuel within ROK borders. This is not entirely a stalemate situation, however, as there is currently a 10-year Joint Fuel Cycle Study (JFCS) being carried out by U.S. and ROK in collaboration to evaluate various aspects of the feasibility of using pyroprocessing for this purpose. One of the most important issues being studied is the safeguardability of the process. A positive finding from this study may pave the path for ROK receiving consent from the United States to build pyroprocessing facilities and operate them

for treatment of their spent fuel. NMA is viewed as being very difficult to implement in a pyroprocessing facility because of (1) lack of input accountability, (2) high accumulation of fissile material in process equipment, and (3) the heterogeneity of the process materials and products [1]. The goal of closing the plutonium mass balance within a significant quantity (8 kg) is daunting at even sub-commercial scales. Technology for PM as a supplemental safeguards tool is being studied both within and outside of the JFCS. A new term, Signature Based Safeguards (SBS), has been used in reference to applying PM for detection of misuse of pyroprocessing systems and perhaps support NMA in a complementary fashion [1, 2]. But there have been little specifics for how SBS would be implemented or how PM would otherwise be used for safeguards. The scope of this paper includes the introduction of an algorithm for using PM to detect misuse in a pyroprocessing system. A simplified system has been constructed that includes three key pyroprocessing unit operations—oxide reduction, electrorefining, and electrowinning. The chemical separations achieved by each unit operation are included in the model. Several sensors based on mature technology are included as well. The model generated time-driven and event-driven dynamics which are analyzed by the algorithm to detect anomalous process events. The model is evaluated up against three hypothetical process anomalies with encouraging results. This paper sets the groundwork for a future study involving analysis of real operational process data.

2. Methodology/Framework for Anomaly Detection

The detection and interpretation of special events or anomalies may be accomplished without much difficulty if sensors for exactly identifying them are available. However, relying on a single sensor to provide the required detection sensitivity is difficult, as it may require costly or unattainable sensor performance. For instance, measuring the concentration level of certain material in a given flow stream may not be possible at the required sensitivity level for facility misuse detection. In these cases, a solution may be to combine the data collected from a system of sensors to attain the necessary level of confidence for declaring the

potential occurrence of an event under consideration. This integration of data collected from multiple sensors may be accomplished in both space and time. Under space-only data integration, data collected from distinct distributed sensors are integrated at each time period and an assessment of the plant condition is determined at that moment. Under time-only data integration, data generated from sensors collected at different time periods are integrated to compute the plant state at that moment. Under either case, models are needed that relate in space, time, or both these distributed information. Such models may be derived from physical or/and logical relationships, for example.

While space data integration has received significant attention for many years, time data integration is a relatively new area of research and development. This is mainly due to the following factor. For implementing space data integration, what is mainly needed is a "correlation" model that indicates how the correlation (or lack) of two or more sensor data sets supports a given hypothesis. This correlation model may be derived from first principles knowledge or statistical analysis, for example. Examples of typical techniques used here include kernel-based regression algorithms, multivariate statistical methods, principal component analysis, and neural computing. On the other hand, for implementing time data integration, a model of the system operations is needed characterizing how events and operations occurred and conducted at different time periods relate among them. Such operations models are often formulated as discrete event system (DES) models (e.g., finite state machines, Markov models, and stochastic automata), where the passing of time and change in plant state is due to the occurrence of "events". An event may indicate "tank is full," "concentration of given component is high," or "unusual pattern of gamma-emitting contaminants in given stream". The area of DES theory and application is a relatively recent development, hence the limited utilization of time data integration. However, as the more information one uses should lead into lessen requirements for sensor performance, a hybrid monitoring solution is suggested in this paper, where not only space but also time data integration is conducted in order to assert the occurrence of special (e.g., abnormal) events and

behaviors. Under a hybrid space-time process monitoring (PM) approach, not only current but also past plant behaviors are effectively integrated and used during decision-making. This information-rich approach for PM and decision-making demands for the implementation of system-centric solutions, where systems of sensors, as opposed to single sensors, are used to assert plant conditions and declare the occurrence of facility misuse or material diversion.

In view of the above rationale, system-centric observation platforms for PM are here suggested to effectively monitor large facilities despite (often unavoidable) sensor unavailability and unreliability issues. An architecture for implementing system-centric integrated observation platforms is illustrated in Figure 1. In order to facilitate its implementation, particularly when addressing complex systems such as nuclear reprocessing facilities, rigorous algorithms are needed for not only synthesizing but also optimizing observation platforms. Under the former, it is assumed that the particular sensor configuration (e.g., set of sensors) is given and the task is to design optimized data analysis algorithms (i.e., diagnosers) to meet the required observability performance. Under the latter, not only optimal sensor configurations but also their corresponding data analysis algorithms are selected to meet the required observability performance, under given constraints.

For large throughput nuclear facilities, such as commercial reprocessing plants, it is difficult to satisfy the IAEA safeguards accountancy goal for the detection of one significant quantity of material diversion within the specified time period. Several recent publications have demonstrated, although via simulations and still tentative, the safeguards benefit from using PM on nuclear facilities, as a complementary safeguards measure to NMA (e.g., [3, 4, 5, 6]). One main interest under PM is to detect operational anomalies as the process progresses rather than detecting them by measuring mass balance inconsistencies resulting from abnormal activities. Anomalies of interest for detection may represent undesired plant operations and are defined by specific signatures, such as patterns of events occurring at different locations of the monitored facility and at different stages of operations. PM

essentially consists of utilizing process control measurements to detect abnormal plant operations. Under PM, sensors are deployed to observe unit operations (UO) comprising the monitored facility. The sensor data are collected by unattended on-line/in-vessel MMS, which may have automated data collection and computer-assisted data analysis capabilities. The data collected and/or analyzed by MMS are then processed at the system level by a Data Integration and Interpretation (DII) module to detect and track the occurrence of anomalies of interest. The objective of a DII module for PM is to satisfy user specified observation requirements of anomalies that are of concern. Because data collected from multiple sensors may be partial and unreliable, a DII module needs to be capable for operation under partial and unreliable information. As anomalies may often be defined as patterns of events occurring at different unit operations and time instances, a DII module also needs to recognize anomalies defined in this manner, where the constituting events may span among several unit operations and occur apart in time. Other challenges include the need to distinguish innocent anomalies (e.g., arising from instrument malfunction) from serious anomalies related to facility misuse or nuclear material diversion.

Figure 1: Monitoring architecture

2.1 *Theoretical Developments*

As the monitoring system considered in this work is designed to accommodate facility operations characterized as discrete-event dynamical systems (DEDS), theoretical work considering the monitoring architecture illustrated in Figure 1 for event occurrence analysis is reviewed next. An observation platform consists of sensors generating observed events and a diagnoser that analyzes them. Various algorithms for implementing this diagnoser have been developed. In particular, the work by Sampath et al. [7] deals with detection of special events assuming that a finite-state automaton describes the DEDS, sensors are reliable, and failures are permanent. Subsequent extensions of this work address this event detection problem under sensor unreliability and stochastic conditions in DEDS [8–15]. Counting of the occurrences of intermittent or non-persistent faults is addressed in several papers [16–22]. In particular, the issue of detecting whether a resetting has occurred is addressed by Cotant et al. [16]. Several notions of diagnosability are introduced in Jiang et al. [17] that capture the various counting capabilities of special events. Event counting assuming a deterministic finite-state automaton with partial observations is addressed by Garcia and Yoo [18]. While a deterministic counting strategy for accommodating stochastic automata with unreliable observations is presented in [19], algorithms are reported in [20–22] that fully utilize the probabilistic aspects of stochastic automata.

The work mentioned above focuses on developing diagnoser algorithms for analyzing the behavior of DEDS (e.g., detecting special event occurrences) observed through sensor configurations. While the development of effective algorithms for data integration and analysis is important, it is equally crucial to investigate techniques for optimizing the sensor configurations used for generating data, as

their relative costs may vary with the number of sensors deployed, their quality, their impact on operation, and difficulty of installation, for example. In particular, sensor configurations that cost more may include more sensors and of sensors with better quality, and, typically, they give rise to better diagnoser performance. Considering the diagnoser introduced in Yoo and Garcia [20], sensor configurations optimization algorithms are reported in Lin et al. [23] that balance sensor cost and diagnoser performance. This work is extended in Lin et al. [24] to optimizing observation platforms using the diagnoser developed in [21]. The optimization problem considered in [23, 24] falls into the category considered in [25], namely, that it is computationally hard, and an observation platform optimization algorithm based on heuristics is consequently used for solving it. The algorithms developed there use two greedy heuristics, one myopic and another that considers projected sensor performances. The latter heuristic is similar to the approach for optimal sensor selection developed for heterogeneous sensor networks reported in [26].

Other work on optimal sensor selection for DEDS should be cited [25, 27–29]. While Yoo and Lafortune [25] address computational issues regarding sensor selection for DEDS modeled as finite automata, Aguirre-Salas et al. [27,28] consider optimal sensor selection for satisfying observability properties in Petri nets. Similarly, Jiang et al. [29] considers the sensor optimization problem assuming that sensors are reliable. Optimal sensor selections for supervisory control are addressed in Rohloff et al. [30]. This important problem of selecting optimal sensor configurations for achieving efficient fault detection and isolation (FDI) performance may be formulated using other mathematical frameworks different from DEDS. For example, synthesis and optimization of a monitoring system based on Bayesian belief networks is reported in [31]. On another hand, a graph-based approach is presented in Li et al. [32] for optimally allocating sensors that ensure fault observability and resolution to a maximum possible extent and applied to fault diagnosis of helical coil steam generator systems in pressurized water reactors.

2.2 Monitoring Architecture

As shown in Figure 2, a monitored facility is defined here as a system consisting of multiple unit operations. Each of these units is monitored by sensors, which generate data. Process data may be time series (e.g., those retrieved from process sensors) or discrete observations (e.g., those generated by portals/seals). To extract information, this data is typically processed using data analysis algorithms (e.g., principal component analysis, kernel- regression, neural computing, hypothesis testing) and mapped into higher level qualifiers or “events” characterizing the particular conditions of the monitored process. For example, a large concentration in a process stream of a certain chemical compound exceeding a given threshold can be interpreted as the occurrence of an anomalous event. Discrete observations (i.e., events, scores, evidences) may also be generated by applying more elaborate approaches relying on data generated from multiple sensors, such as those used in material accountancy and solution monitoring, where a residual is generated from solving given conservation formulations and then statistically analyzed. This characterization can be used to generate a descriptive event of the particular situation (e.g., large mass unbalance).

Figure 2. Implementation of process monitoring-based data integration and interpretation

Mapping data to higher level of abstraction for decision making has numerous advantages including avoiding information congestion at the operator/inspector level and moving decision making about the overall state of the monitored facility to a higher intelligence layer. In particular, a system-centric approach is expected to: (a) allow more efficient integration of distributed information, (b) improve sensitivity of anomaly detection, (c) describe emergent behavior of complex systems, and (d) produce tangible intelligent information for operator/inspector monitoring and decision making. Given these events/scores/evidences, a data analysis module with event-driven dynamics is utilized here to integrate and interpret observed events. Mathematical frameworks that can be used include probabilistic reasoning (e.g., Bayesian Belief Networks) and DEDS, with the latter being the paradigm here considered.

Figure 2 also illustrates the basic idea of a DII, where time series sensor observations of the monitored facility are first processed by a MMS and reported as events to a DII module, which embeds an operations model of the monitored process and information on deployed sensor reliability characteristics. Based on the reported observations collected from a system of sensors, the DII module detects and counts the occurrences of the anomalies of interest. In particular, let i denote the anomaly defined by the event anomaly pattern P_i , thus its occurrence is associated with the diversion of a certain amount of plutonium in a waste or product stream. As the observations, y_i , $i = 1, 2, \dots$, become sequentially available, the DII module estimates the number of times i has occurred. When the m^{th} observation becomes available, this count estimate $c(m)$ is calculated based on the observations y_1 through y_m .

2.3 Methodology for Anomaly Detection

Diagnosers in this paper are designed based on the monitoring architecture shown in Figure 1. The monitored engineering system is viewed as a discrete-event dynamical system (DEDS), which generates true events corresponding to its operations (e.g., flow of a particular material from one unit operation to another). These generated events may be then observed by sensors, accordingly reporting observed events to the diagnoser. The diagnoser processes this information in order to determine the possible occurrence of certain special events of interest, such as faults/failures. In particular, the monitored plant in Figure 1 is modeled as a stochastic automaton,

$$SA = (X, S, a, \rho_0) \quad (1)$$

where $X := \{x_1, x_2, \dots, x_{nx}\}$ is the finite state space, $\Sigma := \{\sigma_1, \sigma_2, \dots, \sigma_{n\sigma}\}$ is the set of events, and $\pi_0 := \{\pi_0(x_i) : x_i \in X\}$ is the initial probability distribution of the system. The state transition probability function a is defined as $a : X \times \Sigma \times X \rightarrow [0, 1]$, where, $a(x_i, \sigma, x_j)$ denotes the conditional probability that, given the system is in state $x_i \in X$, $\sigma \in \Sigma$ occurs and transitions the system to state $x_j \in X$. We also assume that the system is live. The interest is to detect and count the number of occurrences of a special event $f \in \Sigma$, which may represent a fault or an anomaly. As shown in Figure 1, diagnosers are used to accomplish this task using observations from unreliable sensors. Assume that there is a given pool of available sensors, $\{s_1, s_2, \dots, s_p\}$, from which sensors are chosen for observing the monitored DEDS. Let $\Delta := \{y_1, y_2, \dots, y_{ny}\}$ be the set of distinctive observation symbols generated from a sensor configuration $S \subseteq U$. We denote the set of observation symbols at the sensor outputs as $\Delta_* := \Delta \cup \{\epsilon\}$, where the symbol ϵ indicates that an event has been executed but no observation is reported. The event output function $b : \Sigma \times \Delta_* \rightarrow [0, 1]$ satisfies the following: $\forall \sigma \in \Sigma$,

$$b(s, \epsilon) + \sum_{i=1}^{n_y} b(s, y_i) = 1 \quad (2)$$

The functional value $b(\sigma, y)$ is the conditional probability of having output $y \in \Delta_*$ when the system executes event $\sigma \in \Sigma$. For example, assume that the given monitored system generates an event σ . Characterizing $b(\sigma, \epsilon)$, $b(\sigma, \sigma)$, and $b(\sigma, \delta)$ (with $\sigma \neq \delta$) equal to 0.1, 0.7, and 0.2 indicates the probabilities of misdetection, correct classification, and misclassification, respectively, for this sensor. The set of observation symbols, Δ , and the function b both depend on the sensor configuration, S .

A stochastic diagnoser (SC) for the monitoring architecture in Figure 1 is designed based on the model of the DEDS (i.e., the stochastic automaton in Equation (1)), the sensor observation model (i.e., the output function b in Equation (2)) corresponding to the given sensor configuration, and the special event, f , to be counted. As the observations, y_i , $i = 1, 2, \dots, m$ become sequentially available, the SC estimates the number $c(m)$ of times f has occurred. Let $N_m(f)$ denote the random variable of the number of times f has occurred up to the m^{th} available observation and $\{Y^i\}_{i=1}^m = Y^1 Y^2 \dots Y^m$ denote the sequence of random observations. The goal of the SCs developed in references 18 and 20 is to calculate $c(m)$ as a function of $\{Y^i\}_{i=1}^m$ such that the mean squared error,

$$E_{\Theta}^{\Theta} \left(N_m(f) - c(m) \right)^2 \quad (3)$$

is minimized. The conditional expectation,

$$c(m) = E_{\Theta}^{\Theta} N_m(f) | \{Y^i\}_{i=1}^m \quad (4)$$

is a solution to minimizing (3). The SCs utilize a recursive algorithm to calculate Equation (4). Specifically, the following information states are considered:

- $Prob(X_m | \{Y^i\}_{i=1}^m)$: conditional probability of system state just after the m^{th} observation becomes available given the observation history up to the m^{th} available observation;
- $E[N_m(f) | X_m, \{Y^i\}_{i=1}^m]$: conditional mean of the number of times f has occurred given the current system state just after the m^{th} observation becomes available and given the observation history up to the m^{th} available observation;
- $E[(N_m(f))^2 | X_m, \{Y^i\}_{i=1}^m]$: conditional second moment of the number of times f has occurred given the current system state just after the m^{th} observation becomes available and given the observation history up to the m^{th} available observation;

As shown in reference 20, these information states can be calculated recursively as the observations sequentially become available. Regarding the occurrence of the given special event, the conditional variance is the mean squared error between the true count and the count estimated by the SC. It also represents the uncertainty of the estimated special event count, and will be used to determine the performance of the SC for a given $S \subseteq U$. Based on simulations, the convergence of the variance calculated after the DEDS has executed n events holds for the practical cases considered here and in previous work [19–20]. If the true number of special event occurrences is known (e.g., in off-line simulations), it is possible to evaluate the performance of SCs by comparing this number and the number estimated by the SC. In particular, given sensor configuration S , we define the average normalized error (ANE),

$$ANE(S) = \frac{1}{n} \sum_{k=1}^n \frac{|k - \hat{T}(k)|}{k} \quad (5)$$

where n is the number of occurrences of the anomaly at the end of the simulation. Moreover, $\hat{T}(k)$ denotes the estimated number of its occurrence just after the given special event has been executed the k th time. The ANE metric can be used to compare performances of SCs.

The description of the SCs above considers only one special event, f . Extension in the sense of Sampeth et. al. [7] to estimate the occurrences of multiple events is straightforward, where each formulation above is then extended to consider multiple special events f_i . If the anomaly considered is defined as a sequence of events, diagnosers considering part of the sequence of events can be used.

2.4 Methodology for Sensor Optimization

Consider the monitoring architecture in Figure 1 and that there are N_s special events. For a sensor configuration, $S \subseteq U$, its sensor configuration cost is

$$ct(S) = \bigcirc_{s \in S} ct(s) \quad (6)$$

where $ct(S)$ and $ct(s)$ denote the costs of the configuration S and a particular sensor $s \in S$, respectively. Note that $ct(s)$ is a compound measure for sensor s that includes metrics such as monetary cost, vulnerability, and intrusiveness of operation. The sensor optimization problem here is to find a sensor configuration that achieves an optimal balance between cost of the configuration and performance of the SC, while satisfying constraints on the sensor configuration cost and SC performance. For this purpose, the following loss index is considered:

$$I(S) = \bigcirc_{i=1}^{N_s} c_i \times \beta^i(S) + c \times ct(S) \quad (7)$$

for $S \subseteq U$, where $c \geq 0$ and $c_i \geq 0$, $i = 1, 2, \dots, N_s$, are weighting factors that indicate the importance of each term in $I(S)$. Note that optimization measures are given by $\beta^i(S)$ and $ct(S)$, while the optimization criterion is given by Equation (7). The sensor optimization problem is formulated into the following minimization problem for $I(S)$: Find

$$S^* := \arg \min \{I(S) : S \subseteq U\} \quad (8)$$

subject to

$$b^i(S) \in b^{i*}, \text{ for } i = 1, 2, \dots, N_s \text{ and } ct(S) \in ct^* \quad (9)$$

where

- ct^* is the maximum cost desired for S ;
- β^{i*} indicates the maximum normalized variance tolerable for the diagnoser in estimating the number of occurrences of special event f_i .

A greedy algorithm may be used to solve the sensor optimization problem formulated in Equation (8). In particular, this algorithm first sets the current sensor configuration, S_{cur} , to $\emptyset$. Then, sensors are added to S_{cur} in a step-by-step manner. The selection of the sensors in each step is based on the following criterion: Find

$$s^* = \arg \min_{s \in (U \setminus S_{cur})} \left(\sum_{i=1}^{N_s} c_i \times b^i(S_{cur} \cup \{s\}) + c \times ct(S_{cur} \cup \{s\}) \right) \quad (10)$$

subject to

$$ct(S_{cur} \cup \{s\}) \in ct^* \quad (11)$$

The sensor, s^* , is the newly selected sensor to be added to S_{cur} . Note that Equation (10) is solved by searching through $s \in U \setminus S_{cur}$. Furthermore, Equation (11) means that sensors costing more than ct^* are not considered. This criterion is based on the instantaneous performance improvement of the candidate sensor and is myopic. Stopping criteria for this greedy algorithm are based on considering whether the constraints in Equation (9) are satisfied and on whether further reduction of the

loss in Equation (7) can be achieved in the current step. The greedy algorithm described above does not often perform well. In general, performance of the greedy algorithm can be improved if we modify the sensor selection criterion to take into account the projected values of the loss index (obtained by adding fictitious sensors with the same statistical performances and costs as the candidate sensor) and how far we are from satisfying or violating the constraints in Equation (9).

In particular, a heuristic search algorithm is proposed for solving Equation (8) that utilizes a sensor selection criterion that minimizes the projected optimal values of the loss index. To this end, the projected loss index and the projected optimal loss index are described first. Let S_{cur} stand for the current chosen sensor configuration, and consider a candidate sensor, $s \in U \setminus S_{cur}$. Let $S_{cur} \cup \{s_{1:k}\}$, $k \geq 1$, represent a fictitious sensor configuration, $S_{cur} \cup s \cup s \cup \dots \cup s$, where the sensor, s , is added k times to the current sensor configuration. This fictitious configuration represents the situation, where k sensors with the same statistical performances and costs as s are added to S_{cur} . The idea is to minimize the chance of getting stuck at a local minimum by projecting the value of the loss index assuming sensors with the same statistical performances and costs as the candidate sensor are added. Furthermore, let $S_{cur} \cup \{s_{1:0}\}$ denote S_{cur} . The projected cost of $S_{cur} \cup \{s_{1:k}\}$ is given by

$$ct(S_{cur} \cup \{s_{1:k}\}) = \underset{s' \in S_{cur}}{\mathring{A}} ct(s') + k \times ct(s) \quad (12)$$

The projected performance of the diagnoser for $S_{cur} \cup \{s_{1:k}\}$ is denoted by $\beta^i(S_{cur} \cup \{s_{1:k}\})$ for f_i , $i = 1, 2, \dots, N_s$. There is no known closed form expression for calculating the projected performance, and, hence, the heuristic formula,

$$b^i(S_{cur} \cup \{s_{1:k}\}) = \underset{\emptyset}{\mathring{C}} \frac{b^i(S_{cur} \cup \{s\})^{\mathring{D}^k}}{b^i(S_{cur})^{\mathring{D}}} b^i(S_{cur}) \quad (13)$$

is used. Since, $\beta^i(S_{cur} \cup \{s\}) \leq \beta^i(S_{cur})$, Equation (13) indicates that, as more sensors s are used, the projected performance of the diagnoser does not degrade and may improve exponentially. The projected loss index for $S_{cur} \cup \{s_{1:k}\}$ is given by

$$I(S_{cur} \cup \{s_{1:k}\}) = \prod_{i=1}^{N_s} c_i \times b^i(S_{cur} \cup \{s_{1:k}\}) + c \times ct(S_{cur} \cup \{s_{1:k}\}) \quad (14)$$

To calculate the projected optimal loss index for each candidate sensor, $s \in (U \setminus S_{cur})$, let

$$k_u(s) = \arg \max_{k \geq 0} \{ct(S_{cur} \cup \{s_{1:k}\}) \leq ct^*\} \quad (15)$$

i.e., $k_u(s)$ is the maximum number of s that can be added to S_{cur} without the cost of $S_{cur} \cup \{s_{1:k}\}$ exceeding ct^* . The value $k_u(s)$ can be found by searching through $k \geq 0$ using Equation (12). Likewise, let

$$k_l(s) = \arg \min_{k \geq 0} \{\prod_{i=1}^{N_s} b^i(S_{cur} \cup \{s_{1:k}\}) \leq b^{i*}\} \quad (16)$$

i.e., $k_l(s)$ is the least number of s that should be added for the performance of $S_{cur} \cup \{s_{1:k}\}$ to meet all specified β^{*}_i . The value $k_l(s)$ can be found by searching through $k \geq 0$ using Equation (13). If $k_l(s) \leq k_u(s)$, define

$$k^*(s) = \arg \min_{k \in [k_l(s), k_u(s)]} \left\{ \prod_{i=1}^{N_s} c_i \times b^i(S_{cur} \cup \{s_{1:k}\}) + c \times ct(S_{cur} \cup \{s_{1:k}\}) \right\} \quad (17)$$

i.e., $k^*(s) \in [k_l(s), k_u(s)]$ minimizes the projected loss index in (14) given S_{cur} and s . The value $k^*(s)$ can be found by searching through $k_l(s) \leq k \leq k_u(s)$ using Equations (6) and (13). The projected optimal loss index given S_{cur} and s is then

$$I\left(S_{cur} \tilde{E}\{s_{1:k^*(s)}\}\right) = \sum_{i=1}^{N_s} c_i \times b^i\left(S_{cur} \tilde{E}\{s_{1:k^*(s)}\}\right) + c \times ct\left(S_{cur} \tilde{E}\{s_{1:k^*(s)}\}\right) \quad (18)$$

Finally, the projected optimal loss index sensor selection criterion is formulated as

$$s^* = \arg \min_{\hat{s} \in U \setminus S_{cur}} \left[\sum_{i=1}^{N_s} c_i \times b^i\left(S_{cur} \tilde{E}\{s_{1:k^*(s)}\}\right) + c \times ct\left(S_{cur} \tilde{E}\{s_{1:k^*(s)}\}\right) : s \in U \setminus S_{cur} \text{ and } k^*(s) \text{ defined} \right] \quad (19)$$

The sensor s^* is the newly selected sensor to be added to S_{cur} and can be found by searching through $s \in U \setminus S_{cur}$ using Equations (6), (13), and (17). Additional details of the proposed sensor optimization algorithm can be found elsewhere [23, 24, 25, 26].

3. Description of Pyroprocessing Plant Model

3.1 Overall Monitored Facility

A realistic yet simplified representation of a facility for pyroprocessing spent nuclear fuel has been constructed, including the formulation of operational anomalies of interest (e.g., facility misuse and material diversion scenarios) and sensors for demonstration. The implementation of this dynamic model of a pyroprocessing facility has been completed in order to generate (synthetic) data needed for evaluating proposed system-centric ADI and PM. In particular at a high level of attraction, this facility model involves three (3) unit operations, namely,

oxide reduction (OR), electrorefining (ER), and electrowinning (EW), as illustrated in Figure 3. There are no recycle streams in this model, so material always flows in one direction. It also includes facility modeling (FM) tools for integrating these unit operation (UO) models and implementing diverse operational scenarios.

Figure 3. Simplified pyroprocess flowsheet

3.2 Unit Operation Descriptions

Oxide Reduction (OR) is the chemical process by which oxide fuel is reduced to metallic form. This can be carried out in an electrolytic cell at 650°C with molten LiCl with 1 wt% Li₂O as the electrolyte [33]. The oxide fuel is placed into a basket that is cathodically charged, while oxygen gas bubbles form and evolve from an inert anode. For the model used in this study, it was assumed that up to 100% of the oxides were reduced to metals, proportional to the coulombs of charge passed through the oxide reducer system. To regulate the fission product contamination level in the electrolyte, batches of salt can be removed as needed for disposal and replaced with uncontaminated LiCl-Li₂O.

Starting with spent fuel in metallic (reduced) form, electrorefining (ER) can be used to separate actinides from fission products and other constituents in the fuel [34]. When coupled with oxide reduction, the fuel basket containing reduced fuel is transferred from the oxide reduction unit to the electrorefining unit. Similar to the oxide reduction unit, the ER contains a molten salt electrolyte. When using standard

operating conditions, U from the spent fuel is oxidized to U^{3+} at the anode and reduced to U metal on a steel cathode. After removal from the ER, the actinide product is loaded into a cathode processor (CP), where a high temperature distillation sequence results in separation of the salt and melting of the metal into a high-density ingot [35].

Electrowinning (EW) is the process by which actinides (U/TRU) are extracted from the ER salt [36]. During normal ER operations, plutonium and the other TRU elements are oxidized to form chloride salts and accumulate in the ER electrolyte. To recover these actinides, a liquid cadmium cathode (LCC) is immersed into the salt. Current passed between the LCC and an inert anode results in U/TRU co-deposition in the cadmium and Cl_2 gas evolution on the anode. Provided that there is a mixture of different actinide elements in the salt, electrowinning will not result in recovery of pure plutonium, commonly referred to as a safeguards benefit of the electrowinning approach to actinide recycling.

The above-mentioned three main UOs, along with their processes and operations, were thus integrated into the Pyro Dynamic process model shown in Figure 4. It was put together using the described simulation platform following the design concepts presented above. In particular, it models three primary operations: oxide reduction (OR), electrorefining (ER), and electrowinning (EW). In addition, the model includes two cathode processors (CP), which is considered to be a secondary operation. Collectively, these operations generate sensor data that is sent on to a set of diagnosers via the DAS block. The red blocks and red text in the green blocks show the sensors, which are described in more detail below. The unit operations are represented in Figure 4 as green blocks.

Figure 4: Monitored pyroprocessing facility—Pyro Dynamic Model.

3.3 Sensor Descriptions and Models

While the model tracks all compositions, masses, and other pertinent parameters, it is assumed that a limited type and number of measurements can be made and used to identify anomalies. These measurements are as follows.

- Liquid salt density using a double bubbler device
- Off-gas flow rate or cumulative volume using a gas mass flow meter
- Electric charge passed using the electric power supply
- Solid metal density using a combination of an electric balance and a geometric measurement
- Dose of solid product or waste using radiation detection such as gamma spectroscopy, neutron counting, or calorimetry
- Mass of recovered products using electronic balances

A total of twenty-two sensors were included in the model (see Figure 4). How some of the sensors interface with key unit operations is shown in Figure 5.

For demonstration of the full versatility and benefits of system-centric data integration and interpretation of unattended online measurement and monitoring systems, only currently mature sensor technology was considered. These are listed in Table 1 and can be cross-referenced to variables listed in Figures 4 and 5. As it can be seen from Table 1, these are largely very simplistic sensors including mass flow meters, electronic balance readings, and electric power supply output. Many other innovative sensor concepts are currently being developed for application to pyroprocessing systems, including Laser Induced Breakdown Spectroscopy (LIBS), molten salt voltammetry, and potentiometric sensors. Each of these technologies would likely improve the safeguardability of the system by providing substantially more data on systems and materials in the process, but they are in some cases years away from implementation and were not considered in this analysis.

Figure 5. Sensors for the OR, ER, and ER unit operations

Table 1. Deployed sensors directly measuring process variables

Unit Operation	Sensor Variable	Description	Sensor Technology
Oxide Reduction	VO2	Volume of O ₂ gas released into the off-gas system (liters)	Gas mass flow meter
	QOR	Total coulombs of charge passed through the OR cell (C)	Electric power supply
	PORS	Oxide reduction salt density (g/cm ³)	Double bubbler
	MORF	Mass of oxide fuel (kg)	Electronic balance
	DORS	Dose of waste salt leaving facility	Radiation detector
Electrorefining	QER	Total coulombs of charge passed through the ER cell (C)	Electric power supply
	PERS	Density of molten ER salt	Double bubbler
	CS	Cathode presence sensor	Position switch
Cathode Processor #1 (coupled to ER)	MERS	Mass of returned salt on ER product (kg)	Electronic balance
	MERU	Mass of U in ER product (kg)	Electronic balance
	PERU	Density of U metal in ER product (g/cm ³)	Electronic balance and geometric measurement
	DERU	Dose of ER product	Radiation detector
Electrowinner	QEW	Total coulombs of charge passed through the EW cell (C)	Electric power supply
	VCL2	Volume of Cl ₂ gas released into off-gas system (liters)	Gas mass flow controller
	PEWS	Density of molten EW salt (g/cm ³)	Double bubbler
	MEWS	Mass of EW salt (kg); i=initial, f=final	Electronic balance
	MLCC	Mass of liquid cadmium cathode (kg); i=initial, f=final	Electronic balance
Cathode processor #2 (coupled to EW)	PEWM	Density of EW product metal (g/cm ³)	Electronic balance and geometric measurement
	MEWM	Mass of EW product metal (kg)	Electronic balance
ER or EW	PMW	Density of metal waste from ER or EW (g/cm ³)	Electronic balance and geometric measurement

3.4 Time- and Discrete Event-driven Dynamics

For designing diagnosers, the unit operations were modeled under a discrete event dynamic system framework, with their underlined processes being modeled with time-driven dynamics characterizing the main phenomena governing specific material processing mechanisms. The particulars of which are described below. Please refer back to Figure 4 where the specific blocks are included.

Oxide Reduction (OR) Unit: Before inputs are entered into this unit operation, containers may come in from Dirty OR Salt block to remove salt out of the OR unit. The removed salt is then passed to the output port O5 after Sensor MORS measures its property values. This operation may or may not happen depending on operation procedures. To start the operations, Oxide Fuel is passed into the OR unit from input port I1 after being measured by MORF. Sensors QOR and PORS measure material properties as operations are conducted in the unit. The OR unit then outputs waste gas to output port Off-Gas and Metal Fuel into the ER unit. Properties of the waste gas are measured by sensor VO2.

Electrorefining (ER) Unit: Before inputs are entered into this unit operation, containers may come in from Dirty ER Salt block to remove salt out of the ER unit. The removed salt is then passed to the Electrowinning unit after Sensor MEWSi measures its property values. This operation may or may not happen depending on operation procedures. To start the operations, Metal Fuel is entered from the OR unit, while containers enter from the Cathode Storage. During operations, Sensors QER and PERS measure material properties in the unit. The ER unit then produces Actinide Product to the Cathode Processor, which transfers products to output port O1 after measurements at Sensors PERU and MERU. The Cathode Processor also transfers salt into ER unit for the next batch operation after measurements by MERS.

Electrowinning (EW) Unit: To start the operations, Removed ER Salt (measured by MEWSi) as well as material from Cathode Processor (measured by MLCCi) are entered into the unit. During operations, QEW and PEWS measure material properties. The unit outputs waste gas (measured by VCL2) to output port Off-Gas and Salt Waste (measured by MEWSf) to output O3. LCC Product is then measured by MLCCf before being entered to the Cathode Processor, which, in turn, transfers U/TRU Product to output port O4 after measurements by Sensors MEWM and PEWM. Moreover, depending on system operation procedures, some Salt Waste may be measured by MEWSf before being transferred to output port O3.

Each UO considered has been sized as follows.

- Oxide Reduction: Process 200 kgHM/day using 1000 kg salt
- Electrorefining: Process 200 kgHM/day using 1000 kg salt
- Electrowinning: Process 40 kg of ER salt/day (no holdup)

The 1000 kg salt limit is based on a Pu/U ratio of 1 and 10 wt% U+Pu in the salt. That would result in a maximum of 50 kg of Pu in each ER salt pool. This is the projected Pu limit for the INL Mark-V ER, so it was selected for this model as well. During electrowinning, a maximum actinide metal batch size should be about 4 kg for criticality concerns. If the salt contains 10 wt% U+Pu, that sets the maximum salt mass per EW batch at 40 kg. This would remove 2 kg of Pu per day, which is the equivalent of 56 MTHM fuel/year. All of these calculations assume 80% utilization of the equipment and a single batch per day that the equipment is fully operational.

3.5 Normal Operations

Normal pyroprocessing operations involves reduction of the oxide fuel in the oxide reduction system (OR) followed by U metal recovery in the ER. Salt is removed from the OR and replaced with clean LiCl-Li₂O as needed to reduce concentration of fission products. Salt is also removed from the ER as needed to reduce the

concentration of fission products and TRU actinides. The salt removed from the ER is sent to the EW where U/TRU actinides are separated. The U product stream (O1) consists of pure U metal. The metal waste stream (O2) is actinide-free and consists of noble metal fission products from the spent fuel. The salt waste (O3) is actinide-free and contains active metal fission products from the spent fuel. And the U/TRU product (O4) contains a metallic actinide alloy with U, Pu, and minor actinides. This is considered normal operation of the facility based on the design of the process. The waste is actinide-free and not subject to safeguards. Only one product stream contains plutonium.

For the OR unit operation model, it is assumed that 100% reduction of all oxides can be achieved at a maximum. Reduction fraction is a function of current. The model will assume a set current efficiency (nominally 50%) or allow the current efficiency to vary according to a statistical distribution. It is also assumed that based on voltage limits, the process will automatically turn-off once all of the oxide has been converted to metal. The model calculates amp-hours or coulombs of charge passed and compares that to the mass of fuel in any batch. This will allow for the operator to know the actual current efficiency. It should be noted though that the actual current efficiency might vary from run to run. Active metal fission products are converted to chlorides and allowed to accumulate in the salt phase. A batch of waste salt is to be removed once a certain limit of fission products has accumulated in the salt (nominally 10 wt%).

For the ER model, with each batch of metal fuel immersed in the salt, uranium is collected at the cathode while active metals (rare earths and transuranics) form chlorides and go into the salt.

The oxidation reactions cause a decrease in the concentration of UCl_3 in the salt. It is, thus, necessary to periodically add more UCl_3 to the salt. If a processing limit for the salt is reached (excessive Pu, excessive fission products, or excessive volume), a

batch of salt is then removed from the ER and sent to the electrowinning unit operation.

The primary purpose of the electrowinning unit operation is to extract actinides that can be used for fast reactor fuel. A liquid cadmium cathode is used to stabilize the metallic Pu and other TRUs, which would readily oxidize from a solid cathode surface back to the salt phase.

3.6 Abnormal Operations

Various ways to misuse the designed process were identified. From this assessment came three protracted diversion scenarios of SNM—each involving a different unit operation. Because of the sensitive nature of SNM diversion, specific details of these diversion scenarios are not described here. However, they are all considered to be feasible based on current understanding of pyrochemistry. And they cover a variety of potential pathways for diversion, including waste streams, product streams, and equipment hold-up. They are all also considered to be problematic for detection using NMA-based methods alone and are, thus, ideal for demonstrating the application of PM for pyroprocessing safeguards. Three anomalies were considered during initial evaluations and were simulated via DES. The lack of a published database that describes anomalies should not be considered a flaw in this approach. Such information can be collected within a given nation using subject matter experts and then protected from publication via export control laws. The SBS approach can be perpetually improved by regularly updating the database of theoretical anomalies. To publish such anomalies, however, represents a risk in that it could aid potential proliferators achieve their goals. On a case by case basis, nations may decide to share anomaly databases with other nations to benefit the global non-proliferation objective.

3.7 Anomaly Patterns

An abnormal batch moving through each unit operation may be described as generating anomaly patterns consisting of sequences of “true” events generated at different physical locations and time instances. In the description below, the generation of an event “i” at a given location “j” means that the facility generates the true event “i” (e.g., higher than normal density of U metal product) as an abnormal batch is processed at location “j” (e.g., the CP). Whether this true event is adequately observed or not by the PM-based diagnoser depends on whether there is a sensor (e.g., a densitometer) actually installed at location “j” (e.g., the metal product stream from the CP) and on its associated reliability performance for measuring the specific (physical) quantity (e.g., density) defined by the true event “i”. The task of a PM-based safeguards solution is to infer how many times an “anomaly” pattern may have occurred within a given time window, without necessarily observing, correctly or otherwise, each of its constituting events. The anomalies of interest for detection, described below assuming perfect observations of process variable true values, are denoted by P_1 , P_2 , and P_3 and manifest themselves as a diversion of P_u .

Anomaly pattern P_1 : The anomaly pattern P_1 for detection consists of two parts, A and B:

- Part A: For a given batch of input material,
 - a low (L) ratio occurs between the volume of oxygen off-gas released and the mass of oxide fuel entering the OR AND
 - a high (H) rate of change occurs in salt density AND
 - a low (L) ratio occurs between the coulombs of electric charge passed at OR and the mass of oxide fuel entering the OR AND
 - a low (L) ratio occurs between the coulombs of electric charge passed at ER and the mass of oxide fuel entering the OR AND
 - a low (L) value of mass of U in ER product occurs.
- Part B: A high (H) value occurs in the dose of OR waste salt leaving the facility

Anomaly pattern P₂: The anomaly pattern P₂ for detection consists of two parts, A and B:

- Part A: For a given batch of input material,
 - a low (L) ratio occurs between the mass of returned salt on ER product and the mass of U in ER product AND
 - a high (H) value of density of U metal from ER product occurs AND
 - a high (H) value of dose in ER product occurs.
- Part B: a low (L) value in the density of electrowinning product metal OR a high (H) value of density of the molten salt at EW.

Anomaly pattern P₃: The anomaly pattern P₃ for detection consists of two parts, A and B:

- Part A: For a given batch of input material,
 - a high (H) ratio occurs between the coulombs of electric charge passed at EW and the mass of electrowinning product metal AND
 - a high (H) ratio occurs between the volume of chlorine off-gas released and the mass of electrowinning product metal AND
 - a high (H) ratio occurs between the difference in salt density associated with ER and EW and the mass of electrowinning product metal AND
 - a low (L) ratio occurs between the final and initial sums of mass of electrowinning salt plus the mass of liquid cadmium cathode.
- Part B: A high (H) value occurs in the density of metal waste from ER or EW.

Our anomaly detection rule is simple. If Part A for a given anomaly pattern above occurs for two (not necessarily consecutive) input batches, and, subsequently, its associated Part B occurs, an alarm is triggered signaling the occurrence of the anomaly pattern (P₁, P₂, or P₃). Notice that patterns P₁, P₂, and P₃ are defined by the occurrence of six, five, and five events, respectively. To summarize, Tables 2, 3, and

4 illustrate the definitions of P_1 , P_2 , and P_3 , respectively, where the leftmost columns indicate the events comprising Parts A and B. The columns labeled by D_{ij} , $i, j = 1, 2, 3, 4$, are explained in the next section.

Table 2. Definition of Anomaly Pattern P_1 .

Anomaly #1				D_{11}	D_{12}	D_{13}	D_{14}
<i>One or more of the following</i>							
Volume of off-gas from OR / Mass of oxide fuel into OR	VO2DMF	801	Low	X	X	X	X
Density of salt in OR rate of change	PORSRT	811	High	X	X	X	X
Coulombs of charge in OR / Mass of oxide fuel into OR	QORDMF	821	Low			X	X
Coulombs of charge in ER / Mass of oxide fuel into OR	QERDMF	831	Low		X	X	X
Mass of U product from ER	MERU	91	Low			X	X
<i>Followed by</i>							
Dose of OR waste salt leaving facility	DORS	51	High				X
				2	3	5	6

Table 3. Definition of Anomaly Pattern P_2 .

Anomaly #2				D_{21}	D_{22}	D_{23}	D_{24}
<i>One or more of the following</i>							
Mass of salt returned to ER / Mass of U product	MERSDMU	841	Low		X	X	X
Density of U product from ER	PERU	81	High	X	X	X	X
Dose of metal waste	DERU	191	Low			X	X
<i>Followed by</i>							
Density of EW metal	PEWM	171	Low				X

product							
Density of salt in EW	PEWS	32	High				X
				1	2	3	5

Table 4. Definition of Anomaly Pattern P₃.

Anomaly #3				D ₃₁	D ₃₂	D ₃₃	D ₃₄
One or more of the following							
Coulombs of charge in EW / Mass of EW metal product	QEWDDMM	851	High	X	X	X	X
Volume of off-gas from EW / Mass of EW metal product	VCL2DMM	861	High			X	X
(Density of salt into EW – Density of salt in EW final) / Mass of EW metal product	PEWSDMM	871	High	X	X	X	X
(Mass of salt leaving EW + Mass of LCC final) / (Mass of salt from ER to EW + Mass of LCC initial)	MEWFDMEWI	881	Low		X	X	X
Followed by							
Density of metal waste	PMW	121	High				X
				2	3	4	5

From an implementation perspective, this nuclear fuel reprocessing model is configured into four separate but inter-communicating computational layers, namely, process, sensors, data analysis, and operation. The process layer hosts time-driven models of all processes characterizing physics-based formulations including electrochemical phenomena. This layer also includes modeling constructs to simulate process upsets that may produce abnormal values of process variables even under normal operating conditions. The sensor layer hosts models

characterizing measuring devices that may include instrumentation noise and sensor reliability assumptions regarding probability of misdetection, correct detection, and misclassification. The data analysis layer hosts data acquisition and analysis capabilities that integrate and analyze data collected from distributed sensors and report, using diagnosers, whether an anomaly pattern has occurred. Finally, the operation layer hosts event-driven models of all operations characterizing operational procedures including normal and abnormal operations. These layers were constructed under different modeling and simulation (M&S) environments, which are interconnected via inter-process communication capabilities. Although here applied to a pyroprocessing based nuclear facility, this M&S tool was constructed in a modular, flexible, and scalable manner to effectively facilitate not only the development of dynamic simulations of nuclear facilities but also the design and assessment of safeguards strategies.

4. Simulation Results

4.1 Generation of Process Events

As batches of nuclear material move between unit operations throughout the pyroprocessing facility following event dynamics that characterize possible “operations” defined for this facility, these batches also trigger the generation of continuous time series following time-driven dynamics that characterize the particular “processes” associated with unit operations. Sensors are thus optimally installed throughout the facility to measure process variables (e.g., volume, mass, dose, current, density) as material is being processed at unit operations or transferred from one location to another. These continuous time series or measurements of process variables are discretized or symbolized into discrete-event values belonging to specified alphabets (e.g., low, medium, and high) defined for the particular systems of interest. The discrete event information is then integrated and used to analyze system conditions and potentially detect anomalies. In order to illustrate how a given time series may be transformed into discrete event values, Figure 6 shows an instance of the time series generated at the Oxide

Reduction (OR) unit operation for the process variable of salt density rate when processing a particular spent nuclear fuel material. This time series also include a portion of measurements collected during salt volume maintenance operations (i.e., batches 50-100). The continuous time series is translated into event series involving the discrete quantification of low, medium, and high. Threshold values, which delimit the different symbolization states, are selected to best discern between normal and abnormal conditions. Similarly, Figure 7 illustrates the time series generated at the Electrorefiner (ER) unit operation for the process variable of U density when processing nuclear fuel.

Figure 6. Sensor data discrete event generation at Oxide Reduction.

Figure 7. Sensor data discrete event generation at Electrorefiner.

4.2 Sensor Reliability

It is assumed that sensors may be unreliable in that the true magnitude of a quantity being measured may not be reported (e.g., due to sensor misdetection) or may be reported incorrectly (e.g., reported as another event due to sensor misclassification). Sensors' reliability is here characterized by the probability of misdetection, correct classification, and misclassification when considering the PM-diagnoser. These characterizations are not directly associated with common statistical test errors such as false positive and false negative, as sensor misclassifications, for example, may lead into both false negative and positive outcomes. Assuming that at a given time instance the true coulombs of electric charge passed through the OR is "normal", these terms instead quantify the likelihood that the monitoring system would receive from a given sensor a "no sensor" reading (e.g., due to sensor failure), a "normal" sensor reading (as it should) or a "low" or "high" sensor reading (e.g., due to sensor noise or malfunction), respectively. In this work, sensor reliabilities are characterized by a parameter, α , illustrated in Table 5. We consider $\alpha = 1$ (i.e., perfect sensors), 0.91, 0.82, and 0.73 in simulations reported here.

Table 5. Sensor Reliabilities

Sensors	Probability of misdetection	Probability of correct detection	Probability of misclassification
All except CS	$(1-\alpha)/3$	α	$2(1-\alpha)/3$
CS	$(1-\alpha)/2$	α	$(1-\alpha)/2$

4.3 Detection of Anomalies via PM-based Data Integration and Interpretation

The number of sensors that a given diagnoser may have for assessing plant conditions generally determines its performance. The more sensors it has, the better its performance often is. In order to illustrate this general observation, four diagnosers were synthesized to detect each identified pattern for a grand total of twelve diagnosers. To elaborate, for detecting anomaly pattern P_1 , four diagnosers observing two, three, five, and six events were designed and are labeled D_{11} , D_{12} , D_{13} , and D_{14} , respectively, in Table 6. Specifically, the columns corresponding to D_{1j} , $j = 1, 2, 3, 4$ in Table 6, indicate the events that the diagnosers observe. Notice that the first three diagnosers can only report partial occurrence of P_1 , while the last one reports the occurrence of the complete full pattern. Moreover, the confidence of diagnosers D_{11} , D_{12} , and D_{13} are low, medium, and high, respectively. For detecting anomaly pattern P_2 , four diagnosers observing one, two, three, and five events associated with this pattern, respectively, were designed. Similar to the diagnosers for P_1 , these diagnosers are listed in Table 7. Note that D_{21} , D_{22} , D_{23} only report partial occurrence of P_2 , while D_{24} reports the occurrence of the complete full pattern. Finally, for detecting anomaly pattern P_3 , four diagnosers observing two, three, four, and five events, respectively, were designed. Similar to the diagnosers for P_1 , these diagnosers are listed in Table 8. Note that D_{31} , D_{32} , D_{33} only report partial occurrence of P_3 , while D_{34} reports the occurrence of the complete full pattern. Figure 4 illustrates the diagnoser structure.

Table 6. Number of pattern occurrences for anomaly P₁

Scenario	Sensor Reliability			
	100%	91%	82%	73%
Complete actual	14	14	14	14
Complete estimate: 6 sensors	14.00	13.98	14.23	12.85
Partial actual	22	22	22	22
Partial actual: 5 sensors	22.00	19.58	20.03	19.33
Partial actual: 3 sensors	24.00	21.72	23.37	21.62
Partial actual: 2 sensors	30.00	28.00	29.78	28.33

Table 7. Number of pattern occurrences for anomaly P₂

Scenario	Sensor Reliability			
	100%	91%	82%	73%
Complete actual	3	3	3	3
Complete estimate: 6 sensors	3.00	3.42	5.67	6.44
Partial actual	24	24	24	24
Partial actual: 5 sensors	24.00	25.40	23.90	25.07
Partial actual: 3 sensors	28.00	27.22	25.22	26.41
Partial actual: 2 sensors	40.00	37.44	35.36	38.73

Table 8. Number of pattern occurrences for anomaly P₃

Scenario	Sensor Reliability			
	100%	91%	82%	73%
Complete actual	2	2	2	2
Complete estimate: 6 sensors	2.00	2.01	2.01	2.03
Partial actual	7	7	7	7
Partial actual: 5 sensors	7.00	6.70	5.38	5.19
Partial actual: 3 sensors	7.00	6.62	5.31	5.29

Partial actual: 2 sensors	9.00	8.46	8.52	8.21
---------------------------	------	------	------	------

Tables 6, 7, and 8 report the actual and estimated number of partial and full occurrences of anomaly pattern P_1 , P_2 , and P_3 , respectively, under different assumptions for the number of sensors used and their reliabilities as defined in Table 5. Notice that when assuming perfect sensors, the corresponding diagnoser correctly reports the number of complete occurrences of given anomaly pattern (e.g., 14 for anomaly pattern P_1). This is also the case when a diagnoser uses the maximum number of sensors when reporting the number of partial occurrences. However, notice that even when assuming perfect sensors, more false alarms are generated when the corresponding diagnosers decreasingly uses lesser sensors. For example, it can be concluded from Table 7, 91% column, partial estimate rows, that lowering number of sensors increases false alarms. This trend is expected and comes from the existence of process upsets that cannot be filtered out by the use of additional sensors. Tables 6, 7, and 8 also demonstrate that the performance of the monitoring system decreases as the reliability of the sensors and the number of used sensors decreases. For example regarding anomaly pattern P_2 , more false alarms regarding the occurrence of the complete pattern are generated as the sensor reliability is decreased. On the other hand, the exact opposite occurs when considering anomaly pattern P_1 (i.e., less complete patterns are inferred). For example, it can be concluded from Table 7, complete estimate: 5 sensors row, that lowering sensor reliability increases false alarms. This trend was expected.

Interesting results can also be observed when monitoring for partial patterns, which depend on how the partial patterns were defined. In all partial cases studied here, the estimate on the number of partial occurrence of a given pattern increases as the number of sensors used is decreased, independently of the assumed value for sensor reliability. Tables 6, 7, and 8 also suggest that the sensitivity of anomaly detection accuracy to sensor quantity and reliability depends on the particular anomaly under consideration. For example, it can be concluded from Table 8, 91%

column, partial estimate rows, that the accuracy for counting anomaly P_3 is somewhat insensitive to sensor quantity. Likewise, it can be concluded from Table 8, complete estimate: 5 sensors row, that the accuracy for counting anomaly P_3 is insensitive to sensor reliability. Thus, the required number and precision of sensors may vary depending upon the decision that needs to be made. This is an important deduction.

5. Conclusions

The application of process monitoring to safeguarding nuclear facilities has not yet gained widespread traction internationally. However, if new material processing systems, such as pyroprocessing, are to be accepted as safeguardable, an alternative to traditional Nuclear Material Accountancy (NMA) will be needed. While pyroprocessing is attractive for how it handles actinides with compact facility requirements and integrated fuel fabrication and waste treatment capability, it is hampered by a lack of input accountability and multiple locations for fissile material hold-up that are not ideally suited for periodic inventory cleanout. NMA will thus be very difficult and likely costly to implement for commercial scale pyroprocessing facilities. Through a time- and discrete event-driven simulation of a simple pyroprocessing system with readily available sensor options, this study has shown that a variety of diversion scenarios can be readily detected provided the sensor reliability is high and a sufficient number of sensors have been installed. Frequency of false alarms can be minimized by optimizing the number and reliability of sensors. It is important to emphasize the fact that simplistic sensors have been modeled for these simulations—such as densitometers, electronic balances, and flow meters. Current development activities funded by the U.S. Department of Energy into more complex and information rich sensors such as voltammetry probes, potentiometric sensors, and laser spectroscopy will undoubtedly improve the effectiveness of anomaly detection and further reduce the frequency of false alarms. But the result here should be used to initiate the discussion regarding abandoning sole reliance on NMA for nuclear facility safeguards. None of the

sensors considered in this study are capable of measuring amounts of specific actinides such as plutonium or uranium with any level of accuracy. But when combined in the reported anomaly detection algorithm, they can be very effective at detecting misuse—which is really the ultimate objective of nuclear safeguards. The next logical step towards developing a robust safeguards system for a pyroprocessing facility is to generate and analyze real operational process data from either a lab scale or engineering scale system.

Acknowledgement

The research reported in this paper was supported in part by the U.S. Department of Energy under contract DE-AC07-05ID14517.

References

1. Simpson, M.F., Rappleye, D., Blandford, E.D., and Garcia, H.E.; “Signature Based Safeguards Alternative to Nuclear Material Accountancy,” *Proceedings of the 2014 INMM Annual Meeting*, Atlanta; July 20-24, 2014.
2. Lafreniere, P.L., Rappleye, D.S., Simpson, M.F., and Blandford, E.D.; “Demonstration of Signature-Based Safeguards for Pyroprocessing as Applied to Electrorefining and the Ingot Casting Process,” *Nuclear Technology*, vol. 189, No. 2, February 2015, pp. 173-185.
3. Garcia, H.E., Simpson, M.F., Lin, W.-C., Yoo, T., and Carlson, R.B., “Detecting proliferation activities via system-centric integration and interpretation of multi-modal data collected from a system of sensors,” in *Proceedings of the 54th Annual Meeting of the Institute of Nuclear Materials Management*, 2013.
4. Garcia, H.E., Lin, W.-C. , DeMuth, S.F., Carlson, R., and Burr, T., “Evaluating safeguards benefits of process monitoring as compared with nuclear material accountancy,” in *Proceedings of the 53rd Annual Annual Meeting of the INMM*, 2012.
5. Garcia, H.E., Lin, W., Johnson, S., Burr, T., DeMuth, S., Bakel, A., Krebs, J., and Ehinger, M., “Integrated process monitoring based on systems of sensors for

enhanced nuclear safe- guards sensitivity and robustness,” in *Proceedings of the 52nd Annual Meeting of the INMM*, 2011.

6. Garcia, H.E., Lin, W.-C., and Yoo, T., “Process monitoring for safeguards via event generation, integration, and interpretation,” in *Proceedings of the 51st Annual Meeting of the Institute of Nuclear Materials Management*, 2010.

7. Sampath, M., Sengupta, R., Lafortune, S., Sinnamohideen, K., and Teneketzis, D., “Diagnosability of discrete event systems,” *IEEE Trans. on Auto. Contr.*, vol. 40, no. 9, pp. 1555–1575, Sept 1995.

8. Lunze, J. and Schrder, J., “State observation and diagnosis of discrete-event systems described by stochastic automata,” *Discrete Event Dyna. Syst.: Theory Appl.*, vol. 11, no. 4, pp. 319–369, 2001.

9. Thorsley D. and Teneketzis, D., “Diagnosability of stochastic discrete-event systems,” *IEEE Trans. Automatic Control*, vol. 50, no. 4, pp. 476– 492, 2005.

10. Yoo, T. and Garcia, H.E., “Diagnosis of behaviors of interest in partially-observed discrete-event systems,” *Systems & Control Letters*, vol. 57, no. 12, pp. 1023–1029, 2008.

11. Thorsley, D., Yoo, T., and Garcia, H.E., “Diagnosability of stochastic discrete event systems under unreliable observations,” in *Proceedings of the American Control Conference*, 2008, pp. 1158 – 1165.

12. Athanasopoulou, E., Lingxi, L., and Hadjicostis, C., “Probabilistic failure diagnosis in finite state machines under unreliable observations,” in *Proc. of 2006 8th International Workshop on Discrete Event Systems*, Ann Arbor, MI, July 2006, pp. 301 – 306.

13. Athanasopoulou, E., Lingxi, L., and Hadjicostis, C., “Maximum likelihood failure diagnosis in finite state machines under unreliable observations,” *IEEE Tran. on Automatic Control*, vol. 55, no. 3, pp. 579–593, 2010.

14. Carvalho, L.K., Basilio, J.C., and Moreira, M.V., “Robust diagnosability of discrete event systems subject to intermittent sensor failures,” in *Proc. of 10th International Workshop on Discrete Event Systems*, Berlin, Germany, Aug. 2010, pp. 94–99.

15. Lima, S.T.S., Basilio, J.C., Lafortune, S., and Moreira, M.V., “Robust diagnosis of discrete-event systems subject to permanent sensor failures,” in *Proc. of 10th*

International Workshop on Discrete Event Systems, Berlin, Germany, Aug. 2010, pp. 100–107.

16. Contant, O., Lafortune, S., and Teneketzis, D., “Diagnosis of intermittent faults,” *Discrete Event Dynamic Systems: Theory and Applications*, vol. 14, no. 2, pp. 171–202, 2004.

17. Jiang, S., Kumar, R., and Garcia, H.E., “Diagnosis of repeated/intermittent failures in discrete-event systems,” *IEEE Trans. Robotics and Applications*, vol. 19, no. 2, pp. 310– 323, 2003.

18. Garcia, H.E. and Yoo, T., “Model-based detection of routing events in discrete flow networks,” *Automatica*, vol. 4, no. 41, pp. 583–594, 2005.

19. Yoo, T. and Garcia, H.E., “Event counting of partially-observed discrete-event systems with uniformly and nonuniformly bounded diagnosis delays,” *Discrete Event Dynamic Systems*, vol. 19, no. 2, pp. 167–187, 2009.

20. Yoo, T. and Garcia, H.E., “Stochastic event counter for discrete-event systems under unreliable observations,” in *Proceedings of the American Control Conference*, 2008, pp. 1145 – 1152.

21. Lin, W.-C., Garcia, H.E., Thorsley, D., and Yoo, T., “Sequential window diagnoser for discrete-event systems under unreliable observations,” in *Proc. of 47th Annual Allerton Conf. on Communication, Control, and Computing*, Monticello, IL, Sept. 2009, pp. 668– 675.

22. Lin, W.-C., Garcia, H.E., and Yoo, T., “A diagnoser algorithm for anomaly detection in DEDS under partial and unreliable observations: Characterization and inclusion in sensor configuration optimization,” *Discrete Event Dynamic Systems: Theory and Applications*, vol. 23, pp. 61–91, 2013.

23. Lin, W.-C., Yoo, T., and Garcia, H.E., “Sensor configuration selection for discrete-event systems under unreliable observations,” in *Proc. of the 6th IEEE Conference on Automation Science and Engineering*, Toronto, Canada, Aug. 2010, pp. 477–484.

24. Lin, W.-C., Garcia, H.E., and Yoo, T., “Selecting observation platforms for optimized anomaly detectability under unreliable partial observations,” in *Proceedings of the 2011 American Control Conference*, pp. 4470–4477.

25. Yoo, Y. and Lafortune, S., "NP-completeness of sensor selection problems arising in partially-observed discrete-event systems," *IEEE Trans. Automat. Contr.*, vol. 47, no. 9, pp. 1495–1499, 2002.
26. Yoo, T. and Garcia, H.E., "Sensor deployment optimization for network intrusion detection," in *Proc. of 44th Annual Allerton Conf. on Communication, Control, and Computing*, Monticello, IL, Sept. 2006, pp. 996–1004.
27. Aguirre-Salas, L., Begovich, O., and Ramirez-Trevino, A., "Optimal sensor choice for observability in free-choice Petrinets," in *Proc. of IEEE Int. Symposium on Intelligent Control*, Mexico City, Mexico, Sept. 2001, pp. 270–275.
28. Aguirre-Salas, L., "Sensor selection for observability in interpreted Petri nets: a generic approach," in *Proc. of 42nd IEEE Conf. on Decision and Control*, Hawaii, Dec. 2003, pp. 3760–3765.
29. Jiang, S., Kumar, R., and Garcia, H.E., "Optimal sensor selection for discrete event systems under partial observation," *IEEE Trans. Automatic Control*, vol. 48, no. 3, pp. 369–381, 2003.
30. Rohloff, K.R., Kuller, S., and Kortsarz, G., "Approximating the minimal sensor selection for supervisory control," *Discrete Event Dyn. Syst.: Theory and Appl.*, vol. 16, no. 2, pp. 143–170, 2006.
31. Lin, W.-C. and Garcia, H.E., "Synthesis and optimization of a bayesian belief network based observation platform for anomaly detection under partial and unreliable observations," in *Proceedings of the 9th IEEE Conference on Automation Science and Engineering*, August 2013, pp. 51–58.
32. Li, F., Upadhyaya, B.R., and Perillo, S.R.P., "Fault Diagnosis of Helical Coil Steam generator Systems of an Integral Pressurized Water Reactor Using Optimal Sensor Selection," *IEEE Transactions on Nuclear Science*, Vol. 59, No. 2, pp. 403–410, April 2012.
33. Herrmann, S.D. and Li, S.X., "Separation and recovery of uranium metal from spent light water reactor fuel via electrolytic reduction and electrorefining," *Nuclear Technology*, vol. 171, no. 3, pp. 247–265, 2010.

34. Laidler, J.J., Battles, J.E., Miller, W.E., Ackerman, J.P., and Carls, E.L., "Development of pyroprocessing technology," *Progress in Nuclear Energy*, vol. 31, no. 1-2, pp. 131–140, 1997.
35. Westphal, B.R., Marsden, K.C., Price, J.C. and Laug, D.V., "On the development of a distillation process for the electrometallurgical treatment of irradiated spent nuclear fuel," *Nuclear Engineering and Technology*, vol. 40, no. 3, pp. 163–174, 2008.
36. Lee, H.-S., Park, G.-I., Kang, K.-H., Hur, J.-M., Kim, J.-G., Ahn, D.-H., Cho, Y.-Z., and Kim, E.-H., "Pyroprocessing technology development at KAERI," *Nuclear Engineering and Technology*, vol. 43, no. 4, pp. 317–328, 2011.