MASTER

UCRL-87538, Revision #1
PREPRINT

Conf-320656--6 Rev. 1

A KALMAN-FILTER MODEL FOR DETERMINING
BLOCK AND TRICKLE SNM LOSSES *

UCRL--87538-Rev. 1

DE93 000751

NOTICE

R. E. Barlow
M. J. Durst
N. G. Smiriga

**PORTIONS OF THIS REPORT ARE ILLEGIBLE. It
has been reproduced from the best available
copy to permit the broadest possible avail-
ability.**

Symposium on Low Probability,
High Consequence Events
Arlington, VA
June 14-17, 1982

July, 1982

**Lawrence
Livermore
National
Laboratory**

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

* This work was performed under the auspices of
the Department of Energy by Lawrence Livermore
National Laboratory under Contract No. @-7405-Eng-48.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. It contains certain information that may be confidential in nature, and its disclosure could be injurious to the national defense. It is to be understood that neither the United States Government nor any agency thereof, nor any person acting on behalf of the United States Government, is to be held responsible for any errors or for any consequences arising from the use of the information contained herein. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISTRIBUTION OF THIS REPORT IS UNLIMITED

gdy

Abstract

This paper describes an integrated decision procedure for deciding whether a diversion of SNM has occurred. Two possible types of diversion are considered: a block loss during a single time period and a cumulative trickle loss over several time periods. The methodology used is based on a compound Kalman filter model. Numerical examples will illustrate our approach.

KEYWORDS:

1. Introduction

Commercial facilities which handle Special Nuclear Material (SNM) are required by the Nuclear Regulatory Commission to test for diversion of SNM. Although this provides an obvious application for risk analysis, a satisfactory analysis has yet to be done. This is largely due to the difficulty of measuring the various costs (for instance, the expected cost to society of an unnoticed diversion). However, the other component of a risk analysis- the probability, given observed data, that a diversion of SNM has occurred- can be calculated using existing methods. This probability is an essential component of any successful decision procedure, but methods currently in use employ it only in an ad hoc fashion. In this paper we indicate one path to a more realistic procedure, beginning by outlining current practice.

The Current Approach.

In the United States, regulatory practice in the area of safeguards accounting is based on the calculation of the inventory difference (ID) and its standard deviation. The inventory difference may be written

$$ID = BI + A - EI - R \quad (1.1)$$

where BI = measured beginning inventory,

A = measured arrivals,

EI = measured ending inventory,

R = measured removals.

Since all measurements are subject to error, statistical methods must be used to decide whether a significant loss of nuclear material has occurred. Ordinarily the calculated variance of ID, based on equation (1.1), is used to calculate a "Limit of Error for Inventory Difference" (LEID), where

$$LEID = +2 \text{ Variance (ID) } .$$

In practice, limits are placed on the magnitude of LEID. If

- (1) LEID is within allowable limits, and
- (2) $ID < LEID$,

the inventory difference is assumed to be due to normal process measurement variations. If either condition is violated, the process variation is assumed to be abnormal and the cause of abnormality must be found.

As an estimator of loss or diversion, ID is known to be insensitive to repeated small losses, such as those which occur in process holdup. For this reason, the Cumulative Inventory Difference (CID), along with its standard deviation, is often used as a basis for making decisions about whether trickle loss has occurred. For further details on ID and CID, see Jaech³.

Disadvantages of the Current Approach.

The advantage of the current approach is its simplicity. ID and CID are basic accounting statistics, easy to understand and interpret. However, this simplicity has its limitations. For example, experience with a process makes possible the assessment of normal holdup levels, and these can be compared to the actual rate of trickle loss; but such a procedure must take place outside the framework of ID, CID, and their associated Limits of Error. A more critical problem is the meaning of tests involving ID and CID. In most standard practice, the hypothesis of no loss is tested at the 95% confidence level. The result of such a test is either rejection of the hypothesis or its acceptance. Neither result gives a measure of the strength of our conviction regarding a possible abnormality. This results in absurdities such as the declaration by a process manager that the process is normal, when in fact the last dozen ID measurements have been barely within LEID.

Current methods of nuclear material accounting owe much to quality control procedures, so it is not surprising that some of these same criticisms have been voiced regarding quality control procedures in general. The Bell Laboratories Quality Assurance Center has recently implemented a new Quality Measurement Plan² which "replaces the T-rate system, which evolved from the pioneering statistical work of Shewhart and Dodge [in control methods] during the 1920's and 1930's at Bell Laboratories." The new Quality Measurement Plan differs from much earlier quality control work in that it uses a Bayesian approach to decision making. A well written but elementary exposition of the Bayesian approach is provided by Lindley⁴.

The Kalman Filter Approach.

In 1975, Pike, Morrison and Holland^{5,6} advocated a Kalman Filter approach to nuclear material accountability. However, they remain within the normal framework of tests of a null hypothesis; and although they present filters for both block and trickle loss, they do not integrate them into a single decision making system. The Kalman Filter modelling which we pursue in this paper is more closely akin to the work of Harrison and Stevens¹, although their primary interest was in forecasting rather than in determining the present state of a process.

We develop a compound Kalman Filter model which allows a decision maker to decide at each time period whether the data indicates a diversion. A block loss, by definition, will be a substantial amount which we hope to detect at the end of the period in which it occurs. A trickle loss, on the other hand, is a smaller amount which we do not expect to detect in a single occurrence. A trickle loss may consist of a diversion or process holdup (or both), while a block loss is always a diversion. We give two models for the process during each time period; in one, a block loss is assumed to have occurred, while in the other, only the usual trickle loss takes place. Since there are two models at each time period, a fully Bayesian analysis would require 2^n models at the end of n time periods, which is computationally untenable. In such circumstances, Harrison and Stevens recommend approximating the process by a Markov process, which allows one to consider only four models at each time period. We make a simpler

approximation, which rests on the assumption that a block loss is a low-probability event. With this approximation we may consider only two models each period, with all inference conditional on the assumption of no block loss in past periods (which has probability virtually equal to 1 as long as we have never come close to deciding that a block loss has occurred). By comparing these two models, we decide whether a block loss has occurred, and if we decide that it has an investigation is initiated. Since trickle loss, at least in the form of process holdup, is always assumed to occur, we will never decide that no trickle loss has occurred. We will either decide that a trickle diversion has occurred over several past periods, or we will decide that we as yet are unconvinced that a trickle loss beyond the normal holdup has occurred.

Figure 1 - Diagram of possible decision sequences relative to diversion of special nuclear material.

In Figure 1, $\beta(1), \beta(2), \dots$ etc. denote the amount of possible but unknown block losses during their respective time periods. The amount of possible but unknown trickle losses are denoted by $\tau(1), \tau(2), \dots$ etc. In our approach, we shall have two models: one model for block loss, say M_B , and one model for trickle loss, say M_T . We believe that model M_B holds with probability $p(M_B)$ and model M_T with probability $1 - p(M_B)$. Given data D , $p(M_B|D)$ is our updated probability for the block loss model M_B . If our updated probability for the block loss model is too high, then we will decide to investigate the possibility of a block loss. A decision regarding possible trickle loss, on the other hand, is based on the probability that loss beyond the normally expected holdup has occurred over several time periods; i.e.

$$P\{\tau(1) + \dots + \tau(t) > c|D\}$$

where c is the normally expected holdup over t time periods. Thus, as indicated in Figure 1, our decision sequence is the customary one; at each time period we either decide that a substantial block loss has occurred in the most recent period, that an unusually large trickle loss has been occurring in the past few periods, or that no block loss is likely to have occurred and that trickle loss is within acceptable limits. Our decision procedure does not formally permit the conclusion that a block loss has occurred other than within the most recent period, but we will show in Section 3 that certain trickle alarms indicate the presence of an undetected block loss in some past period.

2. A Compound Kalman Filter Model For Block and Trickle Losses

During time period t , a number of measurements related to the process will be made, as illustrated in Figure 2. Decisions will be based on the measurements made during the time period, as well as on knowledge of the process.

In Figure 2, the process variables are:

$\theta_B(t)$ = Beginning inventory level in time period t ;

$\theta_A(t)$ = Arriving inventory in time period t ;

$\theta_R(t)$ = Removed inventory in time period t ;

$\theta_E(t)$ = Ending inventory in time period t .

These are the values which would be observed were there neither measurement error nor diversions.

Figure 2 - Manufacturing process parameters as functions of time.

The numbers actually observed in the absence of any diversion are:

$$\begin{aligned}y_A(t) &= \theta_A(t) + v_A(t), \\y_R(t) &= \theta_R(t) + v_R(t), \\y_E(t) &= \theta_E(t) + v_E(t).\end{aligned}\tag{2.1}$$

In Kalman filter theory, equations (2.1) are called "Observation Equations". The random noise terms $v_A(t)$, $v_R(t)$ and $v_E(t)$ correspond to measurement error.

We will model the evolution of the manufacturing process parameters through the "System Equations":

$$\begin{aligned}
\Theta_B(t) &= \Theta_E(t-1), \\
\Theta_A(t) &= \Theta_A(t-1) + w_A(t), \\
\Theta_R(t) &= \Theta_R(t-1) + w_R(t), \quad (2.2) \\
\Theta_E(t) &= \Theta_B(t) + \Theta_A(t) - \Theta_R(t) \\
&= \Theta_E(t-1) + \Theta_A(t-1) - \Theta_R(t-1) \\
&\quad + w_A(t) - w_R(t).
\end{aligned}$$

These equations model the evolution of a process in a steady state of operation. The random noise terms $w_A(t)$ and $w_R(t)$ are present because we cannot realistically expect the process to be at identical levels from one period to the next, even if it is designed to be in a steady state. Of course, few processes operate for long in a steady state. In a real process, the equations (2.2) would be replaced by equations giving the intended changes in process levels between period $t-1$ and period t . In general, these equations will differ from time period to time period.

Equations (2.1) and (2.2) describe the process completely in the absence of any kind of loss. However, we must realistically always expect some kind of loss, if only due to holdup. Since this loss will affect the evolution of the process parameters, the system equations are unrealistic, and should be as follows:

$$\begin{aligned}
\Theta_B(t) &= \Theta_E(t-1), \\
\Theta_A(t) &= \Theta_A(t-1) + w_A(t), \quad (2.3) \\
\Theta_R(t) &= \Theta_R(t-1) + w_R(t), \\
\Theta_E(t) &= \Theta_E(t-1) + \Theta_A(t-1) - \Theta_R(t-1) + w_A(t) \\
&\quad - w_R(t) - \tau(t), \\
\tau(t) &= \tau(t-1) + w_\tau(t),
\end{aligned}$$

where $\tau(t)$ is the amount of the trickle loss and $w_\tau(t)$ is a random noise component corresponding to random changes of the true level of trickle loss. Under the more realistic assumption of trickle loss, the behavior of the process is completely described by (2.1) and (2.3) instead of (2.1) and (2.2). At each time period we will further consider the possibility that a single large loss, i.e. a block loss, can occur, as described by the following observation equations:

$$\begin{aligned}
y_A(t) &= \Theta_A(t) + v_A(t), \\
y_R(t) &= \Theta_R(t) + v_R(t), \quad (2.4) \\
y_E(t) &= \Theta_E(t) - \beta(t) + v_E(t),
\end{aligned}$$

where $\beta(t)$ is the amount of block loss. The block loss model is completely described by (2.3) and (2.4). The trickle loss model, like the simpler model with no loss, evolves continuously from time period to time period. The block loss model does not evolve this way. We have not included an equation describing the evolution of $\beta(t)$; it is introduced as a new possibility at each time period.

In addition, the block loss term appears in the observation equation rather than the system equation, reflecting this lack of evolution. If we fail to detect a block loss, our modelling in future periods considers it to be subsumed under the trickle loss. We choose, for each time period, a suitable normal distribution for $\beta(t)$.

We shall now show how we use the block and trickle models. We first consider the situation at the initial time period, just after cleanout. Let $p(M_B)$ be our probability that the block model, M_B , is correct. Then $1-p(M_B)$ is our probability that the trickle model, M_T , is correct. The initial distribution of the process parameters has to be specified; we assume that $\theta_{EI}(0)$, $\theta_A(0)$ and $\theta_R(0)$ all have specified normal distributions. Furthermore, the trickle loss $\tau(0)$ is also assumed to be normally distributed. The distribution of the noise terms in the system and observation equations must all be specified. We assume that they are normally distributed with mean zero and variance independent of the time period. In addition, the noise terms are also assumed to be independently distributed, with distributions independent of the distribution of the process parameters, the trickle loss, and the block loss. After the data, say D , has been observed (i.e., $y_A(1)$, $y_R(1)$, and $y_{EI}(1)$ become available), the probability $p(M_B|D)$, i.e., the conditional probability that the block model is correct given the data, is computed (formulas for the computation are given in Harrison and Stevens¹). We can also compute the posterior distribution for the magnitude of $\beta(1)$. Then we must decide whether or not a block loss has occurred. If the decision is that a block loss has occurred, we follow it up by the necessary actions (such as rechecking inventories and/or a cleanout) and restart the whole procedure. If we decide to act as if no block loss has occurred, then this means that we accept the trickle model M_T . As already mentioned, a trickle loss is thought of as one which can only be detected after several occurrences. Therefore, since this is the first time period, we use the M_T model to update the distributions of the process parameters and the trickle loss, expecting no alarm. For the next time period all the posterior distributions for all parameters but the block loss parameter are used as prior distributions. The block loss distribution of $\beta(t)$ is not updated from period to period. Our posterior probability $p(M_B|D)$ may or may not, depending on our belief, be used as our prior probability for the M_B model in the next time period. After the first time period, if we decide no block loss has taken place, then the distribution of the sum of possible trickle losses in the current and previous time period will be calculated. Our decision regarding the possibility of cumulative trickle losses will be based on this distribution (a normal distribution).

In summary, we have two models. The block model M_B

$$y_A(t) = \theta_A(t) + v_A(t),$$

$$y_R(t) = \theta_R(t) + v_R(t),$$

$$y_{EI}(t) = \theta_{EI}(t) - \beta(t) + v_{EI}(t),$$

$$\begin{aligned}
\theta_B(t) &= \theta_E(t-1), \\
\theta_A(t) &= \theta_A(t-1) + w_A(t), \\
\theta_R(t) &= \theta_R(t-1) + w_R(t), \\
\theta_E(t) &= \theta_E(t-1) + \theta_A(t-1) - \theta_R(t-1) + w_A(t) \\
&\quad - w_R(t) - \pi(t), \\
\pi(t) &= \pi(t-1) + w_\pi(t),
\end{aligned}$$

is considered to hold with probability $p(M_B)$, and the trickle model M_T

$$\begin{aligned}
y_A(t) &= \theta_A(t) + v_A(t), \\
y_R(t) &= \theta_R(t) + v_R(t), \\
y_E(t) &= \theta_E(t) + v_E(t), \\
\theta_B(t) &= \theta_E(t-1), \\
\theta_A(t) &= \theta_A(t-1) + w_A(t), \\
\theta_R(t) &= \theta_R(t-1) + w_R(t), \\
\theta_E(t) &= \theta_E(t-1) + \theta_A(t-1) - \theta_R(t-1) + w_A(t) \\
&\quad - w_R(t) - \pi(t), \\
\pi(t) &= \pi(t-1) + w_\pi(t),
\end{aligned}$$

is considered to hold with probability $1-p(M_B)$. The trickle model subsumes what is normally called the no-loss model, as we always expect some trickle loss in the form of holdup.

For each time period we consider first the block model; if we reject it, then we consider the trickle model. If the posterior distribution of the sum of the trickle losses over the present and past periods does not alarm us, then we proceed to the next time period. For that period we choose a distribution for block loss, we may or may not use $p(M_B|D)$, and we use the present posterior distributions for all remaining parameters as the new prior distributions. The next section will illustrate our procedure with respect to a smaller model, for simplicity, and show in detail how our Kalman filter can be used.

3. Numerical Examples for a Simplified Model

In order to clearly illustrate the salient features of these models, we will study closely a simplified model with only one measurement each period. In the absence of any diversions our process model is

$$\begin{aligned}
y(t) &= \theta(t) + v(t), \\
\theta(t) &= \theta(t-1) + w_\theta(t),
\end{aligned} \tag{3.1}$$

where $v(t)$ models our uncertainty with respect to the measurement error and $w_\theta(t)$ models our uncertainty about the difference in the actual amount processed between time period $t-1$ and t . At time t , $\theta(t)$ is the quantity of interest, but we can only observe $y(t)$. We assume that all variables in (3.1) are normally distributed. More specifically we assume that $v(t)$ and $w_\theta(t)$ have distributions which are independent of t , and that they have zero mean and variances $\sigma^2(v)$ and $\sigma^2(w_\theta)$ respectively. In addition $v(t)$ and $w_\theta(t)$ are assumed to be independent of each other, and of all $v(s)$ and $w_\theta(s)$ for $s \neq t$, and of $\theta(0)$. We further have to specify the mean and variance of $\theta(0)$; let $\mu(\theta(0))$ denote the mean of $\theta(0)$ and let $\sigma^2(\theta(0))$ denote the variance of $\theta(0)$. In the examples that follow we shall consider the following numerical values:

$$\begin{aligned}\mu(\theta(0)) &= 100, & \sigma^2(\theta(0)) &= 1, \\ \sigma^2(v) &= 4, & \sigma^2(w_\theta) &= 1.\end{aligned}$$

All values are expressed in pounds. Figure 3 displays simulated output values based on our simplified model without diversion.

Figure 3 - Simulated output -
no diversion or holdup.

The simplified trickle model M_T is:

$$\begin{aligned}y(t) &= \theta(t) + v(t), \\ \theta(t) &= \theta(t-1) - \tau(t) + w_\theta(t), \\ \tau(t) &= \tau(t-1) + w_\tau(t).\end{aligned}\quad (3.2)$$

We will use $\mu(\tau(0)) = 0.5$, $\sigma^2(\tau(0)) = 0.25$, $\sigma^2(w_\tau) = 0.0025$, where $\mu(\tau(0))$ and $\sigma^2(\tau(0))$ are the mean and variance of $\tau(0)$, and $\sigma^2(w_\tau)$ is the variance of $w_\tau(t)$. As with $w_\theta(t)$, $w_\tau(t)$ is normally distributed, with mean zero. Its distribution does not depend on t , on other noise terms (including $w_\tau(s)$ for $s \neq t$), or on initial distributions for $\theta(0)$ and $\tau(0)$. We shall consider several examples for the trickle model (3.2). The first example will consider a small trickle due to expected holdup rather than possible diversion. Let us assume that this trickle is 0.5 lbs per period. The next figure shows the simulated output, assuming a 0.5 lbs holdup per period. More precisely, if $y(t)$, $t = 1, \dots, 10$ are the y -values of figure 3, then the y -values of figure 4 are $y(t) - (0.5)t$, $t = 1, \dots, 10$ (i.e., we assumed in generating the output that the variance $\sigma^2(w_\tau)$ is zero).

Figure 4 - Simulated output - holdup of
0.5 lbs per period.

The next figure shows the distribution of the total trickle loss for $t = 1, \dots, 10$.

Figure 5 - Densities for total trickle loss
($t = 1, \dots, 10$) under model M_T (3.2) with trickle of
0.5 lbs/period (due to holdup).

The next example shows again the densities for the total trickle loss for $t = 1, \dots, 10$ but this time we assume a trickle of 1.5 per period (i.e., namely 0.5 due to holdup and 1.0 due to trickle diversion). The y -values that were used are $y(t) = (1.5)t$, $t = 1, \dots, 10$, with the values for $y(t)$ used in Figure 3.

Figure 6 - Densities for total trickle loss
($t = 1, \dots, 10$) under model M_T with trickle loss of
1.5 lbs/period (0.5 lbs due to holdup, 1.0 lbs due to
diversion)

If we expect a holdup of 0.5 lb per period, then we would consider the behavior in Figure 5 to be normal. After 10 periods the probability that the sum of the trickle losses is greater than 5 lbs is roughly 50%, which is, of course, to be expected. If this probability is higher, then we would be alarmed. Using the data which lead to Figure 6, this same probability is now 83%.

The next example demonstrates how a decrease of the variance of $\pi(0)$ will make the process less sensitive to detecting trickle loss. Figure 7 gives the same information as Figure 6, but extended to cover 20 time periods. Figure 7 is a variant of the box-and-whisker plot of Tukey⁷, with the box extending one standard deviation from the mean in either direction, and the whiskers extending two standard deviations from the mean. Figure 8 uses the same data, but changes the variance of $\pi(0)$ to $\sigma^2(\pi(0)) = 0.0625$.

Figure 7 - Densities for total trickle loss
 (t = 1,...,20) under model M_T , diversion of 1.5
 lbs/period, with $\sigma^2(\pi(0)) = 0.25$.

Figure 8 - Densities for total trickle loss
 (t = 1,...,20) under model M_T , diversion of 1.5
 lbs/period, with $\sigma^2(\pi(0)) = 0.0625$.

The larger variance for $\pi(0)$ gives distributions for the total trickle loss which are much more in agreement with the actual loss than those given by the smaller variance.

Next we consider the simplified block model M_B :

$$y(t) = \theta(t) - \pi(t) + v(t),$$

$$\theta(t) = \theta(t-1) - \pi(t) + w_\theta(t), \quad (3.3)$$

$$\pi(t) = \pi(t-1) + w_\pi(t).$$

For the following examples, we keep the same values for the mean and variance of $\theta(0)$, $v(t)$, $w_B(t)$ and $w_T(t)$ as before. For $\pi(0)$ we will use a mean $\mu(\pi(0)) = 0.5$ and a variance $\sigma^2(\pi(0)) = 0.25$. For the M_B model, let us assume that $\beta(0)$ is normally distributed with mean $\mu(\beta(0)) = 5.0$ and variance $\sigma^2(\beta(0)) = 6.25$. Let $p(M_B) = 0.01$. If we use the simulated data which is displayed in Figure 4 (i.e., just a holdup of 0.5 lbs per period) and if we never update $p(M_B)$ (i.e., as prior probability we always use $p(M_B) = 0.01$) then $p(M_{B|D})$ will fluctuate between 0.0014 and 0.0157 over 10 periods. If we update $p(M_B)$ then $p(M_{B|D})$ will decrease after each period; after the first period $p(M_{B|D}) = 2.5 \cdot 10^{-3}$ and after the tenth period $p(M_{B|D}) = 3.1 \cdot 10^{-7}$. For both procedures, no update and continuous update, we never obtain a probability $p(M_{B|D})$ that will alarm us.

Next we consider an example with the same data used in Figure 4 for $t = 1, \dots, 7$, but for $t = 8$ we use as observed y -value $y(8) = 10$, where $y(8)$ is the value plotted in Figure 4. Below we plot the probability $p(M_{B|D})$ for this scenario.

Figure 9 - $p(M_{B|D})$ versus time: $p(M_B)$ initially 0.01, $p(M_B)$ reset each period, block loss of 10 lbs at $t = 8$.

Table 1 - $p(M_{B|D})$ versus time: $p(M_B)$ initially 0.01, $p(M_{B|D})$ continuously updated, block loss of 10 lbs

Note that in Table I, $p(M_B|D) = 2.1 \cdot 10^{-6}$ after seven time periods and $p(M_B|D) = 1.7 \cdot 10^{-4}$ after eight time periods, i.e. $p(M_B|D)$ has increased roughly 80 times from period 7 to period 8. If $p(M_B|D)$ is continuously updated, a sudden increase will indicate a possible block loss. However, the size of the loss will be more accurately indicated if $p(M_B|D)$ is repeatedly reset. With our present prior distribution, using the reset $p(M_B|D) = 0.4582$ for period 8 ,

$$P(B(0) > 10|M_B) = 0.0228$$

$$P(B(0) > 10) = (0.01)(0.0228) = 2.3 \cdot 10^{-4}$$

and

$$P(B(8) > 10|M_{B,D}) = 0.0379$$

$$P(B(8) > 10|D) = (0.4582)(0.0379) = 0.0174 .$$

The posterior probability of a block loss greater than 10 is roughly 75 times larger than the prior probability. This example shows how easy it is to detect a block loss of 10 lbs or more. A significant increase in $p(M_B|D)$ indicates a possible block loss. Such an increase should be followed by examination of the prior and posterior distributions of the amount of block loss. Specifically, the probability of a loss greater than some critical quantity can be examined to determine whether an alarm is warranted.

The next example shows how a smaller block loss, though not caught by model M_B , may later be detected by model M_T . We use the same data as before except that at period 8 a block loss of 7 occurs. The continuously updated $p(M_B|D)$ decreases from period 1 to period 7 ; at period 7 and 8 we have $p(M_B|D) = 0.2 \cdot 10^{-6}$ and $p(M_B|D) = 1.4 \cdot 10^{-5}$ respectively ; this increase is possibly too small to alarm us. The reset $p(M_B|D)$ varies between 0.0014 and 0.0073 from period 1 to period 7 . At period 8 it changes to $p(M_B|D) = 0.0663$, again an increase which might not be big enough to justify stopping production. If we decide to continue the operation of the plant, then we shall detect at a later period that trickle loss is alarmingly high. Remember that we have assumed that the plant has a holdup of 0.5 lbs per period ; therefore we expect that

$$P(\sum \pi(i) > (0.5)t) \approx 0.5 .$$

At each time period we compute this probability, expecting it to fluctuate around 0.5. With this example, the probability does indeed fluctuate around 0.5 until period 8 and then increases gradually. At period 14 it is 0.71, an amount which can no longer be justified as a fluctuation. Thus the small block loss of 7 lbs at period 8 is identified as a cumulative trickle loss at period 14. This gradual increase of trickle loss per period is the characteristic "footprint" of an undetected block loss in the past, and should guide any investigation once the trickle alarm has been sounded. A retrospective analysis could show that the hypothesis of block loss (at the onset of increase in calculated trickle loss per period) explains the data better than the trickle loss model, thus inducing a successful, if tardy, detection of the 7 lb. block loss. We can even detect smaller block losses, but it will, of course, take longer.

4. Conclusions

The conclusions of this study fall into two categories: technical properties and model criticisms. Certain technical properties of the compound Kalman filter are advanced as information for prospective users of our methods. The model criticisms are evangelical in nature, and by stating them we hope to promote use of these techniques by those as yet unfamiliar with them..

Technical Properties.

The values of distribution parameters, even in our simplest model, must be set carefully. We saw, for example, that too little initial uncertainty about possible trickle loss made the process surprisingly unresponsive to large unexpected losses. A set of distribution parameters can be entirely self-consistent, seem on casual inspection quite sensible, and still produce undesirable behavior of the detection procedure. Thus distribution parameters should not be set arbitrarily or casually, but only after a careful assessment of process and loss uncertainties which takes into account the effect of the parameters on the resulting decision procedure.

Although we have tried to make our examples realistic and representative, we emphasize that this work is not intended as a blueprint for actual applications. The use of these techniques will, of necessity, involve all available information about the dynamic structure of the facility being modelled. A filter model of an actual facility would have hundreds of process parameters, with interrelationships which would have to be considered carefully. Of course, a summary model could be constructed from the full model in the form of equations (2.1)-(2.4), and our results would hold verbatim for that summary model.

Model Criticisms.

The compound Kalman filter model provides a detection process which can compete with currently popular methods. Large block losses are detected handily, while somewhat smaller block losses are often detected later by the trickle model. Trickle losses consistently in excess of the expected holdup are detected rapidly, and smaller trickle losses are detected as the total amount of trickle loss becomes large. In addition, we believe this method has certain advantages not shared by non-Bayesian methods.

In most currently popular methods, decisions are made on the basis of the improbability of observed data under a null hypothesis. By contrast, the compound Kalman Filter uses as a basis for its decisions the ability of different models to explain the data. Although this requires that our uncertainties about possible losses be assessed, it does yield some advantages. For example, the probability of a serious block loss is modelled directly and is not just inferred from the value of ID. This means that decisions about the possibility of a block loss can be made directly, rather than as a side product of analysis under a null hypothesis. Another advantage is that this procedure can be "tuned" in order to respond better to the type of loss which is expected. For instance, if a large block loss is feared, preparations can be made against it by increasing the mean of expected block loss (of course, this will have the effect of reducing sensitivity to smaller block losses). Or, if the amount of holdup is not well known, a very broad distribution can be chosen for that. This flexibility gives us the ability to respond to the situation as it is actually observed, whereas standard methods are incapable of adapting to changing circumstances. Proponents of standard methods sometimes view this inflexibility as an advantage, but we consider it preferable to be able to respond to the situation if it changes.

With standard methods, decisions must be made with a test of fixed significance; otherwise, the frequentist interpretation of the test does not hold. Since we are dealing with probability distributions, we are not limited to setting a critical threshold and a critical probability. In fact, our simulations indicate that it is best to take into account all the information given by the posterior probabilities. The results of a single hypothesis test, although a convenient summary, may be misleading. We encourage the user of our methods to examine the probabilities of multiple critical regions, something which is not possible with standard methods.

In sum, this procedure will be harder to set up than a standard approach to an SNM accounting procedure, because of the problems of assessing uncertainties and finding appropriate prior probabilities. We believe, however, that the effort involved in setting up an integrated analysis of both block and trickle loss is justified by the additional amount of information obtained about the operation of the process, information which can be brought to bear directly on the decision making process.

5. References

1. P. J. Harrison, C. F. Stevens, "Bayesian Forecasting, (with Discussion)," J. Roy.-Statist. Soc., Ser. B. 38, (1976), pp. 205-247.
2. B. Hoadley, "The Quality Measurement Plan (QMP)", The Bell System Technical Journal, Vol. 60, No. 2, (1981), pp. 215-273.
3. J. L. Jaech, "Statistical Methods in Nuclear Material Control", USAEC TID-26298, (1973).
4. D. V. Lindley, Making Decisions, Wiley Interscience, (1971).
5. D. H. Pike, G. W. Morrison, C. W. Holland, "Linear Filtering Applied to Safeguards of Nuclear Material", Trans. American Nuclear Society, Vol. 22, (1975), pp. 143-144.
6. D. H. Pike, G. W. Morrison, C. W. Holland, "A Comparison of Several Kalman Filter Models for Establishing MUF", Trans. American Nuclear Society, Vol. 23, (1976), pp. 267-268.
7. J. W. Tukey, Exploratory Data Analysis, Addison-Wesley Publishing Company (1977).

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government thereof, and shall not be used for advertising or product endorsement purposes.

Time Period	$p(N_B^i D)$
1	0.0025249
2	0.0018269
3	0.0003563
4	0.0000352
5	0.0000229
6	0.0000064
7	0.0000021
8	0.0001775

1/12/1

~~1/12/1~~

100
90
80
70
60
50
40
30
20
10
0

Figure 1: Time-series data showing a sharp peak in the top series around time 20, followed by a gradual decline. The other series show more gradual increases or remain flat.

100
100

7.3.2

Figure 1

Figure 1

Figure 1