18th International Conference on Information Fusion
Washington, DC - July 6-9, 2015

Distributed MHT with Active and Passive Sensors

Stefano Coraluppi, Craig Carthel, Cynara Wu, Mark Stevens, Joel Douglas, Gerard Titi, and Mark Luettgen
Systems & Technology Research
600 West Cummings Park, Suite 6500, Woburn, MA 01801 USA
{stefano.coraluppi, craig.carthel, cynara.wu, mark.stevens, joel.douglas, gerard.titi, mark.luettgen}@STResearch.com

***Abstract*—This paper provides advances in distributed MHT. Our specific challenge is to provide a multi-sensor surveillance capability with disparate sensors: both active and passive sensors with widely varying clutter statistics, update rates, and coverage. Our first contribution is to develop a passive tracking capability that overcomes the lack of range observability with an effective use of virtual sensor measurements that provides efficiency, unbiased estimates, and consistent hypothesis scoring. Our second contribution is to enable high-performance fusion across active and passive sensors through equivalent-measurement processing. This has two benefits: equivalent measurements from passive data are sufficiently informative to score cross-sensor association hypotheses, and sensor update rates are made comparable across the network, thus simplifying downstream track fusion logic. We demonstrate the effectiveness of our distributed MHT solution with simulated air & ground targets in the context of Sense & Avoid for an unmanned aircraft with air and ground sensors.**

Keywords—multi-target tracking (MTT); multiple-hypothesis tracking (MHT); acitve and passive sensor fusion; virtual measurements; equivalent measurements; distributed tracking; Sense and Avoid (SAA).

I. INTRODUCTION

Multiple Hypothesis Tracking (MHT) is generally acknowledged as the most powerful solution paradigm for *multi-target tracking* (MTT) in real-world systems [1-2]. It was first formalized in what is now referred to as *hypothesis-oriented MHT* (HO-MHT) in [3]. Unfortunately, the hypothesis-oriented methodology typically leads to an unmanageable number of hypotheses even for small problems. The *track-oriented MHT* (TO-MHT) was developed in the 1980s by researchers at ALPHATECH and is well-described in [4]. Further advances in distributed MHT are discussed in [5-7] and references therein.

A fundamental assumption in our approach to distributed MHT is that we are interested in *online* track fusion with tracks that *include* detection-level information. This is in contrast to much of the track fusion literature, which focuses on the challenge of batch association of tracks that lack measurements and their associated statistics. The latter challenge is an important one, and is applicable to forensic applications with legacy sensors. It is worth noting that, under linear-Gaussian assumptions, missing sensor measurements may be recovered from target state estimates through a whitening filter.

In many settings, distributed MHT may outperform centralized MHT. In addition to the obvious benefits of reduced communications requirements and reduced computational load at the fusion center, the benefits of distributed MHT include: reduced sensitivity to data registration errors, reduced sensitivity to target fading effects, low-error single-sensor data association that exploits feature information, and deeper hypothesis tree depths in downstream track fusion processing for a comparable computational effort. The benefits of distributed processing have been noted in many application domains [8-9].

This paper provides advances for two important challenges in distributed MHT. The first is to generate localized passive tracks, with the well-known difficulty that passive measurements lack range observability. On the other hand, passive measurements often come with feature measurements that enable data association, such as object dimensions in imaging sensors, or target frequency information in *electronic warfare* (EW) and sonar sensors. In the context of *Sense & Avoid* (SAA), we assume that airborne electro-optical (EO) data provide object size measurements – given as number of pixels in the image plane. Our solution includes virtual measurements for track initiation and track gating, with no use of virtual measurements in track scoring and in track update filtering. This approach enables efficiency, unbiased state estimates, and consistent hypothesis scoring.

The second challenge is the need for effective cross-sensor integration. Passive state estimates will initially be very poor in the spatial dimensions relevant to data association with active tracks. Secondly, EO sensor rates will be higher than airborne radar sensor rates, and much higher than typical ground radar update rates. Effective integration must suppress clutter tracks from high-rate sensor feeds and not give undue emphasis to high-rate sensors. Our solution exploits equivalent-measurement processing for high-rate sensor tracks, a nearly lossless methodology to combine a sequence of measurements into a single, equivalent one. Our approach is closely related to the formation of sensor *tracklets*, as discussed in [10]. Correspondingly, all single-sensor tracks are sent to the fusion sensor at the same update rates, with well-localized tracks. This reduces spurious track association hypotheses and simplifies track fusion logic as equal decision weight is given to all sensors.

An appealing feature of our distributed MHT is its modularity, in the sense that multiple fusion architectures can be studied straightforwardly using a common MHT processing module. In addition to airborne EO data, we consider airborne radar and ground radar sensors. We quantify the performance of single-sensor tracking, centralized track fusion (i.e. fusion of single-sensor tracks), and distributed track fusion (i.e. fusion of fused airborne tracks and ground radar tracks). We demonstrate the effectiveness of our distributed tracking approach that

includes effective passive sensor tracking and high-performance multi-sensor fusion via equivalent-measurement processing.

The paper is organized as follows. Sections II and III provides some background on the SAA challenge and MHT technology, respectively. Section IV describes our technical advances in distributed MHT. The simulation framework and results are in Sections V-VI, and conclusions and directions for future work are in Section VII.

II. The Sense and Avoid Challenge

Detailed information on recent research in SAA technology may be found in a number of references including [11-13]. The focus of our effort is on analysis of integrated *ground-based sense and avoid* (GBSAA) and *airborne sense and avoid* (ABSAA) systems, with the goal of characterizing the tradeoffs in track fusion and data link architectures, and developing applicable track fusion technology.

Integration of track data from GBSAA and ABSAA systems offers the possibility of significantly improved terminal airspace tracks because these systems often have complementary sensing geometries and measurement error statistics. In particular, ABSAA sensors are often overwhelmed by ground clutter and thus, particularly in the terminal areas that contain many moving and stationary ground objects and that are the focus of our effort, ground-based *air traffic control* (ATC) systems offer the potential for an improved track picture to support *unmanned air system* (UAS) SAA maneuvers relative to a track picture based on ABSAA sensors alone.

In a UAS equipped with EO and radar sensors, there is the potential to integrate radar and EO tracks to create tracks with improved continuity characteristics. For example, considering only the first two rows in Figure 1, if the radar track on an intruder breaks during a period in which the EO sensor is able to maintain track, there is the potential to combine the tracks into a single composite ABSAA track. Such track integration, in which continuous track is maintained on the intruder for the entire duration covered by the EO and radar tracks, has already been demonstrated in flight tests by the *Air Force Research Laboratory* (AFRL) under previous and current ABSAA technology development efforts.

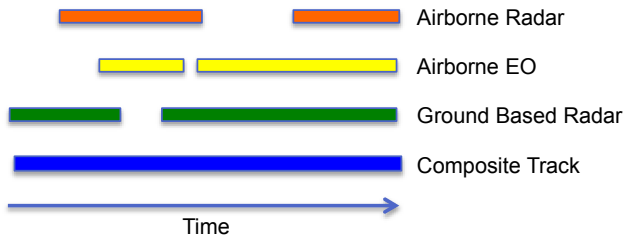

Figure 1. Integrated ABSAA/GBSAA Improves Track Completeness, Continuity, and Accuracy.

III. Multiple Hypothesis Tracking

Track management and data association are the key technical challenges in MTT. One is given a sequence of sets of measurements, and must determine which measurements to associate and which to discard. We have a sequence of sets of contacts $Z^k = (Z_1,...,Z_k)$, and we wish to estimate the state history X^k for all objects present in the surveillance region. X^k is compact notation that represents the state trajectories of targets that exist over the time sequence $(t_1,...,t_k)$. Note that each target may exist for a subset of these times, with a single birth and a single death occurrence, i.e. targets do not reappear. We introduce the auxiliary discrete state history q^k that represents a full interpretation of all contact data: which contacts are false, how the object-originated ones are to be associated, and when objects are born and die. There are two fundamental assumptions of note. The first is that there are no target births in the absence of a corresponding detection, i.e. we do not reason over new, undetected objects. The second is that there is *at most* one contact per object per scan.

We are interested in the probability distribution $p(X^k \mid Z^k)$ for object state histories given data. This quantity can be obtained by conditioning over all possible auxiliary states histories q^k.

$$p(X^k \mid Z^k) = \sum_{q^k} p(X^k \mid Z^k, q^k) p(q^k \mid Z^k). \tag{1}$$

The MHT approach seeks to identify the MAP estimate for the auxiliary state history q^k, and identify the corresponding MMSE estimate for the object state history X^k conditioned on the estimate for q^k.

$$\hat{q}^k = \arg\max_{q^k} p(q^k \mid Z^k), \tag{2}$$

$$\hat{X} = \hat{X}_{MMSE}(Z^k, \hat{q}^k). \tag{3}$$

TO-MHT avoids enumeration of all global hypotheses q^k, though these are implicitly defined in the set of track hypotheses trees. In practice, computational requirements preclude optimal TO-MHT solutions. Practical MHT solutions adopt a number of computation simplifications that includes hypothesis generation logic (both measurement gating and only one of coast & death hypotheses following missed detections) as well as hypothesis pruning (typically n-scan pruning that retains a single global hypothesis with some delay). Hypothesis pruning relies on well-known optimization techniques such as Lagrangian relaxation or linear programming relaxation [5, 16-17].

Track extraction is performed sequentially using logic-based or statistical tests including the *sequential probability ratio test* (SPRT) [1]. We opt for the former approach as in practice the SPRT generally is quite similar to an *M-of-N* confirmation test. Tracks are terminated after *K* missed detections.

For SAA, we model target dynamics with 3-D position-velocity state variables and a nearly constant velocity dynamical model in an *Earth Centered Earth Fixed* (ECEF) frame. We perform gating and calculate association likelihoods

in sensor coordinate frames, i.e. ABSAA processing is in inertial azimuth-elevation and azimuth-elevation-range-rdot space, and GBSAA processing is in inertial azimuth-elevation-range space. We use an *Extended Kalman Filter* (EKF) to capture nonlinearities, and use a sequential EKF that exhibits improved stability characteristics compared to a conventional EKF for processing of radar Doppler measurements [18].

EO feature measurements are given by object size, for which we assume a nearly constant motion model with linear measurements.

IV. ADVANCES IN DISTRIBUTED MHT

For effective single-sensor EO tracking, we use a virtual range measurement, defined by one-half the sensor detection range (with correspondingly large variance), for both track initiation and hypothesis gating. We do not use the virtual measurements for likelihood calculation or track updates. Indeed, the repeated use of virtual measurements introduces a bias towards tracks at one-half the sensor detection range, leading to poor track localization and poor track continuity due to fragmentation.

We use equivalent-frame functionality that allows for a sequence of tracker output frames to be represented by a single frame with little information loss. This item is particularly significant, both in terms of integrated GBSAA/ABSAA processing benefits demonstrated on our baseline encounter scenario, and the broader implications for SAA surveillance technology. Equivalent measurement processing is based on the fact that a sequence of possibly nonlinear measurements may be replaced by a single equivalent linear measurement by a lossless transformation in the sense that the state estimate and state covariance that result from a single prediction step and use of the equivalent measurement match those achieved by a sequence of prediction and update steps. In general, the dimension of the equivalent measurement vector may be as large as the state vector dimension. For example, a sequence of positional measurements may require an equivalent measurement that includes both position and velocity measurements.

The fundamental equations are as follows. Assume that we start with the result of a single prediction step, the predicted estimate X_- and predicted covariance P_-, as well as the result of the *sequence* of prediction-update steps, X_+ and P_+. We seek the measurement Z, the measurement covariance R, and the observation matrix C, that lead from X_- and P_- to X_+ and P_+ with a *single* equivalent measurement update step. This framework is captured in equations (4-6) below:

$$X_+ = X_- + L(Z - CX_-), \tag{4}$$

$$P_+ = (I - LC)P_-, \tag{5}$$

$$L = P_-C^T(CP_-C^T + R)^{-1}. \tag{6}$$

The columns of C^T are the orthonormal eigenvectors of $P_+^{-1} - P_-^{-1}$ corresponding to positive eigenvalues. Defining Λ as the matrix of the positive eigenvalues, the solution is given by:

$$L = (I - P_+P_-^{-1})C^T, \tag{7}$$

$$Z = CX_- + (L^TL)^{-1}L^T(X_+ - X_-), \tag{8}$$

$$R = \Lambda^{-1}. \tag{9}$$

Although the use of equivalent measurements represents a lossless transformation for the purposes of recursive filtering, the same is not true for the multi-target tracking problem for several reasons. First, an equivalent measurement frame requires the same or at least a similar sensor footprint among the frames of interest. This is approximately correct if the frames originate from the same sensor, and is the reason why we only consider equivalent frame formation at the output of single-sensor trackers. Second, for track hypothesis scoring, there is the need for equivalent detection probability and equivalent clutter rate to be defined. We do so with some common-sense but lossy approximations. Finally, the filter innovation associated with the equivalent measurement does not fully capture the information in the sequence of innovations in its impact on track score.

Though lossy, the use of equivalent-frame processing is a powerful computational tool that provides tracking performance benefit in fused SAA tracks. The benefit exists regardless of whether ground information is exploited, and regardless of fusion architecture. The obvious benefits of equivalent measurement processing are due to (i) bandwidth savings in sending data to the fusion center and (ii) reduced computational effort at the fusion center. In addition, there are more subtle advantages.

First, fusion is improved when sensor data rates are normalized to a common rate because clutter tracks in high data rate sensors are not over-weighted relative to negative measurements in lower data rate sensors. This is important in an integrated ABSAA/GBSAA fusion system to enable mitigation of clutter induced tracks in ABSAA sensors by leveraging the reduced clutter in the GBSAA sensor.

Second, we wish to optimize downstream data-association performance. For the purpose of downstream cross-sensor association, what is relevant is the information content in shared state variables and it is beneficial to delay sending data to the fusion center and to send an equivalent measurement when it is sufficiently informative. This is also important in an integrated ABSAA/GBSAA fusion system to enable processing of EO tracks after they have sufficient localization information to improve gating and fused state estimation. The multistage processing may be counterintuitive because, in classical filtering problems, centralized processing always provides an upper bound on distributed processing, and there is no benefit to delayed information transmission. On the other hand, in practical situations such as those we are investigating under this SAA effort in which computation is limited and models never match the physics of real situations, these paradigm shifts may enable robust information fusion.

With respect to data association and track management, we consider multiple association hypotheses between ABSAA and GBSAA tracks, evaluate those using the measurement likelihoods discussed above, periodically prune low likelihood hypotheses, and periodically report a high likelihood global hypothesis that accounts for all incoming tracks.

It is common to simplify MHT processing by separating the data association and track extraction functions. This involves accounting for all measurement data in resolving association hypotheses, and removing spurious tracks in subsequent track extraction. While suboptimal, sequential (multiple-hypothesis) data association followed by sequential track extraction does allow for feedback whereby confirmed tracks are given priority in the data-association stage; this methodology has been shown to work well in practice, and is applied here [19].

V. Simulation and Evaluation Framework

Current *Multiple Intruder Autonomous Avoidance* (MIAA) track fusion processing is based on the architecture shown in Figure 2. We have characterized baseline ABSAA-only performance by emulating this architecture using our distributed MHT architecture. A natural extension of the existing track fusion processing approach is shown in Figure 3: the GBSAA tracks are integrated by adapting the ABSAA track fusion processor to incorporate GBSAA tracks. This architecture has the disadvantage that it requires modification of the existing processing system although this may not be extensive if the processor is modular in architecture. It has the advantage that it processes all tracks simultaneously and thus introduces no latency and has no performance losses due to potential lack of information pedigree at the output of the current ABSAA processor.

Finally, we characterize the architecture in Figure 4 by adapting that same fusion processor again to operate in the two-stage architecture that first emulates the processing on the far right, and then integrates GBSAA report and track data. We find that the specific architecture (option 1 or option 2) is not the main driver of performance; the main driver of performance is the underlying fusion algorithm and adaptation of that algorithm to the sensors of interest, which can then be applied in either architecture.

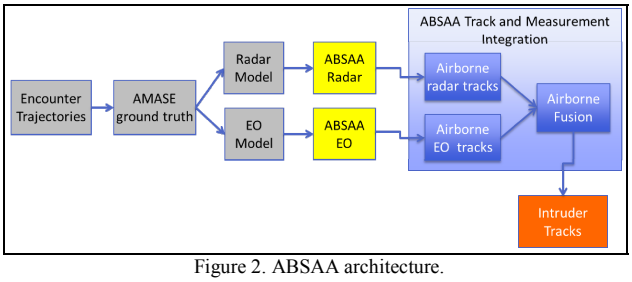

Figure 2. ABSAA architecture.

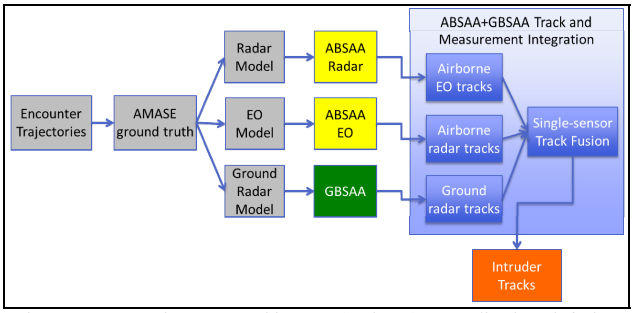

Figure 3. ABSAA/GBSAA architecture (option 1): centralized track fusion.

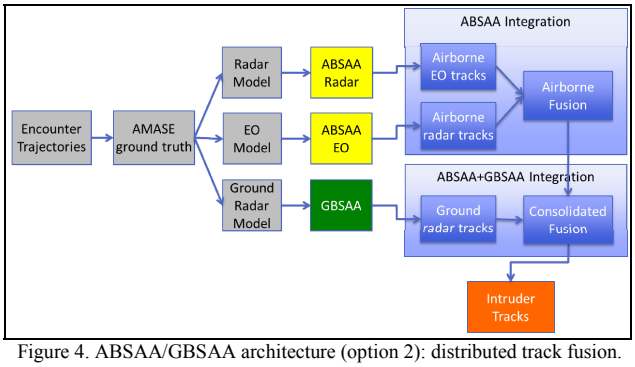

Figure 4. ABSAA/GBSAA architecture (option 2): distributed track fusion.

To provide an evaluation of the three fusion architectures, we need to test using relevant encounter scenarios. We accomplish this by creating encounter trajectories that are consistent with high-fidelity models developed by Lincoln Laboratories, populating the AFRL modeling and simulation toolset with trajectories to generate truth data, modeling the ABSAA and GBSAA sensors to generate EO and radar contacts from the truth data, running single-sensor trackers on the contacts to generate tracks, and sending the tracks to the fusion capabilities to generate multi-sensor tracks.

Figure 5 illustrates a sample scenario. This scenario is based on an AFRL-provided example that was modified to highlight the impact of clutter. It consists of the cyan UAS, which represents the ownship, and the magenta UAS, which represents an intruder. The UASs have been commanded to follow the cyan and magenta routes at 75 meters per second, and they are on a course that results in them being approximately 3 meters apart at their closest point of approach.

This encounter is representative of the uncorrelated encounter models for conventional aircraft developed by Lincoln Laboratory. We added a ground entity to the base scenario to represent the location of the ground radar, along with several moving ground entities to represent ground clutter. These entities and their paths are shown in green in Figure 5.

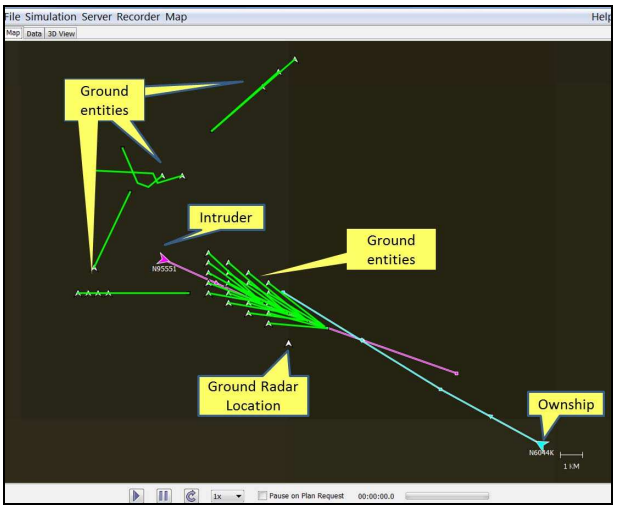

Figure 5. Illustration of Encounter Scenario.

We use the AFRL simulation tool to model sidelobe clutter discretes by creating additional simulation entities. These clutter tracks were based on analysis of a simple but representative azimuth monopulse processor. When a radar is pointed toward the horizon, the sidelobes cover a large area of ground, resulting in potentially significant sidelobe returns. This, in turn, results in angle ambiguities, illustrated in Figure 6. In this figure, the "x" in the center represents the return from the main beam. The "x"s representing the returns from the side lobes have the same apparent spatial coordinates. "Pixels" with the same color and intensity have the same monopulse response and are hence ambiguous with one another.

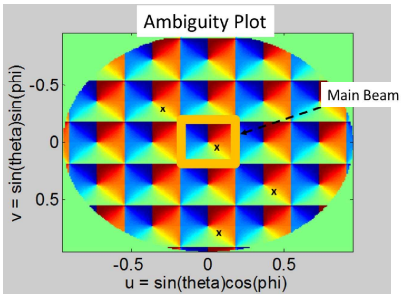

Figure 6. Example in which targets at position "x" all have the same apparent spatial coordinates.

Figure 7 illustrates how sidelobe clutter can result in a false track. We illustrate typical clutter trajectories based on our analysis for an airborne platform in straight and level flight at 40m/sec and 1 km above the ground.

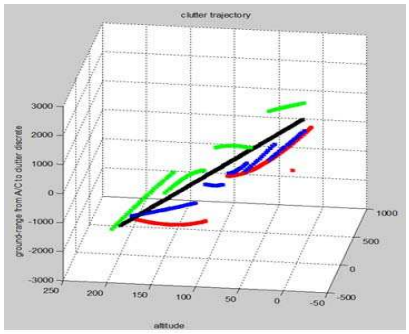

Figure 7. Typical clutter trajectories (ground truth is black).

After we create an encounter scenario in the AFRL simulation environment, we feed the output into the ABSAA and GBSAA sensor models to create contacts for each sensor type. It is important to stress that our sensor parameter settings are notional, and are intended simply to be representative of actual sensor performance. More details descriptions of sensor characteristics may be found in [11-15].

The airborne EO sensor model generates noisy bearing and elevation measurements, with user specified error statistics. We select performance parameters consistent with ABSAA EO sensor performance. In particular, we used +/- 95 degree FOV in azimuth, +/- 20 degree FOV in elevation, Gaussian angular errors of 1 degree standard deviation in both azimuth and elevation, a 20 Hz sampling rate, a 10 nm detection range, a fixed probability of detection of 0.8, and Poisson distributed thermal false alarms with a rate of 1 per frame. Thermal false alarm measurements are assumed uniformly distributed over the field of view. Ground clutter, which impacts both the EO and radar, is discussed further below. All of the ABSAA EO sensor model parameters are user configurable so that it is straightforward to experiment with different parameters that represent the range of sensors that may become available across UAS. We extend the sensor model to include object size measurements with additive Gaussian noise because we anticipate that advanced EO sensor processing will include some form of feature-aided tracking using size, shape, or appearance features to improve data association and false alarm rejection performance, and inclusion of a size-based measurement with user configured feature measurement statistics will allow us to modulate EO tracker performance to account for that feature-aided processing in a simple way. The size measurement is based on a spherical model for objects with entity-specific radius, with size given by the number of object pixels that project into the sensor focal plane given the target-to-sensor distance.

The airborne radar and ground radar sensor models include noisy bearing, range, elevation, and range-rate measurements with additive Gaussian noise. As with the airborne EO sensor, we selected performance parameters consistent with expected ABSAA radar and GBSAA radar sensor performance. In particular, for the airborne radar, we used +/- 100 degree FOV in azimuth, +/- 20 degree FOV in elevation, Gaussian angular errors of 1 degree standard deviation in both azimuth and elevation, a 1 Hz sampling rate, a 10 nm detection range, a fixed probability of detection of 0.9, and Poisson distributed thermal false alarms with a rate of 10 per frame (these are nominal numbers used for the sake of initial experimentation; the framework is set up to enable parametric analysis over these parameters). Thermal false alarm measurements are assumed uniformly distributed over the field of view. False alarms due to clutter are discussed further below. All of the ABSAA radar sensor model parameters are user configurable so that it is straightforward to experiment with different parameters that represent the range of sensors that may become available across UASs. Similarly, for the ground based radar, we used +/- 180 degree FOV in azimuth, a -5 to 90 degree FOV in elevation (under the assumption that the ground based radar is elevated), Gaussian angular errors of 1 degree standard deviation in both azimuth and elevation, a 0.2 Hz sampling rate, a 100 km detection range, a fixed probability of detection of 0.9, and Poisson distributed thermal false alarms with a rate of 10 per frame. Thermal false alarm measurements are assumed uniformly distributed over the field of view (again, these are nominal numbers to enable initial experimentation).

We model false alarms due to thermal and other noise sources as uniformly-distributed in measurement space. Separately, we model ground clutter in two ways. First, we generate ground clutter tracks within the AFRL framework as discussed earlier. Second, we generate uniformly distributed *tentative* ground clutter measurements in bearing and elevation. Each tentative measurement defines a line connecting the sensor to the measurement and, if the line intersects the earth, which we represent using the WGS 84 ellipsoid, and the intersection is inside the sensor footprint, we define a clutter return (in the case of radar, the range measurement corresponds

to the intersection of the line and the ellipsoid with additive Gaussian noise, and the range-rate measurement is uniformly distributed in measurement space); otherwise, the tentative measurement is discarded.

Tracker performance assessment is a challenging task; many metrics are required so as to describe performance adequately. These metrics are generally coupled; further, even exhaustive lists of metrics do not fully characterize all performance characteristics of interest. We seek a compact set of metrics that quantifies aggregate surveillance performance as well as providing a focused assessment of surveillance performance in the vicinity of UAVs. A difficulty with many approaches to performance assessment is that they rely on an uncertain global mapping of ground truth trajectories and target tracks.

Recently, the *Optimal Subpattern Assignment* (OSPA) metric has been introduced that allows for tracker output assessment without requiring global assignments of truth to tracks [20]. A recent attempt to extend the OSPA metric to labeled tracking outputs is reported in [21].

We take an approach that leverages the OSPA assignments, thus avoiding global assignment decisions, while still establishing track-level performance. Ultimately, we must assess the similarity between two sets of (track-level) objects. We do so by computing quality and purity of each set. *Quality* measures the fraction of track that finds an assignment under the OSPA metric. *Purity* measures the fraction of track that is consistent with the mode (i.e. most frequent) assignment. Thus, the quality of the truth trajectories can be thought of as track-level detection probability. The quality of the target tracks is the fraction of true track. Truth purity is inversely proportional to track fragmentation, and track purity is inversely proportional to track swap occurrences.

Our track-level OSPA metrics rely on a track-truth proximity gate that is set to 1 km. In addition to the purity and quality metrics noted above, we also measure track localization RMSE that is necessarily bounded by the proximity gate. These metrics provide a compact and easily interpretable aggregate performance assessment because they provide an intuitive notion of symmetry in that the metrics are unaffected by interchanging truth and tracks.

In addition to aggregate performance, we are interested in a focused assessment of surveillance performance as it impacts UAS operations. Accordingly, we consider truth quality, truth purity, and track localization RMSE for intruder objects that are close to the UAS. A fourth local metric of interest is the average number of tracks in the vicinity of intruder objects. These constitute false tracks to which the UAS must respond.

Our performance assessment includes the following:

- Five aggregate track-level OSPA metrics, evaluated for each of the 9 tracker outputs (airborne EO, airborne radar, ground radar, airborne fusion, centralized fusion, distributed fusion, enhanced airborne fusion, enhanced centralized fusion, enhanced distributed fusion);
- Four local track-level OSPA metrics, evaluated for each of the 9 tracker outputs.

VI. Simulation Results

We evaluate each of the three track fusion architectures represented in Figures 2-4 both with and without the use of equivalent measurement processing. Snapshots from the 3D scenario are in Figures 8-13. The left panel illustrates the 3D track picture, while the middle and right panels illustrate 2D track pictures projected on to *x*-*y* (ground-ground) and *y*-*z* (ground-elevation) axes. Both 3 sigma detection covariances and track estimate covariances are displayed, with virtual measurements applied to all EO detection displays.

Ground truth trajectories are in black. Target 1 in the figures is the UAV, and target 2 is the intruder. Additional objects are on the ground. Note that there are only tracks on the UAV when GBSAA data is included, since airborne EO and radar data does not include the platform itself. More importantly, redundant and fragmented tracks are generally observed on the intruder, except with good performance (e.g. Figures 14-15).

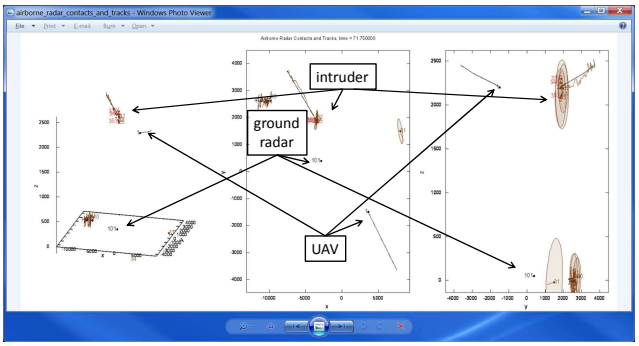

Figure 8. ABSAA radar sensor tracks.

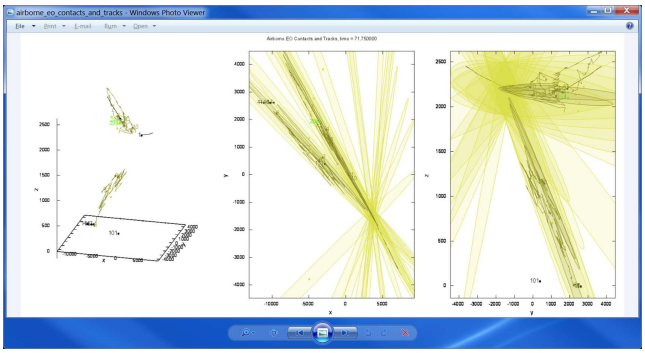

Figure 9. ABSAA EO sensor tracks.

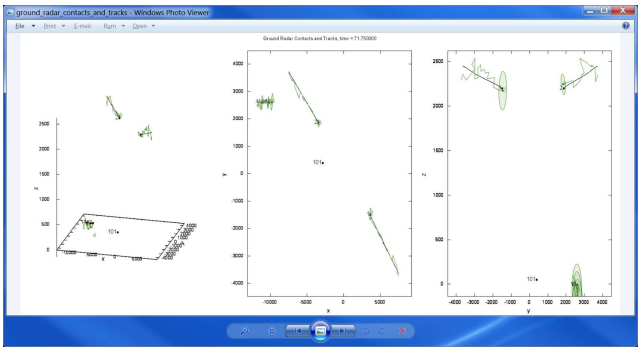

Figure 10. GBSAA radar sensor tracks.

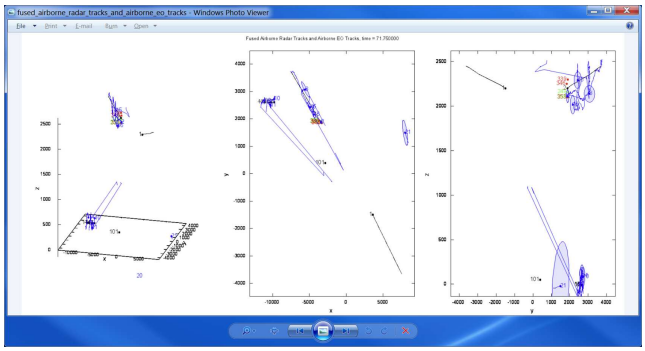

Figure 11. ABSAA fused tracks.

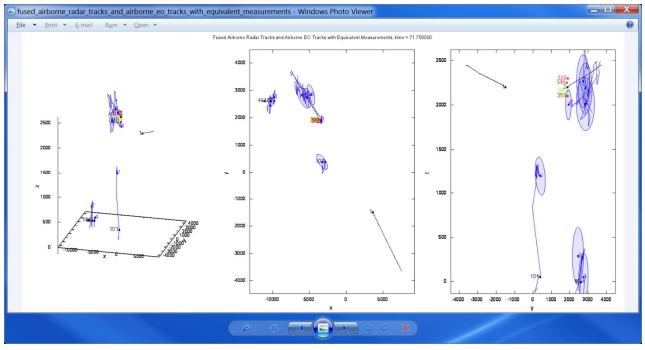

Figure 12. ABSAA fused tracks with equivalent measurements.

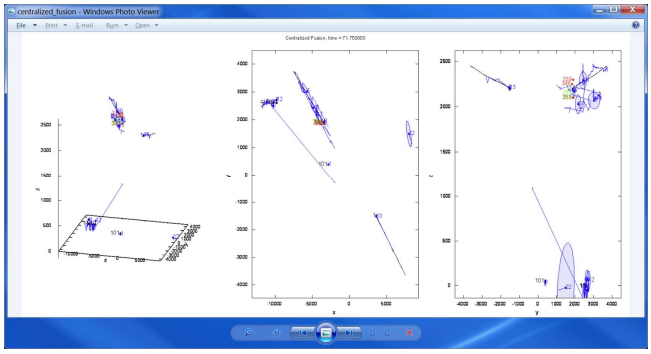

Figure 13. ABSAA/GBSAA centralized (option 1) track fusion.

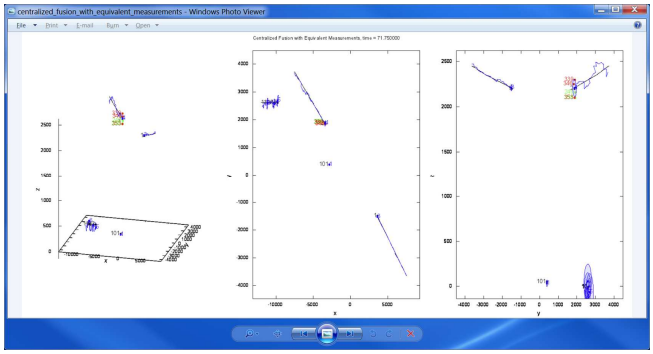

Figure 14. ABSAA/GBSAA centralized (option 1) track fusion with equivalent measurements.

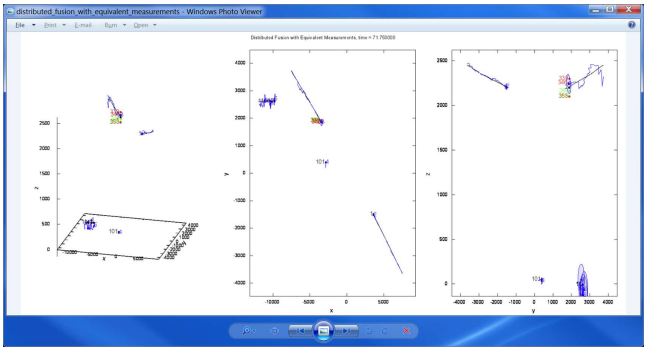

Figure 15. ABSAA/GBSAA distributed (option 2) track fusion with equivalent measurements

Tables 1-2 show the performance according to the quantitative metrics. In these tables, we include all possible tracking architectures, including the use of each of the three sensors alone, ABSAA-only fusion, and ABSAA/GBSAA fusion using each of the two architectures both with and without equivalent measurements. For each metric (shown in the columns), we highlight the best performing architecture. We see that the ABSAA/GBSAA fusion using equivalent measurements is best, with a small improvement in localization error using centralized fusion. We believe that the small decrease in localization accuracy is well worth the more flexible architecture, as this will allow more straightforward integration into existing systems.

Table 1. Aggregate measures of performance.

	Quality Truth	Quality Track	Purity Truth	Purity Track	Localization Error (m)
Airborne EO	0.12	0.14	0.91	0.64	407
Airborne Radar	0.72	0.76	0.98	0.96	99
Standard Airborne Fusion	0.59	0.51	0.99	**1**	97
Airborne Fusion w/Equivalent Measurements	0.51	0.57	0.99	**1**	140
Ground Radar	0.72	**1**	**1**	**1**	81
Option 1 Standard Fusion	0.92	0.65	0.97	1	54
Option 1 Fusion w/Equivalent Measurements	**0.93**	**1**	**1**	**1**	**46**
Option 2 Standard Fusion	0.92	0.66	0.93	**1**	54
Option 2 Fusion w/Equivalent Measurements	**0.93**	**1**	**1**	**1**	57

Table 2. Intruder measures of performance.

	Quality Truth	Purity Truth	Localization Error (m)	Proximity Count
Airborne EO	0.10	0.67	326	4
Airborne Radar	0.86	0.96	88	28
Standard Airborne Fusion	0.97	0.96	93	6
Airborne Fusion w/Equivalent Measurements	0.97	0.96	103	4
Ground Radar	0.72	**1**	46	**2**
Option 1 Standard Fusion	**0.97**	**1**	46	7
Option 1 Fusion w/Equivalent Measurements	0.93	**1**	**38**	**2**
Option 2 Standard Fusion	**0.97**	**1**	39	7
Option 2 Fusion w/Equivalent Measurements	**0.97**	**1**	47	**2**

VII. CONCLUSIONS

The work summarized in this paper demonstrates the applicability of advanced distributed MHT technology to fusion of active and passive sensor data in realistic settings with clear surveillance performance benefits. Results are demonstrated in the context of unmanned aircraft SAA.

Specific conclusions that can be drawn include the following:

- GBSAA sensors are valuable to mitigate ABSAA clutter, which is important to improving landing and takeoff safety;
- There is comparable performance with two fusion architectures, one with single-stage ABSAA and GBSAA sensor integration, and an alternative in which they are integrated in multiple stages;
- Judicious use of virtual measurements in passive sensor tracking provides good-quality passive tracks that are amenable to active-passive fusion;
- Equivalent measurement processing reduces spurious track associations and normalizes sensor update rate for effective clutter suppression from high-rate sensors;
- Sequential EKF processing provides good-performance filtering that is adequate for the nature of SAA measurement nonlinearities.

In future work, we hope to address the following challenges for improved performance:

- Coupled kinematic and feature filtering of EO measurements would improve performance. Indeed, the image-plane object-size process noise is geometry-dependent. Conversely, object size evolution impacts kinematic localization.
- Our use of equivalent measurements can be made more effective by using it on track initiation only. Once a well-localized EO track is established, there is no need to introduce the processing latency associated with equivalent measurement formation. However, appropriate modifications to track fusion logic are needed to avoid undue emphasis on EO clutter tracks. For this, a composite track fusion rule will be beneficial; relevant work along these lines is in [22].
- Closed-loop simulations that leverage representative autonomous avoidance algorithms [11] would provide greater realism and provide a relevant illustration of coupled surveillance and sensor management.
- Our current simulations assume unlimited bandwidth and communications fidelity. System evaluation using representative communications limitations is of interest.

ACKNOWLEDGEMENT

This research was supported by the Air Force Research Laboratory Sensors Directorate under contract FA8650-13-M-1646.

REFERENCES

[1] S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*, Artech House, 1999.

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

[3] D. Reid, An Algorithm for Tracking Multiple Targets, *IEEE Transactions on Automatic Control*, vol. 24(6), December 1979.

[4] 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.

[5] S. Coraluppi, C. Carthel, M. Luettgen, and S. Lynch, All-Source Track and Identity Fusion, in *Proceedings of the MSS National Symposium on Sensor and Data Fusion*, San Antonio TX, USA, June 2000.

[6] C.-Y. Chong, G. Castanon, N. Cooprider, S. Mori, R. Ravichandran, and R. Macior, Efficient Multiple Hypothesis Tracking by Track Segment Graph, in *Proceedings of the 12th International Conference on Information Fusion*, Seattle WA, USA, July 2009.

[7] S. Coraluppi and C. Carthel, Multi-Stage Multiple-Hypothesis Tracking, *ISIF Journal of Advances in Information Fusion*, vol. 6(1), June 2011.

[8] S. Coraluppi and C. Carthel, Advances in Data Fusion Architectures, in *Integrated Tracking, Classification, and Sensor Management: Theory and Applications*, Wiley-IEEE, 2012.

[9] W. Koch, *Tracking and Sensor Data Fusion*, Springer, 2014.

[10] O. Drummond, Track and Tracklet Fusion Using Data from Distributed Sensors, in *Proceedings of the Workshop on Estimation, Tracking, and Fusion: A Tribute to Yaakov Bar-Shalom*, Monterey CA, USA, May 2001.

[11] S. Graham, W.-Z. Chen, J. De Luca, J. Kay, M. Deschenes, N. Weingarten, V. Raska, and X. Lee, Multiple Intruder Autonomous Avoidance Flight Test, in *Proceedings of AIAA Aerospace 2011*, St. Louis MO, USA, March 2011.

[12] Scally et al, Unmanned Sense and Avoid Radar (USTAR), in *Proceedings of AIAA Aerospace 2011*, St. Louis MO, USA, March 2011.

[13] M. Kochenderfer, J. Holland, and J. Chryssanthacopoulos, Next Generation Airborne Collision Avoidance System, *Lincoln Laboratory Journal*, vol. 19(1), 2012.

[14] Stamm et al, "STARS and DASR GBSAA", presented at CNS/ATM conference, Feb. 2012

[15] D. Accardo, G. Fasano, L. Florenza, A. Moccia, and A. Rispoli, Flight Test of a Radar-Based Tracking System for UAS Sense and Avoid, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49(2), April 2013.

[16] A. Poore and N. Rijavec, A Lagrangian Relaxation Algorithm for Multidimensional Assignment Problems Arising from Multitarget Tracking, *SIAM J. Optimization*, vol. 3(3), August 1993.

[17] P. Storms and F. Spieksma, "An LP-based algorithm for the data association problem in multitarget tracking", in *Proceedings of the 3rd International Conference on Information Fusion*, Paris, France, July 2000.

[18] S. Bordonaro, P. Willett, Y. Bar-Shalom, Performance Analysis of the Converted Range Rate and Position Linear Kalman Filter, in *Proceedings of the IEEE Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove CA, USA, November 2013.

[19] S. Coraluppi and C. Carthel, Modified Scoring in Multiple-Hypothesis Tracking, *ISIF Journal of Advances in Information Fusion*, vol. 7(2), December 2012.

[20] 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), August 2008.

[21] Branko Ristic, Ba-Ngu Vo, Daniel Clark, and Ba-Tuong Vo, A Metric for Performance Evaluation of Multi-Target Tracking Algorithms, *IEEE Transactions on Signal Processing*, Vol. 59(7), July 2011.

[22] W. Blanding, P. Willett, Y. Bar-Shalom, and S. Coraluppi, Multisensor Track Management for Targets with Fluctuating SNR, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 45(4), October 2009.