

# The fusion of redundant SEVA measurements

Mihaela Duta, Manus Henry

## Abstract—

The self-validating (SEVA) sensor carries out an internal quality assessment, and generates, for each measurement, standard metrics for its quality, including on-line uncertainty. This paper discusses consistency checking and data fusion between several SEVA sensors observing the same measurand. Consistency checking is shown to be equivalent to the maximum clique problem, which is NP-hard, but a linear approximation is described. A technique called uncertainty extension is proposed which causes a smooth reduction in the influence of outliers as they become increasingly inconsistent with the majority.

**Index Terms**— Sensor fusion, SEVA sensors, Maximum clique problem

## I. INTRODUCTION

This paper proposes a simple method for fusing measurement data from a set of independent self-validating (SEVA) sensors monitoring the same real-time measurand in order to provide the combined best estimate for the value, uncertainty and measurement status of the measurand. The method also provides consistency checking between the measurements.

After a brief description of the concept of the SEVA sensor, the basic issues in sensor fusion will be discussed. A method of fusing a set of measurement values and uncertainties will then be presented, based on the analogy with the maximum clique problem from graph theory. A linear search method for finding the maximum clique, tuned to the specific problem at hand, is then proposed, and the results obtained after applying this method on simulated data are contrasted with those obtained by finding the maximum clique using exhaustive search.

These are followed by simulation results demonstrating how, in a three-sensor system, a variety of sensor faults are dealt with.

## II. THE CONCEPT OF THE SEVA SENSOR

For the purpose of this paper a **sensor** is a device consisting of one or more *transducers* and a *transmitter* which converts transducer signals into a form recognizable by the control or monitoring system.

Increasingly, the availability of local computing power has been exploited to carry out internal diagnostics within so-called 'intelligent' sensors. A SEVA sensor, [1], performs additional processing to generate generic validity metrics for each measurement, as follows (Figure 1):

- The **validated measurement value (VMV)** is the best estimate of the true measurand value, taking all diagnostic information into account. If a fault occurs, then the VMV is corrected to the best ability of the sensor. In the most

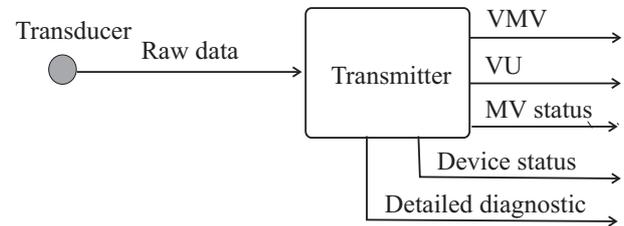


Fig. 1. The SEVA sensor

severe cases (where the raw data is judged to have no correlation with the measurand), the current VMV is extrapolated from past measurement behaviour.

- The **validated uncertainty (VU)** is the uncertainty associated with the VMV. The metrological definition is used here, [2], [3]: the VU gives a confidence interval for the true value of the measurand. For example, if VMV is 2.51 units, and the VU is 0.08, then there is a 95% chance that the true measurement lies within the interval  $2.51 \pm 0.08$  units. The VU takes into account all likely sources of error, including noise, measurement technology and any fault-correction strategy currently being used.
- The **measurement value status (MV status)** is a discrete-valued flag indicating how the VMV has been calculated. The basic categories defined are shown in Table I. The MV status assists users (whether human or automated) to determine whether the measurement is acceptable in the particular application - e.g. BLIND data should never be used for feedback control.

The SEVA sensor concept has become the basis of a British Standard for the reporting of measurement quality in industrial control systems, [4].

In the absence of localised validation, measurement redundancy has frequently been used to ensure that a verified and reliable measurement is provided with high availability. Such redundancy may be implemented through the use of several independent sensors monitoring the same measurand (hardware redundancy), or through a plant model to provide an independent estimate of the measurand (analytical redundancy), [5]. The SEVA model assumes that the designer builds into the sensor techniques to detect the most important fault modes of the device. However, there remains a non-zero probability that a fault may go undetected for a not insignificant period of time. It thus may be desirable in certain applications to use higher level validation to perform consistency checking and data fusion between redundant SEVA measurements.

In this paper a **correct** SEVA measurement  $x$  with uncertainty  $u$  means one that is truly representative of the measurand, i.e. the true value of the measurand lies within  $x \pm u$  with a probability of 95%. An **incorrect** SEVA measurement falsely rep-

SEVA category	Definition
SECURE DIVERSE	The VMV is derived from multiple values of the same measurand at least two of which are CLEAR or better and which are not susceptible to common mode faults
SECURE COMMON	The VMV is derived from multiple values of the same measurement at least two of which are CLEAR or better, but which are susceptible to common mode faults
CLEAR	The VMV is calculated normally
BLURRED	The VMV is based on live data, but is being corrected for a fault
DAZZLED	A transient state: the VMV is based on historical data while the fault is assessed
BLIND	No credible live data is available. The VMV is being projected from past behaviour
REPLACED	The VMV comes from an external agent (operator or higher level modelling program)
OFFLINE	the instrument is off-line The VMV and VU are typically zero
UNVALIDATED	Validation not in operation

TABLE I

resents the true value of the measurand, that is, there is higher than 5% probability that the true value of the measurand lies outside  $x \pm u$ . A correct measurement should not be confused with a measurement that has been ‘corrected’ as part of the diagnostic process within a SEVA sensor.

An incorrect measurement does not mean a SEVA measurement with a status other than CLEAR or SECURE. Indeed such a measurement is likely to be correct: whatever condition is causing the non-ideal status value has been identified by the sensor, which implies that the measurement and uncertainty have been corrected accordingly, so they are likely to still be representative of the measurand. As stated above, an incorrect measurement is likely to be generated if the sensor fails to notice a fault within itself, or if a phenomenon such as available measurement drift, discussed later, occurs. In either case the status is likely to be CLEAR or even SECURE. Of course if the correction algorithms are inadequate then a BLURRED, DAZZLED or BLIND measurement may also be incorrect.

### III. SENSOR FUSION: MAIN ISSUES

In the broadest sense, the process of sensor fusion is the synergistic use of a set of not necessarily consistent measurements from different sources to achieve a specific task. There is a considerable wealth of academic literature devoted to this issue. The applications range from simple problems (e.g. robot

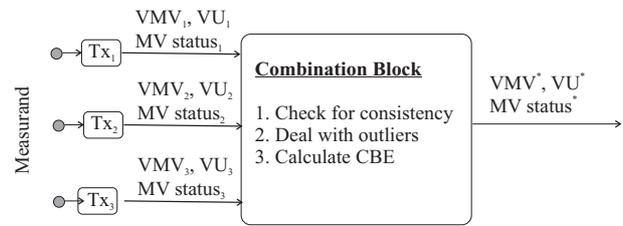


Fig. 2. The combination block for redundant SEVA measurements.

self-location) to 3-D tracking of multiple targets and decoys by fusing data from different sensor systems (e.g. radar, range-finders) in air defence systems. There also exists a wide variety of methods for performing the actual fusion of the measurements, among which the most important are the Bayesian methods ([6] and [7]), Dempster-Shafer methods ([7]) and methods based on robust statistics (see [8]). The reader is directed to these references for further details.

The basic questions posed by sensor fusion are:

- What is the model of the measurand to be tracked?
- What is the model of measurement uncertainty (here used in the broadest sense)?
- How are measurements judged to be consistent?
- What happens to measurements inconsistent with the rest (typically they are discarded)?
- How are consistent measurements fused to provide a combined best estimate of the measurand?

Clearly, the measurand and uncertainty models dictate both the consistency checking and fusion methods. In the literature a range of uncertainty models have been used, based normally on probabilistic or fuzzy models, but no examples of sensor fusion based upon metrological uncertainty have been found.

Consider now the case of  $n$  SEVA measurements  $x_i$ ,  $i = 1, \dots, n$ , with their associated uncertainties  $u_i$ , all estimating the same single valued measurand. This is a relatively simple problem compared with e.g. target tracking. What is to be determined is whether any of the measurements is inconsistent with the rest and, having dealt with any such outliers, what is the combined best estimate of the measurand and its uncertainty. The MV status for the combined measurement is also to be determined, based upon the consistency of the input measurements as well as their individual MV status values. These calculations can take place within a generic Combination Block (CB).

Figure 2 illustrates the scenario. Each measurement may be the product of a different sensor, as shown. Alternatively, all calculations can take place within a single transmitter handling multiple transducers. In either case the calculation of the combined estimate will be largely identical.

The core functions of the CB are:

- checking for consistency for a set of measurements using their uncertainties
- combining the consistent measurements and uncertainties to obtain a best estimate

These issues will be studied in detail in the following sections.

#### IV. COMBINATION OF A SET OF MEASUREMENTS

Given  $n$  measurements  $x_i$  and their associated uncertainties  $u_i$ , and assuming that they are all consistent, a combined best estimate (CBE)  $x^*$  of the measurand is given by, [3]:

$$x^* = \sum_1^n w_i x_i \text{ where } w_i = \frac{(\frac{1}{u_i})^2}{\sum_1^n (\frac{1}{u_i})^2} \quad (1)$$

It follows that the uncertainty  $u^*$  of CBE is:

$$u^* = \sqrt{\frac{\sum_1^n w_i^2 u_i^2}{(\sum_1^n w_i)^2}} = \frac{1}{\sqrt{\sum_1^n (\frac{1}{u_i})^2}} \quad (2)$$

The following example illustrates the use of these formulae. Table II lists five measurements, their uncertainties and corresponding weights. The measurement values were randomly selected from normal distributions  $\mathcal{N}(0, u_i/1.96)$ . The CBE is  $x^* = 0.0328$ , and its uncertainty  $u^* = 0.0876$ . Notice the very high weighting given to the measurement with the smallest uncertainty. Note also that the overall uncertainty is smaller than the smallest individual uncertainty. This is to be expected, as additional information about the measurand is being accumulated, so that the result should have a reduced uncertainty.

$x_i$	0.0192	0.1154	0.1868	-1.2192	-0.2808
$u_i$	0.1	0.2	0.5	1.0	2.0
$w_i(\%)$	76.8	19.2	3.07	0.768	0.192

TABLE II

Let the combination operation be denoted by:

$$x^* = \Lambda_{i=1}^n x_i \equiv \Lambda\{x_1, x_2, \dots, x_n\}$$

An important property of  $\Lambda$  is that it is associative. For example, given variables  $x_1$ ,  $x_2$  and  $x_3$  and defining

$$x_{1,2}^* \equiv \Lambda\{x_1, x_2\} \text{ and } x_{2,3}^* \equiv \Lambda\{x_2, x_3\}$$

it is trivial to show that

$$\Lambda\{x_{1,2}^*, x_3\} \equiv \Lambda\{x_1, x_{2,3}^*\} \equiv \Lambda\{x_1, x_2, x_3\}$$

This implies that the order in which measurements are combined is not important.

#### V. CONSISTENCY CHECKING

##### A. Consistent vs. inconsistent measurements

**Consistent measurements** agree with each other, according to a criterion to be defined later. Inconsistencies can arise for any of the following three reasons.

Firstly, even when all of the redundant measurements are *individually* representative of the measurand, random fluctuations may result in *mutual* inconsistencies occurring from sample to sample.

The other two reasons are more serious, in that they entail a *misrepresentation* of the measurand by one or more measurements:

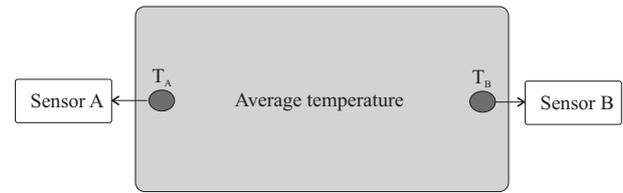


Fig. 3. Sensors A and B both measure the same nominal temperature, but due to local temperature fluctuations, their output  $T_A$  and  $T_B$  may drift apart.

- Each measurement is generated by a SEVA sensor, which should provide detailed and device-specific fault detection. It might reasonably be assumed that a commercial SEVA sensor should be able to detect (say) between 90% and 99.9% of all occurrences of faults within itself, allowing for faults which are inherently difficult to detect and commercial design limitations due to cost/benefit trade-offs. Thus there exists the possibility that a sensor may fail to detect a fault within itself, and so generate an unrepresentative or incorrect measurement. Within a single transmitter handling multiple transducers, consistency checking may be the primary form of validating the raw data, and so inconsistencies may arise more frequently.
- A sensor is only able to measure the *available* value of a process parameter, rather than its ideal or true value [9]. For example (Figure 3), the ‘average’ temperature within a pressure vessel may be the parameter of interest, but in practice only localised temperatures (such as A and B) near the vessel wall are available. A better estimate of the average temperature may be obtained by combining the two measurements. Irrespective of any sensor faults, it is possible for the two measurements to become inconsistent if, for example, a significant temperature gradient develops across the vessel, a phenomenon called available measurement drift.

Of course consistency does not necessarily imply that the measurements are correct, only that they agree with each other. However in practice, just as the true value of the measurand is never known the correctness of a measurement is never known. The fundamental principle of consistency checking, the purpose of which is to identify incorrect measurements to assume that:

- Incorrect measurements are relatively rare.
- Correct measurements are likely to be consistent with one another.

It follows that if one measurement is inconsistent with the rest, it is likely that it is incorrect, as the alternative, that it is correct and all the other measurements are incorrect, is much less probable. Generalising, the principle of majority voting is derived: if a majority of measurements are consistent they are assumed to be correct; any minority of measurements inconsistent with the majority are judged to be incorrect. If there is no majority consensus, then special action will have to be taken. To summarize, the guiding principle is that **inconsistency implies incorrectness**.

### B. The case of two measurements

The detection of inconsistencies between redundant measurements is a well-established academic field of study. Typically, however, the measurements are treated as time series of point values. It is less common for consideration to be given to the uncertainty interval surrounding each measurement, as its magnitude is not usually available.

According to the classical techniques of analytical redundancy [5], given a set of redundant measurements, one or more residual functions are created, each of which is designed to remain 'close' to zero as long as the measurements are consistent. When a fault occurs, a variety of techniques may be applied to determine which sensor (or other plant component) is responsible for the inconsistency. Normally such techniques entail modelling of plant dynamic behaviour and/or sensor fault modes, which can be difficult and/or expensive. Choices must also be made about each decision-making threshold, i.e. the value which, if exceeded by a residual, indicates a significant inconsistency.

The availability of the uncertainty of each measurement provides a richer set of information to work with; on the other hand it adds to the complexity and dimensionality of the problem.

The aim of this paper is not to extend the techniques of analytical redundancy to exploit uncertainty information, but rather to develop an algorithm which does not require any modelling of the measurand  $M$  beyond the simplest possible assumption that the expected value of each measurement is equal to the measurand:

$$E(x_1) = E(x_2) = \dots = E(x_n) = E(M) \quad (3)$$

It is assumed that such an algorithm will be less able than a more elaborate scheme to detect subtle inconsistencies, but could be more widely and readily used in applications which do not justify the expense of detailed modelling. Nor is it proposed here to detect all undiagnosed faults in a SEVA sensor, for example through pattern recognition of time-frequency properties which may be characteristic of specific fault modes, [10]. Rather, the intention is simply to be able to detect faults which cause measurements to be inconsistent with their redundant peers.

Moffat, [11], suggests a method of testing consistency between two measurements  $x_1$  and  $x_2$ , given their uncertainties  $u_1$  and  $u_2$ . Under the hypothesis that the measurements are correct, i.e. they are representative of the same measurand, then the function

$$\phi = x_1 - x_2 \text{ with uncertainty } u_\phi = \sqrt{u_1^2 + u_2^2}$$

should be close to zero. In other words, it is expected that

$$d_{12}^M = \frac{x_1 - x_2}{\sqrt{u_1^2 + u_2^2}} \quad (4)$$

where  $d_{12}^M$  is the Moffat distance, satisfies the following criterion (called here the Moffat criterion):

$$|d_{12}^M| < 1 \quad (5)$$

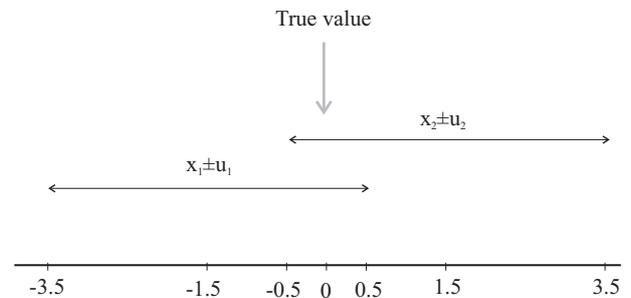


Fig. 4. An example of uncertainty intervals which, though overlapping, are not Moffat consistent.

at the usual, say, 95% probability. The Moffat consistency test can thus be seen as a simple static form of residual function. This definition of consistency is somewhat counter-intuitive, in that uncertainty intervals may overlap and yet still be declared inconsistent, as illustrated in the following example.

Suppose that the true measurand is 0, and that the two sensors generate measurements  $x_1$  and  $x_2$  with a distribution  $\mathcal{N}(0, 1)$ . Then it follows that  $u_1 = u_2 = 1.96$ . Suppose that at a particular instant  $x_1 = -1.5$  and  $x_2 = +1.5$ . Figure 4 illustrates the situation.

It can be seen that the uncertainty intervals overlap the true value and each other, and yet the consistency test fails, for  $\phi = -3$  and  $u_\phi = \sqrt{2} \times 1.96 = 2.77$ . In other words, both measurements are correct, and yet they are not consistent. This is due to the probabilistic nature of the test - the chance of a Type I error is 5%. As the distributions are normal, this can be confirmed analytically:  $x_1$  and  $x_2$  are  $\mathcal{N}(0, 1)$ , so  $\phi = x_1 - x_2$  has distribution  $\mathcal{N}(0, 2)$ . Thus there is a 5% probability that a random sample from this distribution will fall outside the uncertainty interval  $u_\phi = \sqrt{2} \times 1.96$ .

The degree of overlap required for Moffat consistency is maximum when  $u_1 = u_2$ . Suppose  $u_1$  is kept constant and  $u_2$  is increased, then the degree of overlap required for consistency, as a proportion of  $u_1$ , decreases asymptotically to zero. Indeed, as shown in Appendix 1, Moffat consistency ensures that the CBE of the two measurements falls within the uncertainty intervals of each. A logical corollary is that there must be an overlap between the two uncertainty intervals and that the CBE falls within the overlap. A further useful (and intuitively necessary) result is proved in Appendix 2: if  $x_1$  and  $x_2$  are Moffat consistent, then the CBE (n.b. with its reduced uncertainty) is also Moffat consistent with  $x_1$  and  $x_2$ .

The Type I threshold of 5% is presumably acceptable for the analysis of experimental data, the context in which Moffat conceived the test, and of course all the uncertainties themselves are expressed at the 95% probability level. However, for the purposes of on-line monitoring of redundant measurements in an industrial context, this probability is too high, leading to a steady stream of trivial alarms. The alarm frequency could be reduced by modifying the test as follows: use the test criterion  $ku_\phi \leq \phi \leq ku_\phi$  to demonstrate consistency, where  $k$  is a fixed but arbitrary value which controls the probability of a Type I error. The value  $k = \sqrt{2}$  has intuitive appeal, because then two uncertainty intervals of equal magnitude are declared

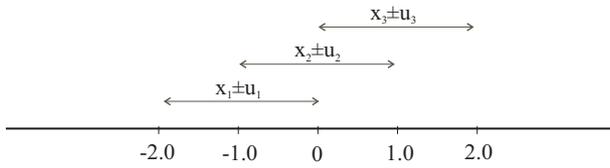


Fig. 5. Three SEVA measurements  $x_1$ ,  $x_2$  and  $x_3$ ,  $x_1 \odot x_2$ ,  $x_2 \odot x_3$ , but  $x_1$  is not consistent with  $x_3$ .

consistent if there is any overlap between them, and has a reduced Type I error of about 0.25%. However, as say  $u_2$  increases relative to  $u_1$ , the counter-intuitive result is derived that two intervals are consistent even if they do not overlap at all, indeed if there is a large gap between them. For example, using  $k = \sqrt{2}$ , all of the following uncertainty intervals pairs are consistent, even where the uncertainty intervals do not overlap:

- $0 \pm 1$  and  $1.99 \pm 1$
- $0 \pm 1$  and  $14 \pm 10$
- $0 \pm 1$  and  $140 \pm 100$

It is concluded that  $k = 1$  is the only acceptable value. There remains the concern that the 5% probability of a Type I error is too high. This will be dealt with in a later section.

## VI. THE FUSION OF 3 OR MORE MEASUREMENTS

For the principle of majority voting to be applicable, a sensor fusion system needs to have 3 or more measurements. When two measurements are found to be inconsistent with each other majority voting cannot resolve the issue.

However, for more than two variables, Moffat's definition of consistency introduces a problem, in that the consistency criterion is not transitive. Using the expression  $x_1 \odot x_2$  to indicate that measurement  $x_1$  with uncertainty  $u_1$  is consistent with measurement  $x_2$  with uncertainty  $u_2$ ,  $\odot$  can be thought of as a binary relation, which is reflexive ( $x_1 \odot x_1$ ) and symmetric ( $x_1 \odot x_2 \Rightarrow x_2 \odot x_1$ ). Unfortunately, the relation is not transitive:  $x_1 \odot x_2$  and  $x_2 \odot x_3 \not\Rightarrow x_1 \odot x_3$ .

For example, as illustrated in Figure 5,  $-1 \pm 1$  is consistent with  $0 \pm 1$  and  $0 \pm 1$  is consistent with  $1 \pm 1$ , while  $-1 \pm 1$  is not consistent with  $1 \pm 1$ .

Also, it has been shown that there is a 5% probability that any two correct measurements of the same measurand are not consistent. Thus, given a set of 3 or more independent measurements that need to be combined, two issues need addressing. First, the maximum subset of mutually consistent measurements must be found and declared the consistent subset. Second, the measurements outside this subset, termed outliers, must be dealt with bearing in mind that inconsistency may be due to probabilistic jitter rather than sensor error.

### A. Consistency checking

It can be shown that the problem of finding the maximum subset of mutually consistent measurements is equivalent to the maximum clique problem in graph theory (illustrated in Figure 6). That is, given a set of nodes and arcs, find the maximum subset of nodes (called the clique) with the property that each node from the subset is connected to every other. If each node

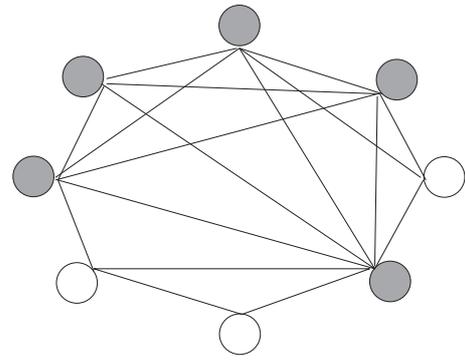


Fig. 6. The maximum clique problem. In this example, the maximum clique consists of the shaded nodes.

### 1 Initialisation

- 1.1 Compute  $d_M^{i,j}$  for all  $i$  and  $j = i + 1, 2, \dots, n$
- 1.2 Build up the measurement graph

### 2 Search for maximum cliques

- 2.1 Start with  $p = n$
- 2.2 While no clique is found and  $p > 1$ 
  - 2.2.1 Build the tree of all possible combinations of  $n$  nodes taken in sets of  $p$
  - 2.2.2 Visit the entire tree and record all cliques
  - 2.2.3  $p = p - 1$ ; goto 2.2

Fig. 7. The exhaustive search algorithm.

is a measurement and each arc is a consistency relation, then this is equivalent to the problem of measurement consistency checking.

The maximum clique problem is known to be NP-hard, so an exhaustive search is required to find the solution. Consider a set of  $n$  SEVA measurements  $x_i$  with uncertainties  $u_i$ ,  $i = 1, 2, \dots, n$ . A prerequisite for the search is the building of the measurement graph. The  $n$  nodes are the values  $x_i$ , while the existence of an arc between  $x_i$  and  $x_j$  is determined by whether they are consistent: i.e. whether  $x_i \odot x_j$ . Let  $p$  be the maximum clique order. The search starts by trying  $p = n$  (i.e. all measurements are consistent) and systematically works down until a clique is found or until  $p = 1$ . When a clique is found the algorithm further searches for any other cliques of the same order. The algorithm is summarised in Figure 7.

### B. Approximation of the maximum clique by linear search

The exhaustive search for the maximum cliques can become extremely onerous as the number of measurements increases and the order of the maximum clique decreases. A method for approximating the maximum clique is proposed which uses overlapping intervals instead of the Moffat criterion to check for consistency. Because this method is linear in the number of measurements it has far less complexity than the exhaustive search. Moffat consistency is ensured within the resulted cliques by a later processing stage called uncertainty expansion, which is described in detail in the next section.

Consider again the set of  $n$  SEVA measurements  $x_i$  with uncertainties  $u_i$ ,  $i = 1, 2, \dots, n$ , and let the uncertainty interval for

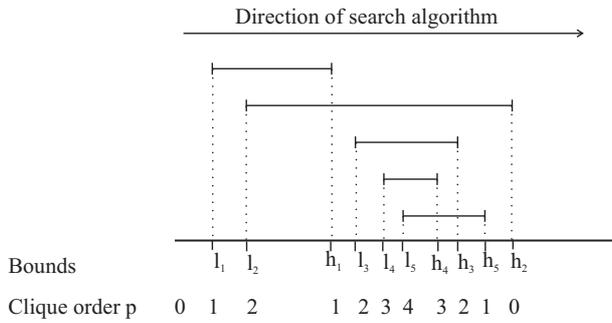


Fig. 8. Illustrative example of overlapping intervals.

the  $i$ th measurement,  $i = 1, 2, \dots, n$ , be  $(l_i, h_i)$ , where  $l = x_i - u_i$  and  $h_i = x_i + u_i$ , the lower and upper bound, respectively. The set of  $n$  measurements can then be described by an ordered bound list containing all  $l_i$  and  $h_i$ . Without loss of generality the  $x_i$  can be assumed ordered so that  $l_1 < l_2 < \dots < l_n$ . Of course the  $h_i$  may occur in any order interleaved through the  $l_i$ , subject only to the constraint that  $h_i > l_i$  (and hence  $h_i > l_k, k = 1..i$ ). The overlapping intervals are readily identified by stepping through the ordered list of bounds. The approximation of the maximum clique(s) is given by the measurements whose uncertainty intervals define the area(s) of maximum overlap.

Figure 8 illustrates the method. The bound list is in this case given by  $l_1 l_2 h_1 l_3 l_4 l_5 h_4 h_3 h_5 h_2$ . The point of maximum overlap involves measurements 2, 3, 4 and 5, which are therefore considered as an approximation of the maximum clique.

The algorithm walks through the bound list in increasing order. When a lower boundary is encountered the corresponding measurement is added to the set of active measurements, whose order  $p$  is thus incremented. When an upper bound is encountered, the corresponding measurement is removed from the set of active measurements whose order is thus decremented. At each stage, if the order of the active measurement set exceeds all previous values, then the active set becomes the new maximum clique. If its order equals that of the current maximum clique then the set is stored as an additional maximum clique. The algorithm is summarised in Figure 9.

### C. Dealing with outliers

Having found or approximated a maximum clique, the obvious next step would be to use it to calculate the CBE using Equations 1 and 2 and to ignore all outliers. This approach has a number of difficulties:

- 1) Given the probabilistic nature of the uncertainty, even if all measurements are correct representations of the measurand, there is only a 95% chance of each pair being consistent. As the number of inputs increases, the probability of all measurements being consistent reduces. For example, with 10 normally distributed measurements of equal variance and mean, there is only a 37.3% chance of all 10 sensors being mutually consistent at any given time.
- 2) If, on average, one measurement is only marginally consistency with the rest, then sample by sample it may reg-

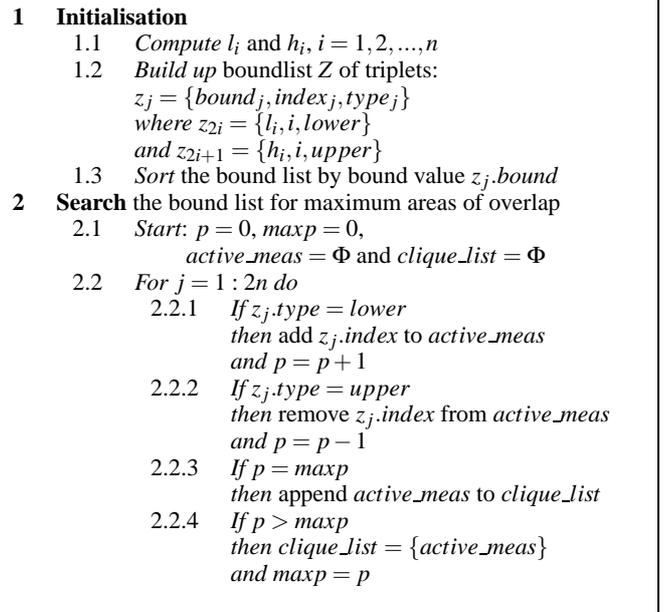


Fig. 9. The approximation of the maximum clique(s) by linear search.

ularly switch between being judged consistent and inconsistent. This will generate undesirable jitter on the CBE.

- 3) It is possible that at any given time there may be more than one maximum clique. For example, with three measurements  $x_1, x_2$  and  $x_3$  such that  $x_1 \odot x_2$  and  $x_2 \odot x_3$ , while  $x_1 \odot x_3$  is not true, then there are two maximum cliques,  $(x_1, x_2)$  and  $(x_2, x_3)$ . It is not obvious which of the maximum cliques to use for calculating the CBE.

A simple strategy can be implemented to tackle these issues. The underlying idea is that any inconsistent measurement can be 'made consistent' by a sufficient increase in its own uncertainty, and that such an increase will cause a reduction in the weight of that measurement in the CBE. This approach is not based on uncertainty theory, but rather is a heuristic approach which has the desirable characteristics of smoothing over probabilistic inconsistency jitter, and providing a smooth reduction of weighting for inconsistent measurements.

In the most general case when there is more than one clique, the measurements are partitioned into two sets:

- the **core set** - the intersection of all the maximum cliques
- the **peripheral set** - the rest of the measurements i.e. those being either in at least one of the maximum cliques, but not in the core set, or those outside any maximum clique

If the maximum cliques were found using exhaustive search, then the mutual Moffat consistency of the measurements inside it is ensured. However, this is not guaranteed to be the case with the linear search, thus, for the core and peripheral sets resulted from the linear search algorithm, additional consistency checking need to be done before the CBE is computed. The maximum Moffat distance  $d_{max}^M$  between pairs of measurements from the core set is computed. If this value is greater than one, then at least one of the measurements pairs is inconsistent. The uncertainty  $u_i$  of all measurements from the core set is then increased to  $u_i' = d_{max}^M \times u_i$ , values which will ensure mutual consistency.

Each measurement from the peripheral set is then considered in turn and the maximum Moffat distance to the measurements in the core set is found. If this value is greater than a specified threshold (for example 3.0), then the measurement is judged to be a true outlier and is ignored. If, however, this distance is less than the specified threshold, then its uncertainty interval is expanded as described to make it consistent with the measurements in the core set. The measurements from the peripheral set thus processed are then merged with those in the core set to obtain the CBE. This technique of uncertainty expansion reduces, but does not eliminate, the influence of the involved measurements on the CBE. In particular, if a measurement slowly drifts into inconsistency with the rest, uncertainty expansion ensures a smooth reduction of influence on the CBE before it is finally labelled as an outlier. This is illustrated in the simulations at the end of the paper.

One circumstance not covered by the above procedure is where there are multiple maximum cliques with no intersection between them. Here the "middle clique" is found as being the maximum clique closest to the mean of the merged values for each maximum clique. The middle clique is then considered to be the core set while the peripheral set contains the remaining measurements.

#### D. The Combination Block

Given a set of  $n$  SEVA measurements  $(x_i, u_i, status_i)$ , the task of the Combination Block is as follows.

- calculate the CBE and its uncertainty using the techniques described above. Normally, the VMV output is set equal to the CBE and the VU to its uncertainty - see below.
- assign the MV status of the output block. As a configuration option, the user can assign the minimum acceptable size of the maximum clique (for example 2 out of 3 or 6 out of 10). If this size is not reached, then the CBE is not used but instead the VMV and VU are projected from past history in the usual way, [1], and the MV status is set to DAZZLED or, if the condition persists, BLIND. If the minimum acceptable size of clique is reached, then the MV status is set to SECURE; COMMON if the sensors are of identical type, otherwise DIVERSE. A further configuration option is the minimum number of CLEAR (or better) consistent measurements required to declare the CBE to be SECURE. If this target is not met, then the CBE is assigned the best status of the consistent measurements (i.e. CLEAR, BLURRED, DAZZLED or BLIND).
- each SEVA measurement is also assigned a consistency flag.

The takes the value 1 if the measurement was found to be in the core or was made consistent with the core by uncertainty expansion, and 0 otherwise. This flag may be used (possibly after further filtering to avoid jitter) to trigger additional diagnostic testing within any SEVA sensors whose measurements were found inconsistent with the majority.

#### E. Exhaustive search vs. linear search approximation

Simulations have been carried out to compare the performance of the two methods for finding the set of mutually consistent SEVA measurements i.e. the exhaustive search for the

	Exhaustive search	Linear search	Theoretical value
<b>3 sensors:</b>			
Mean of CBE	0.002	0.002	0.0
Std of CBE	0.581	0.579	0.577
Mean of unc.	1.136	1.129	1.131
Sets with $k < n$	0	0	
<b>6 sensors:</b>			
Mean of CBE	0.002	0.002	0.0
Std of CBE	0.411	0.410	0.408
Mean of unc.	0.806	0.804	0.800
Sets with $k < n$	0	0	
<b>10 sensors:</b>			
Mean of CBE	0.0	0.0	0.0
Std of CBE	0.320	0.320	0.316
Mean of unc.	0.626	0.624	0.619
Sets with $k < n$	0	0	

TABLE III

NOTE:  $k$  IS THE NUMBER OF CONSISTENT MEASUREMENTS AFTER UNCERTAINTY EXPANSION AND  $n$  IS THE TOTAL NUMBER OF MEASUREMENTS IN THE SET

maximum clique and the approximation of the maximum clique by linear search. In these first studies, fault-free behaviour is considered. It is desirable to have a match between theoretical and simulation results for the following statistics:

- mean of the CBE
- variance in mean of the CBE
- reported uncertainty of the CBE

In addition, it is desirable for the reported uncertainty to be reasonably constant, and for the incidence of reported inconsistencies to be low (as there are no true faults, just random variations).

100000 random sets of 3, 6 and 10 SEVA measurements were generated as follows:

- the true measurand value is 0
- the measurements were randomly generated from a normal distribution with a mean of zero and unit variance. This corresponds to an uncertainty of 1.96.

The theoretical value of the standard deviation of the CBE is then  $\frac{1}{\sqrt{n}}$ . This gives an uncertainty of  $1.96\frac{1}{\sqrt{n}}$ . The means and standard deviations of the reported values of the CBE and its uncertainty has been computed over the 100000 simulations, and they are contrasted with the corresponding expected values in Table III.

In this fault-free simulation, all sensor values were included in the calculation of all the CBEs through the use of the expanded uncertainty weighting technique. By contrast, without this technique, a significant percentage of sets are found to be inconsistent (e.g. in the case of 10 sensors, only 37.3% of the sets were found fully consistent).

At this point it can be concluded that the exhaustive search for the maximum clique and the approximation of the maximum clique with the linear search give very similar results. Given the simplicity and computational efficiency of the lin-

ear search, it may be preferred, certainly for larger numbers of sensors (say  $> 5$ ). Also, the results show a reasonable match between the expected value of the CBE uncertainty, its actual variation, and its reported uncertainty.

## VII. SIMULATION RESULTS

Experiments have been carried out to study the behaviour of the Combination Block when one of the SEVA sensors either signals a fault or gives an incorrect description of the measurand. In view of the results in the previous section, the linear search method was used to generate the following results.

The experiments consisted of simulating the on-line behaviour of 3 SEVA sensors. Two of the SEVA sensors give a correct description of the measurand (as in the previous section), while the third SEVA sensor either signals a fault or generates an incorrect description of the measurand.

In each case, a constant true measurement value of 2 was considered. The simulated faults were as follows:

- **Example 1:** a spike fault - saturation at upper limit occurs at 125s; the fault is permanent; the SEVA sensor detects the fault and first changes the MV status to DAZZLED and then to BLIND.
- **Example 2:** a spike fault - a faulty ramping value is added to the true measurement with the slope of 0.001 units per second; the fault begins at 125s and is permanent; the SEVA sensor detect the fault and changes the MV status to BLURRED.
- **Example 3:** a drift fault - a faulty ramping value is added to the true measurement with the slope of 0.001 units per second; the fault begins at 125s and is permanent; the SEVA sensor does not detect the fault and reports the measured value along with an MV status of CLEAR.

For the cases when the third SEVA sensor gives an incorrect description of the measurand, this was accomplished by ensuring that the VU was of the usual magnitude and the MV Status was CLEAR, while the VMV in fact suffers a drift starting at 125s, with a slope of 0.001 units per second.

The time series of the VMV, the VU and the MV status for a typical sensor (as used in this study) exhibiting fault-free behaviour and generating a correct description of the measurand is given in Figure 10.

Figures 11, 12 and 13 show the outputs of the faulty sensor and the Combination Block for examples 1, 2 and 3 respectively.

In example 1 the third sensor exhibits a permanent saturation fault. Its output is characterised by the usual SEVA response i.e.

- The VMV is projected from past history. In this case the VMV remains reasonably accurate as the process is stationary.
- The MV status changes to DAZZLED and then BLIND when it is deemed that the saturation is permanent.
- The uncertainty increases at a rate learned from past history using conventional SEVA algorithms.

The response of the Combination Block is as follows:

- 1) The MV status of the Combination Block can only remain SECURE COMMON if a configured number of input sensors are CLEAR. In this case the number is three,

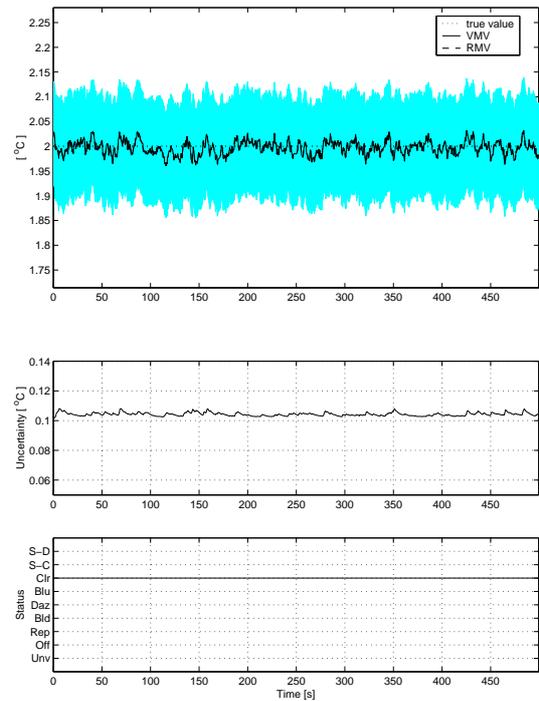


Fig. 10. A typical fault-free sensor with a constant true value.

so as soon as sensor 3 changes MV status the Combination Block output reverts to CLEAR. Note that hysteresis is used to prevent excessive jitter on the combination block MV status.

- 2) The measurements are combined according to their consistency and uncertainty weightings. In both cases the measurement from the faulty sensor remains consistent, but its influence declines rapidly, weighted by the inverse square of its increasing uncertainty. This also results in the rapid increase in the uncertainty of the combined measurement from about 0.06 to 0.075 after the fault.

In example 2 a drift fault occurs in sensor 3, but the sensor detects the fault and attempts to compensate. Thus the raw measurement value (RMV) is seen to drift off quickly, but the SEVA sensor reduces the effect of the fault by internal correction (which still leaves some marginal drift). The sensor declares the measurement BLURRED and increases its uncertainty. In these cases, the slow increase in the VU of the faulty sensor is reflected in a very marginal increase in the uncertainty of the Combination Block. Again, the change in the MV status is also accounted for in the change of MV for the Combination Block.

In examples 1 and 2, since the fault is compensated for inside the SEVA sensor, the reported VMV is a correct representation of the true measurand. Therefore the combination block finds all 3 measurements to be consistent and uses them all to calculate the CBE. The occurrence of the fault is then reflected in the value of the VU for the combination block and in the MV status of the Combination Block output (determined by the change in MV status of the faulty sensor).

Example 3 shows the most important case: when one SEVA sensor does not give a correct representation of the measurand.

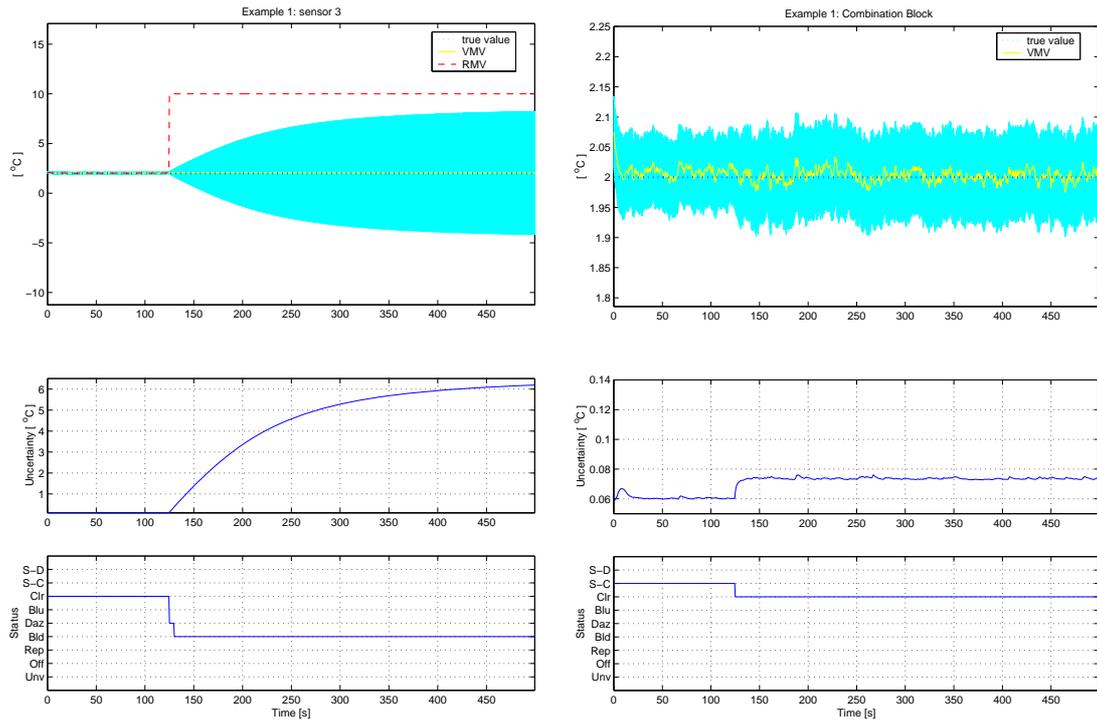


Fig. 11. The faulty sensor and the Combination Block for example 1.

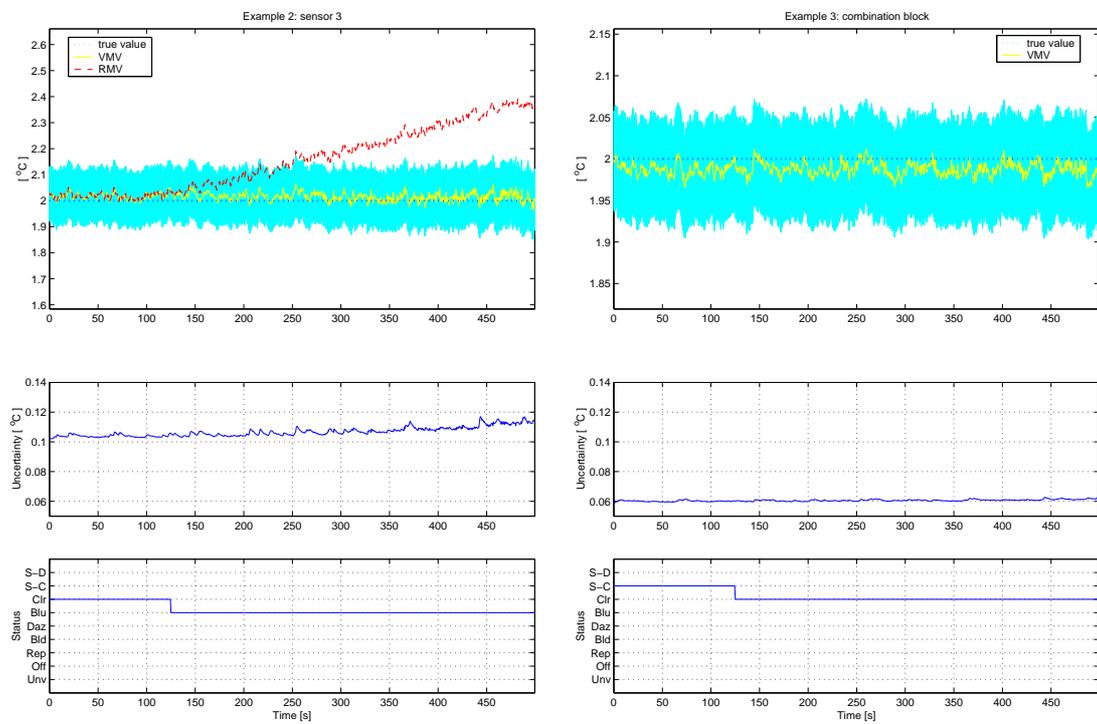


Fig. 12. The faulty sensor and the Combination Block for example 2.

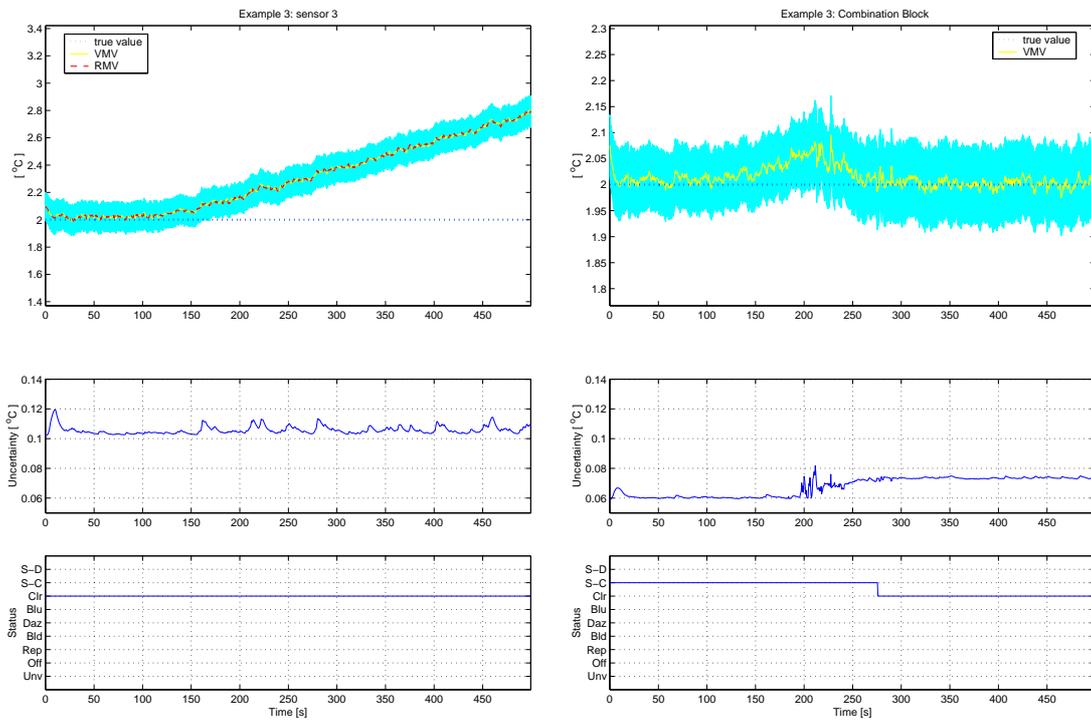


Fig. 13. The faulty sensor and the Combination Block for example 3.

The chain of events is as follows:

- 1) An undetected drift fault begins in sensor 3 at  $t = 125$ s.
- 2) The CBE begins to rise as long as the faulty value remains consistent with the rest.
- 3) From  $t=200$ s to  $t=275$ s sensor 3 becomes increasingly inconsistent with the other two (i.e. its Moffat distance from their combination is between 1 and 3). Accordingly, its influence diminishes, the CBE returns towards the true value and the uncertainty increases as reliance is placed on only two instead of three measurements.
- 4) Finally, at  $t=275$ s sensor 3 is deemed to be persistently inconsistent (Moffat distance  $> 3$ ) and the MV status of the output drops to CLEAR.

The results shown illustrate that the combination block is capable of detecting and compensating for both detected and undetected faults in one of a set of SEVA sensors. The VU of the CBE is increased accordingly to account for faults and, when this is necessary, the faulty sensor is excluded from the calculation of the CBE. The CBE provided by the combination block in these examples remains a correct representation of the measurand, and is smooth, while the MV status is free from jitter.

## VIII. CONCLUSIONS

### APPENDIX 1

*Theorem 1:* For any two measurements  $x_1, x_2 \in \mathfrak{R}$  with uncertainties  $u_1, u_2 \in \mathfrak{R}^+$ , the Combined best estimate (CBE) of the measurand  $\Lambda(x_1, x_2)$  defined as

$$\Lambda(x_1, x_2) = \frac{x_1 u_1^2 + x_2 u_2^2}{u_1^2 + u_2^2}$$

satisfies

$$\Lambda(x_1, x_2) \in [x_1 - u_1, x_1 + u_1] \quad (6)$$

$$\Lambda(x_1, x_2) \in [x_2 - u_2, x_2 + u_2] \quad (7)$$

if the Moffat consistency criterion is satisfied

$$|x_1 - x_2| < \sqrt{u_1^2 + u_2^2}$$

*Proof:*

With no loss of generality and by applying suitable linear transformations (e.g.  $x_i = x_i - x_1$ ,  $u_i = \frac{u_i}{u_1}$ ), it can be assumed that

$$x_1 = 0 \text{ and } u_1 = 1$$

The definition of  $\Lambda$  then becomes

$$\Lambda(x_1, x_2) = \frac{x_2}{u_2^2 + 1}$$

while the Moffat consistency criterion can be rewritten as

$$|x_2| < \sqrt{u_2^2 + 1} \quad (8)$$

and condition (6) becomes

$$|\Lambda| < 1 \quad (9)$$

Define  $q = \sqrt{1 + u_2^2}$  and note that  $q > 1, \forall u_2 \in \mathfrak{R}^+$ . The required proof then becomes: (8)  $\Rightarrow$  (9) and (7).

1) **(8)  $\Rightarrow$  (9):**

If  $|x_2| < \sqrt{u_2^2 + 1} = q$ , this implies  $\frac{|x_2|}{q} < 1$ . Then  $\Lambda$  becomes:

$$|\Lambda| = \frac{|x_2|}{u_2^2 + 1} = \frac{|x_2|}{q^2} = \frac{|x_2|}{q} \frac{1}{q}$$

But  $\frac{|x_2|}{q} < 1$  and  $q > 1$ . Thus, it follows that

$$\Lambda < 1$$

2) **(8)  $\Rightarrow$  (7):**

The condition (7) can be rewritten as

$$x_2 - u_2 < \frac{x_2}{1 + u_2^2} < x_2 + u_2 \quad (10)$$

Multiplying through the above by  $1 + u_2^2$  and manipulating the result

$$x_2 + u_2^2 x_2 - u_2 - u_2^3 < x_2$$

Subtracting  $x_2$  and dividing through by  $u_2$

$$u_2(x_2 - u_2) < 1 \Rightarrow x_2 < u_2 + \frac{1}{u_2} \quad (11)$$

But from Moffat consistency criterion  $x_2 < \sqrt{1 + u_2^2}$ . So, if  $\sqrt{1 + u_2^2} < u_2 + \frac{1}{u_2}$ , then  $1 + u_2^2 < u_2^2 + 2 + \frac{1}{u_2^2}$ , true for  $u_2 \in \mathfrak{R}^+$ . So (11) is proven.

Now take the second inequality of (10). This is trivial if  $x_2 \in \mathfrak{R}^+$ , but this cannot be assumed. Let this inequality be multiplied through by  $1 + u_2^2$ :

$$x_2 < x_2 + u_2^2 x_2 + u_2 + u_2^3$$

Subtract  $x_2$  and divide through by  $u_2$

$$x_2 u_2 + 1 + u_2^2 > 0$$

Which leads to  $x_2 < \frac{1}{u_2} + u_2$ , which is the same as 11, thus proving 10.

*Corollary 1:* If the Moffat consistency criterion is satisfied, then there exists an overlap between  $x_1 \pm u_1$  and  $x_2 \pm u_2$ , and  $\Lambda(x_1, x_2)$  falls within this overlap.

## APPENDIX 2

*Theorem 2:* If  $x_1 \odot x_2$ , then  $\Lambda(x_1, x_2) \odot x_1$  and  $\Lambda(x_1, x_2) \odot x_2$ .

*Proof*

As before, let  $x_1 = 0, u_1 = 1$ .

1)  $\Lambda(x_1, x_2) \odot x_1$  can then be rewritten as

$$\frac{x_2}{\sqrt{u_2^2 + 1}} \pm \frac{u_2^2}{\sqrt{u_2^2 + 1}} \odot 0 \pm 1$$

which is equivalent with

$$\left| \frac{x_2}{\sqrt{u_2^2 + 1}} \right| < \sqrt{1 + \frac{u_2^2}{u_2^2 + 1}} \quad (12)$$

The right-hand side of the inequality can be written

$$\sqrt{\frac{2u_2^2 + 1}{u_2^2 + 1}} > 1 \text{ for } u > 0$$

while the left-hand side is equal to  $\Lambda(x_1, x_2)$  and is less than 1. These prove the inequality 12

2)  $\Lambda(x_1, x_2) \odot x_2$  can be rewritten as:

$$\frac{x_2}{q} \pm \frac{u_2}{q} \odot x_2 \pm u_2$$

where  $q = \sqrt{1 + u_2^2}$ . As before, this is equivalent to

$$|x_2| \left(1 - \frac{1}{q}\right) < \sqrt{u_2^2 \left(1 + \frac{1}{q^2}\right)}$$

But  $u_2^2 = q^2 - 1$ , thus

$$|x_2| < \frac{\sqrt{(q^2 - 1)(q^2 + 1)}}{1 - \frac{1}{q}}$$

Multiply the right-hand side of the inequality with  $q$  and square both sides and obtain:

$$x_2^2 < \frac{(q^1 + 1)(q^2 - 1)}{q - 1}$$

This leads to

$$\frac{x_2^2}{q^2 + 1} < \frac{q + 1}{q - 1}$$

But for  $q > 1, \forall u_2 \in \mathfrak{R}^2$ , the right-hand side is greater than 1, while from 8  $|x_2| < q$ . Then  $\frac{x_2^2}{q^2 + 1} < 1$ . These prove the desired inequality.

## ACKNOWLEDGMENTS

The authors wish to acknowledge the financial support of Invensys PLC. The second author is an EPSRC Advanced Research Fellow.

## REFERENCES

- [1] M. P. Henry and D.W. Clarke, "The self-validating sensor: rationale, definitions and examples," *Control Eng. Practice*, vol. 1, no. 4, pp. 585–610, 1993.
- [2] S.J. Kline and F.A. McClintock, "Describing uncertainties in single sample experiments," *Mech. Eng.*, pp. 3–8, 1853.
- [3] The Quality Management and Statistics Standards Policy Committee(QMS/24), "Vocabulary of metrology part3. guide to the expression of uncertainty in measurement," PD 6461, 1995.
- [4] BSI, "Specification for data quality metrics for intelligent fieldbus instruments," AMT/7, 2000.
- [5] P.M. Frank, "Fault diagnosis in dynamic systems using analytical and knowledge-based redundancy - a survey and some new results," *Automatica*, vol. 26, no. 3, pp. 459–474, 1990.
- [6] J.K Hackett and M. Shah, "Multi-sensor fusion: A perspective," *Proc. 1990 IEEE Int. Conf. Rob. Autom.*, pp. 1324–30, 1990.
- [7] R. C. Luo and M. G. Kay, "A tutorial on multisensor fusion," *16th Annual Conf. of IEEE Industrial Electronics Soc. IECON 90*, pp. 707–722, 1990.
- [8] R. McKendall and M. Mintz, "Using robust statistics for sensor fusion," *SPIE Vol 1198 Sensor Fusion II: Human and Machine Strategies*, pp. 178–191, 1990.

- [9] J.C.-Y. Yang, *Self-Validating Sensors*, Ph.D. thesis, University of Oxford, 1993.
- [10] S. K. Yung, *Signal Processing in Local Sensor validation*, Ph.D. thesis, University of Oxford, 1993.
- [11] R. J. Moffat, "Contributions to the theory of single sample uncertainty analysis," *ASME Journal of Fluid Engineering*, vol. 104, pp. 250–260, 1982.
- [12] Sira Test & Certification Ltd, "The impact of fieldbus and sensor fusion on flow measurements," Tech. Rep., Department of Trade and Industry, NMS Programme for Flow Measurement, 1998.

**Authors biographies**

PLACE  
PHOTO  
HERE