Copy No. _____

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

DEFENCE R&D DÉFENSE

Tactical Picture Compilation Performance Specifications in HMCCS

D. J. Peters

Defence R&D Canada – Atlantic

Technical Memorandum
DRDC Atlantic TM 2007-369
December 2008

Canada

This page intentionally left blank.

Tactical Picture Compilation Performance Specifications in HMCCS

D. J. Peters

Defence R&D Canada - Atlantic

Technical Memorandum

DRDC Atlantic TM 2007-369

December 2008

Author

Original signed by D. J. Peters

D. J. Peters

Approved by

Original signed by F. Desharnais

F. Desharnais

Head, Maritime Information and Combat Systems Section

Approved for release by

Original signed by R. Kuwahara for

C. Hyatt

DRP Chair

Abstract

The HALIFAX Class Modernised Command and Control System (HMCCS) project was a component of the Modernised HALIFAX Class (MHC) project before the latter project was reorganised and renamed in 2007. This document is concerned with a subset of the specifications that were written in order to define the capabilities of the HMCCS; specifically, those specifications that were concerned with the performance of the system in terms of tactical picture compilation.

Each specification was expressed in terms of a threshold value of some Measure of Performance (MOP), to be achieved relative to a predefined scenario and an associated data set. This document details the ten specification statements that were developed, as well as the process used to derive the MOP thresholds. The threshold derivation process was based upon simple data fusion methods, in most cases involving “nearest neighbour” association rules and a Kalman Filter for tracking.

Résumé

Le projet de modernisation du système de commandement et de contrôle de la classe HALIFAX (MSCCH) constituait un élément du projet de modernisation de la classe HALIFAX (MCH) avant que ce dernier soit réorganisé et qu’une nouvelle appellation lui soit donnée, en 2007. Le présent document porte sur un sous-ensemble des spécifications qui avaient été préparées pour définir les capacités du projet MSCCH, plus précisément les spécifications qui avaient trait à la performance du système en termes de compilation de la situation tactique.

Chaque spécification était exprimée en termes d’une valeur de seuil de certains critères de rendement, qu’il fallait atteindre par rapport à un scénario prédéfini et une série de données connexes. Le présent document donne des précisions sur les dix énoncés de spécification qui ont été formulés, ainsi que sur le processus qui a permis de dériver les valeurs de seuil des critères de rendement. Le processus de dérivation des valeurs de seuil était fondé sur de simples méthodes de fusion de données qui, dans la plupart des cas, comportaient des règles d’association « du plus proche voisin » et un filtre de Kalman pour la poursuite.

This page intentionally left blank.

Executive summary

Introduction

Before being reorganised and renamed in 2007, the Modernised HALIFAX Class (MHC) project consisted of several smaller projects. One of these, the HALIFAX Class Modernised Command and Control System (HMCCS) project, was intended to procure a new Command and Control System (CCS).

One of the functions of a CCS is the formation of a tactical picture through the fusion of data from a variety of sources. DRDC Atlantic provided support in the development of specifications for the tactical picture compilation function, formulated in terms of the quality of the tactical picture produced by the system.

Procurement practice required that specifications be stated in terms of objectively testable system capabilities. Given the wide variety of operational situations possible, this constraint posed a formidable task for the writing of specifications for the picture compilation function. Further complicating the task was uncertainty about the characteristics of the sensors that would be providing information to the system.

Results

Ten generic situations involving various combinations of sensor data were designed to represent expected operational situations. For each scenario, quantitative Measures of Performance (MOPs) were defined, and threshold values for acceptable performance were determined using basic picture compilation techniques applied to synthetic data sets. A key feature of this process was the development of reasonable data sets to be treated as input data, thereby enabling any candidate system to create track data that would then be evaluated with respect to the appropriate MOPs.

Significance

This work implemented a process to quantify the extremely complex and vital CSS function of picture compilation. While the time constraints imposed by project deadlines limited the final discriminatory power of the technique, this work could serve as a starting point for future project teams. By improving the sophistication of the baseline fusion and tracking algorithms, it may be possible to provide more robust evaluation criteria.

Future Work

The issue of picture compilation performance is ongoing and is closely related to the development of scenarios for training and for combat system concept experimentation. The general techniques of this study may be further explored in these contexts.

Peters, D. J. 2008. Tactical Picture Compilation Performance Specifications in HMCCS. DRDC Atlantic TM 2007-369. Defence R&D Canada - Atlantic.

Sommaire

Introduction

Avant d’être réorganisé et de recevoir une nouvelle appellation, en 2007, le projet de modernisation de la classe HALIFAX se composait de plusieurs projets de petite envergure. L’un d’entre eux, le projet de modernisation du système de commandement et de contrôle de la classe HALIFAX C (MSCCH), visait l’acquisition d’un nouveau système de commandement et de contrôle (SC2).

L’une des fonctions de tout SC2 est la formation de la situation tactique par la fusion de données provenant de toute une gamme de sources. RDDC Atlantique a accordé son appui à l’élaboration des spécifications de la fonction de compilation de la situation tactique, formulée en termes de qualité de la situation tactique produite par le système.

En vertu des pratiques en matière d’acquisition, les spécifications devaient être formulées en termes des capacités de systèmes pouvant objectivement être mis à l’essai. Étant donné la grande diversité des situations opérationnelles possibles, cette contrainte a posé un défi formidable pour la rédaction des spécifications de la fonction de compilation de la situation. Ce qui compliquait davantage le tout, c’est que la tâche comportait des incertitudes au sujet des caractéristiques des capteurs qui fourniraient des données au système.

Résultats

Dix situations génériques comportant diverses combinaisons de données de capteur ont été élaborées pour représenter des situations opérationnelles escomptées. Pour chaque scénario, des critères de rendement quantitatif ont été définis, et des valeurs de seuil de rendement acceptable ont été déterminées à l’aide de techniques fondamentales de compilation de la situation appliquées à des séries de données synthétiques. Un élément clé de ce processus a été l’élaboration de séries de données raisonnables devant servir de données d’entrée, ce qui permettait donc à tout système potentiel de créer des données de poursuite, qui seraient alors évaluées par rapport aux critères de rendement appropriés.

Portée

Les travaux ont permis de mettre en œuvre un processus de quantification de la fonction extrêmement complexe et vitale de SC2 qu’est la compilation de la situation. Bien que les contraintes de temps imposées par les délais du projet aient limité le pouvoir de discernement final de la technique, les travaux pourraient servir de point de départ pour les futures équipes chargées de projets. En améliorant le raffinement des algorithmes de poursuite et de fusion de référence, il pourrait être possible de fournir des critères d’évaluation plus robustes.

Recherches futures

Le rendement de la compilation de la situation est une question qui demeure, et qui est étroitement rattachée à la formulation de scénarios pour la formation et la mise à l’essai des

concepts des systèmes de combat. Les techniques générales de la présente étude pourraient être étudiées davantage dans ces contextes.

Peters, D. J. 2008. Spécifications du rendement de la compilation de la situation tactique dans le cadre du projet MSCCH. DRDC Atlantic TM 2007-369. Defence R&D Canada - Atlantic.

This page intentionally left blank.

Table of contents

List of figures

Acknowledgements

Thanks go to the participants in the discussions that formed the context for the exercise of scenario creation that is presented in this document – in particular: Bruce McArthur, Nelson McCoy and LCdr Cliff Stroud.

In addition, my hearty thanks go to Mark Hazen, for his very useful comments on the early drafts of this document.

This page intentionally left blank.

1. Introduction

Plans are underway for the modernisation of the HALIFAX class frigates, as embodied by the HALIFAX Class Modernisation (HCM) project, formerly known as the Modernised HALIFAX Class (MHC) project. Originally, the project was organised as several smaller projects, each concerned with a different aspect of the functioning of the ships. The project specifically concerned with the Command and Control System (CCS) was known as the HALIFAX Class Modernised Command and Control System (HMCCS) project [1]. In 2007 the plans were changed, and the renamed HCM project was reorganised as an integrated whole.

In 2005 and 2006, before the change of plans, a significant amount of effort went into defining the specifications for the new CCS. Some of the specifications pertained to the performance of the new system. Such specification statements faced an interesting challenge, as described in Section 1.1.

The performance of the system's tactical picture compilation capability, in particular, is the subject of those specification statements with which the present document is concerned. Ten specification statements of this kind were written in 2006. The purpose of this document is to present these ten specifications, to describe the process that led to them, and to discuss the lessons learned from this process, for the sake of supporting future efforts of this kind. It was originally expected that process underlying these specification statements, suitably refined, might serve as the basis for the performance specification statements that would be included in the Request For Proposals for the HMCCS project.

The approach that was taken in formulating the specification statements is summarised in Section 1.2. Section 1.3 presents an outline of the remainder of this document.

1.1 Constraints on specification statements

Many specification statements were written in order to define the capabilities of the new system. Some specifications dealt with the kinds of data that the system should be able to process, some dealt with the rules that govern the circumstances under which the system tracks should be updated, and some dealt with the way in which the resulting tactical picture should be presented to the user. But it is also important for the tactical picture generated by the system to correspond (in some sense) to reality. This requirement is what we mean by the "performance" of the system's tactical picture compilation capability. The challenge was to find a way to express this requirement in specification statements.

Following standard procurement practice, the specifications were to be written in the form of so-called "*shall* statements" – that is, statements of the form "The HMCCS shall …" – and written with sufficient precision so that the compliance of a proposed system with the specifications could be judged objectively. This compliance was to be viewed in a strictly binary manner: Either the proposal fulfills all the specifications, or it does not.

The “*shall* statement” constraint posed a problem for the definition of the performance specifications. Early drafts of some of the pertinent statements used such words as *accurately* or *correctly*, when describing what the system ought to do. But we cannot specify that the picture must be free of errors, for this ideal is impossible to achieve. In short, it was necessary to find a way to express a requirement for acceptable performance without demanding the impossible.

Performance can often be quantified according to various metrics (known as Measures of Performance, or MOPs). In order to make a performance specification statement sufficiently precise, it is usually necessary to define a MOP for that statement. Since each specification statement was required to be worded in such a way that a candidate system would either pass or fail (objectively), every statement with a MOP needed a threshold value for that MOP, defining the boundary between acceptable and unacceptable performance.

In order to be meaningful, a specification statement of this form – involving a MOP and a corresponding threshold – must be made within the context of a specific scenario or a class of scenarios. Ideally, a MOP threshold statement would be expressed statistically, with respect to a large number of scenarios generated randomly out of some kind of “scenario space”. But to define such a “scenario space” precisely, in a way that (demonstrably) leads to more satisfactory specifications than a single scenario, was judged to be impractical, given time and budget constraints. So the tactical picture compilation performance specification statements in HMCCS were each expressed in terms of a single simulated scenario.

Moreover, a scenario cannot be fully specified without defining the characteristics of the sensors involved, but the HMCCS specifications were to be written without knowledge of the characteristics of the sensor suite of the future (upgraded) HALIFAX class frigates.

At the time of this work, the process of finding an appropriate rationale for setting the performance thresholds was in a very early stage. Thus, for the purpose of this work, the MOP thresholds were set to values that were demonstrably achievable by fairly simple methods.

1.2 Creating the performance specification statements

Ten complete performance specification statements were developed, pertaining to various aspects of tactical picture compilation. For each such statement, a scenario was designed.

Five of the ten scenarios are concerned with tracking accuracy. Scenarios One and Two deal with the simple situation of tracking a single target with a single radar, either 2D (in the case of Scenario One) or 3D (in the case of Scenario Two). Scenario Eight deals with the tracking of a single target in 3D by a suite of heterogeneous sensors. Scenario Nine presents a more challenging situation, wherein three targets are tracked in 3D in the absence of 3D sensors, by fusing data from a 2D radar (which measures bearing and range) and an Infra-red Search and Track (IRST) sensor (which measures bearing and elevation angle). In Scenario Ten, a single target is tracked in 2D by way of triangulation among three bearing-only sensors at different locations.

Two other scenarios are concerned with association performance. Scenario Three uses a 2D radar alone, while Scenario Seven uses a 2D radar plus a bearing-only Electronic Support Measures (ESM) sensor. There are three targets in each scenario.

The remaining scenarios explore other issues: The timely creation of a track for a newly-launched missile (Scenario Four); the timely creation of tracks and the avoidance of spurious tracks in a high-density, high-clutter situation (Scenario Five), and the ability of the system to self-correct after an association error arising when two targets merge and split (Scenario Six).

These ten scenarios arose from the discussions within the HMCCS project in the spring and summer of 2006. It was originally envisaged that further refinement to these scenarios would take place, and that additional scenarios would be developed as required.

The scenarios were coded in Mathematica[1]. Simple models of targets and sensors were used in order to generate contacts, and simple data fusion methods were used to generate tracks for scenarios that require track input.

Wherever possible, simple data fusion methods were applied to the input data in order to find an achievable level of performance. Chapter 2 outlines the main methods used. The level of performance thereby achieved was expressed as a threshold, in terms of the chosen MOP(s) – the MOP definitions having been refined as much as necessary to remove ambiguity.

Several runs were executed for each scenario, each run generating a new set of input data for the same ground truth, in order to get an informal sense of what results would be typical for the simple methods that were used. Whenever there was a clear binary distinction between "good" and "bad" performance (for example, having or not having the correct number of system tracks confirmed over the course of the run), the chosen run was one in which the simple methods gave a good result. In most other cases, the chosen run was one for which the relevant MOP results were close to their mean values over all the runs.

For each scenario, a file representing the input data (derived from the chosen run) was created, with the intent of providing it along with the specifications document. Also, it was intended that the MOPs could be calculated by anyone. Therefore, whenever the chosen MOP was defined with reference to the "truth" underlying the scenario (as it was in most cases), a "ground truth" file was provided as well.

1.3 Outline of Chapters 3 through 12

As indicated briefly in Section 1.2, a specification statement was constructed for each of ten scenarios. The proposed specification statements, intended for the main part of the HMCCS specification document, are presented in the opening paragraphs of Chapters 3 through 12, accompanied in each case by a high-level description of the scenario.

In order to achieve the required degree of precision, each performance specification requires so much detail that it was impossible to present each one fully in a single sentence. Therefore most of the content of the specification definition, in each case, was separated from the

[1] *Mathematica* is a registered trademark of Wolfram Research Inc.

corresponding main specification statement and arranged with the intention of putting it into an annex at the back of the specification document. The first numbered section (Section *n*.1 for Chapter *n*) of each scenario chapter presents the proposed (prototype) text of the corresponding annex, in (essentially) the version that was achieved by August of 2006. Because these sections present the full text of the corresponding prototype annexes, there is a certain amount of repetition among them (especially with respect to the MOP definitions).

The second section (*n*.2) of each scenario chapter presents supplementary information to help the reader in interpreting the scenarios, and the third section (*n*.3) discusses the process by which the performance threshold values were determined.

1.4 Note on terminology

A *contact* is a single state estimate at a single time, derived from a set of sensor data. A package of data that is derived from one or more contacts associated with each other, representing an assertion or a potential assertion about a real (or simulated object), is called a *track*.

Some sensors come with built-in data fusion features, and thus provide the user with track data in place of contact data. Therefore, in the context of a discussion of candidate data fusion systems, it is useful to distinguish between *sensor tracks* (which can be input to the candidate system) and *system tracks* (created and maintained within the candidate system). In the HMCCS project, the word *entity* was used to refer to a system track, while the word *track* was used only to refer to a sensor track (i.e. an input track). The remainder of this document follows this terminology.

Insofar as entities (system tracks) can have varying degrees of confidence, it is assumed that any candidate system will have a threshold of confidence above which any entity is to be interpreted as an assertion about the existence and the characteristics of some real (or simulated) object, while those entities that are below this threshold function as tentative suggestions awaiting further confirmation. The former entities, called *confirmed entities*, are the only entities that are considered when assessing the performance of the system.

2. Common tracking and association methods

A common set of tracking and association methods was used in several of the scenarios – both for creating input track data and for finding an achievable threshold for the system performance. The purpose of this chapter is to describe these methods. (A few of the scenarios have peculiar characteristics that are not covered by this chapter. See the individual scenario descriptions in Chapters 3 through 12.)

It is assumed that each sensor will provide messages consisting of one or more contacts (or track updates). These messages are input to the data fusion system. Each input message triggers an iteration of association and tracking.

2.1 Association rules

The first step in that iteration is to judge the feasibility of association of each new contact (or track update) with each entity. (The following description assumes that the input message consists of contacts. For the scenarios that take track updates as input, see the individual scenario descriptions.)

Sensor characteristics vary from one scenario to another (see the individual scenario descriptions), but in most scenarios the contact positions are given in terms of range and bearing (and possibly elevation angle). The sensor's nominal uncertainties in each given dimension are transformed into Cartesian coordinates (by a simple linearised transformation) and are treated as Gaussian for purposes of judging the likelihood of association of that contact with each track (considered independently). (See Section 2.3 for the cases where a contact was missing one or more positional dimensions.)

The association of an entity to a contact is judged to be feasible if the association hypothesis could not be rejected with a statistical significance of 0.1, based on a model according to which the acceleration of the target is treated as Gaussian white noise with a mean of zero. (This level of significance is arbitrary, but it is based on a tradeoff between the need to reject false alarms and the need to maintain tracking through manoeuvres [2].) For most of the scenarios, the targets are aerial, and the magnitude of the modelled noise (representing random acceleration) corresponds to an increase in velocity variance of $1500 m^2/s^2$ each second.

Association decisions (between contacts and entities) are made according to the following rules:

(1) At most one new contact is associated with each entity.
(2) At most one entity is associated with each new contact.
(3) An association can only be made if it has been judged to be feasible.
(4) The number of associations made is maximised, within the constraints of the first three rules.
(5) The number of associations involving entities that have been updated at least once since they were created is maximised within the constraints of the first four rules.

(6) The chosen combination of associations is the one that maximises the product of the corresponding likelihood densities (calculated from the model described in the previous paragraphs).

2.2 Entity creation, maintenance, and deletion

If all associations with a given contact are rejected, a new entity is created from that contact (see Section 2.3 for exceptions). Its initial position is as given by the contact's position, and its initial velocity is set to zero. The covariance matrix expresses the uncertainty in position as dictated by the sensor's nominal uncertainty in each dimension, using (where necessary) a linearised transformation from polar to Cartesian coordinates, and an arbitrary uncertainty in velocity. For most scenarios, the latter uncertainty was 800m/s in each velocity dimension. Uncertainties in position and velocity are initially uncorrelated.

When an entity is associated with a contact, it is updated by that contact according to a standard Kalman Filter algorithm, based on the same motion model that was described in connection with the association rules.

The rules for initiating and updating entities are modified in the case of contacts that are missing one or more positional dimensions. See Section 2.3.

An entity is deleted if it receives no associations within a certain time limit. A newly-created entity must be selected for association within a time τ of its creation in order to escape deletion, where $\tau = 1.25$ times the shortest rotational period among all the sensors of the scenario. If it receives one such update, it must receive two further updates within a time 3τ of this update in order to escape deletion. If it does so, it becomes a confirmed entity, and from then on it will be deleted only if it receives no update for an interval of 5τ.

2.3 Adapting the association and track update procedures to the case of contacts with missing dimensions

Some of the scenarios in this document – those in which no sensors provide vertical position data – involve tracking in two dimensions, while the rest involve tracking in three dimensions. In both cases, there are scenarios in which some contributing sensors provide data of less than full dimensionality, such as bearing-only ESM data. Sensor data of incomplete dimensionality require special consideration, both in association algorithms and in entity update algorithms. In the work reported here, a simple procedure was developed to handle this problem.

For the purpose of being compared to an entity for the sake of association decisions, or for the purpose of updating an entity, the missing (polar) dimensions of a contact's position are given values equal to the corresponding values of the entity's (extrapolated) position, transformed into polar coordinates. The contact is given a nominal standard deviation in the newly-added dimension that is equal to ten times the standard deviation corresponding to that dimension according to the entity's (extrapolated) covariance matrix. This large amount of uncertainty, in the dimensions that are not actually measured, ensures that the contact will be given very little weight, with respect to those dimensions, in the Kalman Filter update equations. (This method, while not standard, has appeared to work well in informal simulation experiments conducted at DRDC Atlantic [3].)

For the purpose of starting a new entity with an unassociated contact, the rules are modified as follows: A contact that lacks range data is simply discarded; no new entity is created. (See Section 12.3 for an exception.) A contact that lacks elevation angle data, in a three-dimensional tracking scenario, is given a nominal elevation angle of zero with a nominal standard deviation of one radian.

3. Scenario One – 2D tracking performance

The first scenario tests two-dimensional tracking accuracy, in the case of a single manoeuvring aerial target sensed by a two-dimensional radar.

The proposed specification statement read as follows: "The HMCCS, if operating in automatic mode and subjected to the input given in the file `s01input.csv`, interpreted according to Annex X.1[2], shall produce exactly one confirmed entity, whose mean absolute error in position, mean absolute error in velocity, and RMS relative error shall be equal to or better than the thresholds specified in that annex."

The proposed text of the annex is reproduced here in Section 3.1. (An input file and a ground truth file, respectively named `s01input.csv` and `s01gtr.csv`, were intended to be provided with the annex.) Section 3.2 provides additional description of the scenario, while Section 3.3 discusses the design of the scenario and the choice of the performance threshold.

3.1 Specification text

General Comments

This scenario tests the tracking accuracy and consistency of the candidate system in response to 2D radar contacts, representing detections of a single manoeuvring aerial target.

Data Requirements

The following data, from each update of an entity in response to an associated contact, must be stored in order to assess the performance:
(1) The estimated position and velocity, in two dimensions.
(2) The estimated covariance matrix corresponding to the estimated position and velocity, or an alternative expression of uncertainty from which such a covariance matrix may be derived.

Scenario Input

Each line of the Input File consists of three numbers, representing a radar contact. The three numbers are
(1) contact time (in seconds, relative to the beginning of the scenario),
(2) contact range (in metres),
(3) contact bearing (in degrees, clockwise from north).

The contacts shall be interpreted according to the following assumptions:
(1) The contacts are produced by a 2D tracking radar.
(2) The contacts are communicated to the HMCCS, one by one, in the order given in the

[2] "X.1" was merely a placeholder for whatever alphabetical or numerical indicator would have been assigned to that annex of the HMCCS specification document. A similar comment applies to the corresponding statements in all the other scenarios.

Input File.
(3) The radar's uncertainties in range and bearing are given by nominal standard deviations of 40 metres and 1.0 degree, respectively.
(4) All targets are aerial.
(5) We may be dealing with multiple targets and clutter. (In fact, the contacts represent a single target and no clutter, but this background information may not be used by the system.)

Performance Requirements

In order to meet the specification, the following results are required:
(1) Only one confirmed entity shall be produced in response to the data.
(2) The mean absolute error in position (see below) shall be less than or equal to 200 m.
(3) The mean absolute error in velocity (see below) shall be less than or equal to 90 m/s.
(4) The RMS relative error in position and velocity (see below) shall be less than or equal to 2.0.

Definitions

The "absolute error in position", at each entity update, is the scalar magnitude of the position error vector:

$$E_{\text{pos}} = \left\| \begin{pmatrix} \hat{x} - x \\ \hat{y} - y \end{pmatrix} \right\|,$$

where x and y are the true coordinates of the target at that time (as given in the Ground Truth File below), and $\hat{x}$ and $\hat{y}$ are the estimates of these coordinates according to the entity. The "mean absolute error in position" is the mean of E_{pos} for all entity updates (of the confirmed entity), where each update is given equal weight.

Similarly, the "absolute error in velocity", at each entity update, is the scalar magnitude of the velocity error vector:

$$E_{\text{vel}} = \left\| \begin{pmatrix} \hat{v}_x - v_x \\ \hat{v}_y - v_y \end{pmatrix} \right\|,$$

where v_x and v_y are the true velocity components of the target at that time, etc. The "mean absolute error in velocity" is the mean of E_{vel} for all entity updates (of the confirmed entity), where each update is given equal weight.

The "squared relative error in position and velocity", at each entity update, is the quantity

$$D = \begin{pmatrix} \hat{x} - x & \hat{y} - y & \hat{v}_x - v_x & \hat{v}_y - v_y \end{pmatrix} \mathbf{P}^{-1} \begin{pmatrix} \hat{x} - x \\ \hat{y} - y \\ \hat{v}_x - v_x \\ \hat{v}_y - v_y \end{pmatrix},$$

where $\mathbf{P}$ is the 4-by-4 covariance matrix representing the uncertainty claimed by the entity, with respect to its estimates of x, y, v_x and v_y. The "RMS relative error in position and velocity" is the square root of the mean of D for all entity updates (of the confirmed entity), where each update is given equal weight.

Each line of the Ground Truth File gives the "ground truth" corresponding to the matching line number of the Input File. The four numbers in each line are
(1) x position (in metres, east of the radar position),
(2) y position (in metres, north of the radar position),
(3) x-component of velocity (in m/s),
(4) y-component of velocity (in m/s).

For this scenario, it is permissible to interpret the x and y values in the Ground Truth File as "slant x" and "slant y" (that is, as xr/ρ and yr/ρ, where $\rho^2 = x^2 + y^2$ and $r^2 = \rho^2 + z^2$), and to perform tracking in ("slant x", "slant y") space. Equivalently, it may be assumed that all targets are at zero elevation, despite being "aerial" in nature.

3.2 Additional description

Many sets of contacts were generated (see Section 3.3). The set that was selected for the specification statement is illustrated in Figure 1, along with the corresponding ground truth. The target was constructed as if it were flying about at $z = 0$. (Hence the note at the end of Section 3.1.) The target trajectory was designed to include some fairly hard turns (over 4g) in both directions.

3.3 Threshold determination

After defining the target trajectory, fifty runs of a simple simulation were performed. In each run, a random set of contacts was generated. The simulated radar had a rotation period of 2 seconds. Each time its beam swept past the target, it was decided randomly (based on a detection probability of 0.95) whether a contact would be generated. If so, random normally-distributed errors were generated for both range and bearing, with the standard deviations indicated in Section 3.1. The time-tag of each resulting contact was calculated from the bearing measurement, so that the times and bearings would be strictly consistent with the rotational period.

A simple algorithm for 2D association and tracking was implemented in each run. At the start, there are no entities. For each contact received from the simulated radar, an iteration of association and tracking is performed, as described in Chapter 2. (The parameter τ was set to 2.5 seconds.)

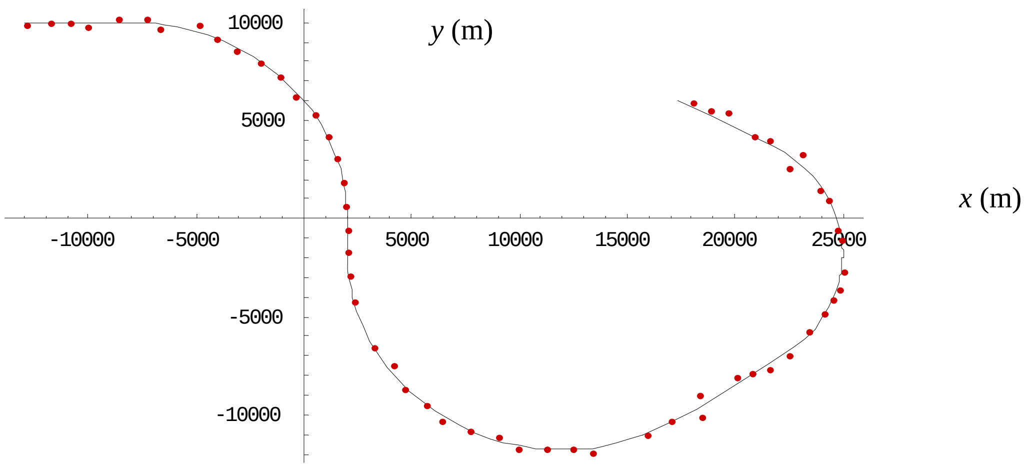

Figure 1. *Ground truth and radar contacts in Scenario One.*

Ground truth is shown in black, radar contacts in red. The target starts in the upper left.

The performance of this algorithm, in terms of the three MOPs defined in Section 3.1, was averaged over the fifty runs. (The average MOP scores were 193.1 m, 93.6 m/s, and 1.98.) Additional runs were then performed, in order to search for a case in which the simple algorithm achieved a score that was close to this average in each of the three MOPs. One such run was selected for the specification scenario. (The MOP scores for this run were 198.5 m, 87.3 m/s, and 1.99.) The input file provided with the specification statement consists of the contacts that were generated in the selected run. The ground truth file gives the true position and velocity of the simulated target at the times corresponding to the contacts in the input file. The threshold values for the specification were set to be slightly worse than the scores achieved in the selected run by the simple algorithm, in accordance with the goal of making sure that acceptable performance would be achievable.

The chosen MOPs for tracking accuracy are commonly used [4]. The absolute error MOPs are self-explanatory, while the relative error in position and velocity provides a measure of the realism of the uncertainty claimed by the data fusion system as manifested in the entity's covariance matrix.

The search for a suitable run for the specification statement was also subjected to the constraint that the simple algorithm must produce only one confirmed entity, which must be maintained throughout the rest of the run. This condition had been met in 48 of the original 50 runs, so it caused no difficulty.

4. Scenario Two – 3D tracking performance

The second scenario tests three-dimensional tracking accuracy, in the case of a single manoeuvring aerial target sensed by a three-dimensional radar.

The proposed specification statement read as follows: "The HMCCS, if operating in automatic mode and subjected to the input given in the file `s02input.csv`, interpreted according to Annex X.2, shall produce exactly one confirmed entity, whose mean absolute error in position, mean absolute error in velocity, and RMS relative error shall be equal to or better than the thresholds specified in that annex."

The proposed text of the annex is reproduced here in Section 4.1. (An input file and a ground truth file, respectively named `s02input.csv` and `s02gtr.csv`, were intended to be provided with the annex.) Section 4.2 provides additional description of the scenario, while Section 4.3 discusses the choice of performance threshold.

4.1 Specification text

General Comments

This scenario tests the tracking accuracy and consistency of the candidate system in response to 3D radar contacts, representing detections of a single manoeuvring aerial target.

Data Requirements

The following data, from each update of an entity in response to an associated contact, must be stored in order to assess the performance:
(1) The estimated position and velocity, in three dimensions.
(2) The estimated covariance matrix corresponding to the estimated position and velocity, or an alternative expression of uncertainty from which such a covariance matrix may be derived.

Scenario Input

Each line of the Input File consists of four numbers, representing a radar contact. The four numbers are
(1) contact time (in seconds, relative to the beginning of the scenario),
(2) contact range (in metres),
(3) contact bearing (in degrees, clockwise from north),
(4) contact elevation angle (in degrees).

The contacts shall be interpreted according to the following assumptions:
(1) The contacts are produced by a 3D tracking radar.
(2) The contacts are communicated to the HMCCS, one by one, in the order given in the Input File.
(3) The radar's uncertainties in range, bearing, and elevation angle are given by nominal standard deviations of 40 metres, 1.0 degree, and 1.0 degree, respectively.

(4) All targets are aerial.
(5) We may be dealing with multiple targets and clutter. (In fact, the contacts represent a single target and no clutter. But this background information may not be used by the system.)

Performance requirements

In order to meet the specification, the following results are required:
(1) Only one confirmed entity shall be produced in response to the data.
(2) The mean absolute error in position (see below) shall be less than or equal to 260 m.
(3) The mean absolute error in velocity (see below) shall be less than or equal to 80 m/s.
(4) The RMS relative error in position and velocity (see below) shall be less than or equal to 2.0.

Definitions

The "absolute error in position", at each entity update, is the scalar magnitude of the position error vector:

$$E_{\text{pos}} = \left\| \begin{pmatrix} \hat{x} - x \\ \hat{y} - y \\ \hat{z} - z \end{pmatrix} \right\|,$$

where x, y, and z are the true coordinates of the target at that time (as given in the Ground Truth File below), and $\hat{x}$, $\hat{y}$, and $\hat{z}$ are the estimates of these coordinates according to the entity. The "mean absolute error in position" is the mean of E_{pos} for all entity updates (of the confirmed entity), where each update is given equal weight.

Similarly, the "absolute error in velocity", at each entity update, is the scalar magnitude of the velocity error vector:

$$E_{\text{vel}} = \left\| \begin{pmatrix} \hat{v}_x - v_x \\ \hat{v}_y - v_y \\ \hat{v}_z - v_z \end{pmatrix} \right\|,$$

where v_x and v_y are the true velocity components of the target at that time, etc. The "mean absolute error in velocity" is the mean of E_{vel} for all entity updates (of the confirmed entity), where each update is given equal weight.

The “squared relative error in position and velocity”, at each entity update, is the quantity

$$D = \begin{pmatrix} \hat{x}-x & \hat{y}-y & \hat{z}-z & \hat{v}_x - v_x & \hat{v}_y - v_y & \hat{v}_z - v_z \end{pmatrix} \mathbf{P}^{-1} \begin{pmatrix} \hat{x}-x \\ \hat{y}-y \\ \hat{z}-z \\ \hat{v}_x - v_x \\ \hat{v}_y - v_y \\ \hat{v}_z - v_z \end{pmatrix},$$

where **P** is the 6-by-6 covariance matrix representing the uncertainty claimed by the entity, with respect to its estimates of x, y, z, v_x, v_y, and v_z. The “RMS relative error in position and velocity” is the square root of the mean of D for all entity updates (of the confirmed entity), where each update is given equal weight.

Each line of the Ground Truth File gives the “ground truth” corresponding to the matching line number of the Input File. The six numbers in each line are
(1) x position (in metres, east of the radar position),
(2) y position (in metres, north of the radar position),
(3) z position (in metres, above the horizontal plane of the radar position),
(4) x-component of velocity (in m/s),
(5) y-component of velocity (in m/s),
(6) z-component of velocity (in m/s).

4.2 Additional description

Many sets of contacts were generated (see Section 4.3). The set that was selected for the specification statement is illustrated in Figure 2, along with the corresponding ground truth, plotted as (x, y). Figure 3 shows z as a function of time, for the same data set. The target manoeuvres both left and right (as in Scenario One), and it also climbs and dives.

4.3 Threshold determination

This scenario is the 3D version of Scenario One. Contacts were generated in the same way, except that they also included an elevation angle measurement with a random Gaussian error (with a standard deviation of one degree). Again, the association and tracking algorithm was that of Chapter 2, but now in three dimensions.

Again, there were fifty test runs over which the tracking accuracy MOPs were averaged, followed by additional runs to search for a case in which all three of these MOPs were close to that average. (The mean MOP scores for the fifty runs were 260.2 m, 82.5 m/s, and 2.06. The MOP scores for the selected run were 260.6 m, 76.9 m/s, and 1.99.) The input file provided with the specification statement consists of the contacts that were generated in the selected run. The ground truth file gives the true position and velocity of the simulated target at the times corresponding to the contacts in the input file. Again, the threshold values for the specification were set to be about the same as the scores achieved in the selected run by the simple algorithm.

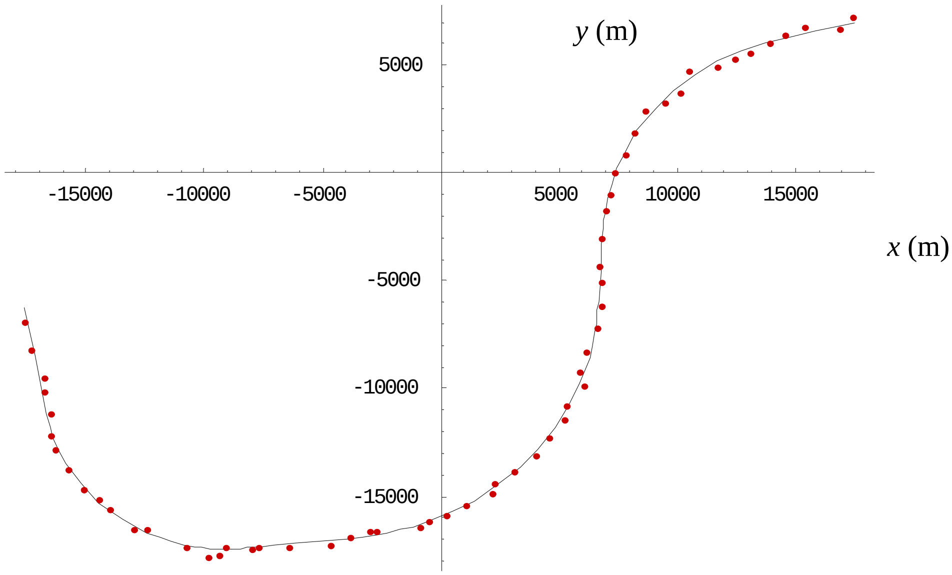

Figure 2. *Ground truth and radar contacts in Scenario Two – 2D positions*

Ground truth is shown in black, radar contacts in red. The target starts at the upper right.

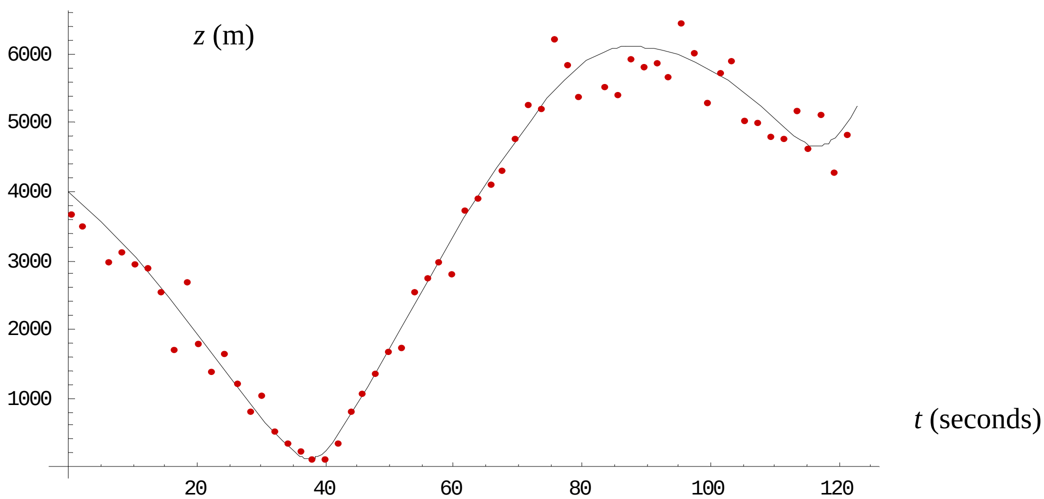

Figure 3. *Ground truth and radar contacts in Scenario Two – altitude (z) versus time (t)*

Ground truth is shown in black, radar contacts in red.

Again, the search for a suitable run for the specification statement was also subjected to the constraint that the simple algorithm must produce only one confirmed entity, which must be maintained throughout the rest of the run. And again, this condition was met in 47 of the original 50 runs, so it caused no difficulty.

5. Scenario Three – Association performance

This scenario tests the data fusion system's ability to make correct associations, in the case of three aerial targets flying close to each other with their paths frequently crossing. The targets are detected by a 2D radar. Contrary to previous scenarios, the radar data includes clutter.

The proposed specification statement read as follows: "The HMCCS, if operating in automatic mode and subjected to the input given in the file `s03input.csv`, interpreted according to Annex X.3, shall produce exactly three confirmed entities, resulting in a weighted average track purity that is equal to or better than the threshold specified in that annex."

The proposed text of the annex is reproduced here in Section 5.1. (An input file and a ground truth file, respectively named `s03input.csv` and `s03gtr.csv`, were intended to be provided with the annex.) Section 5.2 provides additional description of the scenario, while Section 5.3 discusses the choice of performance threshold.

5.1 Specification text

General Comments

This scenario tests the association performance of the candidate system in response to 2D radar contacts, representing detections of three interweaving aerial targets.

Data Requirements

The following data, for each confirmed entity, must be stored in order to assess the performance:
(1) Which contacts were associated with that entity. (Contacts may be identified by their time tags, or by their line numbers in the Input File.)

Scenario Input

Each line of the Input File consists of four numbers, representing a radar contact. The four numbers are
(1) either a "1" or a "0", indicating whether this contact is the first contact in a message or a continuation of a previous message,
(2) contact time (in seconds, relative to the beginning of the scenario),
(3) contact range (in metres),
(4) contact bearing (in degrees, clockwise from north).

The contacts shall be interpreted according to the following assumptions:
(1) The contacts are produced by a 2D tracking radar.
(2) The contacts are communicated to the HMCCS in messages that can consist of up to four contacts each. Each contact that has a "1" in the first column (in the Input File) belongs to a message that arrives later than any message containing any contact that appears earlier in

the file. Each contact that has a “0” in the first column belongs to the same message as the contact that immediately precedes it in the file.

(3) The radar’s uncertainties in range and bearing are given by nominal standard deviations of 25 metres and 1.0 degree, respectively.

(4) All targets are aerial.

(5) We may be dealing with multiple targets and clutter.

Performance Requirements

In order to meet the specification, the following results are required:

(1) Three, and only three, confirmed entities shall be produced in response to the data.

(2) The weighted average track purity (see below) shall be greater than or equal to 0.98.

Definition

Each line of the Ground Truth File contains one number, identifying the target (if any) from which the contact of the corresponding line in the Input File originated. Each number is an integer varying from 0 to 3. A “0” indicates that the contact is a false alarm. Otherwise, the number identifies the target.

Let C_{ij} denote the number of contacts from target j (j varying from 1 to 3) that were associated with entity i (i varying from 1 to 3, representing the three confirmed entities). Let F_i denote the number of false alarms that were associated with entity i. Then the weighted average track purity[3] is

$$WATP = \frac{\sum_i \max_j C_{ij}}{\sum_i \left(F_i + \sum_j C_{ij} \right)}.$$

5.2 Additional description

Ground truth and radar contacts for the selected run of Scenario Three (see Section 5.3) are given in Figure 4. For the sake of this figure, the radar contacts are interpreted as if the elevation angle is always zero. In the vicinity of these targets, the error induced by this assumption is insignificant at the scale of this figure.

[3] See the comments on this MOP in section 5.3.

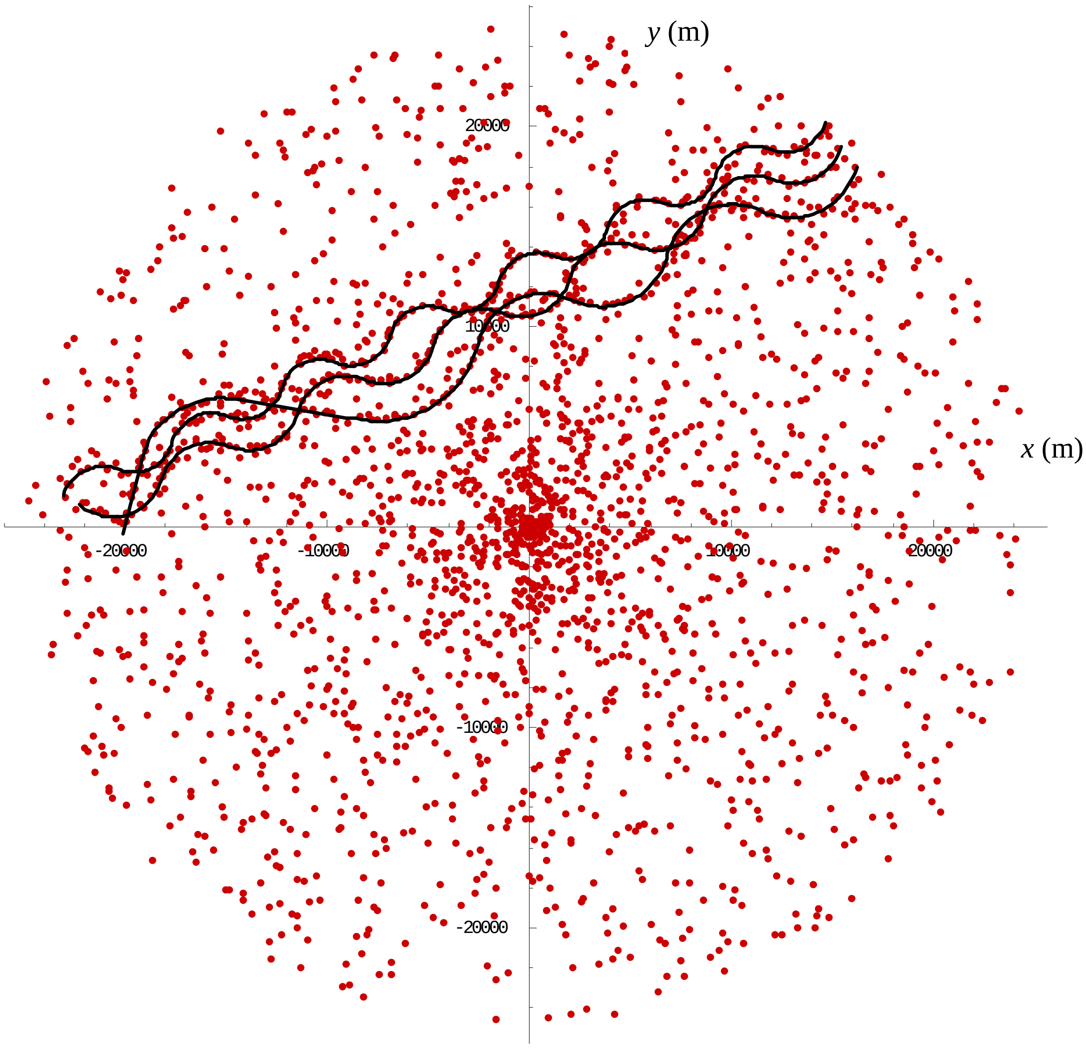

Figure 4. *Ground truth and radar contacts in Scenario Three*

Ground truth is shown in black, radar contacts in red. Numerous false alarms are distributed throughout a circular region of radius 25km, centred on the origin. (A few of the false alarms, in the vicinity of the axis labels, are not plotted.)

5.3 Threshold determination

After defining the three target trajectories, three hundred runs were performed. The purpose of the test runs was to explore the performance of a simple 2D tracking algorithm from Chapter 2 with respect to track purity. (Track purity is a standard tracking MOP, although in the literature it is often defined without the false alarms in the denominator [4].)

It was not hard to find many cases in which the simple algorithm confirmed only three entities, and several of those cases in which the weighted average track purity was greater than 0.98. (The mean over these runs, of the weighted average track purity, was 0.75, but the distribution of this MOP was bimodal, with a second peak approaching the perfect score of 1.) Essentially, the algorithm either got everything nearly right, or failed to do so; and it was deemed desirable to pick a case in which the performance was good. So one of these runs – having only three confirmed entities, and with a *WATP* greater than 0.98 – was chosen for this scenario.

The behaviour of the simulated radar in this scenario differs from that of Scenario One in three respects. First, it has better precision in range (see Section 5.1). Second, and more noticeably, it generates false alarms (that is, spurious contacts). These appear according to a Poisson distribution (in time) at an average rate of 7.2 contacts per second. The range of each false alarm is randomly generated according to a uniform distribution from zero to 25km. The bearing of each false alarm is calculated to correspond to its time tag. Finally, the contacts (both the false alarms and the "real" contacts corresponding to the targets) are grouped together into messages of up to four contacts each. The radar is treated as if it accumulates its contacts in a holding pool, and dumps the contents of the pool into a message whenever four contacts are accumulated, or whenever 0.15 seconds has elapsed since the time of the earliest contact in the pool, whichever comes first.

6. Scenario Four – Timeliness of entity formation

In order to pass the test of this scenario, a candidate system must not take too long to create and confirm an entity. An aircraft, which is first detected by our radar at a time 1.26 seconds after the start of the scenario, fires a missile at a time 100 seconds after the start of the scenario, the missile then being detected 0.58 seconds later.

The proposed specification statement read as follows: "The HMCCS, if operating in automatic mode and subjected to the input given in the file `s04input.csv`, interpreted according to Annex X.4, shall produce exactly two confirmed entities, which shall be confirmed within the time intervals specified in that annex."

The proposed text of the annex is reproduced here in Section 6.1. (An input file named `s04input.csv` was intended to be provided with the annex. Note that there was no ground truth file in this scenario, since the performance specifications make no reference to the ground truth.) Section 6.2 provides additional description of the scenario, while Section 6.3 discusses the choice of performance threshold.

6.1 Specification text

General Comments

This scenario tests the timeliness of entity formation and confirmation in the candidate system in response to 3D radar contacts, representing detections of two aerial targets, one of which is a missile fired by the other.

Data Requirements

The following data, for each confirmed entity, must be stored in order to assess the performance:
(1) The time at which the entity becomes confirmed.

Scenario Input

Each line of the Input File consists of five numbers, representing a radar contact. The five numbers are
(1) either a "1" or a "0", indicating whether this contact is the first contact in a message or a continuation of a previous message,
(2) contact time (in seconds, relative to the beginning of the scenario),
(3) contact range (in metres),
(4) contact bearing (in degrees, clockwise from north),
(5) contact elevation angle (in degrees).

The contacts shall be interpreted according to the following assumptions:
(1) The contacts are produced by a 3D tracking radar.
(2) The contacts are communicated to the HMCCS in messages that can consist of up to

four contacts each. Each contact that has a “1” in the first column (in the Input File) belongs to a message that arrives later than any message containing any contact that appears earlier in the file. Each contact that has a “0” in the first column belongs to the same message as the contact that immediately precedes it in the file.

(3) The radar’s uncertainties in range, bearing and elevation are given by nominal standard deviations of 40 metres, 1.0 degree and 1.0 degree, respectively.

(4) All targets are aerial.

(5) We may be dealing with multiple targets and clutter.

Performance Requirements

In order to meet the specification, the following results are required:

(1) Two, and only two, confirmed entities shall be produced in response to the data.

(2) One entity shall achieve confirmed status less than 10 seconds after the beginning of the scenario, and the other entity shall achieve confirmed status between 100 seconds and 110 seconds after the beginning of the scenario (that is, within 10 seconds after the missile is launched).

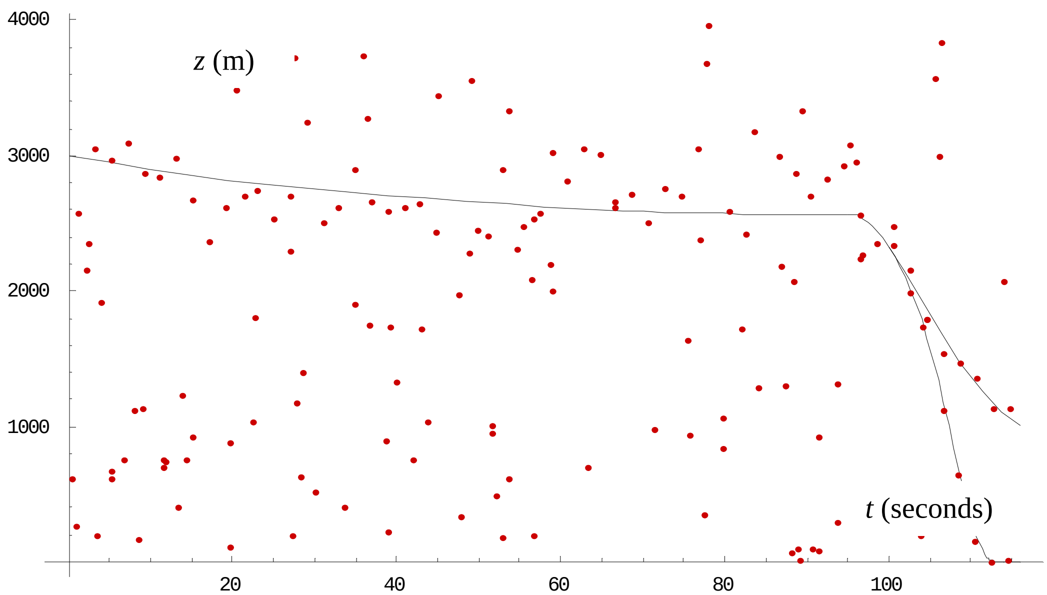

Figure 5. *Ground truth and radar contacts in Scenario Four – altitude (z) versus time (t)*

Ground truth is shown in black, radar contacts in red. The missile is fired at *t* = 100 seconds. The last 3 or 4 seconds of its flight are at a very low altitude, indistinguishable from the horizontal axis at this resolution. About half of the contacts are at altitudes higher than 4000m, and are excluded from the figure.

6.2 Additional description

Ground truth and radar contacts for the chosen run of Scenario Four (see Section 6.3) are shown in Figures 5 and 6. Figure 5 shows *z* as a function of time, while Figure 6 is an (*x*, *y*) plot of the same data.

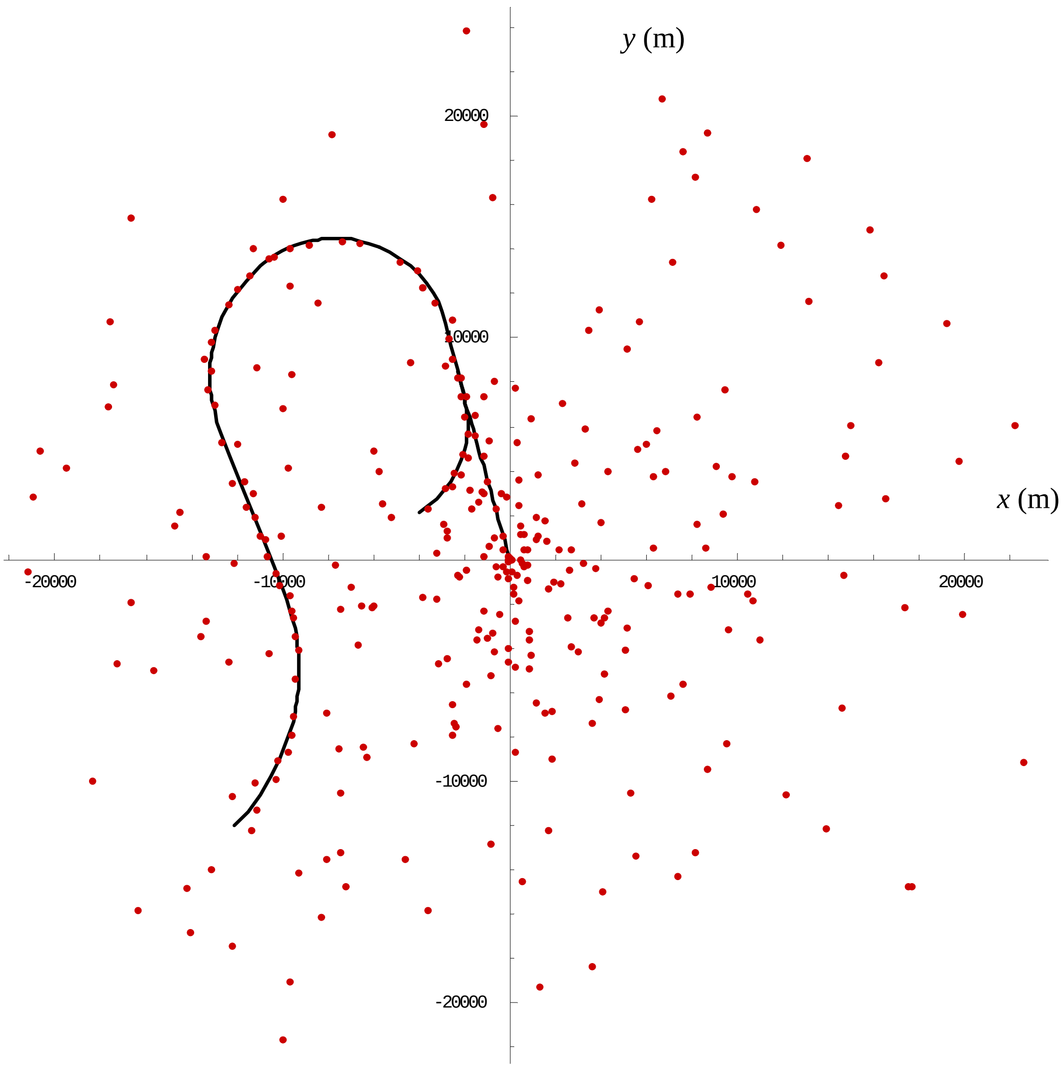

Figure 6. *Ground truth and radar contacts in Scenario Four – 2D positions*

Ground truth is shown in black, radar contacts in red. The aircraft starts at the lower left. The ground truth curve splits as the aircraft fires a missile at the ownship and turns away.

6.3 Threshold determination

The 3D radar of this scenario is much like the 3D radar of Scenario Two, except that its messages contain up to four contacts each (according to the same rules as in Scenario Three) and it produces false alarms. The false alarms are Poisson distributed in time, appearing at an average rate of about 2.6 contacts per second. The range is distributed uniformly from zero to 25km, and distributed from zero to 90 degrees in elevation angle with a probability density that is proportional to the cosine of that angle. The bearing of each false alarm is calculated to conform to its time tag.

The association and tracking methods of Chapter 2 were used. Initially, twenty test runs were performed (without the input data being recorded). Eighteen of the twenty runs achieved two confirmed entities, one corresponding to the aircraft and the other to the missile. The average elapsed times from scenario start to entity confirmation, among these eighteen runs, were 7.3 seconds for the aircraft and 108.3 seconds for the missile. Another run was performed (and its input data recorded in detail), and its results were typical of the test runs, with entity confirmations occurring 7.2 seconds and 106.6 seconds after the start of the scenario.

7. Scenario five – High-density tracking

For this scenario (only), we are dealing with surface targets instead of aerial targets. A large number of small boats are converging on the ownship in a coordinated attack, and the situation is further complicated by a large amount of clutter. This scenario tests the ability of the candidate system to form valid entities under these challenging conditions.

The proposed specification statement read as follows: "The HMCCS, if operating in automatic mode and subjected to the input given in the file `s05input.csv`, interpreted according to Annex X.5, shall produce at least thirty-eight confirmed non-spurious entities and at most six confirmed spurious entities within the first ten seconds of the scenario."

The proposed text of the annex is reproduced here in Section 7.1. (An input file and a ground truth file, respectively named `s05input.csv` and `s05gtr.csv`, were intended to be provided with the annex.) Section 7.2 provides additional description of the scenario, while Section 7.3 discusses the choice of performance threshold.

7.1 Specification text

General Comments

This scenario tests the ability of the candidate system to form entities in a high-density and high-clutter environment. A large number of small surface targets are converging on the ownship.

Data Requirements

The following data, for each confirmed entity, must be stored in order to assess the performance:
(1) Which contacts were associated with that entity. (Contacts may be identified by their time tags, or by their line numbers in the Input File.)

Scenario Input

Each line of the Input File consists of four numbers, representing a radar contact. The four numbers are
(1) either a "1" or a "0", indicating whether this contact is the first contact in a message or a continuation of a previous message,
(2) contact time (in seconds, relative to the beginning of the scenario),
(3) contact range (in metres),
(4) contact bearing (in degrees, clockwise from north).

The contacts shall be interpreted according to the following assumptions:
(1) The contacts are produced by a 2D tracking radar.
(2) The contacts are communicated to the HMCCS in messages that can consist of up to ten contacts each. Each contact that has a "1" in the first column (in the Input File) belongs to a

message that arrives later than any message containing any contact that appears earlier in the file. Each contact that has a “0” in the first column belongs to the same message as the contact that immediately precedes it in the file.

(3) The radar’s uncertainties in range and bearing are given by nominal standard deviations of 40 metres and 1.0 degree, respectively.

(4) All targets are surface vessels.

(5) We may be dealing with multiple targets and clutter.

Performance Requirements

In order to meet the specification, the following results are required:

(1) At least thirty-eight non-spurious entities (see below) shall achieve confirmed status within ten seconds after the beginning of the scenario.

(2) At most six spurious entities (see below) shall achieve confirmed status within ten seconds after the beginning of the scenario.

Definition

Each line of the Ground Truth File contains a pair of numbers, identifying the source target (or targets), if any, for the contact of the corresponding line in the Input File – that is, the target(s) from which that contact originated. Each number is an integer varying from 0 to 50. If both numbers in a line are “0”, the contact in question is a false alarm. If one number in a line is “0” but the other is not, then the non-zero number identifies the source target. If neither of the two numbers in a line is “0”, then the contact is unresolved, and both of the targets indicated by those numbers are sources for the contact.

For the purpose of this performance test, the definition of *spurious* is as follows:

Let **C** be the set of contacts that were associated with entity **E**, including the contact that caused the initiation of that entity.

(1) If there are no false alarms in **C**, then **E** is not spurious.

(2) If there exists a target **T** and a subset **M** of **C**, such that

(a) **M** includes more than half of the contacts in **C**,

and

(b) **T** is the source, or one of the sources, for every contact in **M**,

then **E** is not spurious.

(3) Otherwise, **E** is spurious.

7.2 Additional description

Ground truth and radar contacts for the chosen run of Scenario Five are shown in Figures 7 and 8. Figure 7 represents the entire scenario, while Figure 8 represents only the first ten seconds.

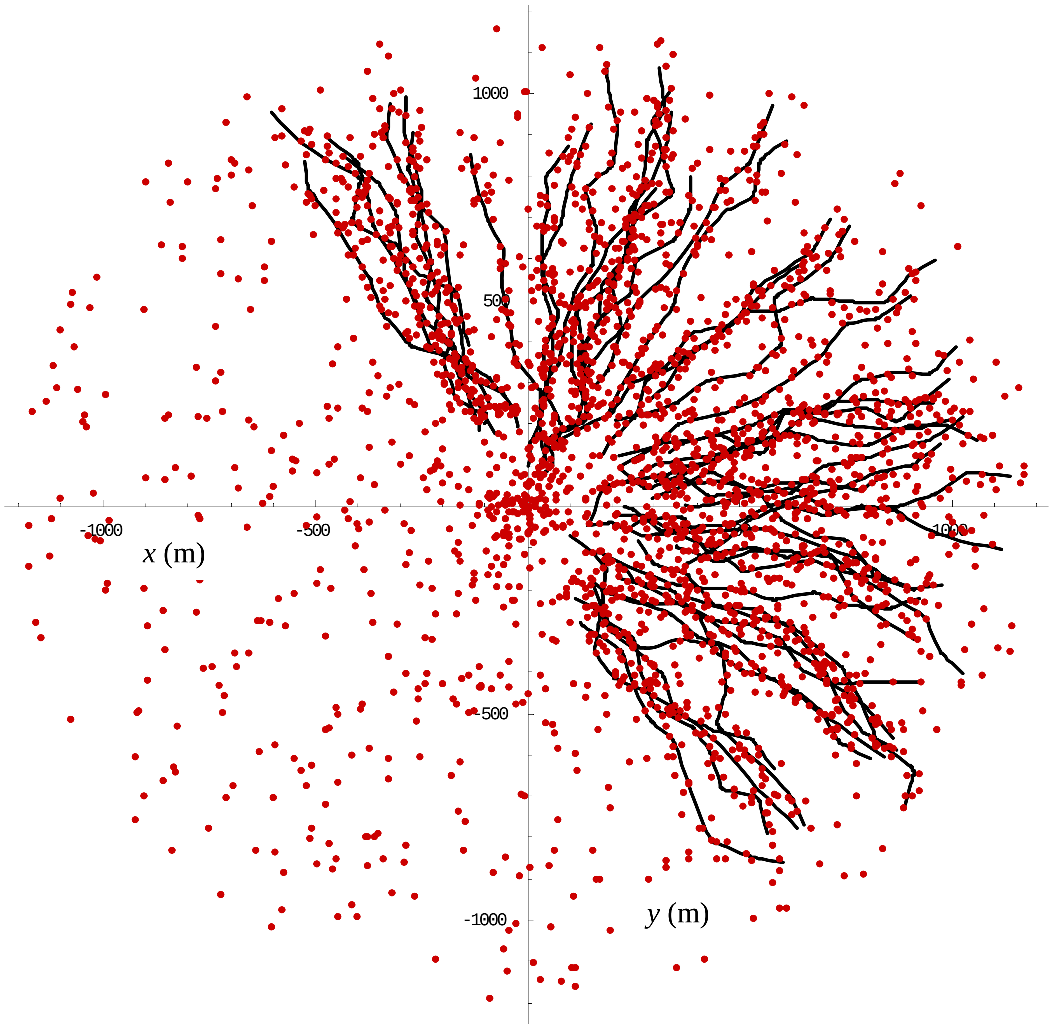

***Figure 7.** Ground truth and radar contacts in Scenario Five – full scenario*

Ground truth is shown in black, radar contacts in red.

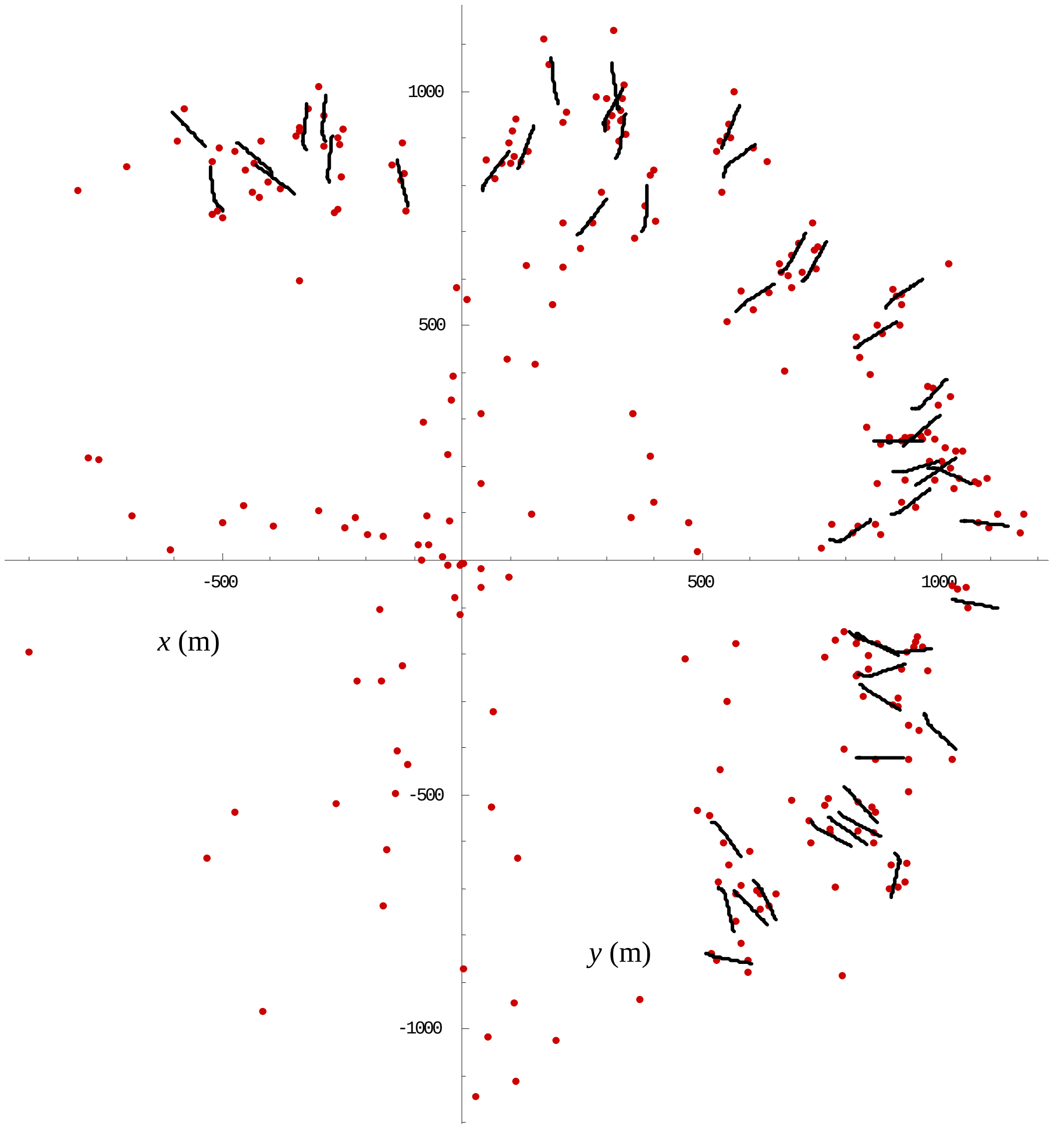

Figure 8. *Ground truth and radar contacts in Scenario Five – first ten seconds*

Ground truth is shown in black, radar contacts in red.

7.3 Threshold determination

This scenario is unusual in that the true target trajectories were not defined precisely from the start. Instead, many targets were generated randomly, within a set of rules that brought them gradually closer to the ownship (that is, to the origin of our coordinate system).

The random initial bearing of each target was based on a uniform distribution within a half-circle of possible bearing values. The random initial range was based on a Gaussian distribution, centred on the value 1000 m, with a standard deviation of about 100 m. The speed of each target was held constant at 10 m/s. The random initial course of each target was based on a Gaussian distribution, centred on a course directly toward the ownship, with a standard deviation of about 20 degrees. The target moved straight along this course for a random amount of time that varied from 6 to 10 seconds, with a distribution that was sinusoidal in shape, with a single peak at 8 seconds. At the end of this segment, the target chose a new random course, again centred on a course directly toward the ownship, with a standard deviation of about 20 degrees. It turned toward its new course with a turning radius of 20 m, then it moved straight along its new course for a random amount of time, based on the same distribution that governed the amount of time spent on its initial course. This cycle (of choosing a new course, turning toward it, and moving straight along it for a few seconds) was repeated until the target's range was less than 100 m.

Targets were generated in sets of fifty, until a set was found for which the population standard deviation of the range of the targets, at the time of expiry of the first target's trajectory definition, was less than 90 m. This set was selected for the scenario. (The scenario was intended to be of a coordinated attack, and more variation in this final range would correspond to less "coordination".)

The radar was the same as in Scenario One, except that the detection probability was 0.9, false alarms were generated, and different rules were followed for grouping contacts into messages. The false alarms appeared at an average rate of nearly 11 contacts per second, and their random ranges were based on a uniform distribution from zero to 1200 m. Each radar message could contain up to ten contacts, with each contact being held for a maximum of 0.2 seconds before being released into a message.

The association and tracking algorithms were the same as in Chapter 2, except for two of the parameter settings. The process noise covariance was smaller by a factor of 150 (suitable for surface targets), and the initial uncertainty in velocity for a newly-created entity was set to 10 m/s in each horizontal dimension.

With this simple algorithm, 44 entities were confirmed within the first ten seconds of the scenario, and five of those were spurious. The performance threshold was set to be slightly worse than these results in order to allow achievability.

8. Scenario six – Recovery from association error

In this scenario, two aerial targets merge and split, while being tracked by a radar (with a built-in tracking system). The radar makes an association error by failing to recognize that the target trajectories actually cross each other. There is also an “Identification: Friend or Foe” (IFF) system in use, and it is assumed that the IFF contacts contain sufficient information to distinguish between the two targets. In order to pass the test of this scenario, the candidate system must correct the association error, using the IFF contacts.

The proposed specification statement read as follows: “The HMCCS, if operating in automatic mode and subjected to the input given in the file `s06input.csv`, interpreted according to Annex X.6, shall produce exactly two confirmed entities, whose identities shall be determined correctly, as defined in that annex, within the time interval specified therein.”

The proposed text of the annex is reproduced here in Section 8.1. (An input file named `s06input.csv` was intended to be provided with the annex.) Section 8.2 provides additional description of the scenario, while Section 8.3 discusses the choice of performance threshold.

8.1 Specification text

General Comments

In this scenario, two aerial targets are being tracked by a 2D radar, but the radar tracks contain an association error. The scenario tests the ability of the candidate system to use IFF contacts to recover from this track-swapping error.

Data Requirements

The following data, from each update of an entity in response to a radar track update or to an IFF contact, must be stored in order to assess the performance:
(1) The track number of the radar track that is associated with that entity.
(2) Whether or not the system has sufficient confidence about the entity’s identity in order to display that identity; and if so, the identity of that entity. (Identity is represented by a “target identifier code” – see below.)

Scenario Input

Each line of the Input File consists of eighteen numbers. If the first number is a “1”, the line represents a radar track update. If the first number is a “2”, the line represents the return from an IFF transponder.

In case of a radar track update, the eighteen numbers are
(1) a “1”, as stated above,
(2) either a “1” or a “0”, indicating whether this track update is the first track update in a message or a continuation of a previous message,
(3) track update time (in seconds, relative to the beginning of the scenario),

(4) radar track number (either “1” or “2”),
(5) *x* position (in metres, east of the radar position),
(6) *y* position (in metres, north of the radar position),
(7) *x*-component of velocity (in m/s),
(8) *y*-component of velocity (in m/s),
(9) variance in *x* position (in m^2),
(10) covariance between *x* position and *y* position (in m^2),
(11) covariance between *x* position and *x*-component of velocity (in m^2/s),
(12) covariance between *x* position and *y*-component of velocity (in m^2/s),
(13) variance in *y* position (in m^2),
(14) covariance between *y* position and *x*-component of velocity (in m^2/s),
(15) covariance between *y* position and *y*-component of velocity (in m^2/s),
(16) variance in *x*-component of velocity (in m^2/s^2),
(17) covariance between *x*- and *y*-components of velocity (in m^2/s^2),
(18) variance in *y*-component of velocity (in m^2/s^2).

In the case of an IFF report, the eighteen numbers are
(1) a “2”, as stated above,
(2) either a “1” or a “0”, indicating whether this contact is the first contact in a message or a continuation of a previous message,
(3) contact time (in seconds, relative to the beginning of the scenario),
(4) target identifier code (either “4” or “6”),
(5) contact range (in metres),
(6) contact bearing (in degrees, clockwise from north),
(7-18) twelve placeholders (all “0”), which can be ignored.

The contacts and track updates shall be interpreted according to the following assumptions:
(1) The track updates are produced by a 2D tracking radar. All radar tracks are considered to be already confirmed.
(2) The IFF reports are produced by an IFF system attached to the same radar.
(3) The IFF reports and track updates are communicated to the HMCCS in messages that can each consist either of up to two IFF reports or up to two radar track updates. Each IFF report or track update with a “1” in the second column (in the Input File) belongs to a message that arrives later than any message containing any data that appear earlier in the file. Each IFF report or track update with a “0” in the second column belongs to the same message as the IFF report or track update that immediately precedes it in the file.
(4) The IFF reports’ uncertainties in range and bearing are given by nominal standard deviations of 100 metres and 2.0 degrees, respectively.
(5) We are dealing with two aerial targets. The target identifier code from an IFF report can be considered as a reliable indication of the target from which the report originated. The two numbers (“4” and “6”) were chosen arbitrarily to represent the two targets.

Performance Requirements

In order to meet the specification, the following results are required:
(1) Two, and only two, confirmed entities shall be produced in response to the data.
(2) On the basis of the first two IFF reports included in the input file, target identifier code “4” shall be assigned to the entity that is associated with radar track “2”, and target identifier

code “6” shall be assigned to the entity that is associated with radar track “1”.
(3) Eventually, the system shall conclude that target identifier code “4” now belongs to the entity that is associated with radar track “1”, while target identifier code “6” now belongs to the entity that is associated with radar track “2”. This conclusion will be reached in less than 45 seconds after the beginning of the scenario, with enough confidence for those identity assignments to be displayed, and will be maintained firmly thereafter until the end of the scenario.

8.2 Additional description

Figure 9 shows ground truth, radar tracks, and IFF contacts for Scenario 6. For the sake of this figure, the radar tracks and IFF contacts are interpreted as if the elevation angle is always zero.

The targets begin to diverge at about 39.5 seconds into the scenario.

8.3 Threshold determination

The simple fusion algorithms from Chapter 2, that were used in earlier scenarios in order to set the performance threshold, provided a convenient means for making simulated radar tracks in the present scenario.

Several test runs were performed, in which contacts were formed by the same simulated radar that was used in Scenario One, and radar tracks were formed by the association and tracking algorithm from Chapter 2. It quickly became clear that exactly two confirmed tracks were made in a large majority of the runs. Test runs were repeated until one was found that suffered from the association error described in Section 8.1.

In the selected run, track update reports from confirmed tracks (only) were sorted by time and grouped into messages, according to the rules that were used in Scenario Three for the grouping of radar contacts. Each radar track update message had either one or two updates.

IFF contacts were then generated in the same manner as radar contacts, but using a perfect probability of detection (for convenience), and larger errors (as described in Section 8.1). Target identity data, represented by the two arbitrary labels “4” and “6”, were included in the IFF contacts. IFF contacts were grouped into messages, in the same manner as the radar contacts of Scenario Three. Each IFF message had two contacts, one from each target. As a final step in the creation of the input file, the IFF

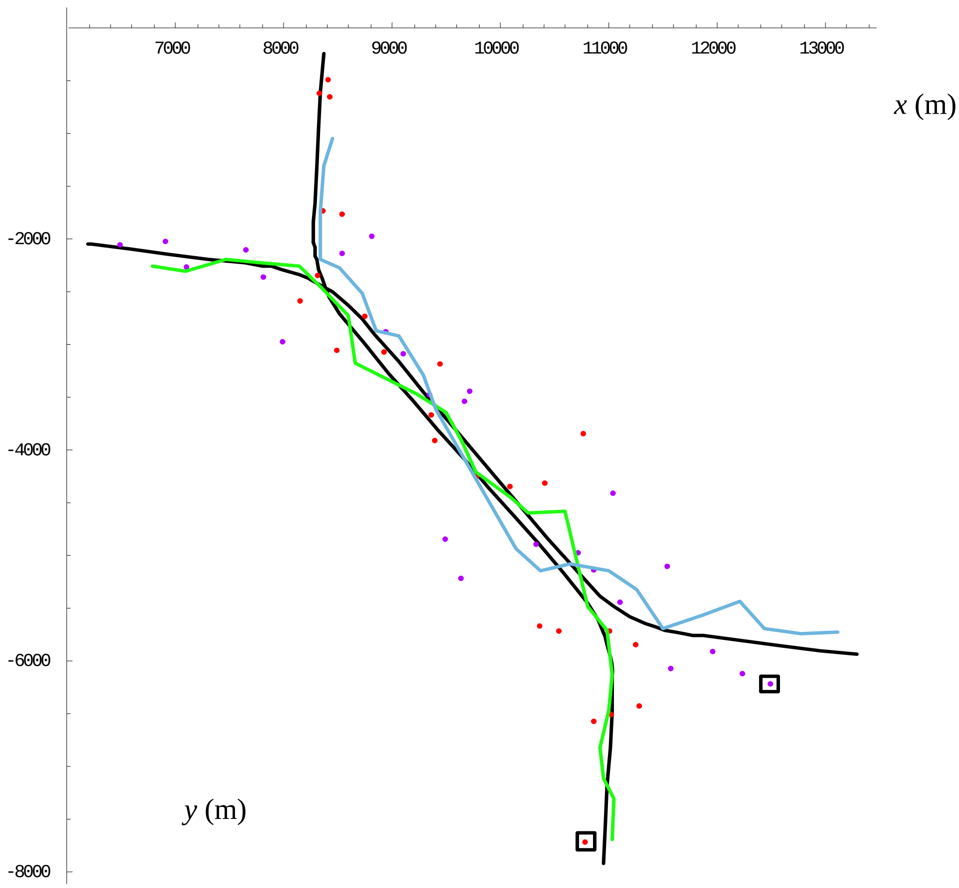

Figure 9. *Ground truth, radar tracks, and IFF contacts in Scenario Six*

Ground truth is shown in black, radar tracks in blue and green, and IFF contacts in purple and red (one colour for each target). The targets start in the lower right. The black squares indicate the first two IFF contacts.

contact messages and the radar track update messages were combined and sorted by time.

In this scenario, no fusion algorithms were implemented to set an achievable threshold. (An algorithm to handle exactly this kind of situation was subsequently developed for DRDC Atlantic [5].) Instead, the threshold was derived from theoretical considerations, as follows.

Based on the traditional Gaussian interpretation of the radar track covariance matrices and the nominal IFF contact uncertainties (transformed to Cartesian coordinates in the usual way), the relative probability of the two contact-to-track association combinations were calculated for each pair of IFF contacts without reference to identity data. In particular, the probability of the association combination that contradicts the natural initial identity assignment (based on the first few IFF messages) was compared to the probability of the association combination that confirms that initial identity assignment. In other words, the hypothesis that the track IDs are correct is compared to the hypothesis that the track IDs are incorrect.

For the pair of IFF contacts that appeared just before 42 seconds had passed from the beginning of the scenario, the probability of the association combination that contradicts the initial identity assignment was about 50,000 times higher than the other combination. For the pair of IFF contacts that appeared just before 44 seconds had passed from the beginning of the scenario, that probability ratio is greater than 10^{16} – surely high enough to remove any doubt that an error had occurred. The specification time limit of 45 seconds is based on these observations.

9. Scenario seven – Association performance with heterogeneous sensors

This scenario is similar to Scenario Three, testing the association performance of the candidate system in response to three aerial targets. But this time there is an ESM sensor (giving bearing-only position measurements along with ID-related data) as well as a 2D radar.

The proposed specification statement read as follows: "The HMCCS, if operating in automatic mode and subjected to the input given in the file `s07input.csv`, interpreted according to Annex X.7, shall produce exactly three confirmed entities, resulting in a weighted average track purity that is equal to or better than the threshold specified in that annex."

The proposed text of the annex is reproduced here in Section 9.1. (An input file and a ground truth file, respectively named `s07input.csv` and `s07gtr.csv`, were intended to be provided with the annex.) Section 9.2 provides additional description of the scenario, while Section 9.3 discusses the choice of performance threshold.

9.1 Specification text

General Comments

This scenario tests the association performance of the candidate system in response to a combination of 2D radar contacts and ESM contacts (i.e., bearing-only contacts with emitter data), representing detections of three aerial targets.

Data Requirements

The following data, for each confirmed entity, must be stored in order to assess the performance:
(1) Which contacts were associated with that entity. (Contacts may be identified by their time tags, or by their line numbers in the Input File.)

Scenario Input

Each line of the Input File consists of five numbers. If the first number is a "2", the line represents an ESM contact. Otherwise, the line represents a radar contact.

In the case of a radar contact, the five numbers are
(1) either a "1" or a "0", indicating whether this contact is the first contact in a message or a continuation of a previous message,
(2) contact time (in seconds, relative to the beginning of the scenario),
(3) a "0", which can be ignored,
(4) contact range (in metres),
(5) contact bearing (in degrees, clockwise from north).

In the case of an ESM contact, the five numbers are
(1) a "2", as stated above,
(2) contact time (in seconds, relative to the beginning of the scenario),
(3) emitter number ("3", "5", or "9"), identifying the target from which the contact originated,
(4) a "0", which can be ignored,
(5) contact bearing (in degrees, clockwise from north).

The contacts shall be interpreted according to the following assumptions:
(1) The radar contacts are produced by a 2D tracking radar. The radar's uncertainties in range and bearing are given by nominal standard deviations of 40 metres and 0.5 degrees, respectively.
(2) The ESM system (which produces the ESM contacts) is located in practically the same position as the radar. Its uncertainty in bearing is given by a nominal standard deviation of 3.0 degrees.
(3) The contacts are communicated to the HMCCS in messages that can each consist either of a single ESM contact, or up to four radar contacts. Each contact that has a "1" or a "2" in the first column (in the Input File) belongs to a message that arrives later than any message containing any contact that appears earlier in the file. Each contact that has a "0" in the first column belongs to the same message as the contact that immediately precedes it in the file.
(4) All targets are aerial.
(5) We may be dealing with multiple targets. Some of the radar contacts may represent clutter.

Performance Requirements

In order to meet the specification, the following results are required:
(1) Three, and only three, confirmed entities shall be produced in response to the data.
(2) The weighted average track purity (see below) shall be greater than or equal to 0.98.

Definition

Each line of the Ground Truth File contains one number, identifying the target (if any) from which the contact of the corresponding line in the Input File originated. Each number is an integer varying from 0 to 3. A "0" indicates that the contact is a false alarm. Otherwise, the number identifies the target.

Let C_{ij} denote the number of contacts from target j (j varying from 1 to 3) that were associated with entity i (i varying from 1 to 3, representing the three confirmed entities). Let F_i denote the number of false alarms that were associated with entity i. Then the weighted average track purity is

$$WATP = \frac{\sum_i \max_j C_{ij}}{\sum_i \left(F_i + \sum_j C_{ij} \right)}.$$

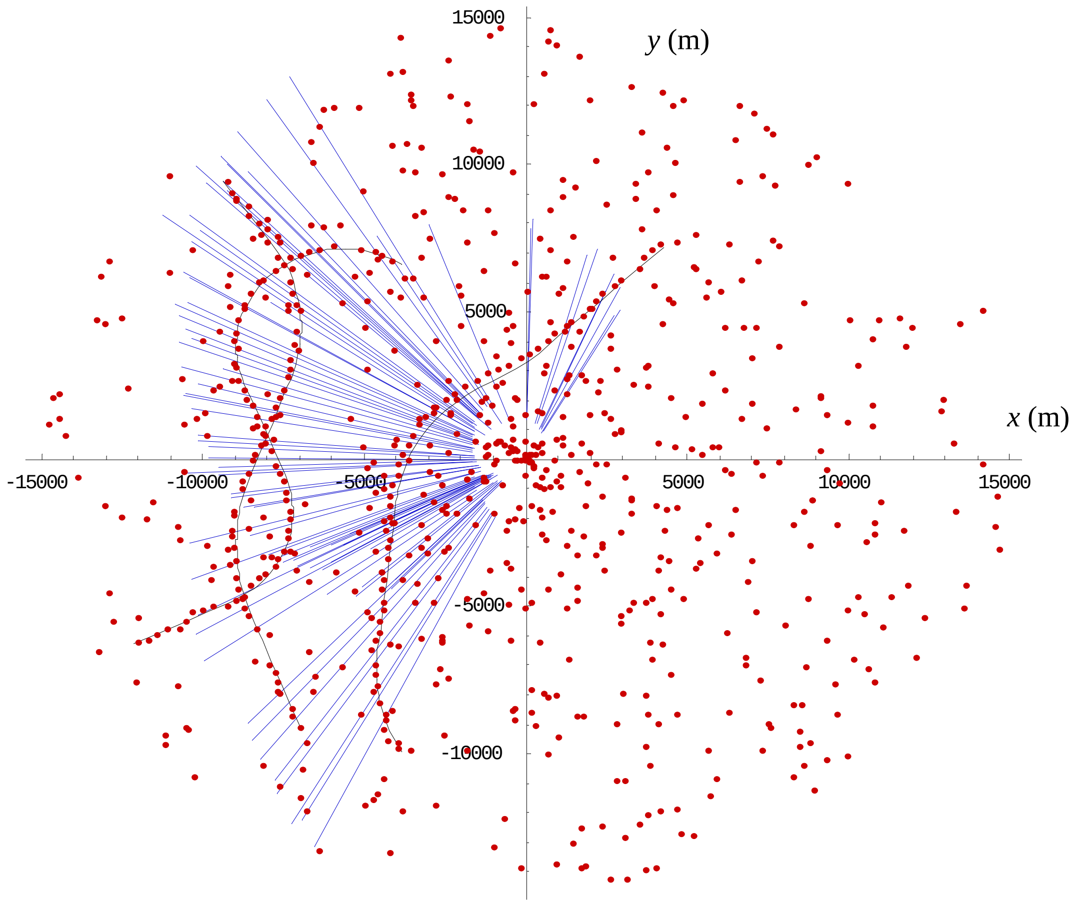

Figure 10. *Ground truth, radar contacts and ESM contacts in Scenario Seven*

Ground truth is shown in black, radar contacts in red, and ESM contacts (bearing lines) in blue. About one third of the radar contacts have a range greater than 15km, and are excluded from this figure. The sequence of ESM bearing lines is indicated by the length of the lines: Those that occur earlier in the scenario are drawn shorter, while those that occur later are drawn longer. The two interweaving targets both start in the lower left, while the nearer target starts in the upper right.

9.2 Additional description

Ground truth, radar contacts and ESM contacts for Scenario Seven are shown in Figure 10. For the sake of this figure, radar contacts are interpreted as if the elevation angle is always zero. (The error induced by this assumption is noticeable in the figure.) ESM contacts are shown as bearing lines.

9.3 Threshold determination

The simulated radar in this scenario was exactly the same as the one used in Scenario Three, including the way in which its contacts were grouped into messages. ESM contacts were generated in a similar way, but with a rotation period of four seconds, no false alarms, range data excluded, and a larger spread in bearing values, as indicated in Section 9.1. Target identity data was included in the ESM contacts, in the form of the arbitrary labels “3”, “5”, and “9”.

The association and tracking methods of Chapter 2 (including the methods of Section 2.3 with respect to the bearing-only sensor data) were used in this scenario, with a small modification to account for the identity data. Whenever considering the association of an ESM contact with an entity that had previously been associated at least once with another ESM contact, the association weight was either decreased or increased (additively) by “2” (an amount chosen for convenience), according to whether the identity label agreed or disagreed with that of the most recent ESM contact to be associated with the entity. This modification to the association weight corresponds to a (multiplicative) increase or decrease in the nominal association probability density by a factor of e^2 (about 7.4).

Ten test runs were generated before selecting (somewhat arbitrarily) a data set to be used for the specification. There were exactly three confirmed entities in four of the ten runs. Among these four runs, the weighted average track purity varied from 0.987 to a perfect score of 1. For the chosen input data, the weighted average track purity achieved by the algorithm above was 0.988.

10. Scenario eight – 3D tracking accuracy with heterogeneous sensors

This scenario, like Scenario Two, tests tracking accuracy in 3D. But this time we have three different sensors: A 3D radar, a 2D radar, and a bearing-only sensor of some unspecified kind.

The proposed specification statement read as follows: “The HMCCS, if operating in automatic mode and subjected to the input given in the file `s08input.csv`, interpreted according to Annex X.8, shall produce exactly one confirmed entity, whose mean absolute error in position, mean absolute error in velocity, and RMS relative error shall be equal to or better than the thresholds specified in that annex.”

The proposed text of the annex is reproduced here in Section 10.1. (An input file and a ground truth file, respectively named `s08input.csv` and `s08gtr.csv`, were intended to be provided with the annex.) Section 10.2 provides additional description of the scenario, while Section 10.3 discusses the choice of performance threshold.

10.1 Specification text

General Comments

This scenario tests the ability of the system to fuse contacts from three different sensors without a proliferation of spurious tracks.

Data Requirements

The following data, for each confirmed entity, must be stored in order to assess the performance:
(1) The time of confirmation of the entity.
(2) Whether the entity was deleted or not, before the time of the final contact of the scenario.
(3) Which contacts were associated with that entity. (Contacts may be identified by their time tags, or by their line numbers in the Input File.)
(4) The estimated position and velocity, in three dimensions, at each update.
(5) The estimated covariance matrix corresponding to the estimated position and velocity, or an alternative expression of uncertainty from which such a covariance matrix may be derived, at each update.

Scenario Input

Each line of the Input File consists of six numbers, representing a sensor contact of some kind. The six numbers are
(1) An integer from 1 to 3, identifying the sensor,
(2) Either a “1” or a “0”, indicating whether this contact is the first contact in a message or a

continuation of a previous message,
(3) contact time (in seconds, relative to the beginning of the scenario),
(4) contact range (in metres), if that information is available,
(5) contact bearing (in degrees, clockwise from north),
(6) contact elevation angle (in degrees) if that information is available.

The contacts shall be interpreted according to the following assumptions.
(1) If the first number in the line is "1", the contact is considered to be produced by a 3D tracking radar. The uncertainties in range, bearing, and elevation angle are given by nominal standard deviations of 40 metres, 0.5 degrees, and 0.8 degrees respectively.
(2) If the first number in the line is "2", the contact is considered to be produced by a 2D tracking radar. The sixth number in the line will be a "0", indicating that the elevation angle has not been measured. The uncertainties in range and bearing are given by nominal standard deviations of 60 metres and 1.0 degree, respectively.
(3) If the first number in the line is "3", the contact is a bearing-only contact. The fourth and sixth numbers in the line will each be a "0", indicating that the range and elevation angle have not been measured. The uncertainty in bearing is given by a nominal standard deviation of 3.0 degrees.
(4) All targets are aerial.
(5) We may be dealing with multiple targets and clutter. (In fact, there is only one target, but this background information may not be used by the system.)

Performance Requirements

In order to meet the specification, the following results are required:
(1) One, and only one, confirmed entity shall be produced in response to the data. This entity shall be confirmed within ten seconds of the beginning of the scenario, and shall not be deleted before the final contacts of the scenario are processed.
(2) The mean absolute error in position (see below) shall be less than or equal to 190 m.
(3) The mean absolute error in velocity (see below) shall be less than or equal to 72 m/s.
(4) The RMS relative error in position and velocity (see below) shall be less than or equal to 2.0.

Definitions

The "absolute error in position", at each entity update, is the scalar magnitude of the position error vector:

$$E_{\mathrm{pos}} = \left\| \begin{pmatrix} \hat{x} - x \\ \hat{y} - y \\ \hat{z} - z \end{pmatrix} \right\|,$$

where x, y, and z are the true coordinates of the target at that time (as given in the Ground Truth File below), and $\hat{x}$, $\hat{y}$, and $\hat{z}$ are the estimates of these coordinates according to the entity. The "mean absolute error in position" is the mean of E_{pos} for all entity updates (of the confirmed entity), where each update is given equal weight.

Similarly, the "absolute error in velocity", at each entity update, is the scalar magnitude of the velocity error vector:

$$E_{\text{vel}} = \left\| \begin{pmatrix} \hat{v}_x - v_x \\ \hat{v}_y - v_y \\ \hat{v}_z - v_z \end{pmatrix} \right\|,$$

where v_x and v_y are the true velocity components of the tracked target at that time, etc. The "mean absolute error in velocity" is the mean of E_{vel} for all entity updates (of the confirmed entity), where each update is given equal weight.

The "squared relative error in position and velocity", at each entity update, is the quantity

$$D = \begin{pmatrix} \hat{x} - x & \hat{y} - y & \hat{z} - z & \hat{v}_x - v_x & \hat{v}_y - v_y & \hat{v}_z - v_z \end{pmatrix} \mathbf{P}^{-1} \begin{pmatrix} \hat{x} - x \\ \hat{y} - y \\ \hat{z} - z \\ \hat{v}_x - v_x \\ \hat{v}_y - v_y \\ \hat{v}_z - v_z \end{pmatrix},$$

where **P** is the 6-by-6 covariance matrix representing the uncertainty claimed by the entity, with respect to its estimates of x, y, z, v_x, v_y, and v_z. The "RMS relative error in position and velocity" is the square root of the mean of D for all entity updates (of the confirmed entity), where each update is given equal weight.

Each line of the Ground Truth File gives the "ground truth" corresponding to the time of the matching line number of the Input File. The six numbers in each line are

(1) *x* position (in metres, east of the radar position),
(2) *y* position (in metres, north of the radar position),
(3) *z* position (in metres, above the horizontal plane of the radar position),
(4) *x*-component of velocity (in m/s),
(5) *y*-component of velocity (in m/s),
(6) *z*-component of velocity (in m/s).

10.2 Additional description

Figure 11 presents the ground truth and the sensor contacts (including bearing-only contacts, shown as bearing lines) for Scenario Eight. For the sake of this figure, 2D radar contacts are interpreted as if the elevation angle is always zero.

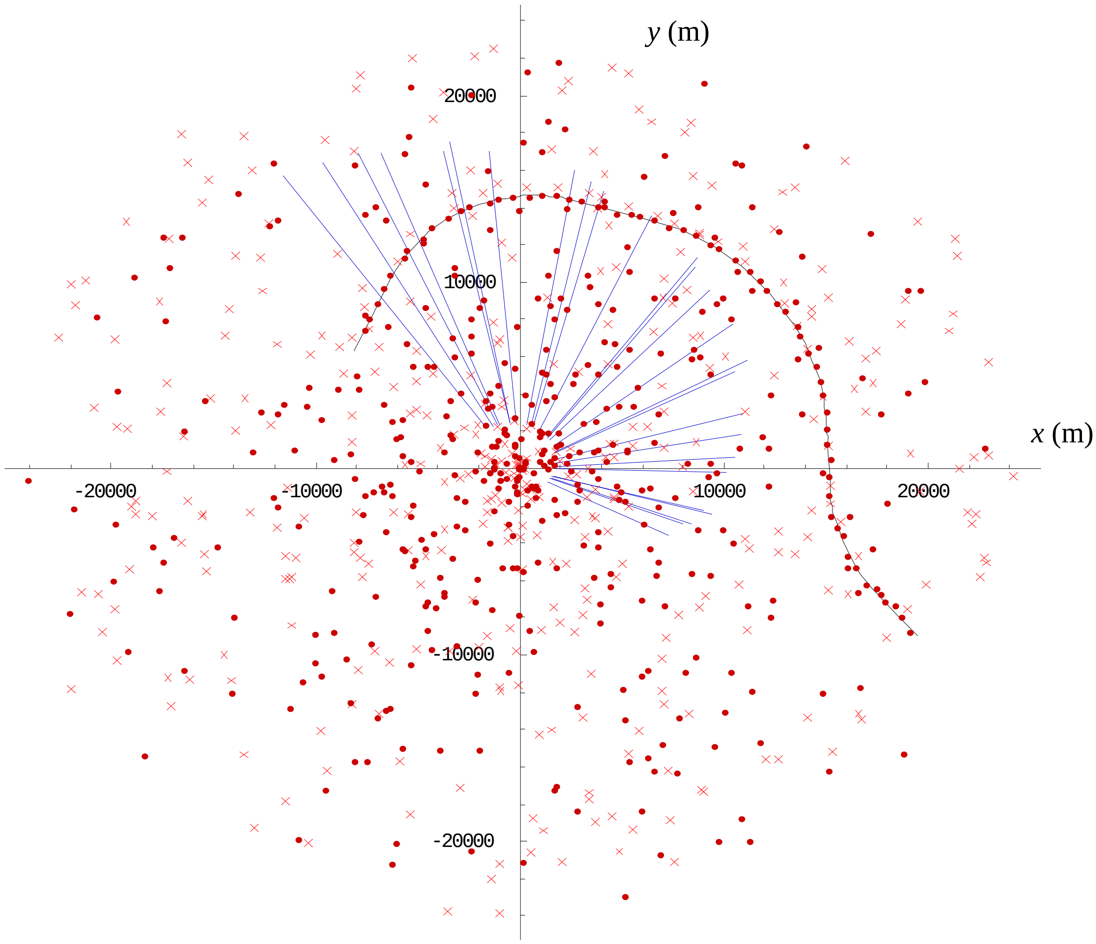

Figure 11. *Ground truth and sensor contacts in Scenario Eight*

Ground truth is shown in black, radar contacts in red, and bearing lines in blue. The red dots indicate 3D radar contacts, while the red crosses indicate 2D radar contacts. The sequence of bearing lines is indicated by the length of the lines: Those that occur earlier in the scenario are drawn shorter, while those that occur later are drawn longer. The target starts in the lower right.

10.3 Threshold determination

Three sensors were used in this scenario. The sensors were modelled in essentially the same way as in previous scenarios, but with slightly different parameters. The Gaussian distributions that were used to generate the random errors in the measured dimensions were as indicated in Section 10.1. The rotational periods of the three sensors were 2 seconds for the 3D radar, 5 seconds for the 2D radar, and 4 seconds for the bearing-only sensor. The probability of detection of the target was 0.95 for each pass of each sensor. The 2D radar generated false alarms in the same way as the radar of Scenario Three (but with a lower average frequency, 3 contacts per second), while the 3D radar generated false alarms in the same way as the radar of Scenario Four (but with a slightly higher average frequency, about 3.2 contacts per second). No false alarms were produced by the bearing-only sensor. Messages from each sensor can contain up to six contacts each, with a maximum elapsed time of 0.15 seconds from the time of the first contact in a message to the time at which the message is sent.

Thirty test runs were performed. The simple algorithm of Chapter 2 (including the methods of Section 2.3 for contacts with missing dimensions) frequently failed (26 times out of 30) to maintain perfect track continuity. One of the four successful runs was chosen arbitrarily for the specification statement, and the threshold values of the tracking accuracy MOPs were set to be slightly worse than those achieved in this run by the simple algorithm.

11. Scenario nine – Building a 3D picture with no 3D sensors

Again, this scenario tests 3D tracking accuracy. But in this case there are no 3D sensors. Instead, we have an IRST, which gives bearing and elevation angle data, plus a 2D radar. The candidate system must build a 3D picture by the association of two kinds of 2D data.

Recognizing the fact that radar systems generally come with a built-in tracking system, the radar data of this scenario consist of track updates rather than contacts[4], while the IRST data consist of contacts.

The proposed specification statement read as follows: “The HMCCS, if operating in automatic mode and subjected to the input given in the file `s09input.csv`, interpreted according to Annex X.9, shall produce exactly three confirmed entities, none of which is associated with any false alarms, and whose collective mean absolute error in position, mean absolute error in velocity, and RMS relative error shall be equal to or better than the thresholds specified in that annex.”

The proposed text of the annex is reproduced here in Section 11.1. (An input file and a ground truth file, respectively named `s09input.csv` and `s09gtr.csv`, were intended to be provided with the annex.) Section 11.2 provides additional description of the scenario, while Section 11.3 discusses the choice of performance threshold.

11.1 Specification text

General Comments

This scenario tests the ability of the system to build up a three-dimensional picture in the absence of 3D sensors. Three aerial targets are being tracked by 2D radar. The resulting tracks are to be combined with IRST contacts.

Data Requirements

The following data, for each confirmed entity, must be stored in order to assess the performance:
(1) Which contacts were associated with that entity. (Contacts may be identified by their time tags, or by their line numbers in the Input File.)
(2) The estimated position and velocity, in three dimensions, at each update.
(3) The estimated covariance matrix corresponding to the estimated position and velocity, or an alternative expression of uncertainty from which such a covariance matrix may be derived, at each update.

---

[4] In retrospect, perhaps it would have been better to use contacts for both sensors. As indicated in section 11.3, the use of radar track data led to a rather convoluted *ad hoc* algorithm being used in place of the simpler algorithm of chapter 2 to set the performance thresholds.

Scenario Input

Each line of the Input File consists of eighteen numbers. If the first number is a “1”, the line represents a radar track update. If the first number is a “2”, the line represents an IRST contact.

In case of a radar track update[5], the eighteen numbers are

(1) a “1”, as stated above,

(2) either a “1” or a “0”, indicating whether this track update is the first track update in a message or a continuation of a previous message,

(3) track update time (in seconds, relative to the beginning of the scenario),

(4) radar track number (“1” or “2” or “3”),

(5) *x* position (in metres, east of the radar position),

(6) *y* position (in metres, north of the radar position),

(7) *x*-component of velocity (in m/s),

(8) *y*-component of velocity (in m/s),

(9) variance in *x* position (in m^2),

(10) covariance between *x* position and *y* position (in m^2),

(11) covariance between *x* position and *x*-component of velocity (in m^2/s),

(12) covariance between *x* position and *y*-component of velocity (in m^2/s),

(13) variance in *y* position (in m^2),

(14) covariance between *y* position and *x*-component of velocity (in m^2/s),

(15) covariance between *y* position and *y*-component of velocity (in m^2/s),

(16) variance in *x*-component of velocity (in m^2/s^2),

(17) covariance between *x*- and *y*-components of velocity (in m^2/s^2),

(18) variance in *y*-component of velocity (in m^2/s^2).

In the case of an IRST contact, the eighteen numbers are

(1) a “2”, as stated above,

(2) either a “1” or a “0”, indicating whether this contact is the first contact in a message or a continuation of a previous message,

(3) contact time (in seconds, relative to the beginning of the scenario),

(4) a “0”, which can be ignored,

(5) contact bearing (in degrees, clockwise from north),

(6) contact elevation angle (in degrees),

(7-18) twelve placeholders (all “0”), which can be ignored.

The contacts and track updates shall be interpreted according to the following assumptions:

(1) The track updates are produced by a 2D tracking radar. All radar tracks are considered to be already confirmed. These tracks have been constructed from range-and-bearing contacts according to the assumption that every target is flying at an altitude of 1000 metres, or at an elevation angle of 60 degrees, whichever is lower.

(2) The IRST system responsible for the IRST contacts is located in practically the same position as the radar.

---

[5] The radar track data are given in Cartesian coordinates because tracking is more conveniently done in Cartesian coordinates. However, a case could be made for using range and bearing, since radar data (even track data) are usually presented this way in practice.

(3) The IRST contacts and radar track updates are communicated to the HMCCS in messages that can each consist either of up to two IRST contacts or up to two radar track updates. Each contact or track update with a "1" in the second column (in the Input File) belongs to a message that arrives later than any message containing any data that appear earlier in the file. Each contact or track update with a "0" in the second column belongs to the same message as the contact or track update that immediately precedes it in the file.
(4) The IRST contacts' uncertainties in bearing and elevation angle are given by nominal standard deviations of 0.5 degrees in each angle.
(5) All targets are aerial.
(6) We may be dealing with multiple aerial targets and clutter. It can be assumed that the radar false alarms have already been processed by the radar tracker. However, some of the IRST contacts may be false alarms.

Performance requirements

In order to meet the specification, the following results are required:
(1) Three confirmed entities (only) shall be produced in response to the data.
(2) No IRST false alarms shall be associated with any confirmed entities, not even before they achieve confirmed status.
(3) The mean absolute error in position (see below) shall be less than or equal to 450 m.
(4) The mean absolute error in velocity (see below) shall be less than or equal to 80 m/s.
(5) The RMS relative error in position and velocity (see below) shall be less than or equal to 8.0.

Definitions

At each update of an entity in response to a contact or a track update, the target from which that contact or track update originated shall be termed "the tracked target" of that entity at that update.

The "absolute error in position", at each entity update, is the scalar magnitude of the position error vector:

$$E_{\mathrm{pos}} = \left\| \begin{pmatrix} \hat{x} - x \\ \hat{y} - y \\ \hat{z} - z \end{pmatrix} \right\|,$$

where x, y, and z are the true coordinates of the target at that time (as given in the Ground Truth File below), and $\hat{x}$, $\hat{y}$, and $\hat{z}$ are the estimates of these coordinates according to the entity. The "mean absolute error in position" is the mean of E_{pos} for all entity updates (of the confirmed entity), where each update is given equal weight.

Similarly, the "absolute error in velocity", at each entity update, is the scalar magnitude of the velocity error vector:

$$E_{\text{vel}} = \left\| \begin{pmatrix} \hat{v}_x - v_x \\ \hat{v}_y - v_y \\ \hat{v}_z - v_z \end{pmatrix} \right\|,$$

where v_x and v_y are the true velocity components of the tracked target at that time, etc. The "mean absolute error in velocity" is the mean of E_{vel} for all entity updates (of the confirmed entity), where each update is given equal weight.

The "squared relative error in position and velocity", at each entity update, is the quantity

$$D = \begin{pmatrix} \hat{x} - x & \hat{y} - y & \hat{z} - z & \hat{v}_x - v_x & \hat{v}_y - v_y & \hat{v}_z - v_z \end{pmatrix} \mathbf{P}^{-1} \begin{pmatrix} \hat{x} - x \\ \hat{y} - y \\ \hat{z} - z \\ \hat{v}_x - v_x \\ \hat{v}_y - v_y \\ \hat{v}_z - v_z \end{pmatrix},$$

where **P** is the 6-by-6 covariance matrix representing the uncertainty claimed by the entity, with respect to its estimates of x, y, z, v_x, v_y, and v_z. The "RMS relative error in position and velocity" is the square root of the mean of D for all entity updates (of the confirmed entity), where each update is given equal weight.

Each line of the Ground Truth File gives the "ground truth" corresponding to the matching line number of the Input File, whenever the contact or track update of that line originates from a target. The six numbers in each line are
(1) x position (in metres, east of the radar position),
(2) y position (in metres, north of the radar position),
(3) z position (in metres, above the horizontal plane of the radar position),
(4) x-component of velocity (in m/s),
(5) y-component of velocity (in m/s),
(6) z-component of velocity (in m/s).
In the case of an IRST false alarm, the corresponding line of the Ground Truth File consists of six zeroes.

11.2 Additional description

Ground truth, radar tracks, and IRST contacts for Scenario Nine are shown in Figure 12. IRST contacts are shown as bearing lines. The error induced by the simulated radar tracker's assumptions about target altitude is noticeable in the figure. In Figure 13, the same ground truth and IRST data are plotted in (θ, α) space, where θ is bearing and α is the elevation angle.

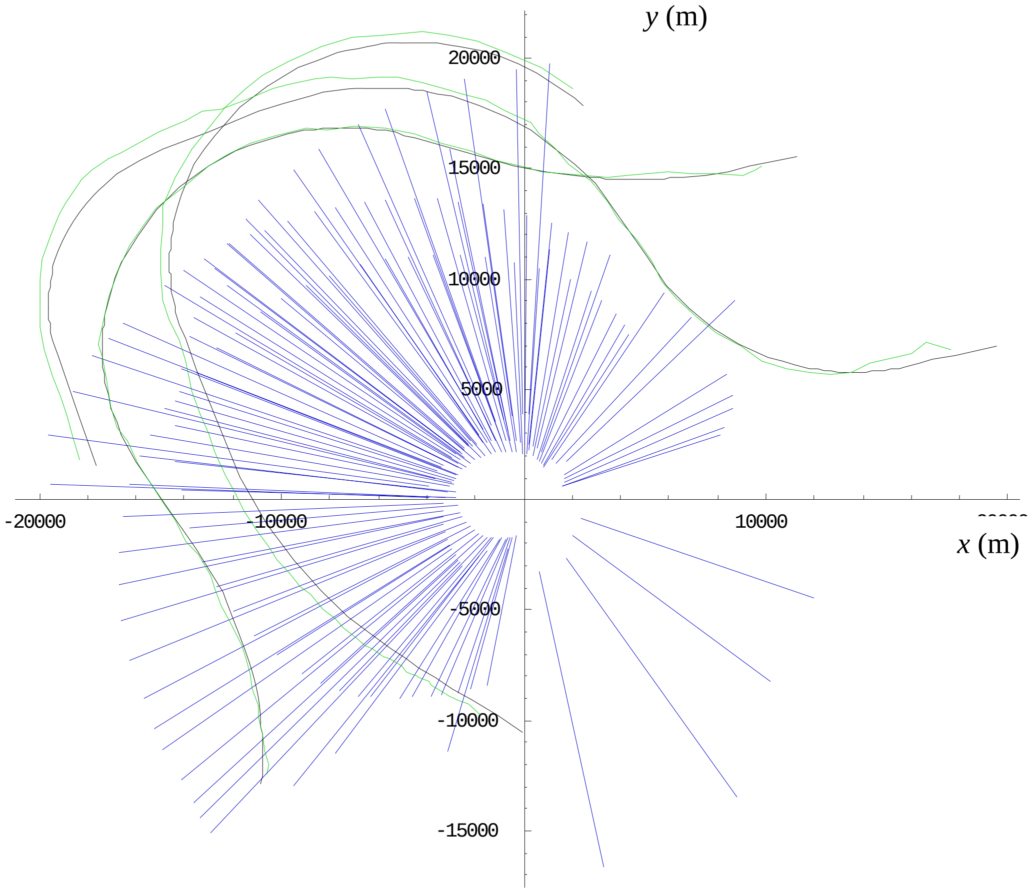

Figure 12. *Ground truth, radar tracks, and IRST contacts in Scenario Nine – 2D positions*

Ground truth is shown in black, radar tracks in green, and IRST contacts in blue. IRST contacts are shown as bearing lines. The sequence of IRST bearing lines is indicated by the length of the lines: Those that occur earlier in the scenario are drawn shorter, while those that occur later are drawn longer. IRST contacts that are excluded from Figure 13 are also excluded from this figure. The target that appears to meet the *y*-axis at about *y* = -10000 m starts at that point, while the other two targets start in the upper right.

11.3 Threshold determination

The radar in this scenario was very similar to that of Scenario Three, but with a lower average rate of false alarms (3.6 contacts per second). Radar tracks were created as in Scenario Six, with the exception of the altitude assumptions as described in Section 11.1. Only a few test runs were performed. In most of them (eight out of ten), exactly three radar tracks were confirmed, and these were maintained throughout the scenario. One of these was chosen arbitrarily for the specification scenario.

The radar track updates that referred to a confirmed track were rearranged as input messages for the fusion system. Track updates were grouped together whenever the original contacts that had prompted those track updates had been grouped together. Each of the resulting radar track update messages contains at most two update reports.

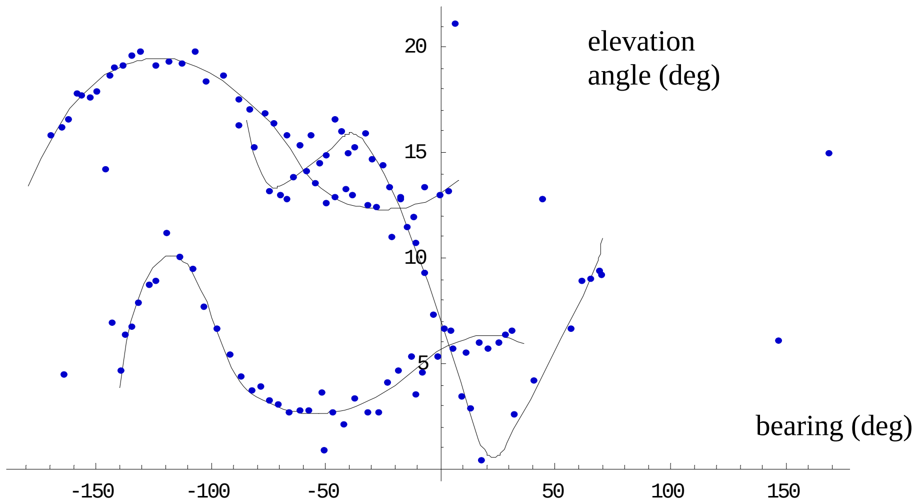

Figure 13. *Ground truth and IRST contacts in Scenario Nine – bearing and elevation*

Ground truth is shown in black, IRST contacts in blue. About a dozen IRST contacts are at an elevation angle greater than 22 degrees, and are excluded from this plot.

IRST contacts were created by a simple modification to the 3D radar simulation used in Scenario Four, with range data deleted. The period of rotation was three seconds, the probability of detection for each target was 0.95, and the false alarms were set to an average rate of about 0.25 contacts per second. Only one set of IRST contacts was generated, and this was used for the specification scenario.

As usual, the specification thresholds were set with reference to the performance achieved by a simple algorithm for association and tracking in three dimensions. The association of IRST contacts with entities, and the updating of entities as a result of association with IRST contacts, were handled according to the method of Chapter 2 (including the method of Section 2.3 to deal with the missing dimensions), except that the vertical component of the velocity was set to zero in the case of an entity that had not previously been associated with any IRST contacts. (This exception was inserted in accordance with the fact that the altitude assigned to a new entity, as in the following paragraph, is arbitrary. Otherwise, the first piece of elevation data would induce a meaningless nonzero vertical velocity to the entity.) Any IRST contact that was not associated with any entity was discarded.

Radar track update reports were associated with entities based on the simple assumption that the radar track numbers were always correct. (This simple expedient was justified by the fact that the targets never come close enough to each other to cause any confusion.) New entities were created from unassociated radar tracks by appending values of 1000 m in the vertical position dimension, with 500,000 m^2 in variance, uncorrelated with the errors in velocity or in horizontal position, and zero vertical velocity, with variance equal to one tenth of the square of the horizontal velocity, uncorrelated with the errors in position or in horizontal velocity. (The latter formula for the vertical velocity variance is based on the assumption that for aircraft in controlled flight, those with faster vertical velocity are usually those that also have faster horizontal velocity.)

When an entity was associated with a radar track update report, the new entity update was created by assuming that the radar track update report is correct with respect to the horizontal components of position and velocity (and the covariance matrix elements involving only horizontal dimensions) and that the predicted (extrapolated) entity is correct with respect to the vertical components of position and velocity. Covariance matrix elements that related a horizontal component of position or velocity to the vertical component of position or velocity were set to zero. The variance in the vertical components of position and velocity was made to agree with either the predicted entity covariance matrix or the initial variances that would have been used if the radar track update report had created a new entity, whichever value was smaller. The covariance between vertical position and vertical velocity was set to zero.

As usual, the performance threshold values were set to be slightly worse than the MOP scores achieved by the method used here. However, it must be admitted that this method is *ad hoc*, and essentially untested outside of the process of designing this scenario.

12. Scenario ten – Triangulation of bearing-only sensors

This scenario, like Scenario One, tests tracking accuracy in two dimensions. But this time all the sensors are bearing-only, and the two-dimensional picture must be built up by triangulation. Three bearing-only sensors are located in three different positions.

In general, the problem of association, in a triangulation scenario, faces special challenges that are not addressed in this scenario. For purposes of this scenario, it is permissible for the candidate system to assume, for simplicity's sake, that there are no false alarms and that all contacts ought to be associated with each other.

The proposed specification statement read as follows: "The HMCCS, if operating in automatic mode and subjected to the input given in the file `s10input.csv`, interpreted according to Annex X.10, shall produce exactly one confirmed entity, whose mean absolute error in position, mean absolute error in velocity, and RMS relative error shall be equal to or better than the thresholds specified in that annex."

The proposed text of the annex is reproduced here in Section 12.1. (An input file and a ground truth file, respectively named `s10input.csv` and `s10gtr.csv`, were intended to be provided with the annex.) Section 12.2 provides additional description of the scenario, while Section 12.3 discusses the choice of performance threshold.

12.1 Specification text

General Comments

This scenario tests the ability of the system to form two-dimensional tracks based on bearing-only contacts from sensors in different positions.

Data Requirements

The following data, from each update of an entity in response to an associated contact, must be stored in order to assess the performance:
(1) The estimated position and velocity, in two dimensions.
(2) The estimated covariance matrix corresponding to the estimated position and velocity, or an alternative expression of uncertainty from which such a covariance matrix may be derived.

Scenario Input

Each line of the Input File contains three numbers, representing a bearing-only contact. The three numbers are
(1) which sensor ("1", "2", or "3") produced the contact,
(2) contact time (in seconds, relative to the beginning of the scenario),
(3) contact bearing (in degrees, clockwise from north).

The contacts shall be interpreted according to the following assumptions:
(1) The uncertainty in bearing of each contact is given by a nominal standard deviation of 2.0 degrees.
(2) Sensor "2" is located 10 km to the north of sensor "1", and sensor "3" is located 10 km to the east of sensor "1". These positions can be treated as reliable.
(3) There is a single aerial target. It can be assumed that there are no false alarms, and that all contacts ought to be associated with each other.

Performance Requirements

In order to meet the specification, the following results are required:
(1) The mean absolute error in position (see below) shall be less than or equal to 880 m.
(2) The mean absolute error in velocity (see below) shall be less than or equal to 130 m/s.
(3) The RMS relative error in position and velocity (see below) shall be less than or equal to 6.5.

Definitions

The "absolute error in position", at each entity update, is the scalar magnitude of the position error vector:

$$E_{\text{pos}} = \left\| \begin{pmatrix} \hat{x} - x \\ \hat{y} - y \end{pmatrix} \right\|,$$

where x and y are the true coordinates of the target at that time (as given in the Ground Truth File below), and $\hat{x}$ and $\hat{y}$ are the estimates of these coordinates according to the entity. The "mean absolute error in position" is the mean of E_{pos} for all entity updates (of the confirmed entity), where each update is given equal weight.

Similarly, the "absolute error in velocity", at each entity update, is the scalar magnitude of the velocity error vector:

$$E_{\text{vel}} = \left\| \begin{pmatrix} \hat{v}_x - v_x \\ \hat{v}_y - v_y \end{pmatrix} \right\|,$$

where v_x and v_y are the true velocity components of the target at that time, etc. The "mean absolute error in velocity" is the mean of E_{vel} for all entity updates (of the confirmed entity), where each update is given equal weight.

The "squared relative error in position and velocity", at each entity update, is the quantity

$$D = \begin{pmatrix} \hat{x} - x & \hat{y} - y & \hat{v}_x - v_x & \hat{v}_y - v_y \end{pmatrix} \mathbf{P}^{-1} \begin{pmatrix} \hat{x} - x \\ \hat{y} - y \\ \hat{v}_x - v_x \\ \hat{v}_y - v_y \end{pmatrix},$$

where $\mathbf{P}$ is the 4-by-4 covariance matrix representing the uncertainty claimed by the entity, with respect to its estimates of x, y, v_x and v_y. The "RMS relative error in position and velocity" is the square root of the mean of *D* for all entity updates (of the confirmed entity), where each update is given equal weight.

Each line of the Ground Truth File gives the "ground truth" corresponding to the matching line number of the Input File. The four numbers in each line are

(1) *x* position (in metres, east of the position of sensor "1"),
(2) *y* position (in metres, north of the position of sensor "1"),
(3) *x*-component of velocity (in m/s),
(4) *y*-component of velocity (in m/s).

12.2 Additional description

Ground truth and sensor contacts for Scenario Ten are shown in Figure 14. Contacts from each sensor are shown as bearing lines, centred on the position of that sensor.

12.3 Threshold determination

The three sensors were derived from a simple modification of the radar simulation used throughout these scenarios. They were each given a rotational period of 4 seconds, a detection probability of 0.95 with each pass, and no false alarms. The standard deviation of the randomly-generated bearing errors was 2 degrees for each sensor, and range data were not reported.

The method used to establish the specification threshold was as follows. The first two contacts (which, conveniently, were reported by different sensors) were used to create a track. The initial position of the track was found by triangulation, and the initial velocity was set to zero. Let r_1 and θ_1 denote the range and bearing of the resulting position relative to the sensor that reported the first contact, and let r_2 and θ_2 denote the range and bearing of the resulting position relative to the sensor that reported the second contact. Then the initial covariance in position (that is, the upper-left-hand
2-by-2 block of the initial full covariance matrix) was set to

$$\frac{2\sigma^2}{\sin^2(\theta_1-\theta_2)}\begin{pmatrix} r_1^2\sin^2\theta_2 + r_2^2\sin^2\theta_1 & r_1^2\sin\theta_2\cos\theta_2 + r_2^2\sin\theta_1\cos\theta_1 \\ r_1^2\sin\theta_2\cos\theta_2 + r_2^2\sin\theta_1\cos\theta_1 & r_1^2\cos^2\theta_2 + r_2^2\cos^2\theta_1 \end{pmatrix},$$

where σ is the standard deviation in bearing for each sensor. (This formula is easily derivable, based on the assumption of small σ and the assumption that we can treat the first bearing measurement as still valid at the time of the second one.) The rest of the covariance matrix was initialised according to the method of Chapter 2.

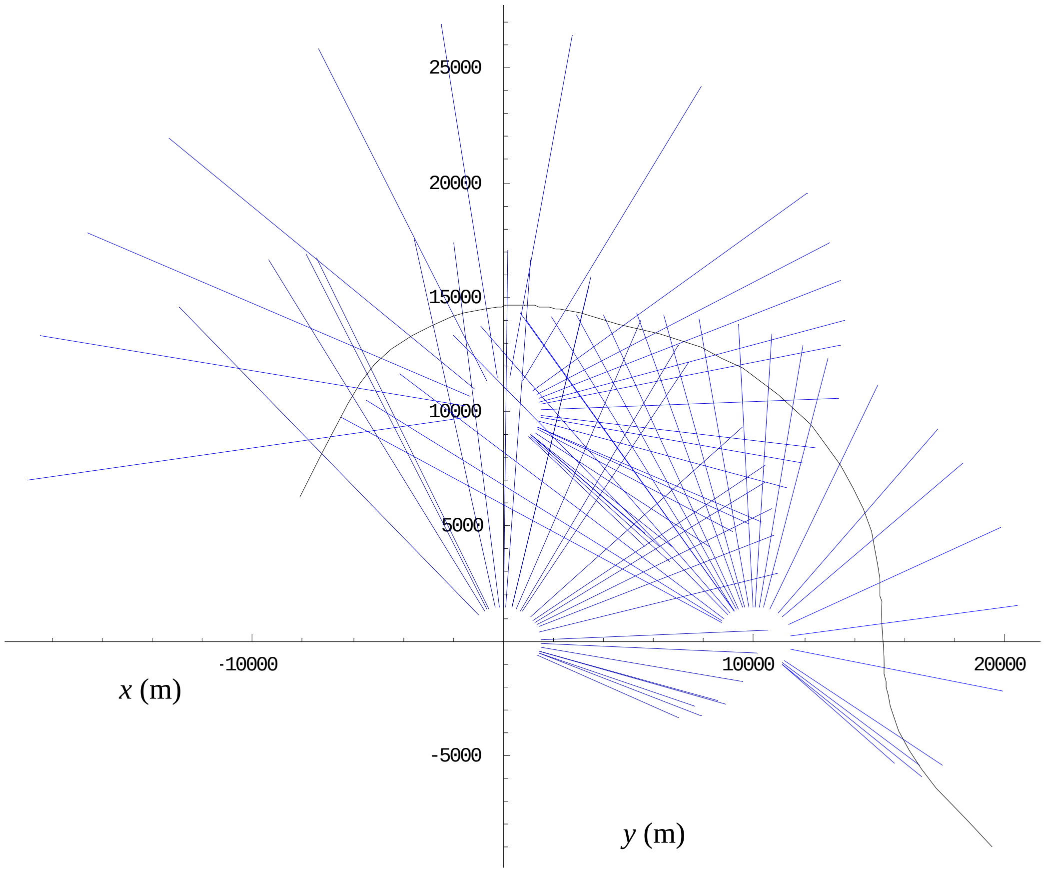

Figure 14. *Ground truth and bearing lines in Scenario Ten*

Ground truth is shown in black, bearing lines in blue. The sequence of bearing lines for each sensor is indicated by the length of the lines: Those that occur earlier in the scenario are drawn shorter, while those that occur later are drawn longer. The target starts in the lower right.

From the third contact onward, the method of Chapter 2 (including the method of Section 2.3 to deal with missing dimensions) was used to update the entity.

Ten test runs were performed. The run that was selected for the specification was one for which the level of performance in all three MOPs was typical of the test runs. As usual, the specification threshold was set to be slightly worse than the level of performance achieved by the simple algorithm.

13. Discussion and Conclusions

Ten scenarios of varying complexity were designed in order to address the task of specifying some of the desired capabilities of the near-future modernised CCS of the HALIFAX Class frigates. Each of these scenarios serves as the basis for a test of some aspect of the performance of a candidate data fusion system. These tests were embodied in draft specification statements for an early phase of the HMCCS project.

Two further scenarios were originally planned as part of this phase of the work, but neither of these was successfully achieved by the time this work stopped. One of them was intended to test the ability of the system to associate system tracks with time-late input data. In order to achieve the performance requirement, some track history would have to be saved. The other planned scenario was intended to test the ability of the system to associate organic (ownship) radar tracks with radar tracks communicated from another platform, in a case where the communicated tracks suffered from an error in registration. (That is, the reported position of the other platform, and/or the alignment of its sensors, would be in error.) The system would be required to detect the registration error and compensate for it.

There are several limitations to the approach of this work that must be acknowledged.

The fusion capabilities embodied in the Mathematica code that was used to generate the thresholds are basic. By using the MOP values generated thereby as the basis of the MOP thresholds for the specification statements, we ensure that the thresholds would be relatively easy to achieve and therefore easy to defend if challenged. However, the resulting specifications would generally not be able to discriminate between a mediocre data fusion system and a very capable one.

Ideally, the MOP thresholds would be based on statistics of the candidate system's performance over a large number of runs, of a large number of scenarios. Such an approach would allow much greater discriminatory power among candidate systems. The disadvantage, of course, is the sheer magnitude of the task of setting up all the scenarios. An examination of how this greater discriminatory power could be achieved efficiently is one possible area of future work.

Because the MOP threshold generation process is primarily based on achievability rather than operational need, those thresholds will be (in a sense) "relative" to the characteristics of the fictional sensors in the scenarios. Intuitively, one expects the relative nature of these thresholds to transfer naturally to any real sensor, thus ensuring the applicability of the tests. A study of the validity of this assumption has not been conducted, and is another possible area for future work.

In any case, the characteristics of the actual sensors that will be used by the upgraded HALIFAX Class frigates are unknown at this time.

One further area of possible future work, inspired by the present work, is the development of algorithms for handling sensor data with missing dimensions.

In conclusion, the set of tests presented in this work provides some discriminatory power for comparisons among candidate systems. A similar set of tests, with some expansion and refinement, could provide considerably more discriminatory power. This work presents a promising approach to the creation of performance tests for a data fusion system.

References

1. HALIFAX Class Modernised Command and Control System (HMCCS) – Statement of Operational Requirement 00002778. Project Sponsor: Chief of the Maritime Staff. January 2002.

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

3. D. J. Peters, unpublished work.

4. C. Gokberk, "Performance Evaluation of Multisensor Data Fusion Systems", Scientific and Technical Report, May 24, 1999, Turkey.

5. A. Sinha, and Y. Li, "Flexible Association Algorithm Development", DRDC Atlantic CR 2008-239.

Distribution list

Document No.: DRDC Atlantic TM 2007-369

LIST PART 1: Internal Distribution by Centre

5 Library

1 B. A. Chalmers

1 F. Desharnais

1 D. G. Hazen

1 M. G. Hazen

1 B. A. McArthur

1 N. McCoy

1 D. J. Peters (author)

---

12 TOTAL LIST PART 1

LIST PART 2: External Distribution by DRDKIM

National Defence Headquarters
MGen George R Pearkes Bldg.
101 Colonel By Drive
Ottawa, ON K1A 0K2
1 DRDKIM

National Defence Headquarters
Louis St Laurent Bldg.
555 boul. de la carriere
Gatineau, QC K1A 0K2
1 Maritime Research Co-ordination Cell
1 DMRS 6
1 DMRS 8
1 DMSS 8
1 PM HCM/FELEX

Defence R&D Canada – Valcartier
2459 Pie XI Blvd North
Val-Belair, QC G3J 1X5
1 Library

Defence R&D Canada – Ottawa
3701 Carling Avenue
Ottawa, ON K1A 0Z4
1 Library

ADM(S&T)
8th floor, Constitution Bldg.
305 Rideau Street
Ottawa, ON K1A 0K2
1 DSTM

Library & Archives Canada
1 Military Archivist, Government Records Branch

---

10 TOTAL LIST PART 2

22 TOTAL COPIES REQUIRED

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).

Tactical Picture Compilation Performance Specifications in HMCCS

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 2008	6a. NO. OF PAGES (total containing information Include Annexes, Appendices, etc). 74	6b. NO. OF REFS (total cited in document) 5

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 2007-369	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).

The HALIFAX Class Modernised Command and Control System (HMCCS) project was a component of the Modernised HALIFAX Class (MHC) project before the latter project was reorganised and renamed in 2007. This document is concerned with a subset of the specifications that were written in order to define the capabilities of the HMCCS; specifically, those specifications that were concerned with the performance of the system in terms of tactical picture compilation.

Each specification was expressed in terms of a threshold value of some Measure of Performance (MOP), to be achieved relative to a predefined scenario and an associated data set. This document details the ten specification statements that were developed, as well as the process used to derive the MOP thresholds. The threshold derivation process was based upon simple data fusion methods, in most cases involving “nearest neighbour” association rules and a Kalman Filter for tracking.

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).

Tactical picture compilation, data fusion, target tracking, HMCCS

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 DÉFENSE

www.drdc-rddc.gc.ca