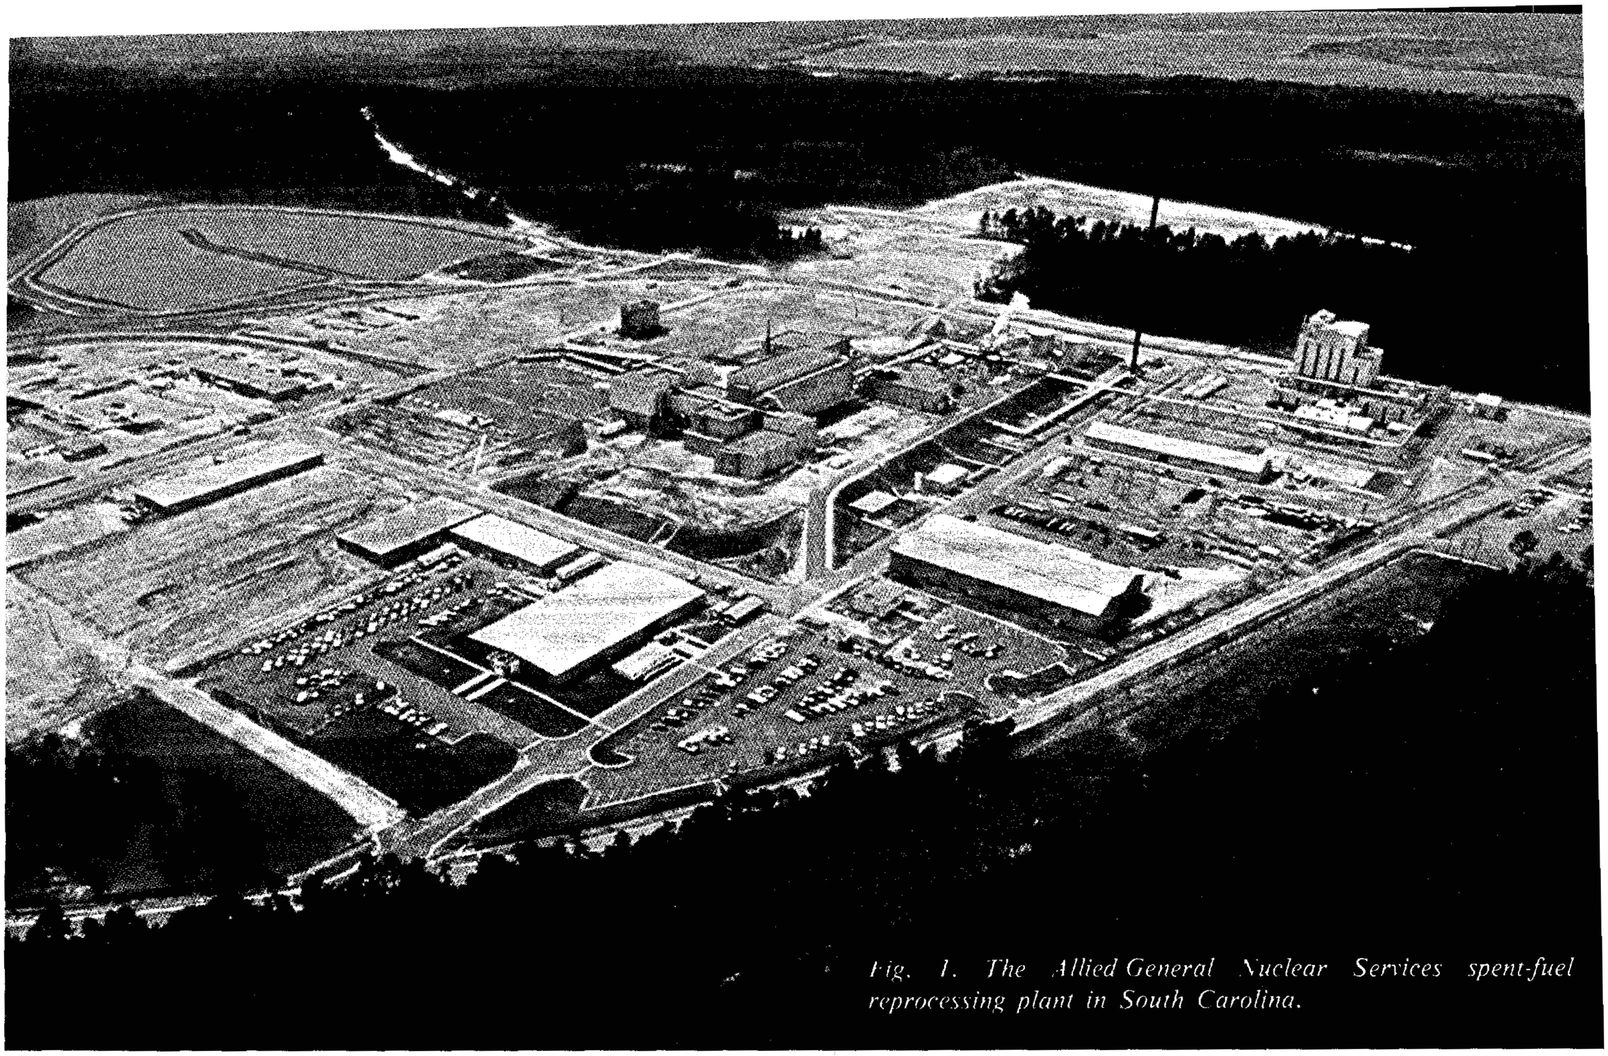

Fig. 1. The Allied General Nuclear Services spent-fuel reprocessing plant in South Carolina.

Dynamic Materials Accounting Systems

by Darryl B. Smith, Dante Stirpe, and James P. Shipley

The problem of maintaining strict control over nuclear material will be made more difficult by the nuclear power demands of the future, which will require large facilities—enrichment plants and reprocessing plants, for example—that process great quantities of high-quality fissile materials. The scale of these operations has forced a reassessment not only of facility design, construction, and process operation, but also of safeguards methods to prevent unauthorized use of the nuclear materials contained in the facilities. A comprehensive domestic safeguards system combines the functions of materials accounting and physical protection.

The Los Alamos Scientific Laboratory has been designated the Department of Energy's lead laboratory for the design and evaluation of materials accounting systems for nuclear facilities of the future. In this article, we examine these systems and the techniques for their design. Nuclear materials accounting systems must keep track of large quantities of materials as they

Systems analysis suggests that near-real-time materials accounting systems designed for future large-throughput nuclear facilities can meet high performance standards.

move through the various processing stages and must keep track of them so well that the absence of even small amounts can be detected. The uncertainties inherent in any measurement process and the difficulties of measuring in high-radiation fields behind heavy shielding complicate this task.

We also illustrate the potential benefits of these systems by describing the development and expected performance of a materials accounting system we have designed for the Allied-General Nuclear Services (AGNS) spent-fuel reprocessing plant at Barnwell, South Carolina (Fig. 1). This plant was designed to process large amounts of irradiated fuel from power reactors. The accounting system was designed after the plant was built and with simulated data because the plant is not yet operating.

The potential of system performance is based on projected measurement capabilities of instruments, some of which are still under development. These projections cannot be tested without access to an operating facility. However, our preliminary evaluations suggest that we can design dynamic materials accounting systems for large bulk-processing facilities that meet detection standards close to those recommended by the IAEA.

The Basis for Materials Accounting

The ultimate aim of nuclear safeguards is to be able to state with confidence, "No significant amount of nuclear material has been diverted." The philosophy underlying the development of materials accounting systems is that the truth of the statement can and should be verified. Thus, materials accounting systems are designed to account for or keep track of the amounts and locations of sensitive nuclear materials by periodic measurements. Materials balances are drawn about suitable areas of the facility according to the equation

$$\begin{aligned}\text{Materials balance} = {} & \text{initial inventory} \\ & + \text{transfers in} \\ & - \text{transfers out} \\ & - \text{final inventory}\end{aligned}$$

defined over a reasonable time interval. In principle, if all nuclear material in each term of the equation has been measured, the materials balance should be zero in the absence of diversion. In practice, however, it is never zero because of the uncertainties inherent in all measurement processes. The measurement uncertainties produce a corresponding uncertainty in the materials balance, so statistical techniques are used to decide whether a balance indicates diversion of material.

At present, materials balances are drawn around an entire plant or a major portion of it after the facility has been shut down and cleaned out to inventory the material present. Although such accounting methods are essential to safeguards control of nuclear materials, they have inherent limitations in sensitivity and timeliness. The sensitivity is limited by measurement uncertainties that may conceal losses of significant quantities of nuclear material in large plants. The timeliness is limited by the frequency of physical inventories; that is, the practical limits on how often a facility can be shut down for inventory and still remain productive.

Both sensitivity and timeliness can be improved by implementation of *dynamic materials accounting*. This approach combines conventional chemical analysis, weighing, and volume measurements with the on-line measurement capability of NDA (nondestructive assay) instrumentation to provide rapid and accurate assessment of the locations and amounts of nuclear material in a facility. Materials balances are drawn without shutting down the plant: in-process inventories are measured, or otherwise estimated, while the process is operating.

To implement the approach, the facility is partitioned into several discrete accounting areas. Each accounting area contains one or more chemical or physical processes and is chosen on the basis of process logic and the ability to draw a materials balance, rather than on geography, custodianship, or regulatory requirements. By measuring all material flows in each area separately, quantities of material much smaller than the total plant inventory can be controlled on a timely basis and any discrepancies can be localized to the portion of the process contained in the accounting area.

Control by dynamic materials accounting is rigorous. It forces a potential divertor to steal nuclear material in quantities small enough to be masked by measurement uncertainties. Thus, to obtain a significant quantity of material, the divertor must commit many thefts and run the concomitant high risk of detection by the accounting system, surveillance instruments, and physical protection system.

Designing a Materials Accounting System

The performance, or diversion detection sensitivity, of a materials accounting system depends on the details of the measurement system, which in turn depend on the details of the process. Because these details vary from one plant to another, the Los Alamos safeguards systems studies focus on specific designs of existing or planned nuclear facilities.

The first step in the development of a facility's accounting system is to determine the flows of nuclear materials through the facility from design data and operator experience. Then, the facility is partitioned into logical accounting areas, and an appropriate measurement system is postulated for each area. Wherever possible, the designer incorporates the measurement processes already in the plant design into the measurement system and augments them with any additional measurements necessary to draw a materials balance.* The final step is to examine the expected performance of the accounting system design.

To develop preliminary designs of materials accounting systems, we model and simulate the in-plant processes and measurement systems by computer because no large fuel-cycle plants are yet in operation. Detailed dynamic models of material flows are based on actual process design data. They include bulk flow rates, concentrations of nuclear materials, holdup of materials in the process line, and the variability of all these quantities. Design concepts for the accounting systems are evolved by identifying key measurement points and appropriate measurement techniques, comparing possible materials accounting strategies, developing and testing appropriate data-analysis algorithms, and quantitatively evaluating the proposed system's capability to detect losses. The use of modeling and simulation allows us to study the effects of process and measurement variations over long operating periods and for various operating modes in a short time.

Computer codes simulate the operation of each model process using standard Monte Carlo techniques. Input data include initial values for all process variables and values of statistical parameters that describe each independent process variable. These data are best estimates obtained from process designers and operators. Each accounting area is modeled separately. When a process event occurs in a particular area, the values of the flows and in-process inventories associated with that part of the process are computed and stored in a data matrix. These data are available for further processing and as input to computer codes that simulate accounting measurements and materials balances.

The flow and inventory quantities from a simulated process model are converted to measured values by applying simulated measurements. Each measurement type is modeled separately; measurement errors are assumed to be normally distributed (Gaussian), and provisions are made for both additive (absolute) and multiplicative (relative) errors. Significant measurement correlations are included explicitly. In most cases the measurement models are derived from the performance of similar instrumentation that has been used and characterized in laboratory and field applications involving similar materials. Simulated measurements are combined to form dynamic materials balances under various accounting strategies.

Data Analysis

We combine the most promising measurement and accounting strategies with statistical techniques in comparative studies of loss-detection sensitivities. One of the major functions of the materials accounting system is to indicate loss, or possible diversion. Diversion may occur in two basic patterns: abrupt diversion (the single theft of a relatively large amount of nuclear material) and protracted diversion (repeated thefts of nuclear material on a scale too small to be detected in a single materials balance because of measurement uncertainties). Protracted diversion usually is the most difficult to detect.

The use of dynamic materials accounting enhances the ability to detect both diversion patterns, but it results in the rapid accumulation of relatively large quantities of materials accounting data. For example, if an area's materials

---

*See "Nondestructive Assay for Nuclear Safeguards.

balance is closed once each 8-hour shift, after 1 month the safeguards operator will have a sequence of 84 materials balances and estimates of their associated uncertainties. Analysis of a single materials balance may be sufficient to detect a large abrupt theft of material, but the entire sequence of data contains the information necessary to detect small protracted diversions. Because small diversions may be masked by measurement uncertainties, they often are difficult to detect, and the operator must use one or more statistical tests of the accounting data to decide whether diversion has taken place.

Decision Analysis

We have developed or adapted a variety of statistical tools for the analysis of materials accounting data that become available sequentially in time. These tools and their implementation are known collectively as *decision analysis*.

A simple, specific example shows what is involved in decision analysis. Suppose we have a sequence of 10 materials balances (47, 2, −109, 76, 2, 40, 62, −20, 34, 18 g)* and an estimate of each balance's standard deviation. The standard deviation σ is a measure of the uncertainty in a materials balance calculated from individual measurement errors. To analyze these data we must select an appropriate *statistical test*, construct a *test statistic* from the materials balance data, and establish one or more *test thresholds*. Then we can compare the value of the test statistic to the threshold(s) and draw a conclusion as to whether material has been diverted.

*"I know there's an 'I.D.'*in here somewhere!"*

*Inventory Difference

In our example, we use the *cusum test*, a statistical test that uses the cumulative summation (cusum) of the materials balances as the test statistic. The cusum is used often because it provides an estimate of the total amount of material diverted during an accounting period. Other test statistics include a single materials balance or a weighted average of the materials balances. Our cusum test will have a single test threshold determined by the false-alarm probability—that is, the probability of concluding (because of measurement uncertainties) that nuclear material has been diverted when, in fact, no diversion has occurred. The false-alarm probability (FAP) is a measure of the significance of the test results. The FAP value used in setting up the test usually depends on the false-alarm rate that can be tolerated in the plant. The rate often depends on the consequences (shutting down the plant, perhaps) of incorrectly concluding that diversion has taken place.

Our cusum test is illustrated in Fig. 2. In the absence of diversion, we would expect the value of our test statistic—the cusum—to be zero. However, because of measurement uncertainties, the cusum value we get from our accounting data will almost never be zero. The curve in Fig. 2 represents the probability distribution of getting various cusum values when no diversion has occurred. The curve is centered at zero—the expected cusum value. The total area under the curve is 1 because the probability is 100% that the cusum has *some* value.

---

*The values are the result of a Monte Carlo simulation: they were obtained from a sequence of 10 normally distributed random numbers having mean zero and standard deviation 1 by using the relationship $MB = RN \times \sigma + D$, where MB is the materials balance, RN is a random number, and D is the diversion.

The width of the curve is determined by the uncertainty (measured by the standard deviation σ) in the cusum, which can be computed from the uncertainties in the individual materials balances. In our example, the standard deviation of the cusum is 100 g. The area under the curve to the left of any cusum value represents the probability that—in the absence of diversion—the cusum will have this value or less.

Now we must set our test threshold. We assume that a 5% false-alarm rate is acceptable. In this case, the test threshold (labeled Z in Fig. 2) is set at 165 g. The 5% of the area under the curve lying to the right of the threshold represents the false-alarm probability, labeled FAP.

We are finally ready to test our materials accounting data for evidence of diversion. The materials balances in our example sum to 152 g. Because this cusum value is less than our test threshold, we conclude that there has been no diversion.

Had our cusum value been greater than 165 g, we would have concluded that material had been diverted, but we recognize that there is a significant chance that this would be an incorrect conclusion. If there were no other considerations, the false-alarm probability could be reduced to any arbitrarily small number by increasing the value of the test threshold—but then what would happen to our ability to detect the diversion of a significant amount of material? The relationship between the test threshold and the detection probability is illustrated in Fig. 3.

Suppose 250 g of material have been diverted. Our probability curve is now centered at 250 g because this is the expected cusum value under the hypothesis that this amount of material has been diverted. The width of the curve has not changed because our cusum still has the same associated uncertainty, and the test threshold has not moved. The area un-

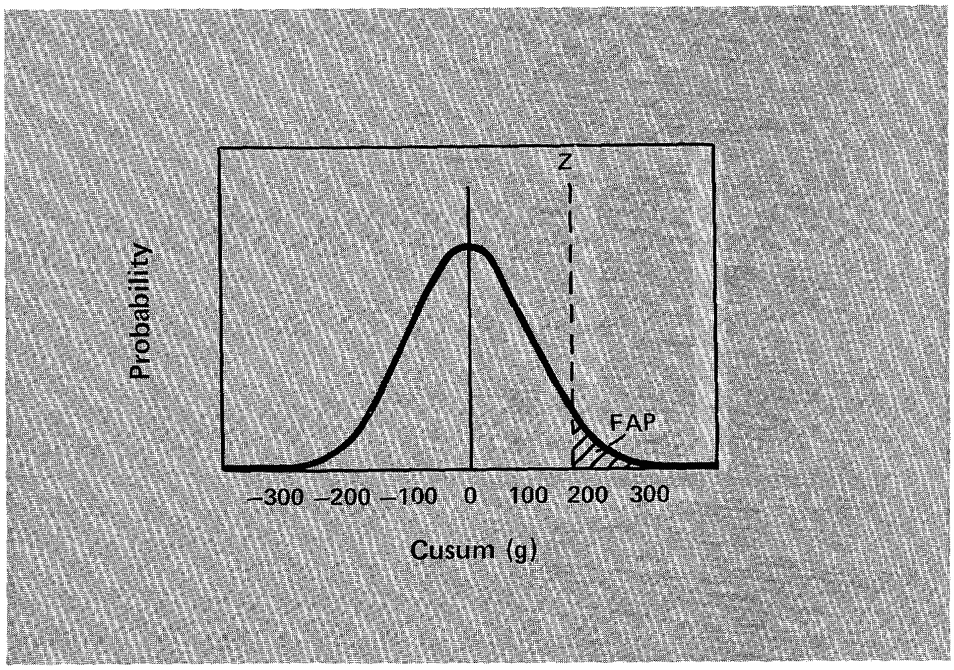

Fig. 2. A cusum test having a single test threshold, Z = 165 g. The value of Z is determined by the false-alarm probability (FAP) desired. The FAP is the area under the curve to the right of Z, and in the example, this area is 0.05.

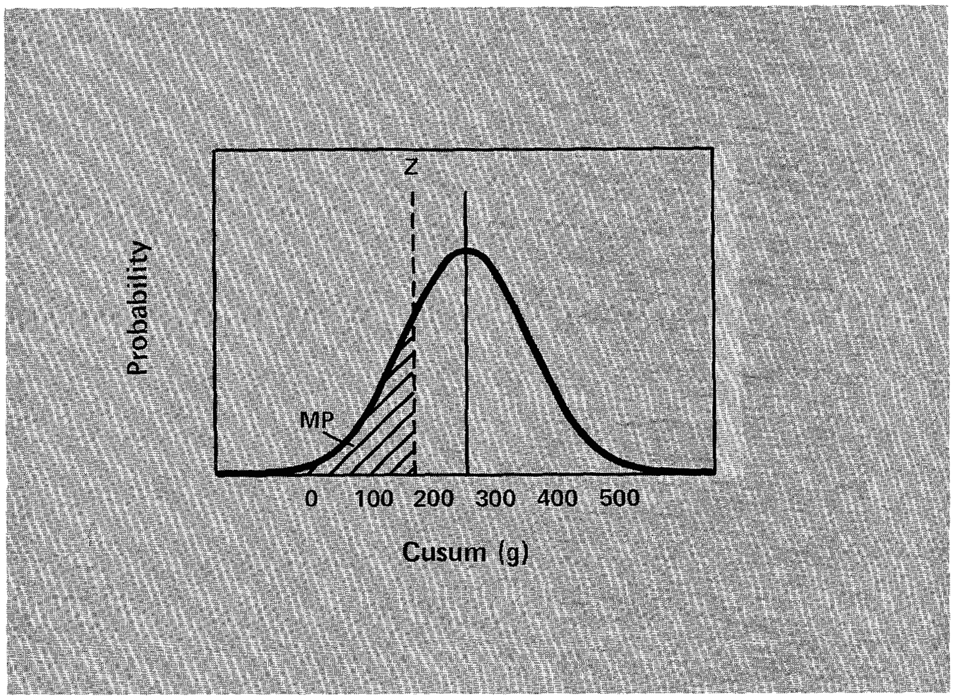

Fig. 3. A cusum test with a total diversion of 250 g and a test threshold kept at Z = 165 g. The area under the curve to the right of Z is the detection probability, which, in this illustration, is 0.80. The miss probability (MP) is 1 − the detection probability.

der the curve to the right of the threshold represents the probability of detecting the diversion of 250 g of material. (In Fig. 3 this area is 0.8.) The shaded area on the left of the test boundary is the miss probability (MP). It is the probability of concluding that there has been no diversion when, in fact, 250 g have been diverted; it is equal to 1 – the detection probability.

In our example, the cusum is 152 g. This value is smaller than the test boundary, and we have already concluded that there was no evidence of diversion in our materials accounting data. However, if material has been diverted (as illustrated in Fig. 3), we have failed to detect this fact. With our test boundary set for a 5% false-alarm probability, we have only an 80% probability of detecting the diversion of 250 g of material. As the amount of material diverted increases, so does our ability to detect the diversion.

In our example, we considered a single cusum test of 10 materials balances and found that we could choose our test threshold based on an acceptable false-alarm rate. In practice, the 10 balances in our cusum would have been accumulated over a period of time: 10 weeks, if a materials balance were drawn at the end of each week. Perhaps we would like to test each materials balance as it becomes available or test the current cusum as each new materials balance is added. However, if we test each cusum, and if the false-alarm and miss probabilities are fixed for each test, the overall false-alarm and miss probabilities become unacceptably large after several such tests.

Sequential Tests

Another kind of test, the sequential probability ratio test (SPRT) is particularly suited for analyzing data that become available sequentially. The SPRT allows us to guarantee that neither the false-alarm probability nor the miss probability will exceed desired values, no matter how long the sequence. The cusum remains an appropriate test statistic.

A sequential test has an upper and a lower threshold. Thus, at any time the test result may be that no diversion has occurred, that diversion has occurred, or that no decision can be made until more data are available. Both test thresholds depend not only on the false-alarm probability but also on the desired detection probability, the average rate of diversion, and the number of materials balances in the cusum.

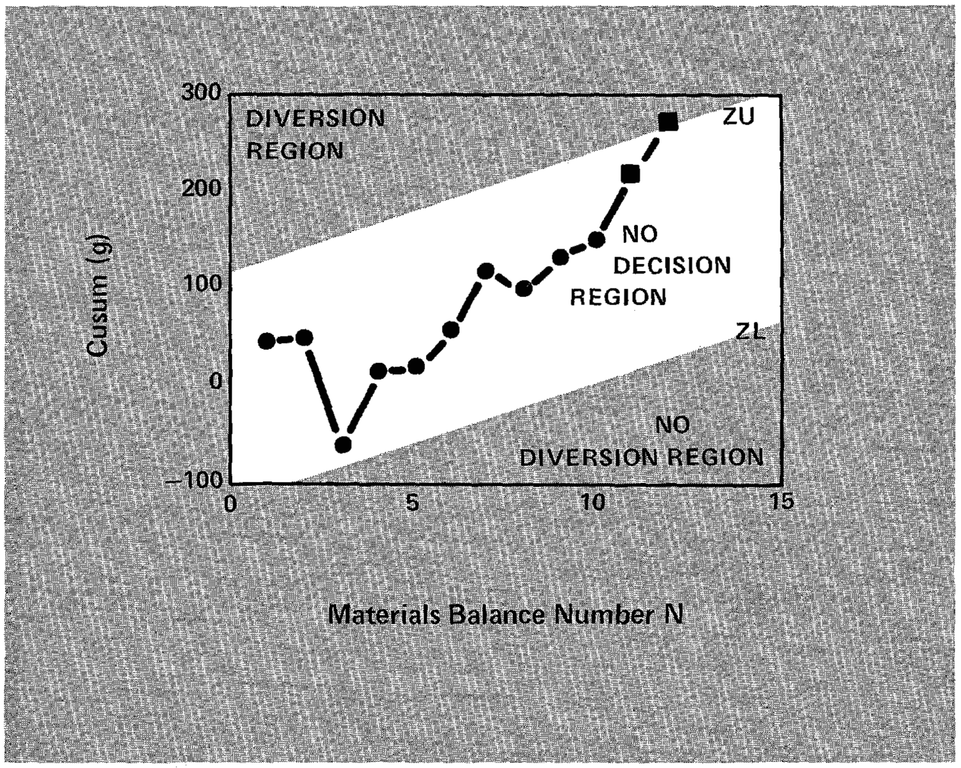

Fig. 4. An example of the Sequential Probability Ratio Test with diversion detected at the 12th materials balance. The test thresholds are established to detect a diversion of 25 g of material during each materials balance period with a detection probability of 95% and a false-alarm probability of 5%. Because the test allows for a no-decision region, an incorrect decision was not made.

A typical SPRT is illustrated by Fig. 4 for our example sequence of materials balance data. In this test, as each new materials balance is drawn, it is added to the previous cusum to obtain a new cusum value; the value is plotted against the materials balance number. The upper and lower test thresholds are the two parallel lines labeled ZU and ZL, respectively, which divide the cusum chart into three regions indicating diversion, no decision, and no diversion. If the current cusum value falls above ZU, we conclude that diversion has taken place. If it falls below ZL, there is no evidence of diversion. If it lies between ZU and ZL, we wait for the next materials balance to be drawn. The thresholds have a positive slope because, if a pattern of protracted diversion is present, the total amount of material diverted increases as the number of materials balances in the cusum increases.

The thresholds in Fig. 4 were set for 5% false-alarm probabilities and 5% miss probabilities and for an average 25-g rate of diversion. The settings mean that we would like to detect, with at least 95% probability, the removal of 25 g of material during each balance period and that we can tolerate a false-alarm rate no greater than 5%. The circular symbols correspond to the 10 cusums computed from our example data sequence. The 10th cusum (N = 10) lies in the region between the test thresholds so, at the time the 10th materials balance was drawn, we were unable to make a decision. Earlier, we saw that the single cusum test applied after this materials balance resulted in the conclusion that no diversion had occurred. Indeed, a similar test applied to each of the nine previous cusums would have resulted in the same conclusion. Such conclusions are incorrect for the simulated process from which our example sequence of materials balances was taken: 25 g had been diverted from each materials balance. On the other hand, the SPRT still has not permitted a decision after the 10th materials balance, and thus no incorrect decision has been made. However, after two additional materials balances are drawn, the cusum exceeds the upper test threshold, resulting in the (correct) conclusion that material has been diverted. The current ($N = 12$) cusum value of 271 g provides an estimate of the total amount diverted: the true quantity was 300 g.

Test Statistics

A variety of test statistics can be formed from the materials accounting data and tested sequentially for indications of diversion. Each statistic is based on a different assumption concerning the state of prior knowledge of the measurement errors and of the diversion strategy. Three of the most useful test statistics are the Shewhart, cusum, and uniform diversion statistics.

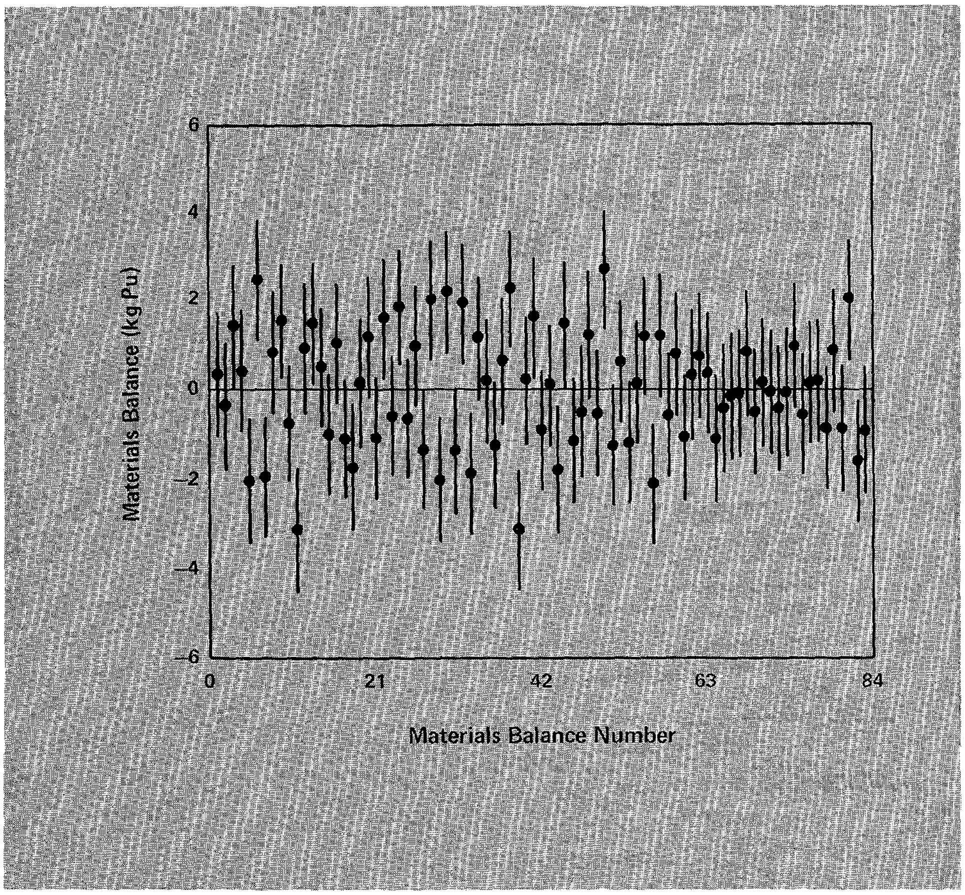

Fig. 5. The Shewhart chart is a graph of sequential materials balance values and their respective materials balance numbers. For each materals balance number, the short, horizontal line gives the materials balance value, and the vertical lines above and below represent the ±1-σ deviations from this value. This chart is rather insensitive to protracted, low-level material diversions, but is sensitive to large, abrupt diversions.

The Shewhart chart (Fig. 5) is the oldest graphical-display tool to be used widely by industry for process control. In the standard form, measured data are plotted sequentially on a chart where 2-σ and 3-σ levels are indicated. In safeguards applications, the Shewhart chart is a sequential plot of the materials balance data with 1-σ error bars. This chart is most sensitive to large, abrupt shifts in the materials balance data.

The cusum statistic is computed after each materials balance period. It is the sum of all materials balances since the beginning of the accounting interval. Cusum charts are sequentially plotted cusum values that are used to indicate small shifts in the materials balance data (Fig. 6). The cusum variance (σ_c^2) is a complex combination of the variances of individual materials balances, because these balances usually are not independent. Correlation between materials balances has two principal sources. The first source is the correlation between measurement results obtained by using a common instrument calibration. The magnitudes of the associated covariance terms depend on the magnitude of the calibration error and the frequency of each instrument recalibration; omission of these terms can cause gross underestimation of the cusum variance. The second source is the occurrence, with opposite signs, of each measured value of in-process inventory in two adjacent materials balances. As a result, only the first and last measurements of in-process inventory appear in the cusum, and only the corresponding variances appear in the cusum variance.

The Kalman filter is a statistical technique applied widely to communications and control systems for signal processing. It is a powerful tool for extracting weak signals embedded in noise.

It has been applied recently to safeguards because dynamic materials accounting systems rapidly generate large quantities of data that may contain weak signals, caused by repeated, small diversions, embedded in the noise produced by measurement errors.

The uniform diversion test (UDT) is designed to detect a small, constant diversion during each materials balance period. Estimates of the average diversion and the inventory at each time are obtained using the Kalman filter. A chart of the UDT is shown in Fig. 7.

The cusum and the UDT are complementary in several respects. The cusum estimates the *total* amount of missing nuclear material at each time step, and its standard deviation is the 1-σ error in the estimate of the total. The UDT, on the other hand, estimates the *average* amount of nuclear material missing from each materials balance, and its standard deviation estimate is the 1-σ error in the estimate of the average. Thus, both the cusum and the UDT search for a persistent, positive shift of the materials balance data—the cusum by estimating the total and the UDT by estimating the average.

Alarm-Sequence Charts

The decision tests examine all possible sequences of the available materials balance data because, in practice, the time at which a sequence of diversions begins is never known beforehand. Furthermore, to ensure uniform application and interpretation, each test is performed at several levels of significance (false-alarm probability). Thus, it is useful to have a graphic display that indicates the alarm-causing sequences, specifying each by its length, time of occurrence, and significance. One such tool is the alarm-sequence chart, which has proven useful in summarizing the results of the various tests and in identifying trends in the materials accounting data.

Fig. 6. The cusum chart is a graph of the sums of all materials balances drawn from the beginning of an accounting period versus the number of materials balances in the cusum. In this chart, the short, horizontal lines give these cusum values, and the vertical lines represent ±1-σ deviations from these values. Because the chart is relatively sensitive to small shifts in materials balance data, it is useful for the detection of protracted, low-level diversion.

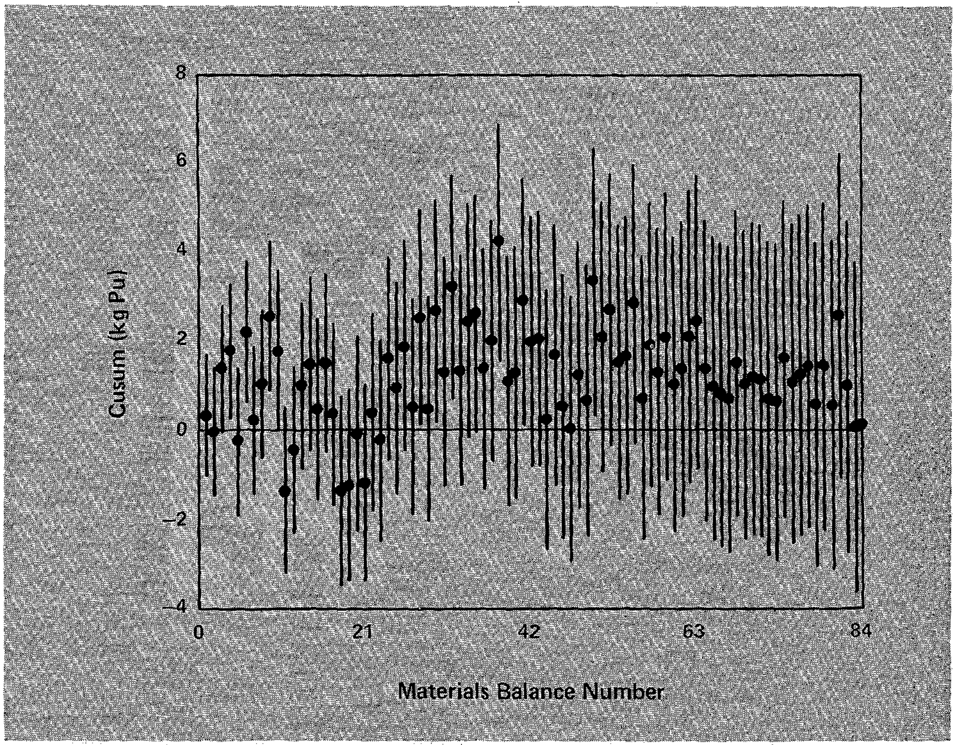

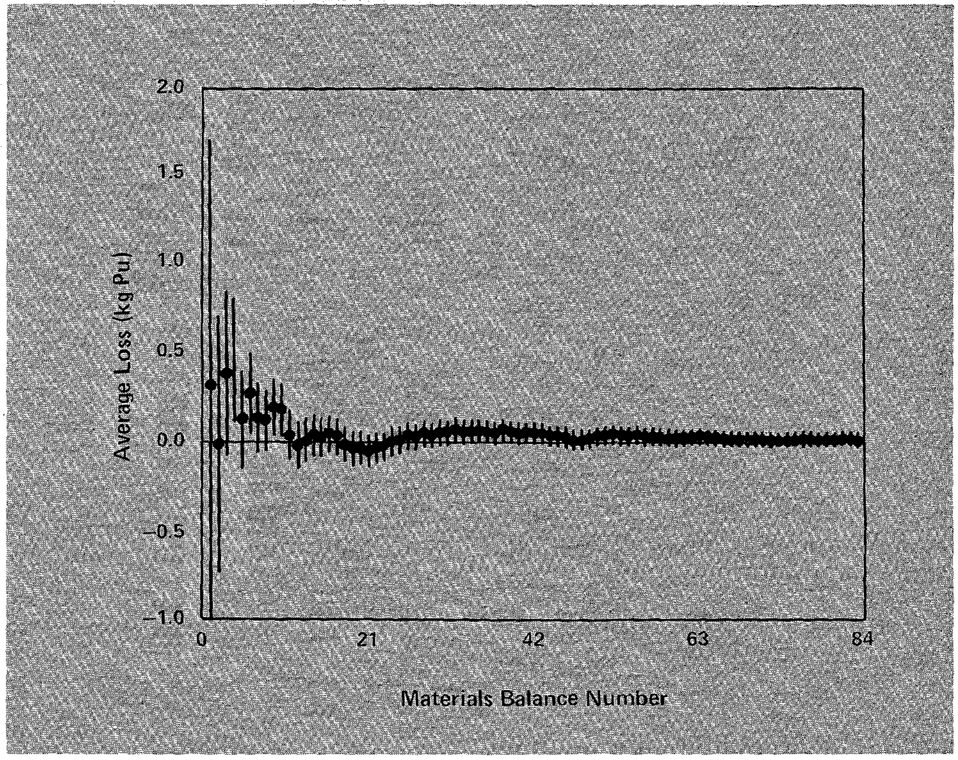

Fig. 7. In this chart of the uniform diversion test, each point on the graph is obtained by a linear combination of previous materials balances. The combination is constructed to provide an estimate of the average amount of material missing per materials balance. Like the cusum test, the UDT searches for a persistent, positive shift in the materials balance data.

An alarm-sequence chart is shown in Fig. 8.

To generate the alarm-sequence chart, each sequence in which the test statistic exceeds the upper boundary ZU and causes an alarm is assigned both a descriptor that classifies the alarm according to its significance (false-alarm probability) and a pair of integers (r_1,r_2) that are, respectively, the indexes of the final and initial materials balances in the sequence. The alarm-sequence chart is a point plot of r_2 vs r_1 for each sequence that caused an alarm, with the significance range of each point indicated by the plotting symbol. One possible correspondence of plotting symbol to significance is given in Table I. The symbol T denotes sequences of such low significance (high false-alarm probabiltiy) that it would be fruitless to examine their extensions; the position of the symbol T on the chart indicates the termination point.

TABLE I
Alarm Classification for the Alarm-Sequence Chart

Classification (Plotting Symbol)	False-Alarm Probability
A	10^{-2} to 5×10^{-3}
B	5×10^{-3} to 10^{-3}
C	10^{-3} to 5×10^{-4}
D	5×10^{-4} to 10^{-4}
E	10^{-4} to 10^{-5}
F	$\leq 10^{-5}$
T	>0.5

It is always true that $r_1 \gg r_2$, so that all symbols lie to the right of the line $r_2 = r_1$ through the origin. Persistent data trends (repeated diversions) cause long alarm sequences ($r_1 \gg r_2$), and the associated symbols on the alarm chart extend far to the right of the line $r_2 = r_1$.

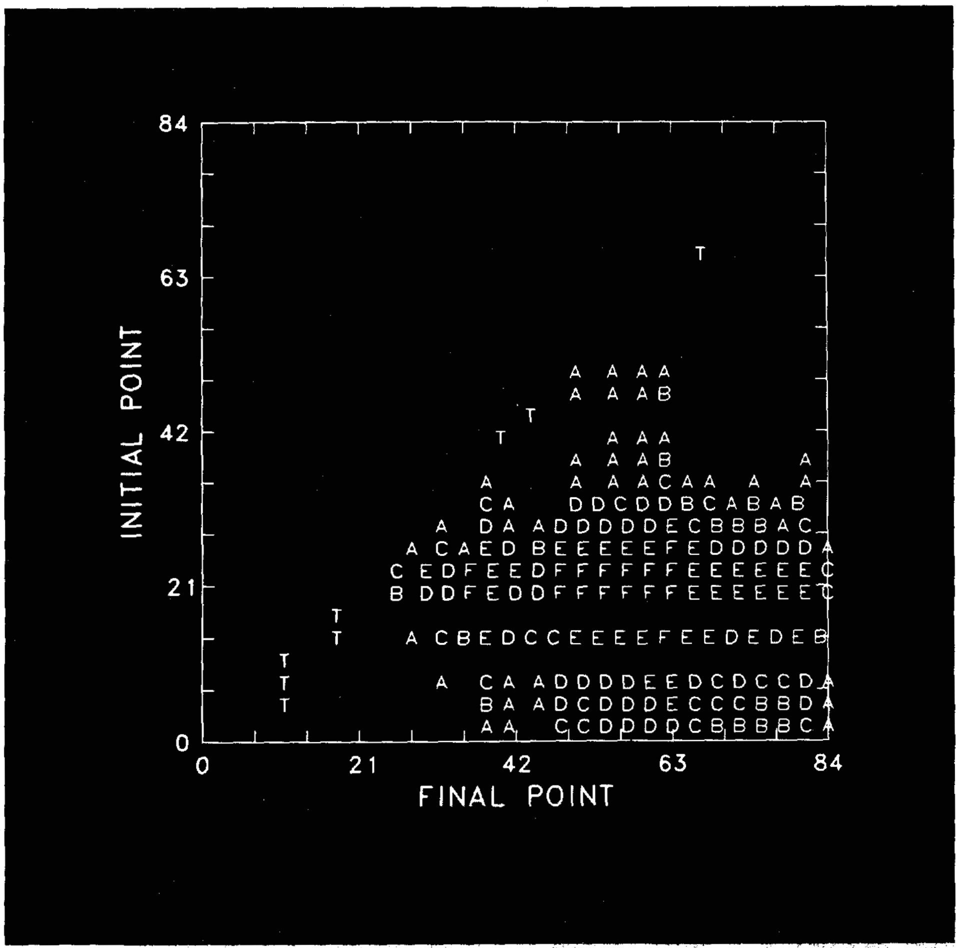

Fig. 8. An alarm-sequence chart. The false-alarm probability associated with each letter is given in Table I. To illustrate how this chart is used, consider a sequence of materials balance data beginning at balance number 21, and suppose that one of the tests gives an alarm with a false-alarm probability of 2×10^{-4} at balance number 30. On the alarm-sequence chart for that test, the letter D would appear at the point (30,21). Because this false-alarm probability is so small, the probability of material diversion commencing with balance number 30 is large.

Systems Performance Evaluation

An analysis of a system's performance in detecting losses of nuclear material is essential to the design of nuclear materials accounting systems. Performance measures must include the concepts of loss-detection sensitivity and loss-detection time. Because materials accounting is statistical, loss-detection sensitivity is described in terms of the probability of detecting some amount of loss while accepting the probability of some false alarms. Loss-detection time is the time required by the accounting system to reach a specified level of loss-detection sensitivity. The loss scenario is not specified in performance measures; whether the loss is abrupt or protracted, the total loss is the measure of performance. The loss-detection time refers only to the accounting system's internal response time.

The performance of any accounting system can be described by a function

$$P\ [L,N,\alpha]\ ,$$

where P is the accounting system's probability of loss detection, L is the total loss over a period of N balances, and α is the false-alarm probability. A convenient way of displaying system performance is a three-dimensional graph of the surface P vs L and N for a specified value of α, called a *performance surface.* A single point (N,L,P) of such a surface is plotted in Fig. 9. The entire surface portrays the expected performance of an accounting system as a function of the three performance measures, loss, time, and detection probability.