Copy No.________

Defence Research and Development Canada
Recherche et développement pour la défense Canada

DEFENCE R&D DÉFENSE

The Simultaneous Availability of Competing Contacts

D. J. Peters

Defence R&D Canada – Atlantic

Technical Memorandum
DRDC Atlantic TM 2005-260
December 2005

Canada

This page intentionally left blank.

The Simultaneous Availability of Competing Contacts

D. J. Peters

Defence R&D Canada – Atlantic

Technical Memorandum

DRDC Atlantic TM 2005-260

December 2005

Author

D. J. Peters

Approved by

J. S. Kennedy

Head, Maritime Information and Combat Systems Section

Approved for release by

Kirk Foster

DRP Chair

Abstract

Algorithms for making association decisions in data fusion systems assume the simultaneous availability of the contacts that compete with each other for association with any given track. A data fusion system that processes all contacts on arrival is flawed, in principle, for it assumes that potentially competing contacts will always be sent in the same message by the data source in question. This assumption leads to errors in association that would not otherwise arise. Such errors have been observed in the COMDAT TD project, thus motivating a more detailed examination of the incidence of such errors.

An experiment was conducted, using simulated aerial tracks and simulated radar contacts, to examine the effects of the fragmentation of contacts (that is, the distribution of contacts among separate messages) on the performance of a single-scan data association algorithm. The odds of correct association are significantly degraded, by more than a factor of two, by highly fragmented data. However, a milder degree of fragmentation, roughly representative of the tracking radar of a HALIFAX class frigate, causes a much milder degradation in the odds of correct association. These results do not lead to any recommended changes to the COMDAT system. However, the fragmentation of contacts must be considered, in general, in the design of any automated data fusion system.

Résumé

Les algorithmes permettant de prendre des décisions d'association dans les systèmes de fusion de données reposent sur l'hypothèse de la disponibilité simultanée des contacts qui entrent en concurrence pour une association avec une piste donnée. Un système de fusion de données qui traite tous les contacts à l'arrivée est en principe imparfait, du fait qu'il suppose que des contacts potentiellement concurrents sont toujours envoyés dans le même message par la source de données en question. Cette hypothèse entraîne des erreurs d'association qui, autrement, ne se produiraient pas. De telles erreurs ont été observées dans le projet de démonstration de la technologie d'aide aux décisions de commandement (COMDAT), ce qui justifie un examen plus poussé de leur incidence.

On a mené une expérience au moyen de pistes aériennes simulées et de contacts radars simulés pour examiner les effets de la fragmentation des contacts (c'est-à-dire la répartition des contacts parmi différents messages) sur la performance d'un algorithme d'association de données à un seul balayage. Les chances d'association correcte sont sérieusement minées, d'un facteur de plus de deux, par des données hautement fragmentées. Un degré moindre de fragmentation, plus ou moins représentatif du radar de poursuite d'une frégate de la classe HALIFAX, cause cependant une dégradation beaucoup plus faible des chances d'association correcte. Ces résultats mènent à la recommandation d’apporter aucune modification au système COMDAT. En règle générale, il faut cependant tenir compte de la fragmentation des contacts dans la conception de tout système automatisé de fusion de données.

This page intentionally left blank.

Executive summary

Introduction

Data fusion systems must decide which incoming contacts belong to which established tracks. Methods for making this decision assume that all contacts that compete for association with a given track will be available at the same time, so that they can be compared with each other. A data fusion system that processes all contacts on arrival is flawed, in principle, for it assumes that potentially competing contacts will always be sent in the same message by the data source in question. This assumption leads to errors in association that would not otherwise arise. Such errors have been observed in the COMDAT TD project, thus motivating a more detailed examination of the incidence of such errors.

The effects of the distribution of contacts among separate messages (here called “fragmentation”) on the performance of a single-scan data association algorithm were investigated by way of an experiment involving simulated aerial tracks and simulated radar contacts.

Results

Statistical tests were conducted to compare the rate of correct association for each of two levels of message fragmentation with the rate of correct association in the case of no fragmentation. Separate tests were conducted for each of three levels of track density.

The milder level of message fragmentation, roughly similar to that imposed by the tracking radar of the HALIFAX class frigate, caused only a mild degradation in the rate of correct association. These results do not lead to any recommended changes to the COMDAT system. The more severe level of fragmentation caused a significant degradation of the odds of correct association, by a factor greater than two.

Significance

The present results do not map straightforwardly onto the degree of degradation of the rate of correct association that is to be expected in any real operational system. However, it has been demonstrated that the processing of all contacts on arrival may be a problem, depending on the messaging behaviour of the data source in question. In any future data fusion system, the potential fragmentation of contact data must be considered. A simple buffer management module, inserted between the data source and the main part of the data fusion system in order to regroup incoming contacts, can eliminate the problem in most cases. The cost of such a solution is a slight delay in the processing of contacts.

Peters, D. J. 2005. The Simultaneous Availability of Competing Contacts. DRDC Atlantic TM 2005-260. Defence R&D Canada – Atlantic.

Sommaire

Introduction

Les systèmes de fusion de données doivent associer les contacts d'arrivée aux pistes établies pertinentes. Les méthodes permettant de prendre ce type de décision supposent que tous les contacts en concurrence pour une association avec une piste donnée sont disponibles en même temps, de sorte qu'ils puissent être comparés entre eux. Un système de fusion de données qui traite tous les contacts à l'arrivée est en principe imparfait, du fait qu'il suppose que des contacts potentiellement concurrents sont toujours envoyés dans le même message par la source de données en question. Cette hypothèse entraîne des erreurs d'association qui, autrement, ne se produiraient pas. De telles erreurs ont été observées dans le projet de démonstration de technologie COMDAT, ce qui justifie un examen plus poussé de leur incidence.

Les effets de la répartition des contacts parmi différents messages (ce qu'on appelle « fragmentation ») sur la performance d'un algorithme d'association de données à un seul balayage ont été étudiés dans le cadre d'une expérience comportant des pistes aériennes simulées et des contacts radars simulés.

Résultats

Des essais statistiques ont été menés pour comparer le taux d'association correcte de chacun des deux niveaux de fragmentation des messages au taux d'association correcte en l'absence de fragmentation. Des essais distincts ont été menés dans le cas de chacun des trois niveaux de densité des pistes.

Le degré moindre de fragmentation des messages, plus ou moins semblable à ce qui est imposé par le radar de poursuite de la frégate de la classe HALIFAX, n'a causé qu'une légère dégradation du taux d'association correcte. Ces résultats mènent à la recommandation d’apporter aucune modification au système COMDAT. Le degré plus élevé de fragmentation a causé une importante dégradation des chances d'association correcte par un facteur de plus de deux.

Portée

Les présents résultats ne correspondent pas de façon pure et simple au degré de dégradation du taux d'association correcte auquel on s'attend dans le cas de tout système opérationnel réel. Il a cependant été démontré que le traitement de tous les contacts à l'arrivée peut poser un problème, selon la façon dont la source de données en question réagit aux messages. Dans le cas de tout futur système de fusion de données, il faut tenir compte de la fragmentation potentielle des données de contact. Un simple module de gestion tampon, inséré entre la source de données et la partie principale du système de fusion de données pour regrouper les contacts d'arrivée, règle le problème dans la plupart des cas. L’inconvénient d'une telle solution est un petit retard dans le traitement des contacts.

Peters, D. J. 2005. Disponibilité Simultanée de Contacts Concurrents. DRDC Atlantic TM 2005-260. Defence R&D Canada – Atlantic.

Table of contents

List of figures

List of tables

This page intentionally left blank.

1. Introduction

Data fusion systems must decide which incoming contacts belong to which established tracks. Methods for making this decision assume that all contacts that compete for association with a given track will be available at the same time, so that they can be compared with each other. But contacts do not always arrive in convenient packages, where those that may compete with each other are grouped together. The word *fragmentation* will be used to refer to the distribution of potentially competing contacts among separate messages.

If a data fusion system is designed to process all contacts on arrival, errors in association will sometimes arise due to the absence of other contacts that ought to have been considered along with the present ones. This paper aims to assess the impact of contact fragmentation on such a system.

1.1 Association and simultaneity

Traditional "level one" data fusion combines two essentially distinct processes: tracking, which is the updating of state estimates in response to new data, and association, which is the decision of which of the new data should be applied to which state estimate. In its typical military applications, the word *tracking* can be taken literally, for the problem in question is that of generating the best possible estimates of the position, behaviour, and identity of each of several platforms of tactical relevance.

The problem of association is to decide which contacts belong to which tracks[1], and which contacts belong to no existing track at all (and thus can be used to initiate a new track). There are many approaches to this problem [1, 2]. We might make an immediate, firm decision based on rules that are related to the distance between the contacts' measured positions and the target positions that are expected, according to the track data, at that time. Or we might defer the association decision, perhaps by allowing our entire tactical picture to be divided into several different hypotheses, according to several different plausible association decisions, hoping that some of the hypotheses will be rendered untenable by later data. Or we might perform tracking according to a weighted combination of the plausible association decisions, and avoid the need to make any specific assignment.

All of these approaches to the problem of association have in common the assumption that the contacts that compete for association with any given track will be available at the same time, in order for the system to perform its comparison of the various

[1] For the purposes of this document, the word *track* will refer to the package of data that contains the fusion system's estimate of the position, behaviour, and attributes of an alleged target, and the word *contact* will refer to a package of data that is available for use in updating a track. That is, the system's input consists of "contacts", while its output consists of "tracks". (In some cases, these "contacts" will actually be tracks in their own right. But in the present context, a more general treatment of the difference between a contact and a track would function only as a distraction.)

association options. However, the simultaneous availability of contacts depends on processes that are outside the mandate of the association and tracking aspects of the fusion system. We are dealing with a question of input management.

A sensing device, of any kind, is more likely to be designed as a self-contained piece of equipment, intended to be used by a human operator, than designed as a data source for an automated data fusion system, with reference to the needs of such a system. We should not expect radar contacts, for example, to be grouped together in a single message by the radar device whenever they are close enough to compete with each other in a subsequent association decision. Thus, it can easily happen that contacts that ought to be considered with reference to each other, due to their proximity, are sent to the data fusion system in separate messages.

The long- and medium-range tracking radar systems on board the HALIFAX class frigates send messages (here called "buffers"), each of which contains either a single contact[2] or several contacts detected in sequence. Two criteria are used for grouping contacts together: Some maximum number of contacts is allowed in each buffer, and some maximum angular separation is allowed between the first and last contact in a buffer. Within these constraints, each buffer consists of as many contacts as possible. Two contacts that are close to each other in range and bearing may arrive in separate buffers, in spite of their proximity. This separation will occur whenever the buffering criteria allow only one of these contacts to be included among those that have been assigned already to the next buffer to be sent.

Suppose that such buffers are sent to a data fusion system. If the system is designed with the assumption that close contacts will always be grouped together in a single message, and it makes irreversible association decisions in immediate response to each group of contacts that it receives, then it will make errors in association that it would not otherwise have made.

1.2 Errors from the assumption of simultaneous availability

Association decisions are based on a distance metric that is defined for each track-contact pair. The distances for all the track-contact pairs that are relevant for a given association decision, calculated according to the chosen metric, can be gathered together into an "association matrix". (An example is given below.)

The simplest approach to association is to make an immediate, firm decision, with the constraint that each contact will be associated with at most one track, and each track will be associated with at most one contact. If we have such a decision to make, involving m tracks and n contacts, then the number of pairs we want to associate is the lesser of m and n (assuming that none of the track-contact pairs has been disqualified from association). Within this constraint, a standard criterion for the optimal

[2] The word *contact* is not strictly correct. The "contacts" sent by the radar system are actually tracks, since some level one data fusion has already been applied to the contacts by the Automatic Detection and Tracking (ADT) system that is included within each radar system. But see footnote 1.

association decision is that the sum of the distances of the associated pairs should be minimised[3].

For example, consider the following association matrix of distances between tracks (rows) and contacts (columns):

$$\begin{bmatrix} 2.3 & 4.7 \\ 3.2 & 9.8 \end{bmatrix}$$

Suppose that our system makes immediate, firm association decisions based on the standard criterion of minimal summed distance. In this case, if the contacts are processed simultaneously, then the second track will be associated with the first contact and the first track with the second contact, for a total distance of 7.9, versus 12.1 for the alternate association decision. But if the first contact is handled alone, it will be associated with the first track, so we have immediately an association decision that is less likely to be correct. The resulting track update will lead to a recalculation of the distance between the first track and the second contact. That distance will probably increase, but if it remains less than 9.8, then the second contact will also be associated with the first track. In that case, we have an error of a kind that could not have occurred if the two contacts had been grouped together in a single message. One track has been updated twice, by two contacts that almost certainly originated from different targets, while the other track has missed an update.

Problems of this very kind were found [3] to occur from time to time in the Cycle 2 sea trials of the Command Decision Aid Technology (COMDAT) project [4], which is a Technology Demonstration (TD) project of Defence R&D Canada (DRDC). This project seeks to demonstrate the feasibility of automated data fusion based upon the sensor suite of the HALIFAX-class frigates. So far, the focus of the project has been on above-water sensors – the long- and medium-range tracking radar, their subordinate "Identification: Friend or Foe" (IFF) systems, and Electronic Support Measures (ESM) – plus inter-platform communication via Link-11 or the Global Command and Control System (GCCS).

Figure 1 shows an example of an association problem, caused by the separation of contacts into separate buffers, that occurred during the sea trial. Figure 1a shows the radar contact positions while figure 1b shows the resulting update positions of the corresponding COMDAT system tracks. The paths followed by the aircraft start from the lower right. A gap in the radar data (between the points labelled "1" and "3" near the centre of figure 1a) leads to the deletion of the corresponding system track (at the point labelled "2" near the centre of figure 1b). When radar contacts from that aircraft resume (at the point labelled "3"), the first such contact arrives as the sole contact in its buffer. This contact falls in the gate of the surviving system track, thus causing seduction of that track. If this contact had been packaged together with the nearest contact (in time) from the other target (at the point labelled "4"), this error in association would not have occurred, for the system would have been able to compare the two contacts and to choose the better one.

---

[3] If we have calculated a likelihood p for each track-contact pair, then a standard distance metric is given by $D = -\log p$. Thus the minimum summed distance represents a maximum likelihood product.

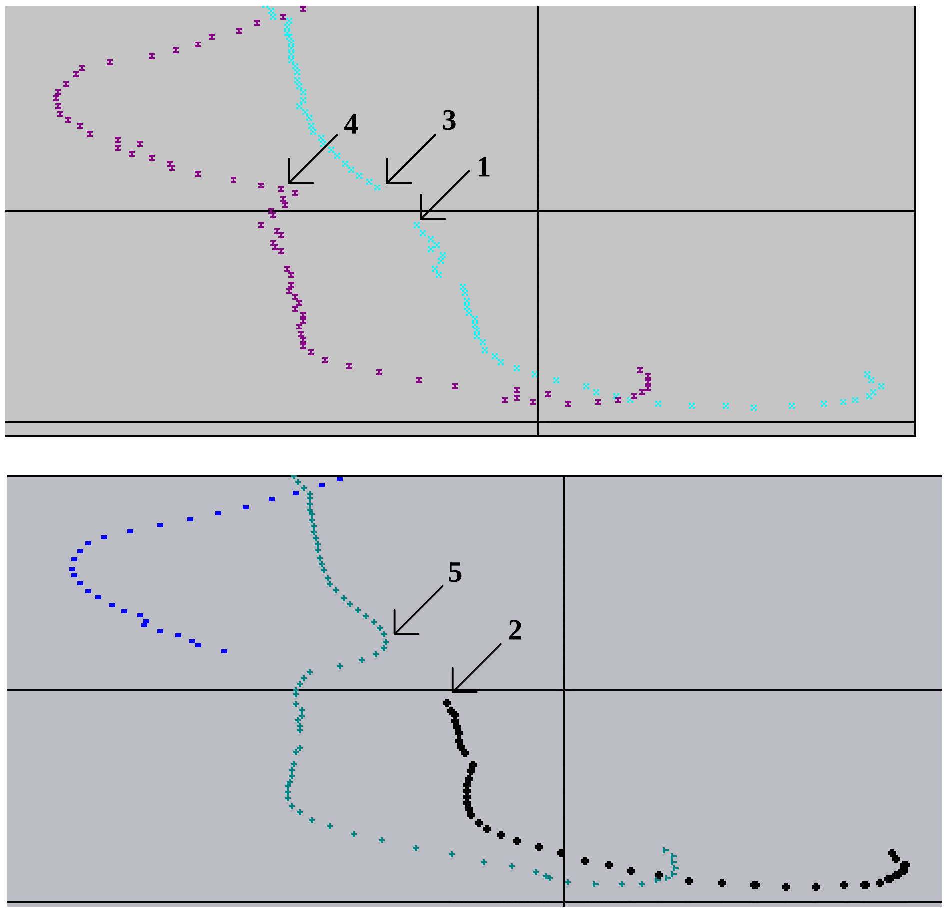

Figure 1. *A misassociation due to contact buffering.*

(a) Radar contacts from two aircraft. (b) COMDAT system tracks. The following events are featured: (1) The radar contacts from one of the aircraft cease. (2) The corresponding system track is dropped. (3) The radar contacts from that aircraft resume. (4) The nearest contact (in time) to that of event 3 arrives in a separate message from the latter. (5) The other system track is seduced by the contact of event 3.

Although the problem is most obvious in those systems that make immediate, firm association decisions, such as COMDAT, other kinds of systems would be affected as well. The separation of contacts into different messages affects the system's ability to discern that those contacts will not have originated from the same target, and thus its judgment of all the relevant probabilities will be skewed.

1.3 Assessing the severity of the problem

The incidence of the problems that result from the separation of contacts can be reduced greatly if we construct a holding pool in which the contacts from the buffers are gathered, and from which they are released to the data fusion system when certain criteria are met. However, the use of a holding pool comes with a price in timeliness. The processing of contacts is delayed.

So we may ask: How often, and under what circumstances, is there a substantial probability of an incorrect association due to the manner in which the contacts are grouped? (To put the problem in perspective, remember that incorrect association decisions will sometimes be made in any case, due to other causes.) It may be that the problem observed in COMDAT was an isolated case, and that we need not be concerned about it. On the other hand, it may be that such problems can be expected to occur frequently. In the latter case, methods to address the separation of contacts (such as the holding pool technique) would be of fundamental importance. Input management would take its place alongside association and tracking as an essential component of level one data fusion.

All else being equal, we are interested in the way in which the association performance of the data fusion system depends on the buffering constraints (that is, the maximum angular separation and the maximum number of contacts per buffer). But of course, this dependence will vary with circumstances. In particular, a set of buffering constraints that causes great problems when the target density is high may have no noticeable effect when targets are sparse. Thus, we are also interested in the effect of target density.

Section 2 describes an exploratory experiment, involving simulated data, that was performed in order to develop a sense of the rate of incidence of problems arising from the separation of contacts, and the dependence of these problems on the target density. Section 3 presents a statistical analysis of the results of that experiment. Section 4 considers briefly some of the holding pool methods that might be used to reduce the impact of the fragmentation of contacts.

2. Data association simulation

Two hundred scenarios were generated, each with two hundred targets, confined to a single quadrant of bearing. In each scenario, a track is created for each target. The tracks are designed to have the form of mature tracks from a real data fusion system, as if they had been updated most recently by a set of contacts from a recent radar sweep. (See subsection 2.1 on the generation of tracks.) Then a contact is created for each target, as if from the radar sweep immediately following that which led to the last track update. (See subsection 2.2 on the generation of contacts.) All the contacts are processed once by a simple data fusion system.

In the first hundred scenarios (the “baseline scenarios”), no data buffering rules are applied, and hence all contacts are compared with all tracks. A global association decision is made. Each contact is associated with at most one track, and each track is associated with at most one contact. For each track, it is noted whether the system made the correct association, an incorrect association, or no association.

As noted in the introduction, we are also interested in the effect of target density on the degree to which contact fragmentation will affect the association performance. In order to capture the effect of target density, the tracks are grouped according to the number of contacts whose hypothetical association with the track in question was considered feasible by the system. (See subsection 2.3 for the association rules.) This number is called the “gate population” of the track. Within each gate population category, the fraction of successful tracks serves as a baseline – where “success” may refer either to a correct association having been made, or to incorrect associations having been avoided (see section 3).

In the remaining hundred scenarios (the “buffer scenarios”), the contacts are grouped into buffers. There is a maximum number of contacts per buffer, and a maximum angular separation from the first contact to the last in any given buffer. Each buffer consists of a sequential set of contacts, the number of contacts in each buffer is maximised within the given constraints, and this maximisation is performed in the order of the time tags of the contacts.

Two different sets of rules are used for the grouping of contacts. Each buffer scenario is executed twice; one run is for the first set of rules for contact grouping, and the other run is for the second set. In the first set of rules (the “mildly fragmented” case), each buffer has a maximum of twenty contacts, with a maximum separation of fifteen degrees from the first contact in the buffer to the last. (This case is roughly representative of the buffering behaviour of the long- and medium-range tracking radar of the HALIFAX class frigates, so the effects of this degree of fragmentation are applicable to the COMDAT system.) The second rule set (the “severely fragmented” case) imposes tighter constraints: a maximum of five contacts per buffer, and a maximum angular separation of 4.8 degrees[4]. (This case is motivated by the very real possibility that a higher degree of contact fragmentation may have to be faced by the designers of future data fusion systems.)

As each buffer is processed, each track that is associated with a contact is updated with reference to that contact. (Subsection 2.4 presents the rule for track updating.) Afterwards, the updated tracks are treated no differently from the others. In particular, they are available to compete for association with the contacts of each subsequent buffer. Thus, a track may be associated with several contacts over the course of the sweep, but with no more than one from any given buffer.

The performance of the system (with respect to correct and incorrect associations) is stored for each run. At the end of all scenarios, we have performance data for three cases – the baseline, the mildly fragmented case, and the severely fragmented case – sorted according to the gate population of the tracks in question.

We want to maintain consistency in the manner in which target density is reflected by each track's gate population score. Therefore, the gate population score for a track in a buffer scenario is defined to be equal to what it would have been if that scenario had been a baseline scenario. That is, all tracks are compared with all contacts before the two runs are performed, and the gate population of each track is the number of contacts of which the hypothetical association with that track would be considered feasible, based on that global comparison.

Within each gate population category, the performance of the system in the baseline case is compared with that in each of the other two cases. See section 3 for the statistical methods.

The simulation was implemented in Mathematica[5], version 5.0.

2.1 Construction of tracks

The track positions are confined to the region with range between five and twenty kilometres, and with bearing between zero and ninety degrees. The simulated radar system is considered to be located at the origin.

The targets are generated sequentially, and are grouped into "clusters" in order to bring about a wide variation in the gate population scores. Each target is given a probability of $(3/4)^n$ to initiate a new cluster, where n is the number of clusters already begun. If a new cluster is initiated, the track is given a random position with a distribution that is uniform (in probability per unit area) over the target zone. Otherwise, the new target

---

[4] The constraints that were used in the simulation were originally conceived in terms of a maximum elapsed time between the first and last contact of a buffer. The mildly fragmented case has a maximum buffer time of 0.125 seconds, and the highly fragmented case has a maximum buffer time of 0.04 seconds. For a radar rotation period of three seconds, these time limits work out to the angular limits described in the text.

[5] "Mathematica" is a registered trademark of Wolfram Research, Inc.

is assigned to an existing cluster, with an equal probability assigned to each cluster. The position of the target, in that case, is randomly generated with a normal distribution, centred on the average position of the existing targets in that cluster, and with a standard deviation of 400 metres in each horizontal Cartesian dimension. (This distribution is truncated in order to keep the targets within the chosen target zone.) Figure 2 gives an example of the resulting distribution of track positions.

The time tag of each track is taken directly from the track bearing, as if the last track update had coincided in position with the last contact used for that update:

$$t = \frac{\theta}{2\pi}\Delta t\,, \qquad (2\text{-}1)$$

where $\theta = \arctan(x/y)$ is the bearing of the track, and Δt is the radar's rotational period, which we set to three seconds. That is, the radar is considered to be rotating clockwise, and to have been pointing due north (for the last time prior to the track updates) at time zero.

The velocity of each track is given a random direction, uniformly distributed over all directions, and a random magnitude, normally distributed with a mean of 150m/s (roughly 300 knots) and a standard deviation of 40m/s. The track state consists of its Cartesian position and velocity, represented as $\mathbf{x} = \begin{pmatrix} x & y & v_x & v_y \end{pmatrix}^{\mathrm{T}}$.

In order to construct a realistic covariance matrix, we imagine that the track has been updated many times by the same kind of Kalman Filter [1, 2] that we use later for track updating within the simulation, as described in subsection 2.4. For convenience, we also imagine that the past updates were periodic, and that the Kalman Filter assumed the same measurement covariance matrix for each contact that was associated with the track. The measurement covariance matrix used for this purpose is derived from the track's current position, with an uncertainty of $\sigma_r = 50$m in range and $\sigma_\theta = 0.01$ radians (about 0.57 degrees) in bearing. This measurement covariance is then

$$\mathbf{R}_{\text{history}} = \frac{1}{x^2 + y^2}\begin{pmatrix} y & x \\ -x & y \end{pmatrix}\begin{pmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{pmatrix}\begin{pmatrix} y & -x \\ x & y \end{pmatrix}, \qquad (2\text{-}2)$$

where $\sigma_1 = \sigma_\theta \sqrt{x^2 + y^2}$ and $\sigma_2 = \sigma_r$.

The process model for the Kalman Filter is the "constant velocity" model, meaning that all the acceleration is contained within the process noise. The transition and process noise matrices are

$$\mathbf{F}(T) = \begin{pmatrix} \mathbf{I}_{(2)} & T\mathbf{I}_{(2)} \\ \mathbf{0}_{(2)} & \mathbf{I}_{(2)} \end{pmatrix} \text{ and}$$

$$\mathbf{Q}(T) = \eta\begin{pmatrix} \frac{1}{3}T^3\mathbf{I}_{(2)} & \frac{1}{2}T^2\mathbf{I}_{(2)} \\ \frac{1}{2}T^2\mathbf{I}_{(2)} & T\mathbf{I}_{(2)} \end{pmatrix}, \qquad (2\text{-}3)$$

where the noise parameter η is set to 500m^2/s^3, $\mathbf{I}_{(2)}$ is the 2-by-2 identity matrix, $\mathbf{0}_{(2)}$ is the 2-by-2 matrix of all zeroes, and T is the time elapsed since the previous update. Here, we assume that $T = \Delta t$ for every update in the track's (imagined) history.

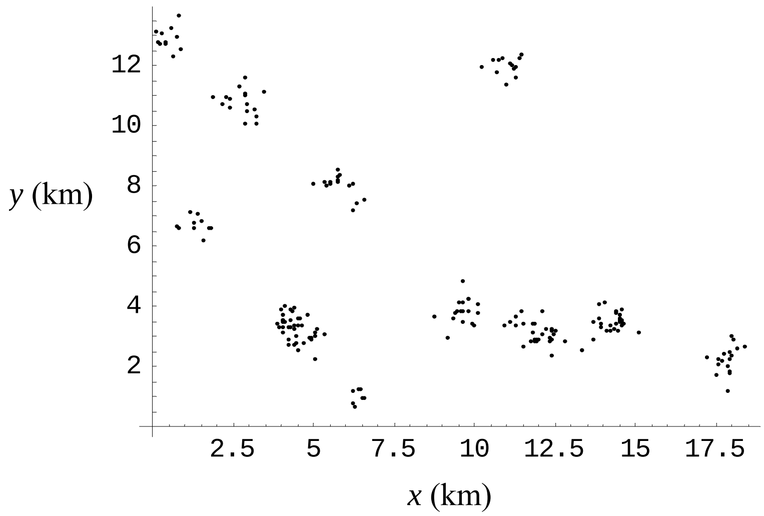

***Figure 2.** A target sample.*

Each dot represents the current position estimate of one track. These two hundred track positions were generated in the same way as those that were used in each scenario in the study. The scale is in kilometres.

The track's covariance matrix is set to the steady-state value[6] that it would have reached under the assumptions given above. That is,

$$\mathbf{P} = \frac{1}{x^2 + y^2}\begin{pmatrix} y & x & 0 & 0 \\ -x & y & 0 & 0 \\ 0 & 0 & y & x \\ 0 & 0 & -x & y \end{pmatrix}\begin{pmatrix} \beta_1 & 0 & \kappa_1 & 0 \\ 0 & \beta_2 & 0 & \kappa_2 \\ \kappa_1 & 0 & \upsilon_1 & 0 \\ 0 & \kappa_2 & 0 & \upsilon_2 \end{pmatrix}\begin{pmatrix} y & -x & 0 & 0 \\ x & y & 0 & 0 \\ 0 & 0 & y & -x \\ 0 & 0 & x & y \end{pmatrix}, \quad (2\text{-}4)$$

where

$$\upsilon_i = \eta\left(\frac{\beta_i}{\kappa_i} - \frac{\Delta t}{2}\right),$$

$$\kappa_i = \sqrt{\eta \Delta t\left(\sigma_i^2 - \beta_i\right)},$$

$$\beta_i = \sqrt{\tfrac{1}{3}\left(2\mu_i + \tfrac{1}{4}\eta\Delta t^3\right)^2 + 2\mu_i\sigma_i^2} - \mu_i - \tfrac{1}{6}\eta\Delta t^3, \text{ and}$$

$$\mu_i = \sqrt{\eta\Delta t^3\left(\tfrac{1}{48}\eta\Delta t^3 + \sigma_i^2\right)}, \quad (2\text{-}5)$$

for $i = 1$ and $i = 2$.

2.2 Construction of contacts

The time tag of each track, from equation (2-1), is based on the notion that the most recent update (in the imagined prior history of the track) was generated from a radar contact belonging to the sweep of the first quadrant that began at the time $t = 0$. The

[6] If we start with equation (2-17) for the updated covariance matrix, set $\mathbf{R} = \mathbf{R}_{\text{history}}$ (given by (2-2)), set $t_c = t + \Delta t$, set $\mathbf{P}_{\text{update}} = \mathbf{P}$, and solve for $\mathbf{P}$, then equations (2-4) and (2-5) present the only positive definite solution.

contacts that are used in the scenarios are considered to come from the following rotation, wherein the sweep of the first quadrant began at the time $t = \Delta t$. One contact is generated for each track. (That is, we are assuming, for convenience, no false alarms and a perfect probability of detection.)

The first step in constructing each contact is to find the time t_p at which the radar will be pointing in the direction of the target in question, under the assumptions that the corresponding track state is correct and that the target will maintain a constant velocity in the interim. Under these assumptions, the projected position of the track would be

$$\mathbf{x}_p = \begin{pmatrix} x_p & y_p \end{pmatrix}^T = \mathbf{HF}(t_p - t_0)\mathbf{x}, \tag{2-6}$$

where $\mathbf{x} = \begin{pmatrix} x & y & v_x & v_y \end{pmatrix}^T$ is the track state, t_0 is the track's time tag, and $\mathbf{H}$ is the "measurement matrix": $\mathbf{H} = \begin{pmatrix} \mathbf{I}_{(2)} & \mathbf{0}_{(2)} \end{pmatrix}$. Let r_p and θ_p denote the projected range and bearing:

$$x_p = r_p \sin\theta_p \text{ and } y_p = r_p \cos\theta_p. \tag{2-7}$$

Then the time t_p is found by solving (numerically) the equation

$$t_p = \left(1 + \frac{\theta_p}{2\pi}\right)\Delta t \tag{2-8}$$

together with (2-6) and (2-7).

The "true" position $\mathbf{x}_t = \begin{pmatrix} x_t & y_t \end{pmatrix}^T$ of the target at time t_p is generated randomly, according to a normal distribution whose mean is the projected position $\mathbf{x}_p$ and whose covariance matrix is

$$\mathbf{P}_p = \mathbf{H}\left(\mathbf{F}(t_p - t_0)\mathbf{PF}(t_p - t_0)^T + \mathbf{Q}(t_p - t_0)\right)\mathbf{H}^T. \tag{2-9}$$

In other words, the track state is projected ahead through a time interval $t_p - t$, and the position-related components of the resulting state and covariance matrix are assumed to represent faithfully the probability distribution of the true position. Let r_t and θ_t denote the "true" range and bearing:

$$x_t = r_t \sin\theta_t \text{ and } y_t = r_t \cos\theta_t. \tag{2-10}$$

The contact range r_c is generated randomly, according to a normal distribution whose mean is the "true" range r_t and whose standard deviation is $\sigma_r = 50$m. The contact bearing θ_c is generated similarly from the "true" bearing θ_t, using $\sigma_\theta = 0.01$ radians.

The covariance matrix

$$\mathbf{R} = \begin{pmatrix} \cos\theta_c & \sin\theta_c \\ -\sin\theta_c & \cos\theta_c \end{pmatrix} \begin{pmatrix} r_c\sigma_\theta & 0 \\ 0 & \sigma_r \end{pmatrix} \begin{pmatrix} \cos\theta_c & -\sin\theta_c \\ \sin\theta_c & \cos\theta_c \end{pmatrix} \tag{2-11}$$

is used in the procedures for association and track updating.

Finally, the time

$$t_c = \left(1 + \frac{\theta_c}{2\pi}\right)\Delta t, \tag{2-12}$$

is used (instead of t_p) as the time tag of the contact.

2.3 Association

Given a track with state estimate **x** and covariance matrix **P** at time t_0, and a contact with position $\mathbf{z}_c = (x_c \quad y_c)^T$ at time t_c, the hypothetical association between that track and that contact is declared feasible if the contact is found in the track's gate, which is to say

$$\nu \mathbf{S}^{-1} \nu^T \leq \gamma , \qquad (2\text{-}13)$$

where

$$\nu = \mathbf{z}_c - \mathbf{HF}(t_c - t_0)\mathbf{x} ,$$
$$\mathbf{S} = \mathbf{H}\left(\mathbf{F}(t_c - t_0)\mathbf{PF}(t_c - t_0)^T + \mathbf{Q}(t_c - t_0)\right)\mathbf{H}^T + \mathbf{R} , \qquad (2\text{-}14)$$

F and **Q** are given by (2-3), **R** is given by (2-11), and $\gamma = 9.21$. (If we imagine ν to be a normally distributed two-dimensional random vector with zero mean, and that **S** is its covariance matrix, then condition (2-13) accounts for 99% of the probability distribution.) If feasible, the association distance for that track-contact pair [2] is

$$\delta = \nu \mathbf{S}^{-1} \nu^T + \ln(\det \mathbf{S}), \qquad (2\text{-}15)$$

for the purpose of comparing that hypothetical association with its alternatives.

For a given scenario and a given contact-grouping case, the buffers are processed sequentially. Every contact in a buffer is compared with every track. Each track-contact pair is assessed for the feasibility of association, and the association distance is computed for all feasible pairs.

A firm association decision is made, according to the standard [2] Nearest-Neighbour criteria:

1. Each contact is associated with at most one track.

2. Each track is associated with at most one contact.

3. A track and contact can be associated with each other only if that association is declared feasible.

4. The number of associations is maximised, within the constraints of the first three rules.

5. The summed association distance of the chosen associations is minimised, within the constraints of the first four rules.

The association routine that was implemented in this simulation was based on Blackman's [2] presentation of the Munkres algorithm, modified [5] in order to obey the third rule consistently.

2.4 Tracking

When the association decision is made for a buffer, the tracks are then updated according to the rules for a Kalman Filter [1, 2]. (An exception is made for the last buffer of a run. After its association decisions are made, the run is complete.) If a track of state estimate **x** and covariance matrix **P** at time t_0 is associated with a contact of position $\mathbf{z}_c = (x_c \quad y_c)^T$ at time t_c, then the updated state estimate and covariance are given by

$$\mathbf{x}_{update} = \mathbf{P}_{update}\left(\left(\mathbf{F}(t_c - t_0)\mathbf{P}\mathbf{F}(t_c - t_0)^T + \mathbf{Q}(t_c - t_0)\right)^{-1}\mathbf{x} + \mathbf{H}^T\mathbf{R}^{-1}\mathbf{z}_c\right) \text{ and} \quad (2\text{-}16)$$

$$\mathbf{P}^{-1}_{update} = \left(\mathbf{F}(t_c - t_0)\mathbf{P}\mathbf{F}(t_c - t_0)^T + \mathbf{Q}(t_c - t_0)\right)^{-1} + \mathbf{H}^T\mathbf{R}^{-1}\mathbf{H}. \quad (2\text{-}17)$$

The track is modified accordingly: $\mathbf{x}_{update}$ replaces **x** as the track's state estimate, $\mathbf{P}_{update}$ replaces **P** as its covariance matrix, and t_c replaces t_0 as its time tag.

The updated tracks are now free to compete with other tracks for association with the contacts of the next buffer.

3. Statistical test

The tracks from all the scenarios were gathered together and grouped according to their gate population, as described in section 2. Gate population scores varied from zero to 27. The sample size for each gate population score from one to fifteen was greater than 100. For each gate population score from one to fifteen, the fraction of the tracks for which the correct association (only) was made is shown in figure 3. These results are shown for all three cases: the baseline, the case of mild fragmentation, and the case of severe fragmentation.

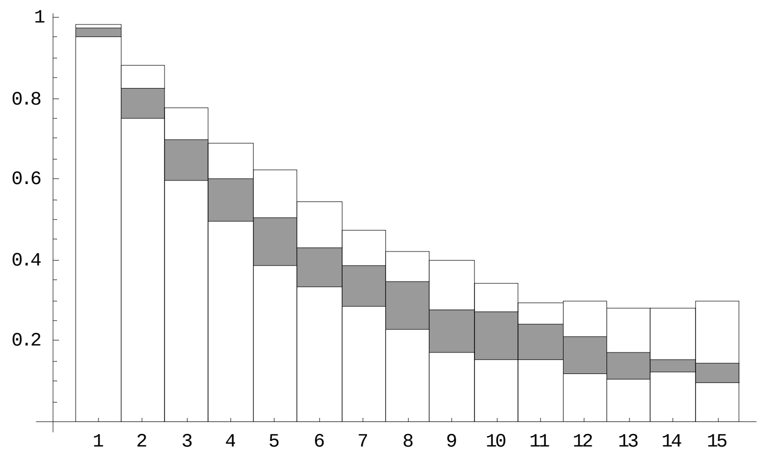

Figure 3. *Rate of perfect association, by gate population.*

For each gate population score from one to fifteen, the full bar shows the success rate (that is, the rate of perfect association) in the baseline case. The top and bottom of each grey box indicates respectively the success rate for the mildly and severely fragmented cases.

For the sake of the statistical tests, the sorting of the complete sample according to the gate population scores was done more coarsely. The tracks are sorted into three broad categories: those with a gate population of two or less, those with a gate population of three to six, and those with a gate population of seven or more. The sample sizes for the three categories were respectively 3427, 9763, and 6810 in the baseline case, and 3371, 10005, and 6624 in the buffered cases.

For each gate population category, the rate of success of the system in the baseline case is compared to the rate of success of the system for each of the other two cases. Three kinds of "success" are considered for this purpose: (a) *correct* association, meaning that the track is associated with the correct contact, (b) *clean* association, meaning that the track is not associated with any incorrect contacts, and (c) *perfect* association (as in figure 3), meaning that the track's association is both correct and clean. The rate of success (of each kind) is shown for each case and each category in table 1.

Table 1. Sample success rates.

	NO FRAGMENTATION (BASELINE)	MILD FRAGMENTATION	SEVERE FRAGMENTATION
gate population 0 to 2	0.92 [0.92, 0.93]	0.88 [0.89, 0.92]	0.83 [0.85, 0.90]
gate population 3 to 6	0.66 [0.66, 0.66]	0.56 [0.60, 0.62]	0.46 [0.55, 0.60]
gate population 7 or more	0.39 [0.39, 0.39]	0.30 [0.35, 0.38]	0.20 [0.32, 0.42]

Success rates are shown in the format: *perfect* [*correct*, *clean*]

In the baseline case, the rate of perfect association equals the rate of correct association, because each track can only be associated with one contact, as explained in section 2. The very slight difference between the rate of perfect association and the rate of clean association (in the baseline case) arises from the rare cases where a track has a gate population of zero. As the degree of fragmentation increases, we encounter an increasing variety of ways in which the system may fail to achieve perfect association, which leads to an increasing discrepancy between the rate of perfect association and the rate of each of the other two kinds of success. The decrease in success rate as the degree of fragmentation increases is particularly noticeable if we focus on perfect association, and less dramatic if we focus on correct association. Counter-intuitively, the rate of clean association is affected least of all by fragmentation – and for the higher gate populations, this sample generated a *higher* rate of clean association under severe fragmentation than in the baseline.

The remainder of this section describes a statistical test that was performed in order to determine whether the grouping of contacts into buffers caused a significant and substantial decrease in the probability of success, relative to the baseline.

3.1 Method

Given two samples of sizes n_1 and n_2, with s_1 and s_2 successes respectively, we wish to compare the probability of success p_1 in the process underlying the first sample to the corresponding probability p_2 for the second sample. (Let the first sample refer to the baseline case and let the second sample refer to one of the other cases.)

Suppose we wish to perform a statistical test wherein the null hypothesis is $p_2 \geq p_1$ and the alternative hypothesis is $p_2 < p_1$. Then the test statistic

$$Z = \frac{\frac{s_1}{n_1} - \frac{s_2}{n_2}}{\sqrt{\left(\frac{s_1 + s_2}{n_1 + n_2}\right)\left(1 - \frac{s_1 + s_2}{n_1 + n_2}\right)\left(\frac{1}{n_1} + \frac{1}{n_2}\right)}} \tag{3-1}$$

is appropriate [6]. If $p_1 = p_2$, and n_1 and n_2 are large, then Z should be (approximately) normally distributed with a mean of zero and a standard deviation of one. For this test

Table 2. Test significance for a decrease in success rate, relative to the baseline.

	MILD FRAGMENTATION	SEVERE FRAGMENTATION
gate population 0 to 2	4.2×10^{-8} [3.6×10^{-6}, *0.065*]	0 [0, 4.0×10^{-4}]
gate population 3 to 6	0 [0, 0]	0 [0, 0]
gate population 7 or more	0 [8.6×10^{-7}, *0.14*]	0 [0, *0.9999*]

Significance scores are shown in the format: *perfect* [*correct*, *clean*]. A zero indicates a result less than 10^{-8}. Results that fail to show a significant decrease are shown in italics.

statistic, we would reject the null hypothesis if the integral of such a normal function from Z to infinity evaluates to 0.05 or less – that is, if $Z \geq 1.645$.

Table 2 shows the value of this integral for all three gate population categories, for all three kinds of success, for both degrees of fragmentation. A significant drop in the rate of correct (or perfect) association is found for all cases. The drop in the rate of clean association, with severe fragmentation, is significant only for the first two gate population categories, while the corresponding drop with mild fragmentation is significant only for the intermediate gate population category.

In the present study, however, the mere claim that $p_2 < p_1$ is not surprising, and it is not clear that such a result is a matter of concern. Rather, what we need to know is whether the difference in probability is large enough to matter.

If we knew the probabilities in question, the natural[7] measure of the resulting difference in performance would be the odds ratio

$$\lambda = \frac{p_1(1-p_2)}{p_2(1-p_1)}. \tag{3-2}$$

We need a statistical test to investigate whether λ is or is not greater than some threshold λ_0 (which is greater than one). Thus, our null hypothesis is $\lambda \leq \lambda_0$, and our alternative hypothesis is $\lambda > \lambda_0$. The value of λ_0 is chosen to reflect the degree to which we are prepared to tolerate the problems that arise from the decrease in the probability of success (with respect to association decisions). Equivalently, we could state the null hypothesis as

$$p_2 \geq \frac{p_1}{p_1 + \lambda_0(1-p_1)} \tag{3-3}$$

and the alternative hypothesis as

$$p_2 < \frac{p_1}{p_1 + \lambda_0(1-p_1)}. \tag{3-4}$$

[7] An absolute difference in probability would not be a natural measure of the difference in performance. Consider, for example, the meaning of an absolute decrease in probability of 0.05, if the baseline success probability is 0.95, or 0.5, or 0.1. In one case the probability of success is halved, in another case the probability of failure is doubled, and in yet another case, both the probability of failure and the probability of success are affected only slightly.

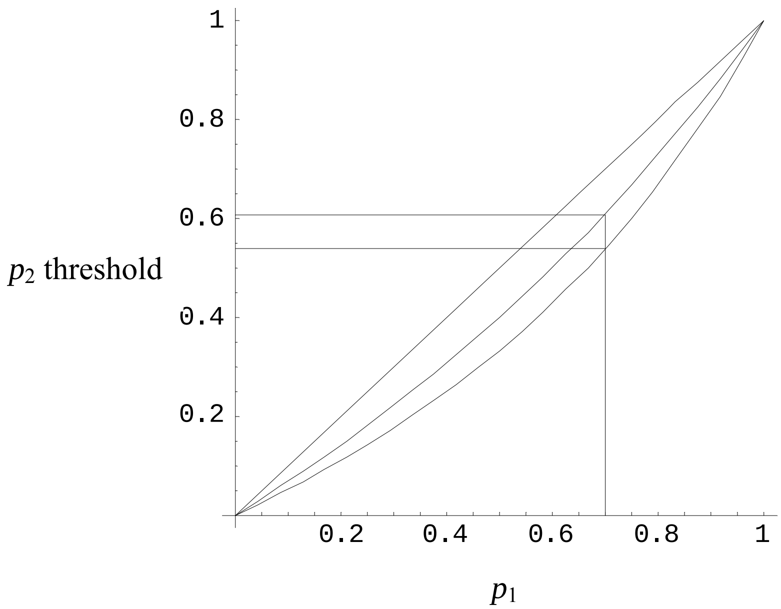

Figure 4. *The probability shift.*

The curves show the threshold of the probability p_2 as a function of p_1, for three different values of λ_0. From top to bottom, the three curves correspond to $\lambda_0 = 1$, $\lambda_0 = 1.5$, and $\lambda_0 = 2$. The vertical and horizontal lines illustrate the example of $p_1 = 0.7$.

We can test these hypotheses using the same test as we used earlier, if we substitute

$$s_1' = \frac{s_1}{s_1 + \lambda_0 (n_1 - s_1)} \tag{3-5}$$

for s_1 in (3-1).

Figure 4 shows the threshold probability as a function of the original probability, in the cases $\lambda_0 = 1$ (no change), $\lambda_0 = 1.5$ and $\lambda_0 = 2$. If $p_1 = 0.7$, for example, then inequality (3-4) becomes $p_2 < 0.61$ for $\lambda_0 = 1.5$,or $p_2 < 0.54$ for $\lambda_0 = 2$. This example is indicated in figure 4 by the vertical and horizontal lines.

What value of λ_0 makes sense? How much does the probability of success have to decrease in order for us to say that we have a practical problem? The best answers would make reference to the overall system effectiveness, defined in terms of the success (and cost) of a mission. But it is not a straightforward matter to map our probability shift onto any such measure of effectiveness. Even if we confine our attention to the lower levels of abstraction, and try to consider the tradeoffs involved in whatever procedures one might adopt in order to cure the problem, we will run into difficulty. For example, it is not easy to pin down the rate of exchange between a decrease in the probability of correct association and an increase in the time it takes to identify an incoming target. Moreover, our chosen value of λ_0 might itself depend on the baseline probability of correct association.

Instead of choosing an arbitrary value for λ_0, and checking for statistical significance given that threshold, we will solve for the value of λ_0 that is required in order to achieve statistical significance.

3.2 Results

Table 3 shows the odds-ratio threshold λ_0 that would be required in order for our sample to justify the rejection of the null hypothesis (4-3) based on the criterion $Z \geq 1.645$, where Z is calculated using formula (4-1) with the substitution (4-5). (The word "none" appears instead of a number if the significance criterion fails even for $\lambda_0 = 1$.) That is, for each number in the table, our sample gives us confidence in the claim that the odds of success is decreased, relative to the baseline, by at least the amount shown, but we cannot claim a statistically significant result for a larger drop in probability.

These results show, again, that the rate of clean association is affected less than the rate of correct association by the separation of contacts. If we choose the value $\lambda_0 = 1.5$ to delineate the magnitude of the change in success rate that is large enough to matter, then we find no statistically significant results in the case of mildly fragmented contacts. The severely fragmented case, with the tighter restrictions on the size of a buffer, brings about a statistically significant decrease in the rate of perfect association, for all gate population categories, and in the rate of correct association, for the first two gate population categories. No such decrease is found in the rate of clean association.

Table 3. *Limits on odds-ratio thresholds required for a statistically significant decrease in success rate, relative to the baseline.*

	MILD FRAGMENTATION	SEVERE FRAGMENTATION
gate population 0 to 2	1.37 [1.27, none]	2.15 [1.81, 1.17]
gate population 3 to 6	1.45 [1.26, 1.14]	2.23 [1.52, 1.26]
gate population 7 or more	1.40 [1.12, none]	2.38 [1.27, none]

Critical values of λ_0 are shown in the format: *perfect* [*correct*, *clean*]

4. Buffer management methods

A very simple way to avoid the problem of multiple associations is to impose a rule according to which any newly updated track is ineligible for further associations with contacts from the same sensor until some appropriate time period has expired. In the case of a track that has just been updated by a radar contact, the duration of the moratorium might be one half or one fourth of the radar's rotational period. This solution does not, however, address the fact that our association method will have been compromised by the artificial removal of some of the contacts that ought to have been considered.

A wide class of other solutions requires the insertion of an additional component – a buffer management module (BMM) – to the fusion system, between the sources of input and the part of the system that makes association decisions. The task of the BMM is to regroup the contacts. It would include a holding pool for each sensor, where incoming contacts from that sensor are retained until some conditions are met, at which time several contacts, grouped together as a new buffer, will be sent from the pool to the rest of the data fusion system.

Whatever rules we choose to govern the release of contacts from the holding pool, we must take care not to hold the contacts for too long. If our chosen conditions allow the possibility that multiple contacts originating from the same target may be in the pool at the same time, and our fusion system is not designed to handle that situation, then we will have created potential problems that are as bad as those we set out to fix. Thus, an upper bound (such as one half of a radar's rotational period) should be imposed on the amount of time that a contact can be held in the pool, and this time limit should take precedence over the usual rules.

The use of a time limit implies that our rules for releasing contacts from the holding pool will not always solve the original problem. However, the time limit would (presumably) not be reached in scenarios where targets are sparse, while the reappearance of the original problem from time to time in a dense scenario may be less important than the other problems that will probably arise in such a scenario.

Depending on the data structures that are used for the holding pool, it may also be necessary to impose an upper bound on the number of contacts that can be held.

The basic rules for a holding pool are based on a measure of proximity between contacts:

1. Contacts that are sufficiently close to each other must be grouped together. (This rule is transitive, of course.)

2. Contacts are not grouped together unless forced by the first rule.

3. A contact is held in the pool as long as there remains any possibility of the appearance of a new contact that would be grouped with it according to the first rule.

4. Whenever the third rule is found not to apply to a group of contacts in the pool, those contacts are sent to the rest of the fusion system, and deleted from the pool.

The rule set can be made easier to implement if we modify the first rule to say: "Contacts that are sufficiently close to each other, or were originally grouped together in the same buffer, must be grouped together." If we use the modified rule, fewer pairwise comparisons are required, and the original buffers can simply be concatenated whenever any contact in one buffer and any contact in the other are found to be sufficiently close. However, this modification would tend to increase the total amount of time any given contact is held in the pool, and thus would increase the danger that the time limit will have to apply.

It remains to discuss what we might mean by "sufficiently close". Perhaps the simplest proximity measure is the time difference of the contacts. If the contacts in the buffer are sorted by time, then each contact would (at most) need to be compared only to the one immediately preceding it and the one immediately following it. If the absolute time difference is less than some threshold value, then the contacts would be considered sufficiently close to be grouped together.

For a rotating radar system, a time difference between contacts is equivalent to a difference in bearing. Notice that a holding pool with a time-based proximity measure is practically equivalent to the transformation of an input source with tight buffering constraints to one with more generous buffering constraints.

A two-dimensional distance would be better, in principle. For example, contacts might be considered sufficiently close if their absolute difference in bearing is less than some threshold, and their absolute difference in range is less than some other threshold. As a way to cut down the number of pairwise comparisons that are required, we could imagine an adaptive system in which the projected track positions are relayed to the BMM from the tracking module, and the proximity measures are applied more strictly in the vicinity of the projected positions of a track.

5. Discussion and conclusions

The performance of a data fusion system, with respect to association, is degraded by contact fragmentation. This effect is statistically significant. The magnitude of this effect, relative to the effect of track density, has been explored.

An ideal statement of the effect of fragmentation would refer to operational measures of effectiveness. For example, it would be useful to quantify the increase in probability of accidentally destroying a friendly target (due to the increased probability of an error in association). However, it is not a straightforward matter to map the degradation in performance that has been studied here onto a degradation in operational effectiveness.

Moreover, even a mere statement about the degradation in association performance must be made with caution, since the association performance of any real operational system will depend on many variables. The problem is both system dependent and situation dependent. A host of design details in any future operational multi-sensor data fusion system will have impact on the extent to which contact fragmentation will affect association, and the nature of that impact is hard to predict.

Intuitively, it seems that most of the impact of the artificial separation of contacts must be mediated by the phenomenon of gate splitting: If all the contacts that can feasibly be associated with a given track (in a given scan) happen to arrive in the same buffer, then the buffer configuration will not cause a problem. If they are split, then a problem may arise. Given a new (hypothetical) data fusion system, we might watch for the typical gate sizes, and watch for the degree to which a sensor's buffering behaviour might split the gates. We will then have a basis for using the results of the present study to make a tentative prediction of the extent to which the grouping of contacts will underlie some of the inevitable association errors.

The “mildly fragmented” case in the present study was inspired by the situation that arose in COMDAT, as described in the introduction. The “severely fragmented” case is offered for the sake of showing how a fusion system would be impacted by a higher degree of gate splitting. There does not appear to be any reason for the designers of a sensing device to anticipate that their device might be used in a fusion system that assumes naively that any contacts that ought to compete with each other will always arrive together in the same message. Thus, the higher degree of gate splitting represented by the “severely fragmented” case deserved some investigation.

It is not clear what value for the odds-ratio threshold λ_0 is large enough to matter. But the degree of degradation in performance caused by the severely fragmented case, at least, appears to be large enough to cause concern. On the other hand, the “mildly fragmented” case appears to be mild in its effects. To the extent that the tracking radar systems are modeled faithfully enough in the present study, there does not appear to be a strong case for the necessity of any modification to the COMDAT system.

It is clear that some attention needs to be paid to this problem. The processing of contacts by an automated data fusion system should not, in general, be designed to occur immediately, whenever any contacts become available. The proper consideration of all hypothetical associations may require the comparison between the newly arrived contacts and some other contacts that are yet to arrive. The messaging behaviour of the source of contact data must be taken into account. If the contacts are too fragmented, then procedures must be implemented to take care of the problem. If nothing else, the simplest kind of buffer management module (as described in section 4), using the time difference as a measure of proximity between contacts, will be practically equivalent to the transformation of a source of highly fragmented data into a source of non-fragmented data. Even with a relatively severe time limit imposed on the retention of contacts in the holding pool, thus avoiding the problem of processing delay, such a module will bring severely fragmented data together enough to be essentially equivalent to the "mildly fragmented" case that was studied here. The results of table 2 suggest that even this small amount of meddling with highly fragmented data can be of significant benefit.

The present study applies directly to fusion systems that make immediate, firm association decisions. The manner in which contacts are grouped will also be relevant to systems of other kinds, for essentially the same reasons, although the impact of contact buffering on those other systems may be of lesser magnitude. A similar study, applied to a system that uses a multiple-hypothesis approach to association and tracking, represents a possible direction for future work.

6. References

1. Y. Bar-Shalom, and X.-R. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*, Yaakov Bar-Shalom, 1995.

2. S. S. Blackman, *Multiple-Target Tracking with Radar Applications*, Artech House, 1986.

3. B. A. McArthur, private communication.

4. B. A. McArthur, A. T. Ashley, É. Bossé, D. Beevis, “The Command Decision Aid Technology (COMDAT) Project”, TTCP Symposium on Battlespace Management, May 1999.

5. D. J. Peters, *A Practical Guide to Level One Data Fusion Algorithms*, DREA TM 2001-201, December 2001.

6. J. L. Devore, *Probability and Statistics for Engineering and the Sciences*, Brooks/Cole, 1982.

List of acronyms

ADT	Automatic Detection and Tracking
BMM	Buffer Management Module
COMDAT	Command Decision Aid Technology
DRDC	Defence Research & Development Canada
ESM	Electronic Support Measures
GCCS	Global Command and Control System
IFF	Identification: Friend or Foe
TD	Technology Demonstration

Distribution list

INTERNAL DISTRIBUTION

6 – Library

1 – J. S. Kennedy

1 – B. A. McArthur

1 – D. J. Peters (author)

EXTERNAL DISTRIBUTION

1 – DRDKIM
NDHQ, MGen. George R. Pearkes Bldg.
101 Col. By Drive
Ottawa, ON K1A 0K2

1 – DMRS 6
NDHQ, MGen. George R. Pearkes Bldg.
101 Col. By Drive
Ottawa, ON K1A 0K2

1 – DMSS 8-9
NDHQ, MGen. George R. Pearkes Bldg.
101 Col. By Drive
Ottawa, ON K1A 0K2

1 – Maritime Research Coordination Cell
NDHQ, MGen. George R. Pearkes Bldg.
101 Col. By Drive
Ottawa, ON K1A 0K2

1 – É. Bossé
DRDC Valcartier
2459 Pie-XI Blvd North
Val-Belair, QC G3J 1X5

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (the name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.)

DRDC Atlantic

2. SECURITY CLASSIFICATION (overall security classification of the document including special warning terms if applicable).

UNCLASSIFIED

3. TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C,R or U) in parentheses after the title).

The Simultaneous Availability of Competing Contacts (U)

4. AUTHORS (Last name, first name, middle initial. If military, show rank, e.g. Doe, Maj. John E.)

Peters, Daniel J.

5. DATE OF PUBLICATION (month and year of publication of document)

December 2005

6a. NO. OF PAGES (total containing information Include Annexes, Appendices, etc).

21

6b. NO. OF REFS (total cited in document)

6

7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered).

Technical Memorandum

8. SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include address).

Defence R&D Canada – Atlantic
PO Box 1012
Dartmouth, NS, Canada B2Y 3Z7

9a. PROJECT OR GRANT NO. (if appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant).

11bg

9b. CONTRACT NO. (if appropriate, the applicable number under which the document was written).

10a ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique to this document.)

DRDC Atlantic TM 2005-260

10b OTHER DOCUMENT NOs. (Any other numbers which may be assigned this document either by the originator or by the sponsor.)

11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

(x) Unlimited distribution
( ) Defence departments and defence contractors; further distribution only as approved
( ) Defence departments and Canadian defence contractors; further distribution only as approved
( ) Government departments and agencies; further distribution only as approved
( ) Defence departments; further distribution only as approved
( ) Other (please specify):

12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected).

13. ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual).

Algorithms for making association decisions in data fusion systems assume the simultaneous availability of the contacts that compete with each other for association with any given track. A data fusion system that processes all contacts on arrival is flawed, in principle, for it assumes that potentially competing contacts will always be sent in the same message by the data source in question. This assumption leads to errors in association that would not otherwise arise. Such errors have been observed in the COMDAT TD project, thus motivating a more detailed examination of the incidence of such errors.

An experiment was conducted, using simulated aerial tracks and simulated radar contacts, to examine the effects of the fragmentation of contacts (that is, the distribution of contacts among separate messages) on the performance of a single-scan data association algorithm. The odds of correct association are significantly degraded, by more than a factor of two, by highly fragmented data. However, a milder degree of fragmentation, roughly representative of the tracking radar of a HALIFAX class frigate, causes a much milder degradation in the odds of correct association. These results do not lead to any recommended changes to the COMDAT system. However, the fragmentation of contacts must be considered, in general, in the design of any automated data fusion system.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title).

data fusion, data association, target tracking

This page intentionally left blank.

Defence R&D Canada

Canada's leader in defence
and National Security
Science and Technology

R & D pour la défense Canada

Chef de file au Canada en matière
de science et de technologie pour
la défense et la sécurité nationale

DEFENCE R&D DÉFENSE

www.drdc-rddc.gc.ca