JRC 85583

ISSN 1977-5296

ESARDA Bulletin

NUMBER 19
MAY 1991

News about ESARDA

The ESARDA Bulletin welcomes the new members and heartily thanks the previous members who have contributed to the success of ESARDA. The names of the members of the different committees are reported in the table on the next page.

This table replaces that published in the ESARDA Bulletin No. 16 with the members on 1st June 1989. Please note that there are several changes, especially in the composition of the Steering Committee.

14th ESARDA Meeting 1992

Lanzarote, Spain 12-15 May 1992

The ESARDA Meeting 1992 will be a **restricted meeting** of ESARDA members, including the Steering Committee, the Coordinators and the members of the ESARDA Working Groups (WGs). In addition a number of experts will be invited to each of the two main items of the Meeting, i.e. the following two technical workshops:

- NDA Techniques Applicable to Safeguarding Nuclear Material in Waste
- C/S Safeguards Techniques Applicable to Intermediate and Long Term Storage of Irradiated Fuel.

These two workshops will extend from Tuesday to Thursday (12-14 May 1992) and will include plenary and subgroup sessions. In the last day (Friday 15 May 1992) the conclusions will be presented to the Steering Committee in an extended plenary session with the participation of all the WGs.

The meeting of the following WGs

- from Tuesday to Thursday: DA
- from Wednesday to Thursday: LEU, MOX, RIV

will be held at the same time of the workshops.

A normal meeting of the Steering Committee will also be organized on Thursday.

The extended plenary session of Friday 15 May 1992 will be attended by the Steering Committee, the ESARDA Coordinators, all the WG members and all the workshop participants (NDA WG and C/S WG members plus the invited experts).

This 14th ESARDA Meeting was planned to be held in Spain and a very attractive place was found. It will be organized in the Canary Islands, where the availability of rooms fulfils our needs completely at a price which is completely acceptable. This is the Hotel Teguise Playa on the Island of Lanzarote.

In addition it should also be considered that the participants will find this place convenient from the financial point of view. Inclusive tours (including travel, transport to the hotel and one week of half-board pension) can be organized from anywhere in Western Europe at a price which in May is definitely less than a normal air ticket to Madrid.

This volcanic island is exceptionally interesting because of the landscapes of the sea, the mountain and the various volcanic areas with the different stages of vegetation depending on the different eras of the volcanic eruptions.

This island escapes all definition, because of the multitude of sensations caused by its untouched beauty. It seems like a primitive world, almost lunar, with strong contrasts of colour, where the cultivated lava fields alternate with plots of earth with subtropical vegetation and the lava rocks contrast with the white of the beach. A fantastic view greets the camel-mounted tourist who climbs the picks of Timanfaya, gigantic "Mountain the Fire". Here the tourism has been developed in complete respect of the countryside allowing the island's integrity and special nature to be maintained. Excursions can be organized at moderate cost.

The 14th ESARDA meeting is the fourth in the series of restricted meetings which in these years alternate with the full ESARDA Symposia with open participation. The preceding restricted meetings were held in 1986 in Copenhagen, in 1988 in Karlsruhe (KfK) and in 1990 in Como. In 1993 a general ESARDA Symposium will be held in a place to be decided.

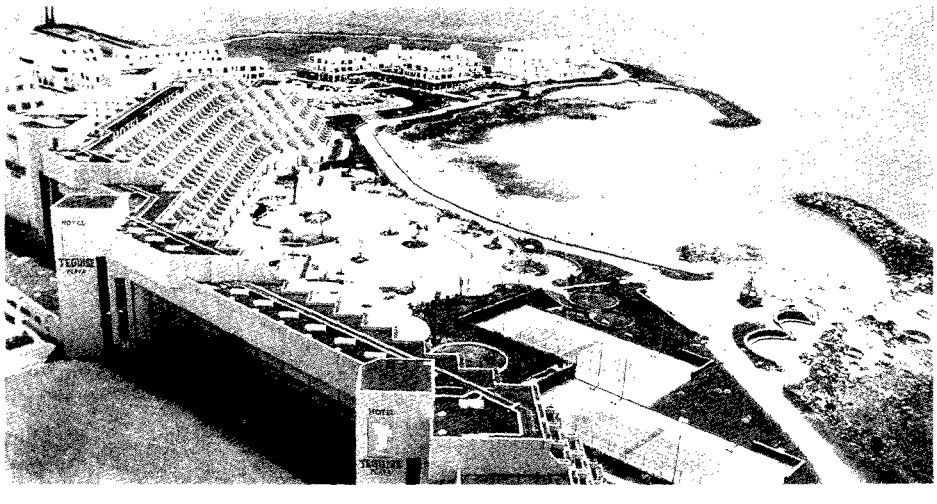

Who's Who in ESARDA?

(as of 21st May 1991)

Chairman 1991 *B.H. Patrick*, AEA Technology, Harwell, U.K.
Appointed chairman 1992
To be appointed from Denmark
Secretary *C. Foggi*, CEC, JRC-Ispra, Italy
Permanent Symposium
Scientific Secretary *L. Stanchi*, CEC, JRC-Ispra, Italy

ESARDA Steering Committee

G. Andrew, Dept. of Energy, U.K.
C. Beets, Ministry of Foreign Affairs, Belgium
C.P. Behrens, Kernkraftwerk Philippsburg, F.R. Germany (Observer)
M.J.C. Charrault, CEC Brussels, Belgium
M. Cuypers, CEC, JRC-Ispra, Italy
G. Déan, CEA Fontenay-axu-Roses, France
P.P. De Regge, CEN/SCK Mol, Belgium
S. Finzi, CEC Brussels, Belgium
C. Fizzotti, ENEA Casaccia, Italy
P. Frederiksen, Risoe, Denmark
A. Gloaguen, EDF, France
W. Gmelin, CEC, Safeguards Directorate, Luxembourg
R. Howsley, BNFL Risley, U.K.
F. Maccazzola, ENEA DISP, Italy
R. Kroebel, KfK, Karlsruhe, F.R. Germany
J.M. Leblanc, Belgonucléaire, Belgium
G. Le Goff, CEA Paris, France
T.T. Nielsen, Ministry of Energy, Denmark
B.H. Patrick, AEA Technology, Harwell, U.K.
F. Pozzi, ENEA Saluggia, Italy
J. Regnier, COGEMA, France
H. Remagen, BMFT, F.R. Germany
J. Sánchez, Ministerio de Industria, Comercio y Turismo, Spain
G. Stein, KFA Jülich, F.R. Germany
A. Velilla, CIEMAT, Spain
A.M. Versteegh, ECN Petten, Netherlands

ESARDA Board

P.P. De Regge, CEN/SCK Mol, Belgium
G. Déan, CEA Fontenay-aux-Roses, France
S. Finzi, CEC Brussels, Belgium
C. Fizzotti, ENEA Casaccia, Italy
C. Foggi, CEC, JRC-Ispra, Italy
W. Klöckner, CEC, Safeguards Directorate, Luxembourg (Observer)
R. Kroebel, KfK Karlsruhe, F.R. Germany
T.T. Nielsen, Ministry of Energy, Denmark
B.H. Patrick, AEA Technology, Harwell, U.K.
J. Sánchez, Ministerio de Industria, Comercio y Turismo, Spain
A.M. Versteegh, ECN Petten, Netherlands

ESARDA Coordinators

W.G. Bahm, KfK, Karlsruhe, F.R. Germany
R. Carchon, CEN/SCK Mol, Belgium
M. Cuypers, CEC, JRC-Ispra, Italy
M. Dionisi, ENEA Casaccia, Italy
P. Frederiksen, Risoe, Denmark
Mrs. F. García, CIEMAT, Spain
R.J.S. Harry, ECN Petten, Netherlands
T.L. Jones, AEA Technology, Dounreay, U.K.
Miss M. Neuilly, CEA Cadarache, France
R. Schenkel, CEC, Safeguards Directorate, Luxembourg

Working Group Convenors

Techniques and Standards for Non-Destructive Analysis (NDA)
S. Guardini, CEC, JRC-Ispra, Italy
Techniques and Standards for Destructive Analysis (DA)
P. De Bièvre, CEC, JRC-Geel, Belgium
Reprocessing Input Verification (RIV)
T.L. Jones, AEA Technology, Dounreay, U.K.
Containment and Surveillance (C/S)
B. Richter, KFA Jülich, F.R. Germany
Low-Enriched Uranium Conversion/Fabrication Plants (LEU)
P.P.A. Boermans, FBFC, Belgium
MOX Fuel Fabrication Plants (MOX)
G. Le Goff, CEA Paris, France

ESARDA Bulletin Editor

L. Stanchi, CEC, JRC-Ispra, Italy

Near Real Time Materials Accountancy How to Begin to Resolve Anomalies

Barry J. Jones

British Nuclear Fuels plc
Risley Warrington UK WA3 6AS

Abstract

To date, work has been aimed towards developing the best possible sequential testing procedure for detecting anomalies in streams of materials accountancy data. When such a test procedure is used as part of a materials control system, it will be the effect which will be observed when the test signals an alarm and the operator will want to know the possible causes. Therefore, before sequential testing procedures can be used for practical purposes of materials control, further development is necessary.

Suppose a test is designed for a particular plant. When a single stream of materials accountancy data is processed and subjected to the near real time materials accountancy test procedure, there may come a time when an anomaly is signalled. There is no easy way of attributing such an effect to a specific cause. The anomaly could have resulted from a number of possible causes or it could be just a false alarm. The problem addressed in the paper is how to begin to resolve such questions.

Introduction

An earlier paper[1] described how a near real time materials accountancy (NRTMA) system can be broken down into a number of modules, each with a clearly defined function. Some modules will need to be custom-designed for the particular plant concerned because of, for example, the need to access the plant data-base. Other modules are versatile and, therefore, capable of being installed as part of any NRTMA system. This paper concentrates on two such modules: - for statistical analysis, and for anomaly resolution.

The literature contains many reports of sequential tests which have been developed for the evaluation of streams of materials accountancy data. Only one procedure, which uses SITMUF[2] and the joint Page's test (the Joint Test), has emerged which is versatile in detecting abrupt and protracted losses, and which has a good overall response. In a series of earlier publications[3-6], it has been clearly demonstrated that the Joint Test is superior to all the other tests which have been published.

Suppose a Joint Test is designed for a particular plant. When a single stream of materials accountancy data is processed and subjected to the Joint Test, there may come a time when an anomaly is signalled. This means that, allowing for the false alarm probability which has been specified when setting up the Joint Test, the data stream appears to be inconsistent with there being no loss of material.

There is no easy way of attributing the alarm from the Joint Test to a specific cause. The anomaly could have resulted from a number of assignable causes or it could be just a false alarm. What is the cause? Any test procedure is insufficient on its own; a follow-up anomaly resolution procedure is essential.

Once the Joint Test has signalled an alarm, subsequent investigation is necessary to try and explain why the alarm occurred. The purpose of the anomaly resolution procedure is to provide evidence of the size and duration of the irregularity which caused the alarm, rather than to indicate the physical form or location of the material involved. Such evidence cannot, in itself, prove the cause of the alarm but rather suggest the most profitable courses of follow-up investigations.

What is Anomaly Resolution?

Anomaly resolution presents another instance of the most common problem that has to be tackled by the applied scientist. A set of data has been observed, and there is a wish to model mathematically the process which might have produced it. The usual approach would be to suggest a plausible form of model for the process and for producing the data, and then, by fitting that model to the data, the model's parameters would be estimated. If the fit is good, the model would probably be accepted. If the fit is poor, there are two possibilities:- either there is a great deal of random variability or measurement error in the data, or the choice of model form is inappropriate. Should the latter be the case, the procedure might be tried several times with different plausible models.

For example, when faced with a series of data on the rate of decay of a radioactive isotope, an exponential curve would probably be proposed as the model, and then the best curve fitted using the least squares method. If the fit is good, the half-life of the isotope could then be calculated from the fitted curve. When fitting calibration data, a straight line might be tried first, and higher order polynomials moved onto if the simple model gave a poor fit. In this case, care must be taken not to go too far. It is always possible to postulate a model which will give a perfect fit; a polynomial of order n will fit exactly through n + 1 points.

In the case of anomaly resolution, the operator, faced with the balances from a number of balance periods and wanting to postulate a model to fit these data, encounters two problems. The first problem is deciding what kind of model to propose. An obvious choice is that there may have been an abrupt loss in period 1, or in period 2, and so on. Another possibility, often considered in the literature, is that a uniform loss has begun in period x and ended in period y. The operator can attempt, in principle, to fit these models which represent events which are possible. In practice, anomalies in the data are more likely to occur because of a gross inventory mismeasurement which will show up as a gain (or loss) in one period followed by a loss (or gain) in the next, or because of a throughput determination having been assigned to the wrong balance period. The second problem is that, even if the choice of model is correct, a series of independent observations is needed to fit the model to the data. The MUF series is not made up of independent items. The SITMUF series does consist of independent items, all with the same variance. Therefore, if plausible models can be thought of for SITMUF, they can be fitted using standard least squares techniques. In other words, there is a need to examine those kinds of SITMUF patterns which can appear under the common scenarios discussed above.

The Plant Model

SITMUF patterns depend on various plant operating parameters, and on materials measurement uncertainties, and will therefore be plant dependent. Fortunately, this does not mean extra work for the operator since the software used for the Joint Test will also calculate these patterns.

Previous work[3-6] has chosen a campaign length of 240 days, divided into 40 balance periods of 6 days. The standard deviation of the throughput measurement error per balance period, T, set at 1 kg gives a standard deviation of the campaign throughput measurement error of 6.325 kg. This, and the standard deviation of the inventory measurement error, I, of 2 kg, is consistent with predictions for the THORP materials accountancy and control system. This plant model will be used throughout the paper.

Characteristic SITMUF Patterns From Various Perturbations

Expected SITMUF Patterns

For any plant which is suffering no loss of material, the expected value of the SITMUF statistic, each period, is zero. Once a loss occurs, the expected SITMUF values are no longer zero. For the

chosen plant model, Figures 1, 2 and 3 illustrate the expected behaviour of the SITMUF statistic in response to an abrupt loss of 10 kg in period 4, in response to a protracted loss of 10 kg spread uniformly over periods 2 to 4, and in response to an inventory measurement error of 10 kg which results in an increased MUF for period 4, followed by a correspondingly reduced MUF for period 5.

Figures 1, 2 and 3 are quite different from one another, and reflect the different types of loss or error. If there is a characteristic pattern in the SITMUF values corresponding to a given loss or error scenario, it seems worthwhile to examine SITMUF values in an attempt to work backwards to the losses or errors which caused them.

Application of the Joint Test

The anomaly resolution procedure is invoked once an alarm has been signalled by the chosen sequential test. The Joint Test is made up of two components, each of which is a Page's test. For these components the two test statistics, S1 and S2, are defined by

$S1_0 = 0$
$S2_0 = 0$
$S1_i = \max(0, S1_{i-1} + Y_i - K1) \qquad i > 0$
$S2_i = \max(0, S2_{i-1} + Y_i - K2) \qquad i > 0$

where $Y_1, Y_2, .., Y_i$ is the series of SITMUF values generated from $MUF_1, MUF_2, ..., MUF_i$. The Joint Test is applied such that an alarm is given if $S1_i > H_1$. Otherwise, no alarm is given unless $S2_i \geq H_2$. For the plant model described above, and a campaign false alarm probability of 5%, it has already been shown[6] that a suitable Joint Test has the following parameters:-

H1 = 0, K1 = 3.4758, H2 = 7.8, K2 = 0.24389

Resolution of Some Anomalies

Three worked examples are now given to demonstrate how the anomaly resolution procedure is applied. In each example, the Joint Test alarms in period 4, but the anomaly resolution procedure leads to quite different conclusions for each example.

Example 1

Table I - Observed SITMUF Values (Example 1)

Period	Observed SITMUF Value
1	0.576
2	1.587
3	2.864
4	3.605
5	1.956

The Joint Test alarms in period 4. Assume, in the first instance, that there has been an abrupt loss in period 4. It has already been pointed out that, for every loss scenario, there is an expected SITMUF sequence. Table II shows, for two sizes of an abrupt loss in period 4, the expected SITMUF values for the first 5 periods.

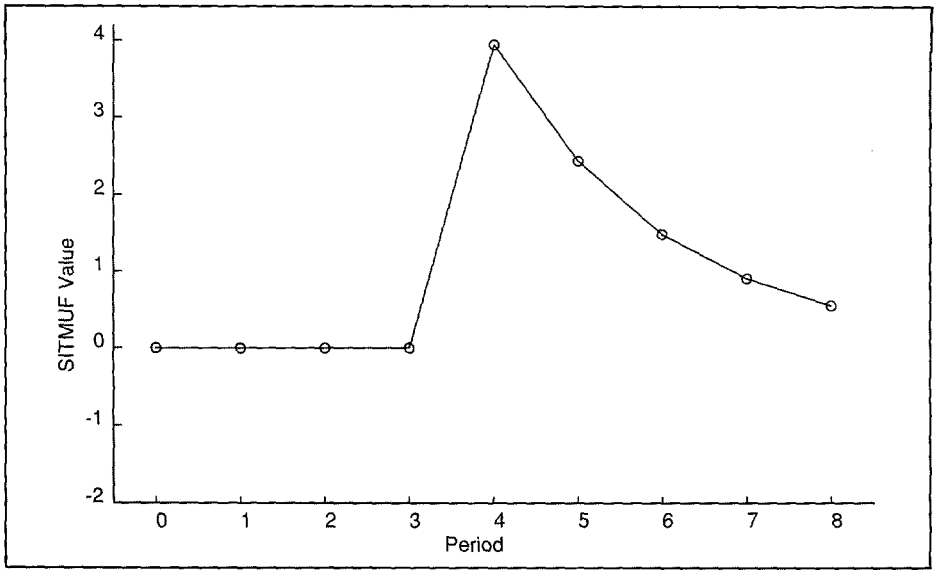

FIG. 1. *Expected SITMUF Values for an Abrupt Loss*

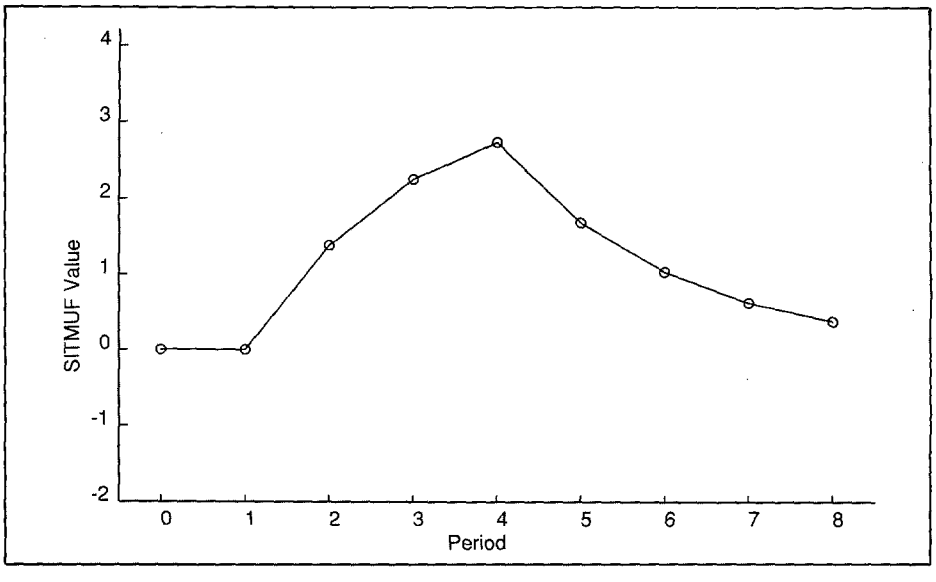

FIG. 2. *Expected SITMUF Values for a Protracted Loss*

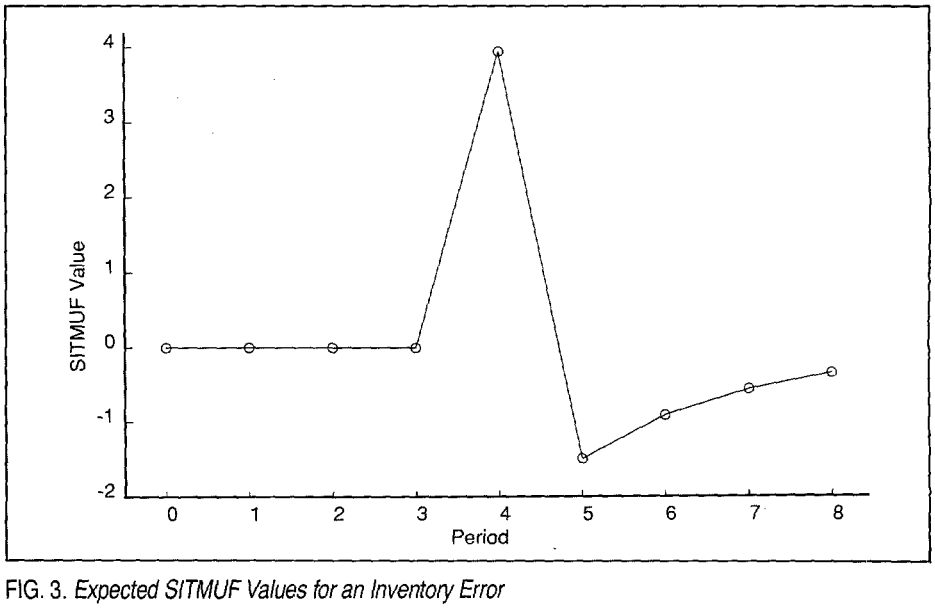

FIG. 3. *Expected SITMUF Values for an Inventory Error*

Table II: Expected SITMUF Values for an Abrupt Loss in Period 4

Period	Expected SITMUF Values	
	Loss of 8 kg	Loss of 10 kg
1	0.000	0.000
2	0.000	0.000
3	0.000	0.000
4	3.153	3.942
5	1.948	2.435

It will be shown later that study of the first 5 periods (rather than only 4) is required in order to be able to distinguish an alarm arising from an abrupt loss from one arising from, for example, an inventory measurement error. Consider, first, the expected SITMUF values corresponding to an abrupt loss of 10 kg in period 4. How well do these values match the observed SITMUF values? The expected values for a loss of 10 kg and the observed values, up to period 5, are plotted in Figure 4. The problem which remains is to find a measure of how well the sequence of expected SITMUF values fits the sequence of observed SITMUF values. One way of doing this is to calculate the differences between corresponding values in the two sequences, and then to calculate the sum of the squares of these differences, S_5, for the first 5 periods.

Then,

$$S_5 = \sum_{i=1}^{5} (SITMUF_i - \widehat{SITMUF}_i)^2$$

For an abrupt loss of 10 kg in period 4,

$$S_5 = (0.576 - 0)^2 + (1.587 - 0)^2 + (2.864 - 0)^2 + (3.605 - 3.942)^2 + (1.956 - 2.435)^2 = 11.394$$

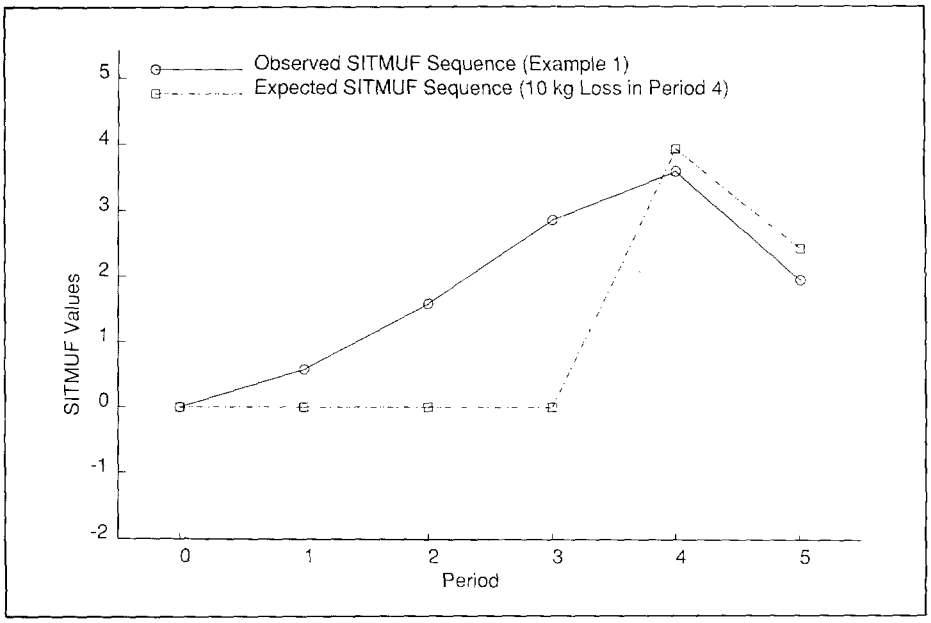

FIG.4. *Comparison of Observed and Expected SITMUF Values*

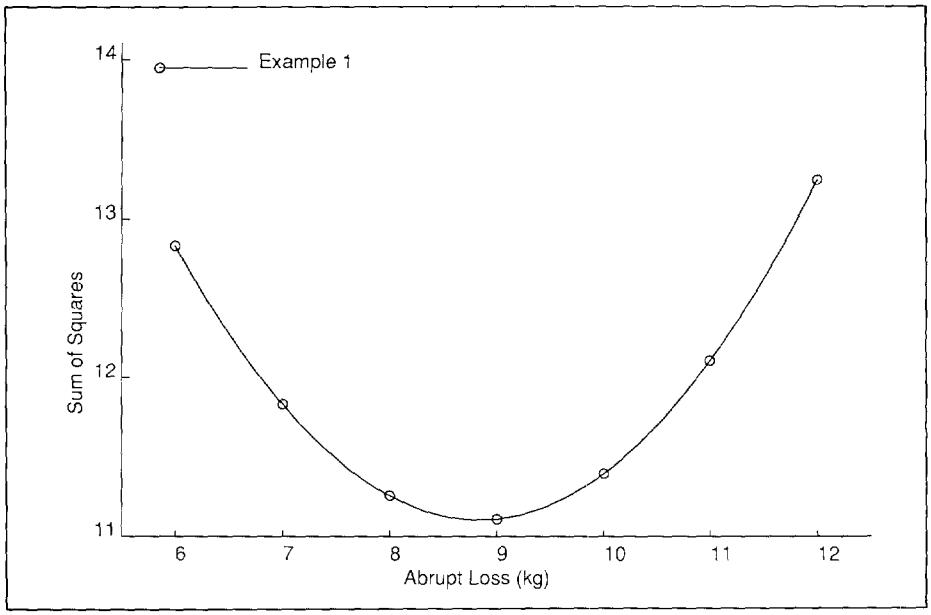

FIG. 5. *Sum of Squares for Abrupt Loss*

Least Squares Fit

Now the value of S_5 will change if the sequence of observed SITMUF values is compared with a sequence of expected SITMUF values for a different size of abrupt loss. Table III shows the sum of squares for various sizes of abrupt loss in period 4.

Table III: Sum of Squares for Various Sizes of Abrupt Loss in Period 4 (Example 1)

Loss (kg)	Sum of Squares (S_5)
6.0	12.835
7.0	11.831
8.0	11.256
9.0	11.110
10.0	11.394
11.0	12.107
12.0	13.249

The same data are plotted in Figure 5 and it can be seen that the sum of squares has a minimum value for a loss of just under 9 kg.

By calculation, the sum of squares is found to have a minimum value of 11.105, corresponding to a loss of 8.839 kg. Consider, next, the possibility that the alarm in period 4 had been caused by a loss which began in an earlier period. Table IV shows, for two examples of a loss spread uniformly over periods 2 to 4, the expected SITMUF values for the first 5 periods.

Table IV: Expected SITMUF Values for a Protracted Loss Spread Uniformly over Periods 2 to 4

Period	Expected SITMUF Values	
	Total Loss of 10 kg	Total Loss of 12 kg
1	0	0
2	1.384	1.661
3	2.254	2.705
4	2.737	3.284
5	1.691	2.029

How well does such a scenario match the observed SITMUF values? The expected values for a total loss of 12 kg and the observed values, up to period 5, are plotted in Figure 6.

For a protracted loss of 12 kg spread uniformly over periods 2 to 4,

$$S_5 = (0.576 - 0)^2 + (1.587 - 1.661)^2 + (2.864 - 2.705)^2 + (3.605 - 3.284)^2 + (1.956 - 2.029)^2 = 0.471$$

Now the value of S_5 will change if the sequence of observed SITMUF values is compared with a sequence of expected SITMUF values for a different size of protracted loss. Table V shows the sum of squares for various sizes of protracted loss spread uniformly over periods 2 to 4.

Table V: Sum of Squares for Various Sizes of Protracted Loss Spread Uniformly over Periods 2 to 4 (Example 1)

Loss (kg)	Sum of Squares (S_5)
9.0	2.639
10.0	1.569
11.0	0.847
12.0	0.471
13.0	0.442
14.0	0.760
15.0	1.425
16.0	2.437

By calculation, the sum of squares is found to have a minimum value of 0.412, corresponding to a total loss of 12.583 kg. There are, of course, a number of plausible models. Consider all the possibilities for abrupt loss, and uniform protracted loss, which end by or before period 4. Such losses are accommodated in Table VI which shows, for each loss scenario, the least squares value and, in parenthesis, the corresponding loss (in kilograms).

Table VI: Least Squares Values, and Corresponding Losses, for Various Plausible Loss Models (Example 1) [Sum of Squares to Period 5]

First Period of Loss	Last Period of Loss			
	1	2	3	4
1	16.580 (5.410)	11.098 (7.668)	5.246 (10.075)	1.596 (12.173)
2		9.857 (7.875)	3.344 (10.551)	0.412 (12.583)
3			5.409 (9.511)	3.295 (11.541)
4				11.105 (8.839)

It can be seen from Table VI that the smallest least squares value (0.412) is obtained for the model of a protracted loss spread uniformly over periods 2 to 4. This means that, from the models which have been considered, this is the one which best fits the observed data. Furthermore, the corresponding estimate of the loss for this preferred model is 12.583 kg.

The data for this example were simulated using a protracted loss of 12 kg spread uniformly over periods 2 to 4.

Example 2

Table VII: Observed SITMUF Values (Example 2)

Period	Observed SITMUF Value
1	-0.287
2	-0.143
3	0.477
4	3.497
5	2.409

Table VIII: Least Squares Values, and Corresponding Losses, for Various Plausible Loss Models (Example 2) [Sum of Squares to Period 5]

First Period of Loss	Last Period of Loss			
	1	2	3	4
1	17.350 (1.618)	16.213 (2.743)	13.828 (4.509)	9.620 (7.020)
2		15.313 (3.238)	11.875 (5.424)	6.206 (8.371)
3			9.984 (5.838)	2.562 (9.252)
4				0.375 (9.154)

As for the first example, consider all the possibilities for abrupt loss, and uniform protracted loss, which end by or before period 4. Such losses are accommodated in Table VIII which shows, for each loss scenario, the least squares value and, in parenthesis, the corresponding loss (in kilograms).

It can be seen from Table VIII that the smallest least squares value (0.375) is obtained for the model of an abrupt loss in period 4. This means that, from the models which have been considered, this is the one which best fits the observed data. Furthermore, the corresponding estimate of the loss for this preferred model is 9.154 kg.

The data for this example were simulated using an abrupt loss of 10 kg in period 4.

Example 3

Table IX: Observed SITMUF Values (Example 3)

Period	Observed SITMUF Value
1	0.477
2	-0.148
3	-0.579
4	4.083
5	-1.540

As for the previous examples, consider all the possibilities for abrupt loss, and uniform protracted loss, which end by or before period 4. Such losses are accommodated in Table X which shows, for each loss scenario, the least squares value and, in parenthesis, the corresponding loss (in kilograms).

There seems to be a problem with this example. In contrast to Examples 1 and 2, there is not a small least squares value for any of the models. In other words, none of the models seems to fit the observed SITMUF values. The observed data (Table IX) show a large positive value in period 4, followed by a large negative value in period 5. This pattern of data is similar to that displayed in Figure 3 for an inventory error at the end of period 4. Table XI shows, for two examples of an inventory error at the end of period 4, the expected SITMUF values for the first 5 periods.

The expected values for an inventory measurement error of 10 kg and the observed values, up to period 5, are plotted in Figure 7.

Table XII shows the sum of squares for various sizes of inventory measurement error at the end of period 4.

By calculation, the sum of squares is found to have a minimum value of 0.585, corresponding to an inventory measurement error of 10.361 kg.

The data for this example were simulated using an inventory measurement error 10 kg at the end of period 4.

Example 3 (reworked)

This last example is now reworked but, this time, only calculating the sum of squares up to period 4, the period of the alarm. Table XIII is analogous to Table X except that the former includes the least squares values, S_4, instead of S_5.

The natural conclusion from Table XIII is that the low least squares value (0.585) is evidence of an abrupt loss in period 4. In practice, such losses of material will be rare. A more likely occurrence is a gross inventory mismeasurement or assignment of a throughput determination to the wrong balance period. Detection of such events is a vital element of good materials control and accountancy and it is important that the operator does not embark on what will turn out to be fruitless follow-up investigations because he has drawn the wrong conclusion from the initial anomaly resolution procedures. One way of reducing this risk, when beginning to resolve anomalies, is by taking account of data in the next period after an alarm.

Table X: Least Squares Values, and Corresponding Losses, for Various Plausible Loss Models (Example 3) [Sum of Squares to Period]

First Period of Loss	Last Period of Loss			
	1	2	3	4
1	19.047 (1.224)	19.070 (1.396)	18.710 (2.027)	17.262 (3.651)
2		19.264 (1.115)	18.741 (2.004)	16.754 (4.069)
3			18.367 (2.251)	15.269 (4.859)
4				12.529 (5.750)

Table XI: Expected SITMUF Values for an Inventory Error at the End of Period 4

Period	Expected SITMUF Values	
	Error of 8 kg.	Error of 10 kg.
1	0.000	0.000
2	0.000	0.000
3	0.000	0.000
4	3.153	3.942
5	-1.187	-1.483

Table XII: Sum of Squares for Various Sizes of Inventory Error at the End of Period 4

Error)kg)	Sum of Squares (s_5)
7.0	2.589
8.0	1.574
9.0	0.914
10.0	0.608
11.0	0.658
12.0	1.062
13.0	1.820
14.0	2.934

Table XIII: Least Squares Values, and Corresponding Losses, for Various Plausible Loss Models (Example 3) [Sum of squares to Period 4]

First Period of Loss	Last Period of Loss			
	1	2	3	4
1	16.286 (1.599)	16.168 (1.984)	15.349 (3.016	12.316 (5.638)
2		16.394 (1.756)	15.207 (3.175)	10.811 (6.669)
3			14.401 (3.568)	7.225 (8.327)
4				0.585 (10.358)

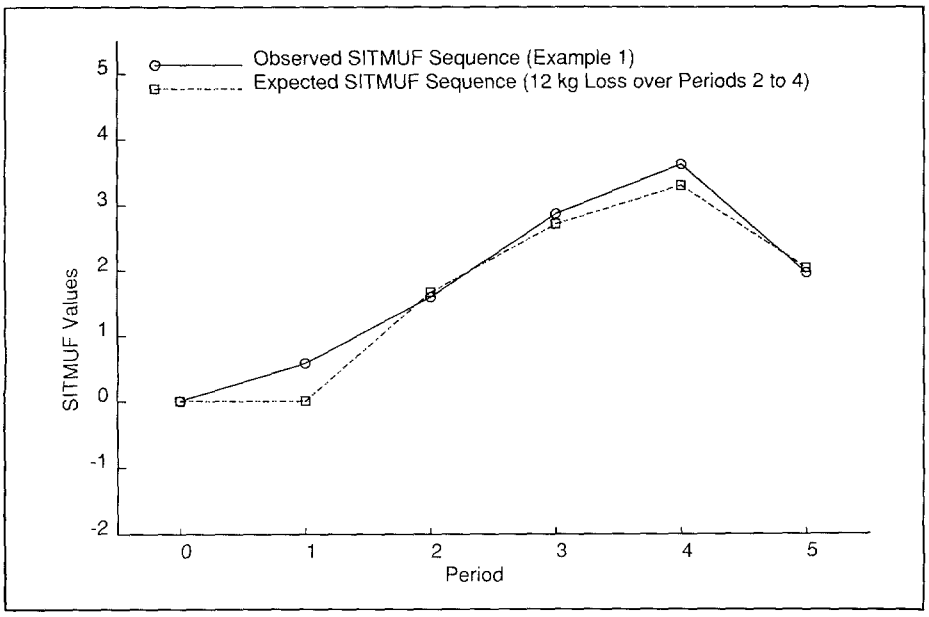

FIG. 6. *Comparison of Observed and Expected SITMUF Values*

Concluding Remarks

The paper concentrates on those modules for sequential testing and anomaly resolution which are versatile and, therefore, capable of being installed as part of any NRTMA system. When the Joint Test gives an alarm, there is no easy way of attributing that alarm to a specific cause. Any test procedure is insufficient on its own; a follow-up anomaly resolution procedure is essential.

The paper approaches the concept of anomaly resolution by treating it as an exercise in mathematical modelling. Expected SITMUF sequences can be derived using a model and incorporating plausible losses or errors. In order to measure how well a particular sequence of the expected SITMUF values fits the sequence of observed SITMUF values, the differences between corresponding values in the two sequences, up to the period after the alarm from the Joint Test, are calculated, squared, and added to one another. For each plausible loss/error model, the parameters can be chosen so that the sum of squares is minimized. The model most likely to match the observed data is the one for which the least squares value is the smallest.

This approach of mathematical modelling is made possible because of the properties of the SITMUF statistic. Firstly, sequences of expected SITMUF values can be calculated for various plausible loss/error models. Secondly, the SITMUF series is made up of independent items with the same variance so that observed and expected sequences can be compared by the method of least squares. The approach is applicable to any plant. SITMUF values are plant dependent but this is taken care of by the Joint Test.

When an alarm has occurred, the anomaly resolution procedure can make good estimates of the total loss, the type of loss pattern, and the periods in which the loss occurred. The procedure can also identify inventory measurement errors and distinguish them from real losses. This has been illustrated by a range of worked examples.

Whilst the paper shows the sound basis for an anomaly resolution procedure, more development is required to produce a comprehensive procedure for operational purposes. Two areas for further work can be already identified. Firstly, those models which best fit the observed data should be ranked. Secondly, standard errors associated with loss/error estimates need to be calculated. The necessary mathematical methods to support these developments have already been identified. The next step is to incorporate these methods into efficient working procedures.

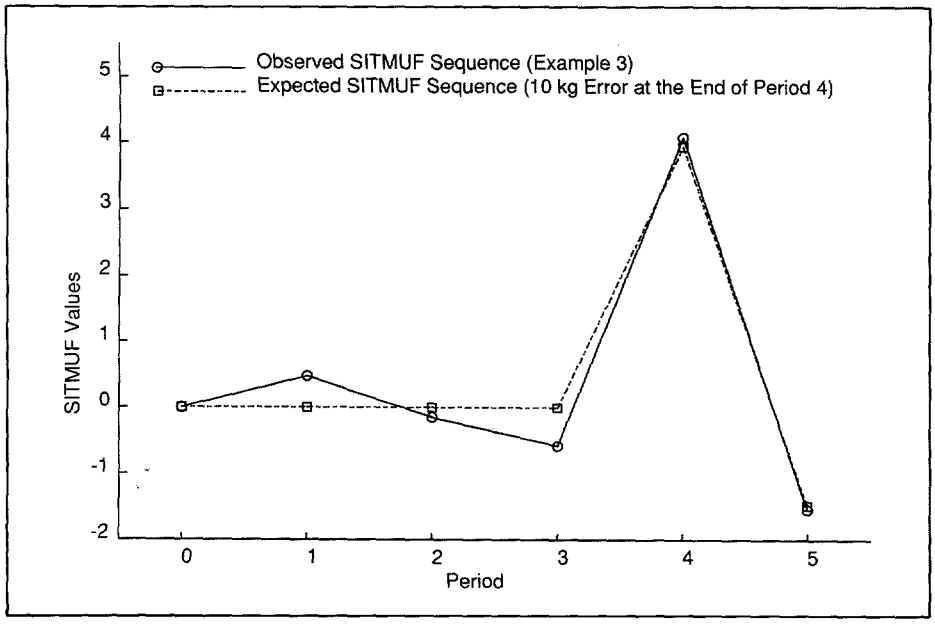

FIG. 7. *Comparison of Observed and Expected SITMUF Values*

References

/1/ BARRY J. JONES – Near Real Time Materials Accountancy: Myths and Misconceptions – Proceedings of the Eleventh Annual ESARDA Symposium on Safeguards and Nuclear Material Management, 357 (1989)

/2/ A J WOODS & D J PIKE – Use of the Standardised ITMUF for Choosing Parameters of Sequential Tests for Protracted and Abrupt Diversions – Proceedings of the Fifth Annual ESARDA Symposium on Safeguards and Nuclear Material Management, 369 (1983)

/3/ BARRY J JONES – Comparison of Near Real Time Materials Accountancy Using SITMUF and Page's Test with Conventional Accountancy – Proceedings of the Ninth Annual ESARDA Symposium on Safeguards and Nuclear Material Management, 255 (1987)

/4/ BARRY J JONES – Near Real Time Materials Accountancy Using SITMUF and a Joint Page's Test: Dependence of Response on Balance Frequency – Proceedings of the Third American Nuclear Society Conference on Facility Operations - Safeguards Interface, 242 (1987)

/5/ BARRY J JONES –Near Real Time Materials Accountancy Using SITMUF and a Joint Page's Test: Comparison with MUF and CUMUF Tests – ESARDA Bulletin, 15, 20 (1988)

/6/ BARRY J JONES –Near Real Time Materials Accountancy Using SITMUF and a Joint Page's Test: Improvement of the Test – ESARDA Bulletin, 16, 13 (1989)

Image Processing Methods for Scene Change Detection and Motion Detection

M.J. Mol

Commission of the European Communities
Joint Research Centre, Ispra Site, Italy

Abstract

The selection of reduced pixel sets and their use in three different algorithms for the detection of local scene changes is considered.

The three algorithms are compared to reveal their specific advantages and disadvantages for practical scene change detectors. The combination of scene change detection with programmable, timer controlled sequences of pixel sets allows for the selective detection of objects moving within certain speed limits into a given direction only.

Introduction

Optical surveillance plays an important role in nuclear safeguards, especially in inaccessible areas or in storage facilities for nuclear materials. Usually, optical surveillance is performed by recording video images from one or more TV cameras looking at safeguards relevant areas.

The increasing amount of recorded surveillance data (presently several million images per year) suggests the use of computerized vision systems either for the partial automation of the reviewing task, or for data reduction during the recording process. Such systems are based on one of two possible operating principles. The first one, to which this paper is related, detects changes inside areas of interest in video images. The second, which is based on the recognition of objects, is not considered here.

With the advent of plug-in image processing boards for the PC/AT bus, it became possible to implement low-cost, PC-based image processing systems from "off the shelf" hardware. Such systems are very flexible because they can easily be adapted to different applications by merely changing the program running on the PC. The image processing hardware consists of a digitizer which samples the video signal, a frame buffer to store the digital samples as individual picture elements (pixels), and a digital to analogue converter which generates a video output signal from the stored pixels. Since a program running on the PC has access to all pixels in the frame buffer, it also has full image processing capabilities. Some popular image processing boards offer frame buffers for multiple video images with resolutions of 512*512 pixels and with 256 grey levels per pixel.

A possible application for such a PC-based image processing system is as a scene change detector or motion detector for the above mentioned optical surveillance in nuclear safeguards. Not only as a reviewing aid, but also as front-end unit to control the selective recording of relevant scenes.

Different image processing techniques can be applied for the detection of scene changes. However, if a detector must be able to detect changes in a single frame of a live video signal, it must be very fast. To obtain the necessary speed with "off the shelf" hardware, simple detection algorithms must be used which need only a limited number of pixels for processing. In the next chapter, the selection of representative sub-sets of pixels from images is described first. Thereafter, the characteristics and performance of three simple detection algorithms are compared. Finally, the use of timer-controlled mask sequences for the selective detection of objects moving into a given direction is explained.

The Selection of Examination Sets

It is often not necessary to process all pixels of an image. Either because only a limited image area is of interest, or because a fast detector response requires a substantial reduction of the number of pixels to be processed. In either case, ordered subsets of pixels must be sampled from an image to form one or more examination sets. Whereas the original image is a two-dimensional function of grey levels, an examination set can often be treated as a one dimensional array of grey levels and be interpreted or visualized as a simple, plain graph. It is also called the profile of the examination set.

The selection of pixels for an examination set is a sampling process based on a logical function which has two possible states for each pixel, TRUE or FALSE. If TRUE, the pixel concerned is included in the examination set, otherwise it is neglected. The following five examples of examination sets use entirely different sampling rules.

- A well known examination set is the rectangular area of interest (AOI). It contains all pixels within a rectangular area defined by two diagonally opposite corner points.
- The "grid line" set contains all pixels along equally spaced, horizontal and vertical grid lines. It can be combined with the definition of a rectangular AOI to cover a limited region only.
- The "grid point" set contains only equally spaced grid points and is a strongly reduced subset of the grid line set.
- A "polyline" is an examination set that contains only the pixels along the straight line sections between neighbouring elements of a vertex sequence which is defined in a vertex table. The profiles of polylines have a real, physical meaning, namely the grey level as a function of the distance along the polyline. Figure 1 is the graphic display of a profile captured during video tape reviewing.
- It is also possible to select pixels via a two-dimensional look-up table. This allows for the creation of examination sets that cannot easily be defined by analytical functions or by simple geometrical rules.

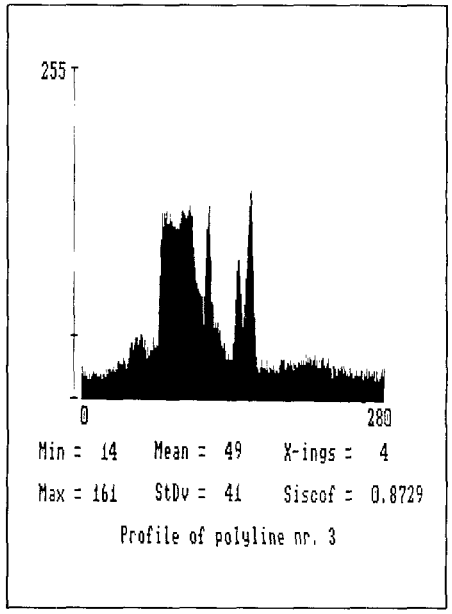

FIG. 1. *Example of a polyline profile*

The first three of the above examples can easily be implemented in hardware. Some earlier motion detectors, realized with special hardware, were based on a few rectangular AOIs with user definable positions and sizes. More recent, micro-computer based detection systems often divide the entire image into a large number of small, rectangular AOIs which can be individually activated or disactivated by the user.

Polylines are most suited for software-based, general-purpose systems which use fast routines to read the pixels along the polylines from the frame buffer into the computer memory. Such systems leave the user completely free to create polylines of any shape in any position and allow

the surveillance of larger image areas without the penalty of a very large number of pixels to be processed. Examples of possible polyline shapes are shown in figure 2. The "scanner" line is interesting because it acts like a reactangular AOI.

In the following text the term polyline will systematically be used to refer to an examination set of any type, at least as long as there is no need to emphasize a different behaviour of a specific type.

Detection Algorithms

Although the three algorithms are described in terms of their use with polylines, the sampling rules for these polylines are not considered here, since they are not important for the understanding of the algorithms proper. The following notations are used to distinguish various polylines in two different images.

If there are k polylines in the image X, they are designated X_1, containing N_1 pixels, to X_k with N_k pixels, and the i-th pixel of the j-th polyline is named $X_{j(i)}$. The corresponding polylines in the reference image Y are named Y_1 through Y_k, and the i-th pixel of the j-th line is $Y_{j(i)}$.

Detection Based on Changes in the Normalized Means

The simplest detection algorithm calculates first the normalized mean intensities of the individual polyline profiles and then their ratios for the corresponding polylines in the two images X and Y. If one of these ratios deviates from 1.0 by more than a given amount, the corresponding polyline enters a warning state. The following expressions show that this concept requires minimal computational efforts.

The sums of the grey levels of the pixels in the j-th polyline in the images X and Y are designated AX_j and AY_j. Thus,

$$AX_j = \sum_{i=1}^{N_j} X_{j(i)} \qquad [1]$$

and

$$AY_j = \sum_{i=1}^{N_j} Y_{j(i)} \qquad [2]$$

and the mean grey levels of these polyline profiles are

$$\bar{X}_j = \frac{AX_j}{N_j} \qquad [3]$$

and

$$\bar{Y}_j = \frac{AY_j}{N_j} \qquad [4]$$

The ensemble sum over all pixels of all polyline profiles in the image X is

$$EX = \sum_{j=1}^{k} \sum_{i=1}^{N_j} X_{j(i)} \qquad [5]$$

and in image Y

$$EY = \sum_{j=1}^{k} \sum_{i=1}^{N_j} Y_{j(i)} \qquad [6]$$

The ratio between the normalized means of the j-th polyline profiles in the images X and Y is therefore

$$RM_j = \frac{\bar{X}_j}{\bar{Y}_j} \cdot \frac{EY}{EY} \qquad [7]$$

If each polyline has its own threshold value $Mthr_{(j)}$, with $0.0 < Mthr_{(j)} < 1.0$, the j-th polyline enters the warning state if

$$RM_j < Mthr_{(j)}$$
or $$1/RM_j < Mthr_{(j)}$$

The simplicity of this algorithm makes it very attractive for special hardware implementations. Except for the final division operation, it only requires adders and a few accumulator registers and, in principle, the summing process could even be performed with gated analogue integrators. Moreover, only the profile sums and ensemble sum of the reference image need to be saved for subsequent computing cycles, not the polyline profiles themselves.

Detection Based on Changes in the Normalized Variance

The second detection algorithm calculates first the normalized variances of the individual polyline profiles and then their ratios for the corresponding polylines in the two images X and Y. If one of these ratios deviates from 1.0 by more than a given amount, the corresponding polyline enters a warning state. The following description explains this concept.

The means Xj and Yj of the individual polyline profiles are again computed as shown in the above expressions [3] and [4], but now also the sums SXj and SYj of the squared deviations are calculated as

$$SX_j = \sum_{i=1}^{N_j} (X_{j(i)} - \bar{X}_j)^2 \qquad [8]$$

and

$$SY_j = \sum_{i=1}^{N_j} (Y_{j(i)} - \bar{Y}_j)^2 \qquad [9]$$

The variance of the j-th polyline profile in image X is

$$VX_j = \frac{SX_j}{Nj - 1} \qquad [10]$$

and the variance of the j-th polyline profile in image Y is

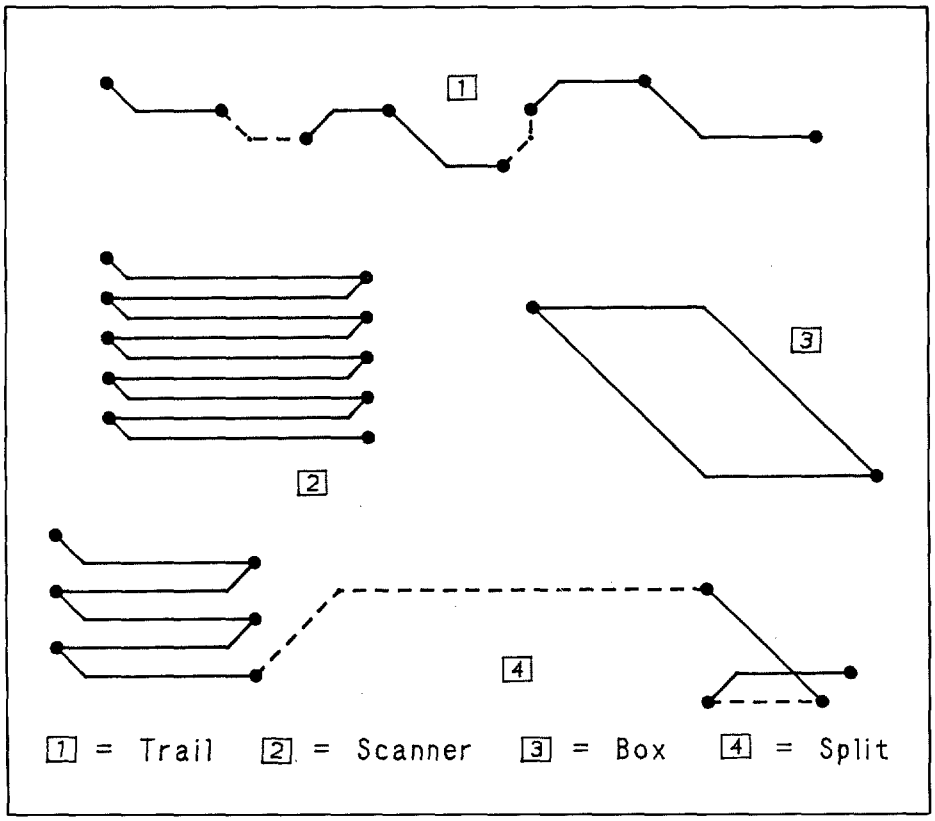

FIG. 2. *Example of typical polylines*

$$VY_j = \frac{SY_j}{N_j - 1} \qquad [11]$$

The ratio between the normalized variances is obtained as

$$RV_j = \frac{SX_j}{SY_j} \cdot \left(\frac{\overline{Y}_j}{\overline{X}_j}\right)^2 \qquad [12]$$

If each polyline has its own threshold value Vthr(j), with 0.0 < Vthr(j) < 1.0, the j-th polyline enters the warning state if

$$RV_j < Vthr_{(j)}$$
or
$$1/RV_j < Vthr_{(j)}$$

This algorithm is less simple than that for the normalized mean. It not only needs the addition but also the squaring of the individual pixel values. Nevertheless, only the means and the sums of squared deviations of the polyline profiles in the reference image need to be saved for subsequent computing cycles, not the reference profiles themselves.

Detection Based on the Cross-correlation of Polylines

The third detection algorithm calculates the correlation factor for each individual polyline profile in image X with its corresponding profile in the reference image Y. If a correlation factor drops below a given threshold, the corresponding polyline enters a warning state. The following description explains this concept.

First, the means X_j and Y_j and the sums of squared deviation SX_j and SY_j are computed as shown in the above expressions [3], [4], [8] and [9]. Thereafter, the sum of cross-products of the polyline profiles X_j and Y_j, which is a measure of their co-variance, is calculated as

$$XY_j = \sum_{j=1}^{N_j} (X_{j(j)} - X_j).(Y_{j(j)} - Y_j) \qquad [13]$$

which leads to the correlation factor

$$RC_j = \frac{XY_j}{\sqrt{(SX_j.SY_j)}} \qquad [14]$$

If each polyline has its own threshold value $Cthr_{(j)}$, with $0.0 < Cthr_{(j)} < 1.0$, the j-th polyline enters the warning state if

$$RC_j < Cthr_{(j)} \qquad [15]$$

This algorithm is the most computation intensive of the three. It not only needs addition and squaring operations on each pixel, but also the calculation of the cross-product for each pixel pair. Moreover, these cross-product calculations require that the entire profiles of the polylines in the reference image remain available for subsequent computing cycles.

If the two terms in expression [15] are interchanged, a polyline enters the warning state if its correlation factor exceeds the threshold value, which means if there is a close match between the profiles from the current and reference image. It is therefore very easy to provide a means for switching between a normal scene change detection mode and an optional scene matching mode. For the latter mode, fixed reference profiles must be memorized for all polylines, and the detector must continuously compare these memorized profiles with the profiles read from the images in the video signal.

The scene matching mode is perhaps a questionable feature for some video surveillance applications because it may fail to generate an alarm if the matching pattern appears in the image at a slightly different position than expected. On the other hand, profile matching is a basis for simple object tracking algorithms.

Common Properties of the Algorithms

The three detection algorithms described above are auto-normalizing. This means that they are insensitive to changes in overall scene illumination as long as there is a linear relation between the amplitude of the video signal and luminosity. Unfortunately, many video cameras have a GAMMA-corrector installed which intentionally produces a nonlinear response. The software in a PC-based scene change detector can compensate for such nonlinearity if it reprograms the input look-up table on the image processing board accordingly. Another interesting feature of a PC-based system is that it requires little programming efforts to include all three detection algorithms as run-time selectable options.

Comparison of the Detection Algorithms

The algorithm using the normalized means is not only the simplest, it is also the easiest understood by non-expert users. It is not based on an intrinsic property of a single polyline profile, but relies on the existence of several profiles to form the normalizing ensemble. Since its detection performance is independent of the image contrast, it can still successfully be applied when, due to a lack of contrast, the two other algorithms might produce too many nuisance alarms. On the other hand, even large changes within a polyline profile will go unnoticed if they do not affect the profile's mean intensity.

The algorithm based on changes in the normalized variance, which is an intrinsic statistical property of a polyline profile, makes the detector sensitive to changes in grey level distribution within a polyline profile, but not to changes outside the polyline. However, changes within a profile that do not produce significant changes in variance, but which might be characteristic for the displacement of an object in front of a contrasting background, will go unnoticed.

The detection method based on the cross-correlation of polyline profiles is definitely the most selective and most severe. It does not check a single, intrinsic statistical property of a profile, but it compares the profiles of corresponding polylines in images X and Y on a pixel by pixel basis in order to detect any significant pattern changes between these profiles. However, a necessary condition for the successful application of the cross-correlation technique is a sufficiently large contrast within the single polyline profiles. A too low contrast leads to many nuisance alarms. A suitable measure for the proper contrast in a polyline profile is its variance, which should be many times higher than the variance of the inevitable image noise.

Figure 3 shows the same rectangular AOI from four different images in a sequence. Only three grey levels are present, 0.0, 0.5 and 1.0. The image pattern were chosen to simplify the calculations of the means, variances and cross-correlations, and to emphasize the different behaviour of the three detection algorithms. In the table of figure 3, a change from one image to the next is marked '+' if the type of pattern change it produces

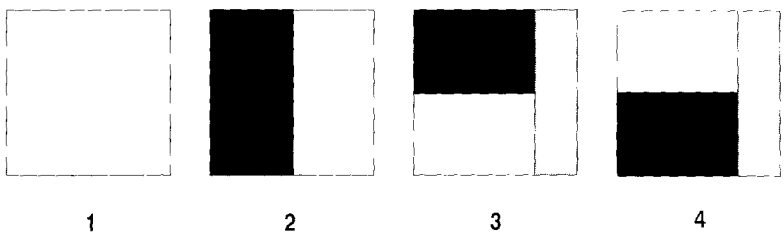

ALGORITHM \ CHANGE	1 2	2 3	3 4
Ratio of normalized means RM: Detection response:	0.5 +	1.0 —	1.0 —
Ratio of normalized variance RV: Detection response:	0.0 +	0.75 +	1.0 —
Correlation factor RC: Detector response:	0.0 +	0.0 +	–1.0 +

FIG. 3. *Response of detection algorithms to three classes of scene changes*

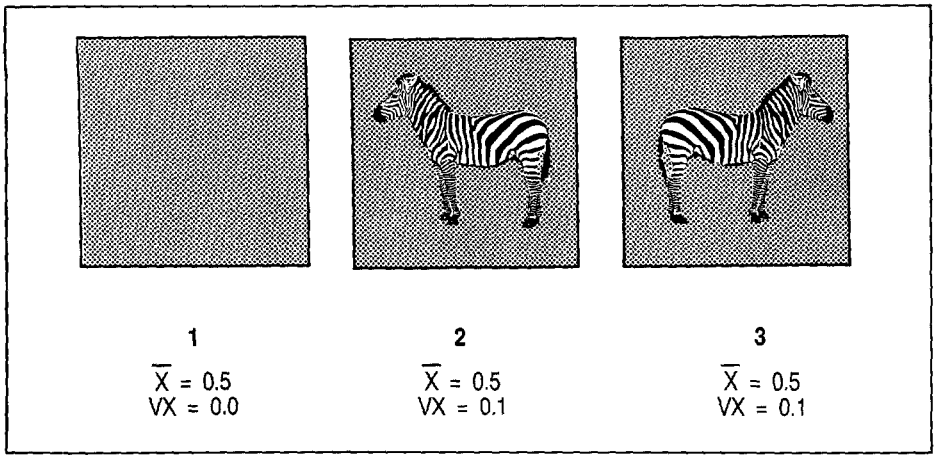

ALGORITHM \ CHANGE	1 → 2	2 → 3
Ratio of normalized means Response:	—	—
Ratio of normalized variances Response:	+	—
Correlation factor Response:	+	+

FIG. 4. *Examples of scene change detection*

can be detected with the algorithm to which the corresponding table row belongs. Else, a '-' is shown.

Figure 4 is another illustration of the differences between the three detection methods. The AOI number 1 covers an area with a homogeneous grey level of 0.5, e.g. part of a concrete wall. When a zebra with equally distributed black and white stripes steps into the picture and stays in front of the wall, a situation like in AOI number 2 arises. Since the mean grey level has not changed, a detector based on the normalized mean is unable to detect this event. On the other hand, the variance changed considerably and detectors based on the normalized variance or on the correlation of profiles will respond with a warning. If the zebra moves inside the AOI, e.g. by turning around, the situation becomes as shown in AOI number 3. In this case, neither the mean grey level nor the variance changes, and only the detector based on the cross-correlation algorithm issues a warning.

Using Detection Masks and Timed Mask Sequences

The inner loop of a scene change detector is rather simple. After a new video image has been acquired, the detection algorithm uses the polyline profiles of that image plus the information from the reference image to establish which profiles have changed and to set the warning pattern accordingly. Thereafter, the new warning pattern is compared with a group of detection masks. Each detection mask is simply a user programmable list of polylines which defines a possible alarm condition by specifying which combination of polylines must be in the warning state to trigger an alarm. Consequently, an alarm is generated if a match between the warning pattern and a mask is found. In the usual scene change detection mode, generating an alarm also causes reference updating by making the current image the new reference.

The clever use of detection masks, in combination with a suitable set of polylines, is the users contribution to a satisfactory detection performance. An often successful technique for detecting an activity is to create a set of polylines with profiles that are all changed by that activity, but which are not subject to the same disturbances, and to combine these polylines in a detection mask.

The concept of detection masks can be extended to make a scene change detector only sensitive to changes that "move" through an image in a given direction within certain speed limits. Such extension, which is independent of the detection algorithm used and enhances the system performance to that of a direction sensitive motion detector, is based on timer-controlled sequences of alarm conditions specified in secondary masks. If a detection mask, which will now be called a "primary detection mask", matches the warning pattern, it does not trigger an alarm any more. Instead, it activates its own sequence of secondary masks by passing control to the first mask in that sequence. Each "secondary mask" consists of the mask proper, which is again a user programmable list of polylines that defines an alarm condition, plus programmable specifiers for a sleep

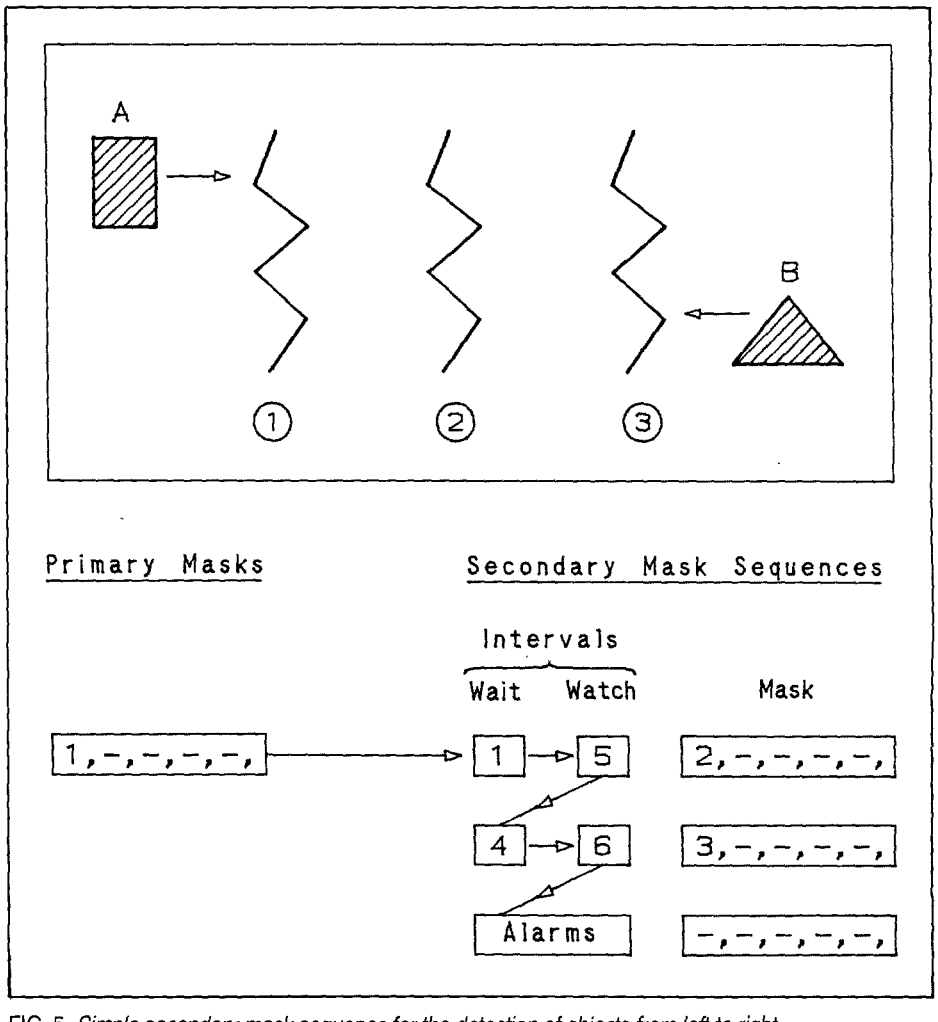

FIG. 5. *Simple secondary mask sequence for the detection of objects from left to right*

interval and a watch interval. When a secondary mask of a sequence gets control, it waits inactively until its sleep interval has expired and then starts its watch interval during which it responds to a match with the warning pattern by passing control to the next mask of the sequence. If the last secondary mask of a sequence matches the warning pattern during its watch interval, it triggers the alarm. If the watch interval of a secondary masks times out without the occurence of a match, the sequence ends prematurely without generating an alarm. A very simple example of direction sensitive motion detection is shown in figure 5, where the polylines numbered 1, 2 and 3 reside in the left, middle and right part of an image. A primary mask specifying only polyline number 1 triggers a sequence of two secondary masks, the first one specifying polyline 2 with sleep and watch intervals of 1 and 5 seconds, the second polyline 3 with sleep and watch intervals of 4 and 6 seconds. This makes the detector only sensitive to objects moving from left to right, like the rectangle A in figure 5, which cover the distance between the polylines 1 and 2 in 1 to 6 seconds and between polylines 2 and 3 in 4 to 10 seconds. The triangular object B, moving from right to left in figure 5, will not be detected. If a system allows for 16 different polylines and for 16 different primary detection masks, each of them leading to a sequence of up to 8 secondary masks, extremely complicated sequences of alarm conditions can be created.

Although the addition of timer controlled mask sequences makes it possible to reduce the number of nuisance alarms by detecting objects moving into a given direction only, it also complicates the installation of a scene change detector considerably. The proper and successful use of polylines, primary masks and sequences of secondary masks requires a sound knowledge of the scenes under surveillance and of the type of changes that should be detected.

Conclusion

The algorithms and image processing methods described in the previous chapters were used in a PC-based prototype scene change detector. Although laboratory tests on its detection performance produced very positive results, true performance evaluation is only possible in field tests with real surveillance data from installations or after an extended period of operation as "front-end" detector.

The successful setting-up of such a scene change detector requires a considerable understanding of its operating principle. Not only for creating the correctly shaped polylines or areas of interest, but also to select the most suitable detection algorithm for the prevailing conditions. However, once the installation dependent parameters have been established and stored in a dedicated data file, the use of the detector for routine reviewing of video tapes becomes rather simple. If the necessary knowhow and experience for the set-up procedure are present, scene change detectors based on the described image processing methods may become powerful tools for video reviewing and for front-end data reduction.

References

/1/ M.J. MOL – Program REVIEW for real-time scene change detection in video images on a PC based vision system.I Report EUR 11677 EN June 1988.

/2/ M.J. MOL – A real-time scene change detector for the review of video images, Proceedings of the INMM 29th Annual Meeting, Vol. XVII. 1988.

/3/ M.J. MOL – Extensions and changes to program REVIEW for real-time scene change detection in video images, Report EUR 12055 EN December 1988.

/4/ M.J. MOL – Considerations on the use of grey pattern correlation for object identification and scene change detection in video images., Report EUR 12829 EN Aoril 1990.

Random Interim Inspections at Power Reactors - A Fable

M. Canty and R. Avenhaus

Part I. The Inspector who got Something for Nothing

Once upon a time there was a safeguards inspector who wanted to spend more time with his family. The inspector was responsible for a power reactor in a far away land, and had to journey there once every year when the reactor was refuelled as well as three times in between because of his timeliness goal.

One day, during a refuelling inspection, the inspector went to the reactor boss and said that he would like, in future, to be allowed to perform his interim inspections on the last day of every month, instead of once every three months as had been the case up until then. The reactor boss frowned and asked the inspector if this was absolutely necessary. The inspector replied that it was, in the interest of increased safeguards effectiveness and efficiency. Upon hearing these words, the reactor boss sighed wearily and agreed to the inspector's request. The inspector then asked the reactor boss if he would be offended if he, the inspector, didn't show up for some of the eleven interim inspections. The reactor boss was puzzled, but said he would most certainly not be offended. The inspector then left the power reactor, rejoicing inwardly, saying "Now I shall only have to make two interim inspections per year, rather than three. I will still attain my timeliness goal and will be able to spend more time with my family!".

Upon arriving back at headquarters, the inspector went to his safeguards bosses and told them of the deal he made at the reactor, and how he intended to save one interim inspection per year. At first the safeguards bosses were very angry, saying that the inspector was mad to think that he could get something for nothing and that his calculations must be incorrect. But then they consulted the literature and found a paper by two obscure but reputable safeguards experts /1/ which confirmed exactly the calculations of the inspector. Then the safeguards bosses laughed and said that the inspector was very wise, and if he got much wiser he wouldn't have to work at all. They rewarded him by making him responsible for a second power reactor.

Thus end the sad tale of the inspector who got something for nothing, but was not able to spend more time with his family.

Part II. The Inspector's Calculation

The inspector, being a man of mathematical inclination, reasoned as follows: The reactor operator agreed to eleven interim inspection opportunities per year. If, as the inspector intended, only two interim inspections actually take place, then there are precisely 55 possible inspection strategies, the number of combinations of 11 things taken two at a time. These strategies can be denoted (1,2), (1,3), ... (10,11). For example, the strategy (5,7) means that interim inspections occur at the end of the 5th and 7th months. The operator, should he wish to divert a spent fuel element, has only 11 sensible strategies, namely to divert at the beginning of the i-th month, i = 11. (He would be foolish to divert in the 12th month, since he knows that the PIV inspection always takes place, and therefore the detection time cannot exceed one month.) For each strategy combination of inspector and operator, there is exactly one detection time. For example for the combination {(5,7),6}, in which the operator diverts at the beginning of the 6th month, the time to detection is 2 months /2/. For the combination {(5,7),8} the detection time is 5 months (because of the PIV after the 12th month). Thus the inspector was able to construct an 11 x 55 matrix which completely described his problem as a finite two-person zero-sum game. He then applied the simplex algorithm to find the saddlepoint of the game, which, by virtue of von Neumann's minimax theorem /3,4/, he knew must exist. The solution told him the optimal mixed strategy, i.e. with which probability he should choose one of the 55 possible interim inspection strategies in any given 12 month period. (The non-zero inspection probabilities for the solution are shown in Table I.) The value of the game, that is the guaranteed payoff to the inspector, turned out to be a detection time of 2.97 months, just within his timeliness goal of three months for irradiated direct use material.

Table I. Optimal mixed inspection strategy

Strategy	Probability
(1,5)	0.035
(1,6)	0.167
(1,7)	0.005
(2,7)	0.195
(2,8)	0.004
(3,8)	0.200
(4,8)	0.046
(4,9)	0.163
(5,9)	0.170
(5,10)	0.015

References

/1/ M.J. CANTY and R. AVENHAUS , "Optimal Randomization Strategies for Timely Verification", to be presented at the 13th ESARDA Symposium, Avignon, 1991

/2/ The inspector was using very modern C/S equipment, made in Germany, so his detection probability was of course 100%.

/3/ J. VON NEUMANN and O. MORGENSTERN, "Theory of Games and Economic Behavior", Princeton University Press, 1947

/4/ M.J. FRYER, "An Introduction to Linear Programming and Matrix Game Theory", Arnold, London, 1978

LB-AB-91-019-EN-N