

Multi-UAV Collaborative Sensor Management for UAV Team Survivability

Craig Stoneking, Phil DiBona, and Adria Hughes
Lockheed Martin Advanced Technology Laboratories
3 Executive Campus, 6th Floor
Cherry Hill, NJ 08002
{cstoneki, pdibona, ahughes}@atl.lmco.com

ABSTRACT

Collaboration among a team of unmanned sensor platforms can provide significant operational advantages through improved situational awareness (SA). Recent work on the Army Aviation Technology Directorate (AATD) sponsored Survivability Planner Associate Rerouter (SPAR) program, as well as separate internally funded research and development (in parallel with the SPAR contract) has provided insights into the challenges related to managing collaborative sensing in support of survivability of a team comprising manned aircraft and multiple sensor-bearing UAVs. This paper will discuss technical challenges related to multi-UAV collaborative sensor management, including sensor resource allocation, sensor platform positioning for collaborative sensing, and integration of collaborative sensing behavior into a comprehensive multi-UAV control system. The paper will also discuss recent, ongoing, and planned investigations into approaches for addressing these challenges.

INTRODUCTION

On a number of externally and internally funded programs, Lockheed Martin Advanced Technology Laboratories (ATL) has studied the problem of dynamic sensor management for situational awareness (SA). In particular, working under subcontract to Lockheed Martin Systems Integration-Owego on the Army Aviation Applied Technology Directorate's (AATD's) Survivability Planner Associate Rerouter (SPAR) program, ATL had the opportunity to examine the problem of managing sensors across multiple unmanned platforms for collaboratively maintaining shared situational

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE AUG 2006		2. REPORT TYPE		3. DATES COVERED 00-08-2006 to 00-08-2006	
4. TITLE AND SUBTITLE Multi-UAV Collaborative Sensor Management for UAV Team Survivability				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Lockheed-Martin Advanced Technology Laboratories,3 Executive Campus, 6th Floor,Cherry Hill,NJ,08002				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 13	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

awareness in support of a single, collective mission. As unmanned vehicles become more common and more autonomous, the drive to develop systems with capabilities, and consequently overcome design challenges, similar to SPAR is sure to increase. Using the SPAR program as an example, the remainder of this paper will discuss a number of challenges related to sensor management system design that we feel will be generic to many future systems like SPAR.

THE SPAR PROGRAM AND THE NEED FOR SITUATIONAL AWARENESS

The aim of the SPAR program is improving survivability, and accordingly, mission effectiveness, of a team of aircraft comprising a single manned rotorcraft and multiple (on the order of four) unmanned rotorcraft. The program focuses on improving team survivability through two major thrusts: improved threat lethality prediction and autonomous collaborative behavior. Figure 1 depicts the benefits of improved threat lethality prediction. In evaluating a route through an area with some number of hostile threats, the risk posed to each aircraft by each threat is calculated along each




Figure 1. Accurate prediction of a threat's lethality envelope allows survivable routes to be planned through areas that would previously have been considered too hazardous.

segment of the route. A simplistic approach of using a standard “lethality radius” around each threat to establish regions of impassibility for the aircraft may too often result in the inability to identify a reasonably survivable route that supports the mission. At the very least, the simplistic approach impairs the ability to accurately balance risk to the aircraft with efficiency in executing the mission. By more accurately modeling the lethality of a given threat to the given aircraft, taking into account the particulars of target aspect, terrain, and other environmental factors, the true risk to the aircraft can more accurately be estimated at any point along a proposed route, providing a capability for making better route planning choices that reduces risk to the aircraft while increasing mission effectiveness.

Figure 2 illustrates the benefits of collaborative behavior, as envisioned in the SPAR program. In the depicted example, two unmanned team members each bring their onboard sensors to bear on a threat and transmit their independent target location estimates to the manned aircraft. By fusing the sensor data from multiple sensor platforms, the manned platform may achieve an engagement-quality location estimate on the target and launch a weapon, from a standoff location, to neutralize the threat.




Figure 2. Coordinated sensing, engagement, avoidance and countermeasures deployment enables a comprehensive, cohesive, team-oriented response to threats with reduced per-platform equipment loading.

By collaboratively applying the sensors, countermeasures, and weapons carried by the group, the survivability and lethality of the team is greater than that of a collection of non-collaborating individuals. Moreover, because one team member may deploy assets on behalf of other team members,

the survivability of each individual may be maintained, while decreasing the necessity for each individual to carry a full suite of sensors, countermeasures, and weapons.

To assess a threat for lethality and then avoid, counter, or engage it, it is first necessary to detect the threat and identify and locate it with sufficient accuracy to support assessment and the response to it. That is, situational awareness is key to enabling the threat lethality prediction and collaborative behavior that are core to SPAR capability.

APPROPRIATE SITUATIONAL AWARENESS AND SENSOR MANAGEMENT

A solid premise of sensor management is that, in any system like SPAR, operational situations will arise (when you need SA the most) in which the set of available sensors is not sufficient to provide total situational awareness. However, a mitigating premise is that *appropriate* SA is likely to be something less than total. For instance, while a SPAR team may require a precise location on a threat that it intends to engage, it may require somewhat less precision for threats that it intends to avoid. Possible threats within surface-to-air missile range must be identified with a certain level of confidence to support threat assessment. Objects beyond maximum threat range may not need to be identified at all. Indeed, too much unneeded SA data can be harmful—consuming computational and communications resources. Good sensor management is needed to ensure that the situational awareness needs of the current operational environment are best satisfied with the available sensors.

The first step in allocating sensor resources to fulfill the situational awareness information needs is to identify what those information needs are and assign to each a priority. An actor's situational awareness information needs derive from that actor's operational context, which may include the mission plan, current stage of plan execution, and current threat picture. Situational Awareness (SA) information needs are associated with either a known, existing battlefield object, or a search volume,

in which additional battlefield objects may be discovered. Battlefield objects may be known *a priori*, specified in the mission plan, or they may be objects that have been discovered as a result of a sensor search. Search volumes may be associated with situational awareness regions that are either plan-centric—specified relative to specific portions of the plan route—or actor-centric—specified relative to the actor and moving with the actor—as plan execution progresses. Figure 3 illustrates plan-centric and actor-centric situational awareness regions.




Figure 3. Situational awareness programs can be actor-centric or plan-centric.

INFORMATION NEEDS ASSESSMENT

In recent work on a number of Department of Defense-funded and internally-funded programs, ATL has developed an automated capability for deriving plan-based information needs from a digital mission plan representation [1]. Using the mission plan, SA regions are established. Different SA regions may be established relative to either the actor or the mission route according to differences in information needs required within each region. For instance, within a general Threat Awareness region, the actor may require knowledge of the existence of threats, with only a general idea of their locations. Within the Threat Awareness region, a Mission Impact Awareness region may be defined, encompassing the mission route, in which the location and identification of threats needs to be known with sufficient accuracy and confidence to allow timely and appropriate mission replanning.

Within the Mission Impact Awareness region, a Survivability Awareness region may be defined, in which the highest accuracy location and highest confidence identification of threats may be immediately necessary for the survival of the actor.

Within each SA region, each type of resource tasking object (RTO) is assigned a set of SA performance objectives that specify the desired accuracy or confidence in sensed parameters such as range, bearing, range rate, crossrange rate, and ID. For purposes of assigning performance objectives, an RTO type takes into account not only the taxonomical identification of the object (tank, truck, air defense unit, etc.) but also the current behavior of the RTO and the actor's current intended response to the RTO. For instance, the SA performance objectives for an air defense radar in search mode that the actor intends to avoid will be different from the same radar, in acquisition mode, that the actor intends to engage. Throughout mission plan execution, the actual values of SA performance parameters, achieved through fusion of available sensor data, are compared to the values of SA performance objectives for each RTO. Wherever actual SA performance fails to meet an SA performance objective, an SA information need is identified.

While the need to establish and use SA regions and SA performance objectives can be identified as abstract requirements of any sensor management solution, the logic for specifying SA regions and performance objective assignments is highly domain specific and reliant on knowledge from subject matter experts. As such, it is not amenable to abstract formalization.

Given that the set of available sensor resources may not be sufficient to satisfy all identified SA information needs once the total set of current information needs has been identified, they must be prioritized to identify where sensor resources would be most profitably spent. The factors taken into consideration for establishing a relative priority for a given SA information need will typically include range, relative bearing, threat type, threat mode, intended response to the threat, etc. For in-

stance, the need for more accurate range and bearing on a nearby air defense unit in acquisition mode is likely to be of higher priority than a need for greater ID confidence on a potential threat that is further away and not emitting. However, the exact logic for assigning priorities to SA information needs is also highly domain specific and reliant on knowledge from subject matter experts. As such, it is also not amenable to abstract formalization.

SENSOR RESOURCE ALLOCATION/TASKING

Once the list of current SA information needs has been prioritized, the available sensor resources can be allocated, or tasked, to satisfy the optimal subset of needs from the list. In searching for the optimal sensor tasking solution, the solution space can be large and highly complex. Different sensors are better at satisfying different information needs; for example, a radar gives better range than azimuth on a metal object, while a Radar Frequency Interferometer gives good azimuth, but no range information on radio emitters. Each sensor may have multiple available modes with different capabilities and characteristics. A single sensor may be able to operate in multiple modes simultaneously or have modes that are mutually exclusive. The time to switch between modes may be significant and variable. Switching modes may cause a loss of data (i.e., tracks). A sensor mode may behave differently when another sensor is present and operating. A sensor may have variable latency and accuracy depending on the number and nature of sensed objects. As the number of resource tasking objects and the number of resources (sensor modes) increases, finding an optimal sensor tasking solution quickly becomes prohibitively costly to compute. If it is necessary to plan sensor tasking over a period of time in excess of a single sense-plan cycle, the problem is compounded even further. Initial investigation suggests that the most realistic approach may be to choose a relatively simple set of

reasonable rules of thumb and iteratively improve on them based on evaluations by subject matter experts through modeling and simulation.

SENSOR PLATFORM POSITIONING

If the SA performance objectives for a given RTO cannot be met with the given set of sensors in the currently planned locations, it may still be possible to satisfy the performance objectives by altering the platform's maneuver plan to place one or more sensors in a better sensing location. A particularly good example is the use of passive, angle-only sensors to geolocate an RF transmitter such as an enemy anti-aircraft radar. Geolocation of the transmitter can be done by fusing multiple bearing measurements, from either a single moving sensor platform, over time, or from multiple platforms separated by some distance (Figure 4). In either case, the achievable geolocation performance for a given RTO is reasonably predictable but is sensitive to the relative geometry of the sensor measurements with respect to the target. Figure 5 illustrates the variables involved in geolocation of a target using angle-only measurements. The Target Location Error (TLE) is highly sensitive to the angle between the measurements (θ), the angular error in each measurement (σ), and the distance from each sensor to the target (r). Consequently, a single sensor platform can get the best results by




Figure 4. Emitter location can be derived from angle-only measurements taken from a single, moving platform over time or from multiple spatially separated platforms.




Figure 5. Variables in target geolocation from angle-only measurements.

making a close pass to the side of the target, rather than directly towards it, at a high speed. This reduces the crossrange error resulting from the measurements' angular errors and the travel time required to achieve the desired measurement angle separation (θ). Multiple platforms would need to be coordinated in their sensing positions to achieve similar geometric sensing advantages.

Adding sensor platform positioning to the sensor management problem dramatically increases the dimensionality of the solution space. With each sensor platform able to maneuver in X, Y, and Z dimensions, the dimensionality of the total solution space is increased by three times the number of sensor platforms. Because positioning the sensor platforms requires planning on a timescale that is large compared to the sensing-planning cycle, including time as a dimension becomes mandatory.

MISSION SYSTEMS INTEGRATION ISSUES

The problem of collaborative sensor management is complex enough on its own, but it is magnified, yet again, when one considers how the sensor management subfunction will be integrated into a larger, comprehensive multi-UAV mission management system. On SPAR, we have been confronted with a number of integration issues that impact sensor management design; we expect these will be generic to any multi-UAV mission management system.

Mediation of subfunctions in conflict is needed. Any team of collaborative autonomous vehicles will have to balance the behavioral needs of multiple subfunctions: threat avoidance, collision avoidance, terrain following, weapons deployment, defensive countermeasures, sensing, communications/datalink management, etc. For example, the goal of satisfying SA information needs is generally furthered by bringing the sensors closer to the threat, while the goals of threat avoidance and countermeasures deployment are generally better served by keeping the sensor platforms farther away from the threat. If any modularity of subfunctions is to be maintained, some common “goodness” valuation/normalization scheme will have to be found that allows tradeoffs between subfunction plans. Any individual subfunction may be expected to evaluate its own subfunction plan for the degree to which the subfunction plan meets its own goals. In order for conflicts to be resolved, however, there will have to be some metric that can be applied across subfunctions that allows one subfunction plan, in whole or in part, to be compared against the plan from another subfunction to determine which should be adopted, and which rejected or modified.

The design of the subfunction mediation algorithm can have a significant impact on the design of the subfunctions. Consider the following three global plan mediation approaches:

- **Global Cost Optimization:** Each subfunction provides a representation of the entire plan solution space in a form that can be combined in a common representation with the plans from the other subfunctions, with costs (using the global mediation metric) associated with each choice represented in the solution space. By overlaying the solution space cost maps from each subfunction, the mediator constructs a global cost map that it can traverse to find a least-cost global plan.
- **Iterative Replan:** The subfunction produces a best-fit or first-fit plan that satisfies its own goals and provides the plan to the mediator. When the plan is in conflict with that of another subfunction, the

mediator may reject the plan, in whole or in part, and direct (hopefully with hints) the subfunction to construct a different plan.

- **Plan Scoring:** The subfunction is little more than a subplan evaluator. The mediator itself constructs a global plan using some gross method and then presents the plan to each subfunction; each subfunction scores the global plan on how well the plan satisfies the subfunction's goals. If the plan scores too low according to one of the subfunctions, then the mediator may make a modification to the plan (with, perhaps, some hint from the subfunction) and ask for another evaluation.

Each of these global mediation strategies places very different demands on the product of the subfunction, which motivates fundamentally different choices in the design of the subfunction algorithm.

One system-level architectural choice which, in particular, has an important influence on the design of each system subfunction is the choice of distributed versus centralized decision making. In the centralized architecture, a single member of the team is chosen to host the executive decision making processing for the entire team. Sensor data collected by each sensor platform is transmitted to the central node. A global plan is produced at the central node, and the relevant portion of the overall plan is sent to each other node for execution. In the distributed approach, each team member hosts the processing that produces the portion of the overall plan that is relevant to that team member. Because interplatform communication is costly, in terms of both time and bandwidth consumption, a distributed architecture discourages collaborative behavior and places emphasis on independent solutions to problems using platform-local resources. If collaboration with another team member is required to solve the problem, the help of the other teammate may be obtained through some form of peer-to-peer negotiation or by appeal to a central mediator. The desire to avoid unnecessary collaboration in the distributed architecture discourages plan optimization across the entire team (since

this requires all information about all team members) but encourages an optimization approach for each individual (since the solution space is much smaller). The potential benefit of the distributed approach is that when little collaboration is required, the overall team solution can be produced very quickly as a number of relatively small solutions computed in parallel. While the centralized architecture is more permissive of a global optimization approach from the perspective of the co-location of team global information, the fact that the optimization problem grows exponentially with the number of team members is a mitigating factor. It is interesting to note that a distributed solution can be adapted to centralized architecture by moving the processing from each team member into a proxy process on the central host and substituting interprocess communication for interplatform communication. However, any solution employing a centralized architecture will be under pressure to optimize for minimum processing time. This is because the communications latency incurred in collecting global team data at the central node and distributing the completed plan to all team members all but guarantees a worse best-case response time as compared to a distributed architecture.

SUMMARY

In our investigations thus far into sensor management design for multi-platform systems, two related major themes arise. First is the recurring discovery that the details of real-life sensors and operational domains make it difficult to design high-performance sensor management algorithms at a level of abstraction that allows them to readily transfer to other sensors and domains. This has become particularly evident in working to design sensor models, information needs generation and prioritization logic, sensor tasking logic, etc. Second is the realization that choices of architectures and algorithmic approaches at the system level have a significant influence on the appropriate choice of algorithmic approach at the lowest level of sensor management subfunction design. This makes it

difficult to design for reuse in future mission systems applications. As a community, we have a great deal of work ahead of us in learning how to properly abstract and modularize the design of sensor management systems for ease of reuse and system integration, but we think it is clear that future demand for increasingly autonomous and collaborative multi-UAV systems will provide both the need and the opportunity to do that work.

ACKNOWLEDGMENT

This material is based upon work supported by the U.S. Army Research, Development and Engineering Command, Aviation Applied Technology Directorate contract number W911W6-04-C-0053.

DISCLAIMER

Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U.S. Army Research, Development and Engineering Command, Aviation Applied Technology Directorate.

REFERENCES

- [1] DiBona, P., Belov, N., Pawlowski, A. (2006). "Plan-Driven Fusion: Shaping the Situation Awareness Process Using Empirical Plan Data," in the *Proceedings of the 9th International Conference on Information Fusion* (Fusion 2006), Florence, Italy, July 10-13.