

Conditions for MHT to be an Exact Bayesian Solution to the Multiple Target Tracking Problem for Target-to-Measurement Association Hypotheses

LAWRENCE D. STONE

This paper finds conditions under which multiple hypothesis tracking (MHT) is an exact Bayesian solution to the multiple target tracking problem for target-to-measurement association hypotheses. The crucial condition is that measurements arrive in scans from one or more sensors, but otherwise the conditions are minimally restrictive. In order to produce a computationally feasible implementation of MHT, some approximations must be made, but this true is for any (existing) method of producing an exact Bayesian solution. Limiting the number of hypotheses considered is an example of such an approximation. This paper is motivated by recent claims that MHT is not theoretically rigorous or “Bayes optimal.”

Manuscript received January 5, 2019; revised November 13, 2019; released for publication 0, 00.

Refereeing of this contribution was handled by Chee-Yee Chong.

The author is with Metron Inc., 1818 Library Street, Reston, VA 20190, USA. E-mail: stone@metsci.com

1557-6418/19/\$17.00 © 2019 JAIF

I. INTRODUCTION

This paper, which is based on [1], considers the question of when multiple hypothesis tracking (MHT) is an exact Bayesian solution to the multiple target tracking problem for target-to-measurement association hypotheses, or, more succinctly, when is it exact Bayesian. By exact, we mean that MHT produces the correct Bayesian posterior distribution on the targets and their states. Recently, there have been claims that MHT is not theoretically rigorous or “Bayes optimal” (see [2, Sec. 10.7.2] and [3]). Ref. [4, Sec. VII] observes that if the random finite set (RFS) version of multiple target tracking is exact Bayesian, then so are certain special cases of MHT that can be derived from the RFS formulation. However, this begs the question as to whether the RFS formulation is exact Bayesian and whether more general versions of MHT are exact Bayesian.

Note, a target-to-measurement association hypothesis is different from a measurement-to-measurement association hypothesis that is used in the standard MHT formulation.

In this paper, we explore the question of when MHT is exact Bayesian by first stepping back a bit and considering a more general definition of MHT than is generally used (see, e.g., [4], [5], and [6]). The plan of this paper is to proceed from general versions of MHT to the specific until we arrive at the most commonly used notions of MHT. This approach has two virtues. First, it shows that the notion and validity of an MHT decomposition (defined below) is more general than the usual notion of MHT. Second, it highlights the special assumptions needed to produce the most common and useful forms of MHT.

In classical multiple target tracking, the problem is divided into two steps: association and estimation. Step 1 associates measurements with targets. Step 2 uses the measurements associated with each target to produce an estimate of that target’s state. Complications arise when there is more than one reasonable way to associate measurements with targets. MHT approaches this problem by forming association hypotheses to explain the source of the measurements. We consider the situation where each hypothesis assigns the measurements to targets or false measurements. For each association hypothesis, MHT computes the probability that the hypothesis is correct and the conditional probability distribution on the joint target state given the hypothesis is correct. The Bayesian posterior is a mixture of the conditional joint target state distributions weighted by the association probabilities. This is the *MHT decomposition* of the multiple target tracking problem.

Theoretically, there are other decompositions that could be used. For example, one could use any set of mutually exclusive and exhaustive conditions for the decomposition. What makes the MHT decomposition special and important is that it is *useful*. Each element (hypothesis) of an MHT decomposition specifies which

measurements are associated with which targets and which are associated with false measurements. Under a hypothesis, the multiple target tracking problem becomes a much more tractable problem. Usually it becomes a set of n single target tracking problems, where n is the number of targets specified by the hypothesis. The MHT decomposition transforms a difficult and daunting multiple target tracking problem into a set of problems we know how to solve. This was Reid’s key insight [5].

While MHT is the most widely used method for solving multiple target tracking problems, multiple target tracking is not limited to the classical case described above. Section II-C includes a brief discussion of multiple target tracking when the notion of associating measurements with targets is not meaningful.

Reid [5] formulated the initial version of MHT that was later generalized by Mori et al. [6]. Since then, many versions and implementations of MHT have been developed (see [4]).

Many of the technical results presented in this paper are based on results from [7, Ch. 4]. However, the emphasis in this paper is on identifying conditions under which MHT is an exact Bayesian solution to multiple target tracking.

We show that the crucial condition ensuring that MHT is exact Bayesian is that measurements arrive in scans as defined below. The additional conditions required for this result are minimally restrictive. Thus, the MHT decomposition is exact Bayesian for a wide class of tracking problems. In order to produce a computationally feasible implementation, some approximations must be made, but this is true of any (existing) method of producing an exact Bayesian solution. In MHT, the number of association hypotheses grows exponentially in the number of measurements, so a typical approximation is to limit the number of hypotheses considered. In addition, it is customary to display only the tracks resulting from the highest probability association hypothesis and treat them as the “tracking solution.”

Section II provides the basic definitions that we use for multiple target tracking. Section III proves the basic result on the validity of the MHT decomposition. The reader will note in this section that we use a more general definition of MHT than is usual. In particular, the various versions of MHT discussed in [4] are all special cases of this definition. Section IV presents additional assumptions that allow the MHT decomposition to be performed recursively, and Section V gives assumptions under which the target state distributions, conditioned on an association hypothesis, are independent. Section VI provides a summary of these assumptions. Section VII provides some conclusions.

II. MULTIPLE TARGET TRACKING

We employ a continuous–discrete formulation of tracking where the target motion takes place in continuous time, but the measurements are received at a dis-

crete sequence $0 \leq t_1 \leq \dots \leq t_K$ of possibly random times. We represent a single target’s state and its motion through the target state space S in terms of a stochastic process $\{X(t); t \geq 0\}$, where $X(t)$ is the target state at time t . The target state can have both continuous and discrete components. In addition to kinematic components, there can be components that correspond to “features” such as color or frequency and source level of an emission. Target motion can include changes in nonkinematic as well as kinematic components.

A. Multiple Target Motion Process

The multiple target tracking problem begins at $t = 0$. The total number of targets is unknown but bounded by \bar{N} , which is known. We assume a known bound on the number of targets because it allows us to simplify the presentation and produces no restriction in practice. It is possible to remove this restriction but that would add complications without adding capability. We add an additional state ϕ to the target state space S . If a target is not present in S , we say that it is in state ϕ . Let $S^+ = S \cup \{\phi\}$ be the augmented state space for a single target and $\mathbf{S}^+ = S^+ \times \dots \times S^+$ be the joint target state space where the product is taken \bar{N} times. This is a vector formulation of the multiple target tracking motion model. Each component (target) can be indistinguishable from the others, or if there is prior knowledge some components can have different motion models. Both are possible but neither is required. In the case where the targets are indistinguishable, the component labels are arbitrary. The notion of including a state such as ϕ to represent target not present in S has precedent in the works of [8], [9], [10], [11], and [12].

Our prior knowledge about the targets and their “movements” through the state space \mathbf{S}^+ is given by a stochastic process $\mathbf{X} = \{\mathbf{X}(t); t \geq 0\}$, where $\mathbf{X}(t) = (X_1(t), \dots, X_{\bar{N}}(t))$ is the state of the system at time t and $X_n(t) \in S^+$ is the state of target n at time t . The term “state of the system” means the joint state of all the targets. If $X_n(t) = \phi$, then target n is not present in S at time t . The motion model can allow for targets to arrive (transition from ϕ to S) and depart (transition from S to ϕ) as time progresses.

B. Multiple Target Likelihood Functions

Definition. A *measurement* is a function of a sensor response.

A sensor response that has crossed a specified threshold and is used to provide an estimate of a target’s position is an example of a measurement. A measurement can be a multivariate function of the sensor response. An example is a peak-picking algorithm that identifies the number of peaks that cross a threshold and their locations. Another example of a measurement is the sensor response itself. This is the identity function applied

to the sensor response to yield the measurement. An example of this is the acoustic times series received at a hydrophone over an interval of time.

Let the random variable $Y(t, j)$ be the measurement from sensor j at time t . Measurements from sensor j take values in the measurement space Ψ_j that may be different for each sensor. We define the *multiple target likelihood function* l_j for sensor j at time t as follows:

$$l_j(t, y|\mathbf{s}) = \Pr\{Y(t, j) = y | \mathbf{X}(t) = \mathbf{s}\} \text{ for } y \in \Psi_j, \mathbf{s} \in \mathbf{S}^+. \quad (1)$$

Note that this likelihood depends on the *system state* \mathbf{s} at time t . The system state gives the state of each target. If a target is not present, its state is ϕ . If false measurements are possible, then a model for these must be defined and used in the calculation of (1). Note that we use \Pr to mean probability or probability density, whichever is appropriate.

Suppose we have obtained measurements at the discrete times $0 \leq t_1 \leq \dots \leq t_K \leq t$. Let the random variable \mathbf{Y}_k be the set of measurements received at time t_k and \mathbf{y}_k denote a value of \mathbf{Y}_k . We extend (1) to define

$$L_k(\mathbf{y}_k|\mathbf{s}) = \Pr\{\mathbf{Y}_k = \mathbf{y}_k | \mathbf{X}(t_k) = \mathbf{s}\} \text{ for } \mathbf{s} \in \mathbf{S}^+. \quad (2)$$

$L_k(\mathbf{y}_k|\cdot)$ is the multitarget likelihood function for the measurement set $\mathbf{Y}_k = \mathbf{y}_k$. If the sensor responses are correlated or there are restrictions such as a target can generate at most one measurement in a set, then these must be taken into account in computing this likelihood function.

Let $\mathbf{Y}_{1:K} = (\mathbf{Y}_1, \mathbf{Y}_2, \dots, \mathbf{Y}_K)$ and $\mathbf{y}_{1:K} = (\mathbf{y}_1, \dots, \mathbf{y}_K)$. These are the measurement sets received at the times $\{t_1, \dots, t_K\}$.

Define

$$\begin{aligned} \mathbf{L}(\mathbf{y}_{1:K}|\mathbf{s}_1, \dots, \mathbf{s}_K) \\ = \Pr\{\mathbf{Y}_{1:K} = \mathbf{y}_{1:K} | X(t_1) = \mathbf{s}_1, \dots, X(t_K) = \mathbf{s}_K\}. \end{aligned} \quad (3)$$

We assume that the distribution of the measurements at the times $\{t_1, \dots, t_K\}$ depends only on the system states at these times. That is,

$$\Pr\{\mathbf{Y}_{1:K} = \mathbf{y}_{1:K} | \mathbf{X}(u), 0 \leq u \leq t\} = \mathbf{L}(\mathbf{y}_{1:K}|\mathbf{s}_1, \dots, \mathbf{s}_K), \quad (4)$$

where $\mathbf{s}_k = \mathbf{X}(t_k)$ for $k = 1, \dots, K$.

Let

$$q(\mathbf{s}_1, \dots, \mathbf{s}_K) = \Pr\{\mathbf{X}(t_1) = \mathbf{s}_1, \dots, \mathbf{X}(t_K) = \mathbf{s}_K\}.$$

Then, the posterior distribution on the multiple target state at time t_K given $\mathbf{Y}_{1:K} = \mathbf{y}_{1:K}$ is

$$\begin{aligned} p(t_K, \mathbf{s}_K | \mathbf{y}_{1:K}) &= \frac{\Pr\{\mathbf{Y}_{1:K} = \mathbf{y}_{1:K} \text{ and } \mathbf{X}(t_K) = \mathbf{s}_K\}}{\Pr\{\mathbf{Y}_{1:K} = \mathbf{y}_{1:K}\}} \\ &= \frac{\int \mathbf{L}(\mathbf{y}_{1:K}|\mathbf{s}_1, \dots, \mathbf{s}_K) q(\mathbf{s}_1, \dots, \mathbf{s}_K) d\mathbf{s}_1 \cdots d\mathbf{s}_{K-1}}{\int \mathbf{L}(\mathbf{y}_{1:K}|\mathbf{s}_1, \dots, \mathbf{s}_K) q(\mathbf{s}_1, \dots, \mathbf{s}_K) d\mathbf{s}_1 \cdots d\mathbf{s}_K}, \end{aligned} \quad (5)$$

where the integral in the numerator of (5) is over the system states at the first $K - 1$ measurement times and the integral in the denominator is over these states at the first K times.

C. Bayes–Markov Recursion

If the motion model is Markovian so that

$$q(\mathbf{s}_1, \dots, \mathbf{s}_K) = \int_{\mathbf{S}^+} q_0(\mathbf{s}_0) \prod_{k=1}^K q_k(\mathbf{s}_k | \mathbf{s}_{k-1}) d\mathbf{s}_0 \quad (6)$$

where

$$q_0(\mathbf{s}) = \Pr\{\mathbf{X}(0) = \mathbf{s}\},$$

$$q_k(\mathbf{s}_k | \mathbf{s}_{k-1}) = \Pr\{\mathbf{X}(t_k) = \mathbf{s}_k | \mathbf{X}(t_{k-1}) = \mathbf{s}_{k-1}\},$$

and the likelihood function in (4) factors so that

$$\mathbf{L}(\mathbf{y}_{1:K}|\mathbf{s}_1, \dots, \mathbf{s}_K) = \prod_{k=1}^K L_k(\mathbf{y}_k|\mathbf{s}_k),$$

then the following Bayes–Markov recursion holds:

$$\begin{aligned} p(t_K, \mathbf{s}_K | \mathbf{y}_{1:K}) \\ = \frac{L_k(\mathbf{y}_K|\mathbf{s}_K) \int_{\mathbf{S}^+} q(\mathbf{s}_K|\mathbf{s}_{K-1}) p(t_{K-1}, \mathbf{s}_{K-1} | \mathbf{y}_{1:K-1}) d\mathbf{s}_{K-1}}{\int_{\mathbf{S}^+} L_k(\mathbf{y}_K|\mathbf{s}_K) \int_{\mathbf{S}^+} q(\mathbf{s}_K|\mathbf{s}_{K-1}) p(t_{K-1}, \mathbf{s}_{K-1} | \mathbf{y}_{1:K-1}) d\mathbf{s}_{K-1} d\mathbf{s}_K}. \end{aligned}$$

Observe that the above recursion does not require the notion of measurement association. The process of performing multiple target tracking with or without measurement association is called *unified tracking* in [13]. Ref. [13, Ch. 5] gives examples where two targets are tracked in a case where the notion of association is not meaningful. In some cases, the targets are indistinguishable and in others not. When association is not meaningful, standard MHT is not applicable to the multiple target tracking problem.

An example where association is not meaningful involves a fixed array of passive omnidirectional acoustic hydrophones. The measurement received at the sensor at time t is the vector of complex amplitudes (as a function of frequency) of the acoustic time series received at the hydrophones of the array at time t . When there is more than one target present, the signals from all targets are received and acoustically summed at each hydrophone so that it does not make sense to associate the measurement with a single target. The maximum posterior probability penalty function (MAP-PF) algorithm, described in [7, Ch. 6], uses the Bayes–Markov recursion above to perform multiple target tracking using these measurements without association or thresholding to produce contacts. By avoiding thresholding, one can utilize more information and provide better tracking solutions than if one is limited to using thresholded data (e.g., called contacts). The MAP-PF algorithm has been applied to a number of operational problems (see [7, Ch. 6, refs. 1–5]).

One can sometimes force the problem into an MHT framework, but the results are suboptimal. In particular, the results are not Bayes optimal. Ref. [13, Sec. 5.3.2]

gives an illustration of this and the resulting degradation of the tracking performance that results.

Ref. [13, Ch. 5] shows that MHT can be derived as a special case of unified tracking that motivates the name because the above recursion provides a unified approach to tracking with or without measurement association.

III. MULTIPLE HYPOTHESIS TRACKING

We take a more expansive definition of MHT than is normally the case—see [4], for example. In particular, we do not require targets to be indistinguishable, or false measurements to be Poisson distributed. We do not require that the target motion processes be independent or Markovian. In addition, Gaussian assumptions are not required. We *do require* that measurements arrive in scans as defined below.

A global measurement association hypothesis, defined more precisely in Section III-A, assigns all measurements received up to a given time to targets or false measurements. These hypotheses are target-to-measurement hypotheses, which are different from measurement-to-measurement association hypotheses used in the standard MHT formulation.

Definition. An *MHT* is a tracker that computes the posterior distribution on system state as follows. It identifies all possible global measurement association hypotheses and calculates their probabilities of being true, computes the conditional target state distributions given each hypothesis, and forms the posterior distribution as a mixture of the conditional target state distributions weighted by the association hypothesis probabilities. This is called the *MHT decomposition*.

In Section III-B, we show that under the conditions assumed in Sections II-A and II-B and the assumption that measurements come in scans, the MHT decomposition produces the exact Bayesian posterior on system state. That is, it is the exact Bayesian solution, theoretically correct, and “Bayes optimal.”

However, the MHT decomposition will be of limited use unless the conditional target state distributions can be computed recursively and the target state random variables are independent given a global measurement association hypothesis. Sections IV and V provide conditions under which these are true.

A. Scans and Global Measurement Association Hypotheses

Definition. A set of measurements at time t_k is a *scan* if each measurement is generated by at most one target and each target generates at most one measurement. We also require that the association of measurements in different scans is independent.

Note that this definition means that not every measurement is a scan by itself.

Assumption. We assume measurements arrive in scans.

Some of these measurements may be false measurements, i.e., not generated by a target, and some targets may not produce measurements on a given scan. Let

G_j = set of measurements in the j th scan,

$G(1 : k)$ = set of measurements in the first k scans

$$= \bigcup_{j=1}^k G_j.$$

Definition. A *global measurement association hypothesis* h on $G(1 : k)$ is a mapping $h : G(1 : k) \rightarrow \{0, 1, \dots, \bar{N}\}$ such that

$h(m) = n > 0$ means measurement m is associated with target n ,

$h(m) = 0$ means measurement m is associated with a false measurement,

and no target has more than one measurement per scan associated with it.

Let $H(k)$ = set of global measurement association hypotheses on $G(1 : k)$. A hypothesis $h \in H(k)$ partitions $G(1 : k)$ into disjoint subsets

$$\Psi_k(n) = \{m \in G(1 : k) : h(m) = n\} \text{ for } n = 0, 1, \dots, \bar{N},$$

where $\Psi_k(n)$ is the subset of measurements associated with target n for $n > 0$ and $\Psi_k(0)$ is the subset of measurements associated with false measurements.

B. MHT Decomposition

MHT calculates the posterior distribution on system state at time t_K given the global measurement association hypothesis h is true and the probability $\alpha(h|\mathbf{y}_{1:K})$ that hypothesis h is true given $\mathbf{Y}_{1:K} = \mathbf{y}_{1:K}$ for each $h \in H(K)$. Specifically, it computes

$$p(t_K, \mathbf{s}_K | h \wedge \mathbf{y}_{1:K}) = \Pr\{\mathbf{X}(t_K) = \mathbf{s}_K | h \wedge \mathbf{Y}_{1:K} = \mathbf{y}_{1:K}\} \quad (7)$$

and

$$\alpha(h|\mathbf{y}_{1:K}) = \Pr\{h | \mathbf{Y}_{1:K} = \mathbf{y}_{1:K}\} = \frac{\Pr\{h \wedge \mathbf{Y}_{1:K} = \mathbf{y}_{1:K}\}}{\Pr\{\mathbf{Y}_{1:K} = \mathbf{y}_{1:K}\}}, \quad (8)$$

where \wedge denotes conjunction. The Bayesian posterior is given by

$$p(t_K, \mathbf{s}_K | \mathbf{y}_{1:K}) = \sum_{h \in H(K)} p(t_K, \mathbf{s}_K | h \wedge \mathbf{y}_{1:K}) \alpha(h|\mathbf{y}_{1:K}). \quad (9)$$

Equation (9) is the MHT decomposition. The validity of this decomposition depends only on the assumptions in Sections II-A, II-B, and III-A. Thus, MHT is an exact Bayesian solution under very general assumptions. The main restriction is that measurements must arrive in scans. However, in most cases, we require more assumptions to compute MHT solutions.

IV. RECURSIVE MHT ASSUMPTIONS

In this section, we add assumptions that allow us to compute the MHT decomposition recursively. We assume the motion model is Markovian in the joint state space and that

$$\mathbf{L}(\mathbf{y}_{1:K}|\mathbf{s}_1, \dots, \mathbf{s}_K) = \prod_{k=1}^K L_k(\mathbf{y}_k|\mathbf{s}_k). \quad (10)$$

Let the Markov transition function be denoted by

$$q_k(\mathbf{s}_k|\mathbf{s}_{k-1}) = \Pr\{\mathbf{X}(t_k) = \mathbf{s}_k \mid \mathbf{X}(t_{k-1}) = \mathbf{s}_{k-1}\} \text{ for } k \geq 1$$

and q_0 be the probability (density) function for $\mathbf{X}(0)$. Then, we can compute the posterior distribution in (5) using the classic Bayes–Markov recursion as follows.

Initial distribution:

$$p(t_0, \mathbf{s}_0) = q_0(\mathbf{s}_0) \text{ for } \mathbf{s}_0 \in \mathbf{S}^+. \quad (11)$$

For $k \geq 1$ and $\mathbf{s}_k \in \mathbf{S}^+$,

$$p^-(t_k, \mathbf{s}_k|\mathbf{y}_{1:k-1}) = \int q_k(\mathbf{s}_k|\mathbf{s}_{k-1})p(t_{k-1}, \mathbf{s}_{k-1}|\mathbf{y}_{1:k-1})d\mathbf{s}_{k-1},$$

$$L_k(\mathbf{y}_k|\mathbf{s}_k) = \Pr\{\mathbf{Y}_k = \mathbf{y}_k \mid \mathbf{X}(t_k) = \mathbf{s}_k\},$$

$$p(t_k, \mathbf{s}_k|\mathbf{y}_{1:k}) = \frac{1}{C} L_k(\mathbf{y}_k|\mathbf{s}_k) p^-(t_k, \mathbf{s}_k|\mathbf{y}_{1:k-1}), \quad (12)$$

where

$$C = \int L_k(\mathbf{y}_k|\mathbf{s}_k) p^-(t_k, \mathbf{s}_k|\mathbf{y}_{1:k-1}) d\mathbf{s}_k.$$

A. Scan and Global Measurement Association Hypotheses

For the k th scan of measurements \mathbf{y}_k , let M_k = number of measurements in the scan.

Definition. A function $\gamma : \{1, \dots, M_k\} \rightarrow \{0, \dots, \bar{N}\}$ is a *scan association hypothesis* if $\gamma(m) = n > 0$ means measurement m is associated with target n , $\gamma(m) = 0$ means measurement m is associated with a false measurement, and no two measurements are assigned to the same positive number (target).

Let

Γ_k = the set of all scan association hypotheses on scan Y_k .

A global measurement association hypothesis $h_K \in H(K)$ is composed of K scan association hypotheses $\{\gamma_1, \dots, \gamma_K\}$, where γ_k is the association hypothesis for the k th scan. The global measurement association hypothesis h_K is an extension of $h_{K-1} = \{\gamma_1, \dots, \gamma_{K-1}\} \in H(K-1)$. That is, h_K is composed of h_{K-1} with γ_K appended. We write this as $h_K = h_{K-1} \wedge \gamma_K$.

1) **Scan Association Likelihood Function:** Define the scan association likelihood function

$$\ell_k(\mathbf{y}_k|\gamma \wedge \mathbf{s}_k) = \Pr\{\mathbf{Y}_k = \mathbf{y}_k \mid \gamma \wedge \mathbf{X}(t_k) = \mathbf{s}_k\} \text{ for } \mathbf{s}_k \in \mathbf{S}^+ \text{ and } \gamma \in \Gamma_k. \quad (13)$$

The conditioning on the right-hand side of (13) means that we are conditioning on the scan association hypothesis γ as well as the system state \mathbf{s}_k .

As a function of \mathbf{s}_k , the likelihood of the scan measurement computed in (13) accounts for the probability of detecting the targets with which measurements are associated, failing to detect the remaining targets, and the false measurements. The likelihood function for the scan $\mathbf{Y}_k = \mathbf{y}_k$ is

$$L_k(\mathbf{y}_k|\mathbf{s}_k) = \sum_{\gamma \in \Gamma_k} \ell_k(\mathbf{y}_k|\gamma \wedge \mathbf{s}_k) \Pr\{\gamma\} \text{ for } \mathbf{s}_k \in \mathbf{S}^+, \quad (14)$$

where on the right-hand side of (14) we assume that the (prior) probability of a scan association does not depend on the system state.

2) **Global Measurement Association Likelihood Function:** From (10), it follows that conditioned on $h \in H(K)$, the likelihood of the measurements received at times t_1, \dots, t_K depends only on the system state values at those times. Specifically, the global measurement association likelihood function l is

$$\begin{aligned} l(\mathbf{y}_{1:K}|h \wedge (\mathbf{s}_1, \dots, \mathbf{s}_K)) \\ &= \Pr\{\mathbf{Y}_{1:K} = \mathbf{y}_{1:K} \mid h \wedge \mathbf{X}(u) = \mathbf{s}_u; 0 \leq u \leq t_K\} \\ &= \Pr\{\mathbf{Y}_{1:K} = \mathbf{y}_{1:K} \mid h \wedge \mathbf{X}(t_k) = \mathbf{s}_k; k = 1, \dots, K\}. \end{aligned} \quad (15)$$

We assume that the scan association likelihoods are independent given $h \wedge (\mathbf{s}_1, \dots, \mathbf{s}_K)$, so that

$$l(\mathbf{y}_{1:K}|h \wedge (\mathbf{s}_1, \dots, \mathbf{s}_K)) = \prod_{k=1}^K \ell_k(\mathbf{y}_k|\gamma_k \wedge \mathbf{s}_k). \quad (16)$$

Finally, we assume that the prior probability of the global measurement association hypothesis h is equal to the product of the prior probabilities of its constituent scan association hypotheses. Specifically,

$$\Pr\{h_K\} = \prod_{k=1}^K \Pr\{\gamma_k\}, \text{ where } h = \{\gamma_1, \dots, \gamma_K\}. \quad (17)$$

Association probabilities: Define

$$C(h_0) = 1 \text{ and } C(h_K) = \Pr\{h_K \wedge \mathbf{Y}_{1:K} = \mathbf{y}_{1:K}\} \text{ for } K \geq 1. \quad (18)$$

Ref. [7, Sec. 4.5.1] shows that

$$\begin{aligned} C(h_K) &= C(h_{K-1}) \Pr\{\gamma_K\} \\ &\times \int \ell_K(\mathbf{y}_K|\gamma_K \wedge \mathbf{s}_K) p^-(t_K, \mathbf{s}_K|h_{K-1} \wedge \mathbf{y}_{1:K-1}) d\mathbf{s}_K \end{aligned} \quad (19)$$

and that the probability of the global measurement association $h \in H(K)$ being correct given $\mathbf{y}_{1:K}$ is

$$\alpha(h|\mathbf{y}_{1:K}) = \frac{C(h)}{\sum_{h' \in H(K)} C(h')}. \quad (20)$$

B. Recursive Calculation of MHT Decomposition

Under the above assumptions, ref. [7, Sec. 4.2.5] shows how the conditional target state distribution in (7) and the association hypothesis probabilities in (8) may be calculated recursively. This allows us to calculate the MHT decomposition in (9) in a recursive fashion.

V. INDEPENDENT MHT

In this section, we give additional assumptions that assure that the conditional target state distributions are independent so that $p(t_K, \mathbf{s}_K | h \wedge \mathbf{y}_{1:K})$ in (9) equals the product of independent probability distributions on the \bar{N} possible targets. In this case, the MHT decomposition in (9) becomes

$$\begin{aligned} p(t_K, \mathbf{s}_K | \mathbf{y}_{1:K}) &= \sum_{h \in H(K)} \alpha(h | \mathbf{y}_{1:K}) p(t_K, \mathbf{s}_K | h \wedge \mathbf{y}_{1:K}) \\ &= \sum_{h \in H(K)} \alpha(h | \mathbf{y}_{1:K}) \prod_{n=1}^{\bar{N}} p_n(t_K, x_n | h \wedge \mathbf{y}_{1:K}) \\ &\quad \text{for } \mathbf{s}_K = (x_1, \dots, x_{\bar{N}}) \in \mathbf{S}^+, \end{aligned} \quad (21)$$

where $p_n(t_K, \cdot | h \wedge \mathbf{y}_{1:K})$ is the marginal distribution on target n . If target n is not present at time t_K under hypothesis h , then $p_n(t_K, \cdot | h \wedge \mathbf{y}_{1:K}) = 1$.

A. Conditionally Independent Association Likelihoods

Definition. The likelihood of a scan $\mathbf{Y}_k = \mathbf{y}_k$ obtained at time t_k is *conditionally independent* if and only if for all scan association hypotheses $\gamma \in \Gamma_k$,

$$\begin{aligned} \ell_k(\mathbf{y}_k | \mathbf{s}_k = (x_1, \dots, x_{\bar{N}}) \wedge \gamma) \\ &= \Pr\{\mathbf{Y}_k = \mathbf{y}_k | \gamma \wedge \mathbf{X}(t_k) = (x_1, \dots, x_{\bar{N}})\} \\ &= g_0^\gamma(\mathbf{y}_k) \prod_{n=1}^{\bar{N}} g_n^\gamma(\mathbf{y}_k, x_n) \end{aligned} \quad (22)$$

for some functions g_n^γ , $n = 0, \dots, \bar{N}$, where g_0^γ can depend on the scan measurements but not \mathbf{s}_k . For $n > 0$, $g_n^\gamma(\mathbf{y}_k, \cdot)$ is typically the likelihood function for the measurement in \mathbf{y}_k that is associated with target n , which may be no measurement, and $g_0^\gamma(\mathbf{y}_k)$ is the probability of receiving the false measurements and measurements generated by targets as specified by the scan association hypothesis γ .

B. Independence Theorem

Under the assumptions of conditional independence of the scan association likelihood functions and independence of the target motion models, MHT decomposes the multiple target tracking problem into \bar{N} independent single target problems by conditioning on a global measurement association hypothesis. Let $q_k^n(s_{n,k} | s_{n,k-1})$ be the transition function at time t_k for target n for $n = 1, \dots, \bar{N}$. The following theorem and proof are from [7].

Independence theorem. Suppose the prior target motion processes are mutually independent so that the multiple target transition function factors as follows:

$$q_k(\mathbf{s}_k | \mathbf{s}_{k-1}) = \prod_{n=1}^{\bar{N}} q_k^n(s_{n,k} | s_{n,k-1}) \quad (23)$$

and the scan association likelihood functions are conditionally independent. Then, the posterior system state distribution conditioned on a global measurement association hypothesis is the product of independent distributions on the targets' states.

Proof. Let $\mathbf{Y}_{1:K} = \mathbf{y}_{1:K}$ be the scan measurements that are received at times $0 \leq t_1 \leq \dots \leq t_K \leq t$. Recall that $H(k)$ is the set of all global measurement association hypotheses on the first k scans. We wish to show for $k = 1, \dots, K$ that

$$\begin{aligned} p(t_k, \mathbf{s}_k | h \wedge \mathbf{y}_{1:K}) &= \prod_{n=1}^{\bar{N}} p_n(t_k, x_n | h \wedge \mathbf{y}_{1:K}) \text{ for } h \in H(k) \\ \text{and } \mathbf{s}_k &= (x_1, \dots, x_{\bar{N}}) \in \mathbf{S}^+, \end{aligned} \quad (24)$$

where

$$\begin{aligned} p_n(t_k, x_n | h \wedge \mathbf{y}_{1:K}) &= \Pr\{X_n(t_k) = x_n | h \wedge \mathbf{Y}_{1:K} = \mathbf{y}_{1:K}\} \\ &\quad \text{for } x_n \in S^+ \text{ and } n = 1, \dots, \bar{N}. \end{aligned}$$

We will prove the theorem by induction.

$k = 1$: We first show that (24) holds for $k = 1$. By the independence of the prior target motion processes,

$$p(0, \mathbf{s}) = \prod_{n=1}^{\bar{N}} p_n(0, x_n) \text{ for } \mathbf{s} = (x_1, \dots, x_{\bar{N}}) \in \mathbf{S}^+,$$

where $p_n(0, \cdot)$ is the initial state distribution on target n . Since the motion models for the targets are independent, the joint distribution at time t_1 before updating for the scan of measurements $\mathbf{Y}_1 = \mathbf{y}_1$ is

$$p^-(t_1, \mathbf{s}_1) = \prod_{n=1}^{\bar{N}} p_n^-(t_1, x_n) \text{ for } \mathbf{s}_1 = (x_1, \dots, x_{\bar{N}}) \in \mathbf{S}^+,$$

where $p_n^-(t_1, \cdot)$ is the motion-updated distribution for target n at time t_1 . A global measurement association hypothesis, $h \in H(1)$, is equal to a scan association hypothesis $\gamma \in \Gamma_1$. By the conditional independence assumption, the likelihood function for the scan \mathbf{Y}_1 factors into functions that depend only on the state of a single target and are independent of the state of the other targets.

To compute the posterior given $\mathbf{Y}_1 = \mathbf{y}_1$ and the association $h = \gamma$, we follow the recursion in (11) and (12) that clearly holds when we condition on a measurement association hypothesis. We multiply the motion-updated multiple target distribution at time t_1 by the likelihood function for $\mathbf{Y}_1 = \mathbf{y}_1$, both conditioned on γ , to obtain

$$\begin{aligned} p(t_1, \mathbf{s}_1 | \gamma \wedge \mathbf{y}_1) &\propto g_0^\gamma(\mathbf{y}_1) \prod_{n=1}^{\bar{N}} g_n^\gamma(\mathbf{y}_1, x_n) \prod_{n=1}^{\bar{N}} p_n^-(t_1, x_n) \\ &\propto g_0^\gamma(\mathbf{y}_1) \prod_{n=1}^{\bar{N}} [g_n^\gamma(\mathbf{y}_1, x_n) p_n^-(t_1, x_n)] \\ &\quad \text{for } \mathbf{s}_1 = (x_1, \dots, x_{\bar{N}}) \in \mathbf{S}^+. \end{aligned} \quad (25)$$

We obtain $p_n(t_1, x_n | \gamma \wedge \mathbf{y}_1)$, the marginal distribution on the state of target n , by integrating the right-hand side of (25) over all components except x_n and normalizing to obtain a probability distribution. The result is

$$p_n(t_1, x_n | \gamma \wedge h_1) \propto g_n^\gamma(\mathbf{y}_1, x_n) p_n^-(t_1, x_n) \\ \text{for } n = 1, \dots, \bar{N},$$

and we see that (24) holds for $k = 1$.

k implies $k + 1$: Suppose that (24) holds for the first k scans. Consider a global measurement association hypothesis $h_{k+1} \in H(k+1)$. Then, $h_{k+1} = \{h_k \wedge \gamma\}$ for some hypothesis $h_k \in H(k)$ and scan hypothesis $\gamma \in \Gamma_{k+1}$. Define

$$p^-(t_{k+1}, \cdot | h_k \wedge \mathbf{y}_{1:k}) = \text{distribution on } \mathbf{X}(t_{k+1}) \\ \text{given } h_k \wedge \mathbf{y}_{1:k}.$$

This distribution is obtained by performing the motion and information updates for the first k scans and the motion update only from time t_k to t_{k+1} . For target n , we define

$$p_n^-(t_{k+1}, \cdot | h_k \wedge \mathbf{y}_{1:k}) = \text{distribution on } X_n(t_{k+1}) \text{ given } \\ h_k \wedge \mathbf{y}_{1:k}.$$

By assumption, the target motion processes are independent. From this and the fact that (24) holds for k , we have

$$p^-(t_{k+1}, \mathbf{s}_{k+1} | h_k \wedge \mathbf{y}_{1:k}) = \prod_n p_n^-(t_{k+1}, x_n | h_k \wedge \mathbf{y}_{1:k}).$$

To obtain the posterior system state distribution at time t_{k+1} , we multiply $p^-(t_{k+1}, \mathbf{s}_{k+1} | h_k \wedge \mathbf{y}_{1:k})$ by the scan likelihood function conditioned on γ_{k+1} to obtain

$$p(t_{k+1}, \mathbf{s}_{k+1} | h_{k+1} \wedge \mathbf{y}_{1:k+1}) \\ = \frac{1}{C} g_0^{\gamma_{k+1}}(\mathbf{y}_{k+1}) \prod_n g_n^{\gamma_{k+1}}(\mathbf{y}_{k+1}, x_n) \prod_n p_n^-(t_{k+1}, x_n | h_k \wedge \mathbf{y}_{1:k}) \\ = \frac{1}{C} g_0^{\gamma_{k+1}}(\mathbf{y}_{k+1}) \prod_n g_n^{\gamma_{k+1}}(\mathbf{y}_{k+1}, x_n) p_n^-(t_{k+1}, x_n | h_k \wedge \mathbf{y}_{1:k}) \\ \text{for } \mathbf{s}_{k+1} = (x_1, \dots, x_{\bar{N}}) \in \mathbf{S}^+.$$

This shows that if (24) holds for k , then it is true for $k + 1$. Since we have shown that (24) holds for $k = 1$, the theorem is proved by mathematical induction.

VI. SUMMARY OF ASSUMPTIONS

In this section, we provide a summary of the assumptions we made to ensure the validity of the MHT decomposition, the recursive computation of the MHT decomposition, and the independence of the posterior distributions on the targets given a global measurement association hypothesis. In each case, MHT produces an exact Bayesian solution. The assumptions are cumulative; e.g., the assumptions in Section VI-B implicitly include those in Section VI-A.

A. Assumptions for Validity of MHT Decomposition

- *Continuous discrete formulation*: Targets move in continuous time but measurements are received at a discrete set of times $\{t_1, \dots, t_K\}$.
- *Target evolution*: Prior knowledge of the motion of the targets in state space is specified by a stochastic process.
- *Scan assumption*: Measurements are received in scans at the discrete times $\{t_1, \dots, t_K\}$.
- Measurements at t_k depend only on system state at t_k , and we can calculate the multiple target likelihood function in (3).

B. Assumptions for Recursive Computation of MHT Decomposition

- Target motion process is Markovian in system state.
- Measurement likelihood functions factor over scans:

$$L(\mathbf{y}_{1:K} | \mathbf{s}_1, \dots, \mathbf{s}_K) = \prod_{k=1}^K L_k(\mathbf{y}_k | \mathbf{s}_k),$$

$$l(\mathbf{y}_{1:K} | h \wedge (\mathbf{s}_1, \dots, \mathbf{s}_K)) = \prod_{k=1}^K \ell_k(\mathbf{y}_k | \gamma_k \wedge \mathbf{s}_k).$$

- Prior on global measurement association hypotheses factors:

$$\Pr\{h\} = \prod_{k=1}^K \Pr\{\gamma_k\}, \quad \text{where } h = \{\gamma_1, \dots, \gamma_K\}.$$

C. Assumptions for Independent Conditional Target State Distributions

- Prior motion models are independent, i.e.,

$$q_k(\mathbf{s}_k | \mathbf{s}_{k-1}) = \prod_{n=1}^{\bar{N}} q_k^n(s_{n,k} | s_{n,k-1}).$$

- Scan likelihood functions are conditionally independent, i.e.,

$$\ell_k(\mathbf{y}_k | \gamma \wedge \mathbf{s}_k = (x_1, \dots, x_{\bar{N}})) = g_0^\gamma(\mathbf{y}_k) \prod_{n=1}^{\bar{N}} g_n^\gamma(\mathbf{y}_k, x_n).$$

D. Comments

Looking at the assumptions for the validity of the MHT decomposition in Section VI-A, we see that only the scan assumption puts any substantial restriction on the class of problems for which MHT is the exact Bayesian solution. The assumptions in Sections VI-B and VI-C make explicit the assumptions made in most multiple target tracking problems to make them more computationally tractable. Most versions of MHT make the further assumption that the false alarm process is Poisson and that the motion and measurement models are linear Gaussian, at least approximately, so that a Kalman filter can be used to calculate target state distributions.

VII. CONCLUSIONS

From the above discussion, we see that the MHT decomposition produces the exact Bayesian solution to the multiple target tracking problem for target-to-measurement hypotheses under quite general assumptions. The main requirement is that measurements arrive in scans. In order to compute the decomposition recursively and to obtain independent conditional target distributions, we add the Markovian and independence assumptions given in Sections IV and V. Section VI summarizes these assumptions. Although the MHT decomposition applies to a great many multiple target tracking problems, it does not apply to all of them. Both unified fusion of [13, Ch. 5] and the RFS approach of [2] deal with problems beyond the purview of MHT.

Because the number of global measurement association hypotheses grows exponentially in the number of measurements received, implementation of an MHT algorithm requires approximations such as limiting the number of global association hypotheses or using a track-oriented approach (see [4]). This problem is not unique to MHT. All (existing) methods of producing an exact Bayesian solution require approximations of some sort to produce a computationally feasible algorithm for even modestly complex multiple target tracking problems.

REFERENCES

1. L. D. Stone
“Conditions for MHT to be an exact Bayesian solution to the multiple target tracking problem,” in *Proc. 21st Int. Conf. Inf. Fusion*, Cambridge, UK, 2018.
2. R. Mahler
Statistical Multisource–Multitarget Information Fusion. Norwood, MA: Artech House, 2007.
3. R. Mahler
“Measurement to track association and finite-set statistics,” arXiv:17101.07078v1, Jan. 5, 2017.
4. C. Y. Chong, S. Mori, and D. R. Reid
“Forty years of multiple hypothesis tracking—A review of key developments,” in *Proc. 21st Int. Conf. Inf. Fusion*, Cambridge, UK, 2018.
5. D. B. Reid
“An algorithm for tracking multiple targets,” *IEEE Trans. Autom. Control*, vol. 24, no. 6, pp. 843–854, Dec. 1979.
6. S. Mori, C. Y. Chong, E. Tse, and R. P. Wishner
“Tracking and classifying multiple targets without a priori identification,” *IEEE Trans. Autom. Control*, vol. 31, no. 5, pp. 401–409, May 1986.
7. L. D. Stone, R. L. Streit, T. L. Corwin, and K. L. Bell
Bayesian Multiple Target Tracking, 2nd ed. Boston, MA: Artech House, 2014.
8. R. E. Bethel and R. G. Rakka
“Optimum time delay detection and tracking,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 26, no. 5, pp. 700–712, Sep. 1990.
9. S. B. Colegrove, A. W. Davis, and J. K. Ayliff
“Track initiation and nearest neighbors incorporated into probabilistic data association,” *J. Elect. Electron. Eng.*, vol. 6, pp. 191–198, Sep. 1986.
10. G. E. Kopec
“Formant tracking using hidden Markov models and vector quantization,” *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 34, no. 4, pp. 709–729, Aug. 1986.
11. D. Musicki, R. Evans, and S. Stankovic
“Integrated probabilistic data association,” *IEEE Trans. Autom. Control*, vol. 39, no. 6, pp. 1237–1241, Jun. 1994.
12. R. L. Streit and R. F. Barrett
“Frequency line tracking using hidden Markov models,” *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 4, pp. 586–598, Apr. 1990.
13. L. D. Stone, C. A. Barlow, and T. L. Corwin
Bayesian Multiple Target Tracking, 1st ed. Boston, MA: Artech House, 1999.



Lawrence D. Stone received the B.S. degree in mathematics from Antioch College, Yellow Springs, OH, USA, in 1964, and the M.S. and Ph.D. degrees in mathematics from Purdue University, West Lafayette, IN, USA, in 1966 and 1967, respectively. He is currently a Chief Scientist at Metron Inc. In 1986, he produced the probability maps used to locate the *SS Central America* that sank in 1857, taking millions of dollars of gold coins and bars to the ocean bottom one and one-half miles below. In 2010, he led the team that produced the probability distribution that guided the French to the location of the underwater wreckage of Air France Flight AF447. He was one of the primary developers of the Search and Rescue Optimal Planning System used by the Coast Guard since 2007 to plan searches for people missing at sea. He continues to work on a number of detection and tracking systems for the United States Navy. He has coauthored the books *Bayesian Multiple Target Tracking*, 2nd ed. (2014) and *Optimal Search for Moving Targets* (2016). He is a member of the National Academy of Engineering and a Fellow of the Institute for Operations Research and Management Science (INFORMS). He was a recipient of the Jacinto Steinhardt Award from the Military Applications Section of INFORMS in recognition of his applications of operations research to military problems, and Lanchester Prize from the Operations Research Society of America in 1975 for his text *Theory of Optimal Search*.