Reformulating Reid's MHT method with generalised Murty *K*-best ranked linear assignment algorithm

R. Danchick and G.E. Newnam

Abstract: The authors reformulate Reid's multiple hypothesis tracking algorithm to exploit a *K*-best ranked linear assignment algorithm for data association. The reformulated algorithm is designed for real-time tracking of large numbers of closely spaced objects. A likelihood association matrix is constructed that, for each scan, for each cluster, for each cluster hypothesis, exactly and compactly encodes the complete set of Reid's data association hypotheses. The set of this matrix's feasible assignments with corresponding non-vanishing products is shown to map one-to-one respectively onto the set of Reid's data association hypotheses and their corresponding probabilities. The explicit structure of this matrix is a new result and leads to an explicit hypothesis counting formula. Replacement of the likelihood association matrix elements by their negative natural logs then transforms the data association matrix into a linear assignment problem matrix and recasts the problem of data association into efficiently finding sets of ranked assignments. Fast polynomial time Murty ranked assignment algorithms can thus replace Reid's original NP-hard exhaustive hypothesis identification, probability evaluation, and branch-and-prune methods and can rapidly determine the maximally likely data association hypothesis, the second most likely, etc. Results from two high fidelity surveillance sensor simulations show the validity of the proposed method.

1 Introduction

1.1 Historical background

It is a well-established tracking community consensus that multiple hypothesis tracking (MHT) methods are of compelling theoretical importance for attacking the closely spaced object multi-target tracking problem. Reid's original algorithm [1] was seminal in its sequential formulation, definitions, notation, derivations, data structure design and overall conceptual elegance. It featured the concepts of clusters and super-clusters for partitioning the global data association problem into smaller sub-problems. Clusters are collections of multi-scan measurements that have been recursively organised into tracks by track gate inclusion. Super-clusters are formed by the union of those clusters which share a common measurement on a new scan.

After processing the set of measurements for a given scan k, the set of all sensor measurements from scan 0 to scan k are contained in the disjoint union of all scan k clusters; no two distinct clusters contain the same measurement. Moreover, under a given scan k hypothesis, no two tracks share the same measurement; no track can have more than one associated measurement per scan.

Reid's algorithm maintains multiple hypotheses as it processes each new scan's measurement set. The retention of only a fixed maximum number, K, of the most likely of such hypotheses from scan-to-scan provides an efficient and accurate approximate solution to the NP-hard multi-scan data association problem first posed by Morefield [2].

Over time, the tracking community has come to understand that Reid's basic technique of hypotheses branching and pruning for a given cluster still left an exponentially growing number of hypotheses to be identified and their probabilities to be computed. This is clearly the case for the large target numbers and high spatial density problems we have in mind. Thus, as originally formulated, Reid's multi-target tracking algorithm is not practical for real-time applications.

Taking off on an algorithmic theme of Nagarajan *et al.* (NCS) [3], Danchick and Newnam [4] showed that collapsing the categories of new targets and false targets into one category enabled a *K*-best ranked assignment attack on the data association problem. Their paper described the new DN *N*-best assignment algorithm and gave comparative timings between the NCS and the DN algorithms on a set of realistic assignment problem test matrices, problem matrix dimensions ranging from 5×5 to 30×30. The timing comparisons showed that the DN algorithm was practical for large-scale real-time applications, whereas the NCS algorithm was not.

Progress on the ranked assignment approach to MHT tracking has been rapid since 1993 [4]. Cox and Miller [5] went back to the operations research literature of the late 1960s and found an extremely efficient ranked assignment algorithm due to Murty [6] and have implemented a highly optimised version of Murtys' algorithm, specialised for MHT [7], based on the Jonker–Volgenant [8] linear assignment problem algorithm.

In a joint paper published in 1997, Cox *et al.* [9] showed the close relationship between the Murty and DN

IEE Proceedings online no. 20050041

doi:10.1049/ip-rsn:20050041

Paper first received 13th May and in revised form 28th July 2005

R. Danchick is a consulting mathematician at 100336 Wilshire Blvd. #101, Los Angeles, CA 90024, USA

G.E. Newnam is with Volt Technical Resources, Orange, CA 92857, USA

E-mail: dodeee@sbcglobal.net

algorithms. They compared the performance of several versions of the two algorithms. The performance dominance of the Murty algorithm was clearly established. The optimised Murty algorithm was on the average about 100 times faster than the DN algorithm over two sets of test cases. The first problem set was a collection of dense assignment problems of dimension 20×20 with independent, identically distributed matrix entries uniformly distributed in the interval, $[-20, 20]$. The second set consisted of four problems whose dimensions were 10×10, 20×20, 30×30, 40×40. This set was obtained from the simulation of a geostationary infrared sensor system for the detection and tracking of ballistic missiles in their boost phase. The authors also noted that an accurate low-order polynomial time version of Reid's original and complete algorithm was now feasible.

The intent of this paper is to return to Reid's original and complete MHT algorithm and show how it can be recast in the context of K-best ranked assignments by using a generalised rectangular matrix Murty algorithm. The generalisation is easily obtained by restricting the Murty algorithm nodes' row indices to be less than or equal to the number of measurements and squaring-off the Murty algorithm's linear assignment problem sub-matrices prior to application of the J–V algorithm by appending blocks of zeros to the bottom of these sub-matrices. By complete, we mean that a given measurement may hypothetically be allocated to an existing target track, it can initiate a new target track, or it may be categorised as a false track.

1.2 Technical background

The Reid algorithm is a multi-target MHT algorithm. It is a sequential formulation of Morefield's batch algorithm which posed the multi-scan data association problem as a $0-1$ integer linear programming problem. The scope of this paper covers: (1) the replacement of Reid's tree data structure for encoding all possible feasible hypothetical assignments of scan k measurements to scan $k-1$ tracks, new tracks or false tracks by a specially structured linear assignment problem matrix; (2) efficiently determining from this matrix a set of the K-best hypotheses ranked in decreasing order of likelihood headed by the maximally likely hypothesis, the next most likely, and so on; (3) the performance of a prototype algorithm that exploits our ranked assignment theory.

Fig. 1 outlines the underlying concept of our K-best data association algorithm. The figure shows how a given cluster and its K-best hypothesis on scan $k-1$ stretches out to K^2 new hypotheses when a new scan's data is input to the algorithm. These hypotheses are then squeezed back down to the K-best for scan k by ranking them in descending order of likelihood. The diagram depicts how, at the end of scan k, each cluster has K data association hypotheses that are the K most likely histories of how the cluster's initiating scan's measurements and previous scan's measurements have been assembled into tracks. Each such cluster hypothesis, h, has a corresponding hypothesis probability, p_{h}^{k}, a set of track state estimates, $\boldsymbol{x}_i$, their error covariance matrices, $\boldsymbol{P}_i$, and a set of multi-scan measurements, $\bigcup_{n=0}^{k} \{\boldsymbol{z}_m^n: m = 1, 2, \ldots, M_n\}$. These measurements constitute the raw multi-scan data that the algorithm has organised into individual tracks per cluster and cluster hypothesis. For each cluster, cluster hypotheses and track belonging to the cluster hypotheses, the algorithm maintains pointers to the individual scan-by-scan measurements that have been assembled into the track. In operational applications, these measurements will also include measurement times and such important ancillary data as multi-band focal plane irradiances, or radar cross-sectional area and their corresponding measurement error covariance matrices.

At the beginning of scan k cluster processing, a subset of the set of all scan k measurements is associated with a given cluster because each of the subset's measurements is contained in the track gate for some track belonging to some cluster hypothesis belonging to the given cluster. These M_k measurements together with the hypothesis N_{TGT} track state estimates and corresponding error covariance matrices (both propagated to the appropriate measurement times) are used to generate a linear assignment problem matrix, $\boldsymbol{L}^*$, which compactly encodes all of Reid's feasible data associations. A generalised rectangular matrix Murty K-best ranked assignment algorithm is applied to $\boldsymbol{L}^*$. This algorithm computes the K most likely hypotheses for each of the K previously retained hypotheses. The resulting K^2 hypotheses are then pruned down by sorting them in descending order of likelihood to the K-best per each cluster.

At the end of scan k cluster processing, updated hypotheses probabilities are recursively computed from Reid's fundamental (16) in [1]. Finally, for each cluster, for each of the newly retained hypotheses per cluster, each of the hypothesis individual track state estimates are updated with the hypothesis-associated measurement or new tracks are initiated for unassociated measurements. We have found that, in practice, $K \leq 10$ suffices to yield overall track purities and track count accuracies in the high 90 percentiles.

The algorithm has been applied to SDI-related geostationary infrared sensor system simulations, to the National Missile Defense System's Ground-Based Interceptor program, and to the HAVE GOLD Early Warning Radar upgrade program. The scale of the applications runs from 10 objects in the radar application to 200 objects in the infrared sensor application.

Fig. 2 shows the ranked assignment approach embedded in the flow diagram of our operational multi-target, MHT algorithm.

In the flow diagram shown in Fig. 2, the block labelled 'FIRST SCAN PROCESSING' initiates a cluster for each first scan measurement. Each such first scan measurement establishes a single hypothesis of unit probability and a single scan-initiated state estimate and its error covariance matrix.

Sequentially, each new scan's measurements are input by 'READ NEXT DATA SET'. Each such measurement is tested for association with each cluster in the set of prior scan clusters in 'ASSOCIATE WITH PRIOR CLUSTER'. A measurement on scan k is associated with a scan $k-1$ cluster if the measurement is contained in the gate of a propagated track under one of the cluster's hypotheses. If the measurement does not associate with any prior cluster, it establishes its own cluster as described in 'FIRST SCAN PROCESSING'.

'COMBINE CLUSTERS WITH LIKE MEASUREMENTS' forms new scan super-clusters from the union of those prior clusters which have at least one common associated new measurement.

For each of the new resulting clusters, cluster hypothesis, the set of associated measurements and hypothesis track state estimates and covariance matrices are used to generate the complete Reid assignment matrix. This rectangular matrix of negative log-likelihood's encodes all possible feasible hypothetical data associations for a given cluster and cluster hypothesis.

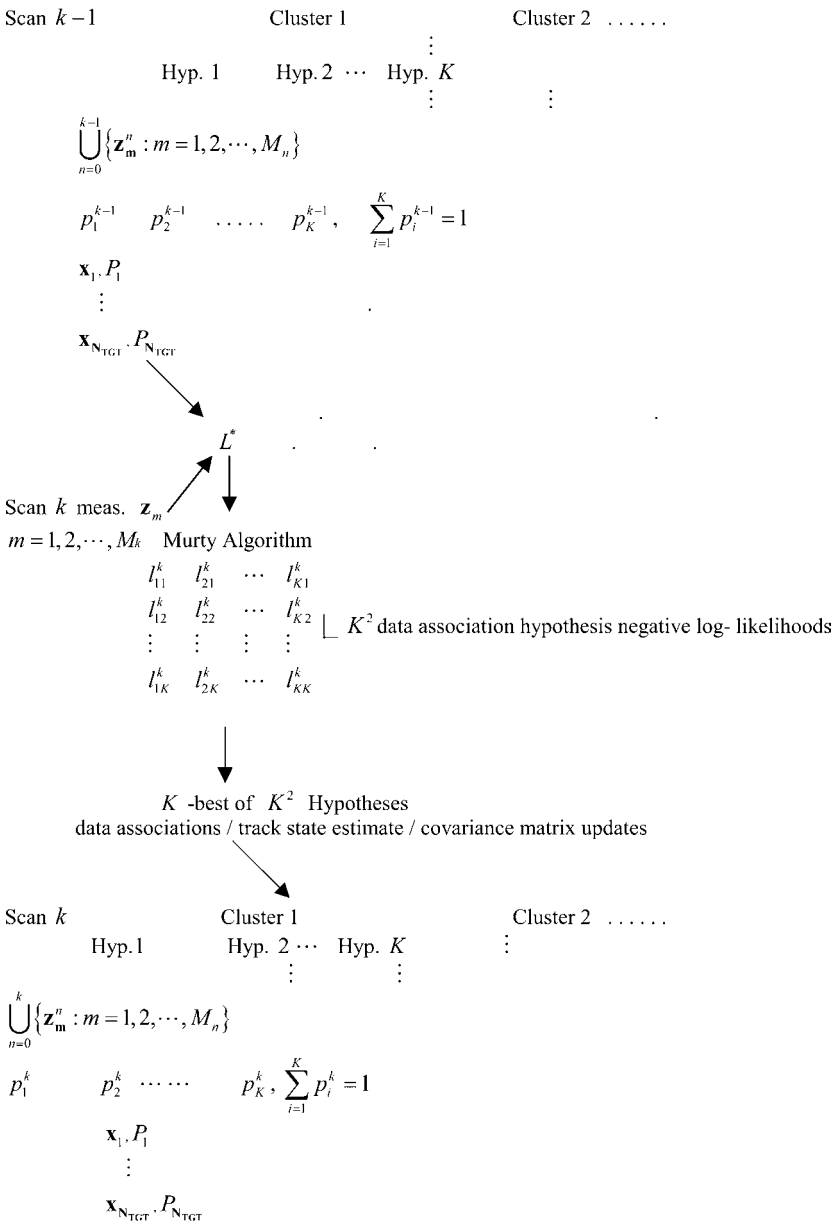

Fig. 1 *K-best hypotheses schema*

The generalised rectangular Murty ranked assignment algorithm is applied to each such matrix in 'GENERATE K-BEST RANKED ASSIGNMENTS'. The resulting K^2 hypotheses are then pruned back down to K as shown in Fig. 1.

For each of the cluster's K-best hypotheses and the hypotheses' propagated state estimates with an associated measurement, the propagated state is updated with the measurement and the covariance matrix of errors in the updated state estimate is computed. Those measurements that have been associated with the given cluster and correspond to a new target arrival under a given hypothesis initiate a new track state and error covariance matrix under the hypothesis. Those measurements that do not fall within the track gate generated by a target state of any existing cluster initiate new one-measurement clusters, a single hypothesis with unit probability and a single one-point initiated state vector estimate and error covariance matrix.

In 'IDENTIFY CONFIRMED TARGETS' those tracks that have the same associated measurement under all K-retained hypotheses are labelled as confirmed as in [1].

Section 2 provides our assumptions and notational setting. Section 3 states and proves the main result of this paper: For each cluster, cluster hypothesis, there exists a scan $k-1$ to scan k likelihood data association matrix such that the set of all its feasible assignments with non-vanishing likelihood products is exactly Reid's set of data association hypotheses and, up to a hypothesis-independent factor, their corresponding hypothesis probabilities. Thus, by replacing the data association matrix elements by their negative natural logs (with zero likelihood's replaced by a suitable strict upper bound), the maximally likely scan k data association hypothesis, the next most likely, and so on, can be determined directly by a generalised Murty ranked assignment algorithm. This ranking of hypotheses eliminates the necessity for the exhaustive hypothesis enumeration and ever-growing hypothesis tree data structures that were features of the original Reid algorithm. The section also contains the derivation of the formula (22) for the number of scan $k-1$ to scan k data association hypotheses as a function of the number of measurements and existing targets for a given scan $k-1$ cluster and hypothesis.

In Section 3.4, we present a two-way table of the total number of $k-1$ to scan k hypotheses for small-scale problems. To provide a concrete instance, we list the symbolic

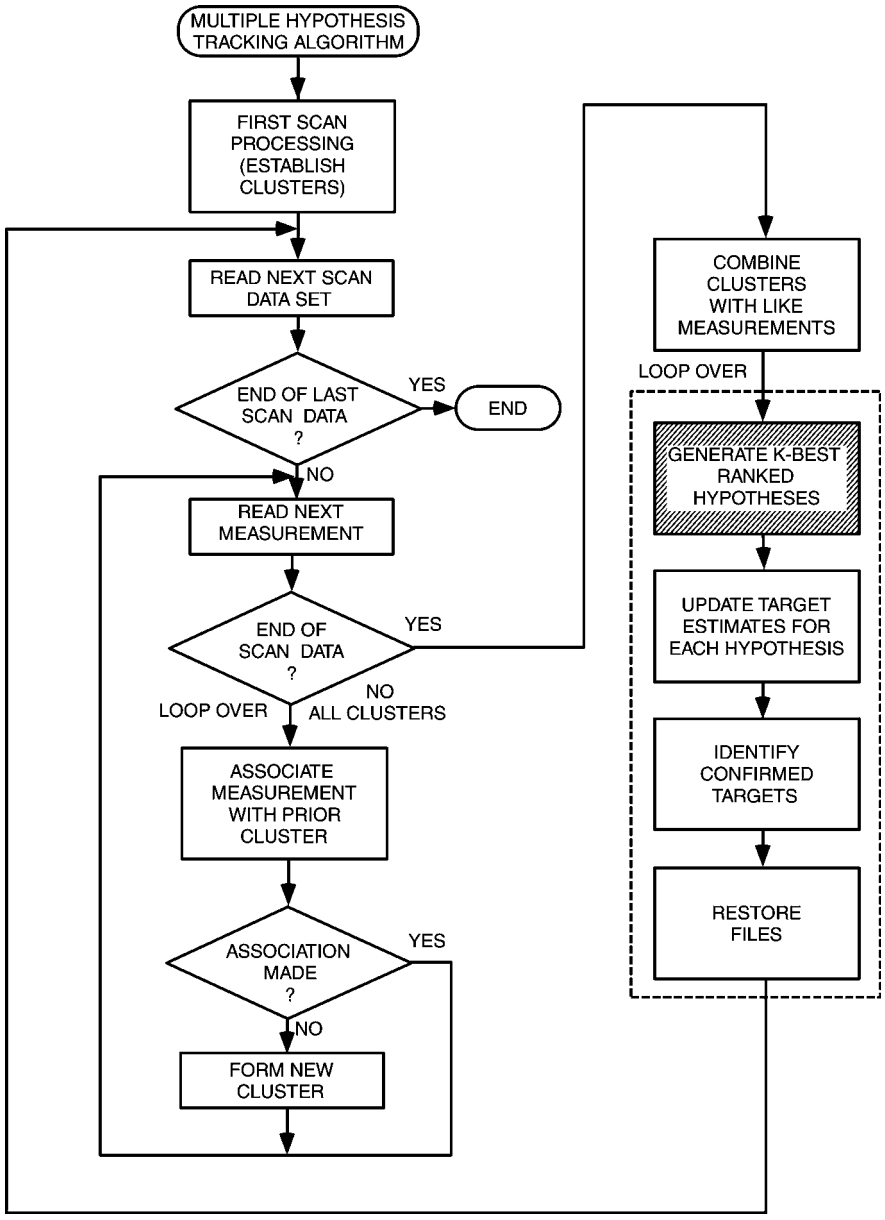

Fig. 2 *MHT algorithm flow diagram*

data association hypotheses and their likelihood products for the smallest scale problem in the table.

In Section 4, we show how our MHT algorithm, which features the generalised Murty algorithm, performs on two simulated SDI-related test cases, a moderate scale, moderately dense scenario of 63 objects and a large scale, dense scenario of 200 objects. Section 5 summarises where we now stand with the ranked assignment version of Reid's MHT technique and identifies some potential current and future applications.

2 Assumptions and notation

In the following, we make the usual Kalman filter-based tracking system assumption about the linearity of state dynamics equations, measurement equation linearity, Gaussian white noise on state estimation process, noise and measurement errors. We also assume Reid's Poisson density-based arrival rate models for new targets and false targets. Although the notational setting is one in which all scan measurements arrive at the same scan time, the algorithm is easily adapted to the situation in which measurements may arise at any time during a given scan. Indeed, we have made exactly this adaptation for the scanning infrared sensor in the simulation described earlier.

Reid's notation is paraphrased as follows: g, an arbitrary prior scan $k-1$ hypothesis for a given cluster; p_{g}^{k-1}, probability of the prior hypothesis g which has been recursively computed from all data assignments in previous scans, 0, 1, 2, ..., $k-1$; k, current scan; h, an arbitrary scan k data association hypothesis given prior hypothesis g; p_{D}, probability of detection; β_{NT}, density of new targets (β_{NT} includes p_{D} as a factor); β_{FT}, density of false targets; dim($\boldsymbol{z}$), dimension of measurement vectors; dim($\boldsymbol{x}$), dimension of target state estimate vectors; M_k, number of measurements associated by gating with the given cluster on scan k; $\boldsymbol{z}_m$, measurement vector m, $m = 1, 2, \ldots, M_k$ on scan k given g; $\boldsymbol{R}$, uniform measurement error covariance matrix; N_{TGT}, number of old targets on scan $k-1$ given g; $\boldsymbol{x}_n(k \mid k-1)$, propagated target state estimate n, $n = 1, 2, \ldots, N_{\mathrm{TGT}}$; $\boldsymbol{P}_n(k \mid k-1)$, propagated target state error covariance matrix $\boldsymbol{n}$; $\boldsymbol{H}$, common dim($\boldsymbol{z}$) $\times$ dim($\boldsymbol{x}$) measurement matrix; $\boldsymbol{r}_{mn} = \boldsymbol{z}_m - \boldsymbol{H}\boldsymbol{x}_n(k \mid k-1)$, the residual given that measurement $\boldsymbol{z}_m$ is associated with the propagated state vector estimate $\boldsymbol{x}_n(k \mid k-1)$; $\Sigma_{mn} = \boldsymbol{R} + \boldsymbol{H}\boldsymbol{P}_n(k \mid k-1)\boldsymbol{H}^{\mathrm{T}}$, the residual covariance matrix;

$l_{mn} = \exp\{-(1/2)\boldsymbol{r}_{mn}^{\mathrm{T}}\Sigma_{mn}^{-1}\boldsymbol{r}_{mn}\}/\sqrt{\{(2\pi)^{\dim(z)}\det(\Sigma_{mn})\}}$, the likelihood of the association between measurement $\boldsymbol{z}_m$ and target state $\boldsymbol{x}_n$, $m = 0, 1, 2, \ldots, M_k$, $n = 1, 2, \ldots, N_{\mathrm{TGT}}$; N_{DT}, the number of old target states, under h, to be associated with measurements; N_{NT}, the number of new targets, under h, to be associated with a second subset of measurements disjoint from the first subset earlier; N_{FT}, the number of false targets, under h, to be associated with a third subset of measurements disjoint from the first two subsets earlier; (N_{DT}, N_{NT}, N_{FT}), the Reid configuration that partitions the measurements set among old targets, new targets and false alarms; $\alpha = \{(m_1, n_1), (m_2, n_2), \ldots, (m_{N_{\mathrm{DT}}}, n_{N_{\mathrm{DT}}})\}$, is under h, a set of feasible measurement-to-target assignment pairs; $\boldsymbol{L} = M_k \times (N_{\mathrm{TGT}} + 2M_k)$, the data association likelihood matrix; and $\boldsymbol{L}^* = M_k \times (N_{\mathrm{TGT}} + 2M_k)$, the linear assignment problem matrix of negative log-likelihoods.

Definitions: Let $\boldsymbol{M}$ be a real $m \times n$ matrix

$$\boldsymbol{M} = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \vdots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{bmatrix}$$

$$p = \min(m, n)$$

A feasible assignment for $\boldsymbol{M}$ is a set of matrix element indices

$$\alpha = \{(i_1, j_1), (i_2, j_2), \ldots, (i_p, j_p)\}$$

in which no two distinct members have the same row index or column index; that is, if (r, s) and (t, u) are distinct members of α, then

$$r \neq t \quad \text{and} \quad s \neq u$$

The linear assignment problem on $\boldsymbol{M}$ is to find a feasible assignment α^* such that

$$V(\alpha^*) = \sum_{(i,j) \in \alpha^*} a_{ij} \text{ is a minimum over all feasible } \alpha$$

The K-best ranked assignment problem on $\boldsymbol{M}$ is to find a sequence of feasible assignments: α^*, α_1, α_2, $\ldots$, α_{k-1}, such that

$$\begin{aligned} &V(\alpha^*) \leq V(\alpha_1) \leq V(\alpha_2) \leq \cdots \leq V(\alpha_{k-1}) \\ &\alpha \neq \alpha^* \quad \text{and} \quad \alpha \neq \alpha_1 \Longrightarrow V(\alpha) \geq V(\alpha_1) \\ &\alpha \neq \alpha^* \quad \text{and} \quad \alpha \neq \alpha_1 \quad \text{and} \quad \alpha \neq \alpha_2 \Longrightarrow V(\alpha) \\ &\quad \geq V(\alpha_1) \quad \text{and} \quad V(\alpha) \geq V(\alpha_2), \text{ and so on.} \end{aligned}$$

For a given scan $k-1$ cluster hypothesis, g, and a scan k data association hypothesis, h, given g, Reid's fundamental equation (16) can be written in our notation as

$$p_{\mathrm{h}}^k = \left\{ \frac{1}{C} \left[p_{\mathrm{D}}^{N_{\mathrm{DT}}} \prod_{(m,n) \in \alpha} l_{mn} \right] (1 - p_{\mathrm{D}})^{N_{\mathrm{TGT}} - N_{\mathrm{DT}}} \times (\beta_{\mathrm{NT}})^{N_{\mathrm{NT}}} (\beta_{\mathrm{FT}})^{N_{\mathrm{FT}}} \right\} p_{\mathrm{g}}^{k-1} \tag{1}$$

In (1) C is a normalising constant that is independent of g and h and $N_{\mathrm{DT}} + N_{\mathrm{NT}} + N_{\mathrm{FT}} = M_k$.

3 Derivation

3.1 The data association hypothesis encoding theorem

The set of all feasible assignments with non-vanishing product indexed over the assignment index pairs for the data association matrix, $\boldsymbol{L}$, defined below corresponds one-to-one with the set of all Reid's hypotheses

$$\boldsymbol{L} \equiv \begin{bmatrix} l_{11} & l_{21} & \cdots & l_{1N_{\mathrm{TGT}}} & \vdots & \\ l_{21} & l_{22} & \cdots & l_{2N_{\mathrm{TGT}}} & \vdots & \left(\frac{\beta_{\mathrm{NT}}(1-p_{\mathrm{D}})}{p_{\mathrm{D}}}\right)\boldsymbol{I} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \\ l_{M_k 1} & l_{M_k 2} & \cdots & l_{M_k N_{\mathrm{TGT}}} & \vdots & \\ \vdots & & & & & \\ \vdots & & \left(\frac{\beta_{\mathrm{FT}}(1-p_{\mathrm{D}})}{p_{\mathrm{D}}}\right)\boldsymbol{I} & & & \\ \vdots & & & & & \\ \vdots & & & & & \end{bmatrix} \equiv \{l_{mn}\}$$

where $m = 1, 2, \ldots, M_k$, $n = 1, 2, \ldots, N_{\mathrm{TGT}} + 2M_k$ and $\boldsymbol{I}$ = the $M_k \times M_k$ identity matrix.

Proof: Following Reid, suppose that, under h, N_{DT} of the M_k measurements are assigned to N_{DT}, old targets and these measurement-to-target assignments are indexed by the set of feasible assignment index pairs

$$\alpha = \{(m_1, n_1), (m_2, n_2), \ldots, (m_{N_{\mathrm{DT}}}, n_{N_{\mathrm{DT}}})\} \tag{3}$$

Further, suppose that, under h, N_{NT} of the M_k measurements are of new targets. Let the set of new target measurements be indexed by the set

$$\beta = \{m_{N_{\mathrm{DT}}+1}, m_{N_{\mathrm{DT}}+2}, \ldots, m_{N_{\mathrm{DT}}+N_{\mathrm{NT}}}\} \tag{4}$$

Finally, suppose that, under h, N_{FT} false targets are indexed by the set

$$\gamma = \{m_{N_{\mathrm{DT}}+N_{\mathrm{NT}}+1}, m_{N_{\mathrm{DT}}+N_{\mathrm{NT}}+2}, \ldots, m_{M_K}\} \tag{5}$$

Noting that a feasible assignment does not permit more than one measurement to be assigned to one target or more than one target to be assigned to one measurement, the set of index pairs

$$\tilde{\alpha} = \alpha \cup \{(r, r + N_{\mathrm{TGT}}) : r \in \beta\} \cup \{(s, s + N_{\mathrm{TGT}} + M_k) : s \in \gamma\} \tag{6}$$

is a feasible assignment for $\boldsymbol{L}$ which accounts for, under h, old targets, new targets and false targets. The product of the matrix elements indexed over $\tilde{\alpha}$ is

$$\prod_{(r,s) \in \tilde{\alpha}} l_{tu} = \left(\prod_{(m,n) \in \alpha} l_{mn} \right) (\beta_{\mathrm{NT}})^{N_{\mathrm{NT}}} (\beta_{\mathrm{FT}})^{N_{\mathrm{FT}}} \left(\frac{1 - p_{\mathrm{D}}}{p_{\mathrm{D}}} \right)^{M_k - N_{\mathrm{DT}}} \tag{7}$$

If both sides of (7) are multiplied by the constant, ρ, where

$$\rho = \frac{p_{\mathrm{D}}^{M_k} (1 - p_{\mathrm{D}})^{N_{\mathrm{TGT}} - M_k} p_{\mathrm{g}}^{k-1}}{C} \tag{8}$$

then

$$\rho \prod_{(r,s)\in\tilde{\alpha}} l_{rs} = p_{\mathrm{h}}^{k} \tag{9}$$

We note that the scaling constant, ρ, above depends only on the prior hypothesis g and is independent of posterior hypothesis h.

Thus, for each Reid data association hypothesis, h, there is a corresponding feasible assignment given by (6) for the data association matrix $\boldsymbol{L}$, which yields the identical allocation of measurements to old targets, new targets and false targets under h. Moreover, the product of elements of $\boldsymbol{L}$ indexed over the members of the assignment is proportional to Reid's positive probability for the data association hypothesis, h.

Conversely, for each feasible assignment set of index pairs, $\tilde{\alpha}$, such that

$$\tilde{\alpha} = \{(i_1, j_1), (i_2, j_2), \ldots, (i_{M_k}, j_{M_k})\}, \quad \prod_{(r,s)\in\tilde{\alpha}} l_{rs} > 0 \tag{10}$$

we can make the partition

$$\tilde{\alpha} = \alpha \cup \beta \cup \gamma \tag{11}$$

in which

$$\alpha = \{(r,s)\colon 1 \le r \le M_k,\ 1 \le s \le N_{\mathrm{TGT}}\} \tag{12}$$

$$\beta = \{(r,s)\colon 1 \le r \le M_k,\ N_{\mathrm{TGT}} + 1 \le s \le N_{\mathrm{TGT}} + M_k\} \tag{13}$$

$$\gamma = \{(r,s)\colon 1 \le r \le M_k,\ N_{\mathrm{TGT}} + M_k + 1 \le s \le N_{\mathrm{TGT}} + 2M_k\} \tag{14}$$

Setting

$$N_{\mathrm{DT}} = |\alpha| \tag{15}$$

$$N_{\mathrm{NT}} = |\beta| \tag{16}$$

$$N_{\mathrm{FT}} = |\gamma| \tag{17}$$

where $|\cdot|$ is the cardinality of the set argument, and by convention, $\prod_{(r,s)\in\phi} l_{rs} = 1$ if ϕ is the null set.

We then have

$$N_{\mathrm{DT}} + N_{\mathrm{NT}} + N_{\mathrm{FT}} = M_k \tag{18}$$

and

$$\prod_{(r,s)\in\tilde{\alpha}} l_{rs} = \left(\prod_{(r,s)\in\alpha} l_{rs}\right)\left(\prod_{(r,s)\in\beta} l_{rs}\right)\left(\prod_{(r,s)\in\gamma} l_{rs}\right) > 0 \tag{19}$$

We can rewrite (19) as

$$\begin{aligned}\prod_{(r,s)\in\tilde{\alpha}} l_{rs} &= \left(\prod_{(r,s)\in\alpha} l_{rs}\right)\left(\frac{\beta_{\mathrm{NT}}(1-p_{\mathrm{D}})}{p_{\mathrm{D}}}\right)^{N_{\mathrm{NT}}}\left(\frac{\beta_{\mathrm{FT}}(1-p_{\mathrm{D}})}{p_{\mathrm{D}}}\right)^{N_{\mathrm{FT}}} \\ &= \left(\prod_{(r,s)\in\alpha} l_{rs}\right)(\beta_{\mathrm{NT}})^{N_{\mathrm{NT}}}(\beta_{\mathrm{FT}})^{N_{\mathrm{FT}}}\left(\frac{1-p_{\mathrm{D}}}{p_{\mathrm{D}}}\right)^{M_k - N_{\mathrm{DT}}}\end{aligned} \tag{20}$$

Equation (20) above is just a restatement of (7). Therefore to the feasible assignment $\tilde{\alpha}$ with non-vanishing matrix element product, there corresponds a Reid configuration (N_{DT}, N_{NT}, N_{FT}), a data association hypothesis, h, which allocates measurements to old targets, new targets and false targets, and a non-vanishing matrix element product over $\tilde{\alpha}$ proportional to Reid's hypothesis probability. □

3.2 Counting the number of hypotheses

We now answer the question of just how many posterior feasible data association hypotheses, h, are there, given prior hypothesis g. For an $M_K \times (N_{\mathrm{TGT}} + 2M_K)$ likelihood assignment matrix, $\boldsymbol{L}$, and a given number of existing target assignments, N_{DT}, such that $0 \le N_{\mathrm{DT}} \le \min(M_k, N_{\mathrm{TGT}})$:

(i) for each such N_{DT} there are $\binom{N_{\mathrm{TGT}}}{N_{\mathrm{DT}}}$ ways to choose N_{DT} existing targets;

(ii) the first of these targets can be assigned to M_k measurement, the next to $M_k - 1$, the next to $M_k - 2$, and so on. Thus, there are

$$M_k(M_k - 1)\cdots(M_k - N_{DT} + 1) = \binom{M_k}{N_{\mathrm{DT}}} N_{DT}!$$

ways of assigning N_{DT} existing targets to M_k measurements.

(iii) For each of the remaining $M_k - N_{\mathrm{DT}}$ measurements there are just two ways to assign each, new target or false alarm. Thus there are $2^{M_k - N_{\mathrm{DT}}}$ such assignments.

From (i)–(iii), it follows that the total number of hypotheses N_{g}^{k}, given prior hypothesis g is given by

$$N_{\mathrm{g}}^{k} = \sum_{N_{\mathrm{DT}}=0}^{\min(M_k, N_{\mathrm{TGT}})} \binom{M_k}{N_{\mathrm{DT}}}\binom{N_{\mathrm{TGT}}}{N_{\mathrm{DT}}} N_{\mathrm{DT}}!\ 2^{M_k - N_{\mathrm{DT}}} \tag{21}$$

3.3 Linear assignment problem matrix $\boldsymbol{L}^*$

Now that the proof of equivalence of the set of Reid's hypotheses to the set of all feasible assignments with non-zero likelihood products for the likelihood association matrix has been established, it is easy to see that if we construct the matrix $\boldsymbol{L}^*$ with the positive non-zero elements of $\boldsymbol{L}$ replaced by their negative natural logs and matrix element zeros being replaced by a strict matrix element absolute value upper bound, ω, according to

$$\boldsymbol{L}^* = \begin{bmatrix} -\ln(l_{11}) & \cdots & -\ln(l_{1N_{\mathrm{TGT}}}) & -\ln(\lambda_{\mathrm{NT}}) & \omega & \cdots & \omega & \omega \\ \vdots & \vdots & \vdots & \omega & \ddots & \omega & \omega & \vdots \\ -\ln(l_{M_k 1}) & \cdots & -\ln(l_{M_k N_{\mathrm{TGT}}}) & \vdots & \omega & -\ln(\lambda_{\mathrm{FT}}) & \omega & -\ln(\lambda_{\mathrm{NT}}) \end{bmatrix} \tag{22}$$

where

$$\lambda_{\mathrm{NT}} = \frac{\beta_{\mathrm{NT}}(1-p_{\mathrm{D}})}{p_{\mathrm{D}}}, \qquad \lambda_{\mathrm{FT}} = \frac{\beta_{\mathrm{FT}}(1-p_{\mathrm{D}})}{p_{\mathrm{D}}}$$

and

$$\omega = \max\left\{\max_{i\le M_k,\ j\le M_k} -\ln(l_{ij}), -\ln(\lambda_{\mathrm{NT}}), -\ln(\lambda_{\mathrm{FT}})\right\} + 1 \tag{23}$$

We can apply the generalised rectangular matrix Murty algorithm to $\boldsymbol{L}^*$ and find the minimum, S, of the sum of

$\boldsymbol{L}^*$ elements taken over all feasible assignments

$$S = \min_{\mu} \sum_{t=1}^{M_k} -\ln(l_{t,\mu(t)}) \tag{24}$$

The set of all feasible assignments are just the one-to-one mappings

$$\mu\colon (1, 2, \ldots, M_k) \longrightarrow (1, 2, \ldots, N_{\mathrm{TGT}} + 2M_k) \tag{25}$$

So, finding the minimising μ, maximises the sum of the natural log likelihoods and thus maximises the product of the likelihood's. In other words, finding this μ produces the maximally likely assignment.

Sequential application of the ranked assignment algorithm gives the second most likely, the third, and so on, down to the K-best. Feasible assignments for $\boldsymbol{L}$ with zero likelihood products will be automatically excluded by the ranked assignment algorithm operating on $\boldsymbol{L}^*$ because their values are greater than any possible feasible assignment with non-vanishing likelihood product. We will call $\boldsymbol{L}^*$ the complete Reid assignment matrix.

What this means is that the 3.5857766×10^8 hypothetical data assignments computed from (21) for an $(M_k, N_{\mathrm{TGT}}) = (10, 10)$ problem do not each have to be algorithmically identified, evaluated, sorted and pruned. The data structure $\boldsymbol{L}^*$ encodes the complete set of such assignments and the generalised Murty ranked assignment algorithm does the identifying, sorting and pruning automatically and efficiently. Note that this encoding depends on $M_k N_{\mathrm{TGT}} + 2$ parameters. Most important, we observe that for each cluster and cluster hypothesis that the data structure $\boldsymbol{L}^*$ is bounded in its dimensions, computed anew on-the fly in each scan, and is only temporary. The temporary nature of $\boldsymbol{L}^*$ and its dimensional boundedness contrasts to the ever-growing hypothesis matrix data tree structure that is shown in Fig. 2 on p. 846 of [1].

3.4 Examples

Table 1 illustrates the rapid growth in the number of hypotheses as a function of problem dimensions, (M_k, N_{TGT}).

As a further more concrete illustration, we list the 14 feasible assignments with non-zero likelihood products for the case above with two measurements and two targets. The symbolic data association matrix is given by

$$\boldsymbol{L} = \begin{bmatrix} l_{11} & l_{12} & \lambda_{\mathrm{NT}} & 0 & \lambda_{\mathrm{FT}} & 0 \\ l_{21} & l_{22} & 0 & \lambda_{\mathrm{NT}} & 0 & \lambda_{\mathrm{FT}} \end{bmatrix}$$

The feasible assignments and their indexed products are shown in Table 2.

4 Results

We have used a high-fidelity geostationary space-based infrared sensor simulation to provide two test cases with different number of objects, object spatial densities and detection engagement time spans. By high-fidelity we mean: the multi-stage missile trajectory equations of motion incorporating the forces of thrust, gravity and drag, were accurately numerically integrated; the satellite orbital positions were computed according to classical Keplerian two-body dynamics; the earth's geoid was approximated by a standard geophysical ellipsoid; the rockets' hot exhaust plumes were characterised as point intensities-at-the source and attenuated by $1/r^2$ and atmospheric absorption effects; the scanning motion of the sensor was realistically simulated; and sensor-to-booster line-of-sight errors were modelled by aggregating focal plane location, sensor boresight and attitude estimation errors.

Table 1: Growth in number of hypotheses

N_{TGT} / M_k	2	3	4	5
2	14	22	32	44
3	44	86	52	248
4	128	304	648	1256
5	352	992	2512	5752

Table 2: Feasible hypotheses for a small example

{(1, 1), (2, 2)}, $l_{11}l_{22}$	{(1, 1), (2, 4)}, $l_{11}\gamma_{\mathrm{NT}}$
{(1, 1), (2, 6)}, $l_{11}\gamma_{\mathrm{FT}}$	{(1, 2), (2, 1)}, $l_{12}l_{21}$
{(1, 2), (2, 4)}, $l_{12}\gamma_{\mathrm{NT}}$	{(1, 2), (2, 6)}, $l_{12}\gamma_{\mathrm{FT}}$
{(1, 3), (2, 1)}, $l_{21}\gamma_{\mathrm{NT}}$	{(1, 3), (2, 2)}, $l_{22}\gamma_{\mathrm{NT}}$
{(1, 3), (2, 4)}, γ^2_{NT}	{(1, 3), (2, 6)}, $\gamma_{\mathrm{NT}}\gamma_{\mathrm{FT}}$
{(1, 5), (2, 1)}, $l_{21}\gamma_{\mathrm{FT}}$	{(1, 5), (2, 2)}, $l_{22}\gamma_{\mathrm{FT}}$
{(1, 5), (2, 4)}, $\gamma_{\mathrm{FT}}\gamma_{\mathrm{NT}}$	{(1, 5), (2, 6)}, γ^2_{FT}

Both test cases were scenarios from the Strategic Defense Initiatives' Global Protection Against Limited Strikes (GPALS). The first moderately stressing test case simulates the boost phase detection and tracking of 63 closely spaced intercontinental ballistic missiles over 25 scans covering 250 s. The second heavily stressing case simulates the detection and tracking of 200 closely spaced boosters over 35 scans covering 350 s. Both simulations were run on a 100 MHz, 128 Mbyte, Silicon Graphics Incorporated (SGI) Indigo II Work Station.

We note that, for this particular application and for general applications in which the probability of detection is close to one and the new target arrival rate is very much larger than the false alarm rate, we have used the device of deleting from $\boldsymbol{L}^*$ the sub-matrix of columns corresponding to false targets to speed-up the performance of the generalised Murty rectangular matrix algorithm. False targets are effectively dealt with using a consecutive missed detection track drop criterion. However, use of the complete Reid linear assignment problem matrix may well prove to be effective in the case of a degraded sensor with lower probability of detection or higher false alarm rate.

Three technical performance measures were used to evaluate the performance of our MHT algorithm specifically and are recommended for tracking algorithm performance assessment in general. These measures are: track assembly purity, track count accuracy and the ratio of actual wall clock scan processing time to simulated sensor scan time.

Track assembly purity measures how well the algorithm globally assembles multi-scan measurements originating from the same real object into the same persistent track. We have instrumented our simulation to do post-run evaluation of track assembly purity by encoding the identity of the actual measurement-generating object in each simulated measurement report. Our observation is that such encoding for post-processing track purity performance evaluation is not a widely followed practice in the tracking community.

Track count accuracy measures an algorithm's ability to keep objects in-track against the tendency to break track and form new ones.

The scan processing time ratio of the wall clock time to process a scan's measurement data to the actual scan time measures the algorithm's throughput performance. The algorithm is real-time if this ratio is uniformly less than 1 over all scans.

For both test cases, we have varied the number of retained hypotheses, K, to illuminate the tradeoff between track assembly purity and track count accuracy on one hand and throughput on the other.

The tracking was performed in a two-dimensional (x, y) line-of-sight rectangular coordinate system. State vector evolution was governed by second-order polynomial kinematics with a pair of decoupled states: $\{[x, \dot{x}, \ddot{x}]^{\mathrm{T}}, [y, \dot{y}, \ddot{y}]^{\mathrm{T}}\}$.

Table 3 shown below illustrates how the performance of our tracking algorithm, as diagrammed in Fig. 2, is post-processed and compiled into a standard performance evaluation report. It captures both the detailed history of scenario evolution and the tracking algorithm's performance over time. The table shows that the sensor did not detect any boosters until its third scan. On that scan, the sensor detected 59 boosting missiles (nd) and the algorithm initiated 59 single point clusters (nc), each with a single cluster hypothesis of unit probability, each hypothesis having a single point-initiated pair of (x, y) track state estimates and error covariance matrices. The number of detections grew to a maximum of 200 in scan 17 as newly launched missiles entered the sensor's field-of-regard (FOR) and then decreased as old tracked missiles began to leave it. The number of tracks (nt) maintained by the algorithm grew smoothly to 200 and then gradually decreased to 3 on the last sensor scan. After an initial transient period of about five scans the algorithm held the number of clusters in the range 12–28 until scan 25. After scan 25, closely spaced objects began to separate in the FOR and form single object clusters while, simultaneously, other objects were leaving the FOR. The maximum number of tracks associated with any cluster (lcsize) grew smoothly from 1 on scan 1 to 112 on scan 13 and was kept fixed from scans 13–25 and then gradually decreased, again a consequence of tracked objects exiting the FOR. This means that the generalised Murty algorithm was finding the 10 best hypotheses from linear assignment problem matrices whose dimensions were on the order of 112×224. The percentage of perfectly assembled tracks (perfect) ranged from a low of 74% on scan 5 to the mid 80s and higher from scan 20 to the last scan. This performance has significant positive implications for the difficult problems of missile typing in the boost phase and warhead discrimination in the mid-course phase of missile trajectories.

Table 3: MHT algorithm output for performance evaluation (200 objects, 10 hypotheses)

Scan	nd	nt	nc	lcsize	perfect	purity	time
3	59	59	59	1	59	1	0.28
4	91	91	9	35	81	0.95333	181.51
5	111	111	11	46	84	0.92337	264.82
6	123	124	10	51	92	0.90365	55.88
7	103	133	26	59	100	0.91992	53.86
8	145	147	18	67	109	0.93671	54.75
9	162	166	15	79	125	0.93829	96.33
10	175	177	17	89	138	0.95648	161.11
11	185	185	15	100	152	0.9712	148.29
12	193	193	14	108	156	0.97311	148.51
13	196	196	13	112	158	0.9745	149.28
14	197	197	13	112	161	0.97864	142.88
15	198	198	13	112	159	0.97824	132.68
16	199	199	12	112	164	0.983	103.29
17	200	200	12	112	162	0.98283	91.68
18	199	200	12	112	167	0.98616	92.88
19	184	200	21	112	164	0.98562	104.45
20	170	199	27	112	167	0.98815	111.65
21	169	184	28	112	157	0.98935	118.02
22	169	170	17	112	148	0.99131	112.26
23	155	169	15	112	146	0.99138	139.71
24	153	169	14	112	148	0.99247	92.16
25	151	155	17	112	134	0.99214	86.19
26	106	154	48	92	134	0.99272	38.95
27	51	151	92	50	133	0.99369	12.15
28	17	106	117	22	92	0.99166	2.87
29	15	51	80	22	43	0.98667	2.1
30	11	17	37	18	14	0.99381	1.29
31	10	15	11	16	13	0.99166	0.68
32	8	11	8	14	10	1	0.38
33	6	10	6	12	9	1	0.32
34	4	8	7	8	7	1	0.31
35	3	6	6	7	6	1	0.28
36	2	4	3	6	4	1	0.12
37	1	3	1	5	3	1	0.09

nd, number of detections; nt, number of tracks; nc, number of clusters; lcsize, largest cluster size; perfect, number of perfectly assembled tracks; purity, average percent of a track's measurements originating from same real object; time, wall clock time to process one scan's worth of data.

Figs. 3–8 illustrate the performance of our MHT algorithm.

Fig. 3 plots global track assembly purity as a function of time (scan number) traded against the number of retained hypotheses for the 63-object test case. We can see that the MHT algorithm ($K = 5$, 10) achieves and maintains 99% purity by scan 11, about a 2% performance gain over $K = 1$. This purity differential is crucial for accurate missile typing and discrimination.

Fig. 4 graphs sensor detections and objects-in-track by scan for the 63 object case. It should be noted that at no time did the number of tracks exceed the true track count maximum number; the algorithm neither broke tracks nor proliferated them. It accurately counted tracks until their individual measurement-generating targets left the sensor FOR. Importantly, track count accuracy was maintained perfectly for each of $K = 1$, 5, 10 retained hypotheses.

Fig. 5 indicates that our MHT algorithm operated essentially in real-time on an obsolete processor for the moderately stressing scenario of the 63 object test case. In the figure, we have plotted $\log_{10}$ of the ratio of wall clock time to scan time against scan number. In only the $K = 10$ case and only for one scan did the wall clock-to-scan time ratio exceed unity. This clearly would not happen with currently available processors that are approximately 30 times faster and have 2–8 times the memory of the SGI Indigo II workstation.

Figs. 6–8 are the performance analogues to Figs. 3–5 for the 200 object case. We can see that with currently available processors, having at least 40 times the throughput of the obsolete SGI machine, that the algorithm provides real-time high performance in terms of track purity and count accuracy with 10 retained hypotheses. We note again that this algorithm is a prototype, neither optimised nor

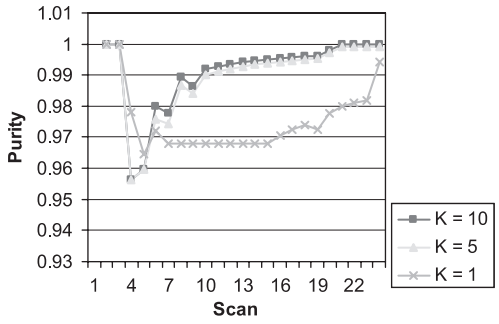

Fig. 3 *Track assembly purity against number of retained hypotheses (63 objects)*

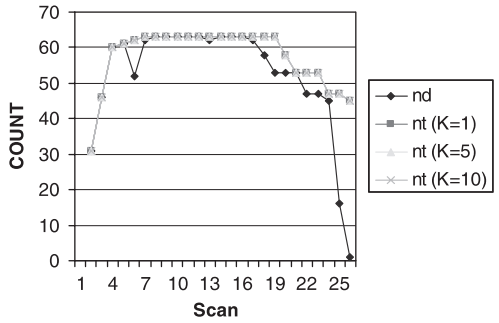

Fig. 4 *Track count against number of retained hypotheses (63 objects)*

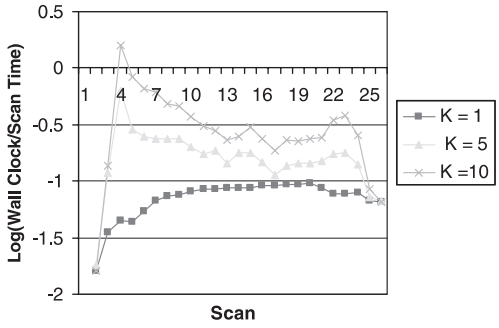

Fig. 5 *Throughput performance (63 objects)*

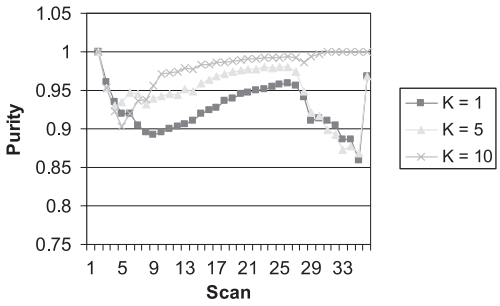

Fig. 6 *Track assembly purity against number of retained hypotheses (200 objects)*

streamlined. Code optimisation and streamlining might conservatively yield another factor of two in throughput performance.

5 Conclusions

We have shown that, for a given scan, cluster, cluster hypothesis and set of cluster hypothesis track state estimates

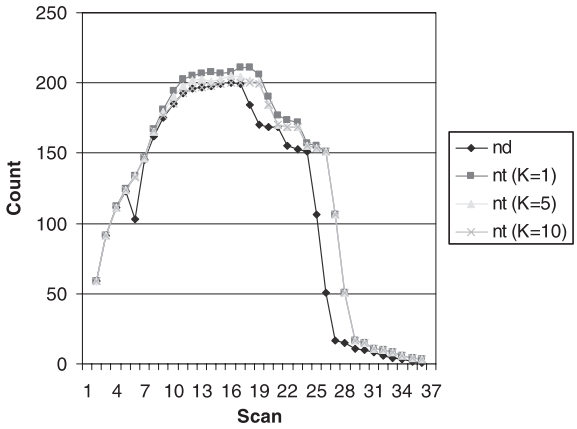

Fig. 7 *Track count against number of retained hypotheses (200 objects)*

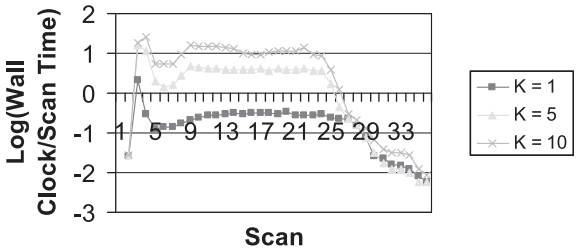

Fig. 8 *Throughput performance (200 objects)*

and corresponding state estimate error covariance matrices, the complete set of Reid's next scan measurement assignment hypotheses can be encoded in a rectangular negative log-likelihood assignment matrix. As in the original and complete Reid algorithm, each such next scan measurement hypothesis allocates the next scan's measurements among existing targets, new targets and false targets. Our algorithmic approach permits the immediate explicit ordering of next scan hypotheses by decreasing likelihood with a fast generalised rectangular matrix Murty ranked assignment algorithm. It does away with the NP-hard problem of explicit next scan hypotheses generation, evaluation, sorting, merging, branching and pruning that made the original algorithm impractical for large-scale real-time applications. It also dispenses with the need to maintain ever-growing hypotheses trees.

Combining the notions of the complete Reid assignment matrix, fast-ranked assignment algorithms, and the efficient gating of Collins and Uhlmann [10] will further enhance the performance of the algorithm for real-time applications. This means that previously intractable MHT problems with hundreds of closely spaced objects can be solved in real-time with today's high-performance gigahertz processing speed and gigabyte memory workstations. Supporting evidence for this assertion is provided by the MHT algorithm performance results in two high-fidelity infrared surveillance sensor system simulations, one of the moderate scale and one of the large scale. Although the current MHT algorithm does not yet incorporate the efficient gating method of Collins and Uhlmann, it nevertheless essentially achieved real-time high purity and track count accuracy in the first case of 63 closely spaced objects with a 10-year-old processor. Conservative extrapolation to currently available processor speed and memory indicates real-time performance in the case of 200 closely spaced targets. Moreover, this performance was attained with a modest number of the *K*-best hypotheses and a

prototype algorithm. Indeed, our results strongly imply the necessity of retaining 10 or more hypotheses for sufficiently large numbers of closely spaced objects, if high track purities and track count accuracies are to be attained in currently envisioned applications.

The military applications of this new 'algorithm technology' are in strategic and tactical ballistic missile defense, air defense and interceptor avionics. Commercial civil applications include air traffic control and airport runway ground traffic control.

6 References

1 Reid, D.B.: 'An algorithm for tracking multiple targets', *IEEE Trans. Autom. Control*, 1979, **AC-24**, pp. 843–854

2 Morefield, C.L.: 'Application of 0-1 integer programming for multitarget tracking problems', *IEEE Trans. Autom. Control*, 1977, **AC-22**, pp. 302–311

3 Nagarajan, V., Chidambara, M.R., and Sharma, R.N.: 'Combinatorial problems in multitarget tracking – a comprehensive solution', *IEE Proc. F, Commun. Radar Signal Process.*, 1987, **134**, (1), pp. 113–118

4 Danchick, R., and Newnam, G.E.: 'A fast method for finding the exact *N*-best hypotheses for multitarget tracking', *IEEE Trans. Aerosp. Electron. Syst.*, 1993, **29**, (2), pp. 555–560

5 Cox, I., and Miller, M.: 'On finding ranked assignments with application to multitarget tracking and motion correspondence', *IEEE Trans. Aerosp. Electron. Syst.*, 1995, **32**, (1), pp. 486–489

6 Murty, K.G.: 'An algorithm for ranking all the assignments in order of increasing cost', *Oper. Res.*, 1968, **16**, pp. 682–687

7 Miller, M., Stone, H., and Cox, I.: 'Optimizing Murty's ranked assignment method', *IEEE Trans. Aerosp. Electron. Syst.*, 1997, **33**, (3), pp. 851–862

8 Jonker, R., and Volgenant, A.: 'A shortest augmenting path algorithm for dense and sparse linear assignment problems', *Computing*, 1987, **38**, pp. 325–340

9 Cox, I., Miller, M., Danchick, R., and Newnam, G.E.: 'A comparison of two algorithms for determining ranked assignments with applications to multitarget tracking and motion correspondence', *IEEE Trans. Aerosp. Electron. Syst.*, 1997, **33**, (1), pp. 295–300

10 Collins, J.B., and Uhlmann, J.K.: 'Efficient gating in data association with multivariate distributed states', *IEEE Trans. Aerosp. Electron. Syst.*, 1992, **28**, (3), pp. 909–916