

Guidance, Navigation, and Separation Assurance for Local-Area UAV Networks: Putting the Pieces Together

Sam Pullen
Stanford University

Jiyun Lee
Korea Advanced Institute of Science and Technology

ABSTRACT

This paper examines the guidance methodology needed to implement networks of autonomously-flown unmanned aerial vehicles (UAVs) controlled by centralized ground stations (GSs). The intended operations would take place within a local area with a diameter of less than 10 km for most applications but potentially as large as 50 - 100 km. UAVs in these networks are envisioned to be potentially quite small and inexpensive but capable of automated flight orientation and stability with guidance updates provided by the GS at 0.5 to 2 second intervals. The GS also provides GNSS differential corrections and integrity information to support sub-meter-level 95% navigation accuracy with 10^{-7} error bounds in the 3 - 10 meter range. Position and timing solutions for each UAV are relayed back to the GS and support both operational (route planning) and tactical (path updating) guidance. This guidance needs to insure safe separation between UAVs within the network and (depending on the airspace used) separation from "out-of-network" UAVs as well as manned aircraft.

The proposed guidance approach centers around "zones of influence" surrounding each UAV that include allowances for navigation error, UAV guidance error, and ground-system guidance error. The amount of error allocated to each error source depends upon the degree of error correlation between each UAV and its neighbors as well as the required probabilities of safe separation that must be maintained. The ground system maintains and updates zones of influence for each UAV within its operational area and guides each UAV it controls to remain within a "zone of operations" to insure that all UAV movements it commands avoid collisions with other vehicles or the ground. This paper provides examples of how this is done and how adjustments are made to reflect changes in navigation performance and the influx of UAVs operating outside the network.

1.0 Introduction to UAV Network Concept

A large number of applications have been proposed for unmanned aerial vehicles (UAVs). Today, some of these applications, such as taking pictures and monitoring particular locations on the ground, have been implemented to a limited degree with remotely-piloted UAVs operated singly in or small groups. However, requiring each UAV to be controlled by a human pilot who must carry out the UAVs mission while monitoring for and maintaining safe separation from other vehicles greatly constrains the number of vehicles that can participate safely and cost-effectively. Obtaining the full benefits of these applications will most likely involve UAVs operating autonomously and coordinating their activities in large groups. This is particularly true for data-collection and monitoring over large areas.

Several examples of applications that are suited for networks of autonomous UAVs are presented in [1]. The simpler of these use the potential cost-effectiveness of autonomous UAV networks to obtain results more quickly and cheaply than existing methods using piloted aircraft or remotely-piloted UAVs. Photography and other forms of passive data collection are good examples, as the data is either not needed in real time or is independent of mission planning. For example, piloted aircraft are used today in major metropolitan areas to monitor road traffic conditions during busy periods. UAV networks could do this job better and more cheaply simply by allowing both denser and more widespread coverage. The outputs of UAV monitoring would be used in near real time to display traffic conditions and warn of accidents and bottlenecks, but the changes to traffic conditions would not require real-time changes in UAV positions or monitoring patterns. Thus, high-level changes to UAV guidance, such as repositioning and UAV air traffic control could be handled by ground personnel, although effective autonomous guidance would be much more cost-effective.

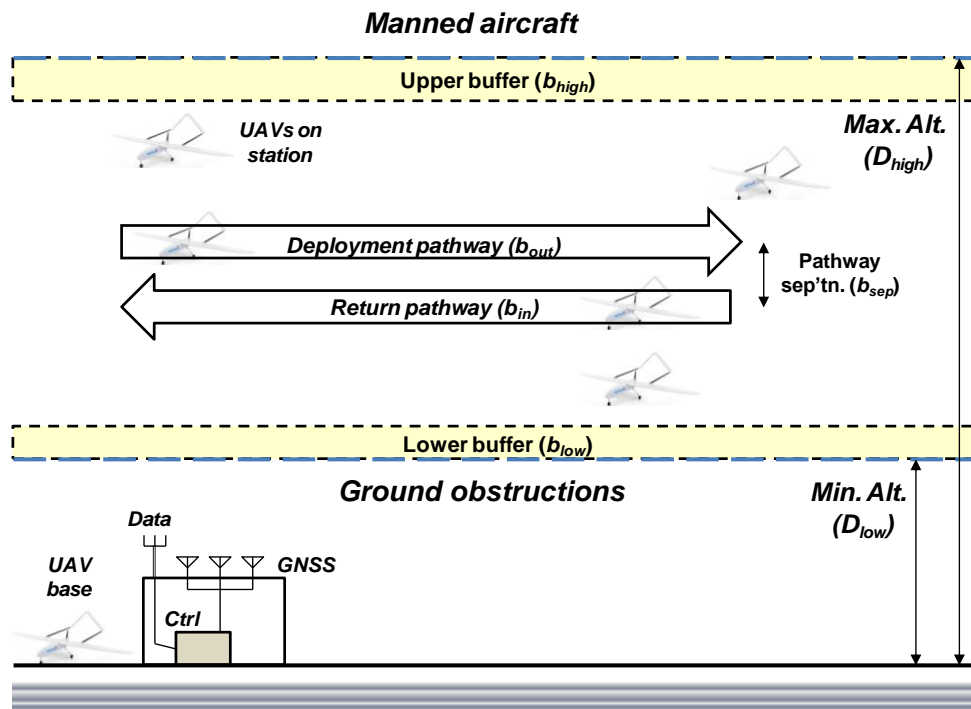


Figure 1: Local-Area UAV Network - Conceptual Diagram [1,2]

Other applications described in [1], such as reconnaissance and surveillance, would place much greater demands on human guidance and would therefore require autonomy guidance for most, if not all, guidance functions. In a surveillance application such as monitoring for potential lawbreaking at a large shopping center or subdivision of houses, UAVs would be spread out in a standard patrol pattern most of the time. When suspicious activity is detected, it would be desirable to reposition several UAVs in real time to get a better view of the area that generated alert to confirm if it is suspicious and, if so, to follow the object of concern until security personnel can arrive. On a small scale, this repositioning could be handled by humans, but as the scale expands to kilometers, most of it will need to be done automatically.

This paper expands on the local-area UAV network concept outlined in [1,2] to explain how autonomous guidance of the type described here can be implemented. Section 2.0 provides an overview of this concept, including the use of local-area differential GNSS (LADGNSS) navigation [1] and the derivation of safe separation standards [2]. Section 3.0 describes how the navigation and guidance error models for individual UAVs are combined with the separation standards to generate multiple "zones of influence" (ZoIs) for each UAV that represent the region around each UAV that must be kept clear to avoid collisions. Section 4.0 describes how the current activities of each UAV are reflected in "zones of operation" (ZoOs) that represent the region within which each UAV is allowed to maneuver while performing a particular activity. Section 5.0

describes how zones of influence and operation work together under nominal ("status-quo") conditions to assure that the overall mission is carried out safely. Section 6.0 describes how these zones are used to manage guidance under off-nominal conditions. Section 7.0 addresses the complications of sharing airspace with other users, including other UAV networks (with separate controllers) and manned aircraft. Section 8.0 summarizes the paper and described the next steps in refining this guidance methodology.

2.0 Local-Area UAV Network Concept

Figure 1 illustrates the concept of local-area UAV network operations developed in [1,2]. The control station shown at the lower left is the source of local-area differential GNSS corrections and integrity information as well as real-time guidance for each UAV in the network. The LADGNSS and guidance information are separate data messages combined in the same outbound transmission. The guidance function also requires feedback from each UAV in real time, including its current position and velocity. Therefore, a two-way datalink is required and is used to relay GNSS information (such as position-domain protection levels) as well as position and velocity from UAVs to the control station. The maximum operational range of a single network of this type is limited by many factors, including the effective range of the datalink and the range beyond which LADGNSS errors grow unacceptably or become too difficult to reliably detect [1].

Most individual UAVs are “on station,” meaning that they are stationary or nearly so and are observing the ground and/or the nearby environment. Under normal conditions, UAVs on station are spread out enough to ensure safe separation among themselves, manned aircraft, and ground obstructions while providing enough density to perform the observation mission assigned to them. Each UAV “flies itself” in the sense that it is responsible for its own attitude control, station-keeping, and path-following (when the GS commands a movement to a new location).

Several events might require one or more UAVs to divert from their normal stations. Because the endurance of each UAV is limited, the ground control station (hereafter simply referred to as “ground station,” or “GS”) must recover, refuel, and re-launch UAVs periodically at a site near the ground station. To support this, specific pathways in space are defined to separate deploying and returning UAVs from those on station. These pathways can also be used to transition UAVs from one “on-station” location to another if called for by normal mission operations or the need to reinforce a region where an alert has been generated. While most recoveries are planned in advance based on GS monitoring of the endurance of each UAV, failures of individual UAVs may require sudden, unexpected recovery. For example, a UAV that suddenly loses partial engine power may be able to be recalled along the recovery path if it can still control itself and fly safely. If all power is lost, the GS must determine where the vehicle can land and guide it to land in the safest possible location that is physically reachable. In general, “landing in place” will be the most common form of recovery when safe, guided flight becomes impossible.

Another need for repositioning arises from unexpected degradation of navigation or guidance quality, both of which are monitored by the GS using the UAV-to-ground datalink. If, for example, the LADGNSS measurement quality suddenly degrades significantly on one or (likely) more UAVs, the distances required for safe separation increase correspondingly, and this may require multiple UAVs to autonomously move in such a way as to avoid each other under new, tighter, geometry constraints. The new separation requirements will be reflected in updated, larger zones of influence (ZoIs) for each affected UAV, as described in Section 3.0. The GS must use these revised ZoIs to quickly determine the maneuver that restores acceptable safety with the least overall disruption to the mission (see Section 6.0).

The largest challenge to automated GS guidance is probably the unexpected intrusion of manned aircraft or UAVs not controlled by the GS (called “out-of-network” UAVs, as opposed to the “in-network” UAVs controlled by the GS). The most complex problem is the arrival of multiple out-of-network UAVs in the same operating region as the in-control UAVs. Without the ability to

control intruding out-of-network UAVs, the GS needs some means to gain information about the location and intentions of out-of-network UAVs so that it can execute maneuvers to maintain safe separation in a reasonable amount of time. The same requirement (with a higher level of criticality) exists with manned aircraft, but usually no more than one manned aircraft will intrude at any given time.

This concept of autonomous UAV operations assumes “uncontrolled airspace” in the terminology of air traffic control. That is, each UAV is free to fly around as it chooses or is guided to without being required to follow the commands of external human controllers. In Figure 1, the airspace where UAVs may operate is separated from manned aircraft (including uncontrolled and controlled varieties of manned airspace) by an upper altitude limit. While it is unlikely that manned and unmanned aircraft can be separated completely in this manner, this division allows us to focus on the challenges of autonomously operating many UAVs in the same airspace. It would be almost impossible (and certainly not cost-effective) to attempt to provide real-time human control of multiple non-cooperating UAV networks in the same region and to replicate this everywhere that UAVs fly. If this were attempted, it would need to address the reality that non-cooperating UAVs may attempt to occupy the same airspace and therefore need external supervision to fairly separate them, as opposed to the operators of each system being tasked to allocate airspace in real time.

3.0 UAV Zones of Influence (ZoIs)

3.1 Definition of Zones of Influence

The concept of a “zone of influence” comes from the networked-UAV separation criteria derived in [2]. Figure 2 (from [2]) shows how the UAV and target (potential collision object) errors are defined. The uncorrelated components of these errors combine to define a buffer zone that must remain unoccupied for collisions to remain sufficiently improbable based on the risk assessment methodology in [2] (for other examples of risk-based approaches, see [3,4]). The error sources relevant to a given in-network UAV are (a) navigation error; (b) path-following error; and (c) guidance error. “Navigation error” represents the deviation of the reported position from the true position. “Path-following error” represents the guidance (control) error due to the UAV’s inability to fly its own desired path (the path determined locally by the UAV to follow the GS guidance). “Guidance error” represents the UAV guidance (control) error due to imperfections in GS guidance, including primarily the error in representing the ideal path determined by the GS to the UAV. In practice, guidance error is small and is hard to distinguish from path-following error, but it is defined separately because it affects all UAVs in a

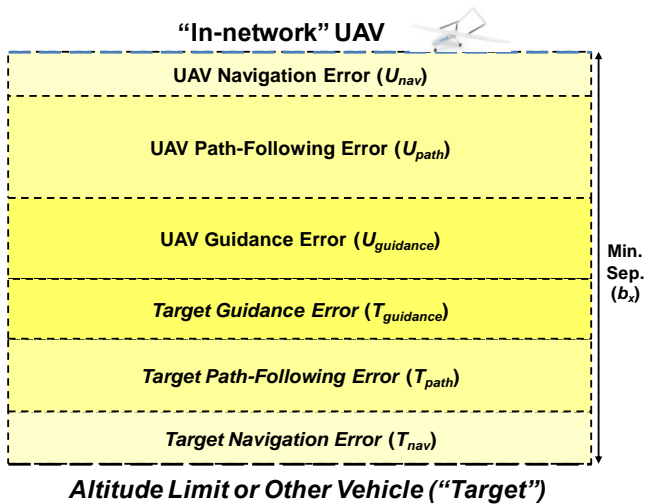


Figure 2: Elements of Required Separation [2]

manner different than path-following error and can be distinguished from path-following error under some circumstances, including sudden maneuvers.

To form a zone of influence for a particular UAV relative to a particular target, the uncorrelated components of the six error sources in Figure 2 are convolved to construct ellipsoidal (3-D Gaussian) error models, which are then extrapolated to the required separation probability for the particular target being considered. These individual errors are typically zero-mean and thus concentric around the estimated position of the UAV, but are not necessarily so. Also, these errors are not represented separately beyond the notional one-sigma level. Instead, conservative one-sigma representations of each error source (i.e., Gaussian distributions that bound rare-event errors using one-sigma values that larger than the nominal one-sigma values) are convolved together to form a single one-sigma bounding ellipse, and this ellipse is extrapolated to bound the separation probability required for a given target.

Since each ZoI reflects collision risk relative to a particular type of target, each UAV actually possesses multiple ZoIs at any given time, and these ZoIs change with time as the UAV and nearby vehicles move. For example, a single UAV "on-station" has a ZoI relative to neighboring in-network UAVs, including those "on-station" and those moving for any reason. It has a separate ZoI with respect to the ground, a separate ZoI with respect to the upper altitude limit, and at least one separate ZoI for each out-of-network UAV that is nearby, if any. These ZoIs are not the same for two reasons [2]. First, the required probabilities of safe separation are generally different for different classes of targets. Second, the error models themselves differ because the degree of error correlation differs. In particular, both navigation and path-following errors are highly correlated

Error Source	Bounding Sigma (meters)
User navigation	0.2
User path-following	0.6
User guidance	0.3
Target navigation	0.2
Target path-following	0.6
Target guidance	0.3
RSS	1.0

Very highly correlated

Highly correlated

Figure 3: Example 3-D Error Budget for Two In-Network UAVs [2]

Error Source	Bounding Sigma (meters)
User navigation	1.5
User path-following	2.0
User guidance	1.0
Target navigation	3.0
Target path-following	4.0
Target guidance	N/A
RSS	5.7

Potentially correlated, but assumed independent

Figure 4: Example Vertical Error Budget for One In-Network UAV Relative to Manned Aircraft [2]

between in-network UAVs in close proximity to each other because GNSS and LADGNSS errors are highly correlated and because disturbances such as wind are also correlated. The degree of correlation that can be assumed with regard to out-of-network UAVs and manned aircraft is lower and depends on the degree to which the GS knows the characteristics of these other aircraft. Finally, non-moving targets such as the upper and lower altitude limits (protecting manned airspace and ground obstructions, respectively) don't have the same error sources and thus require different models (see [2]).

Figures 3 and 4 illustrates this situation using example error numbers from [2] for two scenarios. In Figure 3,

approximate bounding one-sigma errors are given for the separation of two UAVs in the same network. These numbers are very low because each error source is highly correlated between two UAVs that are using the same GNSS, the same differential corrections, the same guidance commands, and are experiencing the same external disturbances (since they need to be near each other for separation to be relevant). Note that the individual ZoIs of in-network UAVs relative to each other are based on the three error sources for a single UAV. The RSS of the three error sources for each UAV in Figure 3 is about 0.7 meters; thus all in-network UAVs have ZoIs relative to each other based on extrapolating from this RSS rather than the RSS of all six error sources (the latter would be "double-counting").

Figure 4, in contrast, shows example correlation-adjusted bounding error sigmas between an in-network UAV and a manned aircraft that might lie just above the upper altitude buffer shown in Figure 1. The errors here are much higher than in Figure 3 because, while there might be significant correlation between UAV and nearby aircraft errors (especially if both use the same GNSS), this cannot be taken for granted; thus error independence must be assumed. Also, the full burden of separation must be assumed by the UAV when the potential target is outside the network; thus the RSS of all five relevant errors shown in Figure 4 must be used to derive the ZoI for the UAV in this case. Since the resulting RSS of 5.7 meters is much larger than the one of 0.7 meters that applies to the case in Figure 3, the resulting ZoI will also be much larger.

Assuming zero-mean Gaussian error models can be relied upon to bound actual errors at the separation probabilities required, the computation of ZoIs for each separation scenario mirrors that of zero-mean (or "fault-free") protection levels for GNSS navigation [6]. In this case, the magnitude of the ZoI in a given position axis can be modeled as:

$$ZoI_x = K_{\text{ffmd}} \Sigma_x$$

where Σ_x represents the bounding one-sigma error in dimension x generated by convolving the individual error sources from Figure 2, and K_{ffmd} represents the scalar multiplier needed to extrapolate a bounding one-sigma value to the required separation probability based on the standard (zero-mean) Gaussian distribution. Applying this equation to the numbers in Figures 3 and 4 and the relevant separation probabilities for these two scenarios from [2] further emphasizes the difference between them. For two in-network UAVs, the suggested probability of 10^{-4} for each UAV gives $K_{\text{ffmd}} \cong 3.8$ from the standard Gaussian distribution, and multiplying this by the 3-D bounding error RSS for each UAV of 0.7 meters gives a ZoI of about 2.7 meters for each UAV in 3-D (meaning

the combination of horizontal and vertical dimensions). The more demanding probability of 10^{-7} for a UAV relative to manned aircraft gives $K_{\text{ffmd}} \cong 5.3$, and multiplying this by the vertical RSS of 5.7 meters from Figure 4 gives a ZoI in the vertical axis of about 30.2 meters (the 2-D horizontal ZoI would be about 25 meters) – an order of magnitude larger than the ZoI between in-network UAVs.

The key to understanding the impact of ZoIs on UAV guidance is that ZoIs between UAVs and all potential collision targets simultaneously and separately apply, rather than the largest ZoI always being the governing one. The ZoI representing separation from manned aircraft or the upper altitude limit (call this "ZoI_{air}") is typically the largest, but it is only a significant limiting factor on UAV guidance for UAVs operating near this upper limit. The same is true of the somewhat smaller value of ZoI with respect to UAVs operating close to the lower limit (call this "ZoI_{gnd}"). The primary constraint on most UAVs not close to either altitude limit will be the ZoI relative to other UAVs ("ZoI_{UAVin}"), particularly out-of-network UAVs ("ZoI_{UAVout}") if any are present.

For any UAV location and position axis, one of the above ZoI will dictate the separation that most constrains the guidance of that UAV. This is called the *limiting ZoI*, and it is the basis for determining Zones of Operation for airspace allocation, as explained in Section 4.0. To clarify this concept, consider the case with no out-of-network UAVs so that three ZoIs apply to each in-network UAV in the vertical dimension: ZoI_{air} = 30 m, ZoI_{gnd} = 20 m, and ZoI_{UAVin} = 2.5 m. If a given UAV is operating near the upper altitude limit, the 30-meter value of ZoI_{air} is the limiting ZoI for vertical positioning. If, instead, the UAV is far from the upper limit but is instead near the lower altitude limit, the 20-meter value of ZoI_{gnd} is limiting. Otherwise, the limiting ZoI is the 2.5-meter vertical ZoI between UAVs in the same network.

3.2 Real-Time Updating of ZoIs

Conservative "floor" values representing each ZoI that applies to in-network UAVs under nominal operating conditions can be computed ahead of time and used in designing the nominal guidance strategy, including the locations of "on-station" UAVs and the best pathways for launch, recovery, and transition of UAVs between stations. "Pre-optimization" of each local-area UAV network installation is vitally important to insure that normal operations support the desired mission objective without presenting any significant guidance burden to the GS. In addition, all likely anomaly and alerting scenarios can be simulated during pre-optimization to ensure that sufficient margin exists to handle them, both in terms of GS and datalink capability and airspace capacity.

The ZoI "floor" values for each target represent conservative bounds on typical UAV behavior, meaning that it is expected (based on simulation, previous experience, and initial on-site test operations) that these values exceed the actual ZoI values most of the time (preferably at least 90% of the time). These values can therefore be used by the GS in real time unless some condition changes that requires one or more ZoIs to be inflated. As noted above, under nominal conditions, the GS should be able to handle all guidance operations routinely, with lots of margin for contingencies.

One status change that affects ZoI is a degradation in the protection levels that are used to derive the navigation-error component of ZoI. A single UAV may experience this when one or more satellites are temporarily masked by the UAV or nearby terrain. If a problem occurs with GNSS itself, such as the unexpected loss or exclusion of signals from a GNSS satellite, all UAVs using that satellite for positioning (likely all UAVs in the network) will experience similar protection level increases and thus ZoI increases. The GS receives protection level updates from each UAV on a regular basis (at least once every 5 s) and receives an out-of-sequence alert when the protection level jumps by more than a certain percentage or grows to exceed a certain threshold; thus the GS can update its ZoIs quickly and then determine if repositioning is required to regain safe separation (this should be rare unless unusually large increases occur). Note that, while GNSS protection levels are normally computed for the rare-event probability to which safe separation must be assured, the one-sigma value of the protection level is used here to allow convolution with the other guidance errors that make up ZoI.

While GNSS navigation error statistics are relatively easy to monitor using protection levels, statistics of estimated guidance errors are also maintained by the GS to monitor for changes that would affect ZoIs. The UAV-to-ground datalink that provides protection levels for each UAV to the GS also provides estimated true position and velocity as well as the commanded position and velocity vectors derived by each UAV for station-keeping or path-following. The GS can therefore keep track of a moving average of the residuals between desired and achieved vehicle states for each UAV, and it uses this to update the path-following error distribution for that UAV if needed. This takes the form of a threshold test, in which the "floor" value for path-following error is used unless the threshold is exceeded, after which the path-following error sigma is progressively increased to track the test statistic. If the test statistic exceeds a certain upper threshold, the guidance quality becomes sufficiently suspect that recovery and replacement of that UAV is triggered by the GS. In addition, test statistics for each in-network UAV based on the difference between UAV-generated paths and the original GS guidance are derived

and combined to generate a running estimate of the overall guidance quality of the GS.

4.0 UAV Zones of Operation (ZoOs)

As stated before, the high-level goal of GS guidance is to maximize mission performance while maintaining sufficiently low risk of in-network UAV collisions with external vehicles or ground objects. The ZoIs computed and maintained for each UAV support the latter objective. To support the former objective, it is natural to define "zones of operation" or "ZoOs" for each UAV corresponding to the current activity being conducted by each UAV and the approximate location of each UAV, which determines which of its ZoIs are limiting. Each ZoO represents a rectangular volume surrounding (and containing) the limiting ZoIs and represents an allocation of airspace by the GS to the UAV included within it. The ZoO extends beyond the ZoI to provide maneuvering room for the UAV, allowing the UAV to move around to some degree without significantly affecting separation.

The amount of maneuvering room or "margin" that lies between the ZoI and ZoO boundaries is an important design parameter. This additional space is used to respond to disturbances or to optimize the performance of the mission. In addition, margin is needed to allow for delayed reaction of the GS to contingencies that require actions to maintain separation, such as another UAV that "blunders" into the ZoO of a given UAV. As with human air traffic controllers, time is needed both for the GS to decide on the appropriate response and to communicate it to affected UAVs. Increasing this margin eases the guidance and datalink requirements on the GS at the cost of lower airspace utilization, as larger ZoOs (give the same ZoIs) provide room for fewer UAVs in the same volume of space. Note that, when UAVs are maneuvered, it is acceptable for ZoOs to overlap as long as these events are infrequent and as long as the ZoIs inside them are not violated.

Figure 5 shows an example of how ZoOs are placed around limiting ZoIs for a network of UAVs focused on ground observations. In this figure, four UAVs are shown "on station" near the lower altitude limit to be as close as possible to the ground. Two other UAVs are shown in the process of deployment to an "on-station" location and recovery back to base. For the "on-station" UAVs, vertical separation is limited by the nearby lower altitude limit. Horizontal separation is theoretically limited by other UAVs, but this limit is not stressed due to the low density of UAVs "on station" (note that the figure is not drawn to scale, or the difference between vertical and horizontal ZoIs would be greater). Because the deployment and return pathways are deliberately located far from the upper and lower altitude limits, UAVs in these pathways have ZoIs limited only by other UAVs.

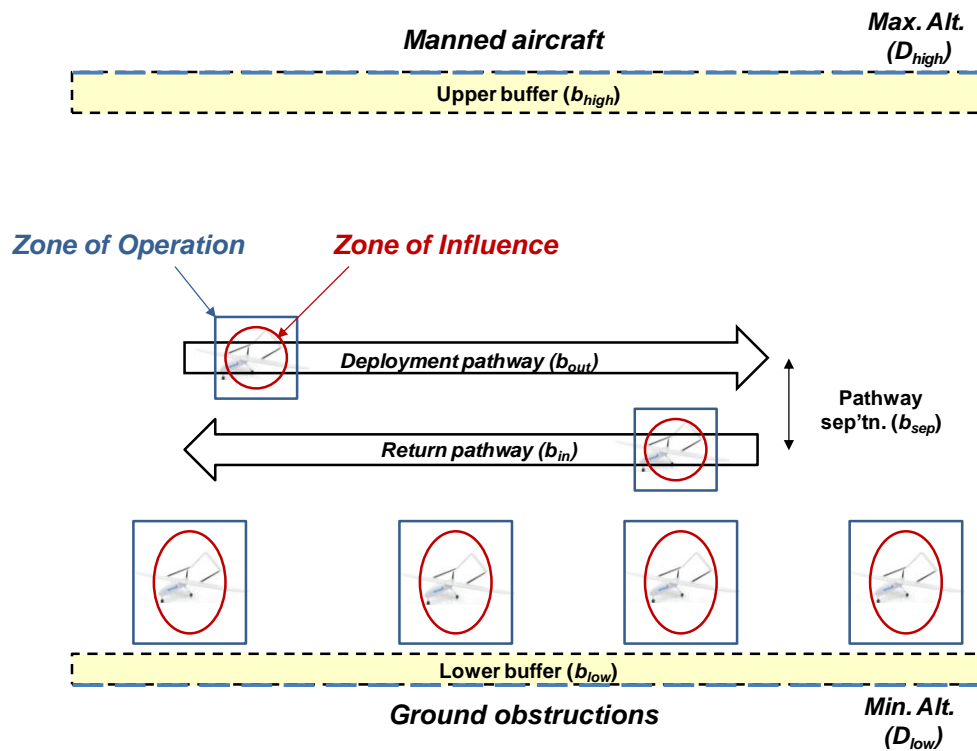


Figure 5: Zones of Operation and Influence for Example Geometry of UAVs under Nominal Conditions

5.0 UAV Guidance Under Nominal Conditions

As noted several times previously, the GS should be able to maintain operations indefinitely under nominal conditions with minimal burden on its algorithms, processing, datalink capacity, and available airspace. The ZoO concept is defined primarily for nominal conditions to make guidance and planning straightforward. Under this concept, as explained in Section 4.0, every individual UAV operation or activity is assigned a ZoO whose volume significantly exceeds the dominant ZoI for the location where the activity takes place. This not only preserves room for maneuvering each UAV around the center of the ZoO without a "repositioning" operation, but it also provides margin in time and space for the GS to react when anomalies or unexpected events occur.

Figure 6 uses the example UAV geometry of Figure 5 and illustrates common vehicle maneuvers under typical conditions. Maneuver (1) is the recall of an "on-station" UAV when it nears the end of its endurance (e.g., battery charge). While sufficient endurance remains, the affected UAV is boosted out of its on-station position and into the return pathway, where it proceeds back to base for refurbishment. In maneuver (2), the remaining UAVs on station are shifted horizontally as needed to cover the gap in observability (if any) caused by the absence of the recalled UAV that was recalled. Maneuver (3) shows the deployment of a recharged UAV from the base to replace the recalled one. Note that, if deployment and recall are

essentially simultaneous, or if the remaining UAVs are sufficient to perform the observation mission, repositioning of the other UAVs is not necessary. In any case, none of these maneuvers should any significant burden on the GS, datalink, or UAV equipment.

6.0 UAV Guidance: Reaction to Unexpected Events

When unexpected (or at least infrequent) events occur, the GS uses the margin built into guidance and planning under nominal conditions to react before the mission or safety are unduly compromised. The burden placed on the GS depends on the severity of the change, whether any advance warning was available, whether one, a few, or many UAVs must be repositioned, and how quickly a response is needed.

As with nominal maneuvers, most foreseeable faults and anomalies should not immediately lead to hazardous conditions. Figure 7 illustrates this with two examples of faults that can almost always be resolved without any threat of loss of separation. The first fault type is an unexpected loss of the primary (GS to UAV) datalink. Unless the UAV has internal telemetry that indicates a fault on-board the UAV, the reason for this loss is unknown at the time that it occurs. As long as the UAVs were positioned appropriately before this datalink loss, the wisest course of action is to "hold position" and wait a short, pre-defined interval for the datalink to be restored. Note that datalink loss prevents the GS from commanding

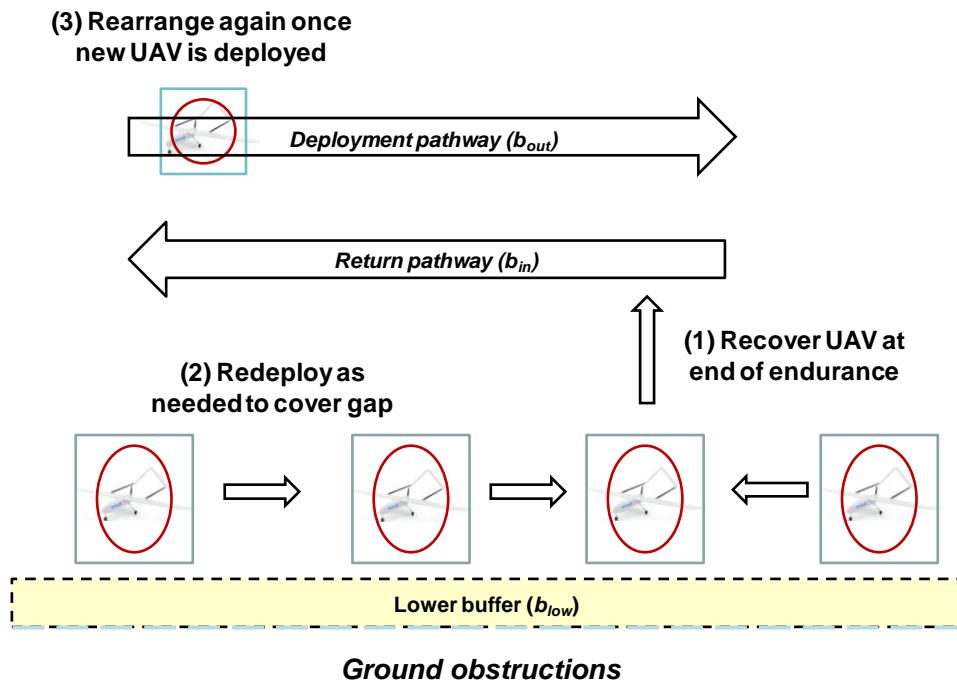


Figure 6: Typical UAV Guidance Actions under Nominal Conditions

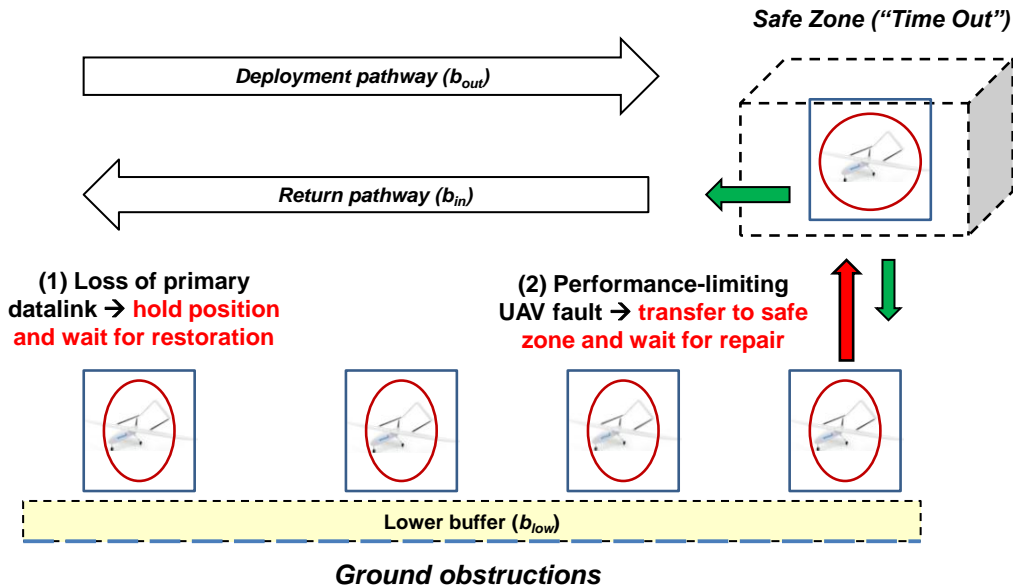


Figure 7: Example UAV Guidance Actions under Anomalous Conditions

the UAV; so the UAV must have this contingency pre-programmed. The most common cause of datalink loss is intermittent RF interference; thus waiting a few seconds for restoration will usually resolve the problem.

The second fault type includes faults on individual UAVs that limit performance such that the pre-existing ZoI and ZoO are no longer valid but do not make the device unflyable or uncontrollable (note that a datalink loss due to a UAV failure is an example). In this scenario, the

primary recourse is to have the GS (or the vehicle itself, if the datalink is down) command the affected UAV to move to a "safe zone" where the UAV can wait for repair to occur. Figure 7 shows this as a designated subset of airspace, but any relatively empty region where the applicable ZoIs are loose enough to provide room for the faulty UAV is fine. The UAV can hold position there for a certain time until the fault disappears, is corrected, or appears to be permanent. If the fault goes away, the UAV can be returned to its former position. If not, the UAV




Figure 8: Example of Severe UAV Fault: Uncontrolled Descent

will eventually be returned to base and replaced with a working UAV.

A few foreseeable fault scenarios are almost immediately threatening and are thus more difficult to handle. One example is shown in Figure 8, where a UAV on station just above the lower altitude limit suffers a crippling failure causing sudden, uncontrolled descent toward the ground. If the GS notices this situation (because the navigation and return datalink functions on the UAV remain functional), it will immediately respond with a command to execute a "guided recovery" to a safe zone or to base. If GS intervention fails but the UAV's local guidance function still works, it will issue this command itself. If both interventions fail, or if the remaining UAV functionality is insufficient to boost the UAV back above the lower altitude limit, whatever guidance is functional attempts to cushion the descent and select a landing site that poses as little danger to people and property on the ground as possible.

In the case of total UAV failure, where the UAV "shuts down" and becomes completely unresponsive, no response will help – the UAV will fall down and impact the ground (or a building) at a location near where the UAV was hovering "on station." This creates a derived reliability requirement: the probability (per unit time) of the UAV entering this unrecoverable state must be well below the allowed probability of violating the lower

altitude limit. Pre-planning by the GS can also mitigate the consequences of this failure by selecting "on-station" positions over areas where UAV crashes would be less dangerous. For example, positioning small UAVs over large buildings with flat roofs likely minimizes the severity of an uncontrolled UAV crash into the ground compared to hitting the ground directly and possibly hitting people or their property.

As illustrated by the sudden-descent scenario in Figure 8, at the onset of an unexpected event, significant separation margin may be lost before the GS can react. The maneuver margin built into the ZoOs is the first line of defense. If it is insufficient, but GS reaction is feasible, the GS will attempt to move the smallest number of UAVs necessary to restore a situation where all ZoOs avoid intersecting (or "overlapping"). If this is not possible, or if the maneuvers required are beyond the computational or coordination capability of the GS, *the ZoOs are discarded* as a basis for emergency planning until it becomes possible to resume nominal operations. Instead, separation guidance is based on a set of *penalty functions* that reflect, in real time, the collision risk implied by the current UAV positions and those forecast to occur under one of several alternative guidance strategies. The development and refinement of these penalty functions is thus key to maintaining safety while allowing operations to proceed and recover to nominal as soon as possible.

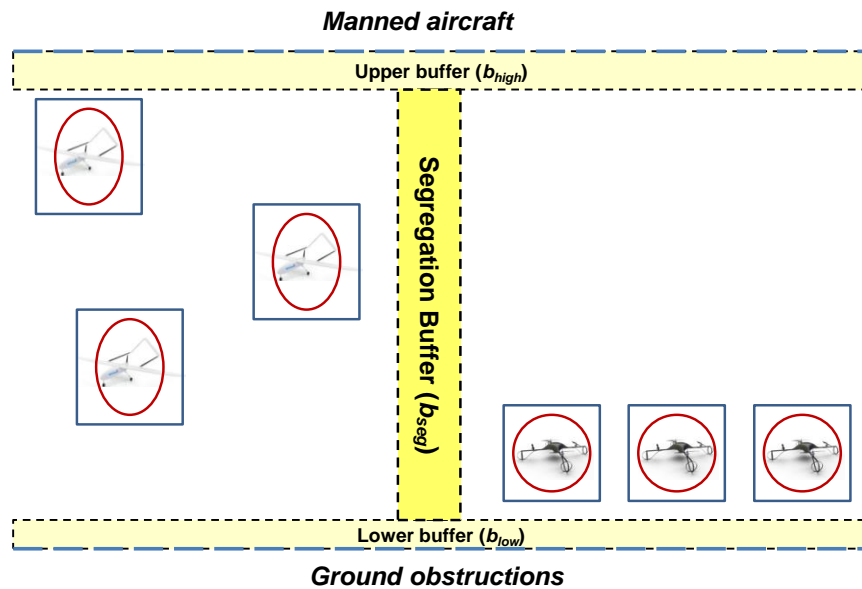


Figure 9: Airspace Segregated Among Different UAV Networks

The concept of "penalty functions" comes from the field of global optimization with "soft" constraints, meaning constraints that can be violated to some degree but at a cost to the objective function that is proportional to the degree of violation (for an introduction, see [5]). Here, the soft constraint is the concept that ZoIs must not overlap or intersect each other, which would imply that the separation standards are (at least temporarily) violated. Since the set of unexpected events that may befall networked UAVs cannot be bounded, it is not possible to ensure that this constraint is met under every conceivable scenario. In other words, as illustrated by the uncontrolled-crash scenario in Figure 8, it is impossible in practice to make the ZoI-non-overlap constraint a "hard" one. Instead, a penalty-function approach (e.g., using polynomials of 3rd to 5th order) properly motivates the GS to avoid violations of the ZoI separation constraint where possible without forcing it to find solutions that prevent all overlaps. Given the multiple operational parameters that need to be evaluated and traded off in real time, no ideal penalty function or set of functions exists. Instead, extensive simulation and (where feasible) testing of foreseeable anomalies and surprises is needed to refine the order and coefficients of each penalty function until the responses close to those that would be chosen by human controllers are achieved.

7.0 Multi-User Airspace Management

7.1 Other UAV Networks ("Out-of-Network" UAVs)

If all UAVs operating in a given region are "in network," meaning that they are controlled and use differential corrections from the same GS, guidance and control are

much simpler than if UAVs operated by multiple, separate controllers are allowed to share the same airspace. Figure 9 shows the situation where this difficulty is resolved by assigning separate sub-spaces of the overall UAV airspace to different users at given times. In it, the right-hand section of the airspace is allocated to UAVs focused on zoomed-in observations of the ground (as in Figures 5-8), while the left-hand section is allocated to a separate network of UAVs spread out more to collect data of a different type. The vertical "segregation buffer" between the two sub-spaces is sized based on a generous (conservative) ZoO that bounds the ZoI for "out-of-network" UAVs relative to each other.

In order to fairly and efficiently sub-allocate the available airspace, external coordination would be required. A simple model would be reserving limited blocks of time in advance, much like tennis courts and golf tee times are assigned. If, for example, 10 different UAV users wanted to access the same airspace "cell" during a given day, an external, independent means would be developed to parcel out the hours of the day among these 10 users so as to maximize overall user utility and "fairness" over the set of users (e.g., minimizing or limiting the "disutility" suffered by the least-favored user). This approach thus gains the relative simplicity of single-user operations in each block of airspace against the restricted benefit that can be obtained if only one user can access a given block of airspace at a time.

Figure 10 shows the more-flexible alternative of sharing the same airspace among multiple UAV users and networks in real time. The short lines coming from each UAV indicate the presence of a new low-power, short-range datalink that allows UAVs to communicate with

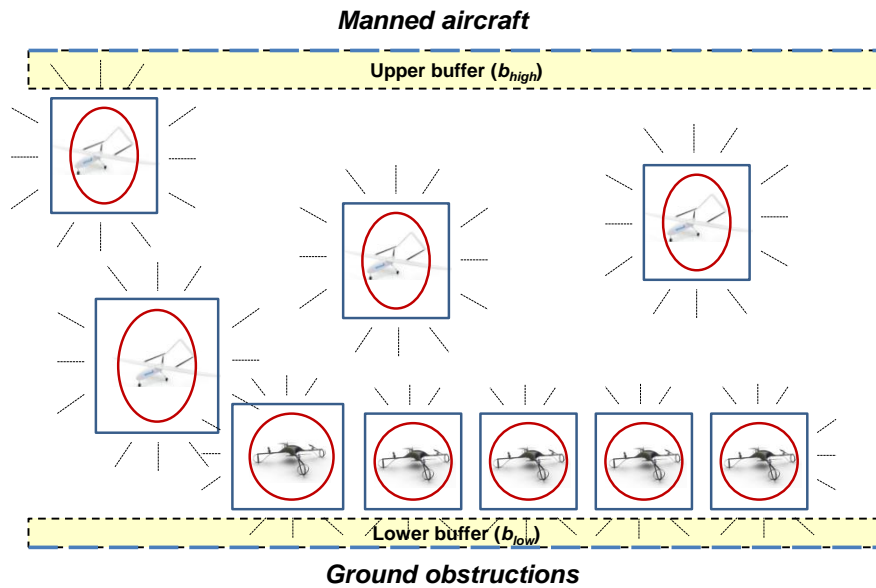


Figure 10: Airspace Shared Among Different UAV Networks




Figure 11: Potential Airspace "Takeover" by Non-Cooperating UAV

each other across networks to share information on position, velocity, and intent (i.e., "I intend to maintain my current velocity for T seconds and then halt."). Information conveyed by this datalink would be downlinked to each GS to allow coordination of UAV maneuvers within the same airspace. In Figure 10, the two networks sharing the airspace are performing different missions (as in Figure 9), which keeps most of their UAVs in separate regions of the airspace and minimizes potential conflicts. However, the lower-left corner of Figure 10 shows two different UAVs operating close to each other and constraining each other's actions. The leftmost UAV of "Type 1" (the type located just above the lower altitude limit) and the nearby UAV of "Type 2" have larger ZoIs and ZoOs than the others because these two UAVs have separation standards dictated by the nearby presence of the other (out-of-network) UAV.

The sharing of airspace shown in Figure 10 is far more desirable than having to operate separately in Figure 9,

but the level of coordination and cooperation required across UAV networks to predictably and fairly share airspace in real time is very challenging. Figure 11 shows an extreme scenario that illustrates the potential problems. In it, a single UAV approaching from the right properly uses the short-range datalink to alert a network of on-station UAVs of its entry into the airspace. Since overlapping ZoOs would result, a cooperative protocol is required to "balance the utilities" of the two UAV networks and thereby determine where the newly-arriving UAV can go (and if it can be allowed to require the pre-existing "on-station" UAVs to move). However, if the arriving UAV is "non-cooperating," it could proceed on its intended course regardless of the guidance determined by the cooperative protocol. The network of on-station UAVs following the protocol would presumably then be forced to vacate their positions to avoid unsafe loss of separation, allowing the arriving UAV to "clear the airspace" for its own use or perhaps to mask activities going on below. Pilots acting in such an unsafe manner would be punished by revocation of their license, but

future UAVs are likely to be more numerous and harder to track and police.

7.2 *Sharing Airspace with Manned Aircraft*

A great deal of research is needed to develop cooperative protocols that are operationally efficient, fair to all parties, and somehow robust to "bad actors" like the intruder in Figure 11. The difficulty of achieving these objectives among different UAV networks illustrates the further difficulty of sharing the same airspace among both UAVs and manned aircraft. Nevertheless, this is the clear goal of the U.S. Federal Aviation Administration (FAA) as it seeks to integrate UAVs into the U.S. National Airspace rather than provide separate or segregated airspace for UAVs [7].

In Figures 10 and 11, a datalink is proposed to allow adjacent out-of-network UAVs to share information. This is partially intended to duplicate the functions of Automatic Dependent Surveillance-B (ADS-B) in manned aircraft without requiring UAVs themselves to be equipped with ADS-B. However, UAVs that operate in the same airspace as manned aircraft will likely need to use ADS-B in some fashion to support the same conflict-avoidance function. The ADS-B equipment available today is too large, too power-hungry, and too expensive for small UAVs, but the gap has been shrinking. One supplier, SageTech, has developed ADS-B transponders intended for UAVs that only weigh 100 grams and have been flight tested on an "Arcturus" UAV [8,9].

For the local-area UAV networks proposed in this paper, the preferred alternative to equipping each UAV with ADS-B would be equipping only the GS with ADS-B. Since the GS is the primary source of guidance, it is the most natural place to receive and decode ADS-B information from manned aircraft ("ADS-B In") and determine if any changes to UAV guidance are required. It can also provide "ADS-B-Out" broadcasts of location information for each UAV in the network. Making the GS the center of ADS-B information avoids the complication of individual UAVs having to transmit ADS-B information back to the GS and wait for a response (if warranted).

8.0 **Summary and Ongoing Work**

Based upon the concept proposed in [1,2], this paper describes an automated guidance methodology for local-area UAV networks controlled by a single ground station. It defines the mission requirements (in terms of continuous observation and real-time response capability) and safety requirements (in terms of maintaining safe separation from other vehicles and ground obstacles) and defines operational procedures to achieve them under nominal and anomalous conditions.

Central to this methodology are the maintenance and updating of zones of influence and zones of operation. Zones of influence (ZoIs) describe the regions around each UAV that must be kept clear to keep the risk of colliding with nearby vehicles and other objects acceptably small. Zones of operations (ZoOs) describe larger regions around each UAV that enclose all ZoIs of that UAV plus additional space that serves as maneuvering room as well as safety margin for failures or other unexpected events.

Under nominal conditions, most UAVs are "on station" making observations, and the GS maintains them with separated (non-overlapping ZoOs) with minimal effort. When anomalies or other unexpected events occur, the GS first attempts to maintain separated ZoOs through a series of procedures that include moving faulty UAVs to "safe zones" where off-nominal performance is temporarily tolerable. For scenarios where this is insufficient, ZoO maintenance is discarded in favor of penalty functions that strongly motivate the GS to keep ZoIs separated (and thus continue to meet separation standards) where possible but allow minor separation violations to occur if necessary. This methodology is illustrated under both nominal conditions and under anomalous conditions that require rapid GS response.

To be practical, this concept must support simultaneous usage by separate UAV networks with different controllers. Two different approaches have been proposed: one in which subsets of airspace are reserved for the use of a single network at a time, and one in which multiple networks use the same space and cooperate in real time to maintain safe operations. The latter is far more flexible but poses many challenges, including handling "non-cooperating" users that may attempt to exploit the safety protocols of the cooperating users to "take over" much of the airspace. Solving these problems for separate networks of UAVs would represent a major step forward to allowing safe sharing of airspace between UAVs and manned aircraft.

The concepts developed here are new and need additional quantification and simulation-based evaluation to judge their practicality. As explained in [1], a key uncertainty is the navigation and path-following performance of different types of small, inexpensive UAVs likely to be used for commercial applications. UAV flight experiments with LADGNSS navigation are planned as a means to evaluate this and thereby improve our estimates of achievable ZoIs for UAVs. Meanwhile, numerical simulations of nominal and unexpected conditions will continue to refine the selection of ZoO sizes (relative to ZoIs) and the penalty function parameters that determine how large of a separation violation is acceptable under various anomaly scenarios. The existing numerical simulation tool will be gradually expanded as new

guidance strategies and complications (e.g., sharing airspace with out-of-network UAVs or manned aircraft) are considered.

ACKNOWLEDGMENTS

The authors would like to thank the John Warburton and his colleagues at the FAA William J. Hughes Technical Center and Jason Burns and his colleagues at the FAA Satellite Navigation Program Office for their support of the authors' research on GBAS. Information on commercial UAV equipment provided by the Unmanned Systems Research Group at KAIST is also greatly appreciated. However, the opinions expressed here are solely those of the authors.

REFERENCES

- [1] S. Pullen, P. Enge, and J. Lee, "High-Integrity Local-Area Differential GNSS Architectures Optimized to Support Unmanned Aerial Vehicles (UAVs)," *Proceedings of ION ITM 2013*, San Diego, CA, Jan. 28-30, 2013. <http://www.ion.org/meetings/abstract.cfm?meetingID=39&pid=119&t=A&s=3>
- [2] S. Pullen, "Managing Separation of Unmanned Aerial Vehicles Using High-Integrity GNSS Navigation," *Proc. EIWAC 2013*, Tokyo, Japan, Feb. 19-21, 2013. http://www.enri.go.jp/eiwac/2013/eng_2013.html
- [3] D. Gebre-Egziabher and Z. Xing, *Analysis of Unmanned Aerial Vehicles: Concept of Operations in ITS Applications*, Final Report (CTS 11-06), Center for Transportation Studies, University of Minnesota, March 2011. <http://www.cts.umn.edu/Publications/ResearchReports/pdfdownload.pl?id=1512>
- [4] C.W. Lum, B. Waggoner, "A Risk Based Paradigm and Model for Unmanned Aerial Systems in the National Airspace," *Proc. AIAA Infotech@Aerospace 2011 Conference*, St. Louis, MO, March 29-31, 2011, AIAA-2011-1424. http://www.aa.washington.edu/research/afsl/publications/lum_risk_for_uavs_2010_2.pdf
- [5] K. Bryan, Y. Shibberu, "Penalty Functions and Constrained Optimization," Dept. of Mathematics, Rose-Hulman Institute of Technology. <http://www.rose-hulman.edu/~bryan/lottamath/penalty.pdf>
- [6] S. Pullen, "Augmented GNSS: Fundamentals and Keys to Integrity and Continuity," Tutorial Presentation, *ION GNSS 2012*, Nashville, TN, Sept. 17-21, 2012. http://www-leland.stanford.edu/~spullen/ION12_tutorial.pdf
- [7] "Unmanned Aircraft (UAS) - Questions and Answers," U.S. Federal Aviation Administration (website), Last Updated March 19, 2013. http://www.faa.gov/about/initiatives/uas/uas_fa/
- [8] SageTech Corporation website, <http://www.sagetechcorp.com/unmanned-solutions>
- [9] Arcturus UAV Aircraft website, <http://www.arcturus-uav.com/aircraft.html>