Greenway, P., Deaves, RH., & Bull, DR. (1996). Communications management in decentralised data fusion systems. In *IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems, Washington DC* (pp. 796 - 805). Institute of Electrical and Electronics Engineers (IEEE).
<https://doi.org/10.1109/MFI.1996.572318>

Peer reviewed version

Link to published version (if available):
[10.1109/MFI.1996.572318](https://doi.org/10.1109/MFI.1996.572318)

[Link to publication record in Explore Bristol Research](#)
PDF-document

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Communications Management In Decentralised Data Fusion Systems

P. Greenway & R.H. Deaves

Advanced Information Processing Dept.,
Sowerby Research Centre,
British Aerospace (Operations.) Ltd.
FPC 267, PO Box 5, Filton,
Bristol, BS12 7QW,
United Kingdom.

D.R. Bull

Centre for Communications Research,
University of Bristol,
Merchant Venturer's Building,
Woodland Road,
Bristol, BS8 1UB.
United Kingdom.

Abstract

We compare two communications management algorithms: one based on an information theoretic approach, which builds on and complements previous work in this area, the other based on a non-information based (round robin) approach. We consider different numbers and combinations of targets. The results are presented as process models which relate the delay in achieving reliable identification to the communications to sensor update ratios for both algorithms. Future research areas highlighted by the work are discussed.

1 Introduction

This paper describes our initial investigation into *communications management* within decentralised multi-sensor systems. Such systems consist of a number of distributed nodes each making local decisions about target tracks and identity, and communicating information about those to other nodes. A key question is 'who should say what, to whom, and when?' [1]. The requirement for the management of such communications arises from two main constraints: the *maximum transmission bandwidth* which is dependent on the technological and physical constraints of the communications hardware and communications medium respectively, and the *available bandwidth* which is dependent on the system specifications such as number of processing nodes etc [2].

Here we consider the effect of such constraints on a probabilistic (Bayesian) target identification system. This system is able to accumulate information locally before communicating to other system processing nodes. Hence the overall effect of the bandwidth constraint is to increase the target identification time when compared with a higher bandwidth implementation. This increase in identification time is dependent on the communication management algorithm employed. This relationship

is a research area which has remained, to date, *relatively unexplored* [3]. The problem we have addressed is: can we demonstrate the benefit of managing a finite communications resource in a decentralised data fusion system.

In this initial investigation we compare two communications management algorithms: one is based on an information theoretic approach, the other is based on a round robin (non-information based) approach. The algorithms are compared on the basis of average and maximum increase in identification delays, which are defined here as the time difference between restricted communications when compared with an unrestricted implementation. These results are obtained from a laboratory set up and are presented as *process models* which relate the delay in identification to the communications to sensor update ratio.

Section 2 provides the relevant background information required for the paper. The experimental set-up used in our analysis is described in section 3. Section 4 provides brief details of the communications management algorithms investigated in this paper. Sections 5 and 6 provides details of the experimental method and results respectively. The conclusions of the paper and future research areas are given in section 7. Note that space prohibits the inclusion of a full description of the initial sensing and identity estimation part of this work (although these have deliberately been kept simple); full references to this material are given in the text.

2 Background

2.1 Bayesian Identification Algorithm

Various methods and algorithms exist for the fusion of identity estimates. These include Dempster-Shafer evidential reasoning, artificial neural networks, voting methods etc [4]. For the purpose of this work a decentralised Bayesian algorithm is implemented [5][6].

Here we are concerned with n distinct object types i.e. $\mathbf{x} = [x_1, x_2 \dots x_n]$, and having made k independent observations we wish to establish the posterior distribution according to Bayes rule:

$$p(\mathbf{x}|\mathbf{Z}^k) = p(\mathbf{x}|\mathbf{Z}^{k-1}) \times \frac{\Lambda(Z(k)|\mathbf{x})}{p(Z(k))} \quad (1)$$

- $p(\mathbf{x}|\mathbf{Z}^k)$ - Identity estimates for targets $\mathbf{x}$ after k readings.
- $\Lambda(Z(k)|\mathbf{x})$ - Likelihood from reading $Z(k)$ given targets $\mathbf{x}$.
- $p(Z(k))$ - Prior probability of reading $Z(k)$.

An empirical value for $p(Z(k))$ maybe difficult to obtain, but $p(\mathbf{x}|\mathbf{Z}^k)$ must sum to 1 over the object set $\mathbf{x}$ (since $\mathbf{x}$ is constrained to be mutually exclusive and exhaustive). Therefore the term $1/p(Z(k))$ in equation 1 may be obtained by normalisation. This normalising constant is dependent on the readings $\mathbf{Z}^{(k-1)}$ and $Z(k)$ i.e. $\lambda_{total}(\mathbf{Z}^{(k-1)}, Z(k))$. In addition, this constant may be resolved to give $\lambda_{total}(\mathbf{Z}^{(k-1)}, Z(k)) = \lambda_{comb}(\mathbf{Z}^{(k-1)}, Z(k)) \times \lambda_{model}(Z(k))$ where $\lambda_{comb}(\mathbf{Z}^{(k-1)}, Z(k))$ is the combination rule constant, and $\lambda_{model}(Z(k))$ is the sensor model constant. Hence, equation 1 becomes:

$$p(\mathbf{x}|\mathbf{Z}^k) = p(\mathbf{x}|\mathbf{Z}^{k-1}) \times [\Lambda(Z(k)|\mathbf{x}) \times \lambda_{model}(Z(k))] \times \lambda_{comb}(\mathbf{Z}^{(k-1)}, Z(k)) \quad (2)$$

We define $[\Lambda(Z(k)|\mathbf{x}) \times \lambda_{model}(Z(k))]$ to be the ‘combined normalised likelihood model’ for the sensor, i.e. $p(Z(k)|\mathbf{x})$. The sensor model used for this work has been determined empirically and is given in figure 1, for the object set $\mathbf{x} = [unknown\ small, car, jeep, unknown\ large]$.

Equation 2 is the recursive form of Bayesian identity (ID) estimation for a single sensor and can easily be distributed over multiple sensors, see [5][6], to give a ‘Local ID Estimate’ and ‘Received ID Estimate’, given in equations 3 and 4 respectively:

$$p(\mathbf{x}|\mathbf{Z}^{k(local)}) = p(\mathbf{x}|\mathbf{Z}^{k(local)-1}) \times p(Z(k(local))|\mathbf{x}) \times \lambda_{comb}(\mathbf{Z}^{(k(local)-1)}, Z(k(local))) \quad (3)$$

where $k(local)$ represents the number of local independent observations from the node’s own sensor, and $p(Z(k(local))|\mathbf{x})$ is obtained from the ‘combined normalised likelihood model’ for that sensor; and

$$p(\mathbf{x}|\mathbf{Z}^{k(received)}) = p(\mathbf{x}|\mathbf{Z}^{k(received)-1}) \times p(\mathbf{x}|\mathbf{Z}^l) \times \lambda_{comb}(\mathbf{Z}^{(k(received)-1)}, \mathbf{Z}^l) \quad (4)$$

Piecewise Linear Combined Normalised Likelihood Model

Figure 1: Sensor Model. This gives the probability that any given sensor reading $Z(k)$ would have been produced by observing an object of each type. Note that the car is less easily identified than the jeep. See [11] for more details.

where $k(received)$ represents the number of received independent observations, from other nodes in the system, and l represents the number of independent observations accumulated at the communicating node since the previous communication. It should be noted that $p(\mathbf{x}|\mathbf{Z}^{l=0})$ is initialised to the least informative, maximum entropy value.

These are combined to give a ‘Global ID Estimate’:

$$p(\mathbf{x}|\mathbf{Z}^{k(global)}) = p(\mathbf{x}|\mathbf{Z}^{k(local)}) \times p(\mathbf{x}|\mathbf{Z}^{k(received)}) \times p(\mathbf{x}) \times \lambda_{comb}(\mathbf{Z}^{k(local)}, \mathbf{Z}^{k(received)}, p(\mathbf{x})) \quad (5)$$

where $p(\mathbf{x})$ is the prior unconditional probability vector for targets $\mathbf{x}$ and $p(\mathbf{x}|\mathbf{Z}^{k(local)=0})$ and $p(\mathbf{x}|\mathbf{Z}^{k(received)=0})$ are initialised to the least informative, maximum entropy value.

2.2 Entropy

Entropy (or average information) has been widely used in the field of communications, where Hartley and Shannon defined its use as an information metric [7][2].

The entropy of a (discrete) posterior distribution after k observations is defined as in [5]:

$$h(k) \equiv - \sum_{j=1}^n p(x_j|\mathbf{Z}^k) \log(p(x_j|\mathbf{Z}^k)) \quad (6)$$

where the dimensionless unit of $h(k)$ is the *nat* when the logarithm base is e . Entropy provides a measure of uncertainty, where a low entropy value indicates low uncertainty of the target type and a high value indicates high uncertainty of the target type.

Figure 2: Decentralised and Centralised Systems.

2.3 Decentralised Multi-Sensor Systems

Here we provide a brief introduction to decentralised systems. Figures 2(a) to (e) show a range of levels of decentralisation and connection between the group of sensors and processing nodes. Some of the main advantages of decentralised systems are:

- **Survivability:** The decentralised system provides graceful degradation in system performance as a communication link, sensor, or processing node fails.
- **Extensability:** The decentralised system has the properties of being scalable (sensors are easily added or taken away), modular (much of the code on the processing node is identical, as is the hardware), and flexible (the number of nodes and how they are connected can easily be varied).

One of the major disadvantages of decentralised systems for practical applications is that a larger number of processors may be required than for a centralised system. Generally the architecture chosen will depend on many factors which include the application, performance and cost.

2.4 Communications Within Decentralised Systems

Figure 3(a) 'Circular Connection' shows the information flow that occurs in a fully connected decentralised system after each of the three sensing nodes has made one observation. For this system six items of data,

Figure 3: Fully and Non-fully Connected Topologies.

$d1 \dots d6$, are communicated. This circular connection can be mapped onto the line connection shown in figure 3(a). Here the sensing node (N) has been modified slightly (N') to deal with this 'broadcast' communications. It should be noted that the number of items of data communicated $d1 \dots d6$ has remained unchanged.

S. Grime and H.F. Durrant-Whyte [3] developed a partially connected decentralised algorithm and topology, which had a reduced communications requirement¹. This is shown in figure 3(b). Here the items of data to be communicated are combined, hence the number of items of data that need to be communicated has reduced from 6 to 4. The research detailed in this paper builds on (and complements) this work by addressing the problem of 'given a limited communications bandwidth i.e. a bandwidth that does not allow all the data to be communicated, what data items should be communicated to achieve some given performance criteria?'

2.5 Dealing With Bandwidth Constraint

Communications systems comprise three individual components, a transmitter, a receiver and a transmission medium. In electronic communications systems the *maximum transmission bandwidth* is dependent on all three components. In addition, the *available bandwidth* may be constrained due to system requirements, such as frequency sharing, number of targets etc. This bandwidth constraint can be overcome by a variety of techniques,

¹S. Utete and H.F. Durrant-Whyte have further developed this work [8].

Figure 4: Experimental Set-up.

see [2]. The system used in this paper uses a *selected data transmission* technique.

Here items from the data set are selected so as to fill the available bandwidth. This is achieved by not transmitting certain items of data. The subset to be transmitted is determined by some decision algorithm i.e. communications management. It should be noted that a different subset of data may be transmitted each time. The advantage of this method is that it is scalable, and may be computationally in-expensive. The main disadvantage of this method is that items of data will be delayed and may be discarded completely.

3 Experimental Set-up

Figure 4 represents the experimental set-up for the work described in this paper. The Sowerby Research Centre (SRC) data fusion test bed ('DFTB Hardware') is a real-time multi-sensor system based on a fully decentralised fully connected architecture, and is housed in an SRC laboratory. A full description of the DFTB is outside the scope of this paper, but is summarised here with appropriate references. The DFTB sensors comprise: CCD cameras, optical barriers, ultrasonic and

infra-red sensors, see [9]. Two decentralised sub-systems provide tracking and identification information on the DFTB. The target detection and tracking algorithms, based on the decentralised Kalman filter, are described in [10], with probabilistic and evidential reasoning methods used for the identification algorithm, see [11]. These algorithms are programmed in OCCAM and run on transputer processors. The DFTB has been used to investigate a variety of data fusion issues, which include information fusion and sensor management [12] [13].

Here the DFTB provides tracking, data association and identification information for the 'Identification Processing Nodes'. In addition, the DFTB sensors may be simulated in order to speed up the development of communications management process models. These can then be verified using the 'DFTB Sensors'. The working area of the DFTB is 7m by 8m, with the field of view (FOV) of Sensor B being within the FOV of Sensor A. The targets $[t_1, t_2 \dots t_N]$ are provided by toy remote control vehicles. Data relating to the identity of targets within the common view of sensors A and B are communicated between the 'Identification Processing Nodes' via the ethernet using a time division multiple access (TDMA) protocol [2]. The frequency of the TDMA can be set in software. For the round robin and information based communications management algorithms considered a *restriction* is placed on the TDMA, that is, data on only one target can be communicated at each time slot. In addition, an *unrestricted* algorithm is employed which allows data on all targets to be communicated at each time slot; this provides a reference with which the round robin and information based communications management algorithms may be compared.

The 'Identification Processing Nodes' in effect implement equations 3, 4 and 5. Full details of this implementation will be reported elsewhere, but are omitted here for brevity.

An unrestricted, round robin, and information based communications algorithm are each 'run' simultaneously. This fact, along with choosing a flat combined normalised likelihood model, see figure 1, provides consistent results for a given scenario. Therefore, the variance in the results obtained for different 'runs' is negligible, hence multiple 'runs' to determine statistical significance are not required.

4 Communications Management

4.1 Round Robin Algorithm

A round robin method is employed as the non-information based communications management algorithm for this work. The algorithm is initialised by establishing a list of targets corresponding to those in common view of the sensors. The order of the list is dependent on

the initial location of the targets. Hence, for N targets there are $N!$ possible sequences. The order of communications then begins with the first item in the list and is incremented after each communication. When communication on the last data item in the list has occurred the communications ‘wraps around’ back to the first item. Therefore the sequence $[t_1, t_3, t_2]$ produces the following communications: $t_1, t_3, t_2, t_1, t_3, t_2 \dots$. The communications rule is given by:

$$c = c + 1 \quad \text{if } c > N \text{ then } c = 1$$

where c is the index of the target whose identity information is to be transmitted. The total ‘Global ID Estimate’ entropy on all targets, H_{global} , at the receiver’s node after the communication is given by:

$$H_{global} = H_{local} + H_{received} + h_r^c \quad (7)$$

where

H_{local}	-	Receiver’s ‘Local ID Estimate’ entropy
$H_{received}$	-	Receiver’s ‘Received ID Estimate’ entropy
h_r^c	-	The entropy contribution to H_{global} of the communication of data on target c

An important characteristic of this method is: *The round robin communications management algorithm does not take into account the information that it is communicating. Hence we refer to this algorithm as ‘dumb’ or ‘non-intelligent’.*

4.2 Information Based Algorithm

The information based communications management algorithm uses the concept of ‘maximising information’. For this strategy the transmitter decides which communicated item of data, with entropy h_r^c , will minimise the overall entropy at the receiver over one communication. This provides other decision making processes associated with the system, for example sensor management, with the maximum information on which to base their choice of action. In order to accomplish this the transmitter needs to be able to accurately *predict* the receiver’s current ‘Global ID Estimate’. This comprises two parts: the receiver’s ‘Received ID Estimate’, and the receiver’s ‘Local ID Estimate’. For our set-up the receiver’s ‘Received ID Estimate’ can be accurately predicted by the transmitter, since it is the only node that transmits to the receiver. The receiver’s ‘Local ID Estimate’ can also be predicted by the transmitter, primarily using the transmitter’s ‘Received ID Estimate’, or (if this is not available) using the transmitter’s own ‘Local ID Estimate’ and a model of the receiving node’s sensor. It should be noted that for this work the receiver’s ‘Local ID Estimate’ can be accurately predicted due to choice of the sensor model (see the ‘flat’ profile of this model as shown

in figure 1). Therefore, the item of data to be communicated is determined from:

$$c = \arg \min_c [\hat{H}_{local} + \hat{H}_{received} + h_r^c] \quad (8)$$

where $\hat{H}_{local}$ is the predicted receiver’s ‘Local ID Estimate’ entropy contribution to H_{global} , and $\hat{H}_{received}$ is the predicted entropy of the receiver’s ‘Received ID Estimate’; here $\hat{H}_{received} \equiv H_{received}$.

The ‘Global ID Estimate’ entropy at the receiver node after the communication is given by:

$$H_{global} = H_{local} + H_{received} + h_r^c \quad (9)$$

where h_r^c is the entropy contribution to H_{global} of the communication of data on target c (*transmit*). If we have accurate prediction (as we do for this investigation) we get $\hat{H}_{local} \approx H_{local}$.

An important characteristic of this method is: *The information based communications management algorithm does take into account the information that it is communicating. Hence we refer to this algorithm as ‘intelligent’.*

5 Experimental Method

5.1 Scenarios

For this work, three key quantities are varied. These are:

1. The ‘communications to sensor update ratio’ i.e. $u = \tau_{communication} / \tau_{sensor}$ where $\tau_{communication}$ is the time interval between two consecutive communications, and τ_{sensor} is the time interval between two consecutive local sensor updates.
2. The ‘total number of targets’ i.e. N ; we set $N = 2, 3$.
3. The ‘mix of targets’ which is defined as the fraction of target types i.e. $m \equiv no_{car} / (no_{car} + no_{jeep})$, where no_{car} is the number of car targets and no_{jeep} is the number of jeep targets in the scenario being investigated.

For each value of N we test every value m , and in each experiment we vary u between 2 and 24.

All the relevant target positions for $N = 2$ are shown in figures 5 (a) to (c), while all relevant target positions for $N = 3$ are shown in figures 5 (d) to (g).

5.2 Production of Results

The results were produced using the following procedure:

Figure 5: Two and Three Target Scenarios.

1. Consider a number of targets being viewed by the sensors, see figure 6(a). Entropy profiles for a given communications to sensor update ratio are represented in figure 6(b). These entropy profiles are analysed by determining the time difference between the round robin (RR) and unrestricted (UN) algorithms, and the entropy (EN) based and UN algorithms, for 40 equally spaced entropy points. This analysis provides two time values: $R_{\{i,j\}}$ and $I_{\{i,j\}}$, see figure 6(b), where i is the target number and j is the entropy value of interest.

The values AR and AI , the average round robin delay and average entropy delay respectively, are given by:

$$AR = \arg \max_i \left[\frac{1}{40 \times s} \times \sum_{j=1}^{j=40} R_{\{i,j\}} \right] \quad (10)$$

$$AI = \arg \max_i \left[\frac{1}{40 \times s} \sum_{j=1}^{j=40} I_{\{i,j\}} \right] \quad (11)$$

where s is the sensor update time interval. In addition, the values MR and MI , the maximum round robin delay and maximum entropy delay respectively, are given by:

Figure 6: Results Analysis.

$$MR = \arg \max_{i,j} \left[\frac{1}{s} \times R_{\{i,j\}} \right] \quad (12)$$

$$MI = \arg \max_{i,j} \left[\frac{1}{s} \times I_{\{i,j\}} \right] \quad (13)$$

Here we are interested in *average* and *maximum* identification delays since, the requirements in such systems are defined in terms of these metrics. It should be noted that these delay values are dimensionless and are represented as the number of sensor updates, τ_{sensor} .

2. Values are gathered for each possible round robin communications pointer initialisation position. The largest values of AR , AI , MR , and MI are selected as *worst case* values. We use the worst case result since we cannot easily predict the target order of the round robin communications management algorithm.
3. Plots are then produced, using these worst case maximum and average values, for different communications to sensor update ratios i.e. $u = 2, 4, 6, 8, 10, 12, 14$ and 24 .
4. The procedure is then repeated for different target combinations i.e. $m = 0.0, 0.5$, and 1.0 for the two target scenario, and $m = 0.0, 0.33, 0.66$, and 1.0 for the three target scenario, see figure 5.

6 Results

6.1 Two Targets

Plots of worst case average and maximum delay against communication to sensor update ratio are given

(a) Average Delay Summary

(b) Maximum Delay Summary

Figure 7: Average and Maximum Delay: Two Target Summary.

in figure 7(a) and (b) respectively. Here a comparison is shown between the results obtained for all the two target scenarios investigated.

6.2 Three Targets

Comparison plots of the results obtained for the three target scenario are given in figures 8 and 9. Here the average and maximum delay results are compared between the scenarios where all the targets are the same (see (a) and (b)) and the scenarios where the targets are different (see (c) to (f)).

Figure 8: Average Worst Case Delay: Three Target Summary.

6.3 Repeatability of Results

The repeatability of the results obtained is dependent on many factors, these include:

- The processing speed of the 'Identification Processing Node' code on the sun workstations, see section 3. This depends on the number of users and processes running on the suns, which can be fairly well controlled.
- The availability of communications bandwidth on the ethernet. This depends on the number of system users and the tasks they are performing. This is not so easily controlled and is time variant.

Taking these factors into account, under typical conditions, the maximum delay produced can vary by up to 1 sensor update intervals and the average delay results by approximately 0.05 sensor update interval. It should be noted that these repeatability values have been derived empirically.

6.4 Development of Process Models

The performance of the round robin communications management algorithm is determined by the least discernible target being viewed. Hence, for the scenarios investigated where a target of type car is present the round robin worst case average and maximum delays follow those of the scenarios where all the targets are of type car. This situation arises since the round robin algorithm provides all targets with an equal share of the communications transmission slots available i.e. the algorithm operates in a non-intelligent fashion. This does not apply to the entropy based algorithm.

The two and three target process models for both communications algorithms, as investigated here, are given in figures 7, 8 and 9 respectively. These show that at low values of communications to sensor update ratio, u , the performance of both the communications algorithms are comparable. This occurs since a low u value indicates high communications frequency which allows high information exchange. In addition, the process models show that at higher values of u (in the region of interest) the information based communications management algorithm out performs the round robin algorithm. At even higher values of u the communications frequency is so low that the targets are identified by local sensor readings only, hence the communications management algorithms become insignificant. It is also evident that as the number of targets, N , is increased the performance margin between the communications algorithms increases. In addition, the process models show that as the proportion of targets, m , that are relatively difficult to discern (car in our investigations) are increased, and $m \neq 0$, the performance margin between the communications algorithms decreases. These observations are

Figure 9: Maximum Worst Case Delay: Three Target Summary.

true for both our performance metrics i.e. average and maximum delay.

7 Conclusions

This *initial* investigation into communications management was based on a fully decentralised Bayesian identification system. Two communications management algorithms have been considered, one based on an information theoretic approach the other based on a round robin algorithm (non-information based). Process models for two and three target scenarios were developed from which the following conclusions were derived:

1. At low values of communications to sensor update ratio, u , the performance of both management algorithms are comparable. At higher values of u (of relevance) the information based algorithm outperforms the round robin algorithm. At even higher values of u the communications management algorithms become insignificant.
2. As the total number of targets, N , is increased the performance margin between the algorithms increases.
3. As the proportion of targets, m , that are relatively difficult to discern is increased, and $m \neq 0$, the performance margin between the algorithms decreases.

These conclusions hold true for both the performance metrics i.e. average and maximum identification delay. However, it should be noted that further investigations to broaden the scope of the scenarios investigated i.e. increase the value of N , resolution of m etc, would be required to generalise these process model conclusions. The other factor that affects these conclusions is the *sensor model*, l , this variable has not been considered here.

In order to gain further understanding and evaluation of communications management a number of research areas need investigating. These include², in no particular order:

1. The effects on the conclusions derived of increasing the total number of targets N .
2. The effects on the conclusions derived of increasing the resolution of the 'mix' proportions.
3. Investigating the effect of changing the sensor model, currently as in figure 1.
4. Investigating the effect of increasing the size of the object set $\mathbf{x}$, currently 4.
5. Investigating the effect of increasing the number of nodes in the system, currently 2.
6. Investigating the effect of the performance of the communications algorithms when coupled with other management strategies e.g. sensor management.
7. Investigating the effect of latency on communications management.
8. Investigating the effect that conflicting nodal information has on communications management.
9. Investigating the effect data association has on the communications management.
10. Investigating the effect prediction has on the performance of information based communications management.
11. Investigating the effect of communications protocols other than TDMA. In addition, the results obtained could be used to develop our own protocols based on adaptive transmission communications management e.g. only communicate data when its information content has obtained a certain value.

Although many research areas have been identified and are certainly worth pursuing, these are 'vertical research areas' i.e. building on the results obtained from the investigation into a Bayesian identification system. It may be prudent at this time to expand our knowledge in 'horizontal research areas' i.e. using the results obtained to see if they are compatible with the needs of other sub-systems, e.g. tracking systems.

8 Acknowledgements

The authors would like to thank Colin Noonan, British Aerospace PLC. for his useful contribution and support for this work.

References

- [1] Edward Waltz and James Llinas. *Multi-Sensor Data Fusion*. Artech House, Boston, MA., 1991.
- [2] H. Taub and D. Schilling. *Principles of Communications Systems*. McGraw Hill, 1987.
- [3] S.H. Grime. *Communications Within Decentralised Sensing Systems*. PhD thesis, Oxford University, 1993.
- [4] L.A. Klein. *Sensor and Data Fusion Concepts and Applications*. SPIE Press, 1993.
- [5] J. Manyika and H.F. Durrant-Whyte. *Data Fusion and Sensor Management a Decentralised Information Theoretic Approach*. Ellis Horwood, 1994.

²This is by no means an exhaustive list!

- [6] B.S.Y. Rao. *Data Fusion Methods for Decentralised Sensing Systems*. PhD thesis, Oxford University, 1991.
- [7] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley, 1991.
- [8] S. Utete and H.F. Durrant-Whyte. Reliability in decentralised data fusion networks. *IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, pages 215–221, October 1994.
- [9] R.H. Deaves. *Evaluation Of The Sowerby Research Centre SKIDS Sensors*. Bristol University M.Sc Thesis, Information Engineering, 1993.
- [10] B.S.Y. Rao, H.F. Durrant-Whyte, and J.A. Sheen. A fully decentralized multi-sensor system for tracking and surveillance. *The International Journal of Robotics Research, Massachusetts Institute of Technology*, 12(1):20–44, February 1993.
- [11] R.H. Deaves and P. Greenway. Experimental Evaluation of Distributed Identity Fusion. *SPIE International Symposium on Photonics for Industrial Applications, Sensor Fusion 7, Boston*, 2355, November 1994.
- [12] P. Greenway and R. Deaves. An Information Filter for Decentralised Data Fusion and Sensor Management. *SPIE, Orlando*, April 1994.
- [13] P. Greenway and R.H. Deaves. Sensor Management Using the Decentralised Kalman Filter. *SPIE International Symposium on Photonics for Industrial Applications, Sensor Fusion 7, Boston*, 2355, November 1994.