Copy No. _____

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

DEFENCE DÉFENSE

Advances in Flexible Association / ID-Aided Tracking

D. J. Peters

Defence R&D Canada – Atlantic

Technical Memorandum
DRDC Atlantic TM 2009-296
December 2009

Canada

This page intentionally left blank.

Advances in Flexible Association / ID-Aided Tracking

D. J. Peters

Defence R&D Canada – Atlantic

Technical Memorandum
DRDC Atlantic TM 2009-296
December 2009

Principal Author

Original signed by D. J. Peters

D. J. Peters

Defence Scientist

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

Chief Scientist, DRDC Atlantic

Abstract

The term “Flexible Association / ID-Aided Tracking” (FA/IDAT) refers to a novel approach to data fusion in which kinematic and non-kinematic data are maintained in separate data structures. In comparison with traditional data fusion systems, FA/IDAT is less vulnerable to errors in data association, because of its ability to represent situations that are not easily represented in a traditional data fusion system. In particular, FA/IDAT allows the existence of a known entity to be asserted without committing that identity to a particular associated kinematic track. The representation of such a situation is more natural in FA/IDAT than in Multi-Hypothesis Tracking (MHT). Thus, FA/IDAT has the potential to achieve the benefits of MHT with far less computational load.

In the form in which it was first introduced, the FA/IDAT algorithm lacked the capability to deal with false alarms and missed detections. Such a capability has now been added.

This document presents the current state of development of the FA/IDAT algorithm, points out several areas where further development is required, and discusses the relationship between FA/IDAT and MHT.

Résumé

Le terme « association souple/suivi assisté par identification », ou FA/IDAT en anglais, désigne une nouvelle démarche de fusion de données dans laquelle des données cinématiques et non cinématiques sont stockées dans des structures de données distinctes. Comparativement aux systèmes de fusion de données classiques, l’algorithme FA/IDAT est moins sensible aux erreurs d’association de données en raison de sa capacité de représenter des situations qui ne sont pas faciles à représenter au moyen des systèmes de fusion de données classiques. Plus particulièrement, l’algorithme FA/IDAT permet que l’on puisse poser comme hypothèse l’existence d’une entité connue sans qu’il soit nécessaire d’affecter cette identité à une piste cinématique connexe particulière. La représentation d’une telle situation est plus naturelle avec un système FA/IDAT, que dans le cas du suivi multi-hypothèses (MHT). L’algorithme FA/IDAT dispose ainsi du potentiel requis pour obtenir les avantages du suivi MHT, mais en exigeant une capacité de calcul très inférieure.

Dans sa première forme, l’algorithme FA/IDAT n’était pas en mesure de traiter les fausses alarmes et les absences de détection. Ces capacités ont été ajoutées.

Le présent document présente l’état actuel de développement de l’algorithme FA/IDAT, expose divers aspects qui exigent d’autres travaux de développement et discute de la relation entre l’algorithme FA/IDAT et le suivi MHT.

This page intentionally left blank.

Executive summary

Advances in Flexible Association / ID-Aided Tracking

D. J. Peters; DRDC Atlantic TM 2009-296; Defence R&D Canada – Atlantic; December 2009.

Introduction: The term "Flexible Association / ID-Aided Tracking" (FA/IDAT) refers to a novel approach to data fusion in which kinematic and non-kinematic data are maintained in separate data structures. In comparison with traditional data fusion systems, FA/IDAT is less vulnerable to errors in data association, because of its ability to represent situations that are not easily represented in a traditional data fusion system. In particular, FA/IDAT allows the existence of a known entity to be asserted without committing that identity to a particular associated kinematic track.

Results: FA/IDAT is not in a mature state of development. Existing implementations are not sufficiently general to be applied to a real tactical setting. In the form in which it was first introduced, the FA/IDAT algorithm lacked the capability to deal with false alarms and missed detections. Such a capability has now been added. This document presents in detail the current state of development of the method. The relationship between FA/IDAT and Multi-Hypothesis Tracking (MHT) is also discussed. MHT is an established method for dealing with some of the same difficulties addressed by FA/IDAT, but its theoretical benefits are compromised by the measures that need to be taken to prevent system overload. FA/IDAT has the potential to achieve the benefits of MHT with far less computational load.

Significance: The occasional misidentification of platforms due to association errors is an inevitable result of any tactical data fusion process, with potentially severe consequences. By going beyond the traditional representation of a tactical picture as merely a dynamic set of tracks, the FA/IDAT algorithm allows an automated system to emulate some very simple forms of human reasoning and thereby recover quickly and gracefully from association errors.

Future plans: A version of FA/IDAT with a realistic model for building and maintaining identity/classification data needs to be implemented, thus enabling a test of FA/IDAT using real data. Eventually, it is hoped that a systematic experimental comparison between FA/IDAT and MHT will be performed, demonstrating differences in performance and in computational requirements.

Sommaire

Advances in Flexible Association / ID-Aided Tracking

D. J. Peters; DRDC Atlantic TM 2009-296; R & D pour la défense Canada – Atlantique; Décembre 2009.

Introduction : Le terme « association souple/suivi assisté par identification », ou FA/IDAT en anglais, désigne une nouvelle démarche de fusion de données dans laquelle des données cinématiques et non cinématiques sont stockées dans des structures de données distinctes. Comparativement aux systèmes de fusion de données classiques, l'algorithme FA/IDAT est moins sensible aux erreurs d'association de données en raison de sa capacité de représenter des situations qui ne sont pas faciles à représenter au moyen des systèmes de fusion de données classiques. Plus particulièrement, l'algorithme FA/IDAT permet que l'on puisse poser comme hypothèse l'existence d'une entité connue sans qu'il soit nécessaire d'affecter cette identité à une piste cinématique connexe en particulier.

Résultats : Le développement de l'algorithme FA/IDAT n'a pas atteint la maturité. Ses mises en oeuvre actuelles ne sont pas suffisamment générales pour être appliquées à un environnement tactique réel. Dans sa première forme, l'algorithme FA/IDAT n'était pas en mesure de traiter les fausses alarmes et les absences de détection. Ces capacités ont été ajoutées. Le présent document présente en détail l'état actuel de développement de la méthode. De plus, on discute de la relation entre l'algorithme FA/IDAT et le suivi MHT. La méthode MHT est une méthode établie de résolution des difficultés traitées par FA/IDAT, mais ses avantages sont compromis par les mesures qui doivent être prises afin de prévenir la surcharge du système. L'algorithme FA/IDAT a le potentiel de présenter les avantages de MHT en imposant une charge de calcul nettement inférieure.

Signification : Les erreurs d'identification de plateforme causées par des erreurs d'association constituent un résultat inévitable de tout processus de fusion de données tactiques et elles peuvent avoir des conséquences graves. En allant au-delà de la représentation classique du tableau tactique, représenté simplement par un ensemble dynamique de pistes, l'algorithme FA/IDAT permet à un système automatisé d'émuler quelques formes très simples de raisonnement humain et ainsi de récupérer rapidement et sans heurts des erreurs d'association.

Plans futurs : Il est nécessaire de mettre en oeuvre une version de FA/IDAT au moyen d'un modèle réaliste afin d'élaborer et de tenir à jour les besoins en matière de données d'identification et de classification qui doivent être mis en oeuvre, ce qui permettrait de mettre à l'essai l'algorithme FA/IDAT au moyen de données réelles. On espère qu'il sera éventuellement possible d'effectuer des comparaisons expérimentales systématiques de FA/IDAT et de MHT, ce qui permettra de montrer leurs différences sur les plans du rendement et des besoins de calcul.

Table of contents

This page intentionally left blank.

1 Introduction

A tactical picture is traditionally constructed as a dynamic set of tracks. A track is a set of data expressing what we know (or believe, or guess) about a single alleged entity of interest. The set of tracks is modified in response to new data (or perhaps in response to the absence of expected new data). Ideally, we aim for a one-to-one correspondence between tracks and entities of interest.

A track consists of self-referential data (track management data) and of data pertaining to the corresponding alleged entity. The latter kind of data can be subdivided further into kinematic and non-kinematic data. Non-kinematic data includes all data pertaining to the identity (including allegiance and platform type) of the alleged entity.

Arguably, this structure for a picture – a dynamic set of tracks – is too limiting, no matter how sophisticated the algorithm for updating the track set. In particular, there is no way, within such a scheme, for the existence of a firmly-known identity to be asserted without attaching that identity to a particular set of kinematic data. The example given in section 3 elaborates on this difficulty.

This limitation is addressed, to some extent, by the Multi-Hypothesis Tracking (MHT) [1, 2] method, wherein difficult association decisions are deferred for some time, ideally until better data become available. However, MHT is highly computationally intensive, and its theoretical benefits are compromised by the measures that need to be taken to prevent system overload.

The Flexible Association / ID-Aided Tracking algorithm (FA/IDAT) [3] aims to have the practical benefits of MHT but with much less computational load. Like MHT, FA/IDAT allows a target tracking system to recognize cases of association ambiguity, and to defer certain decisions until better data become available, meanwhile maintaining multiple hypotheses for the association of data. Thus, FA/IDAT could be viewed as a simplification of MHT, in a sense. However, the overall arrangement of tactical picture data in FA/IDAT is quite different from that in any existing MHT implementation, and it corresponds more straightforwardly to the needs of the task of tactical picture compilation.

The essential feature of FA/IDAT is the separation of kinematic and non-kinematic entity data into different data structures. The word "track" will be applied to a package of kinematic data, while a package of non-kinematic data will be called an "ID". Along with these data structures, there is also a dynamic set of association hypotheses linking the tracks to the IDs.

The primary aim of this document is to present the current state of development of the FA/IDAT algorithm, and to discuss the areas where further development is needed.

2 The FA/IDAT Algorithm

The algorithm was previously presented [3] in a restricted form: it was assumed, for the sake of simplicity, that there were no false alarms and no missed detections. Moreover, very little was said about the creation of tracks and IDs, and nothing was said about their deletion. Related to this restriction was the assumption that the total number of tracks and the total number of IDs would always be equal. The latter assumption is arguably inappropriate, since the destruction of a platform during a time of association ambiguity would not then be representable.

This presentation of the algorithm will generalize beyond those restrictions. In particular, the algorithm now explicitly handles false alarms and missed detections, and the number of tracks maintained at any given time can now differ from the number of IDs maintained at that time.

2.1 Notations

Let $\boldsymbol{T}(k)$ denote the set of tracks, and let $\boldsymbol{C}(k)$ denote the set of IDs, at time k. (The phrase "time k" should be taken as shorthand for "the k^{th} iteration", as k is dimensionless and integer-valued.) There will always be at least as many IDs as tracks. (Some IDs may be "null", indicating an assertion that *something* exists to correspond to that ID, but no assertion about the nature of that platform.) Let $\boldsymbol{L}(k)$ denote the table of currently possible global track-to-ID association hypotheses, along with the corresponding probabilities. A global track-to-ID association hypothesis $\boldsymbol{K}$ is a set of specific hypotheses, each associating a track with an ID, subject to the constraint that each track is associated with exactly one ID and each ID is associated with at most one track.

The set of measurements up to time k is denoted Z^k, and the set of measurements at time k is denoted $\mathbf{z}(k)$. The latter set can be decomposed into a set $\mathbf{z}_{\mathrm{d}}(k)$ of kinematic measurements and a set $\mathbf{z}_{\mathrm{g}}(k)$ of feature measurements.

Let $p_{\mathrm{D}}(j)$ denote the probability of detection of the target corresponding to the j^{th} ID. In general, detection probability can be platform-dependent. Assumptions need to be made about what values to assign to the detection probability for different kinds of platforms, and for what detection probability to use in the case of platform whose type has not been resolved. Let λ_{FA} denote the expected number of false alarms (and/or new targets) in each iteration. Note that all the measurements belonging to a given iteration are assumed to come from a single sensor.

In order to process the set $\mathbf{z}(k)$ of measurements at each iteration, the global measurement-to-track hypotheses and global measurement-to-ID hypotheses will have to be considered. A global measurement-to-ID hypothesis $\boldsymbol{I}$ is a set of specific hypotheses, each associating a measurement with an ID, such that at most one ID is associated with each measurement and at most one measurement is associated with each ID. Specific hypotheses within such a global hypothesis will be denoted as (for example) I_b^j for the hypothesis that measurement j originated from the target corresponding to ID b, or I_b^0 for the hypothesis that no measurements originated from that target. Similarly, a global measurement-to-track hypothesis $\boldsymbol{J}$ is a set of specific hypotheses, each associating a measurement with a track, such that at most one track is associated with each measurement and at most one measurement is associated with each track.

Let $\delta_{\boldsymbol{I}}(j)=$ either 0 or 1, depending on whether the j^{th} ID is associated with a measurement or not, according to the hypothesis $\boldsymbol{I}$ – that is, 1 if that ID is associated with a measurement and 0 if it is not. Let $\phi(\boldsymbol{I})$ be the number of false alarms or new targets implied by $\boldsymbol{I}$ (i.e., the number of measurements not associated with any ID according to $\boldsymbol{I}$) and let $\mu(\boldsymbol{I})$ be the number of missed detections implied by $\boldsymbol{I}$ (i.e. the number of IDs not associated with any measurement according to $\boldsymbol{I}$). Note that the difference $\phi(\boldsymbol{I})$ - $\mu(\boldsymbol{I})$ will be equal for all global hypotheses $\boldsymbol{I}$ pertaining to a given set $\mathbf{z}(k)$ of measurements.

Similarly, let $\delta_{\boldsymbol{J}}(j)=$ either 0 or 1, depending on whether the j^{th} track is associated with a measurement or not, according to $\boldsymbol{J}$. The quantities $\phi(\boldsymbol{J})$ and $\mu(\boldsymbol{J})$ are also defined the same way for global measurement-to-track hypotheses as they are for global measurement-to-ID hypotheses.

2.2 Projecting the global track-to-ID association hypotheses

Let $p\left(\boldsymbol{K}(k) \mid Z^{k}\right)$ be the probability for the global track-to-ID hypothesis $\boldsymbol{K}$ according to the table $\boldsymbol{L}(k)$. In section 2.5, the probabilities $p\left(\boldsymbol{K}(k) \mid Z^{k}\right)$ will be derived from the probabilities from the previous iteration, i.e. $p\left(\boldsymbol{K}(k-1) \mid Z^{k-1}\right)$. But in order to get there, we will need the projected probabilities $p\left(\boldsymbol{K}(k) \mid Z^{k-1}\right)$. These projected probabilities will also be used in the update procedures for the tracks and the IDs. We have

$$p\left(\boldsymbol{K}(k) \mid Z^{k-1}\right)=\sum_{\forall \boldsymbol{K}'} p\left(\boldsymbol{K}(k) \mid \boldsymbol{K}'(k-1)\right) p\left(\boldsymbol{K}'(k-1) \mid Z^{k-1}\right) \tag{1}$$

where we approximate $p\left(\boldsymbol{K}(k) \mid \boldsymbol{K}'(k-1)\right)$ according to the amount of kinematic overlap among the tracks, as follows. If the hypothesis $\boldsymbol{K}$ leaves unassociated an ID that is associated with a track according to the hypothesis $\boldsymbol{K}'$, or vice-versa, then $p\left(\boldsymbol{K}(k) \mid \boldsymbol{K}'(k-1)\right)=0$. Otherwise, let $\lambda(t, t')$ denote the likelihood density for track t to be measured at the predicted position of track t' (assuming successful detection). That is, if track t has a projected state estimate $\mathbf{x}_{t}(k \mid k-1)$ and a projected covariance matrix $\mathbf{P}_{t}(k \mid k-1)$, then we let $\lambda(t, t')$ be the value taken by a Gaussian probability distribution function having that covariance and centred on that state estimate, restricted to the measured (position) dimensions, evaluated at the position of the projected state estimate of track t' (i.e., at the position $\mathbf{H}\mathbf{x}_{t'}(k \mid k-1)$, where $\mathbf{H}$ is the measurement matrix). Now

$$p\left(\boldsymbol{K}(k) \mid \boldsymbol{K}'(k-1)\right) \propto \prod_{l} \lambda\left(t_{l}, t_{l}'\right) \tag{2}$$

then normalized so that $\sum_{\forall \boldsymbol{K}} p\left(\boldsymbol{K}(k) \mid \boldsymbol{K}'(k-1)\right) = 1$, where t_l and t'_l denote the tracks that are associated with ID l according to hypotheses $\boldsymbol{K}$ and $\boldsymbol{K}'$ respectively. The IDs included in the product, in equation (2), are only those with which a track is associated according to the hypotheses $\boldsymbol{K}$ and $\boldsymbol{K}'$.

2.3 Updating the IDs

The probability of the specific measurement-to-ID hypothesis I_b^j is given by

$$p\left(I_b^j \mid Z^k\right) = \sum_{\forall \boldsymbol{I} : I_b^j \in \boldsymbol{I}} p\left(\boldsymbol{I} \mid Z^k\right) \tag{3}$$

where the probability for the global measurement-to-ID hypothesis I is given by

$$p\left(\boldsymbol{I} \mid Z^k\right) \propto \frac{\lambda_{\mathrm{FA}}{}^{\phi(\boldsymbol{I})} e^{-\lambda_{\mathrm{FA}}}}{\phi(\boldsymbol{I})!} \left(\prod_{\forall j} \left(p_{\mathrm{D}}(j) \delta_{\boldsymbol{I}}(j) + \left(1 - p_{\mathrm{D}}(j)\right)\left(1 - \delta_{\boldsymbol{I}}(j)\right)\right)\right) p\left(\mathbf{z}(k) \mid \boldsymbol{I}, Z^{k-1}\right) \tag{4}$$

then normalized so that $\sum_{\forall \boldsymbol{I}} p\left(\boldsymbol{I} \mid Z^k\right) = 1$. (Here we assume a Poisson distribution for the false alarms and/or new targets.) Note that j can be zero in equation (3), in which case we are referring to the probability that no measurements are associated with the target corresponding to ID b. If we assume that the probability of detection is constant for all targets, then equation (4) reduces to

$$p\left(\boldsymbol{I} \mid Z^k\right) \propto \frac{\lambda_{\mathrm{FA}}{}^{\phi(\boldsymbol{I})} e^{-\lambda_{\mathrm{FA}}}}{\phi(\boldsymbol{I})!} p_{\mathrm{D}}{}^{n(k)-\mu(\boldsymbol{I})} \left(1 - p_{\mathrm{D}}\right)^{\mu(\boldsymbol{I})} p\left(\mathbf{z}(k) \mid \boldsymbol{I}, Z^{k-1}\right) \tag{5}$$

where $n(k)$ is the number of IDs being considered at time k (that is, the number of IDs in the system after time k-1).

The final term of equations (4) and (5) is given by

$$p\left(\mathbf{z}(k) \mid \boldsymbol{I}, Z^{k-1}\right) = \sum_{\forall \boldsymbol{J}} p\left(\mathbf{z}(k) \mid \boldsymbol{I}, \boldsymbol{J}, Z^{k-1}\right) p\left(\boldsymbol{J} \mid \boldsymbol{I}, Z^{k-1}\right) \tag{6}$$

where $p\left(\boldsymbol{J} \mid \boldsymbol{I}, Z^{k-1}\right)$ is equal to the sum of the probabilities $p\left(\boldsymbol{K}(k) \mid Z^{k-1}\right)$ over all global track-to-ID hypotheses $\boldsymbol{K}(k)$ that are compatible with $\boldsymbol{I}$ and $\boldsymbol{J}$ together. The probability density $p\left(\mathbf{z}(k) \mid \boldsymbol{I}, \boldsymbol{J}, Z^{k-1}\right)$ is approximated as a product of probability densities dictated by the existing tracks and IDs. By decomposing the measurements $\mathbf{z}(k)$ into kinematic measurements and feature measurements, we can express it as

$$p\left(\mathbf{z}(k) \mid \boldsymbol{I},\boldsymbol{J},Z^{k-1}\right) = p\left(\mathbf{z}_{\mathrm{d}}(k) \mid \boldsymbol{J},\boldsymbol{T}(k-1)\right)p\left(\mathbf{z}_{\mathrm{g}}(k) \mid \boldsymbol{I},\boldsymbol{C}(k-1)\right) \tag{7}$$

Now $p\left(\mathbf{z}_{\mathrm{d}}(k) \mid \boldsymbol{J},\boldsymbol{T}(k-1)\right)$ is itself a product of probability densities. If we denote the i^{th} kinematic measurement in $\mathbf{z}_{\mathrm{d}}(k)$ as $\mathbf{z}_{\mathrm{d}}^{(i)}(k)$, then

$$p\left(\mathbf{z}_{\mathrm{d}}(k) \mid \boldsymbol{J},\boldsymbol{T}(k-1)\right) = \prod_{\forall i} p\left(\mathbf{z}_{\mathrm{d}}^{(i)}(k) \mid \boldsymbol{J},\boldsymbol{T}(k-1)\right) \tag{8}$$

If the hypothesis $\boldsymbol{J}$ associates $\mathbf{z}_{\mathrm{d}}^{(i)}(k)$ with an existing track of $\boldsymbol{T}(k-1)$, then we compute $p\left(\mathbf{z}_{\mathrm{d}}^{(i)}(k) \mid \boldsymbol{J},\boldsymbol{T}(k-1)\right)$ by evaluating a Gaussian probability distribution function, given by that track's state estimate and covariance matrix (and restricted to the measured dimensions), at the position of the measurement. Otherwise $p\left(\mathbf{z}_{\mathrm{d}}^{(i)}(k) \mid \boldsymbol{J},\boldsymbol{T}(k-1)\right) = 1/V_{\mathrm{d}}$, where V_{d} is the total volume of the region of interest (restricted to the measured dimensions).

Similarly,

$$p\left(\mathbf{z}_{\mathrm{g}}(k) \mid \boldsymbol{I},\boldsymbol{C}(k-1)\right) = \prod_{\forall i} p\left(\mathbf{z}_{\mathrm{g}}^{(i)}(k) \mid \boldsymbol{I},\boldsymbol{C}(k-1)\right) \tag{9}$$

where $\mathbf{z}_{\mathrm{g}}^{(i)}(k)$ is the i^{th} feature measurement in $\mathbf{z}_{\mathrm{g}}(k)$. How to proceed from here depends on the way in which ID-related data are represented. So far, all work on the FA/IDAT algorithm has treated platforms as having a true value in an abstract Euclidean "feature space", and has treated all feature measurements as independent estimates of that true value. The ID data structure, then, consists of a feature space estimate and a corresponding covariance matrix. Evaluation of $p\left(\mathbf{z}_{\mathrm{g}}^{(i)}(k) \mid \boldsymbol{I},\boldsymbol{C}(k-1)\right)$ is then analogous to evaluation of $p\left(\mathbf{z}_{\mathrm{d}}^{(i)}(k) \mid \boldsymbol{J},\boldsymbol{T}(k-1)\right)$: If the hypothesis $\boldsymbol{I}$ associates $\mathbf{z}_{\mathrm{g}}^{(i)}(k)$ with an existing ID from $\boldsymbol{C}(k-1)$, then $p\left(\mathbf{z}_{\mathrm{g}}^{(i)}(k) \mid \boldsymbol{I},\boldsymbol{C}(k-1)\right)$ is given by evaluating a Gaussian probability distribution function, given by the ID's state estimate and covariance matrix (and restricted to the measured dimensions), at the "position" (in feature space) of the measurement. Otherwise $p\left(\mathbf{z}_{\mathrm{g}}^{(i)}(k) \mid \boldsymbol{I},\boldsymbol{C}(k-1)\right) = 1/V_{\mathrm{g}}$, where V_{g} is the total volume of the feature space (restricted to the measured dimensions).

Alternative calculations of $p\left(\mathbf{z}_{\mathrm{g}}^{(i)}(k) \mid \boldsymbol{I},\boldsymbol{C}(k-1)\right)$ are possible, based on different models of the space of possible feature values or classifications.

Each ID is updated by only one feature measurement: that for which the probability $p\left(I_b^j \mid Z^k\right)$ is greatest. (If $p\left(I_b^j \mid Z^k\right) < p\left(I_b^0 \mid Z^k\right)$ for all $j > 0$, the ID is not updated at all.) Given the way IDs have been represented in this work so far, this update procedure is merely the standard procedure for sequential estimation of a constant vector using observations which are perturbed by additive Gaussian noise. (This is equivalent to a Kalman Filter in the case of a stationary process model and no process noise.) However, we make one modification to this procedure for the sake of caution (in order to avoid contamination among the IDs): For each ID, the specific hypothesis probabilities $p\left(I_b^j \mid Z^k\right)$ are sorted in decreasing order, $\{p_1, p_2, \ldots\}$. Before applying the update to the feature estimate, the measurement covariance is inflated by the factor $1/\left(p_1 - p_2\right)$. (This seemingly *ad hoc* factor is based on finding the set of basic belief assignments that is, in a sense, the "least committed" [4] among those from which our specific hypothesis probabilities could be derived, and weighting the measurement by an amount equal to the weight of the belief assignment that applies to that measurement alone.)

As an alternative, it may be better to ignore the specific measurement-to-ID hypotheses, and to update each ID instead according to the measurement that is associated with it according to the single best *global* measurement-to-ID hypothesis. The sorted hypothesis probabilities would then be for the global hypotheses rather than the specific hypotheses.

The most conspicuous gap in the development of FA/IDAT is in the description of the IDs. Treating an ID as an estimate of a point in feature space is surely an inadequate approach for a real, practical system. We might instead treat an ID as a set of belief assignments, in the sense of Dempster-Shafer theory, but in that case further work is required to find an appropriate way to compute the quantity $p\left(\mathbf{z}_{\mathrm{g}}^{(i)}(k) \mid \boldsymbol{I}, \boldsymbol{C}(k-1)\right)$ in equation (9).

2.4 Updating the tracks

The probability for the global measurement-to-track hypothesis $\boldsymbol{J}$ is given by

$$p\left(\boldsymbol{J} \mid Z^k\right) \propto \frac{\lambda_{\mathrm{FA}}{}^{\phi(\boldsymbol{J})} e^{-\lambda_{\mathrm{FA}}}}{\phi(\boldsymbol{J})!} \left(\prod_j \left(p'_{\mathrm{D}}(j) \delta_{\boldsymbol{J}}(j) + \left(1 - p'_{\mathrm{D}}(j)\right)\left(1 - \delta_{\boldsymbol{J}}(j)\right) \right) \right) p\left(\mathbf{z}(k) \mid \boldsymbol{J}, Z^{k-1}\right) \tag{10}$$

then normalized so that $\sum_{\forall \boldsymbol{J}} p\left(\boldsymbol{J} \mid Z^k\right) = 1$. Here $p'_{\mathrm{D}}(j)$ is the probability of detection of the target corresponding to the j^{th} track, which is given by

$$p'_{\mathrm{D}}(j) = \sum_{\forall \boldsymbol{K}} p\left(\boldsymbol{K}(k) \mid Z^{k-1}\right) p_{\mathrm{D}}(i_j^{(\boldsymbol{K})}) \tag{11}$$

where $i_j^{(\boldsymbol{K})}$ indexes the ID that is associated with track j according to the hypothesis $\boldsymbol{K}$. If we assume that the probability of detection is constant for all targets, then equation (10) reduces to

$$p\left(\boldsymbol{J} \mid Z^k\right) \propto \frac{\lambda_{\mathrm{FA}}{}^{\phi(\boldsymbol{J})} e^{-\lambda_{\mathrm{FA}}}}{\phi(\boldsymbol{J})!} p_{\mathrm{D}}{}^{n_{\text{track}}(k) - \mu(\boldsymbol{I})} \left(1 - p_{\mathrm{D}}\right)^{\mu(\boldsymbol{I})} p\left(\mathbf{z}(k) \mid \boldsymbol{J}, Z^{k-1}\right) \tag{12}$$

where $n_{\text{track}}(k)$ is the number of tracks being considered at time k.

The final term of equations (10) and (12) is given by

$$p\left(\mathbf{z}(k) \mid \boldsymbol{J}, Z^{k-1}\right) = \sum_{\forall \boldsymbol{I}} p\left(\mathbf{z}(k) \mid \boldsymbol{I}, \boldsymbol{J}, Z^{k-1}\right) p\left(\boldsymbol{I} \mid \boldsymbol{J}, Z^{k-1}\right) \tag{13}$$

where $p\left(\boldsymbol{I} \mid \boldsymbol{J}, Z^{k-1}\right)$ is equal to $p\left(\boldsymbol{J} \mid \boldsymbol{I}, Z^{k-1}\right)$ (see previous section). The probability density $p\left(\mathbf{z}(k) \mid \boldsymbol{I}, \boldsymbol{J}, Z^{k-1}\right)$ was discussed in the previous section.

Having found the hypothesis probabilities, the tracks can be updated according to a Kalman Filter, with each track using the measurement that is associated with it according to the single best global measurement-to-track association hypothesis. Alternatively, the tracks could be updated by a weighted average of hypotheses, as in the Probabilistic Data Association Filter (PDAF) [1].

2.5 Updating the global track-to-ID association hypothesis probabilities

The global track-to-ID association hypothesis probabilities are given by

$$p\left(\boldsymbol{K}(k) \mid Z^{k}\right) \propto p\left(\mathbf{z}(k) \mid \boldsymbol{K}(k), Z^{k-1}\right) p\left(\boldsymbol{K}(k) \mid Z^{k-1}\right) \tag{14}$$

then normalized so that $\sum_{\forall \boldsymbol{K}} p\left(\boldsymbol{K}(k) \mid Z^{k}\right) = 1$. See section 2.2 for $p\left(\boldsymbol{K}(k) \mid Z^{k-1}\right)$.

The term $p\left(\mathbf{z}(k) \mid \boldsymbol{K}(k), Z^{k-1}\right)$ is given by

$$p\left(\mathbf{z}(k) \mid \boldsymbol{K}(k), Z^{k-1}\right) = \sum_{\forall \boldsymbol{I}} \sum_{\forall \boldsymbol{J}} p\left(\mathbf{z}(k) \mid \boldsymbol{I}, \boldsymbol{J}, Z^{k-1}\right) p\left(\boldsymbol{I}, \boldsymbol{J} \mid \boldsymbol{K}(k), Z^{k-1}\right) \tag{15}$$

where $p\left(\mathbf{z}(k) \mid \boldsymbol{I}, \boldsymbol{J}, Z^{k-1}\right)$ was dealt with in section 2.3.

Note that the hypothesis $\boldsymbol{J}$ is determined, given $\boldsymbol{I}$ and $\boldsymbol{K}(k)$, because a measurement is associated with a given track if and only if it is also associated with the ID that is associated with that track. Thus, the probability $p\left(\boldsymbol{I}, \boldsymbol{J} \mid \boldsymbol{K}(k), Z^{k-1}\right)$ in equation (15) is zero if the hypotheses $\boldsymbol{I}$ and $\boldsymbol{J}$ together disagree with $\boldsymbol{K}(k)$. Otherwise, it can be expressed without reference to $\boldsymbol{J}$:

$$p\left(\boldsymbol{I}, \boldsymbol{J} \mid \boldsymbol{K}(k), Z^{k-1}\right) = C \frac{\lambda_{\mathrm{FA}}{}^{\phi(\boldsymbol{I})} e^{-\lambda_{\mathrm{FA}}}}{\phi(\boldsymbol{I})!} \left(\prod_{j} \left(p_{\mathrm{D}}(j) \delta_{\boldsymbol{I}}(j) + \left(1 - p_{\mathrm{D}}(j)\right)\left(1 - \delta_{\boldsymbol{I}}(j)\right)\right)\right) \tag{16}$$

where C is a constant that will cancel itself out in the normalization of equation (14). If all detection probabilities are equal, equation (16) can be written as

$$p\left(\boldsymbol{I},\boldsymbol{J}\mid\boldsymbol{K}(k),Z^{k-1}\right)=C\,\frac{\lambda_{\mathrm{FA}}{}^{\phi(\boldsymbol{I})}e^{-\lambda_{\mathrm{FA}}}}{\phi(\boldsymbol{I})!}\,p_{\mathrm{D}}{}^{n(k)-\mu(\boldsymbol{I})}\left(1-p_{\mathrm{D}}\right)^{\mu(\boldsymbol{I})}. \tag{17}$$

2.6 Creation and deletion of tracks and IDs

A measurement (contact) that does not contribute to the updating of an existing ID, and either does not contribute to the updating of an existing track or (in the case of a PDAF-like track update method) has a very small weight (less than some pre-set threshold) for the updating of every existing track, is used to create a new track and a new ID. This new track and this new ID are associated with each other, so every hypothesis in the table $\boldsymbol{L}(k)$ is extended to express this link. The kinematic values of the new track are based on standard procedures for initiating a track from a single measurement. If the measurement has no ID-relevant content, the new ID is null valued.

In order to reflect our uncertainty that the new track-ID pair represent a real object, the detection probability for a newly-formed ID should be suppressed by some factor in the following iteration. That factor would be derived from our prior expectations about the relative frequency of new (real) targets versus false alarms.

A track is deleted if is not updated after some pre-set number of consecutive iterations. (The deletion criteria could depend on the history of the track, with a well-established track being treated more generously than a newly-formed one.) Every hypothesis in the table $\boldsymbol{L}(k)$ is then modified so that an ID which was associated with the now-deleted track is now left unassociated.

An ID is deleted if the total probability of being associated with some track, according to all the hypotheses in the table $\boldsymbol{L}(k)$, drops below some pre-set threshold, and stays there for some pre-set number of consecutive iterations. The hypotheses that associate the now-deleted ID with a track are all deleted, while the remaining hypotheses are re-normalized (and modified to remove references to the now-deleted ID).

This section alludes to several unspecified parameters. Further study is needed to determine appropriate values thereof.

3 Comparison with MHT

In MHT [1, 2], the full tactical picture consists of a set of hypotheses with associated probabilities, each hypothesis being a dynamic set of tracks (like a traditional tactical picture), together with a probability score. Each hypothesis is modified in each iteration, and a hypothesis is split into sub-hypotheses corresponding to the track-to-contact association combinations that are deemed to be feasible. MHT allows difficult decisions to be deferred until better data become available. However, there is a tendency for hypotheses to proliferate exponentially, potentially swamping any computing system. In order to be practical, procedures for pruning and/or merging of hypotheses must be adopted.

Typical implementations of MHT impose a limit on the time during which an association decision can be deferred. The time limit may take the form of a fixed number n_d of system iterations – so that a firm decision is made about the associations of the k^{th} iteration on the basis of the hypothesis probabilities calculated for the $(k + n_d)^{th}$ iteration. (The branch of the hypothesis tree at the k^{th} iteration whose descendants at the $(k + n_d)^{th}$ iteration have the greatest summed probability is kept, while other branches at the k^{th} iteration are dropped.) This time window technique keeps the hypothesis tree from becoming unmanageable. However, the advantages of MHT may be compromised whenever a period of ambiguity exceeds n_d in duration.

In order to compare FA/IDAT with MHT, let us consider a very simple example of association ambiguity (in fact, the very example that originally motivated the development of FA/IDAT). Two targets are being tracked, each with a known identity. The identities are different in some important way, such as one being friendly and the other hostile. The two targets approach each other, closely enough to bring about likely association errors. Then, after some time, they diverge, and the divergence is followed by an ID-relevant measurement which we can associate confidently (on kinematic grounds) with one of the two targets. See Figure 1. (For the sake of this discussion, let us assume that there are few or no false alarms among the contacts.)

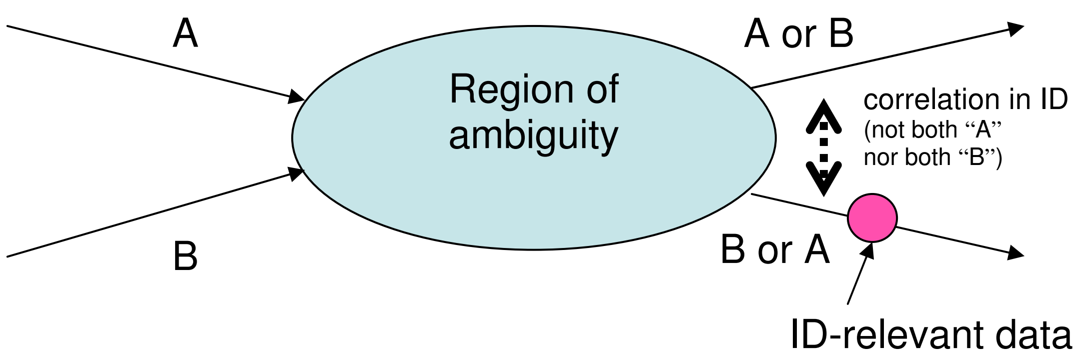

Figure 1: Two targets converge and diverge. We do not know, at first, which is which. If the data represented by the red circle are sufficient to distinguish “A” from “B”, then both targets ought to have their identities restored.

What happens in this scenario? In a simple traditional data fusion system (using firm association decisions), there may easily be an error in association. Suppose that the new data identify the track as friendly, while that track was previously identified as hostile. This clash between old and new data, pertaining to a single track, creates an identity conflict. Although that target is in fact friendly, the system is reluctant to label it as such, until several more ID-relevant measurements confirm that identity. Meanwhile, the other track, previously identified as friendly, now actually corresponds to a hostile target, thanks to the accidental track switch. That error remains (until further evidence shows that target to be hostile), since the system has no way of representing the correlation between the two tracks.

A FA/IDAT system handles this scenario smoothly. (The first simulated scenario in [3] was exactly like the scenario considered here.) The system recognizes the difficulty of the association in the region of ambiguity, and creates a second association hypothesis which links the IDs to the tracks in the opposite way from the original hypothesis. When the targets diverge, both hypotheses will have a substantial probability score. The ID-relevant measurement will then help the system establish easily which track goes with which ID. The (kinematic) tracks may still have traded places, but that association error is rendered irrelevant. From an ID-centric point of view, we could say that an error was avoided, while from a track-centric point of view, we could say that the system recovered gracefully from its error.

Ideally, an MHT system also handles this scenario easily, putting aside limitations in computing power. The number of hypotheses grows exponentially during the time interval of ambiguity, with each hypothesis branching at each iteration. When the targets diverge, these hypotheses will naturally fall into two classes: Those with an even number of association errors, and those with an odd number of association errors. The new ID-relevant measurement will favour the class of hypotheses with an even number of association errors while suppressing the other class of hypotheses.

Notice that each of the two hypotheses in the FA/IDAT system corresponds to one of the two classes of hypotheses in the MHT system. A sufficiently sophisticated MHT system may recognize that the differences among the hypotheses within a single class are minor, and would (perhaps) then merge them into a single hypothesis. But no such merging is required in FA/IDAT, since the hypotheses do not suffer from exponential proliferation.

Now let us consider a small modification to the scenario. Suppose that among the contacts being used to track the targets during the time interval of ambiguity, there are some with ID-relevant data. Presumably, these contacts will be used to update the ID content of the tracks within each MHT hypothesis. Given that some (indeed, most) of the hypotheses will contain some association errors, the ID content of the tracks will be compromised. With the ID modifications included, the differences among hypotheses belonging to the same class (i.e., having either an odd or an even number of errors) may now be significant. The tracks in many of the hypotheses will suffer from ID contamination, since their ID components will have been built up from a mixture of contacts originating from target “A” and contacts originating from target “B”, and this contamination will tend to persist for some time after the divergence. By contrast, a FA/IDAT system will maintain only two IDs, and its cautious approach to ID update will tend to keep the two IDs pure.

Finally, let us consider the implications of the time window method that is commonly used to keep the number of hypotheses manageable in MHT. Suppose that the duration of the time interval of ambiguity is longer than the time window for pruning the hypothesis tree. It may easily happen that the wrong branch of the tree will be selected as the one to survive the pruning. Taking this observation together with the discussion in the previous paragraph of ID contamination, it seems that the time-window-based pruning of hypotheses in MHT can, in some cases, severely compromise the theoretical benefits of the method.

4 Future Work

The existing rules for the creation and deletion of tracks and IDs need to be generalized further.

- The amount and the duration of the suppression of the detection probability of a new ID (see the second paragraph of section 2.6) should be formalized.

- A measurement that does not contribute to the updating of a track, but does contribute to the updating of an ID, should perhaps be used to initiate a new track in some cases, such that some of the global track-to-ID hypotheses will associate the new track with the just-updated ID. The criteria for this kind of track creation need to be made precise. The implications of this kind of track creation also need to be examined, with particular attention given to the possibility of some global track-to-ID hypotheses leaving some tracks unassociated.

- If an ID is associated (according to some hypotheses) with a track that is lost, it is possible for an additional ID to be created with reference to the same target. One possible response is to allow IDs to be merged in some cases. However, unlike the association of a track to an ID, the merging of IDs would be irreversible. So such a step would have to be made with great care.

A realistic model for building and maintaining IDs needs to be worked into FA/IDAT, with particular attention given to the calculation of $p\left(\mathbf{z}_{\mathrm{g}}^{(i)}(k) \mid \boldsymbol{I}, \boldsymbol{C}(k-1)\right)$ (see the final comments of section 2.3). The ID aspects could be based on the approach used in the Command Decision Aids Technology (COMDAT) [5] Technology Demonstration project, for example. A test of FA/IDAT using real data, such as the data from the COMDAT sea trials, would also be worthwhile in order to uncover hidden weaknesses of the algorithm.

Finally, a systematic experimental comparison between FA/IDAT and MHT, demonstrating differences in performance and in computational requirements, would be enlightening.

5 Conclusions

The FA/IDAT algorithm has much in common with MHT: it allows a target tracking system to recognize cases of association ambiguity, and to defer certain decisions until better data become available, meanwhile maintaining multiple hypotheses for the association of data. However, FA/IDAT avoids the tremendous proliferation of hypotheses that is typical of MHT, and the resulting challenges of implementation, by structuring the data differently.

The FA/IDAT algorithm has been extended beyond the basic form in which it was originally introduced: it now accounts for false alarms and missed detections. Also, the number of tracks (packages of kinematic data) is no longer constrained to be equal to the number of IDs (packages of non-kinematic data), thus allowing the system to handle track loss during a period of association ambiguity. The current version of the FA/IDAT algorithm was presented here in detail, and the areas of required further development were described.

References

[1] Y. Bar-Shalom and X.-L. Li, Multitarget-Multisensor Tracking: Principles and Techniques, YBS Publishing, Storrs, CT, 1995.

[2] S. S. Blackman and R. Popoli, Design and Analysis of Modern Tracking Systems, Artech House, Norwood, MA, 1999.

[3] A. Sinha and D. J. Peters, "Flexible ID Association-based Tracking Algorithm," DRDC Atlantic SL 2009-055, Proceedings of the Twelfth International Conference on Information Fusion, July 2009.

[4] B. Ristic and P. Smets, "Target Identification using Belief Functions and Implication Rules," IEEE Trans. On Aerospace and Electronic Systems, Vol. 41, No. 3, pp. 1097-1103, July 2005.

[5] B. A. McArthur, A. T. Ashley, É. Bossé, and D. Beevis, "The Command Decision Aid Technology (COMDAT) Project", TTCP Symposium on Battlespace Management, May 1999.

Distribution list

Document No.: DRDC Atlantic TM 2009-296

LIST PART 1: Internal Distribution by Centre

3	Library
1	F. Desharnais
1	M. Hazen
1	B. McArthur
1	G. Mellema
1	D. Peters (author)
8	TOTAL LIST PART 1

LIST PART 2: External Distribution by DRDKIM

1	Library and Archives Canada, Atten: Military Archivist, Government Records Branch
1	É. Bossé DRDC Valcartier 2459 Pie-XI Blvd North Val-Belair, QC G3J 1X5
1	A. Benaskeur DRDC Valcartier 2459 Pie-XI Blvd North Val-Belair, QC G3J 1X5
1	DRDKIM National Defence Headquarters MGen George R Pearkes Bldg. 101 Colonel By Drive Ottawa, ON K1A 0K2
4	TOTAL LIST PART 2

12 TOTAL COPIES REQUIRED

This page intentionally left blank.

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.) Defence R&D Canada Atlantic 9 Grove Street P.O. Box 1012 Dartmouth, Nova Scotia B2Y 3Z7	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 or U) in parentheses after the title.)

Advances in Flexible Association / ID-Aided Tracking

4. AUTHORS (last name, followed by initials – ranks, titles, etc. not to be used)

D. J. Peters

5. DATE OF PUBLICATION (Month and year of publication of document.) December 2009	6a. NO. OF PAGES (Total containing information, including Annexes, Appendices, etc.) 26	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
9 Grove Street
P.O. Box 1012
Dartmouth, Nova Scotia 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.) 11bs	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 2009-296	10b. OTHER DOCUMENT NO(s). (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.)

Unlimited

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.))

Unlimited

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 term "Flexible Association / ID-Aided Tracking" (FA/IDAT) refers to a novel approach to data fusion in which kinematic and non-kinematic data are maintained in separate data structures. In comparison with traditional data fusion systems, FA/IDAT is less vulnerable to errors in data association, because of its ability to represent situations that are not easily represented in a traditional data fusion system. In particular, FA/IDAT allows the existence of a known entity to be asserted without committing that identity to a particular associated kinematic track. The representation of such a situation is more natural in FA/IDAT than in Multi-Hypothesis Tracking (MHT). Thus, FA/IDAT has the potential to achieve the benefits of MHT with far less computational load.

In the form in which it was first introduced, the FA/IDAT algorithm lacked the capability to deal with false alarms and missed detections. Such a capability has now been added.

This document presents the current state of development of the FA/IDAT algorithm, points out several areas where further development is required, and discusses the relationship between FA/IDAT and MHT.

Le terme « association souple/suivi assisté par identification », ou FA/IDAT en anglais, désigne une nouvelle démarche de fusion de données dans laquelle des données cinématiques et non cinématiques sont stockées dans des structures de données distinctes. Comparativement aux systèmes de fusion de données classiques, l'algorithme FA/IDAT est moins sensible aux erreurs d'association de données en raison de sa capacité de représenter des situations qui ne sont pas faciles à représenter au moyen des systèmes de fusion de données classiques. Plus particulièrement, l'algorithme FA/IDAT permet que l'on puisse poser comme hypothèse l'existence d'une entité connue sans qu'il soit nécessaire d'affecter cette identité à une piste cinématique connexe particulière. La représentation d'une telle situation est plus naturelle avec un système FA/IDAT, que dans le cas du suivi multi-hypothèses (MHT). L'algorithme FA/IDAT dispose ainsi du potentiel requis pour obtenir les avantages du suivi MHT, mais en exigeant une capacité de calcul très inférieure.

Dans sa première forme, l'algorithme FA/IDAT n'était pas en mesure de traiter les fausses alarmes et les absences de détection. Ces capacités ont été ajoutées.

Le présent document présente l'état actuel de développement de l'algorithme FA/IDAT, expose divers aspects qui exigent d'autres travaux de développement et discute de la relation entre l'algorithme FA/IDAT et le suivi MHT.

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 is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

flexible association, target tracking, data fusion

This page intentionally left blank.

Defence R&D Canada

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

R & D pour la défense Canada

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

DEFENCE R&D DÉFENSE

www.drdc-rddc.gc.ca