Aerial MANETs: Developing a Resilient and Efficient Platform for Search and Rescue Applications

William H. Robinson[1] and Adrian P. Lauf[2]
[1]Department of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, TN, USA
[2]Department of Computer Engineering and Computer Science, University of Louisville, Louisville, KY, USA
Email: william.h.robinson@vanderbilt.edu; adrian.lauf@louisville.edu

Abstract—The ability of first-responders to react to the aftermath of natural disasters depends heavily on receiving accurate, real-time data about the structures that may have been affected. Because transportation infrastructure may be unusable, aerial assessments are the gold standard by which such assessments are performed. The advent of mobile ad-hoc networks (MANETs) and autonomous aircraft represents a unique opportunity to allow for rapid response, while minimizing the cost of deployment and increasing reliability and operator safety. This paper describes the key challenges to implement fault-tolerant and efficient deployments of collaborative autonomous aircraft to increase operational reliability and performance when performing aerial sensing and assessment. Some challenges are affected by mobility, such as wireless communication, group navigation, and data collection. Security also represents a challenge during the operation of the MANET. We consider the effects of limited resources (e.g., real-time processing power, battery packs) available on the aircraft. By understanding both the application context and the resource availability, networked aircraft can reorganize to ensure resiliency for the mission if a resource failure occurs within the network.

Index Terms—Mobile ad-hoc networks; Fault tolerance; Robust and secure wireless communication; Energy efficiency; Unmanned Aerial Vehicles (UAVs)

I. Introduction

The advent of mobile ad hoc networking within the scope of embedded system devices has brought a new realm of opportunity in the field of aerial search and rescue applications. Typically in response to natural disasters (e.g., floods, tornadoes, fires, and earthquakes), search and rescue applications rely upon up-to-the-second data gathered from as large an area as possible, utilizing a broad range of sensory equipment and data-gathering techniques [1]-[6]. Furthermore, in an effort to understand and predict natural phenomena, detailed assessments of affected areas can be used by the scientific research community to understand the formation, scope, and progression of disasters affecting human populations.

The response time of search and rescue personnel to a natural disaster is a key in saving the lives of those in the affected areas. Because the structural integrity of buildings and utility lines (e.g., electrical or gas) may threaten those who are either injured or displaced by weather phenomena, the quick assessment of a post-storm situation is critical to saving lives and reducing injury. Traditionally, search and rescue is performed as a cooperative effort between manned aircraft, such as helicopters, and ground-based personnel with vehicles and search animals. Coordinating such attempts can be costly and difficult to accomplish. Such difficulties can hamper the effectiveness of search and rescue crew, putting individuals in danger, and challenging efficiency.

A significant obstacle to existing methods of assessment, particularly in the aerial component, is the presence of elements that are: (a) prone to human error and (b) composed of elements that present the potential for single points of failure [7]-[9]. In the case of tornado disaster assessment, in which a helicopter is often used as a primary flight tool, the vehicle presents a single point of failure that may lead to inaccurate or incomplete information, which can affect the outcome of search and rescue applications. Reducing this single-point-of-failure component is vital to ensuring reliable aerial assessments and reducing the operational costs (e.g., helicopter-based assessments can costs many thousands of dollars to accomplish, per event). By implementing resilient (i.e., fault-tolerant), distributed assessment methods, the single-point-of-failure issue can be effectively resolved, adding the capability for faster assessments due to the inherent advantages of parallelism.

Fault-tolerance is found in multiple aspects of MANET systems. For example, the implementation of the Multi-Path Transmission Control Protocol (MPTCP) allows compliant clients and servers to issue multiple connections that permit fault-tolerant communication in the case of link failure, allowing data to seamlessly operate between interfaces [10]. In such cases, wireless communications may see improved handling of adverse operating conditions, improving recovery time and throughput. However, we still do not see direct benefits from MPTCP at the application layer, although transmission is indeed improved. For this reason, we propose the introduction of a fault-tolerant platform that leverages protocol-based resiliency with a framework that

Manuscript received March 12, 2013; accepted April 18, 2013
adrian.lauf@louisville.edu
This work was supported in part by the NSF under Grant No.CCF-0424422
doi:10.12720/jcm.8.4.216-224

uses a system-wide perspective to provide a fault-tolerant assessment platform.

The organization of this paper is as follows. Section II gives an overview of challenge of aerial reconnaissance technology. Section III describes key principles of mobile ad-hoc networks (MANETs). Section IV presents several key challenges for incorporating MANETs into search and rescue applications. In Section V, a potential research testbed is outlined and evaluated. Section VI summarizes the paper.

II. Background

MANETs have been proposed for deployment in a variety of settings [11], [12]. However, the deployment of airborne networks represents an opportunity to assist with search and rescue operations in times of a disaster. Aerial communications are still subjected to the potential for faults that can limit the performance of the network.

A. Assisting First Responders

Currently, the fastest and most effective way to assess a large-scale impact produced by natural phenomena is from the air. Because an aerial perspective provides breadth as well as depth, the employment of a helicopter is the most-commonly-used method by local disaster response agencies in the United States [1]. Aerial assessments allow first-responders to understand a variety of situations that may be occurring on the ground, including: (1) the number of structures that have been affected by the event, (2) the extent of the damage to these structures, (3) the state of access roads, and (4) the potential number of people that may be affected. Understanding the extent of damage can allow first-responders on the ground to prioritize their efforts based on those structures that exhibit the greatest need of investigation to determine how many survivors may be present. It also allows them to assess whether or not the structure poses a hazard to individuals nearby. Lastly, aerial assessment provides first-responders with the ability to understand whether or not an area presents a risk to the rescuers themselves, which could compromise the efficiency and effectiveness of their intended operation.

B. Origin of Anomalies Within the Network

A fault stems often from a natural occurrence (or an unintentional, man-made one) that affects networking capabilities of one or more nodes, the operation of hardware and/or software on one or more nodes, or a combination of the two [13]. Such faults can be induced by environmental conditions, such as a source of radio frequency (RF) noise, or a naturally-occurring barrier (e.g., mountains) that prevents transmission and reception. Weather-related signal attenuation can also influence the ability of networked nodes to communicate with each other, depending on the distance needed for signal traversal.

Alternately, faults may originate internally, from software programming errors, hardware defects in processing and networking components, non-RF environmental concerns, (e.g., heat affecting the cooling of a critical system component that subsequently fails), and power disruptions, to name a few examples. Both externally and internally-induced fault categories cause many of the same disruptions in networked task execution. For this reason, we are less concerned with the specific nature of the disruptive force, (e.g., which memory cell has failed on a unit of dynamic memory) provided that we can understand that a failure has happened on the network.

III. Mobile Ad Hoc Networks (MANETs)

A key differentiation between mobile ad-hoc networks and infrastructure-based networks is whether or not systems are communicating through a centralized system (in the case of infrastructure). By contrast, ad-hoc networks have no centralized communications system, but instead rely on the ability of nodes (i.e., agents) on the network to disseminate information to the rest of the network devices, if needed [14-16]. Several key issues must be managed effectively within the MANET [17], such as node discovery, route management, and security. The presence of a discovery protocol is critical to the establishment of a communications network in this instance. Routing ensures that all nodes can be reached to communicate the service capabilities. During discovery, nodes advertise their presence and nearest-neighbor relationships are formed. Following discovery and routing, the establishment of security between network connections is necessary to ensure that the connectivity between nodes remains unaffected by external influences. In the event of node failure, task reallocation must be performed to ensure that the network maintains the capability to complete the assigned mission.

A. MANET Node Discovery

Discovery protocols permit networked nodes to broadcast and receive messages pertaining to the presence of a new node on the network, and the capabilities or resources that are available to it. For the sake of coordinating the activities of nodes on our proposed distributed, collaborative assessment platform, discovery is responsible for three main tasks: (1) establishing the presence of a networked node, (2) performing necessary authentication steps to ensure trusted communication, and (3) relaying to neighboring nodes (in our case, aircraft) the resources and associated capabilities with which the individual aircraft may be equipped.

The discovery process itself involves two components: the actual broadcast of nodes to their neighbors to determine their presence, and the subsequent advertisement of available resources, such that distributed processes and tasks may use the resources present on the nodes. Discovery protocols generally employ one of two techniques to address all of the nodes on a MANET. Gossip-based protocols require constant communication between networked nodes to populate lists of available

resources [14], [15], [18]. While this has benefits of no surges in communications when information is required about neighboring nodes, it can cause increased overall bandwidth utilization, and also risks data staleness based on update frequency. In contrast, on-demand methods perform identification only when required, though the increase in network traffic can be exponential if not regulated properly [19], [20]. On-demand discovery does have advantages in ensuring fresh data is made available at the time the request is initiated.

B. Routing

In ad-hoc networks, routing must be determined either on-demand, or by using a polled or gossip-based update method, allowing nodes to understand how traffic should be routed. Routing protocols can be classified as data-centric, hierarchical, or location-based [21]. Both on-demand and active routing discovery methods are significantly more complex than those found in infrastructure networks. Such methods include reactive methods like Ad-hoc On-demand Distance Vector routing (AODV) [19], [20], which focuses on finding optimal routes in highly dynamic systems, as well as table-based routing protocols such as the Ad-hoc Wireless Distribution Service (AWDS) and the Destination-Sequenced Distance Vector Routing (DSDV), which utilize routing tables between nearest-neighbor nodes that are periodically refreshed as needed. Regardless of the method used, routing is central to effective data distribution. It represents a fundamental transmission method, and is second in importance only to physical connectivity such as wireless radios when considering MANET connectivity. For aerial communications, routing must be maintained with the dynamic changes of aircraft flight patterns.

Emerging discovery and routing protocols can now facilitate the implementation of MANETs using simple ARM-based computing platforms. An example of this is the Better Approach To Mobile Ad-hoc Networks (BATMAN) protocol [22]. Because these protocols are easily enabled in mobile Linux distributions, such as OpenWRT, the convergence of MANETs with embedded systems is becoming increasingly prominent and accessible to many different implementations. Section V demonstrates how we choose to integrate BATMAN into a testbed platform.

Because effective routing depends on the algorithm selected, in addition to various parameters that specify how the algorithms are tuned (e.g., timeouts, number of route origination and discovery messages), each implementation must choose its routing and discovery methods carefully. When considering constrained resources, such as battery power on mobile systems, tradeoffs in performance and battery life must be considered. In addition, as longer-range networks may limit available bandwidth on MANET platforms, the overhead required by gossip-based protocols may present a problem when channel utilization nears the theoretical maximum of the channel's capacity.

An additional concern that is manifested when communicating with a MANET is the possibility of multipath routing and communication. Because it is possible to install multiple network interfaces per node, such as IEEE 802.11, Bluetooth, IEEE 802.15.4, and 3GPP-LTE, the notion of using multiple interfaces for routing data implies that multiple interfaces permits additional redundancy in communications. Typically, communication sessions, such as those using the Transmission Control Protocol (TCP), are stateful and maintain their ability to communicate as long as the connection remains stable. Should the connection become unstable, the nodes must then re-establish a new stateful connection, costing time and throughput. A way of solving this problem is to include multi-interface routing techniques, such as those implemented by Multi-Path TCP (MPTCP) and similar protocols [10]. As MPTCP can assume multiple redundant link paths, with a variety of configurations, battery life is leveraged along with the cost of route establishment and overhead.

However, MPTCP does not account for the available bandwidth, latency, and power consumption requirements of the interfaces available on each node. For instance, we cannot send image data over a link with only enough bandwidth for text-based telemetry. To alleviate this particular problem, we can once again turn to cost functions that maximize link availability while reducing the likelihood that data will need to be sent on a link that has insufficient bandwidth.

C. Network Security

Attacks and node failures can disrupt the MANET, causing a degradation or cessation of functionality. An attack challenges some or all of the five basic tenets of network security: (1) data integrity, (2) confidentiality, (3) availability, (4) authenticity, and (5) non-repudiation [23-27]. Because network security that relies on cryptography can only protect against intrusion in a passive manner (i.e., encryption is a deterrent, rather than an active means of protecting a system), encryption serves as a first line of defense against attacks. To say that a breach in communications security might be disastrous to the operation of a group of aircraft that can influence each other's behaviors in flight is an understatement, and thus we must seek additional means to secure the network collective.

To address the shortcomings of passive security, additional layers of detection and response can be added to determine whether or not an attack has taken place, and what must be done to keep the system operational if possible. An Intrusion Detection System (IDS) for resource-constrained, embedded system platforms has been developed with the goal of establishing whether or not statistically-anomalous behaviors were observed [28, 29]. By forming statistical models over time, and using cross-correlation, an IDS can understand whether or not

behaviors (established by requests from device to device) are considered normal for the operation of a system, or deviant from what should be occurring. For instance, a newly-developed standard of inter-aircraft communications and collision avoidance is being tested by the Federal Aviation Administration (FAA) that uses timed beacons to broadcast aircraft position information. Because the interval timing is periodic, a behavioral pattern can be established for the broadcasts that can be assumed to be normal behavior, provided that the information presented within the beacon is unaffected by external influences.

D. Task Reallocation

While the IDS mechanism has the capability of identifying network threats on an ad-hoc network, it is the response to such threats that ultimately preserves the uptime of the system. To this end, redundant or suitable resources must be identified that are present on other networked nodes to be used in place of a failed or compromised network node. When choosing a reallocation strategy, there traditionally exist two differing methods of performing resource discovery – gossip, and on-demand. The two methods differ greatly in their implementation, but can achieve similar results; their unique characteristics make them suitable towards different types of systems.

Gossip-based resource discovery operates on the premise that information about resources, when continuously shared with nearest neighbors, can reduce the need for traversing an entire network in search of a resource [14, 15, 18]. Instead of requiring end-to-end traversal, clustering of information is facilitated by timed, periodic updates between nearest neighboring nodes of available resources. Such a procedure has a primary advantage in that it mitigates flooding – an explosion of network traffic – required by a resource discovery request. Its primary DISADVANTAGE is the inverse of its advantage – the CONTINUOUS networking overhead required by a system that issues periodic broadcasts. In addition, it is more difficult to assure freshness of resource information in a gossip-based system. One approach to task reallocation utilizes a structure called a resource fitness cache [30]. This structure can reduce the number of messages required to find a suitable replacement resource among remaining nodes on the network, assuming that redundancy is present in among the resources. The resource fitness cache has been evaluated within the context of aerial assessment using a configuration of five autonomous aircraft [31].

IV. Challenges for Aerial MANETs

We have identified several key challenges that could hinder the deployment of resilient and efficient autonomous aircraft that collaborate on a mission to perform aerial search and rescue. As with MANETs in general, the mobility of the agents is a challenge, but the speed and distance of the aircraft could pose unique challenges to maintaining the network. Security is also a concern, given recent news reports of military unmanned aerial vehicles (UAVs) that have transmitted data in the clear. In addition, the search and rescue mission must be cost effective; smaller aircraft (i.e., hobby scale) are used which limits the potential payload for sensors, processing, and power for flight time.

A. Mobility

One characteristic of all MANETs is derived from its name: mobility. Mobility adds an increased requirement for maintaining network cohesion, as nodes may be drifting in and out of communication range at all times. This stresses a number of key components; for instance, a node that is barely within range may cause wireless transmission rates to drop drastically in order to maintain communications with an increased presence of noise. The change in data rates can, in turn, cause a failure of dependent tasks on the MANET, since the task may no longer be able to execute in real time with slower peer connectivity. One consequence is the increased requirement for radio transmission power necessary to maintain communications cohesion; this requirement creates a problem for any system reliant on a limited power source. Data collection can be affected when the wireless communication becomes less reliable. If the surveillance relies upon a coordinated effort to cover a grid, then the failure of the group's navigation can also impact the effectiveness of the system.

In an example test configuration using Gumstix Overo embedded computer systems equipped with 802.11 radios, connected nodes were able to communicate to a range of approximately 350 meters in a line-of-sight configuration. Should radio-frequency interference or other conditions affect the signal, then the communications range may be unknown during actual deployment. This necessitates a design strategy that permits nodes to operate within a given limit that would take into consideration a variety of operating conditions.

B. Security

Within the context of a MANET, an attack constitutes an intentional disabling or disruption of software, services, or hardware of one or more nodes on a network. This causes a degradation or cessation of the MANET's functionality. An attack challenges some or all of the five basic tenets of network security [13]: (1) data integrity, (2) confidentiality, (3) availability, (4) authenticity, and (5) non-repudiation. Before a search and rescue mission, security keys can be distributed to each aircraft. However, jamming and spoofing are two methods that could challenge the security of the aerial communication. A jamming attack would manifest itself as a radio transmission on a similar frequency spectrum as that used to communicate among the aircraft nodes. At sufficient strength, this rogue signal is capable of overwhelming the receiving radio transceivers on the nodes by lowering the

signal-to-noise ratio to unacceptable levels. At these levels, data cannot be distinguished from the generated noise, and the transmission is lost. A spoofed aircraft node could misdirect information about the position, speed, and objectives of other aircraft on the network, deliberately causing a collision or failure of the joint objective. In contrast to the jamming attack, which is more random in nature (even though a particular node may be targeted, the effect of cutting off communications may not be immediately known to the attacker), a spoofing attack has direct and well-defined consequences, as the attacker is responsible for causing changes at a fundamental level of the network.

C. Effects of Resource Limitations

Because of limited resources, the aerial MANET must be efficient when performing its search and rescue application. These limitations and characteristics are, in order of design importance, are: (1) power supply limitations, (2) processing power limitations, and (3) weight and size (physical) limitations of payload. In the case of a restricted power supply, such as a battery pack for aircraft flight, the node must be designed with tradeoffs of computing power and energy efficiency in mind. The processing capabilities of the system are typically dictated by power requirements, and are a direct consequence of power supply constraints. A corollary to the processing power is the need for cooling of the microprocessor. If the processor operates at a high thermal design point, it may require a cooling solution that affects the physical limitations of the payload. In addition, oversized components such as sensors (e.g., cameras) and radio communications equipment can cause significant challenges for mobility, especially if they also dramatically affect the requirements for the power source.

D. Navigation

The navigation of a civilian aircraft, such as an airliner, is governed by a large set of rules and regulations, both internal to the aircraft as implemented by pilots and flight crew, as well as external, such as air traffic controllers and navigational beacons, satellites, and visual identification markers. When considering Unmanned Aerial Systems (UAS), navigation takes on a new degree of complexity. Because a distributed UAS has multiple nodes that operate collectively and cooperatively, navigation must incorporate flight paths, operating conditions, and flight objectives for multiple nodes. While a simple solution to this may be based on a hierarchical scheme, we must ultimately understand how to accomplish this in a truly decentralized method. By decentralizing a navigation control scheme, we permit the operation of a resilient and redundant array of nodes that can be used in search and rescue applications, allowing any node or set of nodes to establish flight paths and navigational objectives.

As an example, consider the case where a ground control station initiates a request for a group of aircraft to fly to a target location along a specified path. Typically, this would require explicit commands from the ground control station with respect to organization and flight organization. Alternately, an *a priori* scheme may have been loaded into the flight control software. However, both these schemes suffer from the fact that they are inflexible, and may not offer ideal configurations for any given set of objectives.

Thus, decentralized navigation represents a spatio-temporal optimization problem. Each aircraft has a potential flight path, available resources, and information about nearest neighbors, if not the entire group of nodes. Cost optimization functions are therefore needed in order to allow group decisionmaking, in which shared information can lead to an optimal grouping and ultimately, navigation, of the MANET.

V. A Potential Testbed for Aerial MANETs

To further the research efforts in airborne networks [32], a hardware/software testbed can be constructed [33]. The Aerial MANET requires the following: (1) the establishment of an ad-hoc discovery protocol that allows the aircraft to become aware of neighbors, thereby establishing data mesh routing, and resource capability maps, (2) the establishment of a secured authentication method, using methods similar to Internet Protocol Secure (IPSec) in order to allow aircraft to ensure that communications are valid and encrypted between nodes, and (3) the establishment of a cooperative navigation scheme, permitting nodes to select optimal flight paths. In this setup, each aircraft is equipped with the ability to communicate with others using a shared secret key, and a hash-based machine authentication code (HMAC) that is salted with temporal data. Once the nodes have begun their interaction and discovery process, they begin securing their data transmissions with the specified secure communications protocol.

The aircraft that we propose to use is a native-electric high-wing trainer aircraft with a 1.8 meter wingspan called the 6-ft Telemaster Electro, from Hobby Lobby. Widely recognized for having high wing loading capacity, this aircraft design allows the transport of instrumentation, power sources, and onboard computer equipment without noticeably altering aircraft's flight characteristics. The power plant on the aircraft is a 3-phase alternating current motor, capable of generating over one horsepower. The power supply is generally implemented as a rechargeable lithium-ion polymer battery, capable of sustaining flight for over 35-40 minutes at continuous speeds of 55 km/h. Furthermore, the aircraft's high lift propensities, low takeoff velocity, and self-righting design ensure that the aircraft will yield stable flight for optimal still and moving image transmission.

To assist with autonomous navigation, each aircraft is equipped with an autopilot control board from the open-source Ardupilot project, designed to allow the aircraft to perform waypoint navigation given global position system (GPS) coordinates. Through the use of sensors onboard

the aircraft, such as GPS, magnetometer, barometer and accelerometer, the autopilot board is designed to fly the aircraft stably and reliably. It is also designed to interface via a serial connection with the onboard Gumstix Overo embedded computer system (i.e., a popular small-scale embedded Linux platform), which can be used to perform the necessary networking, reallocation procedures, and group behavior coordination.

We have investigated a number of ad-hoc discovery protocols, including the Optimized Link State Routing protocol (OLSR), BABEL, and the Better Approach to Mobile Ad hoc Networking (BATMAN) protocol. Our results have favored the use of BATMAN for its ease of implementation with the embedded Linux networking stack, fast discovery times, and robust path discovery. BATMAN is implemented as part of the Open-WRT framework, an embedded Linux operating system designed and optimized for MANET-based systems.

Through the use of the Gumstix Overo boards and their integrated wireless communications hardware (802.11), we established an initial maximum communications range of approximately 350 meters from node to node [3]. Fig. 1 shows an example aircraft configuration.

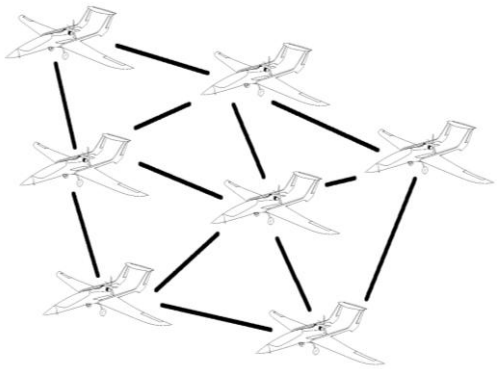

Figure 1. An example network architecture of the distributed aircraft platform

A. Testbed Performance

The testbed configuration used to validate the topics suggested in this paper consisted of five nodes, each powered by a Gumstix Overo embedded Linux computer system. Powered by an ARM Cortex A8 CPU, the Gumstix contain sufficient computing power to process data, such as telemetry information and live video data, and then send it over an array of mobile links. Each Overo was connected by a TTL RS-232 serial interface to the ArduPilot APM 2 autopilot control board. The Gumstix board is tasked with control operations, and must perform functions such as spatio-temporal optimization, group decision-making, data collection and transmission (e.g., collection of video and other sensory data, afterward using the ad-hoc network's configuration to route the data to a consuming node on the ground), and other high-level functions. Because the Gumstix does not run a real-time operating system (it uses a modified OpenWRT Linux system), the execution of Proportional-Integrative-Derivative (PID) control required for autopilot navigation must be offloaded to a microcontroller board. The ArduPilot APM 2 is responsible for flight control, and receives updated directives from the Gumstix. Flight control and telemetry variables are returned from the ArduPilot to the Gumstix so that they can be passed to other aircraft, and to the ground control station.

Node communications is established through discovery and routing, supported by the BATMAN protocol, which has been significantly modified to allow for connectivity parameters that are more appropriate to mobile UAS deployments. For instance, the connection timeout, which is normally set in the kernel module as 2 minutes, has been modified to only 2-10 seconds, depending on the configuration. The actual discovery protocol itself was a custom-written set of embedded programming that complemented the operation of the modified BATMAN protocol.

BATMAN relies on originator messages (OGMs) to propagate route information to the node mesh. OGMs require a certain amount of bandwidth, depending on the number of nodes and update frequency. In the case of our testbed, we found a logarithmic relationship between the number of OGMs, available bandwidth, and the desirability of a particular OGM frequency (Fig. 1).

In this figure, we can see that the throughput falls off rapidly as the interval between OGMs decreases. There exists a relationship between OGM frequency and flight update information, so OGMs must be tweaked carefully to balance throughput with updated information on node formations. As the OGM intervals increase, the aircraft formation may not update routes in time to prevent collisions in the case that aircraft routes or new nodes are not detected correctly. However, because of the nature of the data that is being transmitted, namely video data, throughput remains sensitive. It is possible to constrain BATMAN data to a slower but more reliable interface that is dedicated to telemetry and route information. Such a method may be implemented through a modified MPTCP scheme.

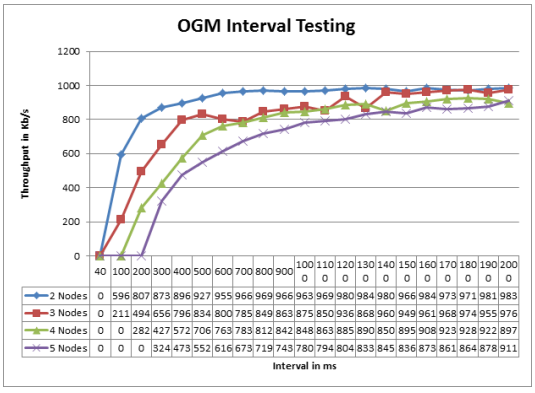

Figure 2. Network throughput vs. OGM frequencies

The time required by the system to discover nearest nodes depends on configuration parameters specified in the modified BATMAN networking architecture. Once

optimized, the network is capable of detecting new aircraft and configuring routing parameters in 6 seconds or less after boot time. Table I shows the results of adding a new node to an existing network architecture, and the discovery time needed to add it.

TABLE I. NODE DISCOVERY TIME

Nodes in Network	Discovery time (s)
Discover 2	**5.8**
Add 1 (total 3)	**5.7**
Add 1 (total 4)	**6.0**
Add 1 (total 5)	**6.0**

B. Resource Utilization

We have repeatedly stressed the resource-constrained nature of our aircraft groups. Although the Gumstix contain high-powered mobile processors and supporting hardware, they do have memory and processor limitations that could potentially impact data gathering and transmission from the onboard sensors. Ultimately, handling video and image streams are the most processor-intensive functions that the Gumstix will face. Fortunately, the selection of the modified BATMAN protocol means that memory and CPU utilization are minimal. Table II demonstrates the average memory utilization under five trial runs for several different node configurations.

TABLE II. MEMORY UTILIZATION

Number of Nodes	Percentage Memory Utilization
1	**0.14**
2	**0.14**
3	**0.14**
4	**0.14**
5	**0.15**

TABLE III. CPU UTILIZATION

Number of Nodes	Percentage CPU Utilization
1	**<1**
2	**<1**
3	**<1**
4	**<1**
5	**<1**

The impact on CPU performance of the networking standard must also be minimal. We evaluated the percentage of processor utilization required by BATMAN during its routing discovery and update procedures. Table III demonstrates the CPU requirements for different node configurations.

C. Overall Performance Impressions

Based on these results, we can see that the modified BATMAN architecture allows the assessment platform to perform without any significant overhead. Because of fast reconnection times and fast discovery and route establishment times, the modified BATMAN provides the testbed platform the ability to perform its work efficiently and reliably. Each of the components used in each aircraft are commonly available and low in cost; the computing platforms incur $200-300, the airframes and avionics an additional $600, and radios and sensory payloads cost an additional $200-300. Each unit can be used indefinitely if properly maintained, and the fault-tolerant aspects should amount to a testbed that offers a wise investment for the study of networked, fault-tolerant aerial assessment platforms for search and rescue applications and the subsequent development of implementations.

The performance impact of the security layer must still be understood. However, as it uses industry-standard symmetric encryption techniques, which can easily be supported in the mobile hardware of the Gumstix platform, the overall performance penalty is projected to be low. This is a topic for further study, as the security layer is implemented and key distribution and generation is addressed appropriately.

VI. SUMMARY

With the development of an autonomous, fault-tolerant means to perform search and rescue, we aim to increase our understanding of how MANET technologies can be made more resilient, and how they can be used to assist first responders in the task of assessing the results of severe weather phenomena on communities. By removing single points of failure, parallelizing (i.e., speeding up) the assessment process, and reducing the required investment by a community, aerial MANETs can have a direct impact in saving lives when aerial assessment information is needed the most.

ACKNOWLEDGMENT

This work was supported in part by TRUST (The Team for Research in Ubiquitous Secure Technology), which receives support from the National Science Foundation (NSF award number CCF-0424422) and the following organizations: AFOSR (#FA9550-06-1-0244) Cisco, British Telecom, ESCHER, HP, IBM, iCAST, Intel, Microsoft, ORNL, Pirelli, Qualcomm, Sun, Symantec, Telecom Italia and United Technologies.

REFERENCES

[1] A. P. Lauf, "Distributed sensing with fault-tolerant resource reallocation for disaster area assessment," Ph.D. Dissertation, Electrical Engineering and Computer Science, Vanderbilt University, Nashville, 2010.

[2] J. Casper and R. R. Murphy, "Human-robot interactions during the robot-assisted urban search and rescue response at the World Trade Center," *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 33, pp. 367-385, 2003.

[3] G. Kantor, S. Singh, R. Peterson, *et al*., "Distributed search and rescue with robot and sensor teams," in *Field and Service Robotics*, S. i. Yuta, H. Asama, E. Prassler, T. Tsubouchi, and S. Thrun, Eds., Heidelberg: Springer Berlin, vol. 24, 2006, pp. 529-538.

[4] I. R. Nourbakhsh, K. Sycara, M. Koes, M. Yong, M. Lewis, and S. Burion, "Human-robot teaming for search and rescue," *IEEE Pervasive Computing*, vol. 4, pp. 72-79, 2005.

[5] M. A. Goodrich, B. S. Morse, D. Gerhardt, *et al.*, "Supporting wilderness search and rescue using a camera-equipped mini UAV," *Journal of Field Robotics*, vol. 25, pp. 89-110, 2008.

[6] P. Doherty and P. Rudol, "A UAV search and rescue scenario with human body detection and geolocalization," in *AI 2007: Advances in Artificial Intelligence*, M. Orgun and J. Thornton, Eds., Heidelberg: Springer Berlin, vol. 4830, 2007, pp. 1-13.

[7] J. Carlson and R. R. Murphy, "Reliability analysis of mobile robots," in *Proc. IEEE International Conference on Robotics and Automation*, 2003, pp. 274-281.

[8] J. Carlson and R. R. Murphy, "How UGVs physically fail in the field," *Robotics, IEEE Transactions on*, vol. 21, pp. 423-437, 2005.

[9] J. Carlson, R. R. Murphy, and A. Nelson, "Follow-up analysis of mobile robot failures," in *Proc. IEEE International Conference on Robotics and Automation*, 2004, pp. 4987-4994.

[10] C. Paasch, G. Detal, F. Duchene, C. Raiciu, and O. Bonaventure, "Exploring mobile/WiFi handover with multipath TCP," in *Proc. ACM SIGCOMM Workshop on Cellular Networks: Operations, Challenges, and Future Design*, 2012, pp. 31-36.

[11] M. Gerla, "From battlefields to urban grids: New research challenges in ad hoc wireless networks," *Pervasive and Mobile Computing*, vol. 1, pp. 77-93, 2005.

[12] J. L. Burbank, P. F. Chimento, B. K. Haberman, and W. T. Kasch, "Key challenges of military tactical networking and the elusive promise of MANET technology," *Communications Magazine, IEEE*, vol. 44, pp. 39-45, 2006.

[13] A. Avizienis, J. C. Laprie, B. Randell, and C. Landwehr, "Basic concepts and taxonomy of dependable and secure computing," *IEEE Transactions on Dependable and Secure Computing*, vol. 1, pp. 11-33, 2004.

[14] S. Jian and G. Wei, "A survey of service discovery protocols for mobile ad hoc networks," in *International Conference on Communications, Circuits and Systems*, 2008, pp. 398-404.

[15] A. N. Mian, R. Baldoni, and R. Beraldi, "A survey of service discovery protocols in multihop mobile ad hoc networks," *Pervasive Computing, IEEE*, vol. 8, pp. 66-74, 2009.

[16] D. D. Perkins, H. D. Hughes, and C. B. Owen, "Factors affecting the performance of ad hoc networks," in *IEEE International Conference on Communications*, 2002, vol. 4, pp. 2048-2052.

[17] G. C. Hadjichristofi, L. A. DaSilva, S. F. Midkiff, U. Lee, and W. D. Sousa, "Routing, security, resource management, and monitoring in ad hoc networks: Implementation and integration," *Computer Networks*, vol. 55, pp. 282-299, 2011.

[18] Z. Gao, Y. Yang, J. Zhao, J. Cui, and X. Li, "Service Discovery Protocols for MANETs: A Survey," in *Mobile Ad-hoc and Sensor Networks*, Heidelberg: Springer Berlin, 2006, vol. 4325, pp. 232-243.

[19] H. A. Amri, M. Abolhasan, and T. Wysocki, "Scalability of MANET routing protocols for heterogeneous and homogenous networks," *Computers & Electrical Engineering*, Corrected Proof.

[20] A. Patwardhan, J. Parker, M. Iorga, A. Joshi, T. Karygiannis, and Y. Yesha, "Threshold-based intrusion detection in ad hoc networks and secure AODV," *Ad Hoc Networks*, vol. 6, pp. 578-599, 2008.

[21] K. Akkaya and M. Younis, "A survey on routing protocols for wireless sensor networks," *Ad Hoc Networks*, vol. 3, pp. 325-349, 2005.

[22] Open-mesh. (2012). *Better Approach to Mobile Ad-hoc Newtorking*.

[23] P. Brutch and C. Ko, "Challenges in intrusion detection for wireless ad-hoc networks," in *Proc. Symposium on Applications and the Internet Workshops*, 2003, pp. 368-373.

[24] Y. Zhang and W. Lee, "Intrusion detection in wireless ad-hoc networks," presented at the 6th annual International Conference on Mobile Computing And Networking, Boston, Massachusetts, United States, 2000.

[25] Y. Zhang, W. Lee, and Y.-A. Huang, "Intrusion detection techniques for mobile wireless networks," *Wireless Networks*, vol. 9, pp. 545-556, 2003.

[26] A. Mishra, K. Nadkarni, and A. Patcha, "Intrusion detection in wireless ad hoc networks," *IEEE Wireless Communications*, vol. 11, pp. 48-60, 2004.

[27] Y.-A. Huang and W. Lee, "Attack Analysis and Detection for Ad Hoc Routing Protocols," in *Recent Advances in Intrusion Detection*, vol. 3224/2004, Heidelberg: Springer Berlin, 2004, pp. 125-145.

[28] A. P. Lauf, R. A. Peters, and W. H. Robinson, "A distributed intrusion detection system for resource-constrained devices in ad-hoc networks," *Ad Hoc Networks*, vol. 8, pp. 253-266, 2010.

[29] A. P. Lauf, R. A. Peters, and W. H. Robinson, "Embedded intelligent intrusion detection: A behavior-based approach," in *4th International Symposium on Embedded Computing*, Niagara Falls, Canada, 2007, pp. 816-821.

[30] A. P. Lauf and W. H. Robinson, "Fault Tolerance in MANETs Using a Task-to-Resource Reallocation Framework," in *Proc. International Conference on Computational Science and Engineering*, 2009, pp. 753-758.

[31] A. P. Lauf and W. H. Robinson, "Fault-tolerant distributed reconnaissance," in *Proc. IEEE Military Communications Conference*, San Jose, CA, 2010, pp. 1812-1817.

[32] K. Namuduri, Y. Wan, M. Gomathisankaran, and R. Pendse, "Airborne network: a cyber-physical system perspective," in *Proc. First ACM MobiHoc Workshop on Airborne Networks and Communications*, 2012, pp. 55-60.

[33] W. H. Robinson and A. P. Lauf, "Resilient and efficient MANET aerial communications for search and rescue applications," presented at the International Conference on Computing, Networking and Communications, San Diego, California, 2013.

William H. Robinson received his B.S. in electrical engineering from the Florida Agricultural and Mechanical University (FAMU) in 1996 and his M.S. in electrical engineering from the Georgia Institute of Technology (Georgia Tech) in 1998. He received his Ph.D. in electrical and computer engineering from Georgia Tech in 2003. In August 2003, Dr. Robinson joined the Department of Electrical Engineering and Computer Science at Vanderbilt University as an Assistant Professor, and was promoted to Associate Professor in 2010.

Dr. Robinson leads the Security And Fault Tolerance (SAF-T) Research Group at Vanderbilt University, whose mission is to conduct transformational research that addresses the reliability and security of computing systems. He collaborates with both the Institute for Space and Defense Electronics (ISDE) and the Institute for Software Integrated Systems (ISIS) at Vanderbilt University. In addition to his research activities, Dr. Robinson serves as the Director of Undergraduate Studies for Computer Engineering. He also participates with the Team for Research in Ubiquitous Secure Technology (TRUST), an NSF Science and Technology Center, where he serves as the Outreach Director.

Dr. Robinson's major honors include selection for a National Science Foundation (NSF) Faculty Early Career Development (CAREER) Program Award and the Defense Advanced Research Projects Agency (DARPA) Computer Science Study Panel, both in 2008. Dr. Robinson is a Senior Member of both the Institute of Electrical and Electronics Engineers (IEEE) and the Association for Computing Machinery (ACM); he has additional memberships in the American Society of Engineering Educators (ASEE) and the National Society of Black Engineers (NSBE).

Adrian P. Lauf received his B.S. in computer engineering along with a second major in music performance (cello) in 2005 at Vanderbilt University in Nashville, TN. He received his M.S. (2007) and Ph.D. (2010) in electrical engineering from Vanderbilt University under the direction of Dr. William H. Robinson. In August 2011, Dr.

Lauf joined the Department of Computer Engineering and Computer Science at the University of Louisville as an Assistant Professor.

Dr. Lauf is the director of the Aerial Robitics Lab (ARL) that the University of Louisville. His research work seeks to integrate emerging embedded computing, networking and security applications with airborne robotics. Such work includes using fault-tolerant distributed behaviors and networking for Mobile Ad-hoc Network (MANET) aircraft arrangements to be used in aerial assessments, as well as the application of image processing and computer vision to equip smaller micro unmanned air vehicles with the ability to fly autonomously in indoor environments.

Dr. Lauf is the chair of the Louisville chapter of the IEEE Computer Society.