

SensorFlock: An Airborne Wireless Sensor Network of Micro-Air Vehicles*

¹Jude Allred, ¹Ahmad Bilal Hasan, ¹Sarooh Panichsakul, ²William Pisano, ²Peter Gray, ¹Jyh Huang,
¹Richard Han, ²Dale Lawrence, ²Kamran Mohseni
¹Department of Computer Science, University of Colorado, Boulder.
²Department of Aerospace Engineering Sciences, University of Colorado, Boulder.
{Sarooh.Panichsakul, William.Pisano}@colorado.edu

Abstract

An airborne wireless sensor network (WSN) composed of bird-sized micro aerial vehicles (MAVs) enables low cost high granularity atmospheric sensing of toxic plume behavior and storm dynamics, and provides a unique three-dimensional vantage for monitoring wildlife and ecological systems. This paper describes a complete implementation of our SensorFlock airborne WSN, spanning the development of our MAV airplane, its avionics, semi-autonomous flight control software, launch system, flock control algorithm, and wireless communication networking between MAVs. We present experimental results from flight tests of flocks of MAVs, and a characterization of wireless RF behavior in air-to-air communication as well as air-to-ground communication.

Categories and Subject Descriptors

J.2 [Computer Applications]: Physical Sciences and Engineering—*Earth and atmospheric sciences*; C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless Communications*

General Terms

Measurement, Performance, Design, Experimentation

Keywords

Wireless Sensor Networks, MAVs, Applications, Deployments

*This work was supported by the National Science Foundation (NSF) project: *ITR: Loosely Cooperating Micro Air Vehicle Networks for Toxic Plume Characterization* (grant number 0427947)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
SenSys'07, November 6–9, 2007, Sydney, Australia.
Copyright 2007 ACM 1-59593-763-6/07/0011 ...\$5.00

1 Introduction

Large wireless networks scaling to hundreds of low cost airborne vehicles are largely still a vision today rather than a reality. In our SensorFlock airborne wireless sensor network (WSN), our research goal is to make a substantial leap forwards towards this vision of hundreds of inexpensive, semi-autonomous, and cooperating airborne vehicles that sense and relay data over a wireless communication mesh network. We present in this paper our progress towards this goal, namely the design of our micro-air vehicle (MAV), the semi-autonomous flight control algorithm capable of hovering individual MAVs in loiter circles, flight validation tests of the MAVs, and an in-depth study of the RF characteristics of air-to-air and air-to-ground communication between MAVs. The benefits that will accrue to the research community from the SensorFlock project include the ability to enhance scientific applications with fine-granularity three-dimensional sampling, the distribution of the MAV aircraft design and software to the wider community, the eventual creation of airborne testbeds that scale up to hundreds of MAVs, and increased understanding of wireless propagation characteristics and networking connectivity behavior between large numbers of MAVs that are rapidly banking and rolling in flight. The latter measurement results will aid the computer science community in developing more realistic RF models for in situ air-to-air and air-to-ground communication, thus leading to improved simulation and design of more robust protocols for practical airborne sensor networks.

An airborne WSN provides the capability to enhance many applications of interest to the scientific community by providing finer granularity three-dimensional sampling of phenomena of interest than would otherwise be feasible. One such class of applications is chemical dispersion sampling. As shown in Figure 1, a deployment of a flock of MAVs sensing and communicating their data back to a network of ground stations enables scientists to study the rate of dispersion of a toxin, pollutant, or chemical, natural or man-made. In another example, MAV flocks may provide the ability to study the distribution of CO_2 concentrations in the atmosphere and its relation to global warming. In all these cases, a flock of MAVs enables accurate sampling of

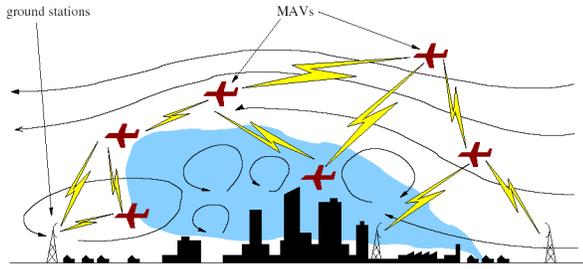


Figure 1. An airborne WSN for 3-D sensing of toxic plumes.

the parameter of interest simultaneously over large regions of a volume. In addition, since MAVs are independently controllable, they can be targeted to track the toxic plume to study the rate of dispersion, fly towards the source of the plume if unknown, and re-distribute to map the boundaries of the plume.

Another class of applications that would benefit from an airborne WSN are those involving atmospheric weather sensing. A flock of MAVs - each MAV equipped with temperature, pressure, humidity, wind speed/direction, and/or other sensors - can provide detailed in-situ mapping of weather phenomena such as hurricanes, thunderstorms, and tornados, and return data that would be useful in improving storm track predictions and the understanding of storm genesis and evolution. Other such examples include: modeling the local weather produced by wildfires to better predict their evolution and improve the deployment of firefighting resources; sensing and modeling of thermodynamic plumes over open ice leads in polar regions to better understand interactions between sea, ice, and atmosphere which contribute to climate change; and improved characterization of heat islands above cities and their impact on local weather patterns.

While many technologies exist that can contribute to the airborne WSN vision, they are either too costly, too restricted, or too limited to fully achieve by themselves a low cost and retargetable airborne sensor network. Passive sensors such as weather balloons and dropsondes cannot be re-targeted to phenomena of interest. Large Unmanned Air Vehicles (UAVs) of 2-3 meter wingspan or more pose a hazard to conventional air traffic and ground personnel. Small, bird sized sensor vehicles have the potential to reduce the consequences of failure to levels that are considered of equivalent safety to FAA approved manned vehicles.

We have therefore pursued a vision of building an airborne WSN composed of many low cost small bird-sized MAVs on the order of a half a meter in wingspan. An example of the MAV that we have built is shown in Figure 2. This MAV would pose little danger to personnel and property on the ground or other air vehicles. They do not need specialized take-off or landing facilities or runways. They are reusable, and could be produced in large numbers at low cost. With a few enhancements to our current prototype's airframe and propulsion system, such small vehicles could potentially remain in flight for periods of about 90 minutes, sufficient to provide highly accurate data for decisions in the critical

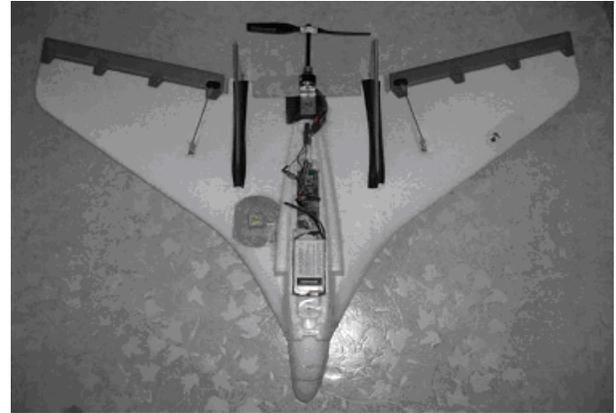


Figure 2. Micro Air Vehicle (MAV) designed and built for flight tests of our airborne WSN.

initial period after a toxin release event. Subsequently, fewer numbers might be used to monitor dispersions over longer periods.

Deploying an airborne WSN with large numbers of vehicles (e.g. tens or hundreds) raises unique research challenges in command and control. Each MAV carries very limited on-board power and computing resources. Flight control, toxin sensing, information processing, communication, and decision making must be extremely simple and decentralized. Yet rather sophisticated aggregate behavior is desired, so that the flock can semi-autonomously seek out plumes, guided by supervisory human operators and real-time models of plume evolution.

This paper describes our SensorFlock solution to large-volume atmospheric sensing, that takes the approach of using a "minimal" autopilot combined with a globally stable and convergent vector field guidance system on each vehicle. This provides a small, low mass, and low cost autopilot system that requires very little human interaction in the form of flight control or path planning. This combination provides a semi-autonomous capability for each MAV, where the operator or an overseeing algorithm can provide the desired center of loiter coordinates and a loiter radius infrequently, decentralizing the vehicle control by moving the management of the flock to a higher level in the control hierarchy. Once the plane reaches its destination it will fly loiter circles around the target point until it is told to do otherwise. With this approach many vehicles can be controlled by a single operator without the threat of failure due to a lack of command or loss of communication. In this manner, the airborne WSN can scale to large number of MAVs launched and overseen by a small number of human operators. The sections below describe how this system is implemented and show experimental results of the system in action.

2 Related Work

Detailed measurements of wireless behavior are essential to understanding real-world performance of wireless networks. For 802.11 2.4 GHz networks, measurements have been performed on static WiFi LANs [1, 4, 2, 9, 8], and an

entire workshop has been devoted to the topic [10]. Mobile vehicular 802.11 ground networks have recently been studied [7].

In wireless sensor networks, detailed measurements have been performed on testbeds of static wireless motes at 900 MHz [11, 27, 24, 29]. Experimental papers studying 802.15.4 radios at 2.4 GHz in static sensor networks have also been reported [19, 21, 28].

Prior work has explored using large UAVs for toxin dispersion characterization [17], though this is in simulation only. Other prior work have simulated UAV networks [26, 25, 12].

Practical airborne systems of wireless networked planes are largely in their infancy. A system with several small helicopters has been reported [3]. The AUGNET project reports results for two of the larger UAVs [5, 6, 14] networked via 802.11, not for the smaller bat-sized MAVs.

As far as we are aware, there is no prior work studying the network dynamics of an airborne WSN composed of bat-sized MAVs. MAVs have attracted significant attention since the mid-1990's for both civilian and military applications. Pioneering work in this area was conducted by AeroVironment [13] and the University of Florida [15] among others. MAVs are by definition small (by weight or size) aircrafts which fly at relatively low speeds. Control of such small lightweight aircraft is especially challenging given changing wind vectors, although small birds and insects have been flying under these conditions for quite some time. Our work in managing an airborne WSN focused on developing control algorithms to manage a flock of MAVs [18], wherein the location of each MAV was governed by a control law that was a function of the sensing phenomena of interest, the concentration of MAVs in a given area, communication requirements, and energy.

Cooperative control of a team of UAVs has also been investigated [16]. This work explores task assignment and trajectory planning on real-world UAVs. The focus of the work is not on wireless characterization.

3 MAV System

The design of the MAV system faced three significant challenges: building a sufficiently lightweight plane capable of flight; designing algorithms into software to achieve semi-autonomous flight; and integrating disparate subsystems such as propulsion, flight control, and wireless networking into a fully functional airborne WSN solution.

3.1 MAV Plane

As shown in Figure 2, the small size of the MAVs designed at the University of Colorado at Boulder has the advantage that in the event of an accident, the potential for damage to property and personnel on the ground is minimal. By keeping the vehicle mass under 500 grams and the maximum speed, even in failure, under 20 m/s, these vehicles fall within NASA's "Inert Debris" range safety classification. Adding that the plane is made of polypropylene foam, and the propeller is in the rear, the potential for collision damage is minimal. Combine this with a production cost of the entire aircraft, including the autopilot, which is less than \$600, and it is clear that the cost of failure of one of these vehicles is



Figure 3. CUPIC Autopilot board. This view shows the top of the board which houses the CPU, pressure sensor, radio, and rate gyro. The integrated GPS receiver is on the bottom of the board.

minimal. It is for this reason that we do not require redundant systems such as those present on larger aircraft autopilot systems.

3.2 Avionics

The CUPIC avionics board is structured around the Microchip PIC18F8722 8-bit microcontroller. The addition of an RC class receiver for testing and fail-safe purposes, and an antenna for the GPS receiver completes the flight package. Analog flight sensors consist of a single roll rate gyro and an absolute pressure sensor, while the GPS sends binary navigation information to the CPU. The control system runs on the PIC and outputs commands to the motor and servos via the on-board PWM interface. An Xbee Pro Zigbee class 2.4 GHz radio is used to support wireless communication and mobile networking. The system is capable of being flown entirely through the Xbee, but a backup RC link is maintained for network testing when there is a possibility of losing contact with an aircraft.

3.3 Fail-safe Operation

An important consideration when operating any aircraft is concern for safety. Several fail-safes are present in the autopilot system to cope with failure contingencies in flight. A watchdog timer is present that will restart the CPU if the timer is not periodically serviced. This avoids the case where a software bug causes a system crash, preventing the control system from executing regularly. The current setup uses two RF links, where the second RF link is a modified Pulse Code Modulated Remote Control radio, commonly used in the RC modeling community. The range of this radio is approximately 1.5 kilometers. Current operating procedure dictates that the aircraft be within visual range at all times. Visual range is about 1/2 kilometer, which is much less than the range of the radio, thus loss of contact due to range should not be an issue. In the potential case where the plane starts flying away from the operator, there is a fail-safe in place that will turn the motor off when the plane loses the RC link. With the motor off, the aircraft will gently glide down.



Figure 4. Plane-A-Pult Automatic Aircraft Launcher performing a fully autonomous launch.

In our testing, though the planes are flying autonomously, an RC pilot is constantly monitoring each plane and can switch back to manual control of the aircraft at any time. This is useful to keep the aircraft from coming down if there is a minor failure, such as loss of GPS. This is also useful if something enters the airspace nearby, such as a manned aircraft or a flock of birds. The aircraft are always operated at an altitude of less than 150 meters to avoid potential conflict with larger aircraft and to maintain line of site for the backup-pilots.

3.4 Launcher

To further make the MAV system operable with as little human interaction as possible, an automatic launching system was developed. The design makes use of a rugged aluminum frame propelled by a constant force coil spring. A "V" shaped aircraft carriage is used that holds the aircraft by the wings while allowing the propeller in the rear of the plane to spin-up prior to launch. A release servo is mounted to the Plane-A-Pult, and connected to the avionics board with a 3-wire interface. The avionics board then sends a signal to the actuator to initiate release. When released the 3-wire connecting plug is pulled disconnected, thus leaving the release actuator behind. Using this method, no human interaction is required for the plane to take off. The avionics can determine when all of its sensors are on-line and determine when to release itself. The autopilot is in control of the aircraft the entire time, thus no pilot is required. Any number of aircraft could be launched simultaneously using this method and a single operator sending a "launch all" command.

3.5 Control Subsystem

The control subsystem takes in input readings from GPS to determine location and ground track heading, and compares the measured heading to a desired heading determined from a pre-determined vector field which is a function of vehicle position [22]. The error between the current and desired heading is fed back in a control loop to generate a roll command and alter the MAV's trajectory, e.g. to loiter in a circle above a given position with a particular radius. Trajectory adjustment happens at a rate of 10 Hz, though an inner feedback loop adjusts roll rate instantaneously to keep the plane aloft, e.g. due to wind, at a rate of 100 Hz. The entire system is capable of autonomous loitering. In our flight tests, we were able to release the RC control of the MAVs and observe the autopilot system engage and fly the plane independently of direct human control. The trajectories generated by the autopilot system are remarkably circular, as will be shown in the experimental section.



Figure 5. Our first flock of 5 MAVs. The design of each MAV is relatively simple, so scaling to produce larger flocks can be quickly achieved.

4 Wireless Networking

Our goal in the networking implementation was to design and build the software and experiments necessary to characterize RF performance and networking connectivity of air-to-air and air-to-ground behavior. This approach will improve the accuracy of new models developed to simulate realistic airborne communication and movement. We designed a series of experiments to collect received signal strength indicators (RSSI), GPS-based position and time, path information, and packet loss and throughput statistics. From these sets of data, we were able to evaluate for example how RSSI varied with various factors such as distance and antenna orientation for air-to-air and air-to-ground scenarios. We conducted many flight tests, and present the results of our five-plane tests. Figure 5 shows the full five-plane set of MAVs that we have built as part of our airborne WSN.

The MAVs are equipped with 802.15.4-compliant radios, namely the XBEE Pro Zigbee radios from Maxstream. These radios were chosen for their combination of light weight, long transmission range, serial interface compatibility with the PIC processor, and packet interface. The radios offer a range of over 1 mile at 60 mW, though we measured at times ranges of 2 miles or more in early balloon-based testing. This also suited the anticipated maximum spacing between neighboring MAVs deployed in our toxic plume scenario. The radio weighs just 4 grams. The packet interface offered by the XBEE Pro substantially simplifies packet processing, as the PIC is not burdened by processing each bit that arrives, as would be the case in more primitive bit-interface radios, such as the Chipcon CC1000 common to MICA2 motes. The packet interface reduces the memory cost and CPU time required for networking on the PIC, and is important for enabling the PIC to sense, process and actuate flight control in real-time. Power consumption of the radio is relatively modest compared with powering the aircraft's propeller propulsion system.

The software controlling the plane's flight via the PIC also is used to route packets through the multi-hop wireless network. In order to meet these real-time constraints of flight control, it was necessary to design the networking software to complete its time quantum in less than 10 ms, where the autopilot controller needs to operate at 100 Hz, though in



(a) The team preparing to fly the flock of MAVs.

(b) Three MAVs in flight simultaneously.

Figure 6. Pictures of the flight tests.

reality 40 Hz should suffice. To maintain simplicity, we implemented a simple software control loop that first consulted the controller and then invoked the networking subsystem, alternating between these two subsystems in an indefinite loop. The networking code was written to quickly execute in its time quantum of 10 ms, and we estimate it did not take more than 0.5 ms in practice. Our networking code only processes one packet per invocation, so as to relinquish control back to the autopilot as soon as possible. This constrained the networking code from taking over the processor. The entire code uses about 1 KB out of the total of 4 KB RAM for the PIC, with about 150 bytes devoted to the networking. We found that this approach was acceptable for maintaining flight while also simultaneously reporting and routing networking data on the PIC processor.

We designed the following experiment to characterize the wireless network to the fullest extent practicable. Each MAV periodically flooded data packets 5 times per second throughout the network. Zigbee routing was disabled on the XBEE radios to permit our custom routing, though we continued to use the 802.15.4 MAC layer. Each data packet contained the originating source ID, the GPS location of the source, the GPS time, the hop count, the sequence number, and the local sender's ID, among other fields. Each node would forward the data packet only if its sequence number exceeded the current highest recorded sequence number from the originating source. In addition, each node would append its ID and the RSSI of the packet to the packet's payload. This approach allowed us both to evaluate the connectivity of every link in the mobile network, and to trace the multi-hop paths of packets through the mobile network.

The base station collected all data packets that it received from each MAV, thus enabling a rich characterization of all paths to the basestation, and the link strengths along the way. However, this only captures part of the behavior of the network, since it does not capture paths to each MAV. Since bandwidth limitations prevented us from reporting the path taken by each data packet received by each MAV, as was the case at the base station, then our approach was to have each MAV collect aggregate statistics of paths to it and periodically report summaries of network connectivity to it.

Each MAV computed aggregate statistics of the packet loss, RSSI, and hop count distribution from each of the other originating sources, as well as from each of its current direct neighbors. Every received data packet would be used to update the aggregate statistics, whether that packet was forwarded or not. Aggregate source statistics and neighbor statistics were collected and reported separately in alternating non-overlapping ten second intervals. Nodes used the same flooding mechanism to route these statistics reports back to the base station.

5 Wireless Experiments

Our goal in performing flight test experiments was to develop a realistic understanding of the airborne RF link's behavior, as well as multi-hop networking performance in such a mobile airborne environment. We performed experiments that evaluated the performance of air-to-air, air-to-ground, and ground-to-ground wireless links. We present in this section detailed analyses of RSSI, path loss exponents, packet loss, symmetry, and multi-hop behavior.

The experiments were performed by flying a flock of five MAVs. Figure 6(a) illustrates the team preparing to fly our flock of MAVs. Figure 6(b) shows in one frame three of the MAVs in flight. To satisfy safety concerns during our tests, each MAV had a human RC pilot acting as backup for the autonomous flight software. This logistical complication, in addition to uncooperative weather, hardware, and software system development factors, raised significant challenges in the collection of the data presented here.

The networking code sourced packets at a rate of 5 packets/sec. Each packet contained at least the following: GPS (x,y,z) coordinates; GPS time; packet sequence number; source ID; hop count; and a field for substituting the RSSI upon reception. Application packets are of fixed length of 50 bytes. In each MAV, the antenna is a quarter wave whip with a very small ground plane, so that the pattern can be approximated by the donut/toroid-shaped antenna pattern of the half wave dipole. The antenna is oriented vertically, pointing upwards when the plane is level. Transmit power was set to the lowest value of 10 dBm. The 802.15.4 radios were configured for a rate of 56 kbps. API mode was enabled for the radios, so that data could be sent and retrieved in a pre-defined

packet format. API mode also conveniently returns the RSSI of received packets. Packets are only returned by the radio if they pass the 802.15.4 error detection check. MAVs were programmed to loiter around a chosen point with a radius of 50 m. During our experiments, we verified using the WiSpy spectrum analyzer that there were no other sources of wireless interference in our immediate area from transmitters at the same 2.4 GHz frequency.

5.1 A Five-Plane Flight Test

To provide a sense of the dynamic behavior of MAVs while in flight, Figure 7 shows the actual trajectories of five MAVs obtained from their GPS coordinates during one of our flight tests. As the trajectories are color coded, these figures are best viewed on a color viewer or printout. The MAVs were programmed to loiter over different points. The trajectories shown capture about 30 minutes of flight, or about 10K packets from each MAV. The base station is located near (700,700). The MAVs were launched from a point about 200 m away from the base station and then flown to their loitering point, at which point autonomous loitering was enabled. The farthest distance in the (x,y)-plane that a MAV flew from the base station was around 500 m. The highest altitude flown during the test was about 125 m.

As can be seen from Figure 7(a), the top-down view of the flight patterns reveals that four of the MAVs achieve nearly circular loitering during part of their trajectories, confirming that the loitering software was operating correctly when enabled. MAV 3 (green) was manually remote controlled, and was piloted for part of its flight to create some straight line flight patterns. Figure 7(b) provides a more three-dimensional perspective on the relative trajectory of each flight.

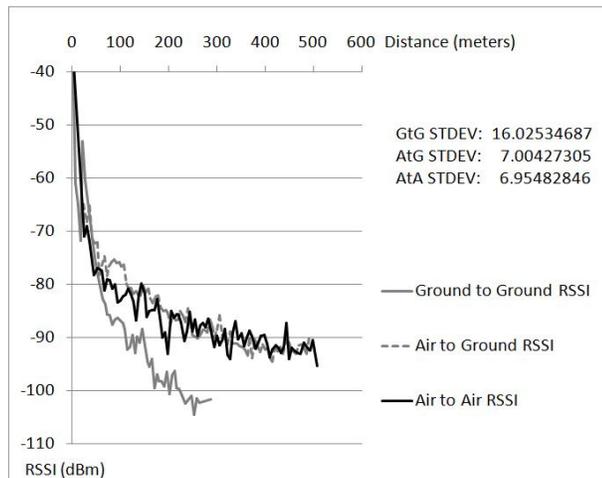


Figure 8. Average RSSI vs. distance for air-to-air, air-to-ground, and ground-to-ground wireless links.

5.2 Characterization of Air-to-Air Received Signal Strength

Our next step is to characterize the airborne RF link’s behavior. In particular, we seek to understand in this section how air-to-air received signal strength is impacted by a variety of factors such as distance and antenna orientation of

both the transmitter and receiver. After gaining an understanding of what factors influence RSSI, we will cover in the next subsection the behavior of packet loss.

5.2.1 Path Loss Exponent

Figure 8 shows the average RSSI vs. distance for air-to-air (AtoA), air-to-ground (AtoG), and ground-to-ground (GtoG) wireless links. The RSSI for the AtoA link falls off more gradually with distance than GtoG links, and is similar in behavior to AtoG links. We expect GtoG links to suffer more from shadowing, lack of RF line of sight, and multipath than AtoA links, which have a direct line of sight. Also, the deviation of RSSI for the AtoA link is substantially smaller than the GtoG link. We believe the same ground effects mentioned earlier will also lead to a wider variation in RSSI for GtoG links than AtoA links under similar circumstances.

The rate of falloff of RSSI vs. distance determines the *path loss exponent*. Figure 9 shows the path loss exponents that we measured for each of the three types of wireless links. The path loss exponent was calculated by finding the slope of the line of the least squares fit to the scatter plot of RSSI (dB) vs. distance (log). For the distances that we measured up to about 500 m, we observe that the path loss exponent for AtoA links is around the value 1.9, which is close to the free space path loss exponent of 2. For the AtoG wireless links, the path loss exponent is only slightly higher at 2.1. In contrast, the GtoG path loss exponent is much larger at 3.5. The ground measurements were taken immediately following the flight tests at the same outdoor RC airfield with the base station at the same position and orientation. The ground was flat and free from any obstructions between the base station and the grounded MAV. For comparison, urban cellular radio experiences path loss exponents from 2.7-3.5 [23].

Calculation of the path loss exponent requires that a close-in reference distance be carefully selected, which is commonly taken as 1 meter [20]. Since it was undesirable to fly the planes within 1 meter of each other, to avoid collision, we used ground-to-ground data points at about 1 meter distance as a substitute. We believe our estimates of the path loss exponent can be slightly improved by obtaining more distance measurements that are close in (< 20 m) for AtoA and AtoG links.

5.2.2 Antenna orientation

The antenna pattern is predicted to be in the shape of a toroid. The transmit antenna gain should depend primarily on the angle between the transmitting MAV’s vertical antenna and the vector from the transmitting MAV to the receiving MAV. We call this the transmit orientation angle θ_t . For example, if $\theta_t = 0^\circ$ or 180° , then the received signal should be weakened because it is transmitted through the hole of the donut, whereas if $\theta_t = 90^\circ$, the RSSI should be relatively strengthened because it is transmitted through the thickest part of the donut. The azimuth should not factor heavily into the RSSI because of the symmetry of the toroid. Similarly, the receiver antenna gain should depend primarily on the angle from the vertical antenna of the receiving MAV to the vector from the receiving MAV to the transmitting MAV. We call this the receiver orientation angle, θ_r .

Figure 10 shows a scatter plot of RSSI vs. transmit orientation angle for a single AtoG link.

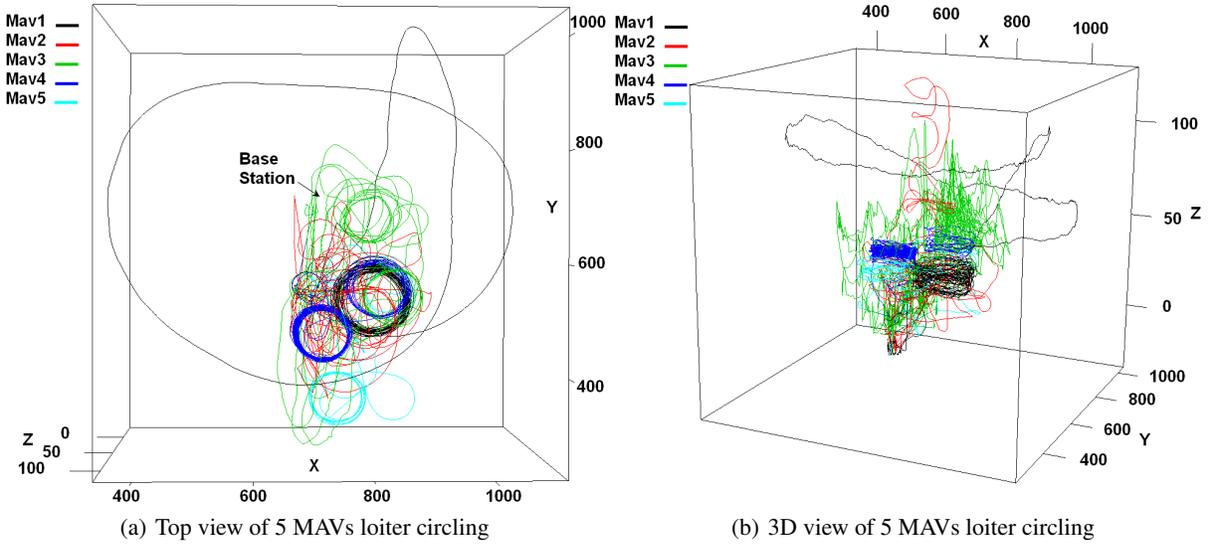


Figure 7. Color-coded trajectories of five-MAV flight test, including loiter circles (distances given in meters).

Note that Figure 10 does not control for a constant distance, and instead aggregated RSSI from all distances. Our belief was that there was a sufficient distribution of distances across each band of orientation angles such that distances would not distort the effect of transmit angle and an average trend would emerge. Indeed, the curve fit reveals a modest dependence of RSSI on transmit antenna orientation. The curve forms a slight bow, with a maximum around 90° , and minimum on the edges of the bow around 0° and 180° , with both minima about 2-3 dBm weaker than the peak. This result roughly confirms our intuition, namely that the antenna gain of the transmitter should be maximal around a 90° angle from the vertical antenna and minimal around 0° and 180° .

For AtoA links, Figure 11 shows the RSSI vs. transmit orientation angle. We have controlled for constant distance, selecting only those data points when the MAVs were within 150-160 m of each other, so distance should not substantially affect the results. The trend from the figure is that transmit angles less than 90° , i.e. the receiving MAV is above the 90° plane formed by the wings and fuselage, result in weaker RSSI than the angles below the 90° plane. Another way to describe this is that the RF toroid is “thicker” below the 90° plane than above it. This was somewhat surprising to us, as we expected the toroid to be of approximately equal strength above and below the plane, which the AtoG results seemed to show. The trend shown here was consistent when calculating the receiving angle as well, which is essentially the transmit angle just from the perspective of the other MAV. One effect we have not controlled for in this figure is the receiver’s orientation, namely we’d like to keep both distance and receiver orientation constant, to study the effect of just the transmitting MAV’s orientation. The AtoG analysis allowed us to hold the receiver orientation constant, resulting in a different bow-shaped trend.

5.2.3 ANOVA analysis

To quantify the impact of distance on RSSI, we performed linear regression on the flight data and employed a statistical technique to assess the relative importance of this factor.

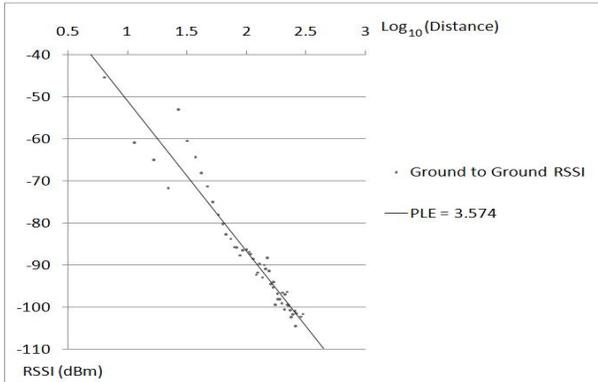
The following equation emerged from our linear regression of the loitering paths data:

$$\begin{aligned}
 \text{RSSI}(dBm) &= f(\log \text{ of distance, residuals}) \\
 &= \beta_1 * X_1 + \epsilon \\
 &= -36.23 * \log \text{ of distance} - 7.17
 \end{aligned}$$

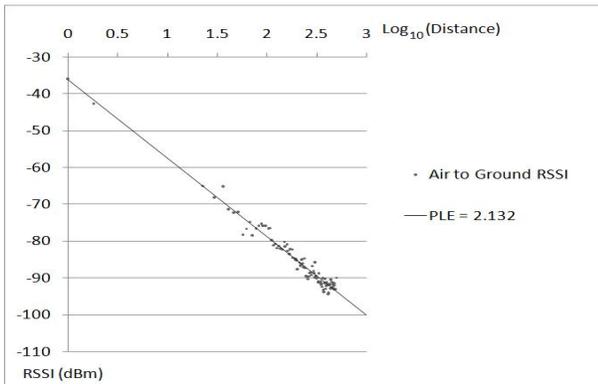
The errors are normally distributed and hence the model is assumed valid. The F-test can be used to check if the concerned variation is more significant than the errors. The F-test shows that the main factor is more significant than the errors with 95% confidence. The fraction of the variation that is explained determines the “goodness” of the regression and is called the coefficient of determination, R^2 . The R^2 for our model is 57%, meaning that our regression model explains 57% of the RSSI value. Thus, whether the RSSI can be linearly modeled as a function of distance is debatable.

We performed the test called ANOVA (ANalysis Of VARIance) that statistically analyzes the significance of how much each factor contributes to the response. The factor is considered important if its contribution to variation is high compared to that of the errors.

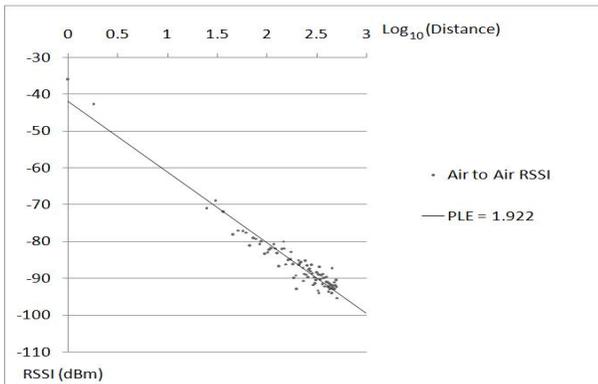
Table 1 indicates for AtoA links the percentage of variation of RSSI due to distance. When both MAVs are loitering in circular paths, the dependence on distance is 57%. When one MAV is flying in a more irregular path, while the other is circling, the dependence on distance decreases to 46%. In both cases, distance is a strong factor in influencing the received signal strength. We hypothesize that other factors such as antenna orientation contribute to the residuals. We employed ANOVA analysis on the transmit and receive orientation angles, but did not observe a strong linear depen-



(a) Linear regression to obtain path loss exponent for ground-to-ground.



(b) Linear regression to obtain path loss exponent for air-to-ground.



(c) Linear regression to obtain path loss exponent for air-to-air.

Figure 9. Path loss exponents for the three types of wireless links.

	log of distance	errors/residuals
loitering paths	57.23	42.77
irregular paths	46.31	53.69

Table 1. Percentage of variation of RSSI due to distance, from Analysis Of Variance (ANOVA), for AtoA wireless links.

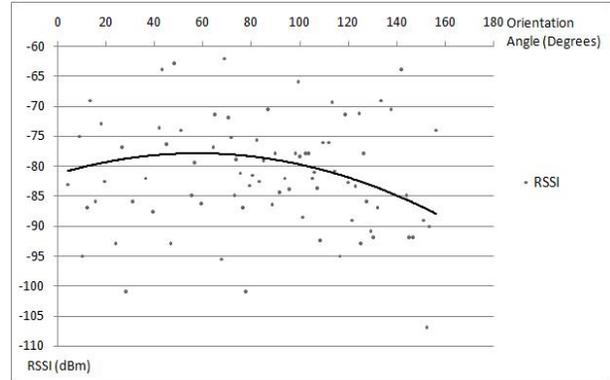


Figure 10. RSSI vs. transmit orientation angle, with curve fit, for an AtoG link.

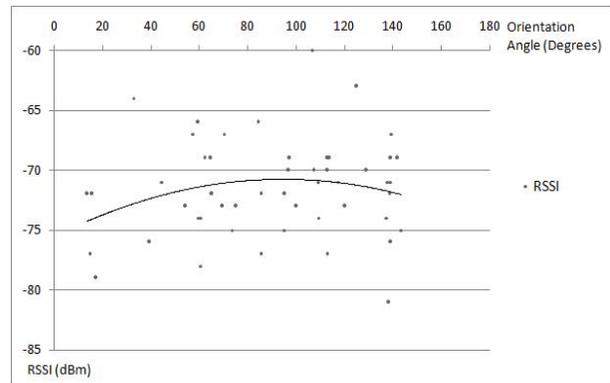
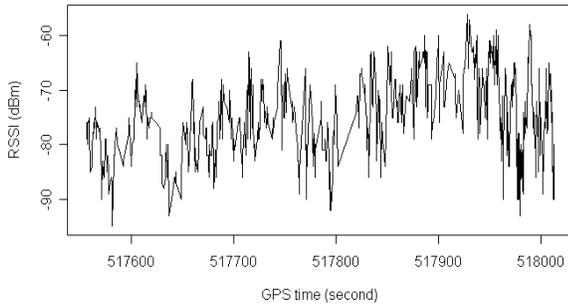
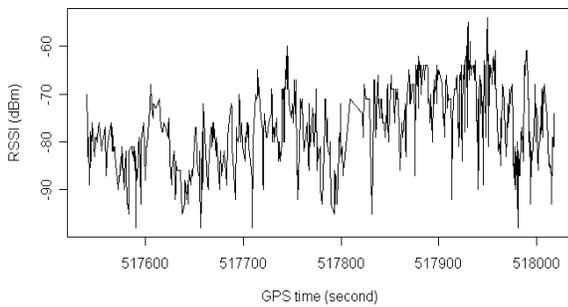


Figure 11. RSSI vs. transmit orientation angle, with curve fit, for an AtoA link.



(a) RSSI versus time for link MAV-1-to-MAV-2



(b) RSSI versus time for link MAV-2-to-MAV-1

Figure 12. RSSI vs. time for the forward and reverse links between two MAVs.

dence. However, as Figure 10 indicated, RSSI has a non-linear dependence on the orientation angle.

5.2.4 Symmetry

We observed that the AtoA RF links exhibited strong symmetry, which can be demonstrated in a variety of ways. First, Figure 12(b) illustrates the “reverse” direction MAV 2→1 RF link in terms of RSSI versus time. We observe that there is a strong resemblance in behavior between the two airborne links 2→1 and 1→2. Compared to Figure 12(a), whenever the trend in one link’s RSSI rises(falls), the other link’s RSSI trend correspondingly rises(falls). Packet loss gaps such as near 517800 are also correlated.

Figure 13 illustrates the percentage difference in RSSI between the forward and reverse air links as a function of time. For each packet sent in the forward direction, the closest packet in time was found in the reverse direction, and vice versa, and their RSSIs were compared. There was no substantial difference in time between the received signal strengths on the forward and reverse links. The average percentage difference in RSSI was only about 2%. This shows that the forward and reverse air links are fairly symmetric.

Figure 14 depicts the percentage difference in RSSI between forward and reverse links as a function of distance. Again, the RSSI measured at an instant in time in the forward direction of a link is compared to the nearest RSSI in time in the reverse direction, and vice versa. The resulting

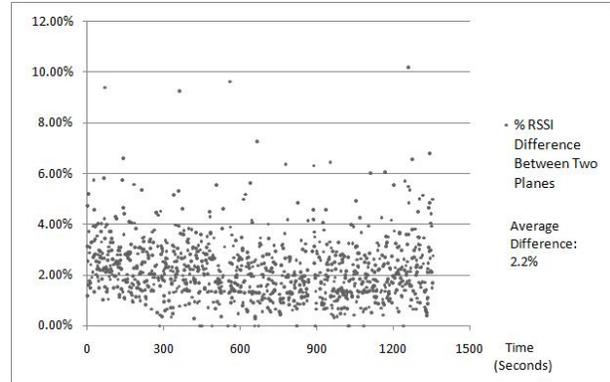


Figure 13. Symmetry of AtoA links: percentage difference in RSSI between forward and reverse links vs. time.

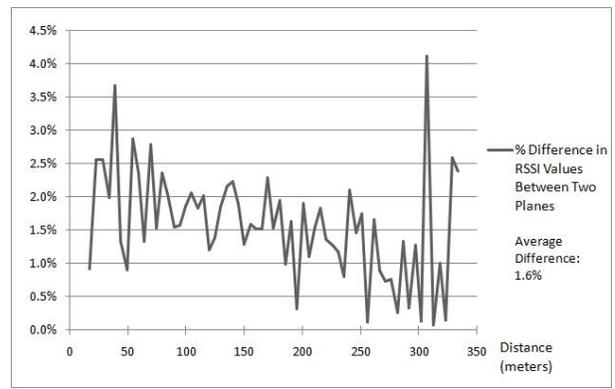


Figure 14. Symmetry of AtoA links: percentage difference in RSSI between forward and reverse links, as a function of distance.

differences are grouped by distance. The percentage difference is relatively small at all measured distances, averaging less than 2%.

5.3 Characterization of Packet Loss

Figure 15 shows the absolute number of packets lost as a function of RSSI and time for one AtoG wireless link. The x-y plane plots RSSI vs time, and shows a periodic behavior due to loiter circles. However, whenever a sequence of packets is lost, we plot the size of the sequence or gap on the z-axis, drawing a vertical line up from the last correct RSSI value by an amount equal to the number of consecutively lost packets. These appear as spikes in the figure. The figure indicates that the frequency and height of the spikes roughly increases as the RSSI decreases.

To further understand where packet losses occur, we show in Figure 16 scatter plots of absolute packet loss vs. both distance and RSSI for both AtoG and GtoG wireless links. For GtoG links, Figure 16(a) shows that as distance increases, the size of packet loss gaps increases. Note that the y-axis is logarithmic. Similarly, Figure 16(b) shows that as RSSI decreases, the size of packet loss gaps increases. The average RSSI for successfully received packets was -75 dBm. As a result, the vast majority of packet losses occur below the average RSSI. For AtoG links, a similar though less pronounced

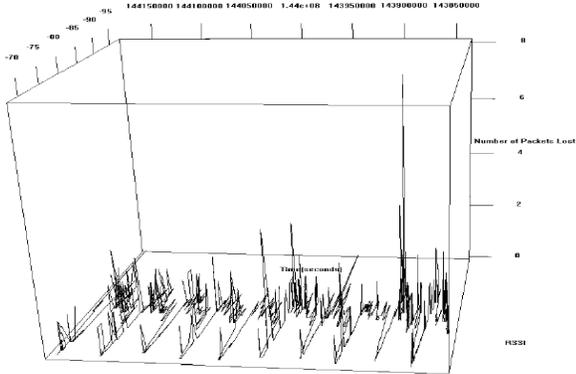


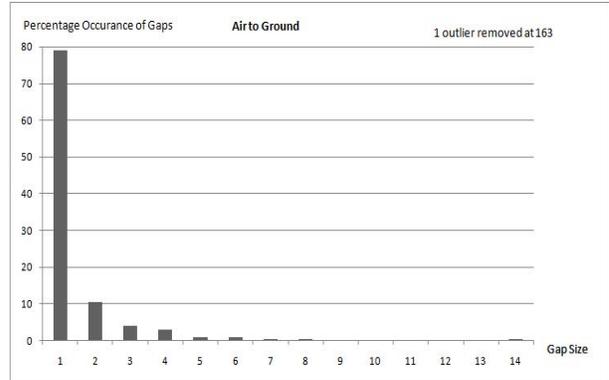
Figure 15. Packet loss for one circling MAV. RSSI vs. time is plotted in the x-y plane, while z-height measures the number of lost packets for each gap, plotted from the last RSSI point before the loss/gap.

nounced pattern emerges. Figure 16(c) shows a slight trend towards more severe packet losses as distance increases, while Figure 16(d) shows a slightly clearer trend towards more severe packet losses as RSSI drops. We believe the trends are not as clear in these AtoG links because the planes were relatively close to the base station during AtoG communication. The average AtoG RSSI for successfully received packets was -80 dBm.

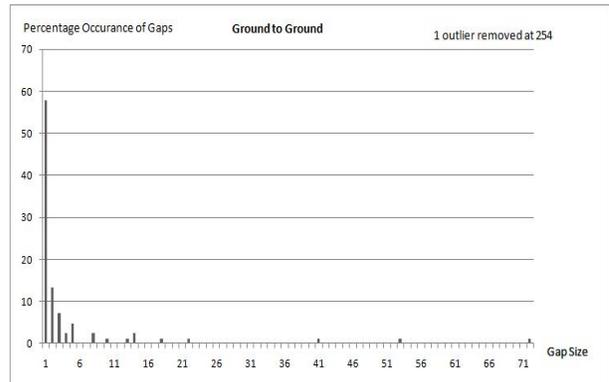
Figure 17 plots the distribution of the packet loss gap sizes for both AtoG and GtoG links. The vast majority of losses are 1-3 packets in length. However, there are some outlier gaps that are quite large, and both graphs have been truncated for clarity, each removing one large outlier (@163 for AtoG and @254 for GtoG).

The *percentage* packet loss vs. distance is shown in Figure 18. Percentages were calculated by binning the distance every 10 m and counting the percentage of lost packets in each bin. A gap of N packets was estimated to have occurred at an approximate distance equal to the average of the two successfully received packets defining the start and end of the gap. An average RSSI was similarly calculated for each gap. In this way, each gap was associated with a given distance, hence a bin. The resulting percentage losses show that our GtoG links began to lose the majority of their packets at about a range of 250 m, while our AtoG links are much more reliable over a longer range. In particular, 20% packet loss occurs at about 400 m for AtoG links vs. 210 m for GtoG links. The corresponding RSSIs at these distances, from Figure 8, are about -92 dBm for AtoG links and -100 dBm for GtoG links.

We would expect AtoA links to experience even less packet loss than AtoG links, given the relative performance numbers for RSSI seen in the previous subsection between AtoG and AtoA links. However, we experienced AtoA packet loss ranging from 20-60%, rising with distance. We believe this was due to the way in which we were collecting AtoA link quality data, which relied on forwarded packets in multi-hop flooding. There was only one packet buffer on each MAV in our implementation, so it is possible that each routing MAV would lose a large fraction of forwarded



(a) AtoG distribution of packet loss gap size.



(b) GtoG distribution of packet loss gap size.

Figure 17. Distribution of packet loss gap sizes.

packets at each hop due to buffer overflow. The same packet loss does not occur for AtoG data because packets reach the base station directly without being relayed by a MAV. We are attempting to determine the source of this packet loss and improve the flooding mechanism. Because these implementation effects exaggerate AtoA packet loss, and do not illustrate the impact of distance and other factors on AtoA packet loss, we have not shown the AtoA data for absolute packet loss gap size in Figures 16 and 17. We hope to provide clearer AtoA packet loss data in a journal version of this paper.

However, Figure 18 does afford us an opportunity to plot AtoA packet loss % provided we can normalize for the effect of relay-induced packet loss. We conducted a 2-plane ground experiment, with both planes within 1 m of each other and the base station, to measure the amount of packet loss introduced by relaying. Under these circumstances, we determined that baseline packet loss due to congestion/relaying averaged about 43%. The normalized AtoA packet loss % after subtracting out the baseline is plotted in Figure 18 as a function of distance. We observe that AtoA packet loss % behavior is similar to AtoG links, rising as the distance approaches 300 m.

Figure 19 shows how packet loss percentage is distributed vs. RSSI. Overall, this figure normalizes for transmit power, whereas the prior figure of packet loss % vs. distance was

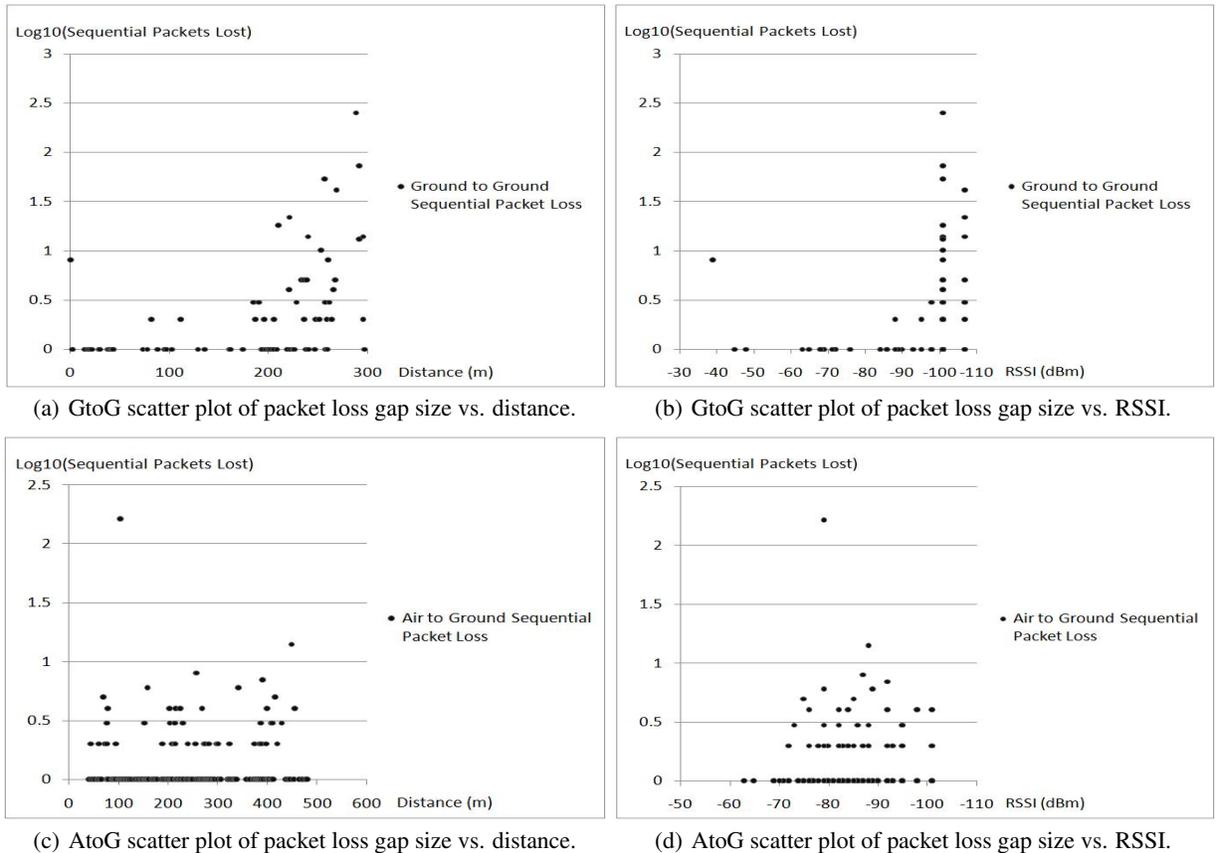


Figure 16. Scatter plots of packet loss gap sizes vs. distance and RSSI.

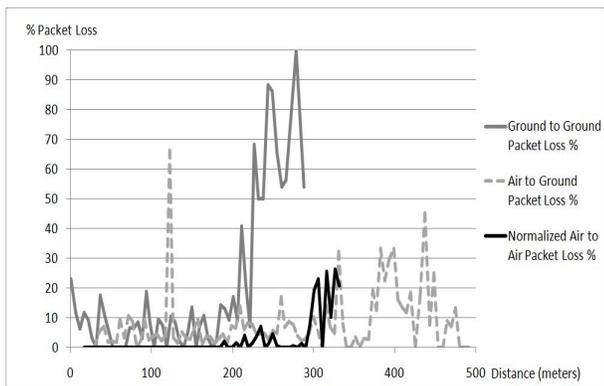


Figure 18. Packet loss percentage vs. distance.

dependent on the initial transmit power. For GtoG wireless links, packet loss % generally increases as RSSI decreases. There were some anomalous results for the highest RSSIs, likely due to how we initiated the GtoG measurements, with one of us standing over the base station while not holding the MAV perfectly level, before walking away with the MAV. For AtoG links, the trend of packet loss % increasing as RSSI decreases is clearer. Again, AtoA links, normalized for baseline congestion, exhibit behavior that is similar to AtoG links. We hope that future experiments will provide

more AtoG and AtoA data points at the lowest received signal strengths.

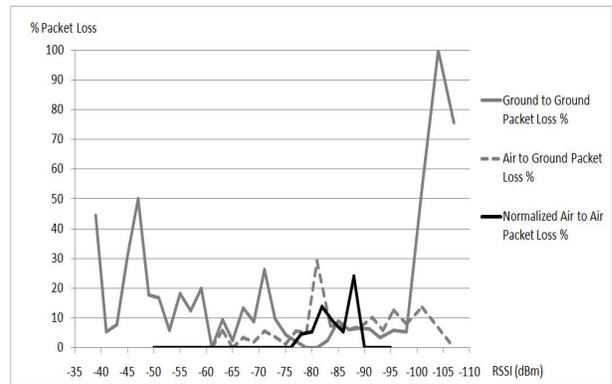


Figure 19. Packet loss percentage vs. RSSI.

5.4 Airborne Multi-Hop Network Characterization

Though the transmit power was lowered to its minimum value, all MAVs were still within range of the base station and each other during the entirety of our flight tests. As a result, true multi-hop routing could not be evaluated in these tests.

Nonetheless, flooding within a single collision domain

does introduce multi-hop forwarding between the MAVs. Figure 20 shows the distribution of hop counts received at the base station. The vast majority of hop counts were observed at 1-2 hops. Since the flooding protocol only forwards packets that it has not heard before, then most MAVs should hear the first broadcast of a packet and then relay it one more hop to the ground. Three or more hops would only occur if the first broadcast was not heard by at least one MAV, which should occur with less frequency, as shown in the figure.

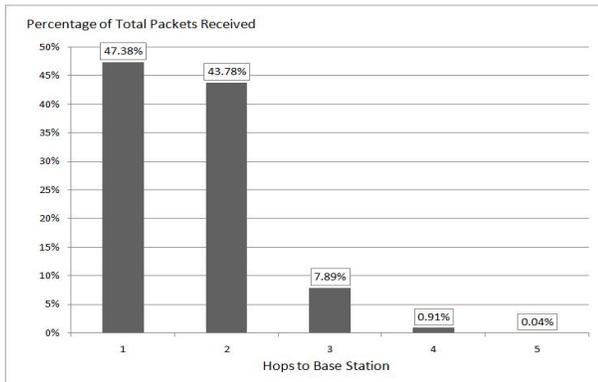


Figure 20. Distribution of received hop counts at the base station.

6 Future Work

We would like to conduct a multi-hop networking flight test in which the MAVs are not all in the same collision domain. This requires almost kilometer separation between the MAVs, which was not available at our RC airfield. Also, our flight tests collected aggregate neighbor and source statistics and periodically flooded these back to the base station. We have not yet analyzed these aggregate data. As our experience grows, we hope to scale SensorFlock to a much larger number of MAVs. Finally, we hope to equip SensorFlock with lightweight chemical sensors to support a real-world application for monitoring toxic plumes.

7 Conclusions

We have presented SensorFlock, our airborne wireless sensor network of micro-air vehicles. We have described the design of MAVs, including the avionics, the flight control software, and the launch system. We have shown the capability to perform controlled loiter-circle hovering of MAVs. We have provided an in-depth characterization of the RSSI behavior for ground-to-ground, air-to-ground, and air-to-air wireless links, showing how AtoA links fall off in signal strength with distance much less severely than GtoG links. In particular, we have quantified the path loss exponents of GtoG (3.57), AtoG (2.13), and AtoA (1.92) wireless links. We have further conducted statistical ANOVA analysis of RSSI dependence on distance, revealing a strong dependence of RSSI on distance over various flight paths. Airborne links were shown to be roughly symmetric. The distribution of packet loss gap sizes showed that the vast majority of packet losses were of size one or two packets in length. Packet loss percentage was shown to increase with increasing distance

and decreasing RSSI for AtoA, AtoG, and GtoG links, with AtoA and AtoG sharing similar behavior.

8 Acknowledgements

We wish to thank Cory Dixon, Jack Elston, James Mack, Michael Buettner, Carl Hartung, Wang-Ting Lin, Todd Mytkowicz, and Richard Messick. We also wish to thank our shepherd Thiemo Voigt and our anonymous reviewers for their comments.

9 References

- [1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *SIGCOMM*, August 2004.
- [2] A. Akella, G. Judd, S. Seshan, and P. Steenkiste. Self-management in chaotic wireless deployments. In *ACM MobiCom*, 2005.
- [3] B. Bellur, M. Lewis, and F. Templin. Tactical information operations for autonomous teams of unmanned aerial vehicles (UAVs). In *IEEE Aerospace Conference Proceedings*, volume 6, pages 2741–2756, 2002.
- [4] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *ACM MobiCom*, 2005.
- [5] T. Brown, B. Argrow, S. D. R.-G. Thekkekkunnel, and D. Henkel. Ad hoc uav ground network (augnet). In *Proc. AIAA 3rd “Unmanned Unlimited” Technical Conference*, Chicago, IL, September 2004.
- [6] T. Brown, S. Doshi, S. Jadhav, and J. Himmelstein. Test bed for a wireless network on small uavs. In *Proc. AIAA 3rd “Unmanned Unlimited” Technical Conference*, Chicago, IL, September 2004.
- [7] V. Bychkovsky, B. Hull, A. Miu, H. Balakrishnan, and S. Madden. A measurement study of vehicular internet access using in situ wi-fi networks. In *ACM MobiCom*, 2006.
- [8] K. Chebrolu, B. Raman, and S. Sen. Long-distance 802.11b links: Performance measurements and experience. In *ACM MobiCom*, 2006.
- [9] Y.-C. Cheng, J. Bellardo, P. Benko, A. C. Snoeren, G. M. Voelker, and S. Savage. Jigsaw: Solving the puzzle of enterprise 802.11 analysis. In *ACM SIGCOMM*, 2006.
- [10] ACM E-WIND (Experimental approaches to wireless network design and analysis) Workshop. <http://acm.org/sigcomm/sigcomm2005/w1-e-wind.html>.
- [11] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker. Complex behavior at scale: An experimental study of low-power wireless sensor networks. In *UCLA Computer Science Technical Report UCLA/CSD-TR 02-0013*, February 2002.
- [12] M. Gerla, Y. Yi, and K. Xu. Team communication among airborne agents. In *Proc. Of American He-*

- licopter Society 59th Annual Forum*, pages 163–167, 2003.
- [13] J. Grasmeyer and M. Keennon. Development of the Black Widow micro air vehicle. AIAA paper 2001-0127, Reno, NV, 2001. 39th AIAA Aerospace Sciences Meeting & Exhibit.
- [14] D. Henkel, C. Dixon, J. Elston, and T. Brown. A reliable sensor data collection network using unmanned aircraft. In *Second International Workshop on Multihop Ad Hoc Networks: from theory to reality (REALMAN)*, Florence, May 2006.
- [15] P. Ifju, D. Jenkins, S. Ettinger, Y. Lian, W. Shyy, and M. Waszak. Flexible-wing-based micro air vehicles. AIAA paper 2002-0705, American Institute of Aeronautics and Astronautics, January 2002. 40th Aerospace Sciences Meeting & Exhibit, Reno, Nevada.
- [16] E. King, M. Alighanbari, Y. Kuwata, and J. P. How. Coordination and control experiments on a multi-vehicle testbed. In *Proceedings of the IEEE American Control Conference*, pages 5315–5320, April 2004.
- [17] M. Kovacina, D. Palmer, G. Yang, and R. Vaidyanathan. Multi-agent control algorithms for chemical cloud detection and mapping using unmanned air vehicles. In *Proc. IEEE/RSJ Conf. Intelligent Robots and Systems*, pages 2782–2788, Lausanne, Switzerland, October 2002.
- [18] D. Lawrence, K. Mohseni, and R. Han. Information energy for sensor-reactive uav flock control. In *3rd AIAA Unmanned Unlimited Technical Conference, Workshop and Exhibit*, September 2004. AIAA paper 2004-6530.
- [19] G. Lu, B. Krishnamachari, and C. Raghavendra. Performance evaluation of the ieee 802.15.4 mac for low-rate low-power wireless networks. In *Workshop on Energy-Efficient Wireless Communications and Networks (EWCN '04), held in conjunction with the IEEE International Performance Computing and Communications Conference (IPCCC)*, April 2004.
- [20] K. Pahlavan and A. Levesque. *Wireless Information Networks*. John Wiley & Sons, 1995.
- [21] M. Petrova, J. Riihijarvi, P. Mahonen, and S. LaBell. Performance study of ieee 802.15.4 using measurements and simulations. In *Proceedings of IEEE WCNC*, April 2006.
- [22] W. J. Pisano and D. A. Lawrence. Autonomous uav control using a 3-sensor autopilot. In *AIAA InfoTech@Aerospace Conf., Sonoma, CA, May, 2007*.
- [23] T. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall, 2002.
- [24] A. Woo and D. Culler. Taming the challenges of reliable multihop routing in sensor networks. In *ACM SenSys*, November 2003.
- [25] K. Xu, X. Hong, and M. Gerla. Landmark routing in ad hoc networks with mobile backbones. *Journal of Parallel and Distributed Computing (JPDC), Special Issues on Ad Hoc Networks*, 2002.
- [26] K. Xu, X. Hong, M. Gerla, H. Ly, and D. Gu. Landmark routing in large wireless battlefield networks using UAVs. In *IEEE MILCOM*, pages 230–234, 2001.
- [27] Y. J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *ACM SenSys*, 2003.
- [28] J. Zheng and M. J. Lee. *A comprehensive performance study of IEEE 802.15.4*. IEEE Press, Wiley Interscience, 2006.
- [29] G. Zhou, T. He, S. Krishnamurthy, and J. Stankovic. Impact of radio irregularity on wireless sensor networks. In *ACM Mobisys*, June 2004.