Howell, J., Bevan, G., and Burr, T. (2012) Inspector interfaces to facilitate subjective reasoning about quantities in trends. Computers and Chemical Engineering. ISSN 0098-1354

Inspector Interfaces to Facilitate Subjective Reasoning about Quantities in Trends

John Howell[1*], Geraint Bevan[2], Tom Burr[3]

[1]University of Glasgow, Glasgow, UK
[2]Glasgow Caledonian University, Glasgow, United Kingdom,
[3]Los Alamos National Laboratory, Los Alamos, NM, USA

Corresponding author: John.Howell@Glasgow.ac.uk; UK 141 330 4070

Keywords: nuclear safeguards, inspector interfaces, monitoring, trend analysis, data analysis, diagnosis

Abstract

Nuclear safeguards inspectors view historical trends on a regular basis to reassure themselves that a plant is operating as declared. Other types of inspector are likely to perform similar activities. Nuclear safeguards are founded on materials accountancy, and hence nuclear safeguards inspectors often want to relate what they see to laws of mass conservation. The interfaces discussed in this paper facilitate a synergy between qualitative reasoning regarding trends with quantitative reasoning about simple models that are driven by forcing functions, which describe materials flows through a plant. The focus in the paper is on the assessment of data trends that pertain to a standard 3-tank arrangement in a nuclear fuel reprocessing plant, and on how inspectors might wish to interact with evaluations performed.

1. Introduction

Chemical process plants are often inspected by various bodies including health and safety directorates. An additional inspection regime is imposed on nuclear process plants, because of the need to safeguard their nuclear material under international agreements. Inspectors of all kinds visit plants to gather and digest evidence. In sifting through this evidence they need to decide if certain aspects require further investigation. They are constrained by both time and knowledge limitations. They only visit at certain times and do not have the same depth of understanding about the state of the plant as the operators do. Although an operator might be prepared to divulge all, the inspector is unlikely to have the resources to assimilate all. Inspector overall objectives differ from those of the operators. The International Atomic Energy Agency (IAEA) publishes its current perspective of the role of nuclear safeguards inspectors regularly. Importantly Nackaerts (2011) emphasises that their goal is to ask the simple question: “Is this a safeguards related issue or isn’t it?” Health and safety directorates will ask a similar question related to health and safety. A considerable amount of effort might be required before this question can be answered with a reasonable level of confidence.

The interfaces discussed in this paper focus on assessing data trends in nuclear fuel reprocessing plants. Nuclear safeguards inspectors view such trends on a regular basis to reassure themselves that a plant is operating as declared. Other types of inspector might perform similar tasks, because graphical trends are produced in many applications including industrial, medical and other areas. This paper focuses on a domain whose processes can be modelled on the basis of laws of conservation. The aim is to reason about how laws relate to trends. Such reasoning might also have relevance during incidents such as that at Three-Mile-Island (Walker, 2004) and at Fukushima (Miller et al, 2011).

Quantitative nuclear materials safeguards are founded on accountancy: laws of mass conservation applied to special nuclear material (SNM), such as plutonium or uranium, form the basis of regular material balance evaluation. Statistical approaches for mass balance evaluation require an estimate of standard deviation, commonly denoted σ_{MB}, because of measurement errors. It is difficult to satisfy the International Atomic Energy Agency (IAEA) accountancy-based diversion detection goal for large throughput facilities, because nuclear material accountancy measurement uncertainty increases as facility throughput increases, despite sustained efforts to reduce measurement uncertainty, Detailed knowledge is also required of the material located throughout the plant. Safeguards inspectors often augment this objective approach with so-called additional assurances. Of relevance here is the need to be assured that the plant is operated as declared.

Any safeguards related issue will involve the undeclared and inappropriate transport of special nuclear material. Such an activity must always result in an unexpected reduction in SNM mass somewhere in the plant that might be visible above alarms that are based on σ_{MB}. So such actions might be apparent to the inspector. However, not all the material will be under close observation all of the time and σ_{MB} can be very large in large facilities, particularly if it is cost prohibitive to close the mass balance very frequently. Every day there might be numerous movements of material through pipes and process units, whose inventories are not that transparent. In addition instrument performance might not provide the fidelity to detect small movements. Any computer tools that might support the nuclear safeguards inspector in detecting and isolating these activities will have to relate to actual quantities, because it is the actual quantities that matter.

Currently a number of different computer tools are used to monitor the flow of liquids (solutions) in various large-scale nuclear reprocessing plants (Burr et al, 2008). In the 1990s, Euratom installed a system at La Hague (Dekens et al, 1995), which has become the basis for several later systems, including that used to evaluate data collected at the Rokkasho Reprocessing Plant (RRP) in Japan. Here data is evaluated from 300 sensors on 92 vessels or other process equipment (Ehinger et al, 2004). System outcomes in all these systems are limited simply to the output of red-amber-green kinds of signals: they detect non-compliance with signatures; the inspector is then left to reason about the cause manually, with the help of graphed trends. A portable version of the package also exists (Janssens-Maenhout & L. Dechamp, 2004). These computer tools are commonly referred to as solution monitoring evaluation systems.

Certain plants, however, do not operate in such a prescribed way, to enable the generation of signatures. In these situations an inspector needs to reason about the flows through the plant. In so doing, reasoning is founded on understanding, which here largely revolves around the quantitative conservation of material. Reasoning supports both detection and subsequent evaluation. A similar approach might be applicable to other domains where operators or inspectors might wish to reason in a hybrid quantitative / qualitative manner about e.g. energy flows. This might well be pertinent during the minutes, hours and days following a major accident. The IAEA has developed a **Ta**nk **M**onitoring **E**valuation **S**ystem (TaMES) family of computer tools (Sirajov & Wang, 2008; Howell et al, 2009) to help inspectors reason about flows through certain parts of a nuclear fuel reprocessing plant. Tanks are by far the most common type of plant component in a reprocessing plant and hold most of the nuclear material. Usually instrumentation is available in every tank, so its bulk contents can be estimated and various trends can be evaluated. Reasoning about the transport of material around a plant must be relevant to many plants, nuclear or otherwise.

The TaMES family is founded on a simple representation known as the sub-event (Scothern & Howell, 1997; Howell & Scothern, 2000), which can be viewed as any extended episode in a trend. Rules are applied to these sub-events to relate them to descriptions of known and acceptable activities (known as events). A typical activity might relate to tank sampling where the operator homogenises tank contents before withdrawing a small quantity. Such an activity should generate signatures in trends that can be represented as sequences of sub-events. Any safeguards related issue over a relatively short time frame should generate one or more sub-events in mass trends that cannot be associated with these acceptable activities. With any such unresolved event, a TaMES would first identify and then eliminate acceptable activities to reveal these un-resolved sub-events. Unfortunately measurement anomalies (beyond typical measurement errors) and operational misunderstandings can lead to the creation of other un-resolved sub-events. There is then a need to reason about the various unexplained sub-events. Current operational versions of TaMES simply list unresolved sub-events, leaving it up to the inspector to think about what might have caused the patterns observable in the trends. In these situations, an inspector would be guided by the knowledge that all the material that has entered a particular part of a plant should still be located somewhere, unless it has left the plant already. However an inspector is unlikely to have either the time or inclination to trace the transport of material through the plant. The computer tool presented here performs this function explicitly. Based on measurement data, these calculations are unlikely to be perfect, especially if certain data are anomalous. Anomalous data are handled explicitly relieving the inspector of potential distractions. *A key aim of the research is to explore how these automatic evaluations should be presented, and how an inspector might wish to query the system about its conclusions. A key assumption is that measurement anomalies beyond usual measurement errors occur in relative isolation, meaning that common-mode instrument malfunctions over neighbouring tanks are extremely unlikely.* A set of model forcing functions can be derived on this basis, which can estimate mass flows correctly even when an instrument is faulty. The sub-event can then be marked as non-safeguards related. The authors believe that if there is a common-mode instrument fault in neighbouring tanks, then continued operation would be in question and the event would remain unresolved.

In the context of process supervision, Gentil & Montemain (2004) argue that there is a natural predisposition to qualitative reasoning, especially when "qualitative reasoning reflects a human's aversion to calculations." Gentil, Montmain & Combastel (2004) have proposed a combined fault detection and isolation/artificial intelligence (FDI/AI) approach, in which qualitative reasoning is applied to patterns of residuals that are generated quantitatively from local numerical sub-models. Safeguards inspectors, however, need to relate back to the original quantitative data as much as possible. TaMES has some similarity to the work of Charbonnier et al (2005) who extract semi-qualitative temporal episodes on-line, from any univariate time series. Three primitives are used to describe the episodes: {Increasing, Decreasing, Steady}. Their method uses a segmentation algorithm, a classification of the segments into seven temporal shapes and a temporal aggregation of episodes. Other examples of the application of segmentation algorithms can be found in Miller & Howell (1999) and Janssens-Maenhout & Dechamp, (2004); the context for the latter can be found in Dekens et al (1995). An important difference from the Charbonnier et al approach is that in our approach a sub-event is represented by the state of the tank (mass, volume, density, temperature) at the start and end of the specified time period, and not just by the single variable that the trend specifically pertains to.

This paper describes improved reasoning processes proposed to help the inspector decide if further investigation is required. Based on considerable breadth and depth of knowledge, inspector reasoning processes are difficult to capture in a single paper. Effort is made in Section 2 to encapsulate some of the aims and pitfalls with clarity

2. Background

Worldwide many facilities process nuclear material in various stages of a civil nuclear fuel cycle, including fuel fabrication and reprocessing. These facilities are usually located in countries whose governments are signatories of the Treaty on the Non-Proliferation of Nuclear Weapons (IAEA Information Circular 140). Agreements (IAEA Information Circulars 153 and 540) often then provide for the inspection of these facilities to confirm that their nuclear material is used as declared. Nuclear process plants are often large heterogeneous facilities that provide many potential opportunities for misuse, some more practicable than others. Although direct, indisputable observation of misuse is desirable, it is possible that misuse would be detected instead through the collection and analysis of evidence, particularly accountancy data. Any conclusions drawn from the analysis of evidence are then likely to invoke further investigation.

It is clear that any inference system that is poorly designed and implemented will create considerable work for all involved. Indeed, detection based on misinterpretations (so-called false alarms) is of major concern to inspectors, operators and governments (Friend, 2008). Importantly, when designed and implemented correctly, successful, robust handling of evidence can contribute to misuse deterrence. Traditional nuclear safeguards agreements focus on the verification of the correctness of declared nuclear material inventories (Demuth et al., 2010) and on mass balance evaluation. Nuclear material accounts are formed and evaluated typically every month.

Accountancy can be augmented by the assessment of data collected more regularly from the plant, for instance every minute (Ehinger et. al., 2004). Such assessments can provide greater fidelity, because they suffer much less from the accumulation of measurement errors that can arise when accounting material movements over weeks of operation. However these assessments are sensitive to day-to-day operational activities and hence might generate false alarms. In many ways there are similarities between fault detection and diagnosis (Patton, Frank and Clark, 2000) and inspection, in that both seek to detect and diagnose anomalies. A key difference with inspection though, is the reduced access to knowledge.

It is important to appreciate that the inventory of plutonium within a plant cannot be measured directly in-line with existing measurement technology. Procedures that provide accurate estimates are both lengthy and expensive to perform and form the basis for relatively infrequent mass balance calculation and evaluation. In between such assessments the inspector looks for unexplainable deviations from acceptable practice, in those variables that can be measured. These variables are typically mass, volume, density and temperature. To do this the inspector has to have a good idea about what is acceptable practice. Of course, in normal operation most of these variables change with time, as tanks are filled and emptied and so on. The inspector has to look for deviations against this changing background. Fortunately any unexplainable deviations in these variables are either likely to have a consequential effect downstream, or have been caused by an incident upstream, or both. Parts of the plant will be opaque to the inspector creating difficulties when trying to make assessments of causes and effects.

All current solution monitoring evaluation systems provide basic auto-correlation and cross correlation facilities that provide a capability to detect disagreements or issues. Current systems base their detection and correlation procedures on what is known as event marking. Event marking is a form of data compression in which trend displays (mass or volume) are marked to highlight key features like the start and stop points of a filling operation. Burr et. al. (2011, 2012) describe the process of 'event marking' and learning to recognize events of interest, and to archive events such as sparging, recirculation, sampling so that we can learn typical behaviour. In the Euratom family of computer tools data compression, auto correlation and cross-correlation functions are configured, implicitly, when a cycle (a form of signature) is constructed from icons [Janssens-Maenhout & Dechamp, 2004]. The simplest cycle pertains to a batch-operated tank: here a cycle of icons would specify those significant ups and downs that would be deemed to be normal operation; the evaluation system would then step through the icons as each up and down was observed in the trend; start and end points would be marked by the observing icon; a disagreement would be raised if any change was unexpected. This is known as auto-correlation. Other tanks are either fed continuously, and emptied quickly, or fed quickly and emptied slowly. Here the allowable range of rate of continuous fill or empty is specified in an icon and a disagreement is raised if the rate deviates out of the specified range. Points when the quick empty or quick fill starts and ends are marked. Cross-correlation is then performed by comparing the mass/volume change between these marks with that change observed in the corresponding buffer tank. Again a disagreement is raised if they do not agree to within specified tolerances.

Currently TaMES (Sirajov & Wang, 2008) helps IAEA inspectors evaluate measurements collected by data acquisition systems installed in batch operated tanks under safeguards. These systems provide assurance that plutonium nitrate carrying solutions travel through a batch-operated set of tanks as expected. The notion of acceptable practice should be transparent with such tank sets, because the bulk mass should remain constant during the period when a tank is filled. Typically, a TaMES will process in the order of 100,000 data records per tank per month to produce time histories pertaining to mass, volume, density, and temperature. Pertinent points are marked in one of these histories (usually mass) to generate *compressed* data. The change from one marked point to the next is known as a sub-event. Events are generated by applying sets of rules to the sequences of sub-events. Typical events include imports, exports, inter-tank transfers, recirculation and acid addition. The construction of the rule-base is specific to a particular facility and individual rules can encapsulate plant operation in detail (Howell et al., 2009).

In the ideal world, all of this would be straightforward. In the real world, the quality of the data might be far from ideal. Difficulties can arise with the instrumentation. Examples of measurement anomalies include dip-tubes becoming blocked or the air supply to the dip-tubes becoming disconnected and so on. The number of possible conclusions is increased, giving the inspector more options to consider. The inspector interface becomes an even more important part of the overall system.

Figure 1 shows a version of an inspector panel for a plutonium nitrate storage area TaMES application. An example of such a facility can be found in Suzuki et al, (2009). Typically the inspector might view about a month of data in one session. The inspector performs these evaluations in three stages by clicking on 'Load Raw Data', then 'Pre-process' and finally 'Analyse'. At each stage results are viewed in windows that fly-out when various options are selected. For instance an editor is provided to manually adjust markings, because accurate data marking can be difficult in certain situations. Figure 2 gives an example of what might be seen when 'Plot markings and events' is selected. There are three distinct areas in this window: the two graphs, a table of start/stop times that pertain to events categorised as 'Out to ST2' and a table of values that relate to the small red square that has been selected by clicking on the yellow mark, which preceded it. The bottom graph shows a time segment of the upper graph with red, circular markings, instead of yellow squares, to highlight the event shown in blue in the 'Out to ST2' part of the window. This was achieved by clicking on that particular row. The pull-down window, which currently displays 'Out to ST2', lists the event categories that might be found. Of particular interest are the two categories, 'Unresolved' and 'User Defined'. The list of 'Unresolved' events are likely to be of most interest to the inspector as these are the ones that could not be explained automatically. It is intended that the inspector would satisfy himself that each event had a reasonable explanation. Having done this, he would remove that event from the list by declaring it as a 'User Defined' event. To do this, he would click on the 'User-define events' button (Figure 1), which would lead to an interactive display where the event could be named, then transferred from the 'Unresolved' list to the 'User Defined' list.

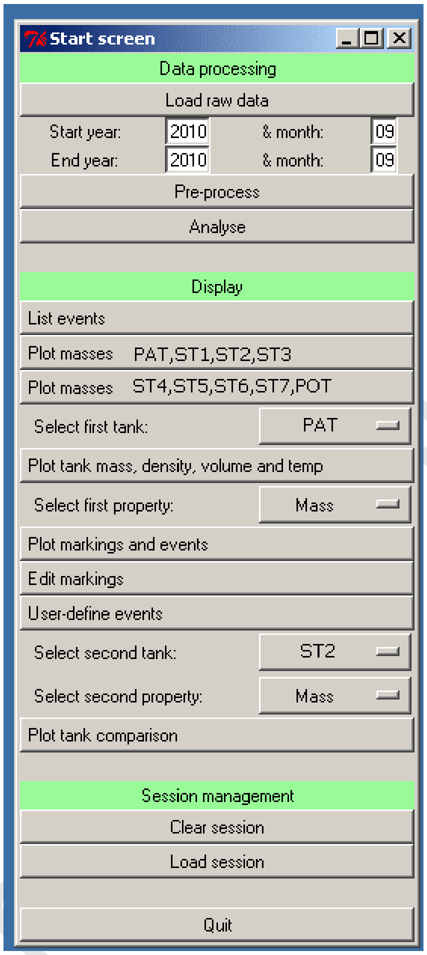

Figure 1 A TaMES Inspector Panel

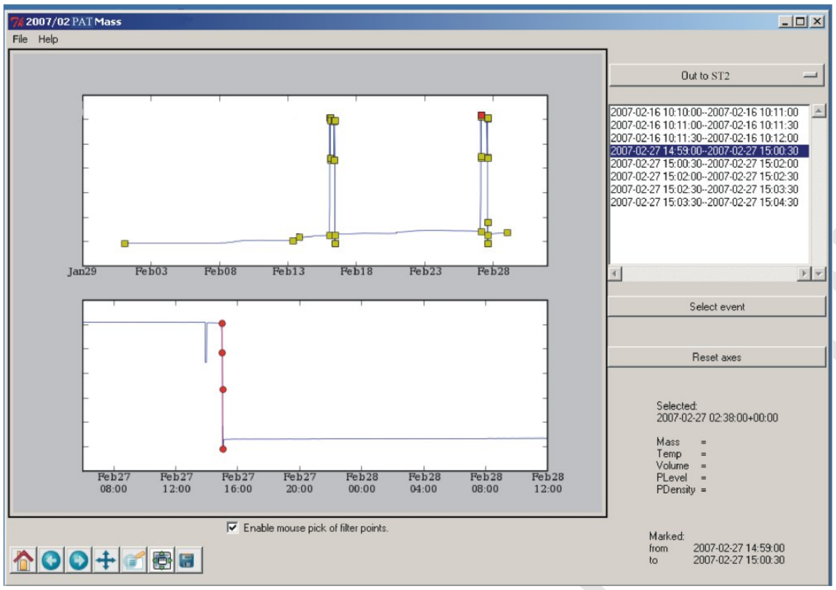

Figure 2 'Plot markings and events'

Of course, reprocessing plants contain a great deal more than sets of batch-operated tanks. As far as the plutonium path is concerned, a reprocessing plant can be viewed in three parts, once spent fuel has been chopped up, dissolved in nitric acid and clarified. The first part passes clarified solution along a set of batch operated tanks until it reaches a feed tank. The second part starts from this feed tank where solution is imported in batches, but exported continuously to a separation process. The first stage of the separation process exports to a standard 3-tank-set, which exports to a polishing plant, which exports to a second standard 3-tank-set, which exports to an evaporator that exports to a receipt tank. The third part starts from this receipt tank where solution is imported continuously, but exported as batches down a batch-operated tank set. Each standard 3-tank set consists of a continuous-import/batch-export receipt tank, a batch-import/batch-export buffer tank and a batch-import /continuous-export feed tank.

The inspector's ability to decide what is acceptable practice now depends very much on the inspector's knowledge of the plutonium flow through the plant. Acceptable practice is that which conserves plutonium mass. Unfortunately, and has been said before, plutonium flows can only be inferred, since they cannot be measured directly with today's technologies and budgets. In addition, measurements can be anomalous, although there is usually more than one way to infer plutonium flows, because of analytical redundancy in the data. This might lead the inspector to a three pass evaluation of trend data, before making a commitment to carry out further investigations at the plant.

Pass 1: reason about TaMES graphs and unresolved sub-events in a semi-qualitative fashion, largely relying on experience-based template matching.

Pass 2: assess whether plutonium flows are conserved, perhaps by eliminating certain suspected-to-be-anomalous measurement sources.

Pass 3: determine that amount of solution that must be removed to explain the plutonium flows observed.

Three examples are discussed here to elaborate on this procedure. To assist with this discussion, Figure 3 shows the key unexplained deviations in their mass trends: figures (a) & (b) pertain to a buffer tank in a standard 3-tank arrangement, whilst (c) relates to a receipt tank in the same arrangement

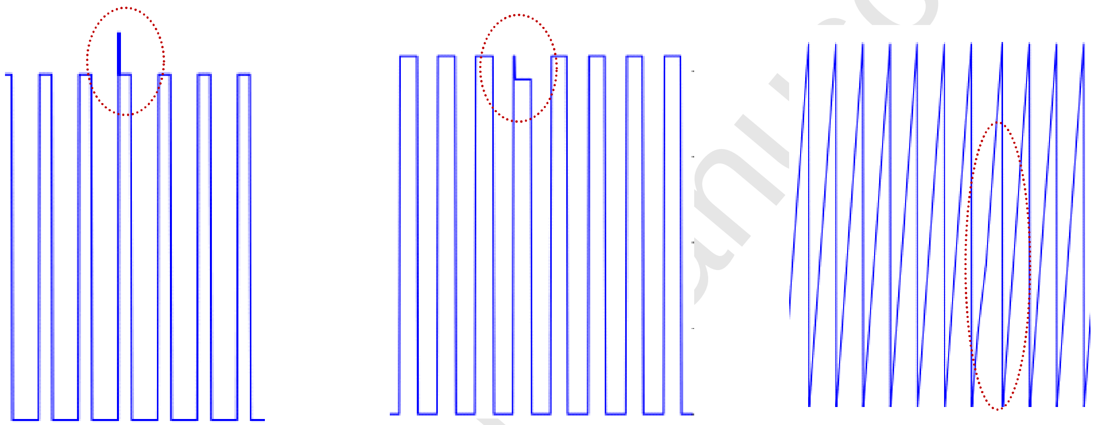

(a) (b) (c)

Figure 3 Mass versus time trends for three separate circled incidents

Example (a) A measurement anomaly in the buffer tank; perhaps this was caused by the introduction of air to inject solution into the tank. Pass 1: the inspector would be familiar with this signature. Pass 2: the inspector might still wish to confirm that the solutions exported from the tank upstream, and imported downstream were appropriate. In other words he would ignore buffer measurements that were collected during that brief period and test whether flows now balanced.

Example (b) A removal of approximately 0.49 m^3 of solution from the buffer tank. Pass 1: the inspector would likely think that the sudden step down looked odd, and consider a number of possibilities. It is possible that an instrument technician had knocked the calibration off; the desk operator had seen this and asked the technician to re-adjust, which had then been done after the tank had been emptied. A small amount of solution might have been exported either downstream, or to its recycle tank. There might have been an undeclared removal. Pass 2: flows along 2 paths require confirmation; flow calculations would either be based on these buffer measurements, or would be based on flows based on neighbouring tanks. Pass 3: would be required if none of these alternatives correlated with the data. Flow calculations would now be repeated, but with a quantity of solution removed at the time of the sudden step down. This scenario would now be correlated with the data.

Example (c): A removal of approximately 2.3 m^3 from the receiving tank: from either the receiving tank directly, or from the pipe leading into it. Pass 1: the inspector might think that the deviation in rate of fill was abnormal, and consider a number of possibilities. It is possible that the throughput of the process upstream had been reduced, temporarily. A quantity might have been exported to the buffer tank. There might have been an undeclared removal. Pass 2: flows into the upstream process unit would be examined to establish if a temporary change had been made; flows to the buffer tank would be correlated with the deviation in fill rate. Pass 3: would be required if none of these alternatives correlated with the data. Flow calculations would now be repeated, but with a quantity of solution removed at the time of the deviation. This scenario would now be correlated with the data.

It can be seen that the procedures are founded on assessments of how materials propagate through a plant. Materials propagation calculation can be viewed as a form of simulation. It can also be seen that the inspector would benefit from a computer tool that enables him to perform these assessments. The authors' intention is that this would be in the form of an enhancement to current TaMES developments.

These three examples are about the measurement or removal of bulk material, and are therefore relatively easy to see in trends. There are many examples that are more about the individual components within flows, and are more opaque in trends. Similar procedures apply, but are more subtle. For ease of explanation what follows refers to reasoning about bulk materials.

3. Overall Goals

The overall aim is to help the inspector decide whether any unresolved event presented by the evaluation system requires a major investigation. In each case the inspector has to decide between the 3 alternatives:

- Measurement anomaly
- Transfer related incident that is acceptable from a safeguards perspective
- Safeguards related incident

The last two suggestions involve the unexpected movement of material, while the first doesn't. In every case the inspector must be able to justify the choice made. Some cases might require some form of verification, even if they are thought to be of little safeguards significance.

It is proposed to provide the inspector with three probabilities, $P\left[M\middle|\text{evidence}\right]$, $P\left[T\middle|\text{evidence}\right]$ and $P\left[S\middle|\text{evidence}\right]$ for each unresolved event, to help with this decision, where event *M* denotes a measurement anomaly, event *T* denotes an acceptable transfer and *S* denotes a safeguards related incident. The probabilities are subjective and cannot be interpreted in any absolute sense; however, they do provide a basis for comparing the relative likelihoods of alternatives M, T, and S. The primary system output is then a simple display such as that shown in Figure 4. A number of sliders, buttons and displays are also provided to enable the inspector to explore both the sensitivity of these three probabilities and how they were come by. All of the sliders are preset at values obtained from the calculations. Their functions will be discussed later.

Once the inspector has selected one or more sub-events that are likely to represent a single event, the system establishes probabilities that the TaMES outputs (θ_T : graphs and un-resolved sub-events) implicate either a safeguards related event or a non-safeguards related event i.e. the system would seek probabilities $P\left[S\middle|\theta_T\right]$ and $P\left[\overline{S}\middle|\theta_T\right]$. At this stage inspector confidence in reaching a definitive decision is likely to be low. This confidence could be improved if the inspector were to have a tool that can assess some form of archived database. If such a tool existed, then the system could start its evaluation with default a priori probabilities $P(S)$ and $P(\overline{S})$, whose values would remain hidden from the inspector. The system would also have defaults for the probabilities that the TaMES outputs (θ_T) were caused by either a safeguards related event $P\left[\theta_T\middle|S\right]$ or a non-safeguards related event $P\left[\theta_T\middle|\overline{S}\right]$. Sliders could be provided to enable the inspector to alter these suggested values.

Then $P\left[S\middle|\theta_T\right]=\frac{1}{\lambda_T}P\left[\theta_T\middle|S\right]P\left[S\right]$ and $P\left[\overline{S}\middle|\theta_T\right]=\frac{1}{\lambda_T}P\left[\theta_T\middle|\overline{S}\right]P\left[\overline{S}\right]$ (1)

where $\lambda_T \equiv P\left[\theta_T\middle|S\right]P\left[S\right]+P\left[\theta_T\middle|\overline{S}\right]P\left[\overline{S}\right]$

However such a tool to estimate the required likelihoods such as $P\left[S\middle|\theta_T\right]$ from an archived database does not exist currently. The approach proposed for now is to ask the inspector to consider the odds in favour of $P\left[\overline{S}\middle|\theta_T\right]$ and against $P\left[S\middle|\theta_T\right]$. The odds in a favour of a non-safeguards related explanation would be displayed and a slider would be provided so that the odds could be changed, within sensible bounds.

The inspector would now move on to the first of 2 windows, where the inspector would explore the possibility that these sub-events related to a measurement anomaly. This would be the first time that the inspector was presented with the 3 probabilities displayed in Figure 4, which would be the main outcome of this exploration. The inspector would move on to the second window, if the safeguards related outcome were still thought to be possible. The actual contents of these 2 windows will be discussed in Section 5. Section 4 first describes what has to be performed to generate their content.

Measurement related	Transfer related	Safeguards related
98.666 %	0.319 %	0.015 %

Figure 4

4. Reasoning about Quantities

Event Generation

A language has evolved in safeguards to describe what is seen in trends: process activities that an inspector or operator can relate to are described as *events*. Each event can be distinguished by a number of deviations in the trends. Each deviation is known as a *sub-event*, so that an event is composed by one or more sub-events. Those deviations (sub-events) that cannot be attributed to known events are then deemed to be unresolved. Sub-event extraction is the process of identifying these features in the data. In its simplest and most explicit form it can derive from ‘vertex marking’. The science, or more appropriately, the skill, of vertex marking is relatively general. The reference to ‘skill’ alludes to the observation that there are many choices that have to be made when specifying procedures for vertex marking. These choices are influenced by the form of the data themselves. Vertices are often masked by process disturbances, perhaps associated with transient effects that might be attributed to motion or to thermal stratification. Other effects might also be observed in the data, like the regular automatic agitation of the liquid (i.e. sparging). Marking the most appropriate points can often be a challenge (Burr et. al, 2011). Further details on this representation can be found in Howell & Scothern, 2000 and in Howell, Ehinger and Burr, 2007. Further details on trend marking can be found in Miller and Howell (1999) and in Combastel (2004).

Once marked, a trend can be outlined by a sequence of markings. A single sub-event denotes the change from one mark to the next. If this trend is of bulk mass and the mass at mark i is $\tilde{M}_i$ then the change in mass b_i that occurred during the i^{th} sub-event SE_i will be estimated as:

$$\hat{b}_i = \tilde{M}_i - \tilde{M}_{i-1} \qquad (2)$$

Tildes ~ denote measurements (either direct or indirect), while hats ∧ denote predictions. A mark relates to a particular data point, there is no data filtering, which is important from an accountancy perspective, because filtering introduces ambiguities.

If $\tilde{M}_0$ is the mass measured at the start of the sequence

$$\tilde{M}_i = \tilde{M}_0 + \sum_{j=1}^{i} \hat{b}_j : \tilde{M}_i \geq 0 \qquad (3)$$

Events are identified by applying lists of rules, R_k : $R_k = [r_{k,1}, r_{k,2}, \ldots]$, to the tables of sub-events (Howell et al., 2009). A separate table of events, $E_{t_i} = \{e_{t_i,1}, e_{t_i,2}, \ldots\}$, is generated for each tank, $t_i \in Tanks$. A separate sub-event table, $SE_{t_i} \in SE$ is associated with each tank so that the first row in SE_{t_i}, $se_{t_i,1}$, has 2 elements $[p_1, p_2]$, the second row, $se_{t_i,2}$, has 2 elements $[p_2, p_3]$ and so on, where p_1 and p_2 are the first 2 marked points for that particular tank and so on. Each marked point has a set of properties associated with it, like those shown in the bottom right-hand table of Figure 2. A bulk mass change $\hat{b}$ can be calculated for each sub-event.

For a specified tank, t_i, each rule is written so that it operates on the first, and where necessary subsequent elements, in SE_{t_i}. When needed sub-events associated with other tanks will also be referenced. This might arise, for instance, when the rule focuses on an inter-tank transfer. The application of any rule might lead to: a) the generation of an event, $e_{t_i,k}$; b) the removal of associated sub-events, both locally, i.e. from SE_{t_i}, and more widely. Formally the l^{th} rule in the k^{th} list, $r_{k,l}$, might be described as:

$$r_{k,l}\left(t_i, SE\big|_j, E_{t_i} = \left\{e_{t_i,1}, \ldots, e_{t_i,k}\right\}\right) \rightarrow E_{t_i} \cup e_{t_i,k+1}, SE\big|_{j+1} \subset SE\big|_j$$

The overall process is then as follows:

$\forall\ t_i \in Tanks$,

Select the most appropriate rule-set, k:

$\forall\ se_{t_i,j} \in SE_{t_i}$,

$$r_{k,p}\left(\ldots r_{k,3}\left(r_{k,2}\left(r_{k,1}\left(t_i, SE\big|_1, \{\}\right)\right)\right)\right) \rightarrow E_{t_i}, SE\big|_p$$

The construction of the rule-base is specific to a particular facility, with individual rules encapsulating plant operation in detail. Its construction can become fairly involved when data lack consistency. The ordering of the rules within each rule base R_k can matter, because the order in which sub-events are eliminated can affect the outcomes. This explains why R_k is denoted by a list [ ] as opposed to a set { } in the above. Examples of rules can be found in Howell et al. (2009). Of particular importance to the discussion here, are those rules that identify a batch transfer between two tanks: these rules cross-correlate that shipped with that received, or more formally, they find a set of sub-events (an event) in both tanks that correlate both in magnitude and in time. To elaborate on this, suppose that an export event E_X has been found in one tank that is immediately upstream of another tank that has a corresponding import event E_I. For simplicity suppose that both events are formed from single sub-events with associated mass changes $\hat{b}_X$ (X for export) and $\hat{b}_I$ (I for import). Then the rule would form and test the residual | $\hat{b}_X$ - $\hat{b}_I$ | < T for some alarm threshold T. Evidence θ_T would then only contain the sub-events associated with E_X and E_I if this test failed.

Having applied all the rules, what remains in $SE\big|_p$ are then the 'unresolved' sub-events.

These sub-events are likely to appear in separate temporal groups so that the inspector can easily focus on both the tank set of interest, and on the analysis window that defines the time period of interest. At this stage the most likely cause of any remaining unresolved sub-events is one of a measurement anomaly, either directly or indirectly where material is transferred into a recycle tank that is so rarely used that its instrumentation has not been maintained.

Material Propagation

Material propagation calculations help distinguish measurement anomalies from transfers, because any mass transport must obey the Laws of Conservation, which dictate that all the material that has entered a particular part of a plant must still be located somewhere, unless it has left the plant already. There are four parts to the calculations: flow estimation, mass prediction, residuals generation and testing. Flow estimation will be erroneous if a measurement anomaly is taken as a transfer instead. Residuals will be erroneous if material simply disappears. The time period over which the analysis is performed must be sufficiently long to enable the observation of material propagation effects.

Alternative flow estimates are examined to isolate measurement anomalies. The simple example with an import and export above is considered again to explain this. To conserve mass either the transfer estimate can be based on mass change $\hat{b}_X$ giving contents predictions

$$\text{tank j-1:}\ \hat{M}_i = \hat{M}_{i-1} - \hat{b}_X \quad ; \quad \text{tank j:}\ \hat{M}_i = \hat{M}_{i-1} + \hat{b}_X \qquad (4)$$

or on mass change $\hat{b}_I$ giving contents predictions

$$\text{tank j-1:}\ \hat{M}_i = \hat{M}_{i-1} - \hat{b}_I \quad ; \quad \text{tank j:}\ \hat{M}_i = \hat{M}_{i-1} + \hat{b}_I \qquad (5)$$

or the average $\hat{b}_I$ and $\hat{b}_X$ can be taken instead. Residuals as defined by

$$r_i \triangleq \hat{M}_i - \tilde{M}_i \qquad (6)$$

can be generated for all of these cases. Suppose that $\hat{b}_X = b_X = 0$, but $\hat{b}_I$ is non-zero, then residuals based on equation 4 will be zero if measurement noise is zero, while the sequence of residuals based on equation 5 will not be zero, from *i* onwards in the upstream tank and from (*i+1*) onwards in the downstream tank. These differences in residual histories provide additional evidence (θ_C) with which to reason.

The approach adopted here is to reason about these alternative calculations (propagations) in parallel. Thus there are propagations *forwards* that are based on b_Xs and *backwards* that are based on b_Is. If unresolved sub-events derive from mass trends, these sub-events must either relate to changes in mass or to measurement anomalies. If the performance of a particular measurement source is under suspicion, then materials propagation can be based on exports from the tank upstream together with imports to the tank downstream:

$$\text{tank j-1: } \hat{M}_i = \hat{M}_{i-1} - \hat{b}_X \quad ; \quad \text{tank j: } \hat{M}_i = \hat{M}_{i-1} + \hat{b}_X - \hat{b}_I \quad ; \quad \text{tank j+1: } \hat{M}_i = \hat{M}_{i-1} + \hat{b}_I \qquad (7)$$

This is akin to constraint suspension (Davis, 1984) if tank j instrumentation is ‘suspended’. In particular the following alternative propagations have been found to be useful for the standard 3-tank set described before: *forwards*, *buffer dominates* and *buffer suspect*. These alternatives are termed *rationales*. In *forwards*, each forcing function is formed from sub-events associated with exports; in *buffer dominates*, where possible each forcing function is formed from sub-events associated with the buffer tank; in *buffer suspect*, where possible forcing function formation avoids sub-events associated with the buffer tank.

The consideration of material propagation across process stages would be slightly different, because unmetered flow rates mean that only key quantities such as the masses of uranium and plutonium are conserved. This extension is not described here.

Time histories can be predicted on the basis of these material propagation calculations. These histories can then be marked and residuals can be generated by comparing these markings with those obtained from the plant data. There is no guarantee that the predicted histories will be marked in exactly the same way as the real data were marked, so care must be taken to ensure that the most appropriate markings are compared.

Probability Determination

Mass balances are structurally correct. If all transfers, b_i, into and out of the balance area are known, without measurement error then

$$M_i = M_{i-1} + b_i : M_i \geq 0 \qquad (8)$$

If M_0 is the true mass at the start of the analysis window:

$$M_i = M_0 + \sum_{j=1}^{i} b_j : M_i \geq 0 \qquad (9)$$

Equations (8) and (9) can be viewed as alternative forms of a structurally correct model that is driven by a forcing function defined by the sequence $\left[b_i\right]_{1\le i\le N}$. An estimate for each b_i can be obtained from the i^{th} sub-event in the analysis window (equation 2).

Uncertainty in the mass trend therefore derives from the predicted transfers $\left[\hat{b}_i\right]_{1\le i\le N}$, whose values will depend on the rationale adopted. A sequence of errors [γ_1, γ_2, γ_3, ...] can therefore be defined for each rationale:

$$\hat{b}_i = b_i + \gamma_i \quad (10)$$

so that

$$\hat{M}_i = \hat{M}_{i-1} + \hat{b}_i : \hat{M}_i \ge 0$$

$$\Rightarrow \hat{M}_i = \hat{M}_0 + \sum_{j=1}^{i} \hat{b}_j : \hat{M}_i \ge 0$$

$$\Rightarrow \hat{M}_i = \hat{M}_0 + \sum_{j=1}^{i} b_j + \sum_{j=1}^{i} \gamma_j : \hat{M}_i \ge 0 \qquad (11)$$

The initial condition $\hat{M}_0$ is estimated as $\hat{M}_0 = \tilde{M}_0$ and in general mass measurements will have errors, $\varepsilon_i : \tilde{M}_i = M_i + \varepsilon_i$

The i^{th} residual, r_i, is then

$$\begin{aligned} r_i &\triangleq \hat{M}_i - \tilde{M}_i \\ &= \tilde{M}_0 + \sum_{j=1}^{i} b_j + \sum_{j=1}^{i} \gamma_j - M_i - \varepsilon_i \\ &= \sum_{j=1}^{i} \gamma_j + \varepsilon_0 - \varepsilon_i \qquad (12) \end{aligned}$$

If there are no measurement anomalies, and errors are additive, then we can assume that $\gamma_i = N\left(0, \sigma_\gamma^2\right)$ and $\varepsilon_i = N\left(0, \sigma_\varepsilon^2\right)$. In practice the true value of γ_i would be smaller than this on those occasions when the condition $\hat{M}_i \ge 0$ has to be imposed.

A transformation **T** can be introduced to relate the vector of residuals (vector **r**) to the vector of error sources (vector **e** where **e** = [γ_1, γ_2 , γ_3 …, ε_1, …]):

$$\mathbf{r} = \mathbf{Te}, \mathbf{r} \in \mathbb{R}^N, \mathbf{e} \in \mathbb{R}^M \qquad (13)$$

The covariance matrix for **r** is then $\mathbf{T\Sigma T'}$ where Σ is the error vector covariance matrix.

Every element of the residuals vector $\mathbf{r}_r$ obtained for rationale r can be normalised thus:

$$\left[\mathbf{T}\Sigma\mathbf{T}'\right]^{-\frac{1}{2}}\mathbf{r}_r \qquad (14)$$

The sum of the squared, normalised residuals is then

$$\mathbf{r}_r'\left(\left[\mathbf{T}\Sigma\mathbf{T}'\right]^{-\frac{1}{2}}\right)'\left[\mathbf{T}\Sigma\mathbf{T}'\right]^{-\frac{1}{2}}\mathbf{r}_r = \mathbf{r}_r'\left[\mathbf{T}\Sigma\mathbf{T}'\right]^{-1}\mathbf{r}_r \qquad (15)$$

Assuming **r** is approximately Gaussian distributed, this sum is approximately Chi-squared distributed with N degrees of freedom, so that the probability that this residuals vector is associated with an anomaly free model can be obtained from the associated cumulative probability distribution.

A simple, configuration specific, rule-base now compares the probabilities obtained for each of the rationales to decide which rationale produces the smallest normalised residuals and on that basis to assign probabilities: $P\left[\theta_C\middle|\left(M\middle|\theta_T\right)\right]$, $P\left[\theta_C\middle|\left(S\middle|\theta_T\right)\right]$ and $P\left[\theta_C\middle|\left(T\middle|\theta_T\right)\right]$ where θ_C denotes the evidence based on the Laws of Conservation as described above. These values are preset on sliders so the inspector can explore other values.

The way that residuals are aggregated depends on both the focus and on the rationale selected. Normally a sequence of residuals would be generated for the focus. Each residual would be squared and summed to obtain a total squared error. However this would not be the case if the measurements for the focal tank were hypothesised as suspect. Residuals would then be generated for both receiving and feeding tanks: the total squared error would then be the combined values.

Probabilities $P\left[S\middle|\theta_T\right]$ and $P\left[\bar{S}\middle|\theta_T\right]$ are then updated as follows:

$$\beta \equiv \frac{P\left[\theta_C\middle|\left(M\middle|\theta_T\right)\right]}{\left(P\left[\theta_C\middle|\left(M\middle|\theta_T\right)\right]+P\left[\theta_C\middle|\left(T\middle|\theta_T\right)\right]\right)} \qquad (16)$$

The probability of a measurement anomaly/acceptable transfer is:

$$P\left[M\middle|\theta_T\right]=\beta P\left[\bar{S}\middle|\theta_T\right] \qquad (17)$$

$$P\left[T\middle|\theta_T\right]=\left(1-\beta\right)P\left[\bar{S}\middle|\theta_T\right] \qquad (18)$$

An additional hypothesis F is introduced to represent the possibility that the event occurred elsewhere:

$$P\left[\left(S\middle|\theta_T\right)\middle|\theta_C\right]+P\left[\left(T\middle|\theta_T\right)\middle|\theta_C\right]+P\left[\left(M\middle|\theta_T\right)\middle|\theta_C\right]=1-P\left[\left(F\middle|\theta_T\right)\middle|\theta_C\right]\equiv\gamma_C \qquad (19)$$

where “Locality” γ_C is set to 0.99, but can be altered by adjusting a slider.

Then

$$P\left[S\middle|\theta_C \wedge \theta_T\right] = P\left[\left(S\middle|\theta_T\right)\middle|\theta_C\right] = \frac{\gamma_C}{\lambda_C} P\left[\theta_C\middle|\left(S\middle|\theta_T\right)\right] P\left[\left(S\middle|\theta_T\right)\right] \tag{20}$$

and so on where

$$\lambda_C \equiv P\left[\theta_C\middle|\left(S\middle|\theta_T\right)\right] P\left[S\middle|\theta_T\right] + P\left[\theta_C\middle|\left(T\middle|\theta_T\right)\right] P\left[T\middle|\theta_T\right] + P\left[\theta_C\middle|\left(M\middle|\theta_T\right)\right] P\left[M\middle|\theta_T\right]$$

Generating Additional Material Movements

The focus here is on the possibility that the unresolved sub-event is a direct consequence of an undeclared material movement and hence has safeguards relevance. That is a mass δ was moved, unexpectedly, during the period of time marked by the unresolved sub-events, and from that particular tank. Mass δ is inferred, the forcing function is adjusted to accommodate this additional flow from the tank, and material is propagated to generate residuals, which are then minimised by adjusting δ. The procedure above is now repeated for each rationale, resulting in the generation of the following conditional probabilities, where θ_R denotes the additional information obtained by removing material:

$P\left[\theta_R\middle|\left(S\middle|\left(\theta_C \wedge \theta_T\right)\right)\right]$, $P\left[\theta_R\middle|\left(T\middle|\left(\theta_C \wedge \theta_T\right)\right)\right]$ and $P\left[\theta_R\middle|\left(M\middle|\left(\theta_C \wedge \theta_T\right)\right)\right]$.

Defining "Locality" γ_R here as

$$\gamma_R \equiv P\left[\left(S\middle|\theta_C \wedge \theta_T\right)\middle|\theta_R\right] + P\left[\left(T\middle|\theta_C \wedge \theta_T\right)\middle|\theta_R\right] + P\left[\left(M\middle|\theta_C \wedge \theta_T\right)\middle|\theta_R\right] = 1 - P\left[\left(F\middle|\theta_C \wedge \theta_T\right)\middle|\theta_R\right] \tag{21}$$

Then

$$P\left[S\middle|\theta_R \wedge \theta_C \wedge \theta_T\right] = P\left[\left(S\middle|\theta_C \wedge \theta_T\right)\middle|\theta_R\right] = \frac{\gamma_R}{\lambda_R} P\left[\theta_R\middle|\left(S\middle|\theta_C \wedge \theta_T\right)\right] P\left[\left(S\middle|\theta_C \wedge \theta_T\right)\right] \tag{22}$$

where

$$\lambda_R \equiv P\left[\theta_R\middle|\left(S\middle|\theta_C \wedge \theta_T\right)\right] P\left[\left(S\middle|\theta_C \wedge \theta_T\right)\right] + P\left[\theta_R\middle|\left(T\middle|\theta_C \wedge \theta_T\right)\right] P\left[\left(T\middle|\theta_C \wedge \theta_T\right)\right]$$
$$+ P\left[\theta_R\middle|\left(M\middle|\theta_C \wedge \theta_T\right)\right] P\left[\left(M\middle|\theta_C \wedge \theta_T\right)\right]$$

and so on.

A brief comparison of the various pieces of evidence

There is one mass conservation model that can be evaluated on the basis of forcing functions.

θ_T is actually mass conservation based on the forwards rationale , although the residuals are not summed together

θ_C compares various forcing functions that are derived from the measurements

θ_R is based on the forwards forcing function when successful, but with additional sub-events introduced to relate to the removal.

5. A Prototype Evaluation System

This section describes a prototype evaluation system developed specifically to investigate inspector interface issues. The system focuses on the standard 3-tank arrangement that is located normally between the solvent extraction and polishing cycles 1, or between polishing and concentration. To do this the system also has to take boundary activities into account. The system assumes that process stages, both upstream and downstream, would operate as 'black boxes' (i.e. inventory histories would not be a available), so boundary activities have to be inferred from tanks located immediately upstream and downstream of these process stages. There would also be a recycle tank making 6 tanks in all. The system model used in the materials propagation studies is kept simple. Both the uranium and fission product content in this part of the plant would be low, so the mass content in each tank can be represented by three components (states): bulk, plutonium content and acid content. The minimisation process needed to determine optimum removals of material is performed by the Scientific Python optimisation routine fmin, which applies the Nelder-Mead (1965) simplex algorithm and does not require gradient evaluations.

The system takes its data from a detailed simulation of this part of a reprocessing plant. The system lacks tools to filter out process and measurement fluctuations, so comparisons that are made between plant and calculated histories will be somewhat peculiar to the implementation. A different interface would have to be developed before the computer package could be applied to data collected from an operational plant.

Sliders are provided so that most system parameters would be transparent to the inspector. There are a few exceptions: estimation error variances, σ_γ^2 and measurement error variances σ_ε^2. Inspectors are well versed in the characterisation of random measurement error sources, so these are likely to remain unchanged also. A discussion of the estimation error variances is given below. The odds in favour of a non-safeguards related event $P\left[\theta_T \middle| \overline{S}\right]$ as opposed to a safeguards related event $P\left[\theta_T \middle| S\right]$ was pre-set at a very pessimistic 400:1. Although a number of parameters are contained in the simple, configuration-specific rule-base that assigns probabilities: $P\left[\theta_C \middle| \left(M \middle| \theta_T\right)\right]$, $P\left[\theta_C \middle| \left(S \middle| \theta_T\right)\right]$ and $P\left[\theta_C \middle| \left(T \middle| \theta_T\right)\right]$, the outcomes are set on sliders, which the inspector can adjust.

Figure 5 shows the measurement anomaly window produced when Example (a) data were evaluated. The results are given in a small table with headings, Rationale, Residual and Probability. The Buffer_Suspect rationale is the only one with a non-zero probability (0.635), so this is selected by the buttons shown immediately below. The simple rule-base now presets the first slider with this probability and divides what remains between the other 2 sliders (i.e. 0.18 each). The 4th slider is now preset at the ratio of "(slider_1 / slider_1 + slider_3)" and locality is preset at 99%. A 91.6% measurement related conclusion is drawn. There is no need to move to the second window (i.e. to simulate a removal). One of the points of providing sliders is to demonstrate to inspectors the sensitivity to the variation of certain aspects on the safeguards related probability. Adjusting either the measurement, transfer or a priori split sliders now just shifts the emphasis from Measurement related to Transfer related without impacting on Safeguards Related. In other words the system is quite confident that the error has no safeguards relevance and so in this case is insensitive to certain aspects of the safeguards related probability. Reducing the odds in favour of a non-safeguards related event $P\left[\theta_T \middle| \overline{S}\right]$ from 400:1 to 60:1 increases the safeguards related probability from 0.084% to just 1.921%. Reducing it further to 23:1 gives 6.66% which might be sufficiently large to lead the inspector to assess the possibility of a removal, where no sensible removal would be found. Because the prior odds are subjective, the intent is to provide the simple slider tool to encourage the inspector to experiment with a range of prior odds as just illustrated.

Figure 6 shows the removal window produced when Example (b) data were evaluated. A removal of 0.493 m^3 of solution is required to compare the model with the data when the forwards rationale is selected. The correlation is high resulting in a probability of 99%, so this is selected by the buttons shown immediately below. The simple rule-base now presets the second slider with this probability and divides what remains between the other two sliders (i.e. 0.05 each). The table suggests an alternative, that of 'Recycle_suspect'. This simply encourages the inspector to check that recycle tank instrumentation is in order. The safeguards related probability is significant (32.9%). A very high number at this stage could escalate the investigation prematurely. Moving the measurement slider has no effect. Moving the transfer slider all the way to 1 shifts the emphasis away from safeguards related, reducing it to 2.46%. This might still be significant and require further investigation. Reducing the Locality slider merely scales all probabilities, again making the inspector aware that something might have happened elsewhere. Reducing the odds in favour of a non-safeguards related event $P\left[\theta_T \middle| \overline{S}\right]$ from 400:1 to 38:1 increases the safeguards related probability from 32.9% to 83.1%; a probability of 96.7% is obtained with odds of 23:1.

When Example(c) data were evaluated, a removal of 2.35 m^3 of solution was required to correlate the model with the data when the forwards rationale was selected. This single piece of information would be sufficient to alert the inspector. The correlation however had a very low probability, although visually the revised model predicted history correlated with the trend very closely. This issue relates to the way batch transfers are estimated in continuously fed or emptied tanks. Safeguards related probabilities of 2.9% and 10.7% were obtained when the estimation error standard deviation was increased by factors of 4 and 5 respectively. Examples (a) and (b) were re-evaluated to examine the effect of increasing the estimation error standard deviation by 4. Safeguards probabilities of 99% and 31.5% respectively were returned, suggesting that these examples were relatively insensitive to this parameter. In other words a common estimation variance would suffice for all examples. However it does raise the question whether a flow rate estimator should be implemented to generate a flow rate trend, which the system could reason with. This would represent an enhancement to the approach described here, and hence is outside the scope of this paper.

6. Conclusions

Over the past decade, experience has built up into the development of computer evaluation tools for inspectors. Transferring an approach from an outline to a field implementation is both time-consuming and expensive. Real (usually confidential) data have to be introduced into the development at an early stage. The end-user has to be involved also. This paper gives an initial approach, which is founded on such experiences. Some of the subtleties in the computer interface reflect this. The concept must be shown to have a realistic chance of success before funding can be committed towards field implementation. The approach described here has a well-founded structure and very few hidden parameters, which the inspector can relate to. However, it is still at its first stage of development. A prototype has been developed for implementation on the standard three-tank configuration (Bevan et al., 2011). This awaits inspector evaluation.

The proposed system represents a considerable extension to all current systems, which have limited reasoning capabilities, A brief comparison of current user interfaces can be found in Howell et al., 2010. Event marking combined with sequential testing of resulting residuals (Burr et al. , 2011) is an alternative approach that has yet to be developed to any realistic extent. Such an approach does look across events, but has probably relative low detection probability for any specific scenario, because it's not tuned to any particular scenario. Any comparison with the proposed methodology would be difficult, because it would force the described approach (which is intended to facilitate subjective reasoning) into an objective/quantitative mode.

Material might be removed from certain parts of the plant that are relatively opaque to the inspector. In such situations it is likely that there would be secondary effects that would be more visible, especially if the inspector has a means of comparison with what is normal. Howell (2009) has explored this, and proposed that material propagation can represent a reference to normality again. This represents an extension to the approach discussed here. Other extensions have been discussed in the text.

Quantitative reasoning here is based on simple structurally correct models that are driven by forcing functions, which can be estimated in a number of different ways. The paper suggests one application specific approach to the formulation of these different ways, to their presentation and to their integration into the evaluation process. As such the end user might find the underlying theory opaque. It might be worth considering alternative presentations that are more explicit in their approaches, whilst remaining attractive to the nuclear safeguards inspector. Safety inspectors might wish to consider whether a similar approach might be of use when faced with incidents like Three Mile Island and Fukushima. Anecdotal evidence suggests the need to reason quantitatively.To a certain extent the approach is similar to what one might envision, for example, in recognizing unnatural disease outbreaks or in modern quality control applications. In monitoring for unnatural disease outbreaks (Burr et al. 2007) one must use observed patient symptoms that are measured with error, account for the background incidence of possible diseases (the "prior" probability of each candidate disease), and somehow aggregate estimated probabilities of candidate diseases by patient over space and time. Burr et al (2007) allowed the user to vary the prior disease probability for the same reason our safeguards sliders are provided. In modern quality control, described for example in (Hines et al., 2007), there is a need to monitor both for product quality drift and for instrument calibration drift, so the task of learning normal behaviour is challenging for reasons that are similar to the reasons it is challenging to learn normal behaviour in our described safeguards application. In both examples, the overall approach is likelihood based, relying on a probability model for residuals associated with events that is learned from training data that is assumed to not contain the event of concern. In our case, the event of concern is any safeguards related event. In monitoring for unnatural disease outbreaks, the event of concern is biological attack. In modern quality control setting, one event of concern is product quality drift, but instrument drift is a nuisance noise source.

Acknowledgements

The Next Generation Safeguards Initiative of the National Nuclear Security Administration part funded this work.

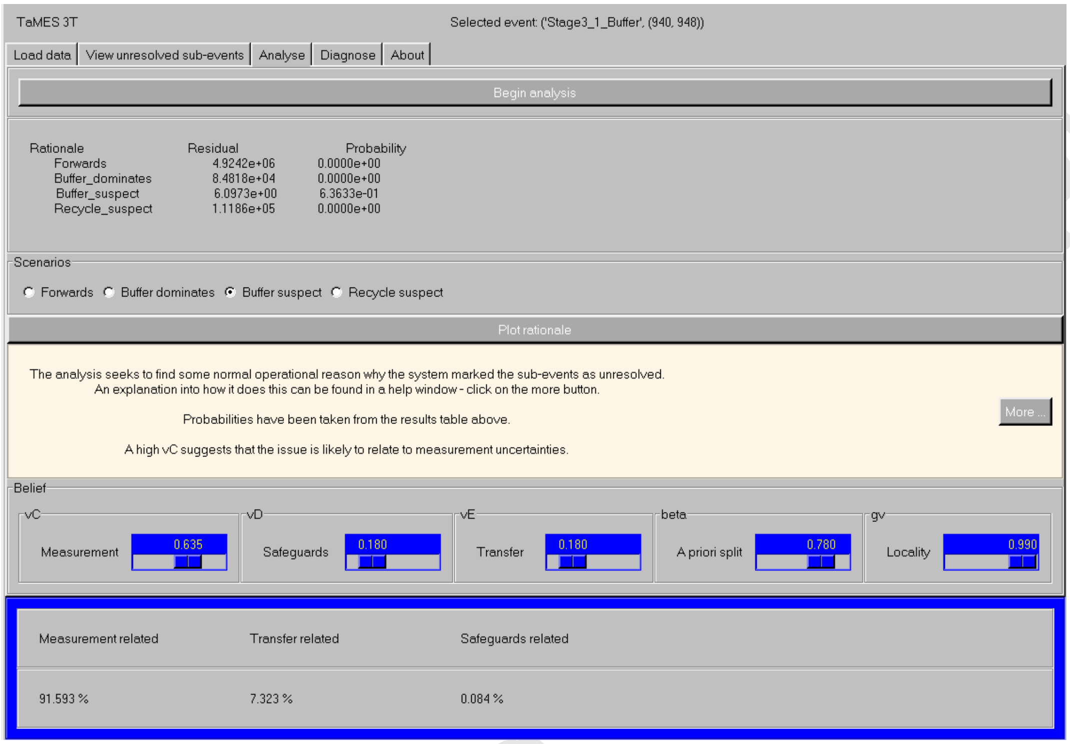

Figure 5 Measurement anomaly window

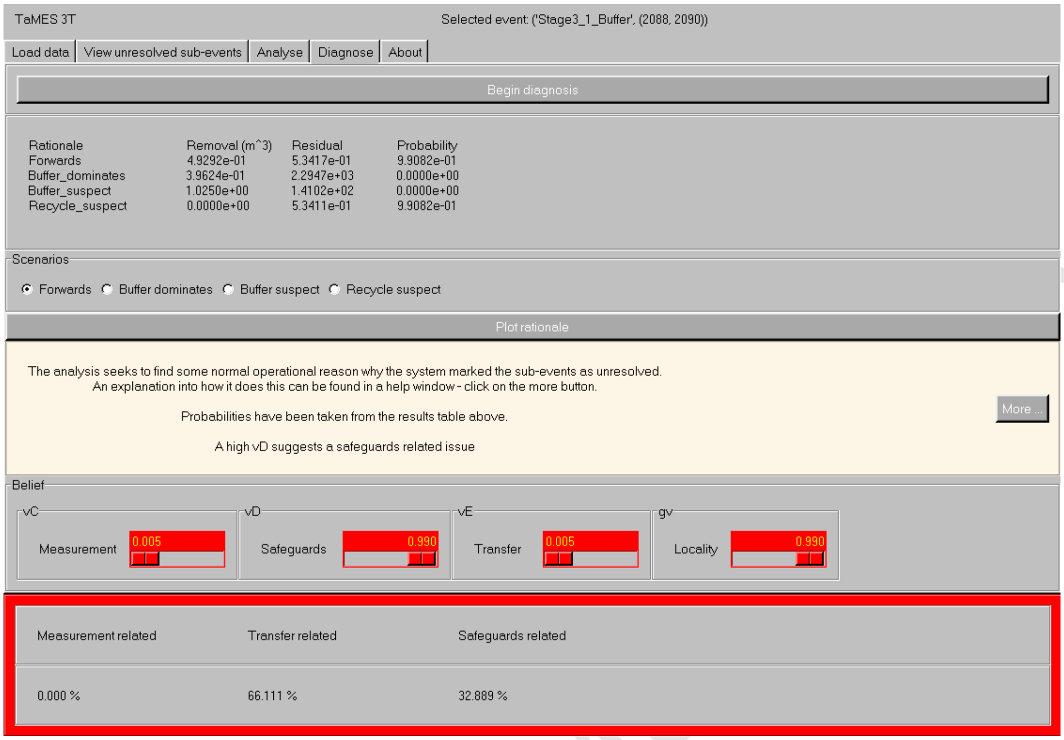

Figure 6 Removal window

References

Bevan, G., J. Howell, T. Burr (2011). A Tank Monitoring Evaluation System for a Standard Three-Tank Arrangement, INMM52, Palm Desert, USA.

Burr T., F. Koster, R. Picard, D. Forslund, D. Wokoun, E. Joyce, J. Brillman, P. Froman, J. Lee (2007). Computer-Aided Diagnosis with Potential Application to Rapid Detection of Diseases Outbreaks, Statistics in Medicine 26: 1857-1874.

Burr, T., M. Ehinger, J. Howell (2008). A Review of Process Monitoring for Safeguards, 8th Int Conf on Facility Operations - Safeguards Interface, Portland, Or.

Burr, T., M. Suzuki, M., J. Howell, J., M. Hamada, C. Longo, C. (2011). Signal Estimation and Change detection in Tank Data for Nuclear Safeguards, Nuclear Instruments and Methods in Physics Research 640 200–221.

Burr T., M. Suzuki, J. Howell, M. Hamada (2012). Loss Detection Results on Simulated Tank Data Modified by Realistic Effects, Journal of Nuclear Science and Technology 49(2), 209-221.

Charbonnier, S., C. Garcia-Beltan, C. Cadet & S. Gentil (2005). Trends Extraction and analysis for complex system monitoring and decision support, Engineering Applications of AI, vol. 18, pp. 21-36.

Davis, R. (1984). Diagnostic Reasoning Based on Structure and Behaviour, Artificial Intelligence, vol. 24, pp 347-410.

Dekens, J. P., Goerten, J., Lahogue, Y., & Wagner, G. (1995). An Integrated Safeguards System for Large Reprocessing Plants, 5th International Conference on Facility Operations-Safeguards Interface, Jackson Hole, Wy.

Demuth, S., B. Boyer, T. Burr, J. Howell, M. Schanfein & R. Stevens (2010). Proliferation Resistance and Safeguards, Handbook of Nuclear Engineering, Vol. 4, Cacuci, D. G. (Ed.) Springer.

Ehinger, M.H., B. Chesnay, C. Creusot, J.Damico, S. Johnson, J. Wuester, S. Masuda, M. Kajii (2004). Solution Monitoring Applications for the Rokkasho Reprocessing Plant, 7th ANS/INMM International Conference on Facility Operations-Safeguards Interface, Charleston, SC.

Friend, P. (2008). Urenco's Views on International Safeguards Inspection, *8th* ANS/INMM Int. Conf. on Facilities Operations, Portland, Or.

Gentil, S. & J. Montmain (2004a) . Hierarchical representation of complex systems for supporting decision making, Advanced Enginneering Informatics, 18, 143-159.

Gentil, S., J. Montmain & C. Combastel (2004b). Combining FDI and AI Approaches Within Causal-Model-Based Diagnosis, *IEEE Trans. SMC-B*, 34(5).

Hines, W. & D. Garvey (2007). Process and Equipment Monitoring Methodologies Applied to Sensor Calibration Monitoring, Quality and Reliability Engineering International, 23, 1123-135.

Howell, J., R. Binner, G. Bevan, and B. Sirajov (2009). Tank Monitoring Evaluation Systems: Methods and Algorithms, in 50th INMM Meeting, Tucson, Az.

Howell, J. & S. J. Scothern (2000). An Explicit Model-based Diagnostic Approach in a Plutonium Nitrate Storage Facility, Control Engineering Practice, vol. 8, pp. 645-656.

Howell, J. (2009). Differentiating Between Normal and Abnormal Situations With Solution Monitoring Systems in 50th INMM Meeting, Tucson, Az., 2009.

Howell, J., Ehinger, M., Burr, T. (2010). The Inspector Interface in Solution Monitoring Systems, 51st INMM Meeting, Baltimore.

IAEA (1970).Treaty on the Non-Proliferation of Nuclear Weapons (NPT), Information Circular INFCIRC/140.

IAEA (1972). The Structure and Content of Agreements Between the Agency and States Required in Connection with the Treaty on the Non-Proliferation of Nuclear Weapons, Information Circular INFCIRC/153.

IAEA (1997). Model Protocol Additional to the Agreements(s) Between State(s) and the International Atomic Energy Agency for the Application of Safeguards, Information Circular INFCIRC/540.

Janssens-Maenhout G. & Dechamp, L. (2004). Process monitoring appropriate for near-real-time accountancy, JNMM 32(3), pp 10-16.

Miller, E.C. & J. Howell (1999). Tank Measurement Data Compression For Solution Monitoring, JNMM, vol. 27, no. 3, pp. 25-32.

Miller, C., A. Cubbage, D. Dorman, J. Grobe, G. Holahan & N. Sanfilippo (2011). Recommendations for Enhancing Reactor Safety in the 21st Century. US NRC.

Nackaerts, N. (2011). A Changing Nuclear Landscape: Preparing for Future Verification Challenges. International Forum on Peaceful Use of Nuclear Energy and Nuclear Non-Proliferation. Vienna.

Nelder, J.A. & R. Mead (1965). Simplex Method for Function Minimization, Computer Journal, 7(4), pp308—313.

Patton, R. J., P.M. Frank & R. N. Clark (eds) (2000). Issues of Fault Diagnosis for Dynamic Systems, Springer.

Scothern, S. J., & Howell, J. (1997). A physical-model-based diagnostic aid for safeguarding nuclear material in a liquor storage facility. Journal of Nuclear Materials Management, 25(4), 20-29.

Sirajov, B., X. Wang (2008). Tank Monitoring Evaluation System, in 8th ANS/INMM International Conference on Facility Operations, Portland, Or.

Suziki, M., Hori, M., Nagaoka, S. & Kimura, T. (2009). Study on Loss Detection Algorithms Using Tank Monitoring Data, Journal Nuclear Science Technology Vol. 46(2), pp.184-192

Walker, J. S., (2004). Three Mile Island: A Nuclear Crisis in Historical Perspective**,** Berkeley: University of California Press.

Highlights

- Nuclear safeguards inspector interfaces for reasoning about abnormal events.
- Semi-qualitative reasoning based on sub-events.
- Quantitative reasoning about simple models that are driven by forcing functions.
- Fusion through Bayesian updating.
- Application to a standard 3-tank arrangement in a nuclear fuel reprocessing plant.
- Examples that show responses to a measurement anomaly and two different removals.