

# Implementation of a Wireless Mesh Network of Ultra Light MAVs with Dynamic Routing

Alberto Jimenez-Pacheco\*, Denia Bouhired\*, Yannick Gasser†,  
Jean-Christophe Zufferey†, Dario Floreano† and Bixio Rimoldi\*‡

\*Laboratory of Mobile Communications and †Laboratory of Intelligent Systems,  
École Polytechnique Fédéral de Lausanne, Switzerland

**Abstract**—This paper describes the implementation and characterisation of a mobile ad-hoc network (MANET) of ultra-light intelligent flying robots. The flying nature of the network makes it suitable to collect or disseminate content in urban areas or challenging terrain, where line-of-sight connection between the Micro Air Vehicles (MAVs) allows for more efficient communication. Dynamic routing in the network enables the nodes to intelligently establish multi-hop routes to extend the communication range or to overcome obstacles. The presented MANET relies on the IEEE 802.11n WiFi standard for data communications and uses the OLSR routing protocol. Routing decisions based on dynamic link quality measurements allow the network to cope with the fast variability of the wireless channel and the high mobility of the MAVs. The implementation of such a system calls for the integration of advanced communication and control technologies in a very restrictive platform, be it in terms of weight, power consumption or availability of suitable off-the-shelf hardware. A detailed description of the system design is presented, and its performance is characterised based on in-flight network measurements. To the best of our knowledge, this is the first report of OLSR successfully tested in a MANET with such fast dynamics. We verify the trade-off between achievable throughput and the number of hops, and we report on the sensitivity of communication performance and routing behaviour to MAV orientation and flight path. Mitigation of such dependencies and improvements to the routing algorithm are discussed along with future research directions.

## I. INTRODUCTION

In the last decade, the development of flying robots has made enormous progress, driven by a rising demand, not only in military applications, but especially thanks to new civilian applications, both scientific and commercial. Among many other uses, unmanned aircraft systems have been successfully employed for climate observation, pollution measurements, fire control, traffic monitoring, surveillance, or mapping [1].

The system of MAVs presented in this work is used to autonomously establish wireless networks in outdoor areas, such as urban environments or disaster sites of difficult access, and employs ultra-light MAVs to achieve the high manoeuvrability and fast deployment essential for such scenarios. The reduced dimensions and weight (under 500 g.) of the MAVs employed make them safe for third parties in case of accident,

so in many countries they can be flown without specific authorisation. However, one trade-off for such MAVs is that large batteries or heavy radio equipment cannot be mounted on-board, limiting flight autonomy and communication range. Furthermore, our implementation aims at keeping costs to the minimum by using only commercial off-the-shelf (COTS) components, which further restricts the design choices.

Our system employs standard 802.11 WiFi cards for wireless communication, mostly due to the availability of inexpensive components with very small form factor and weight. However, 802.11 was mainly designed targeting indoor environments, where the most significant contribution to Doppler effect arises from people moving around essentially static devices. Therefore, WiFi performance can degrade significantly when used in mobile environments [2]. The fact that WiFi operates in the freely usable ISM (Industrial, Scientific and Medical) frequency bands makes it attractive, but also comes with negative implications, such as interference and a very reduced maximum transmit power, limiting the achievable communication range in free space to around 500 m. On the other hand, aerial robots are often required to stay within communication range of a ground station as a safety mechanism.

In order to support high data rate applications (such as video streaming, VoIP or file transfers), and thereby increase the areas of application of WiFi-equipped MAVs, it is essential to increase the communication range offered by WiFi. To this end, multiple MAVs can be employed in a mesh network, where each node is not only responsible for capturing and disseminating its own traffic, but also acts as a router, relaying data for other nodes as needed by the network. The goal is to build a self-healing network, with routing algorithms that allow reconfiguration around broken links.

Traditional routing algorithms such as shortest-path (minimum hop count) are designed for cabled networks under a strong link quality assumption (i.e., either there is a very good connection, or there is no connection at all). Thus, they are not well suited for wireless networks, as they may ignore longer paths that could nonetheless offer higher throughput [3].

A common approach to design routing algorithms for MANETs is to split the problem in two stages: firstly, get a measure of the quality of individual point-to-point links; and then, use that information to find optimum routes with the lowest combined cost [4]. Several metrics can be considered

This work has been supported by armasuisse, competence sector Science+Technology for the Swiss Federal Department of Defense, Civil Protection and Sport.

‡Corresponding author, e-mail: bixio.rimoldi@epfl.ch

to characterise individual links, based on signal-to-noise ratio (SNR) [5], expected transmission count (ETX) [6] or expected transmission time (ETT) [7], [8], offering different trade-offs between implementation cost and accuracy.

The selected routing algorithm for our system is OLSR (Optimized Link State Routing) [9], a proactive routing algorithm that seeks to maintain a constantly updated understanding of the network topology, well suited for MANETs with very dynamic topology changes. Although OLSR can be used with different metrics, we focus solely on the ETX metric, which aims to find paths with high throughput by taking into account the packet loss ratios of the links.

The rest of this paper is organised as follows. Section II provides a description of the complete system, followed by an in-depth discussion of the communication aspects. Dynamic routing for MANETs is discussed in Section III, describing the functioning of OLSR and justifying its choice for our system. In Section IV we present and analyse the experimental results obtained during in-flight tests. We draw our conclusions and discuss future research directions in Section V.

## II. SYSTEM DESCRIPTION

The presented system is an heterogeneous MANET which integrates both MAVs and on-ground stations (laptops). Each MAV navigates autonomously, using distributed control and on-board sensors. For wireless communications all nodes are equipped with WiFi cards supporting the 802.11 ad-hoc mode. The set-up of the network can be initiated by any node, and nodes may join or leave seamlessly at any time.

### A. Flying platform

We use the fixed-wing MAV developed at the Laboratory of Intelligent Systems at EPFL, a basic diagram of which is shown in Fig. 1. With a wingspan of 80 cm and built on expanded poly-propylene (EPP), its total weight is approximately 420 g. The drone is propelled by an electrical motor mounted on the rear end, and a lithium polymer battery provides a flight autonomy of 60 minutes. The drone cruise speed is 10 m/s.

The drones are equipped with two main electronic sub-systems: an autopilot and an embedded computer. Both are integrated in the EPP body. The autopilot uses a dedicated DSP to implement flight control strategies. It integrates a GPS unit, pressure sensors and inertial sensors. The autopilot enables autonomous take-off, followed by way-point navigation at preset altitudes, and autonomous landing when triggered [10].

The embedded PC is a Gumstix Overo Tide computer-on-module (COM) running a distribution of Linux targeted for embedded devices. While the autopilot is responsible for flight control, the COM is responsible for mission control, including data logging, WiFi communications, camera control, etc. Its small dimensions ( $58 \times 17 \times 4.2$  mm) and weight (4.3 g) make the Overo Tide ideal for integration in our MAVs. The Gumstix COM sits atop a small expansion board that provides the necessary connectivity to attach the USB WiFi card and other peripherals to the COM.




Fig. 1. Flying platform.

An interface exists for exchanging commands and data between the autopilot and the embedded PC. Further details about the MAVs used in our system can be found in [10].

### B. Communication Systems

The MAVs are equipped with two communication systems:

- The first one, based on an XBee Pro 802.15.4 radio from Digi International operating in the 2.4 GHz band, is used exclusively for *control traffic*. This includes navigation instructions exchanged between the control ground station and each of the MAVs. This link can meet the requirements of small bandwidth (up to 250 kbps), long range (up to 1.2 km) and limited delay required by control traffic. This control link follows a simple point-to-multi point topology, i.e., MAVs cannot communicate between them using IEEE 802.15.4.
- The second one, used for *data transmission*, is based on WiFi and follows a mesh topology. WiFi can offer higher data rates, as required by multimedia applications, but with a reduced communication range.

The decision to have two separate communication systems is motivated by the orthogonal requirements imposed by *control* and *data* traffics, and by the lack of robust mechanisms in WiFi for traffic prioritisation. For the rest of the paper we discuss only the WiFi communication system.

#### IEEE 802.11 WiFi communications for MAVs

WiFi has been selected as the technology for wireless data transmission due to the availability of inexpensive COTS components with minimal weight and reduced power consumption.

The IEEE 802.11 standard offers support for an ad-hoc mode, where two or more peer nodes can set up a mesh network without relying on any infrastructure. However, the flexibility of this ad-hoc mode is rather limited, since it offers no support for routing: nodes in the network can only communicate directly, if they are within each other's communication range. Nevertheless, support for intelligent routing can be added on top of the 802.11 link layer by using OLSR or other routing protocols, as described in Section III.

A multi standard (802.11 a/b/g/n) WiFi USB card from Sparklan with the Ralink 3572 chipset has been selected for the flexibility of its Linux driver in configuring the physical layer and because it is dual-band, being able to operate in the 2.4 GHz and, more interestingly, in the 5 GHz band. In the current implementation it was decided to operate the WiFi

system exclusively at 5 GHz, to avoid possible interference to/from the XBee control link and from other sources that affect the more crowded 2.4 GHz band.

We use the most recent 802.11n variant of WiFi, which employs OFDM and can exploit multiple transmit/receive antennas using MIMO techniques. Indeed, the WiFi card selected integrates two planar antennas. Still, with its reduced dimensions ( $65 \times 25 \times 2$  mm) and weight (7 g), it has a minimal impact on the aerodynamics of the MAV. The WiFi card is integrated on the left wing of the MAV (see Fig. 1).

Given the harsh communication channel that the MAVs will experience, the WiFi configuration has been driven by robustness criteria. Of all the available Modulation and Coding Schemes (MCS) defined in IEEE 802.11n, the one with the second lowest physical data rate (13 Mb/s) has been selected, which employs a robust QPSK constellation and a rate 1/2 channel code (lowest spectral efficiency but strongest error correction). This data rate is fixed in our setup.

The MIMO configuration selected uses Alamouti coding [11] to provide spatial diversity and hence additional protection against fading.

All nodes in the network operate in the same channel (5.24 GHz) and share exactly the same WiFi settings.

### III. DYNAMIC ROUTING

#### A. Dynamic routing in ad-hoc networks

The IEEE 802.11 WLAN standard only specifies the physical and link layers of the OSI reference model. Hence, it makes no provision for routing in its ad-hoc mode. Static routing tables can be used to partly overcome this limitation, although this approach is inherently inflexible for a MANET of MAVs, where topology changes will be frequent and unpredictable.

Dynamic routing takes the concept of mesh network one step further: not only should each node be able to receive, store and forward packets to its neighbours to effectively build multi-hop routes, but nodes should also be able to take routing decisions in real time, reacting to the changing conditions of the network. The main aim of any dynamic routing protocol is to find the *best* possible route and inform the affected nodes about this decision, trying to keep the delay and the control traffic induced by this process to the minimum. Depending on how the building and maintenance of the routing tables is handled, routing protocols can be classified as:

*a) Reactive:* Only when a request for transmission is initiated, does the routing protocol send the necessary control packets to explore the network and find the best possible route to the destination. Although the traffic overhead is reduced, this approach implies long delays at the beginning of a communication instance. Reactive routing protocols have been dominant in wireless sensor networks (with no mobility), and a widely used example of such protocol type is AODV (Ad-hoc On-demand Distance Vector) [12].

*b) Proactive:* Nodes periodically generate control traffic in order to determine the existence of routes between all the nodes present in the network. Every node announces its connectivity state to the entire network and maintains a

complete understanding of its topology. Although this method incurs in additional traffic overhead, it also enables faster response when a new route is requested. An example of proactive protocol is OLSR, discussed in more detail below.

#### B. Optimized Link-State Routing (OLSR)

For our application OLSR has been selected as routing algorithm, specifically the OLSRd (OLSR Daemon) implementation [13]. This choice has been motivated by several factors; most importantly, the need of a proactive algorithm, given the high mobility of the MAVs. Secondly, OLSRd is physical and link layer agnostic, leaving complete freedom for the selection of the MAC communication protocol or hardware, without requiring any driver modifications. OLSRd simply modifies the entries of the kernel routing tables in a way that is transparent to the user. Last but not least, OLSRd is an open source project, distributed under a BSD-style license.

OLSR relies on link-sensing techniques to establish or dismiss possible network paths. In a collaborative approach, nodes inform one another about all the neighbours in their respective ranges, resulting in the *continuous update* of the routing tables maintained in each node. OLSR uses Dijkstra's algorithm to select minimum cost routes. Important is that every node keeps a table with the next hop for the routes to *all* other nodes in the network. This implies a nearly instantaneous ability to determine the best path to the destination when a link is requested or when there are changes in the network.

OLSR generates control traffic in order to explore the network and to update its members about the current topology. From the various types of control packets, two should be highlighted:

- The task of link-sensing is performed by the periodic transmission of HELLO packets, which results in the building of the *link set* for each node (a list of its neighbours and the associated link qualities).
- The task of topology discovery is maintained through the emission of TC (topology control) messages. These are used by nodes to declare their list of neighbours, thereby propagating the topology information of the network to all member nodes. OLSR designates specific nodes to act as Multi-Point Relays (MPR). Only MPRs are responsible for forwarding control traffic intended for diffusion in the entire network, providing an efficient mechanism for controlled flooding which avoids redundancy and reduces the number of transmissions required [9].

#### Link-state metrics

In our MANET, OLSRd with the link quality extensions has been adopted to take routing decisions based on the ETX metric. The ETX of a link is the expected number of MAC layer transmissions needed to successfully deliver a packet over that link [6]. Each node broadcasts HELLO packets at regular intervals; the percentage of them correctly received from a neighbour during a sliding time window constitutes the Link Quality (LQ) of that link. The symmetrically equivalent value, denoted as NLQ (Neighbour Link Quality), represents

the probability that a HELLO message that we send is correctly received by that particular neighbour. For a measure of the *round-trip* link quality, the two values are combined to represent the probability  $p$  that a transmission is successfully received and correctly acknowledged,  $p = LQ \times NLQ$ . The number of trials before successful transmission is thus a geometric random variable of parameter  $p$ , and its mean is the expected transmission count,  $ETX = 1/(LQ \times NLQ)$ . In practice, an exponential averaging can be employed to give more weight to the most recent HELLO packets.

For multi-hop routes, the aggregate ETX is simply the sum of the ETX of each link in the route. The summation reflects the fact that a relay node in a multi-hop route receives and transmits data in a sequential manner, in order to avoid interference between successive hops of the same path (*intra-flow* interference). On the other hand, nodes do not only contend for bandwidth with other nodes in the same path, but also with nodes in geographic proximity belonging to other paths. In order for this *inter-flow* interference to be considered for route selection together with link quality, the OLSR protocol ought to be extended to incorporate geographical information.

Finally, it must be mentioned that for high mobility to be taken into account, OLSRd shall be configured to propagate routing information with the necessary urgency. Ideally, OLSRd configuration parameters such as the interval between HELLO packets and the exponential ageing factor should be adjusted depending on the network size, node speed and expected mobility patterns to achieve the best performance. Still, there will always be a trade-off between the accuracy of the link measurements and the responsiveness to mobility.

### C. The new 802.11s amendment for mesh networks

While not considered for this work, we must briefly mention the recently ratified (Q4 2011) 802.11s amendment to the IEEE 802.11 standard, which specifically deals with mesh networks. Changes to the PHY layer are not required, since 802.11s inherently depends on one of 802.11a/b/g/n to carry the actual data. 802.11s extends the MAC standard defining an architecture and protocols to support self-configuring multi-hop topologies. The standard specifies a default routing protocol known as HWMP (Hybrid Wireless Mesh Protocol), which combines the flexibility of reactive routing with proactive topology tree extension. The reactive mode of HWMP is known as RM-AODV (Radio Metric-AODV), and is essentially a modified version of AODV that uses the Air-time Link Metric (ALM) instead of hop count for path selection. This metric takes into account the packet loss probability as well as the bit rate of the link. The amendment further defines Radio-aware OLSR as an optional routing protocol, identical to OLSR as described above, but adapted to use the new ALM metric instead of ETX. We refer the reader to [14] for more details about 802.11s and the routing algorithms specified in it.

Although support for 802.11s in COTS components and software is still limited, the promising benefits of the hybrid protocol will certainly be evaluated in future work.

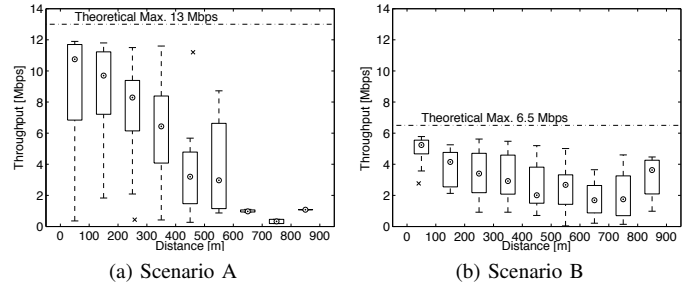


Fig. 2. Box plots of the end-to-end throughput as a function of the distance between GND and the destination MAV for (a) one-hop scenario, (b) two-hop scenario (OLSRd disabled, relaying enforced with static routing).

## IV. SYSTEM EVALUATION AND PERFORMANCE

A series of field tests have been carried out to assess the functionality and performance of the described system. All tests were conducted in an open-air, flat, rural environment, without any major obstructions in vicinity. In all setups a ground operator conducted the measurements using a Linux laptop equipped with the same WiFi card model installed on the MAVs. The configuration of the WiFi parameters is as described in Section II-B.

### A. Throughput analysis

We use the standard tool `iperf` working in TCP mode to measure the error-free throughput available at the application level, including the effect of retransmissions. We consider two different scenarios:

- *Scenario A.* The network consists of two nodes, the fixed ground station (GND) and a drone (MAV<sub>1</sub>) that flies on a straight line, getting away from it or back towards it.
- *Scenario B.* The network consists of three nodes: the ground station and two MAVs. The second MAV (MAV<sub>2</sub>) is used as a relay between GND and the destination MAV (MAV<sub>1</sub>). We characterise the end-to-end throughput of the path between the ground station and the destination MAV by letting the relaying MAV describe circular waypoints with a radius of 50 m, centred halfway between the ground station and the destination MAV.

In both scenarios, the MAVs fly at a constant altitude of approximately 100 m above the ground.

*Throughput vs. Distance:* In Fig. 2 we present the measured throughput as a function of the distance between the ground station and the destination MAV. In these box plots, distances have been quantised with a bin width of 100 m. The circle inside each box represents the median throughput, and the lower and upper ends of each box represent the 25th and 75th percentiles, respectively. The whiskers represent the smallest and largest observations for each distance bin. For Scenario A, Fig. 2a shows a clear inverse dependency between the median achievable throughput and distance. Although the physical bit rate is 13 Mb/s, even in the best conditions we can only expect a throughput slightly below that limit, due to the MAC layer overhead. Beyond 600 m, very few measurements achieve a significant fraction of the maximum 13 Mb/s. Despite this




Fig. 3. Trajectory followed by the MAV in the one-hop scenario. The size of the circles is proportional to the throughput measured at each location.

general trend, even at reduced distances, the throughput measured in different trials shows a big spread, as reflected by the height of the boxes. For instance, around 100 m we recorded many points achieving close to 12 Mb/s, but more than 25% of the points are below 7 Mb/s. This already points out to a more complex dependency of the throughput on the flight parameters of the MAVs, not only on distance, even if that is the variable that shapes the general behaviour.

For Scenario B, Fig. 2b shows the maximum achievable throughput in the two-hop configuration *when relaying is enforced*: OLSRd was temporally disabled and static routing set up so that MAV<sub>2</sub> continually acted as a relay between GND and MAV<sub>1</sub> during this test. The plot shows that the maximum achievable throughput has been reduced by 50% compared to Scenario A. This is consistent with the fact that ours is a single-frequency network, so the relaying MAV cannot receive packets from the ground station and at the same time forward traffic to the destination MAV without collision in the channel. This *intra-flow* interference effectively halves the maximum throughput, unless we used multi-radio MAVs. However, this would reduce the autonomy of the MAVs, increase their cost, and complicate the optimisation of dynamic routing.

Comparing Fig. 2a and Fig. 2b, we observe that the addition of the relaying MAV allows to considerably increase the distances for which we have a consistent link. While the maximum throughput has been reduced, there are many more successful measurements beyond 600 m, many of them achieving a significant fraction of the 6.5 Mb/s theoretical maximum. In this two-hop configuration, the dependency of throughput with the total distance is more difficult to assess because each hop behaves better, and we would have to send the destination MAV further away (more than 1 km) to observe more clearly the decreasing trend. In any case, it is apparent that adding the intermediate relay node helps to extend the range from 400 m to about 800 m.

*Throughput vs. Trajectory*: Figure 3 analyses in more detail the dependency of throughput on the trajectory followed by the MAV in the one-hop scenario. This plot represents the position of the MAV in a two-dimensional reference system with origin in the fixed position of the ground station (GND). For each *iperf* measurement a circle with diameter proportional to the measured throughput is displayed at the point where the




Fig. 4. Link Cost (ETX) measured by OLSRd as a function of distance (Scenario A).

measurement was taken. Instead of following a straight line, this time the MAV described circular trajectories of  $\approx 50$  m radius, with the centre of the circles located progressively further away from the origin. With our reference system, waypoints are described in counterclockwise direction. The figure shows a marked dependency of throughput on distance, but also on MAV orientation: after a certain distance, the link behaves better when the plane is facing GND than when it is flying away from it. This behaviour could be explained by the antenna radiation pattern, or by the location of the WiFi card on the left wing of the plane, with the motor and electronic components shadowing the communication path depending on the relative position and angle between the MAV and GND. Even at short distances the orientation of the WiFi antenna plays an important role, with significantly different performance depending on the part of the circular waypoint the MAV was describing when the measurement was performed. In view of these results, the possibility of finding a different location for the WiFi card, or the use of external antennas should be further studied and its placement optimised for a more uniform performance.

### B. OLSRd behaviour analysis

A monitoring script was written to evaluate the fluctuations of the ETX metric in the two scenarios defined, and to correlate the results with the previous throughput measurements. The aim was to study the stability of the routing protocol as well as its behaviour relative to distance and MAV trajectory.

*OLSRd Link Cost in Two-Node Scenario*: An initial measurement of the link quality is conducted in the settings of Scenario A, with the MAV flying in a straight line between GND and a distant waypoint at about a 1000 m distance. The link cost as a function of the distance is plotted in Fig. 4 for both directions of the flight. We note that  $ETX = 1$  corresponds to a perfect link (as observed for distances below 300 m), and a value of infinite means that in the sliding time window that OLSRd considers for computation of the metric, no correct HELLO packets were received. Values of ETX greater than 4–5 make the link hardly usable (starting approximately from 500 m distance). We observe that, except for very small distances, one direction (when the plane is flying back towards the ground station) behaves systematically better than the other one. This is consistent with the dependency




Fig. 5. Relay status and trajectory followed by destination MAV ( $MAV_1$ ) in Scenario B. ( $MAV_2$  described circular way-points of radius 50 m, with center half-way between GND and the center of the way-points described by  $MAV_1$ ).

of the throughput on antenna orientation and position already observed.

*OLSRd Link Cost in Three-Node Scenario:* In Scenario B, the destination MAV ( $MAV_1$ ) describes 11 circular way-points with an approximate radius of 50 m and centre progressively further from the the ground station, following a straight line, up to a distance of 1000 m. Figure 5 shows the flight path described by  $MAV_1$ . The trajectory followed by the relay MAV ( $MAV_2$ ) is not shown in this figure, but it also described circular way-points of radius 50 m, with centre located at the mid-distance point between the origin and the centre of the way-point described by  $MAV_1$  at each moment<sup>1</sup>. With OLSRd active again, whether  $MAV_2$  acts as a relay between GND and  $MAV_1$ , or it is ignored by the routing protocol in favour of a direct path between GND and  $MAV_1$ , will depend on the link quality values measured.

Figure 5 shows the relay status with different marker symbols. It can be seen that the point at which relaying starts also depends on the aircraft orientation: while relaying starts at about 400 m when  $MAV_1$  is flying in the outbound direction, direct communication begins at about 600 m in the inbound direction, showing that the direct communication link with the ground station is more robust with the plane flying towards it. This behaviour is consistent with our previous observations and it also shows that relaying is not an on/off state. Instead, the dynamic nature of the algorithm can make the chosen path change relatively quickly. There is an intermediate zone (between 300–600 m) where the link quality is not consistently good or consistently bad, but it changes around a single way-point, hence the relay status is *jittery*. Outside this zone the direct link is either consistently good (robust one-hop communication) or consistently bad (robust two-hop communication).

## V. CONCLUSIONS

In this paper we have presented the design of a single-frequency MANET composed of ultra-light MAVs and fixed on-ground nodes. Completely based on standardised wireless technologies in the ISM band, using exclusively commercial

<sup>1</sup>The angular position of both MAVs in their respective circular way-points are not synchronised in any way.

off-the-shelf hardware components and open-source software, it is ideal for the economic and quick deployment over areas without any existing communication infrastructure. We have demonstrated in a small system the feasibility of using dynamic routing with OLSRd to cope with the high mobility of the MAVs and the impairments of the wireless channel.

The communication performance of the system has been characterised through in-flight measurements. We have described the effect of distance and aircraft orientation on the end-to-end achievable throughput and on routing decisions.

Further research is needed aiming at reducing the sensitivity of the performance to MAV orientation and antenna placement. In the future, the interface between the autopilot and the embedded computer should be exploited to guide flight decisions based on communication needs. Having access to the GPS information from the autopilot would allow the position of the relays to be optimised and enable the introduction of new functionality, such as ferrying of information in sparse networks. This would require, however the modification of the routing protocol. This opportunity should be taken to also modify the ETX metric used by OLSRd to account for *inter-flow* interference, which could be an issue in larger networks.

## REFERENCES

- [1] E. W. Frew and T. X. Brown, "Airborne communication networks for small unmanned aircraft systems," *Proc. IEEE*, vol. 96, no. 12, pp. 2008–2027, Dec. 2008.
- [2] P. Mahasukhon, H. Sharif, M. Hempel, T. Zhou, W. Wang, and H.-H. Chen, "IEEE 802.11b based ad hoc networking and its performance in mobile channels," *IET Commun.*, vol. 3, no. 5, pp. 689–699, May 2009.
- [3] D. S. J. de Couto, D. Aguayo, B. A. Chambers, and R. Morris, "Performance of multi-hop wireless networks: Shortest path is not enough," in *Proc. First Workshop on Hot Topics in Networks (HotNets-1)*, Princeton, New Jersey, Oct. 2002.
- [4] G. Parissidis, M. Karaliopoulos, R. Baumann, T. Spyropoulos, and B. Plattner, "Routing metrics for wireless mesh networks," in *Guide to Wireless Mesh Networks*. Springer London, 2009, pp. 199–230.
- [5] C. Dixon and E. W. Frew, "Maintaining optimal communication chains in robotic sensor networks using mobility control," *Mobile Networks and Applications Journal*, vol. 14, no. 3, pp. 281–291, 2009.
- [6] D. S. J. de Couto, D. Aguayo, J. Bicket, and R. Morris, "A high-throughput path metric for multi-hop wireless routing," in *Proc. ACM Int. Conf. on Mobile Computing and Networking*, San Diego, Sep. 2003.
- [7] R. Draves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks," in *Proc. ACM Int. Conf. on Mobile Computing and Networking (MobiCom'04)*, Sep. 2004, pp. 114–128.
- [8] D. Aguayo, J. Bicket, and R. Morris, "SrcRR: a high throughput routing protocol for 802.11 mesh networks (draft)," MIT, Tech. Rep., 2006. [Online]. Available: <http://pdos.csail.mit.edu/rtrm/srcrr-draft.pdf>
- [9] T. Clausen and P. Jacquet, "Optimized link state routing protocol (OLSR)," RFC 3626, Oct. 2003. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc3626.txt>
- [10] S. Leven, J.-C. Zufferey, and D. Floreano, "A simple and robust fixed-wing platform for outdoor flying robot experiments," in *Proc. International Symposium on Flying Insects and Robots*, 2007, pp. 69–70.
- [11] S. M. Alamouti, "A simple transmit diversity technique for wireless communications," *IEEE J. Sel. Areas Commun.*, vol. 16, no. 8, pp. 1451–1458, Oct. 1998.
- [12] C. Perkins, E. Royer, and S. Das, "Ad hoc on-demand distance vector (AODV) routing," RFC 3561, Jul. 2003. [Online]. Available: <http://tools.ietf.org/html/rfc3561>
- [13] OLSRd, An adhoc wireless mesh routing daemon. [Online]. Available: <http://www.olsrd.org>
- [14] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications. Amendment 10: Mesh Networking*, IEEE Std. 802.11s, Oct. 2011.