Multi-Target Surveillance: Distributed Multiple-Hypothesis Tracking, Graph-Based Methods, and Context Exploitation

Stefano Coraluppi
Systems & Technology Research (STR)
600 West Cummings Park
Woburn, MA 01801
USA

stefano.coraluppi@str.us

ABSTRACT

This manuscript offers a readable introduction to some aspects of multi-target tracking. Mathematical details are mostly deferred to the references. We start with a discussion of existence and evolution modelling of targets, leading to the classical track-oriented multiple-hypothesis tracking recursion. We then discuss some limitations in multiple-hypothesis tracking, and the performance and robustness benefits that may be achieved with distributed multiple-hypothesis tracking solutions. This leads to further advances with graph-based tracking and distributed architectures for context exploitation.

DEDICATION

This manuscript is dedicated with friendship and respect to the memory of Craig Alan Carthel, a brilliant applied mathematician and software architect who left us too soon in 2022. For 24 years, Craig developed advanced solutions to address important US and NATO surveillance problems in many domains. The author had the great fortune to be a close colleague and friend to Craig throughout this time. More information about Craig is at https://www.tributearchive.com/obituaries/25473378/craig-alan-carthel. Craig is greatly missed.

1.0 THE PROBLEM OF INTEREST

The interest in multi-target surveillance – generally referred to as *multi-target tracking* (MTT) – is an unknown number of objects that emerge on the scene, evolve in state space (which may include both spatial and feature dimensions), and may depart at some time. There are many models that have been proposed for target existence and dynamics. Here we discuss models that have some reasonableness, and that lend themselves to algorithmic solutions. Like all models, they are applicable to some but not all surveillance problems.

Models for both targets and sensors are generally stochastic, as this offers a mathematically convenient approach to modelling uncertainty. Targets are often modelled as evolving in continuous time, with discretization to a time sequence that includes sensor observation times. The sensors may collect data over a time continuum, though it is customary to consider a discrete sequence of observation times. This assumption is most appropriate for active sensors. For passive sensors, we often introduce time quantization to have the same modelling framework.

1.1 Target Existence

The existence of multiple targets is captured effectively via the Poisson process. This model emerges naturally if we assume a time t_0 at which there are no targets. This is the initial time at which we have an empty universe. A first object emerges at a time governed by an exponential distribution with rate λ_b. Subsequent births ensue, with the assumption of exponentially distributed inter-arrival times (with the same rate λ_b). Once in existence, targets remain alive for an exponentially distributed time with rate λ_χ.

The simplicity of this model is appealing. Of course, one may offer many objections to the model. What if the number of objects that may possibly emerge is finite? What is object birth is based on a spawning process from existing objects? What if object lifetimes exhibit a hard bound?

If we appeal to the reasonableness of the model and to judicious choices of λ_b and λ_χ, a tractable formulation follows readily. We are interested in discretized birth and death statistics at a sequence of times t_1, t_2, ... The number of births in a time interval is Poisson distributed, and the number of births in non-overlapping time intervals are independent. Defining $\Delta t_k = t_{k+1} - t_k$, we have the discrete-time birth rate $\mu_b(\Delta t_k)$ and death probability $p_\chi(\Delta t_k)$.

$$\mu_b(\Delta t_k) = \int_{t_k}^{t_{k+1}} \lambda_b e^{-\lambda_\chi (t_{k+1}-\tau)} d\tau = \frac{\lambda_b}{\lambda_\chi}\left(1 - e^{-\lambda_\chi \Delta t_k}\right) \quad (1)$$

$$p_\chi(\Delta t_k) = \int_{t_k}^{t_{k+1}} \lambda_\chi e^{-\lambda_\chi \tau} d\tau = 1 - e^{-\lambda_\chi \Delta t_k} \quad (2)$$

If we let $t_0 \to -\infty$, stationary statistics emerge; the state-state number of objects in existence is $\frac{\lambda_b}{\lambda_\chi}$. Note the subadditivity property $\mu_b(\Delta t_1 + \Delta t_2) < \mu_b(\Delta t_1) + \mu_b(\Delta t_2)$. This can b e und erst ood as an infa nt-mortality phenomenon. In discrete time, we only account for births of those objects that survive to the end of the discretization interval. Indeed, note that the birth rate is bounded from above by the time-integration of births: $\mu_b(\Delta t_k) < \lambda_b \Delta t_k$.

There exist weaker forms of independence in the MTT literature. For instance, we might bound the number of objects, and allow independent transitioning between existence and non-existence states. Such a model does have a form of independence across targets, but the number of births in non-overlapping intervals are no longer independent.

1.2 Target Evolution

As with target existence, we rely on stochastic models defined in continuous time, for which time discretization yields the equations of interest. Here we limit the discussion to unconstrained motion in Cartesian coordinate systems. Further, we assume uncorrelated motion across dimensions. Hence, we may focus on models for one-dimensional target evolution.

The most employed models are those that rely on linear Gaussian assumptions. First-order models include *nearly constant position* (NCP) and the *Ornstein-Uhlenbeck* (OU) process. Second-order models are most commonly adopted and include *nearly constant velocity* (NCV), the *Integrated Ornstein-Uhlenbeck* (IOU) process, and the 2nd order OU process. Third-order models include nearly constant acceleration, also known widely as the Singer model. Note that a 2nd order or higher model is necessary to model Doppler or velocity measurements. These dynamical models are illustrated in Fig. 1 using standard engineering notation with integral blocks. Note that OU and IOU models revert to NCP and NCV, respectively, in the small-feedback limit.

Continuous-time expressions of all these models take the following classical form, using the Dirac delta function and orthogonal (uncorrelated) noise assumptions. The variation of constants formula enables a methodology for time discretization that may be adopted systematically for all linear Gaussian models of interest. Thus, all time-discretized models take the form in eqn. (4), with appropriate values for the state transition matrix A_k and the process noise covariance matrix Q_k. Typically, stochastic motion models are unstable, in the sense that uncertainty grows unbounded as a function of time. On the other hand, the OU process and 2nd order OU process admit stability and stationarity. This can be useful in long-horizon target simulations as well as in certain context-sensitive applications.

$$\dot{X}(t) = FX(t) + Bw(t), X(0) \sim N(0, \Sigma_0), w(t-\tau) \sim N\big(0, q \cdot \delta(t-\tau)\big), \text{ and } X(0) \perp w(t) \tag{3}$$

$$X_{k+1} = A_k X_k + \int_{t_k}^{t_{k+1}} \exp\big(F(t_{k+1} - t)\big) Bw(t)dt \Rightarrow X_{k+1} = A_k X_k + w_k \tag{4}$$

$$A_k = \exp(F \Delta t_k) \tag{5}$$

$$w_k \sim N(0, Q_k) \tag{6}$$

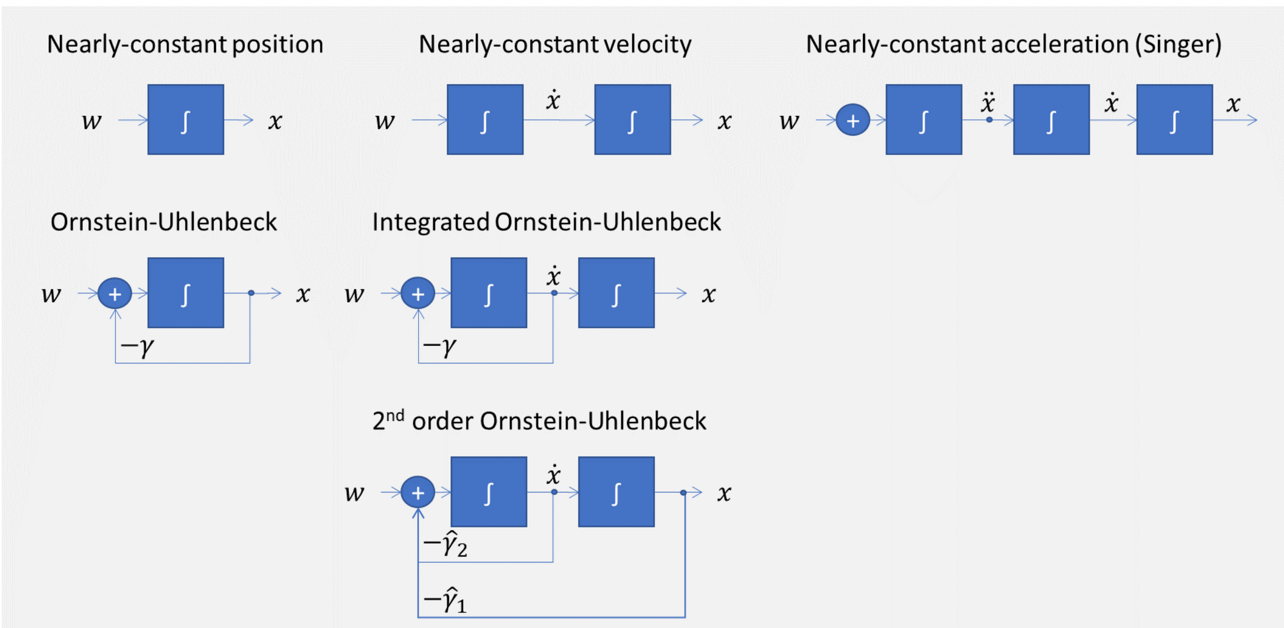

Figure 1. Common linear Gaussian target dynamical models.

1.3 Sensors

In MTT, we generally assume a frame-based sensor model whereby, at a discrete sequence of times, a sensor will observe a (known) field of regard defined in sensor measurement space. The subsequent mathematics simplify if we assume no merged measurements that originate from more than one target, no repeated measurements per target, and a target-state independent detection probability. The target observation, if it exists, is generally modelled as a nonlinear function of the target state that may include range, bearing angle, elevation angle, (bistatic) Doppler, and feature measurements that may include SNR, object classification, etc. The typical assumption is to have additive Gaussian noise for all kinematic measurements. These noises may be correlated in the case of derived measurements; as an example, in bistatic radar or sonar, we may compute a derived range measurement based on uncorrelated time and bearing measurements. The resulting range and bearing measurements are thus correlated.

For object classification measurements, we generally model misclassification via a confusion matrix normalized across rows. The columns of this matrix constitute the likelihood vectors on object type. In many settings, authors will refer to identity measurements when really what is intended is object type measurements. The distinction is that identities are unique, whereas object types are not. Knowledge of object identities is necessarily coupled across objects: if we know that one object is A, then nothing else can be A. On the other hand, for object type we generally have a prior distribution, and our knowledge of the type of one object tells us nothing about another. This form of type independence is important in MTT solutions, much like the independence assumptions in existence and evolution.

1.4 Performance Metrics

Classical metrics (target/track completeness & purity, and state estimation error) are illustrated in Fig. 2. They rely on solution to a sequential 2D assignment problem. In recent years, the *Optimal Subpattern Assignment* (OSPA) metric has become popular, particularly for label-free tracking methods. It combines detection and estimation performance into a single measure, that is a metric in the mathematical sense. Some anomalous aspects of OSPA are addressed in the *Generalized OSPA* (GOSPA) metric. Track-level extensions have also bene proposed recently.

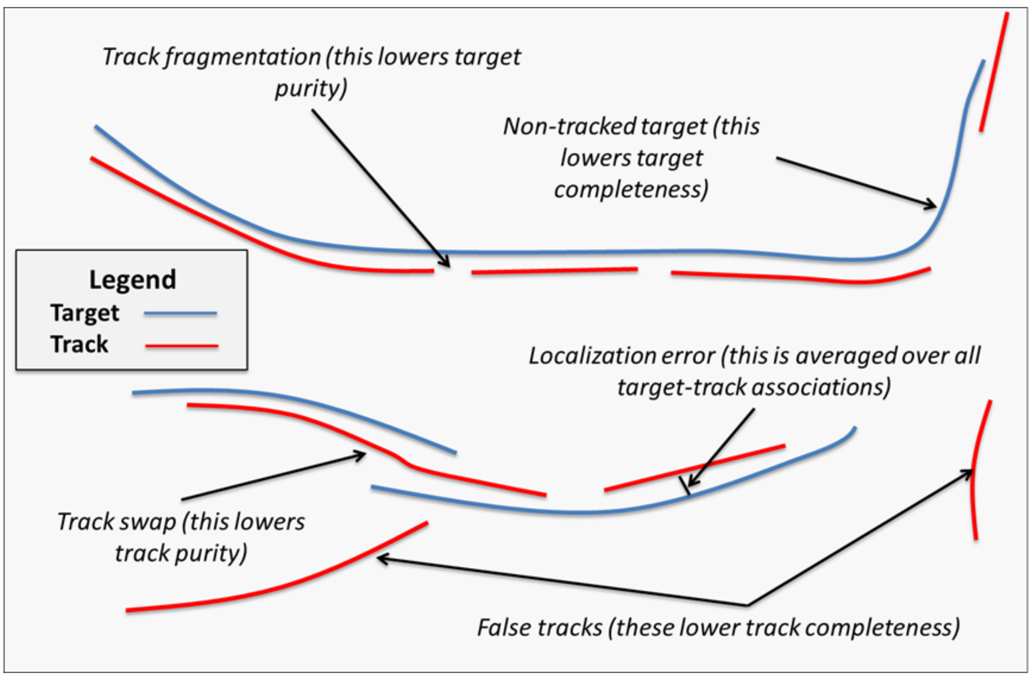

Figure 2. Classical MTT performance metrics.

For decision-level performance, it is perhaps most natural to consider the *probability of correct association* (PCA). This measures the number of correct association decisions relative to the total number of association decisions. On the other hand, this number can be improved by limiting association decisions to high-confidence decisions, with a detrimental impact on overall performance. Thus, we may compute as well the *fraction of correct association* (FCA), which measures the number of correct association decisions relative to the potential number of correct association decisions. Cautious stitching will tend to favor PCA, but degrade FCA.

The PCA and FCA may be defined locally or globally. For local metrics, we assess each pairwise association decision directly. This is easily understood in the illustration of Fig. 3. Consider seven tracks that are to be correctly associated, in the presence of an additional track that is due to another target. In Example 1, the extraneous track temporally follows Track 1, while in Example 2 the extraneous track precedes Track 1. Associations are noted in green (correct) and red (incorrect). Interestingly, though the same set of objects are associated, the local PCA and FCA are sensitive to temporal ordering. These metrics are captured in Table 1.

Alternatively, one may adopt a global PCA and global FCA metrics. With this approach, one considers fully connected sets of associated objects, based on association transitivity. This yields time-insensitive metrics; note in Table 1 that both Example 1 and Example 2 yield the same global PCA and global FCA. Under global reasoning, there are $\binom{5}{2}$ correct association decisions, 5 incorrect association decisions, and $\binom{7}{2}$ potential correct association decisions. These yield the global PCA and global FCA numbers captured in Table 1.

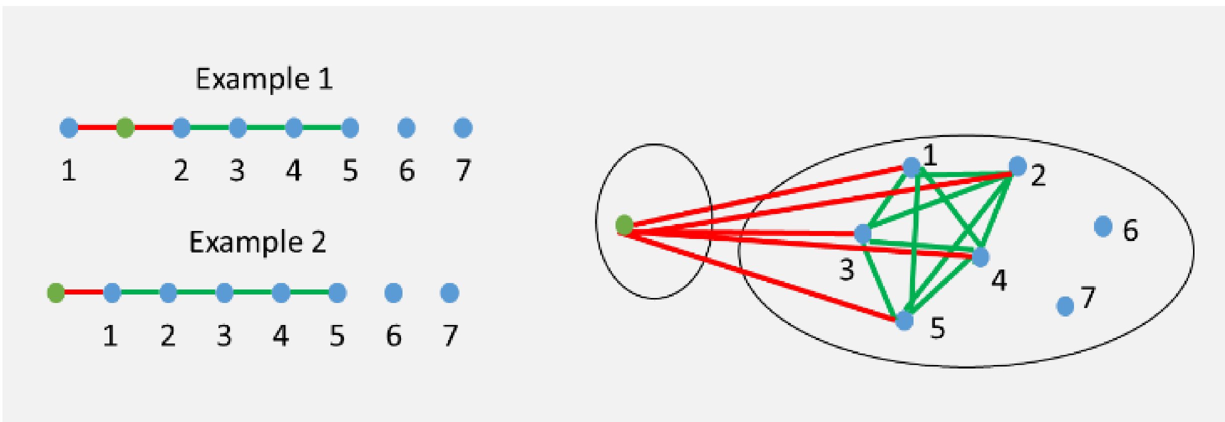

Figure 3. Two examples showing that local scoring (left) and global scoring (right) lead to different PCA and FCA values. The global values are insensitive to data ordering.

Table 1. Association performance for the example in Fig. 3.

Metric	Local PCA	Local FCA	Global PCA	Global FCA
Example 1	3/5	3/6	2/3	10/21
Example 2	4/5	4/6	2/3	10/21

1.5 Comments on theLiterature

A general discussion of multi-target surveillance may be found in [1-2] and references therein. The development of some advanced linear Gaussian motion models is in [3-4]. The key references on modern OSPA and GOSPA metrics are [5-6]. The primary international venue for discussion of developments in MTT technology and fusion algorithms across multiple domains is the ISIF/IEEE FUSION conference [7].

2.0 MULTI-TARGET TRACKING

Most approaches to *multi-target tracking* (MTT) contend explicitly with data association and include track-labelling. The *symmetric measurement equation* (SME) approach is an example of an *association-free* method that recasts the problem as one of nonlinear filtering., while the *probability hypothesis density* (PHD) filter is an example of a *label-free* method.

The SME approach recasts the problem with derived (symmetric measurements) with: (i) association-independent zero-mean noise, (ii) invertible (non-linear) observation function $y = h(x)$; (iii) full-rank Jacobean matrix H. As an example, consider measurements given by eqn. (7). Symmetric measurements may be defined as in eqns. (8-9); note that interchanging the measurements does not change the derived measurements

$$z_i(k) = x_i(k) + v_i(k), E[v_i(k)] = 0, E[{v_i}^2(k)] = \sigma^2, i = 1,\ 2 \tag{7}$$

$$y_1(k) = c_1\big(z_1(k) + z_2(k)\big) \tag{8}$$

$$y_2(k) = c_2(z_1^2(k) + z_2^2(k) - 2\sigma^2) \tag{9}$$

2.1 Multiple-Hypothesis Tracking

We observe a sequence of sets of measurements $X^k = (X_1, \dots, X_k)$ over a time sequence $t^k = (t_1, \dots, t_k)$. The full posterior probability distribution $p\big(X^k|Z^k\big)$ is rarely computed in practice. Even if we could do so, it is not well posed to consider the *maximum a posteriori* (MAP) estimate. A potentially useful decomposition is given in eqn. (10). However, note that the summation is over the large space of global hypotheses. Each global hypothesis specifies which measurements are discarded as not target originated, which are maintained and how these are associated. Additionally, there is a subtle distinction between the space of data association hypotheses, and the larger set of hypotheses that specifies additionally the time of birth of each target and (potentially) the time of death.

$$p\big(X^k|Z^k\big) = \sum_{q^k} p\big(X^k|Z^k, q^k\big)p\big(q^k|Z^k\big) \tag{10}$$

Multiple-hypothesis tracking (MHT) addresses both issues noted above: the meaningfulness of MAP optimization, and the large space of competing global hypotheses. We seek the MAP association hypothesis, and condition on this to determine a multi-target tracking solution without contending with a weighted combination over alternatives: see eqns. (11-12). Further, this solution can be expressed recursively as given by eqn. (13).

$$\hat{q}^k = \arg\max_{q^k} p\big(q^k|Z^k\big) \tag{11}$$

$$\hat{X}^k = \arg\max_{X^k}\ p\big(X^k|Z^k, \hat{q}^k\big) \tag{12}$$

$$p\big(q^k|Z^k\big) = \frac{p\big(Z_k|Z^{k-1}, q^k\big)p\big(q^k|q^{k-1}\big)p\big(q^{k-1}|Z^{k-1}\big)}{p\big(Z_k|Z^{k-1}\big)} \tag{13}$$

Notwithstanding the advances offered by the formulation above, the space of global data association hypotheses is generally large. Hence, we need to avoid explicit enumeration of alternatives. This can be achieved with track-oriented MHT, which factors global hypothesis scores into local track hypotheses. This factorization is lossless under the target independence assumptions discussed previously. Identifying the optimal global hypothesis can be posed as the solution to a binary linear program with inequality constraints.

MHT theory prescribes a batch solution approach, with coupled data association and track extraction. Due to online and computational constraints, a sliding-window approach is preferable. Most implementations decouple data association and track extraction, as these typically operate on disparate-duration time windows. In this approach, the data association module operates with equality constraints (as noted by "=" in Fig. 4), which amounts to including all detection data in track formation. Subsequently, only tracks that pass sufficient scoring and consistency criteria are extracted and confirmed. Further, we often favor confirmed tracks in the data-association process to mitigate the impact of sliding-window processing; this is indicated graphically with the feedback loop in Fig. 4. This feedback also enables significant computational savings, since tentative track hypotheses that will never achieve extraction criterion can be excised from further consideration.

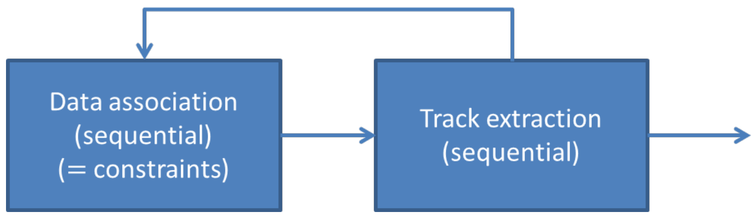

Figure 4. An effective MHT processing configuration.

2.2 Advanced and Distributed MHT

Centralized MHT using the basic track management structure in Fig. 4 can be effective, with good performance and robustness characteristics. With simulations of high target density and with pronounced clutter, track extraction outperforms significantly what the naked eye can achieve. One such realization is illustrated in Fig. 5. Further, performance improves monotonically, albeit with diminishing returns, as the data association reasoning window is increased.

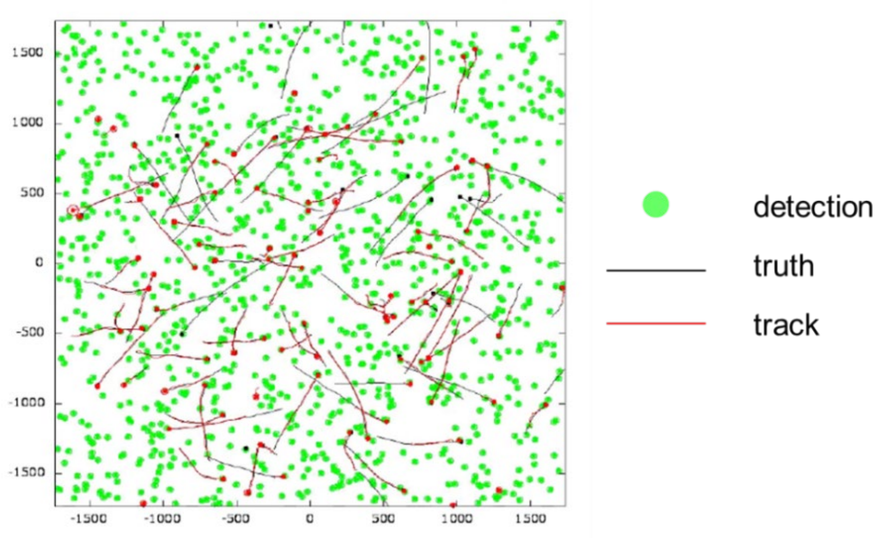

Figure 5. Highly effective baseline MHT processing.

There are many challenges in operational domains that often degrade performance. The point-target detection assumptions may fail, requiring explicit reasoning over merged or repeated measurements. Upstream detection and estimation processing may exhibit slowly time-varying trends, including fading-target phenomena (typically induced by sensor-target-receiver geometry) and bias errors, further complicating the idealized sensor assumptions. Passive sensors suffer limited state observability, typically due to unknown range to the target. High rate sensors provide information, but complicate processing and may render MHT temporally myopic.

Many advances to conventional MHT have been proposed to circumvent such challenges, including *feature-aided tracking* (FAT) methods that exploit non-kinematic target information. Hypothesis aggregation over similar hypotheses can offer great computational savings, enabling temporally deeper hypothesis exploration. Perhaps the most successful paradigm for high performance and robustness in multi-target surveillance is to rely on a divide-and-conquer methodology that decouples the challenges of measurement clustering, localization, clutter suppression, bias correction, and high-continuity tracking with a sequence of processing stages. Distributed processing has proven to be remarkably effective in both single-sensor and multi-sensor domains. At its essence, the approach enables asynchronous data association decision-making, with higher confidence decisions performed first. Another important aspect is the ability to perform lossless distributed estimation since measurements (and not state estimates) are passed to downstream modules. When appropriate, equivalent measurements are passed downstream to simplify downstream nonlinear filtering requirements.

A potential limitation of distributed MHT is that hard association decisions are performed in upstream modules, and these decisions are not revisited downstream. In fact, this too can be addressed by enabling track-breakage logic that allows for graceful recovery from inevitable upstream or in-stage association errors. More generally, distributed MHT need not be viewed as an inflexible feed-forward processing configuration. Enhanced connectivity between modules in the distributed architecture allows for both tracked and untracked measurements to be send to downstream processing modules. Depending on the application, what is of interest may include the tracked measurements (this is the usual case), the untracked measurements, or both.

As one example of effective distributed MHT, we mention that of *ground moving target indicator* (GMTI) radar tracking of ground targets with a low power radar system. Fig. 6 illustrates track extraction that significant outperforms what competing, less advanced MHT and non-MHT solutions can achieve. There are four MHT processing stages: the first forms detection clusters, the second tracks (and discards) data artifacts, the third considers untracked measurements and forms high-confidence, low-maneuver, high-fragmentation tracks while suppressing clutter, and the fourth achieves high-continuity maneuvering tracks.

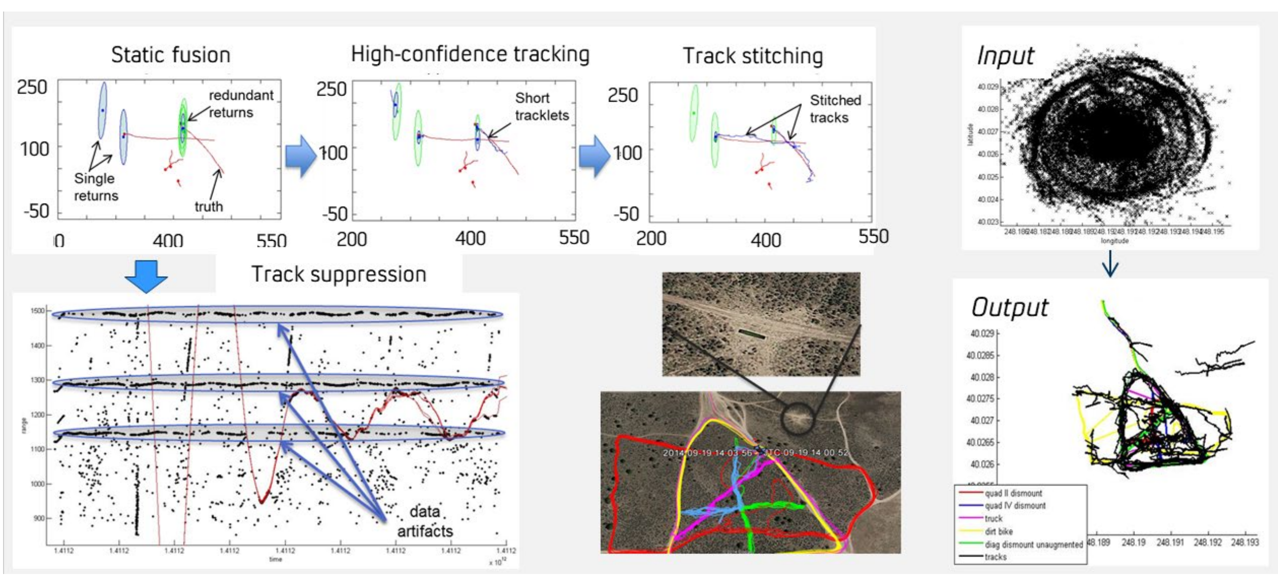

Figure 6. Not only multi-sensor tracking, but also single-sensor tracking can benefit greatly from an advanced distributed MHT processing solution.

2.3 Graph-Based Tracking

Despite the efficiency, performance and robustness that distributed MHT brings to challenging MTT problems, computational challenges persist particularly in disparate-sensor settings. As an illustration, consider the notional example in Fig. 7. We have an unknown number of objects giving rise to three identity tracks and four kinematic tracks, indicating that there are red and green objects present.

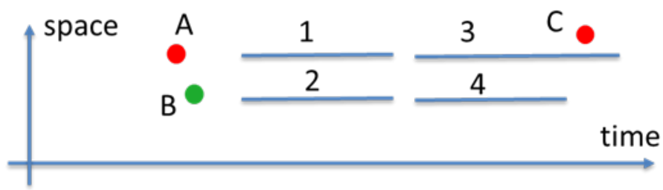

Figure 7. A disparate-sensor track fusion problem.

For simplicity of exposition, we assume a forensic surveillance problem, in the sense that all data has been received at the processing center. This enables a compact representation of the association spaces associated with competing solution approaches, with nodes representing input track. Regardless of online or forensic problems, the data-association process will necessarily rely on sliding-window processing for computational tractability.

Under the simplifying assumption of no never-observed objects (the usual assumption in MHT), there are at most seven objects present, and there are at least two targets (one red, one green). The MAP solution will depend on target and sensor statistical assumptions, and on the measured data itself. The corresponding data structure associated with track-oriented MHT processing may be represented as illustrated in Fig. 8. Note that, for simplicity, we have expressed each path in the MHT track forest as a sequence of tracks. This is not fully reflective of the actual processing sequence since data is ingested and processed in proper time order. As an example, in the leftmost path, some measurements associated with track 3 follow track C. Track filtering and scoring is performed in proper time sequence. The MAP solution associated with MHT processing will be that set of paths that accounts for all the data, while minimizing the sum of negative log likelihood track scores.

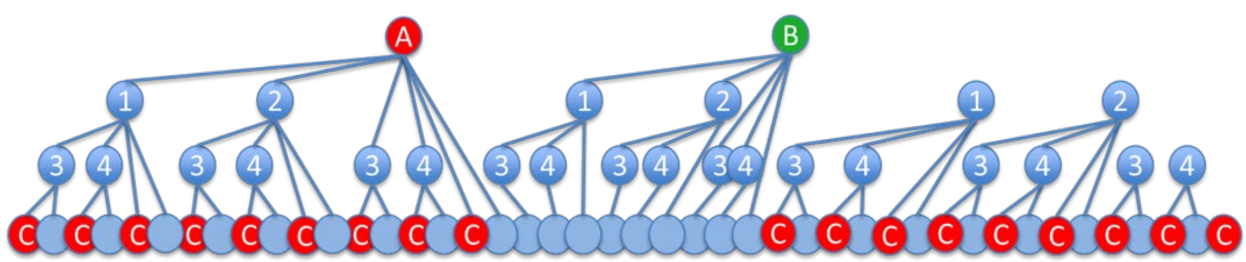

Figure 8. MHT track forest for the example in Fig. 7.

Fig. 9 illustrates a more efficient *graph-based tracking* (GBT) representation. By default, all edges are *directed* (downwards), except where explicitly denoted, i.e., the second edge between track 3 and node C. The graph-based structure is more compact than that of MHT, due to a simplifying path-independence assumption. Thus, for instance, each node associated with track 4 is a sufficient representation of kinematic information on the target, without the need for expressing from whence the target originates. At the same time, it is crucial to maintain graph layers (or subgraphs) associated with distinct object types. No flow is permitted between subgraphs, except for flow from the null subgraph to the object-type subgraphs.

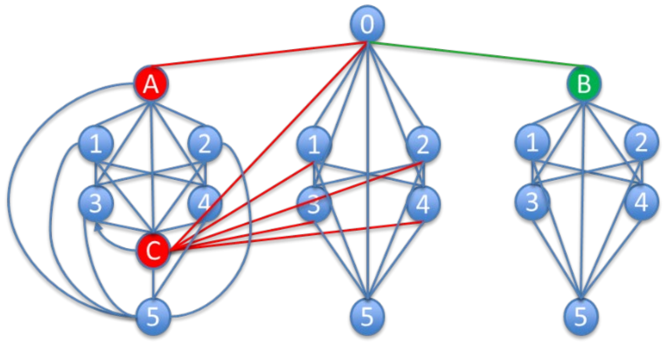

Figure 9. The GBT topology for the scenario in Fig. 7. Constraints in the binary linear program ensure that equivalent vertices only be used once.

GBT processing does rely on a simplifying Markovian assumption that the most recently observed track segment provides a kinematic sufficient statistic. This enables a dramatic reduction in the binary linear program that must be solved. In turn, this enables an improved complexity vs. performance tradeoff with respect to conventional MHT processing since much deeper hypotheses depths may be applied. This concept is illustrated in Fig. 10 and is supported by empirical evidence on disparate-sensor tracker benchmark problems.

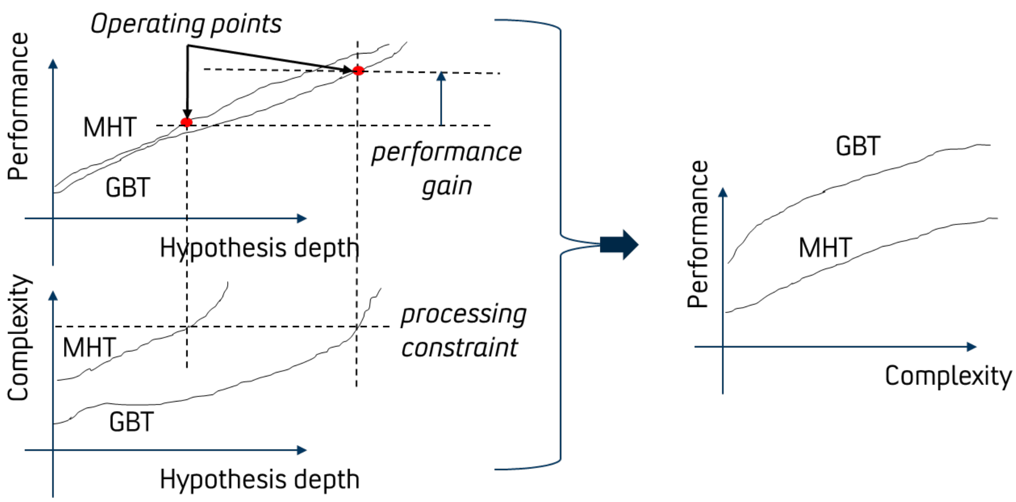

Figure 10. Improved performance-complexity tradeoff with GBT over MHT for some classes of problems.

2.4 Comments on the Literature

See [8] for a discussion of label-free tracking methodologies and [9] for a discussion of an association-free tracking approach. Key references in the early development of MHT include [10-12]. Extensions to account for more complex sensor models, distributed architectures, and graph-based methods include [13-15] and reference therein.

3.0 CONTEXT EXPLOITATION

Effective MTT generally requires a good sensor update rate relative to the target motion uncertainties. Note, for example, that with NCV recursive filtering one observes cubic growth in positional standard deviation in each dimension, as a function of time. Unfortunately, in some settings sensing assets are limited and one must seek to maintain target custody across nontrivial temporal coverage gaps.

Many authors have proposed context exploitation to overcome the limited-data problem. In some cases, such context is quite informative, as when we know that targets move on roads and good road network information is available. When context data is abundant, a common approach is to account for such information by modifying sensor data directly (e.g., by projecting measurements to the road network), or by introducing virtual context-induced sensor measurements (e.g., measurements placed on road segments). Even when context is good, one must contend with the multiplicity of ways that the context can be introduced: to which road segment should a measurement be projected? This mapping is more robustly achieved at the track level, following a first stage of context-free tracking. This is yet another example of distributed tracking, whereby context information is best introduced in a downstream track-stitching stage.

Context data may be less informative than what a good road network can provide. We may have patterns-of-life informing us of aggregate target behavior; examples of such context from various domains are illustrated in Fig. 11. In such cases, an alternative methodology is to embed context information in target dynamical models, without sensor data modification or enhancement. Such methods are often more robust and on firmer conceptual ground. We discuss one such class of methods below.

Figure 11. Examples of aggregate target motion information from air (left), ground (center), and maritime (right) domains.

3.1 Generalized OU Approach

When a nominal trajectory for a target is available, one may consider this as the deterministic component of motion to which a bounded stochastic perturbation is applied, the latter based on a 2nd order OU process. This is illustrated in Fig. 12. With this approach, one may continue to use NCV-based processing in local surveillance where data is available, and resort to generalized OU based processing to associated tracks across coverage regions. This approach is illustrated in Fig. 13 and represents yet another application of the distributed MHT surveillance paradigm.

The nominal trajectory may not be known. One might instead have a family of potential trajectories. In such cases, one may apply a variable-structure multiple-model OU filter, whereby an NCV mode is utilized when data is available, and a set of competing modes is utilized otherwise. This approach is particularly effective in some applications, such as blue-force custody where flight plan information is available. Recent extensions allow for dynamic updates to plans, as will occur when blue forces must replan and are able subsequently to re-establish communication. An illustration of a simulated testing environment for this application is in Fig. 14. Table 2 shows the performance gains that are achieved as one exploits multiple sensing modalities, in combination with dynamic context exploitation. It is worth noting that context exploitation increases target hold times (as reflected in target completeness), while average localization accuracy may be lower than in limited context-free tracking. Nonetheless, localization accuracy does not degrade in those track portions for which data is available. Further details of improved target hold for a particular target are in Fig. 15.

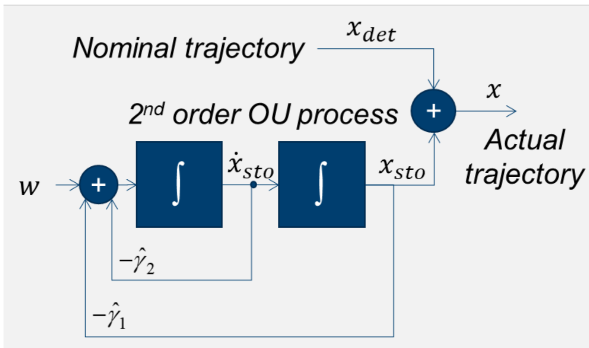

Figure 12. A context-aware target dynamical model.

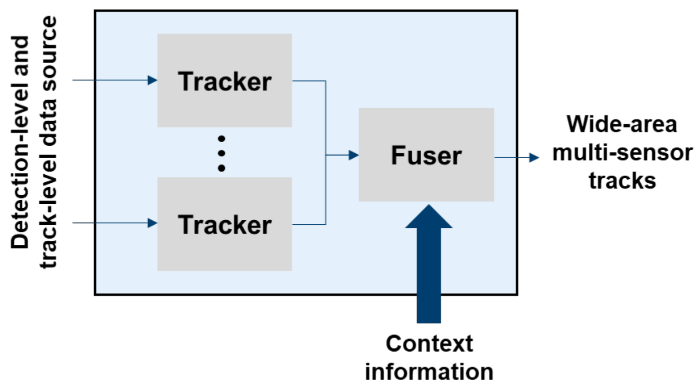

Figure 13. Distributed architecture for context-aware tracking.

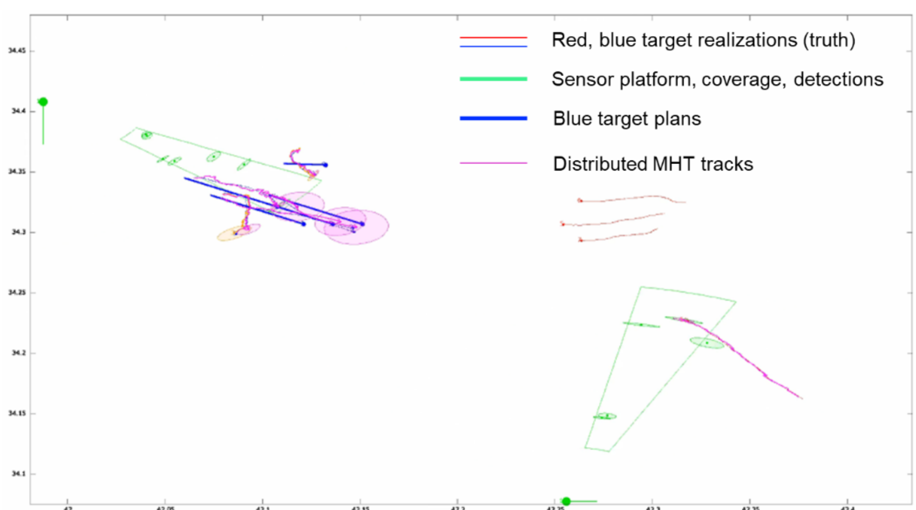

Figure 14. An air surveillance scenario with two scanning radars exhibiting coverage gaps.

Table 2. An example of Monte Carlo performance assessment quantifying the value of multiple sensors and of dynamic context exploitation. Note that targets may deviate from all reference paths, either inside or outside of sensor coverage. The context-enhanced solution is robust to this.

Sensors	Plan	Target completeness	Target purity	Track completeness	Track purity	Localization error [m]	Normalized OSPA
Radar	None	0.38	0.57	0.98	1.00	257	0.69
Radar	Static	0.54	0.75	0.75	0.95	408	0.67
Radar	Dynamic	0.61	0.91	0.93	0.98	452	0.59
Radar & EO	None	0.80	0.83	0.76	0.99	108	0.41
Radar & EO	Static	0.90	0.99	0.90	1.00	77	0.23
Radar & EO	Dynamic	0.95	1.00	0.95	1.00	91	0.15

Plan-Aware Tracking

1. Target exits radar coverage
2. Track is maintained despite unplanned deviation
3. Dynamic plan update is received
4. Target enters coverage 5. Target exits coverage

Baseline Tracking

1. Target exits radar coverage
2. Track is coasted briefly and terminated
3. Target enters radar coverage; track is reacquired and then terminated again

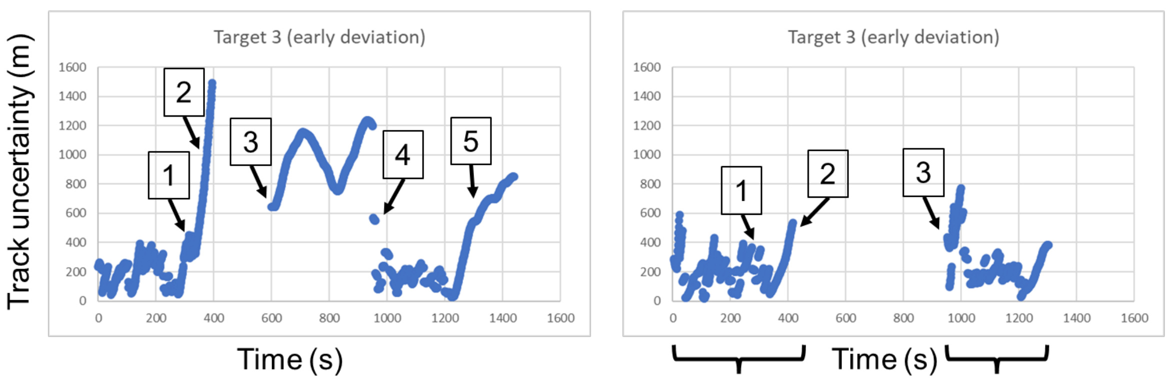

Single track persists throughout scenario (error exceeds 1500m during times 400-600)

Baseline tracker fails to hold target, forming two track fragments

Figure 15. An example of improved hold for a specific air target. Note that larger positional errors for the context-aware solution are at times during which the baseline solution fails to maintain track.

3.2 Sensor Management

When sensor coverage is limited, an effective strategy to maintain custody in addition to context exploitation is to adapt the use of sensor assets under our control. Sensor policies that reason over a time horizon are preferred, consistent with results in optimal control theory. As a notional example, consider a 1D target under NCP motion. Our knowledge of target location is given by a distribution over two modes, each with distinct state estimate, uncertainty, and process noise parameter. We select the target mode to observe, with a detection probability that decreases with the mode-specific area of uncertainty. We assume no false alarms.

We seek no maximize the probability of target reacquisition and consider three policies over a fixed time horizon: (i) randomized (randomly choose mode to view at each time), (ii) myopic (choose mode that maximizes the one-step reacquisition probability), and (iii) optimal (choose mode sequence to maximize reacquisition probability over the horizon). As shown in Fig. 16, empirical results confirm that the optimal policy is preferable over the myopic policy, and the latter is in turn preferable to the randomized policy.

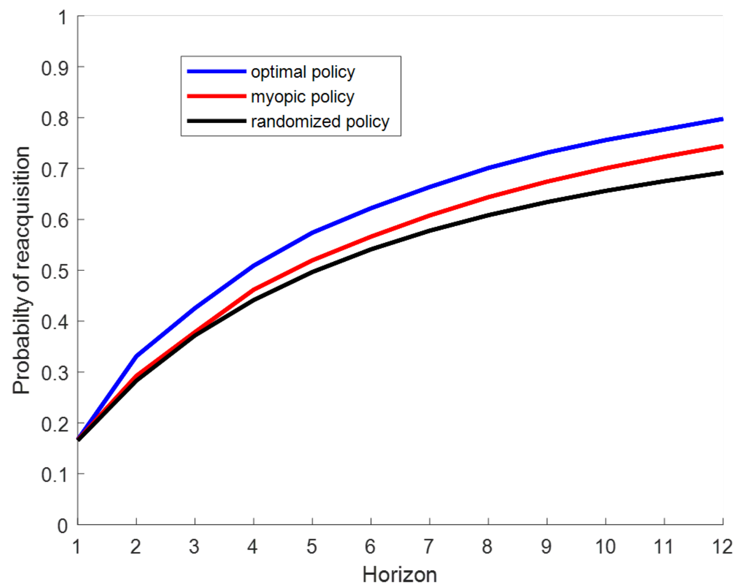

Figure 16. Sensor management policies for target reacquisition, with multi-modal state uncertainty.

3.3 Comments on the Literature

Exploitation of road information using the Viterbi algorithm is proposed in [16]. Our use of the generalized OU approach to context exploitation is developed in [17-19]. A good reference for sensor management algorithms for surveillance applications, including dynamic programming and entropy methods, is [20].

4.0 CONCLUSIONS

MHT is a leading paradigm for effective MTT. Notable extensions include distributed, graph-based, and context-aware processing. These innovations are important to enhance flexibility, increase robustness, reduce complexity, and enable limited-data target custody. A current area of research focus is to extend the context exploitation approach to allow for arbitrary uncertainty distributions, thus relaxing the Gaussian mixture assumption. Another area of investigation is to reconcile the MHT paradigm with identity management methods as has been proposed in soft data fusion settings.

4.0 REFERENCES

[1] Y. Bar-Shalom, P. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*, YBS Publishing, 2011.

[2] B.-N. Vo, M. Mallick, Y. Bar-Shalom, S. Coraluppi, R. Osborne, R. Mahler, and B.-T. Vo, Multitarget Tracking, in *Wiley Encyclopedia of Electrical and Electronics Engineering,* Wiley-IEEE, 2015.

[3] R. Singer, Estimating Tracking Filter Performance for Manned Maneuvering Targets, IEEE Transactions on Aerospace and Electronic Systems, Vol. 6(4), pp. 473-483, July 1970.

[4] S. Coraluppi and C. Carthel, Stability and Stationarity in Target Kinematic Modeling, in *Proceedings of the IEEE Aerospace Conference*, Big Sky MT, USA, March 2012.

[5] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, A Consistent Metric for Performance Evaluation of Multi-Object Filters, *IEEE Transactions on Signal Processing*, vol. 56(8), pp. 3447-3457, August 2008.

[6] A. Rahmathullah, Á. García-Fernández, and L. Svensson, Generalized optimal sub-pattern assignment metric, in *Proceedings of the 20th International Conference on Information Fusion*, Xi'an, China, July 2017.

[7] https://isif.org/events/conference/concluded

[8] R. Mahler, Advances in *Statistical Multisource-Multitarget Information Fusion*, Artech House, 2014.

[9] E. Kamen, Multiple Target Tracking Based on Symmetric Measurement Equations, *IEEE Transactions on Automatic Control*, Vol. 37(3), pp. 371-374, March 1992.

[10] C. Morefield, Application of 0-1 Integer Programming to Multitarget Tracking Problems, *IEEE Transactions on Automatic Control*, vol. 22(3), pp. 302-312, June 1977.

[11] D. Reid, An Algorithm for Tracking Multiple Targets, *IEEE Transactions on Automatic Control*, vol. 24(6), pp. 843-854, 1979.

[12] T. Kurien, "Issues in the Design of Practical Multitarget Tracking Algorithms," Chapter 3, in *Multitarget-Multisensor Tracking: Advanced Applications*, Y. Bar-Shalom (Ed.), Artech House, 1990.

[13] S. Coraluppi and C. Carthel, Multiple-Hypothesis Tracking for Targets Producing Multiple Measurements, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54(3), pp. 1485-1498, June 2018.

[14] C.-Y. Chong, S. Mori, and D. Reid, Forty Years of Multiple Hypothesis Tracking, *ISIF Journal of Advances in Information Fusion*, vol. 14(2), pp. 131-153, December 2019.

[15] S. Coraluppi, C. Carthel, and A. Willsky, Multiple-Hypothesis Tracking and Graph-Based Extensions, *ISIF Journal of Advances in Information Fusion*, vol. 14(2), pp. 152-166, December 2019.

[16] M. Bshara, U. Orguner, F. Gustaffson, and L. Van Biesen, Robust Tracking in Cellular Networks Using HMM Filters and Cell-ID Measurements, *IEEE Transactions on Vehicular Technology*, Vol. 60(3), pp. 1016-1024, March 2011.

[17] S. Coraluppi, C. Carthel, J. LeNoach, and B. Bale, The Ornstein-Uhlenbeck Process in Multi-Target Tracking, in *Proceedings of the IEEE Aerospace Conference*, Big Sky MT, USA, March 2021.

[18] C. Carthel, S. Coraluppi, J. LeNoach, B. Bale, N. Pioch, and J. Durek, Advances in Multi-Target Custody via Distributed Multiple-Model Multiple-Hypothesis Tracking, in *Proceedings of the MSS National Symposium on Sensor and Data Fusion*, Atlanta GA, USA, December 2021.

[19] C. Carthel, S. Coraluppi, B. Bale, N. Pioch, and J. Durek, Multi-Target Custody via Distributed Multiple-Model Multiple-Hypothesis Tracking, in *Proceedings of the IEEE Aerospace Conference*, Big Sky MT, USA, March 2022.

[20] A. Hero, D. Castañón, D. Cochran, and K Kastella (Eds.), *Foundations and Applications of Sensor Management*, Springer, 2008.