Multipath Track Association for Over-the-Horizon Radar using Lagrangian Relaxation

Mark G. Rutten[a,b], Simon Maskell[c], Mark Briers[c] and Neil J. Gordon[a]

[a]Intelligence Surveillance and Reconnaissance Division, Defence Science and Technology Organisation, PO Box 1500, Edinburgh, SA 5111, Australia

[b]Cooperative Research Centre for Sensor Signal and Information Processing, Department of Electrical and Electronic Engineering, The University of Melbourne, Parkville, Victoria 3010, Australia

[c]QinetiQ Ltd., Malvern Technology Centre, St Andrews Rd, Malvern, Worcestershire, UK

ABSTRACT

Over-the-horizon radar (OTHR) uses the refraction of high frequency radiation through the ionosphere in order to detect targets beyond the line-of-sight. The complexities of the ionosphere can produce multipath propagation, which may result in multiple resolved detections for a single target. When there are multipath detections, an OTHR tracker will produce several spatially separated tracks for each target. Information conveying the state of the ionosphere is required in order to determine the true location of the target and is available in the form of a set of possible propagation paths, and a transformation from measured coordinates into ground coordinates for each path. Since there is no a-priori information as to how many targets are in the surveillance region, or which propagation path gave rise to which track, there is a joint target and propagation path association ambiguity which must be resolved using the available track and ionospheric information.

The multipath track association problem has traditionally been solved using a multiple hypothesis technique, but a shortcoming of this method is that the number of possible association hypotheses increases exponentially with both the number of tracks and the number of possible propagation paths. This paper proposes an algorithm based on a combinatorial optimisation method to solve the multipath track association problem. The association is formulated as a two-dimensional assignment problem with additional constraints. The problem is then solved using Lagrangian relaxation, which is a technique familiar in the tracking literature for the multidimensional assignment problem arising in data association. It is argued that due to a fundamental property of relaxations convergence cannot be guaranteed for this problem. However, results show that a multipath track-to-track association algorithm based on Lagrangian relaxation, when compared with an exact algorithm, provides a large reduction in computational effort, without significantly degrading association accuracy.

Keywords: OTHR, multipath, track-to-track association, Lagrangian relaxation

1. INTRODUCTION

Due to the structure of the ionosphere, a single target may produce several resolvable detections at an OTHR receiver. When processed by a OTHR tracker[1] this will result in several tracks for each target. Since measurements of the target are made through the ionosphere, these tracks are not estimates of the target position on the surface of the earth (*ground coordinates*), rather they are in *radar coordinates*, related to the measured distance through the ionosphere. This relationship is shown diagrammatically in Figure 1. Information about the ionospheric propagation conditions[2] is required to map a particular track from radar coordinates into ground coordinates. This information is provided by the OTHR coordinate registration (CR) system. The CR data gives transformations from radar coordinates into ground coordinates and a probability of propagation for a set of possible propagation paths. While this information is necessary to map a track onto the ground, there is little *a priori* indication as to which propagation path produced which track. This gives rise to a track-to-propagation path association ambiguity. In addition, in a multi-target environment there is an ambiguity in the origin of a set of closely separated tracks; they could be due to the same target via separate propagation paths or they could be due to separate targets. This gives rise to a track-to-target association ambiguity. These associations must be made for each track to determine the number of targets and the location of these targets in ground coordinates.

 For correspondence please contact: Mark.Rutten@dsto.defence.gov.au.

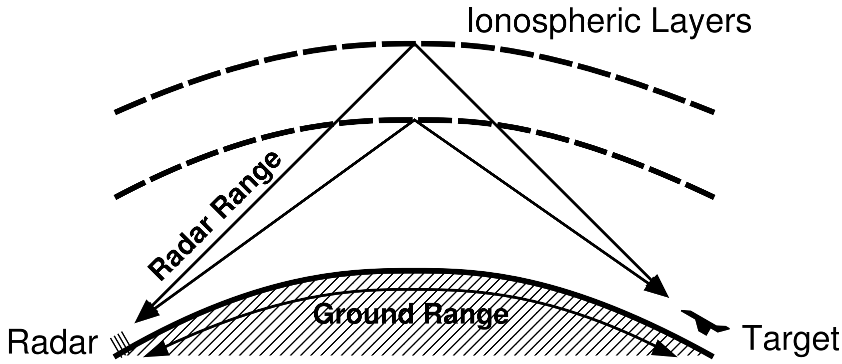

Figure 1. A simplified illustration of the OTHR propagation process, showing the relationship of the distances to the target in *radar coordinates* and *ground coordinates.*

Previous work in resolving the multipath track association ambiguity has used a multi-hypothesis approach,[3–5] evaluating the probability of every feasible combination of all of the tracks and possible propagation paths and extracting the highest probability hypothesis as the solution. A hypothesis tree is formed by recursively adding information about each track in turn. At the first level of the tree, possible targets from the first track form root nodes. At the second level of the tree, all combinations of track-to-target and track-to-propagation path associations are formed which include both the first and second tracks. The process continues until all tracks are incorporated. Approaches to reducing the computational complexity of this algorithm[4] have been based on a heuristic approach, keeping a fixed number of the highest probability hypotheses at each level in the tree before making the next set of branches. While this method will find a set of good, feasible hypotheses there is no guarantee of optimality or indication of the solution quality.

This work poses the association problem in a slightly different way. Rather than forming complete hypotheses containing all tracks, a set of single target hypotheses is formed. A target hypothesis is a set of tracks (with corresponding propagation paths) which could have originated from a single target. The multipath track association problem is then to combine *target hypotheses* together to form a complete hypothesis, aiming to maximise the combined likelihood, while ensuring that each track appears exactly once in the complete hypothesis. This problem can be posed as a 0/1 integer program. The structure of the problem can be used to reformulate the program as an two-dimensional assignment problem with extra constraints.[6] By relaxing the extra constraints into the objective function, using *Lagrangian multipliers*, the problem becomes one of iterative solution of a solvable problem (an assignment problem) and Lagrangian update. This process, known as *Lagrangian relaxation*,[7,8] is a common combinatorial optimisation technique for solving hard problems.[9] In particular this technique has become popular in the tracking literature for solving the multi-dimensional assignment problem.[10–12] Because of the complexity of the problem the resulting solution method is approximate and cannot guarantee an optimal solution. However, the nature of the algorithm can give an estimate of the accuracy of any solution.

For the purposes of this paper it is assumed that the algorithm has no prior information related to the track or ionospheric states. For an operational system this would be relevant in a case where the algorithm is receiving track information for the first time. Since all data is considered at the same point in time no variables have been given time indices. As a consequence of this simplification the tracks are assumed independent in the calculation of fused target states and target association hypotheses. It is recognised that this is a major simplifying assumption, with the inclusion of prior information part of ongoing research. In terms of demonstrating the applicability of the algorithm developed in this paper this assumption will have little or no effect on results and the conclusions drawn from them.

This paper is organised as follows: Section 2 outlines the generation of target hypotheses and the target likelihood calculation. Section 3 gives details of the association algorithm using an assignment algorithm with additional constraints, solved using the Lagrangian relaxation technique. A comparison between the efficiency and accuracy of the proposed algorithm against an exact solution for some simulated examples is given in Section 4, followed by a summary of the work and proposed further research in Section 5.

2. MULTIPATH TRACK ASSOCIATION

2.1. Available Information

The inputs to the multipath track association algorithm are a set of N tracks in radar coordinates with state $s_i \in \mathcal{R}$, $i = 1, \ldots, N$, supplied by the OTHR tracking system and an ionospheric state $w \in \mathcal{I}$, supplied by the OTHR CR system. The transformation, $T^{(m)} : \mathcal{R} \times \mathcal{I} \rightarrow \mathcal{G}$, also supplied by the CR system describes the mapping of the track state in radar coordinates, $\mathcal{R}$, into ground coordinates, $\mathcal{G}$, via propagation path $m \in \{1, \ldots, M\}$ through the ionosphere described by w. The state of track i in ground coordinates via path m is then given by

$$x_i^{(m)} = T^{(m)}(s_i, w). \tag{1}$$

It is assumed that the tracker output, s_i, is Gaussian distributed with mean $\hat{s}_i = \mathrm{E}\{s_i|Z_i\}$ and covariance $P_{s_i} = \mathrm{cov}(s_i|Z_i)$ and that the ionospheric state, w, has mean $\hat{w} = \mathrm{E}\{w|I\}$ and covariance $P_w = \mathrm{cov}(w|I)$. $Z_i \subset Z$ represents the set of all radar measurements associated with track i and $Z = \bigcup_{i=1}^{N} Z_i$ is the set of all associated radar measurements, where it is assumed that the measurement sets for each track do not intersect.[5] I represents the set of all ionospheric measurements. It is also assumed that the transformation from radar coordinates into ground coordinates can be closely approximated by a linearisation around the track and ionospheric state estimates $\hat{s}_i$ and $\hat{w}$, such that

$$x_i^{(m)} \approx T^{(m)}(\hat{s}_i, \hat{w}) + T_{s_i}^{(m)} \left[s_i - \hat{s}_i\right] + T_w^{(m)} \left[w - \hat{w}\right], \tag{2}$$

where

$$T_{s_i}^{(m)} = \left. \frac{\partial T^{(m)}(s_i, w)}{\partial s_i} \right|_{\substack{s_i = \hat{s}_i \\ w = \hat{w}}} \tag{3}$$

$$T_w^{(m)} = \left. \frac{\partial T^{(m)}(s_i, w)}{\partial w} \right|_{\substack{s_i = \hat{s}_i \\ w = \hat{w}}}, \tag{4}$$

are the Jacobian matrices of the transformation. Hence $x_i^{(m)}$ can also be assumed Gaussian distributed with mean and covariance

$$\hat{x}_i^{(m)} = T^{(m)}(\hat{s}_i, \hat{w}) \tag{5}$$

$$P_{x_i}^{(m)} = T_{s_i}^{(m)} P_{s_i} [T_{s_i}^{(m)}]^T + T_w^{(m)} P_w [T_w^{(m)}]^T. \tag{6}$$

The event that track i is associated with propagation path m is denoted by $\theta_i^{(m)}$. The prior probability of this event is supplied by the ionospheric estimation process, which is given by

$$\beta_i^{(m)} = P(\theta_i^{(m)}|I). \tag{7}$$

2.2. Track Associations

A target hypothesis t is defined by set of track-to-propagation path associations

$$\tau_t = \bigcup_{i \in \Phi_t} \theta_i^{(m_i^t)}, \tag{8}$$

where $\Phi_t \subset \{1, \ldots, N\}$ is the set of tracks associated with target hypothesis t and m_i^t is the propagation path assigned to track i under target hypothesis t. So, for example, a two-track target hypothesis containing track 1 via path 4 and track 5 via path 3 would be written as $\tau = \{\theta_1^{(4)}, \theta_5^{(3)}\}$. The physical restriction on path associations

Table 1. The total number of possible target hypotheses, N_τ, with N tracks and M propagation paths.

	M									
N	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
1	1	2	3	4	5	6	7	8	9	10
2	2	6	12	20	30	42	56	72	90	110
3	3	12	33	72	135	228	357	528	747	1020
4	4	20	72	208	500	1044	1960	3392	5508	8500
5	5	30	135	500	1545	4050	9275	19080	36045	63590
6	6	42	228	1044	4050	13326	37632	93288	2e+05	4e+05
7	7	56	357	1960	9275	37632	1e+05	4e+05	1e+06	3e+06
8	8	72	528	3392	19080	93288	4e+05	1e+06	5e+06	1e+07
9	9	90	747	5508	36045	2e+05	1e+06	5e+06	2e+07	6e+07
10	10	110	1020	8500	63590	4e+05	3e+06	1e+07	6e+07	2e+08

is that each target can have at most one track propagating along any one path. Hence the cardinality of the set Φ_t is

$$1 \leq |\Phi_t| \leq \min(M, N), \tag{9}$$

that is there must be at least one track in each target hypothesis and the maximum number of tracks in each target hypothesis is limited by the smaller of the total number of tracks, N, or the number propagation paths, M. Under this constraint, the total number of possible target hypotheses, N_τ can be calculated as

$$N_\tau = \sum_{i=1}^{\min(N,M)} \mathrm{C}_i^N \mathrm{P}_i^M, \tag{10}$$

where C_i^N denotes the combinations of i elements from N tracks and P_i^M denotes the permutations of i elements from M paths. A table of N_τ for a representative set of N and M is given in Table 1. Note the symmetry that arises, since $\mathrm{C}_i^N \mathrm{P}_i^M = \mathrm{P}_i^N \mathrm{C}_i^M$.

A collection of target hypotheses forms a complete hypothesis, λ_h, called a *path-dependent track-to-target association hypothesis* in the literature.[3] The restriction on a complete association hypothesis is that each track must appear exactly once. The complete hypothesis is defined by the set of target hypotheses, Λ_h, from which the hypothesis is formed, that is

$$\lambda_h = \bigcup_{t \in \Lambda_h} \tau_t. \tag{11}$$

2.3. Fused Target State

Under a given target association hypothesis, τ_t, the distribution of the fused target state can be calculated. The state estimates of the tracks associated with each target hypothesis can be fused to form a target state estimate. As discussed in the introduction, the track states are assumed independent and thus the fused target state, x_t, can be expanded as

$$p(x_t|\tau_t, Z, I) \propto \prod_{i \in \Phi_t} p(x_t|\theta_i^{(m_i^t)}, Z_i, I). \tag{12}$$

Each density, $p(x_t|\theta_i^{(m_i^t)}, Z_i, I)$, describes the target state from the set of radar measurements of track i mapped into ground coordinates via propagation path m_i^t, that is it corresponds to $x_i^{(m_i^t)}$ as defined by (5) and (6). Due to the assumptions of linearity and Gaussianity the fused target state can be simplified as Gaussian distributed with covariance and mean

$$P_t^{-1} = \sum_{i \in \Phi_t} \left[P_{x_i}^{(m_i^t)}\right]^{-1} \tag{13}$$

$$\hat{x}_t = P_t \sum_{i \in \Phi_t} \left[P_{x_i}^{(m_i^t)}\right]^{-1} \hat{x}_i^{(m_i^t)}. \tag{14}$$

Hence the target state estimate is a weighted sum of the state estimates of each of the associated tracks mapped into ground coordinates by the propagation paths specified through the target hypothesis' path associations, τ_t.

2.4. Target Hypothesis Likelihood

The probability of a complete path dependent association hypothesis can be expanded as follows

$$
\begin{aligned}
P(\lambda_h|Z,I) &= P(\tau_{t_1},\ldots,\tau_{t_n}|Z,I) &(15)\\
&\propto \int P(Z|x_{t_1},\ldots,x_{t_n},\tau_{t_1},\ldots,\tau_{t_n},I)P(x_{t_1},\ldots,x_{t_n},\tau_{t_1},\ldots,\tau_{t_n}|I)dx_{t_1},\ldots,dx_{t_n} &(16)\\
&= \prod_{t\in\Lambda_h}\int P(\cup_{i\in\Phi_t}Z_i|x_t,\tau_t,I)P(x_t,\tau_t|I)dx_t, &(17)
\end{aligned}
$$

where $\Lambda_h = \{t_1,\ldots,t_n\}$ is the set of all target hypotheses in the complete hypothesis h. Equation (17) shows that the complete hypothesis probability can be written as a product of terms related to each of the constituent target hypotheses. These terms are referred to as *target likelihoods*, which are denoted

$$\ell(\tau_t|Z,I) = \int P(\cup_{i\in\Phi_t}Z_i|x_t,\tau_t,I)P(x_t,\tau_t|I)dx_t. \quad (18)$$

The likelihood of a target hypothesis relates to how likely the set of tracks, mapped into ground coordinates by their specified propagation paths, are to have originated from a single target.

Equation (18) can be expanded in terms of known quantities. The first term in the integrand of (18) gives

$$
\begin{aligned}
P(\cup_{i\in\Phi_t}Z_i|x_t,\tau_t,I) &= \prod_{i\in\Phi_t}P(Z_i|x_t,\theta_i^{(m_i^t)},I) &(19)\\
&= \prod_{i\in\Phi_t}p(x_t|Z_i,\theta_i^{(m_i^t)},I)\frac{P(Z_i|\theta_i^{(m_i^t)},I)}{P(x_t|\theta_i^{(m_i^t)},I)}, &(20)
\end{aligned}
$$

in order to write the probability in terms of the track state densities in ground coordinates. The second term in the integrand of (18) gives

$$
\begin{aligned}
P(x_t,\tau_t|I) &= P(x_t|\tau_t,I)\prod_{i\in\Phi_t}P(\theta_i^{(m_i^t)}|I) &(21)\\
&= P(x_t|\tau_t,I)\prod_{i\in\Phi_t}\beta_i^{(m_i^t)}. &(22)
\end{aligned}
$$

Hence a target hypothesis likelihood can be written as

$$
\begin{aligned}
\ell(\tau_t|Z,I) &= \left[\prod_{i\in\Phi_t}\beta_i^{(m_i^t)}P(Z_i|\theta_i^{(m_i^t)},I)\right]\int P(x_t|\tau_t,I)\prod_{i\in\Phi_t}\frac{p(x_t|\theta_i^{(m_i^t)},Z_i,I)}{P(x_t|\theta_i^{(m_i^t)},I)}dx_t &(23)\\
&= \gamma^{-1}\prod_{i\in\Phi_t}\beta_i^{(m_i^t)}\int\prod_{i\in\Phi_t}p(x_t|\theta_i^{(m_i^t)},Z_i,I)dx_t, &(24)
\end{aligned}
$$

where γ^{-1} is a per-target constant related to the prior target density, which can be calculated from previously observed target densities in the surveillance region. Under linear Gaussian assumptions, the target hypothesis *log-likelihood*, $L(\tau_t|Z,I) = \log(\ell(\tau_t|Z,I))$, can be written as

$$
\begin{aligned}
L(\tau_t|Z,I) = -\log(\gamma) + \sum_{i\in\Phi_t}\log\left(\beta_i^{(m_i^t)}\right) + \frac{1}{2}\log\left(|2\pi P_t|\right) - \frac{1}{2}\sum_{i\in\Phi_t}\log\left(\left|2\pi P_{x_i}^{(m_i^t)}\right|\right)\\
-\frac{1}{2}\sum_{i\in\Phi_t}\left(\hat{x}_t-\hat{x}_i^{(m_i^t)}\right)^T\left[P_{x_i}^{(m_i^t)}\right]^{-1}\left(\hat{x}_t-\hat{x}_i^{(m_i^t)}\right), \quad (25)
\end{aligned}
$$

where $\hat{x}_t$ and P_t are the fused target state mean and covariance, given by (14) and (13) respectively and $\hat{x}_i^{(m_i^t)}$ and $P_{x_i}^{(m_i^t)}$ are the mean and covariance of track i mapped into ground coordinates by propagation path m_i^t, as given by (5) and (6) respectively.

3. A RELAXATION ALGORITHM FOR COMBINING TARGET HYPOTHESES

Having calculated the set of likelihoods for each target hypothesis and the corresponding fused target states, the aim is to find the complete hypothesis with the highest probability. The problem to solve is thus deciding which target hypotheses should be collected together to form an association which maximises the complete hypothesis probability, under the constraint that each track must appear exactly once in the collection of targets. The problem of finding the collection of target hypotheses which form the highest probability complete hypothesis can be written as a 0/1 integer program, where y_t are indicator variables (either 0 or 1), which define the presence of target t in the collection, and where $p_t = L(\tau_t|Z, I)$, the "profit" of target hypothesis t, is the target hypothesis log-likelihood.

$$\begin{aligned} &\max_y \left[\sum_{t=1}^{T} p_t y_t\right] \\ &\text{subject to:} \\ &\sum_{t=1}^{T} s_{tj} y_t = 1, \quad j = 1, \ldots, N, \\ &y_t \in \{0,1\},\ s_{tj} \in \{0,1\} \quad \forall t, j \end{aligned} \tag{26}$$

where s_{tj} is another indicator variable, which is 1 if target hypothesis t contains track j and 0 otherwise.

The structure of the problem can be exploited though rearranging the integer program (26) as an assignment problem. This is achieved by dividing every target association τ_t into two parts, the first track and the remaining tracks. This gives a set of possible single-track targets, τ_u^*, $u \in \{1, \ldots, U\}$, and a set of sub-targets $\tau_v^\dagger$, $v \in \{-1, 1 \ldots, V\}$ such that $\{\tau_u^*, \tau_v^\dagger\} = \tau_t$. To allow single track targets to appear in the association, the special index -1 is included in the set v, which corresponds to an empty set of sub-targets, $\tau_v^\dagger = \emptyset$. Using this indexing scheme the corresponding target profits and indicator variables can be written as $p_{uv} = p_t$, $y_{uv} = y_t$ and $s_{uvj} = s_{tj}$. The result of this reformulation is the following 0/1 integer program

$$\begin{aligned} &\max_y \left[\sum_{u=1}^{U}\sum_{v=1}^{V} p_{uv} y_{uv}\right] \\ &\text{subject to:} \\ &\sum_{u=0}^{U} y_{uv} = 1, \quad v = 1, \ldots, V \\ &\sum_{v=-1}^{V} y_{uv} = 1, \quad u = 1, \ldots, U \\ &\sum_{u=1}^{U}\sum_{v=1}^{V} s_{uvj} y_{uv} = 1, \quad j = 1, \ldots, N. \end{aligned} \tag{27}$$

Hence τ_u^* and $\tau_v^\dagger$ form the row and columns, respectively, of a two-dimensional assignment problem. Note that the summations over u and v include 0. A standard assignment problem would demand that every u be assigned to some v. The dummy index 0 ensures that as many u as necessary can be left unassigned. The problem is then in the form of an assignment problem (which can be easily solved using, for example, the auction algorithm[13]) with N extra constraints, which ensure that the complete hypothesis refers to each track exactly once.

In the context of the Lagrangian relaxation algorithm, program (27) is referred to as the *primal problem.* In general this is a hard problem to solve, so an approximation technique is used to simplify it. A *dual problem* can be formed by relaxing the track constraint into the objective function, using an unconstrained *Lagrangian multiplier* for each track, μ_j, $j = 1, \ldots, N$. The dual program can be written as

$$\begin{aligned} &\min_{\mu} \left\{ \max_{y} \left[\sum_{u=1}^{U} \sum_{v=1}^{V} p_{uv} y_{uv} - \sum_{j=1}^{N} \mu_j \left(\sum_{u=1}^{U} \sum_{v=1}^{V} s_{uvj} y_{uv} - 1 \right) \right] \right\} \\ &\text{subject to:} \\ &\sum_{u=0}^{U} y_{uv} = 1, \quad v = 1, \ldots, V \\ &\sum_{v=-1}^{V} y_{uv} = 1, \quad u = 1, \ldots, U. \end{aligned} \tag{28}$$

The term to be maximised in the objective function of the dual problem can be simplified

$$\sum_{u=1}^{U} \sum_{v=1}^{V} p_{uv} y_{uv} - \sum_{j=1}^{N} \mu_j \left(\sum_{u=1}^{U} \sum_{v=1}^{V} s_{uvj} y_{uv} - 1 \right) = \sum_{u=1}^{U} \sum_{v=1}^{V} \left(p_{uv} - \sum_{j=1}^{N} \mu_j s_{uvj} \right) y_{uv} + \sum_{j=1}^{N} \mu_j. \tag{29}$$

The dual problem thus becomes

$$\begin{aligned} &\min_{\mu} \left\{ \max_{y} \left[\sum_{u=1}^{U} \sum_{v=1}^{V} \left(p_{uv} - \sum_{j=1}^{N} \mu_j s_{uvj} \right) y_{uv} \right] + \sum_{j=1}^{N} \mu_j \right\} \\ &\text{subject to:} \\ &\sum_{u=0}^{U} y_{uv} = 1, \quad v = 1, \ldots, V \\ &\sum_{v=-1}^{V} y_{uv} = 1, \quad u = 1, \ldots, U. \end{aligned} \tag{30}$$

The maximisation over y is a generalised assignment problem, where the profit associated with a target t has been reduced by the sum of the Lagrangians of each track in the target. The minimisation over μ is a non-smooth optimisation problem. Several approaches to solving this problem by iteratively updating the Lagrangians, have appeared in the multi-dimensional assignment literature,[10–12] where the straightforward subgradient method of Blackman and Popoli[11] has been used for the algorithm presented in this paper.

The target collection algorithm proceeds by attempting to solve the dual problem. All N Lagrangians are initialised to a constant and the two-dimensional assignment problem is solved for the modified profits. However, the solution to the two-dimensional assignment may not necessarily satisfy the track constraint. Hence a feasible solution is constructed, which satisfies all of the primal constraints and is in some sense close to the dual solution. Those assigned targets containing tracks that violate the constraint are passed to an algorithm which picks the set of targets that meet the track constraint and maximise the combined profit. For the purposes of this paper a branch-and-bound algorithm finds the optimal set of violating targets that satisfy the constraint.

The profit of a solution to the dual problem is guaranteed to be greater than or equal to, that is an upper bound for, the profit of any primal problem, with equality at the optimal solution. This gives two stopping conditions for the optimisation algorithm. The difference between the dual profit for some set of multipliers μ and the primal profit for some feasible y is known as the *duality gap* and gives an indication of the quality of the primal solution. Therefore as the dual solution converges the duality gap should shrink, so that when the gap falls below some threshold the algorithm can terminate. Furthermore, since the relaxed constraints are equality constraints, the complementary slackness property[8] holds automatically and therefore if the dual solution, y, for multipliers μ is feasible, then y is the optimal solution. Thus if the dual solution is feasible the algorithm can terminate.

If the problem is well-behaved, then an optimal solution to (30) will be an optimal solution to the primal problem (26). However, this problem has the *integrality property*.[7] The integrality property describes a class of problems where the solution to the Lagrangian relaxed problem always has integer solutions even if the integer constraint ($y \in \{0, 1\}$) is removed. This is the case for (30), since it is well known that the integer constraint can be dropped when solving assignment problems.[14] The consequence of the integrality property is that the solution to the dual problem can be no better than the solution to the *linear programming relaxation* (LP relaxation). The LP relaxation is the primal problem with the integer constraint on y removed. This property divides the set of possible problems into "easy" problems and "hard" problems. If the solution to the LP relaxation is feasible, then the solution is optimal and the Lagrangian relaxation will also converge to the optimal solution, making the problem "easy". However, if the LP relaxation is not feasible, then the Lagrangian relaxed problem will converge to the solution of the LP relaxed problem, rather than the optimal. This makes the problem "hard," since neither of the stopping conditions described above will occur and there is no way of knowing how close the primal solution is to the optimal. In addition, since the LP relaxation is not trivial to calculate, it cannot be known whether the problem is "easy" or "hard" before attempting to find a solution.

Despite having the integrality property there are several reasons for persevering with a solution using Lagrangian relaxation. The main reason is that, while the dual problem may not converge to the optimal solution, reconstruction of a primal solution from an Lagrangian dual is much more straightforward than reconstructing from a LP relaxed solution. Secondly, a few iterations of the Lagrangian relaxation algorithm will, in general, be quicker than attempting to solve the LP relaxation.

In summary, the target collection algorithm using Lagrangian relaxation consists of the following steps:

1. Populate a profit matrix, P_1, with an entry for each target.
2. Initialise a Lagrangian multiplier to a constant for each track.
3. Form a modified profit matrix, P_2, by subtracting the Lagrangian multiplier for each track from the appropriate target profits.
4. Solve the two-dimensional assignment problem, using P_2 as the profit matrix.
5. Reconstruct a feasible solution from the resulting assignment, ensuring that each track is referred to exactly once.
6. If the dual solution is feasible then the solution is optimal and the algorithm terminates. If the duality gap is below some threshold, then the algorithm also terminates.
7. Otherwise, the Lagrangians are updated and the algorithm returns to step 3.

Note that since convergence of the algorithm cannot be guaranteed, in practice the algorithm runs for a fixed number of iterations and if convergence has not been reached, the final solution is the best feasible solution found.

4. SIMULATION AND RESULTS

For the purposes of testing the algorithm, the target hypothesis likelihoods were generated randomly. Propagation paths were also randomly selected for each track in each target hypothesis, while ensuring that all tracks were assigned different paths. One example of each type of target hypothesis is generated, where the type of target hypothesis refers to the set of associated tracks. Since the algorithm is attempting to find the complete hypothesis with maximum likelihood, there is no need to consider more than one target hypothesis of the same type, because the highest probability hypothesis will only ever contain the highest profit target from among those with the same type. This implies that

$$\sum_{i=1}^{\min(N,M)} \mathrm{C}_i^N \tag{31}$$

Table 2. Comparison of the LR algorithm versus an exact algorithm for "easy" problems.

	Time		Quality (%)			
Tracks	B&B	LR	True	Estimated	Conv. (%)	Iterations
4	1	1	100	99.73	100	6.7
6	6	2.6	100	98.86	80	12.4
8	41.5	5.7	100	97.66	60	16.5
10	758	13.3	99.44	97.27	60	19.7
12	11892	25.3	98.54	96.95	40	21.7
14	–	75.1	–	79.30	10	24.9
16	–	139	–	86.15	20	24.0

Table 3. Comparison of the LR algorithm versus an exact algorithm for "hard" problems.

	Time		Quality (%)			
Tracks	B&B	LR	True	Estimated	Conv. (%)	Iterations
4	1.5	3.4	100	67.75	0	25
6	6.6	4.9	91.21	58.14	0	25
8	43.2	8.5	88.53	64.90	0	25
10	755	17.3	94.22	78.44	0	25
12	12098	31.7	89.43	72.99	0	25
14	–	68.1	–	75.09	0	25
16	–	148	–	66.79	0	25

target hypotheses are simulated, rather than the much larger number given by (10).

A set of examples was generated with a varying number of tracks and with the number of propagation paths fixed at 4. Random examples were generated for each number of tracks and then sorted into "easy" problems, where the LP relaxation gave an integer solution, and "hard" problems. Ten examples of each were taken and the results averaged. The results are compared with a naïve branch-and-bound algorithm, designed to find the optimal solution. Tables 2 and 3 list the results obtained for the two algorithms in "easy" and "hard" cases respectively. Since each algorithm is programmed in MATLAB, the running times are normalised to the fastest average time (the 4-track "easy" examples), to reduce the effect of any coding quality differences between the algorithms.* The solution quality is simply the ratio of the best profit given by the relaxation algorithm to the profit of the true solution, where the solution profits are the products of the target likelihoods, not the sum of the log-likelihoods. The estimated quality is the ratio of the best profit given by the relaxation algorithm to the minimum dual profit. The percentage of examples where the algorithm converged and the average number of iterations of the relaxation algorithm is also given. Note that the true solution was not calculated for more that 12 tracks because of the computation time required.

It can be seen from Tables 2 and 3 that the computation times for both the exact algorithm and the Lagrangian relaxation algorithm increase exponentially, but that the times required for the relaxation algorithm increase a much slower rate. An important feature is that, especially for the "easy" problems, the benefit of increased speed does not degrade the accuracy of the solution significantly. Comparing the tables for the "easy" and "hard" problems shows that, as expected, the relaxation algorithm converges consistently for the "easy" problems and never converges to the optimal solution for the "hard" problems. Since the algorithm never meets its stopping condition, it always runs for the maximum number of iterations, set to 25 in these simulations. It is also apparent that the duality gap, which is related to the estimated quality shown in the table, is good for the "easy" set of

*The average time taken for the 4 track "easy" examples for branch-and-bound was measured at 15 mS and for the relaxation algorithm 77 mS, giving an average computation time for the 12 track "easy" case of 178.4 S for branch-and-bound and 1.95 S for relaxation.

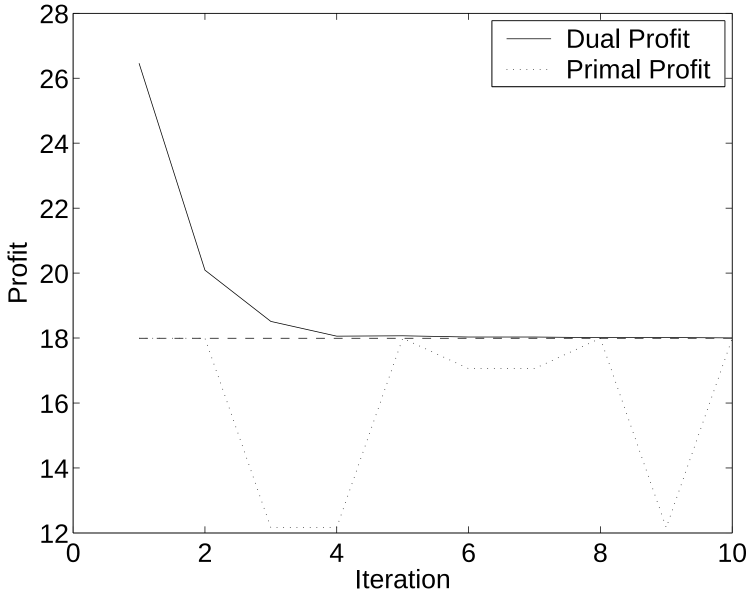

Figure 2. An example of the convergence of the primal and dual profits for an "easy" problem.

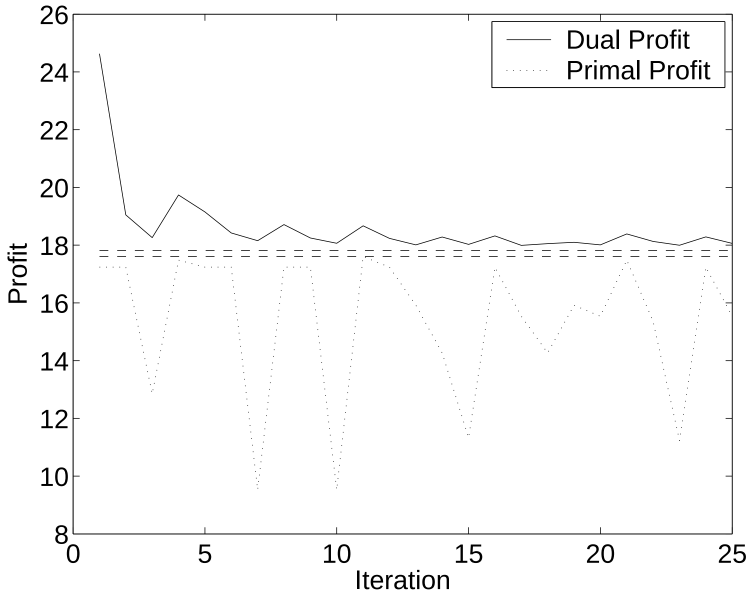

Figure 3. An example of the convergence of the primal and dual profits for a "hard" problem, showing separation between the optimal and LP relaxation profits.

problems, but is poor for the "hard" problems, since the dual is converging to the LP relaxed profit, rather than the optimal. Despite the lack of convergence, the solution found to the "hard" problems is often within 90% of the optimal.

Figures 2 and 3 show the iterative convergence of the primal and dual profits for an "easy" and a "hard" example respectively, with 8 tracks and 4 propagation paths. Note that with the Lagrangian update and primal reconstruction scheme used here neither the primal nor the dual will necessarily monotonically approach the optimal profit. The horizontal line in the easy example shows the optimal profit (which is also the LP relaxation profit in this case), while Figure 3 shows the optimal profit as the lower horizontal line and the LP relaxation

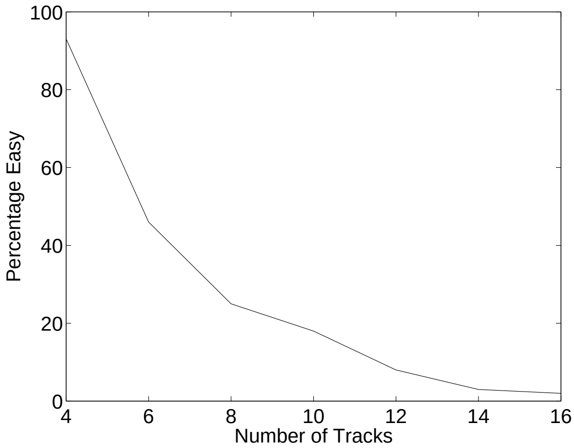

Figure 4. The percentage of "easy" problems as a function of the number of tracks.

profit as the upper line. It is obvious from this example that the dual profit does not approach the optimal profit, but rather approaches that of the LP relaxation. Figure 4 shows the percentage of easy problems from a random selection of 100 problems for different numbers of tracks. The figure shows that unfortunately the proportion of "easy" problems decreases with the number of tracks, and hence the difficulty of the problems increases. Therefore when using the Lagrangian relaxation algorithm there is less knowledge about how good a prospective solution is as the number of tracks increases.

The bulk of the computation time in the Lagrangian relaxation algorithm is spent reconstructing a feasible solution in the first few iterations. During the first few iterations, where the duality gap is quite large, the dual solution has many assignments that violate the track constraint, which requires a substantial effort from the algorithm attempting to find a feasible solution. As the duality gap reduces, the dual solution becomes better and less work is required to reconstruct a feasible solution. The increase in computation time as the number of tracks increases is also related to the time taken for the solution to the two-dimensional assignment. As the number of targets increases, the size of the profit matrix increases, which increases the time required to solve the two-dimensional assignment accordingly.

Since the simulation presented in this paper randomly selects association profits, the calculation of target likelihoods, as outlined in Section 2 is not simulated. Hence the time required to compute the likelihoods has not been taken into account in these results. Given the promising performance of the Lagrangian relaxation algorithm for solving the target collection problem it could be expected that the calculation of target likelihoods will take more time than finding a good association hypothesis.

5. CONCLUSIONS AND FURTHER RESEARCH

This paper has shown that the OTHR multipath track association problem can be reformulated as a two-dimensional assignment problem with an extra constraint for each track. The resulting association algorithm uses a combinatorial optimisation technique known as Lagrangian relaxation.[8] Due to the integrality property[7] of this problem, convergence of the algorithm cannot be guaranteed in general. However experimental results on simulated data show that the algorithm can provide a solution orders of magnitude faster than a guaranteed optimal solution, with little loss in solution quality. It was found that in many of the cases considered, the solution given by the Lagrangian relaxation algorithm was in fact optimal.

Further research is planned to extend the technique to search for several of the highest probability hypotheses.[15] It is also planned to modify the algorithm to restrict solutions to those having a specified number of targets. For example the algorithm would select the highest probability hypothesis with 5 targets in it. This is a much more difficult problem, since in the current problem a feasible solution is guaranteed, due to the inclusion of the single track targets, whereas there may be no feasible solution for the problem of finding a solution with a fixed number of targets.

ACKNOWLEDGEMENTS

Mark Rutten would like to acknowledge the useful discussions with Professors Bill Moran (Department of EEE, The University of Melbourne) and Rob Evans (CSSIP and the Department of EEE, The University of Melbourne).

REFERENCES

1. S. B. Colegrove and S. J. Davey, "PDAF with multiple clutter regions and target models," *IEEE Transactions on Aerospace and Electronic Systems* **39**, pp. 110–124, Jan. 2003.
2. K. Davies, *Ionospheric Radio*, Peter Peregrinus Ltd., London, UK, 1990.
3. D. J. Percival and K. A. B. White, "Multipath track fusion for over-the-horizon radar," in *Proceedings of the SPIE Conference on Signal and Data Processing of Small Targets 1997*, O. E. Drummond, ed., *Proceedings of the SPIE* **3163**, pp. 363–374, (San Diego, CA, USA), July 1997.
4. P. W. Sarunic, M. G. Rutten, D. J. Percival, and K. A. B. White, "Hypothesis management for over-the-horizon radar multipath track fusion," in *EuroFusion99: Proceedings of the International Conference on Data Fusion*, pp. 217–224, (Stratford-upon-Avon, UK), Oct. 1999.
5. M. G. Rutten and D. J. Percival, "Joint ionospheric and target state estimation for multipath OTHR track fusion," in *Proceedings of the SPIE Conference on Signal and Data Processing of Small Targets 2001*, O. E. Drummond, ed., *Proceedings of the SPIE* **4473**, pp. 118–129, (San Diego, CA, USA), Aug. 2001.
6. M. Briers, S. Maskell, and M. Philpott, "Two-dimensional assignment with merged measurements using Langrangrian relaxation," in *Proceedings of the SPIE Conference on Signal and Data Processing of Small Targets 2003*, O. E. Drummond, ed., *SPIE* **5204**, pp. 283–292, (San Diego, CA, USA), Dec. 2003.
7. A. M. Geoffrion, "Lagrangean relaxation for integer programming," *Mathematical Programming Study* **2**, pp. 82–114, 1974.
8. M. Guignard, "Lagrangean relaxation," *Top* **11**(2), pp. 151–228, 2003.
9. M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, 1979.
10. A. B. Poore and A. J. Robertson, "A new Lagrangian relaxation based algorithm for a class of multidimensional assignment problems," *Computational Optimization and Applications* **8**, pp. 129–150, 1997.
11. S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*, Artech House, Norwood, MA, 1999.
12. K. R. Pattipati, R. L. Popp, and T. Kirubarajan, "Survey of assignment techniques for multitarget tracking," in *Multitarget-Multisensor Tracking: Applications and Advances, Volume III*, Y. Bar-Shalom and W. D. Blair, eds., ch. 2, pp. 77–159, Artech House, Norwood, MA, 2000.
13. D. P. Bertsekas, "Auction algorithms for network flow problems: A tutorial introduction," *Computational Optimization and Applications* **1**, pp. 7–66, 1992.
14. D. Blumenfeld, *Operations Research Calculations Handbook*, CRC Press, 2001.
15. I. J. Cox, M. L. Milller, R. Danchick, and G. E. Newnam, "A comparison of two algorithms for determining ranked assignments with application to multitarget tracking and motion correspondence," *IEEE Transactions on Aerospace and Electronic Systems* **33**, pp. 295–301, Jan. 1997.