

Now or Later ? - Delaying Data Transfer in Time-Critical Aerial Communication

Mahdi Asadpour, Domenico Giustiniano, Karin Anna Hummel, Simon Heimlicher, Simon Egli
ETH Zurich
Communication Systems Group, Zurich, Switzerland
{mahdi.asadpour,domenico.giustiniano,karin.hummel,heimlicher,simon.egli}@tik.ee.ethz.ch

ABSTRACT

Search and rescue missions are entering a new era with the advent of small scale unmanned aerial vehicles (UAVs) with communication capabilities and embedded cameras. Yet, delivering high resolution images of the supervised surface to rescuers is *time-critical*. To help resolving this problem, we study how UAVs can take advantage of their controlled mobility to derive the optimum strategy for data transmission. Driven by real-world aerial experiments with both airplanes and quadcopters equipped with 802.11n technology, we show that the UAV should not necessarily transmit as soon as a wireless link is established. Instead, it should wait until it reaches a suitable distance to the receiving UAV, only to transmit when the time to move to the new location and transmit is minimal. We then apply the principle of *delayed gratification*, where the UAV attempts to solve the tradeoff between postponing until it reaches this minimum and the impatience to deliver as much data as soon as possible, before any physical damage on-the-fly may occur. Our empirical-driven simulations demonstrate that the optimal distance of transmission greatly depends on the interplay of actual throughput, data size, UAV cruise speed, and failure rate, and that state-of-the-art UAVs can already benefit from our approach.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless communication*

Keywords

Unmanned Aerial Vehicles, Delayed Gratification, IEEE 802.11n, Measurements

1. INTRODUCTION

Search and rescue (SAR) missions are known to be time-critical as human lives are affected. Here, technological advances could be put to great use. Among these potential technologies, small scale unmanned aerial vehicles (UAVs) and wireless communication have the potential to greatly improve quality and quantity of data available to human rescuers. A fundamental challenge in this

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CoNEXT'13, December 9-12, 2013, Santa Barbara, California, USA.
Copyright 2013 ACM 978-1-4503-2101-3/13/12 ...\$15.00.
<http://dx.doi.org/10.1145/2535372.2535409>.

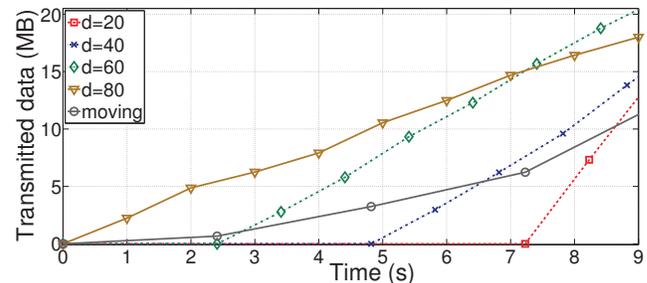


Figure 1: Experimental measurements of transmitted data. One UAV is originally 80 m away from another hovering UAV. It may either immediately send 20 MB of data (case ‘ $d = 80$ m’), or transmit while moving closer (‘moving’), or move closer to the hovering UAV and transmit only after reaching the new position (d). Here, waiting to transmit at a distance of $d = 60$ m outperforms other strategies.

context is the transfer of recordings taken by UAVs’ on-board cameras to SAR headquarters in a timely manner. This communication problem differs from other research areas such as mesh networks or mobile ad hoc networks in important ways: (i) UAVs are robots, and as such the mobility of UAVs is at least partially *controllable*, *i.e.*, if beneficial, their position or trajectory can be adjusted to improve performance, (ii) the period during which UAVs remain in action is limited by battery capacity and harsh environments.

A mobile aerial network that provides control over some of the nodes’ mobility poses unique challenges to radio communications due to cruising at high speeds and positioning dynamics which entail an unreliable communication channel. To illustrate how controlled mobility can influence this communication, consider two UAVs, one hovering, the other carrying data to deliver and just coming in communication range. The UAV may transmit immediately. Yet, transmitting as soon as possible does not necessarily mean to transmit in the shortest time. The UAV may follow a different approach, and move closer to the target UAV, *without transmitting* until it reaches the desired distance and better communication conditions. Figure 1 shows the impact of this decision in a real experiment: as shown by the crossing point between the decision to transmit at a distance of $d = 80$ m and to move and transmit at a distance of $d = 60$ m, the latter strategy reduces the overall time (communication delay) needed to transfer all the data, as long as the total data size to transfer is larger than ≈ 15 MB. The intuition is that the time saved transmitting at the new location exceeds the time spent to reach this position. Moving closer than $d = 60$ m does not help the UAV to deliver its data faster. Finally, because of the unfavorable channel conditions under mobility, transmitting

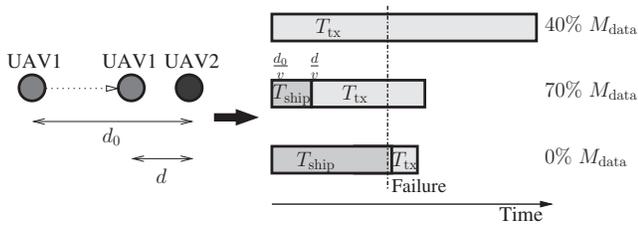


Figure 2: Delivery of collected data (M_{data}): In the example, we compare three strategies: (i) on the top, UAV1 transmits the data immediately at a distance of d_0 ; only transmission time incurs, (ii) in the middle, UAV1 ships the data to distance d , and then transmits; shipping and transmission times incur, (iii) in the bottom, UAV1 moves even closer to reduce the communication delay. Yet, a failure occurs and therefore the best strategy is case (ii), where most of the data is transmitted (70 %).

data while moving towards the target UAV is outperformed by the previous strategies.

Still, delaying the transmission may increase the probability of any interruption of the delivery of data, e.g., caused by hazardous conditions to operate. It may be better for the UAV to transmit its carried data sooner than required to minimize the communication delay. Delaying until the best conditions are encountered is known as *delayed gratification*, where the decision maker holds on to the desire to obtain an immediate gratification for the sake of a higher reward at a later point in time, a concept widely used in economics, physical health, and social competence [1, 2]. We express it as the temptation to deliver the data as soon as a link becomes available to reduce the probability of any operational failure versus the gratification it would obtain waiting until the best rendezvous point has been reached in terms of short communication delay.

Based on the above foundations, our contributions are summarized as follows:

- Motivated by real-world experiments, we introduce a model to express the delayed gratification problem. We derive the essential parameters that the UAV must factor in to compute the optimal strategy (Section 2).
- We characterize aerial communication with various experimental tests, using two heterogeneous types of UAVs, i.e., airplanes and quadcopters (Section 3).
- Driven by the experimental data, we study our strategy for delivering the carried data as a function of the data size, cruise speed, and the failure rate caused by any damage (Section 4).

Our simple model is not the final solution to the delayed gratification problem faced by UAVs. Yet, it provides significant insights and directions for investigations. The main conclusion of this contribution is that there exists a sweet spot where the UAV obtains the highest gratification, according to the interplay of actual throughput, failure rate, and system parameters. Finally, to the best of our knowledge, we are the first that provide an in-depth experimental comparison of different classes of aerial vehicles in terms of their communication capabilities.

2. DELAYED GRATIFICATION

In this section we model the strategy to deliver data. There is a fundamental three-way tradeoff:

- The UAV should select a path that is close to optimal for its sensing task.

- Because the mission is time-critical, the UAV should attempt to *minimize the delay* of delivering all of the gathered data, which is composed of the time required to (i) move to a specific location and (ii) transmit the data.
- UAVs have some probability to fail due to, for example, bad weather conditions or physical collisions. Therefore they would like to *deliver data as soon as possible*.

In this paper, we focus on the second and third aspects of this trade-off. In particular, we consider that the UAV has collected a batch of data (e.g., all the images of some area). The tradeoff we consider is depicted in Figure 2.

In order to address these conflicting needs, we propose a delayed gratification model to deliver the batch of gathered data. We formalize the problem by introducing a utility function $U(d)$ to drive the decision to the optimum, where d is the distance between two UAVs at the time of transmission. Further, we define the communication delay $C_{\text{delay}}(d)$ as the time necessary to ship and transfer the data at a distance of d (either to another UAV or to the ground station).

Based on delayed gratification theory [2], we express the utility $U(d) = \delta(d) \cdot u(d)$, where $u(d)$ is the *instantaneous utility* and $\delta(d)$ the *discount function* of the delayed gratification problem:

- $u(d) = 1/C_{\text{delay}}(d)$ represents the benefit of a UAV to transmit at a distance of d . With infinite lifetime, the UAV achieves the highest gratification when $C_{\text{delay}}(d)$ is minimal.
- With finite lifetime, there exists a reward discounting $\delta(d)$ at a distance of d , that reflects the failure probability due to any physical cause. This will cause the UAV to transmit *sooner* than in ideal conditions.

Given the lack of empirical measurements in the literature for modeling the failure rate of operational micro UAVs, we assume that the failure probability is exponentially distributed with the distance traveled [3]. It follows that the probability of being functional while moving to some location is $\delta(d) = e^{-\rho(d_0-d)}$. Here, d_0 is the distance at which the UAV comes in range of another UAV and it is ready to transfer, and $\rho \geq 0$ is the failure rate that determines how likely it is that the UAV will fail. This yields:

$$U(d) = \delta(d) \cdot u(d) = \frac{e^{-\rho(d_0-d)}}{C_{\text{delay}}(d)}. \quad (1)$$

Note that the failure rate is distance-independent, meaning that the optimal strategy to send the data is stationary.

2.1 Formulation

The highest gratification is obtained by waiting to transmit until distance d_{opt} , where $U(d)$ is maximized:

$$U(d_{\text{opt}}) = \max_d U(d), \quad (2)$$

s.t. $0 \leq d \leq d_0, v > 0, M_{\text{data}} > 0.$

Eq.(2) maximizes the utility function by finding the optimal distance d_{opt} for transmission given the constraints of the initial distance d_0 , the UAV speed v , and the traffic demand M_{data} (that is, the total size of the data collected).

2.2 Modeling Communication Delay

Our attempt is to abstract the essential features of the communication delay. Once M_{data} has been collected, communication between two UAVs may occur at any time (note that collection and subsequent communication can happen multiple times before the

mission ends). Two basic strategies can be followed, i.e., ‘move and transmit’, where the UAVs move and their distance decreases while transmitting, and ‘hover and transmit’, referring to the case that the distance remains constant while transmitting. Finally, mixed strategies combining the basic strategies can be derived. Based on the insights of Figure 1, strategies based on ‘move and transmit’ are outperformed by ‘hover and transmit’. While mixed strategies could further reduce the communication delay, in order to simplify the treatment of the model, in the following we solely consider the strategy ‘hover and transmit’.

We divide the area of interest into *sectors* of size A_{sector} , where one UAV is exclusively responsible to sense and gather data, e.g., taking images. Each picture snapped by the UAV covers a much smaller area of A_{image} ¹. We decompose $C_{\text{delay}}(d)$ into two additive parts $C_{\text{delay}}(d) = T_{\text{ship}} + T_{\text{tx}}$ (see also Figure 2).

Shipping time. T_{ship} : the time it takes the UAV to move into a suitable position $d \leq d_0$ for transmission², at a speed of v . This is given by:

$$T_{\text{ship}} = \frac{d_0 - d}{v}.$$

Transmission time, T_{tx} : the time between the first packet of the acquired batch of images is in the head of the queue ready for transmission and the time when the last packet has been successfully delivered. It is given by:

$$T_{\text{tx}} = \frac{M_{\text{data}}}{s(d)},$$

where $s(d)$ is the throughput in Mb/s at a distance of d , and speed v close to zero. The optimal distance d_{opt} depends on:

- the amount of M_{data} acquired. M_{data} must be selected as a trade-off between detection possibility and network load caused. The sector is scanned by taking a total of $A_{\text{sector}}/A_{\text{image}}$ pictures, with each picture of size M_{image} . This results in $M_{\text{data}} = \frac{A_{\text{sector}}}{A_{\text{image}}} \cdot M_{\text{image}}$.
- the cruise speed v of the UAV when approaching either another UAV or the ground station: The cruise speed v is controlled and locally measured by the UAV sensors.
- the rate $s(d)$ at which M_{data} can be transferred as a function of the distance between sender and receiver and the interference present at the wireless medium. $s(d)$ reflects the quality of the aerial link, which is experimentally studied in the next section.

3. EMPIRICAL STUDY

The basic determinant for delayed gratification is how the bandwidth between the two nodes varies as a function of their distance. In this section we investigate this correlation and examine the main factors that characterize aerial communication. Given the need for high-speed communication to deliver a batch of collected data, we study how practical it is to use 802.11n for air traffic between small scale UAVs. In the tests, a control channel between the ground station and every UAV is maintained, based on XBeePro 802.15.4 operating in the 2.4 GHz frequency band. This channel provides low bandwidth (up to 250 kbps) but long range (up to 1.5 km), and it is reserved for (i) light-weight telemetry data consisting of UAV

¹A picture is rectangle with an aspect ratio k , and the field of view (FOV) is the diagonal of the rectangle. Hence $A_{\text{image}} = (k \cdot \frac{\text{FOV}}{\sqrt{k^2+1}}) \cdot (\frac{\text{FOV}}{\sqrt{k^2+1}})$.

²It is straightforward to demonstrate that it is never convenient for a UAV to move further away from another UAV, and therefore we neglect the case $d > d_0$ in the rest of the analysis.

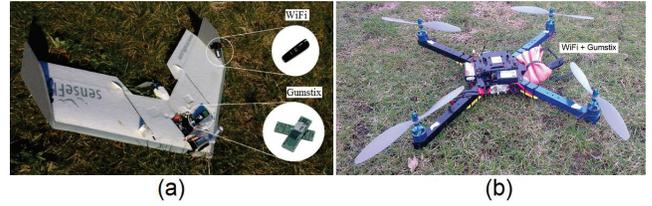


Figure 3: Our flying platforms, (a) airplanes, and (b) quadcopters.

status information (GPS coordinates, speed, etc.) sent to the central planner that resides at the ground station and (ii) new waypoints from the planner to the UAVs.

Flying Platforms. We employ two types of flying platforms, airplanes and quadcopters, which come with different characteristics in terms of hovering capability, cruise speed, weight, etc. Our airplane platform – *Swinglet* [4] – is shown in Figure 3(a). An electrical motor drives its single propeller. Our quadcopter platform – *Arducopter* [5] – has electronics similar to the Swinglets. Four electrical motors drive its propellers and the main electronic system is based on an Arduino board.

The general advantage of quadcopters over airplanes is that they can hover at a waypoint, which helps to perform controlled tests. Airplanes normally cannot hover and have to circle around a waypoint to mimic the behavior, in our case, circle with a radius of at least 20 m. However, our quadcopters travel at lower altitudes, speeds, and are heavier (due to the heavy chassis, it might not be as safe as the swinglets, even at low altitudes). A shortcoming of both types of UAVs is the limited battery life. The main features of the two flying platforms are summarized in Table 1.

Wi-Fi 802.11 Communication. We integrate an embedded system with a Wi-Fi module in our flying platforms. To guarantee fair conditions for our aerial communication tests, we use the same communication package in both the airplane and the quadcopter platform. This package consists of a Gumstix Overo computer-on-module running a Linux distribution and an expansion board to connect the Wi-Fi interface. Among the multiple Wi-Fi interfaces we evaluated, we present the results using a Linksys 802.11n USB with the Ralink 3572 chipset, which comes with two integrated planar, omni-directional antennas and shows higher performance and flexibility than others. We enable features such as channel bonding, A-MPDU frame aggregation, and block ACK. We use a channel bandwidth of 40 MHz and a guard interval of 400 ns. The default number of frames for aggregation is 14. If the physical rate is too high, the embedded system may not fill the buffer fast enough, resulting in a lower number of A-MPDU sub-frames.

At boot up, the UAVs establish an ad hoc network with one another and the laptop/ground station at channel 40. A 5 GHz frequency channel is selected to avoid interferences with XBeePro at 2.4 GHz as it is reserved for critical messages.

Table 1: Main features of our flying platforms.

	Airplane	Quadcopter
Hovering	No	Yes
Size	Wingspan: 80 cm	Frame: 64 cm by 64 cm
Weight	500 g	1.7 kg
Battery autonomy	30 minutes	20 minutes
Cruise speed	10 m/s	4.5 m/s in auto mode
Maximum safe altitude	300 m	100 m

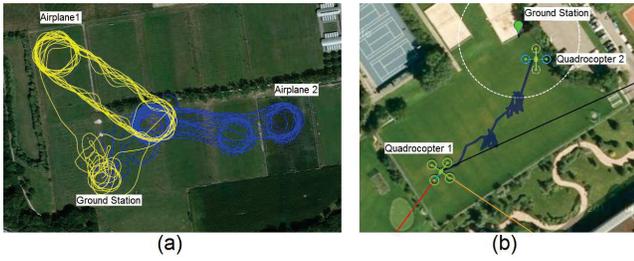


Figure 4: GPS traces of (a) two airplanes, moving between two waypoints with relative distances from 20 to 400 m, altitudes around 80-100 m, and (b) two quadcopters, hovering at altitudes of 10 m and different relative distances from 20 to 80 m.

UAVs’ Waypoints. The main electronic system that controls the UAV is the autopilot, which integrates a GPS unit, pressure sensors, and inertial sensors. The autopilot enables it to take-off and land autonomously, as well as to navigate through waypoints. We configure the mentioned airplanes to fly between two far waypoints in a way that they approach each other, exchange data, and leave (see Figure 4(a) for their GPS traces), resulting in relative speeds between 15 and 26 m/s. Since physical collision avoidance is not implemented, we set the waypoints of the airplanes to two different altitudes, 80 and 100 m, respectively. The quadcopters are used in a separate group of tests where they are configured to be/reside/move to a waypoint, and we investigate various relative distances from 20 to 80 m (see Figure 4(b)). We also test other patterns of waypoints, as described later in the section. Because of safety constraints, we limit the altitude to 10 m.

3.1 Experimental Results with Airplanes

To investigate aerial UAV-to-UAV communications, we start with a study of the performance of the wireless link for airplanes.

Throughput. We study the throughput between two flying airplanes, measured using UDP traffic and the iperf tool for different distances between the airplanes, where the distance is calculated applying the Haversine formula to GPS coordinates. The resulting data are visualized as boxplots in Figure 5. The median shows a throughput degradation with increasing distance. At shorter distances, the throughput (≈ 20 Mb/s) represents more the one expected of 802.11g (≈ 24 Mb/s). For comparison, in indoor lab test using 802.11n, we could get ≈ 176 Mb/s, similar to [6].

The lack of sufficient spatial diversity of the aerial channel impedes to effectively utilize the multiple antennas for MIMO communication. However, we expect higher aerial throughput than the 802.11g-like observed, thanks to 802.11n features such as channel bonding, A-MPDU frame aggregation, and block ACK. Next, we discuss potential causes.

Fixed PHY Rate. We now study the impact of the rate adaptation scheme on the throughput by fixing the PHY rate and using

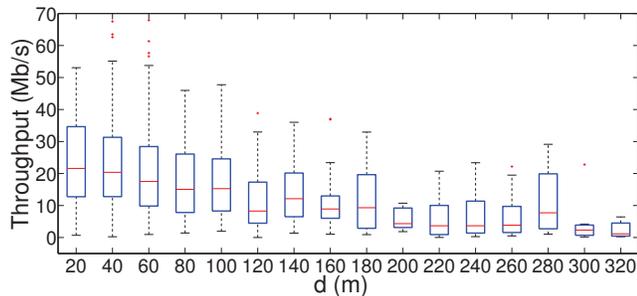


Figure 5: Throughput vs. distance between two airplanes.

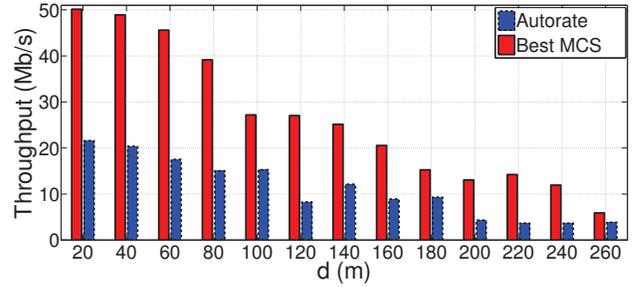


Figure 6: Comparing the best *median* throughput achieved with different MCS values and the result with auto PHY rate, between the two airplanes.

the same waypoints as above. In Figure 6, we compare the best throughput we achieve using different fixed PHY rates. Thus, we select modulation schemes and coding rates, which are represented by a modulation and coding scheme (MCS) index value such as MCS1, MCS2, MCS3 and MCS8, with PHY rates up to 60 Mb/s. Studying the median, we observe that the throughput obtained using the best among the set of MCS rates outperforms PHY auto rate adaptation (with 100% or more higher throughput at each distance). In the figure, MCS3 gives the best throughput for distances between 20 – 160 m; MCS1 for distances between 180 – 220 m; and MCS8 for distances between 240 – 260 m. We note that space time block coding (STBC) (used in MCS1, MCS2 and MCS3) outperforms spatial division multiplexing (SDM) mode (used in MCS8) for distances up to 220 m. Not shown in the figure, the packet loss rate is greatly reduced by simply fixing the rate.

We conclude that a strong component of our losses is caused by the disability of the auto-rate algorithm to adapt to the highly dynamic aerial channel between our airplanes. In order to factor out other sources of noise, we use our quadcopters as shown in the next section.

3.2 Experimental Results with Quadcopters

We configure the quadcopters to hover at different relative distances. The tests are performed with auto PHY rate. Dictated by safety requirements in our field (but not necessary in real missions), we consider lower altitude and closer distances than airplanes. Figure 7 (on the left) plots throughput vs. relative distance for two quadcopters. We observe higher throughput and smaller variability than in the airplanes tests shown in Figure 5 (see the upper and lower quartiles of the boxplot as well as the ends of the whiskers).

Impact of Relative Speed. We take advantage of the high level of control of quadcopters to analyze how the speed affects the

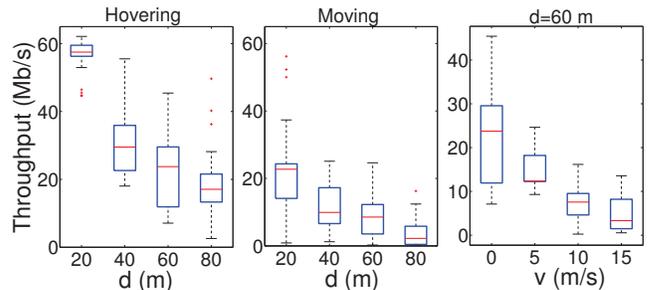


Figure 7: Quadcopters tests. On the left and the center, throughput versus distance while hovering and moving, respectively. On the right, comparison of the throughput achieved traveling at different cruise speeds.

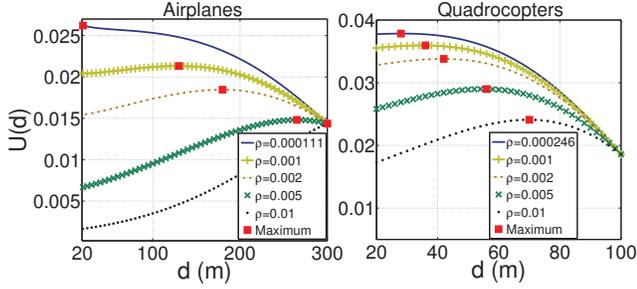


Figure 8: $U(d)$ for various ρ , baseline scenarios.

throughput. We perform several sets of tests where a quadrocopter moves towards a second (hovering) one with an average speed of ≈ 8 m/s and transmits while approaching the target UAV. The results are summarized in Figure 7 (plot in the center) and show a clear drop in the throughput. We conclude that hovering of quadrocopters induces more stable channel conditions, which makes strategies of transmission under little relative motion perform better than strategies based on ‘move and transmit’. The impact of speed is further stressed in Figure 7 (plot on the right) that summarizes the results of transmitting at different speeds at a distance of around 60 m. Clearly, the throughput varies and drops significantly with the speed. Consequently, mixed strategies containing ‘move and transmit’ would require a further dimension (the speed) to empirical-driven throughput estimation, leading to an interesting extension of our model. Strategies based on ‘hover and transmit’ can already be supported by our simpler and tractable characterization of the channel.

4. STRATEGY STUDY

We now provide an analytic formulation of the throughput (fitted to real testbed measurements) and investigate the strategy of delayed gratification under varying conditions by means of empirically-driven simulations based on Matlab. We consider the following two baseline scenarios and parameter settings following the problem formulation in Section 2.2:

- **Airplane scenario:** $M_{\text{data}} = 28 \text{ MB}^3$, $v = 10 \text{ m/s}$, $\rho = 1.11 \cdot 10^{-4} \text{ m}^{-1}$, $A_{\text{sector}} = 500 \text{ m} \times 500 \text{ m}$, $d_0 = 300 \text{ m}$. Since airplanes move continuously, we assume that airplanes can synchronize their trajectory in a way that their relative speed is close to zero at the time of transmission.
- **Quadrocopter scenario:** $M_{\text{data}} = 56.2 \text{ MB}^4$, $v = 4.5 \text{ m/s}$, $\rho = 2.46 \cdot 10^{-4} \text{ m}^{-1}$, $A_{\text{sector}} = 100 \text{ m} \times 100 \text{ m}$, $d_0 = 100 \text{ m}$.

Above, the failure rate ρ is chosen as the inverse of the distance that the UAV could travel at its nominal cruise speed v before the battery will be completely depleted. We solve the optimization problem of Eq. (2), given the empirical input values of $s(d)$. We fit a logarithmic function to the empirical median throughput (auto PHY rate) for different distances [8]: (i) $\hat{s}_{\text{airplane}}(d) = 10^6 \times (-5.56 \times \log_2(d) + 49)$ and (ii) $\hat{s}_{\text{quadrocopter}}(d) = 10^6 \times (-10.5 \times \log_2(d) + 73)$,

³With a resolution image of 1280×720 pixels ($k = 16/9$), flying at an altitude of 70 m, and camera lens angle of 65° , we have $\text{FOV} = 90 \text{ m}$ and $A_{\text{image}} = 3432 \text{ m}^2$ [7]. Accordingly, using $A_{\text{sector}} = 0.25 \text{ km}^2$, we have $M_{\text{image}} = 0.39 \text{ MB}$ in JPG100 format (100% quality, 24 bit/pixel) (<http://web.forret.com/tools/megapixel.asp>) and $M_{\text{data}} = 28 \text{ MB}$.

⁴Flying at an altitude of 10 m gives $\text{FOV} = 12.7 \text{ m}$ and $A_{\text{image}} = 69.4 \text{ m}^2$. With $A_{\text{sector}} = 0.01 \text{ km}^2$, $M_{\text{image}} = 0.39 \text{ MB}$ and $M_{\text{data}} = 56.2 \text{ MB}$.

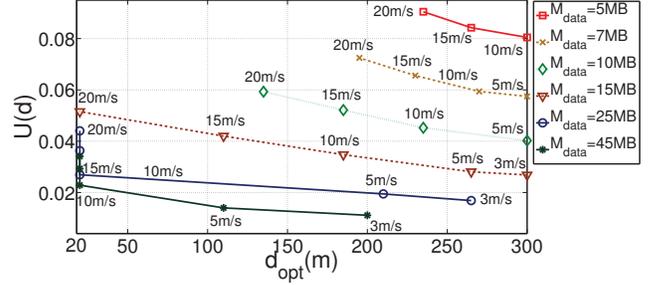


Figure 9: Delayed gratification for different data sizes M_{data} (curve for each setting of M_{data}) and speeds v (sample points with respective labels); airplane scenario.

with coefficient of determination $R^2 = 0.9$ for the airplane scenario and 0.96 for the quadrocopter one. We consider a minimum distance of 20 m between two UAVs to avoid physical collisions.

Figure 8 depicts $U(d)$ versus d as a function of ρ . We observe that the optimal distance d_{opt} of Eq. (1) increases with the failure rate ρ . We note that $U(d)$ can be approximated with a concave function for $\rho \ll 1$, and thus the formulation in Eq. (2) can be approximated as an unconstrained concave maximization problem. However, this result does not hold for higher ρ and may not hold for other $s(d)$ functions. Not shown in the figure, we further observe that d_{opt} does not change having smaller d_0 at which the data is ready. This occurs as long as d_0 does not reach d_{opt} . Once $d_0 = d_{\text{opt}}$, it becomes beneficial to transmit immediately. Different results are expected, e.g., for a non-stationary failure rate. Nevertheless, we expect that realistic failure models will give the same qualitative conclusion that there exists a tradeoff between the minimization of the delay and the amount of data that can be delivered before a failure occurs.

We finally study how changes in data size M_{data} and speed v influence $U(d)$ in the airplane scenario. We vary M_{data} from 1 MB to 45 MB and v from 1 m/s to 20 m/s. Figure 9 shows an excerpt of the resulting $U(d)$ versus d_{opt} . The plot confirms that by increasing the speed it is better to move closer and closer for a given M_{data} . Once the minimum distance is reached, higher speeds even increase the gratification to delay the transmission, as observed for M_{data} of 25 and 45 MB, and speeds above 10 – 15 m/s. Finally, looking at the instances of same speed, we can conclude that having larger M_{data} makes it more advantageous for a UAV to move closer to another UAV but at the cost of reduced $U(d)$ (caused by longer communication delay) once it reaches the optimal distance d_{opt} .

5. DISCUSSION OF OUR APPROACH

Our UAV networking approach depends on design decisions concerning the implementation and use of *controlled mobility*. The approach is mainly driven by the aim to provide a *practical* solution. We assume a centralized system (central planner), which controls the mission and is aware of the positions and trajectories of the UAVs and, thus, of their distances d . For reasons of tractability, we assume aerial links with line-of-sight behavior. This allows us to correlate the Euclidean distance between two communicating nodes and the radio signal quality and, hence, to study the gradient of the transmission time T_{tx} as a simple function of d . To account also for walls and other obstacles, our model requires an extension.

The central planner may send both the receiving and the transmitting UAV to new positions for improving transmission. Thus, UAV movement is determined by the mission and communication optimization. To simplify this task, we make an artificial division

between these two movement purposes. It is expected that a holistic planning approach integrating both movement types will lead to better performance – but this is a potentially complex problem to solve and not yet integrated in our model.

6. RELATED WORK

UAVs are employed in surveillance and SAR operations, where they scan areas but also act as networking relays to transmit data to a ground station. Most related works focus on the transfer of small-sized sensor data [8, 9]. The delivery of large-sized images with time-critical requirements is, on the one hand, hampered by the wireless channel dynamics due to mobility, and, on the other hand, supported by the ability to control mobility.

UAVs provide additional connectivity for the ground nodes, and allow to extend the range of mesh networks [10]. In our preliminary study on link quality assessment (solely on airplanes), we observed high losses in two test scenarios [11]. Channel measurements of 802.11a UAV-to-ground links were presented in [12] to characterize the impact of heights, yaws, and distances. [13] measured a throughput of up to 13 Mb/s from ground to one UAV, and half of the throughput using another UAV as relay. One of the major challenges for aerial links is the antenna orientation of highly mobile nodes. [14] studied the impact of the antenna orientation using a fixed wing UAV, while an antenna extension to 802.11a was proposed in [15]. These extensions improved the communication of UAV-to-ground links at 5 GHz. Our work complements these related studies by an experimental study of 802.11n channel dynamics leveraging two heterogeneous flying platforms, and results originating from PHY rate investigations.

In mobile robot and, in particular, UAV networks, the mobility of the nodes can be controlled [16–19], first, to accomplish a mission, and, second, to create networking opportunities. Controlled mobility [8, 17, 20, 21] differentiates here between *relaying*, i.e., placing a relay node, which remains static, and *ferrying* of information from one stationary node to another stationary node. Although termed differently, the concept of ferrying is well-known in store-carry-forward and delay-tolerant networks (DTNs). It has been shown that the capacity of a wireless multi-hop network can be increased by the mobility of nodes [22]. A consequence of these results is to aim for short channel occupancy times, which we address by delaying transmission on the single link.

Rather than introducing controlled mobility only for a subset of nodes (dedicated ferry nodes), our view is that the scarce number of UAVs flying in the area requires that any mission-oriented UAV can become a ferry. Different to previous works, we introduce a way to decide when to transmit, expressed with the concept of delayed gratification, i.e., a tradeoff between the desire of the ferry to transmit as soon as possible and minimizing the overall communication delay.

7. CONCLUSION

UAVs may contribute to SAR missions by gathering and sending large-sized image data for mapping a supervised area. In an extensive experimental study, we found that the otherwise promising 802.11n technology is limited when exposed to dynamic aerial channels. Consequently, UAVs should control their mobility such that the transmission can be scheduled at the time and location optimal for achieving high performance with little risk of UAV operation failure. By applying the concept of delayed gratification, we computed this best position as a tradeoff between the UAV's desire to transmit as soon as possible and to position itself best for transmission. Yet, to be tractable, the model contains a few sim-

plifications and assumptions which can be subject to extensions of this work, namely, introducing a specific failure model, exploiting new dimensions of the optimization problem, and studying the cost of re-positioning during the planned mission.

Acknowledgments

The authors are grateful to Prof. Douglas Leith for the fruitful discussions on delayed gratification, our shepherd Prof. Lachlan Andrew for detailed and insightful comments, and our reviewers for their constructive feedback. This research has been partially funded by the Swiss National Science Foundation (SNSF) Sinergia project SWARMIX, project CRSI22_133059 and by the Commission of the European Union under the FP7 Marie Curie IEF program contract PIEF-GA-2010-276336 MOVE-R.

8. REFERENCES

- [1] W. Mischel, Y. Shoda, and M. L. Rodriguez, "Delay of gratification," *Choice over time*, p. 147, 1992.
- [2] T. O'Donoghue and M. Rabin, "The economics of immediate gratification," *Journal of Behavioral Decision Making*, vol. 13, no. 2, pp. 233–250, 2000.
- [3] A. Blum, S. Chawla, D. R. Karger, T. Lane, A. Meyerson, and M. Minkoff, "Approximation algorithms for orienteering and discounted-reward TSP," *SIAM Journal on Computing*, vol. 37, no. 2, pp. 653–670, 2007.
- [4] S. Leven, J. Zufferey, and D. Floreano, "A simple and robust fixed-wing platform for outdoor flying robot experiments," in *International Symposium on Flying Insects and Robots*, pp. 69–70, 2007.
- [5] "Arducopter: Arduino-based autopilot." <http://code.google.com/p/arducopter/>.
- [6] K. Pelechris, T. Salonidis, H. Lundgren, and N. Vaidya, "Experimental characterization of 802.11n link quality at high rates," in *ACM WiNTECH '10*, pp. 39–46, 2010.
- [7] L. Duplay, "Target detection using flying robots," tech. rep., Laboratory of Intelligent Systems (LIS), EPFL, Jun. 2012.
- [8] D. Henkel and T. Brown, "Delay-tolerant communication using mobile robotic helper nodes," in *IEEE WiOPT 2008*, pp. 657–666, 2008.
- [9] A. Jenkins, D. Henkel, and T. Brown, "Sensor data collection through gateways in a highly mobile mesh network," in *IEEE WCNC 2007*, pp. 2784–2789, 2007.
- [10] T. X. Brown, B. Argrow, C. Dixon, S. Doshi, R.-G. Thekkekkunnel, and D. Henkel, "Ad hoc UAV ground network (AUGNet)," in *AIAA 3rd Unmanned Unlimited Technical Conference*, 2004.
- [11] M. Asadpour, D. Giustiniano, and K. A. Hummel, "From ground to aerial communication: dissecting WLAN 802.11n for the drones," in *ACM WiNTECH '13*, pp. 25–32, 2013.
- [12] E. Yanmaz, R. Kuschnig, and C. Bettstetter, "Channel measurements over 802.11a-based UAV-to-ground links," in *IEEE GLOBECOM Workshops 2011*, pp. 1280–1284, 2011.
- [13] A. Jimenez-Pacheco, D. Bouhired, Y. Gasser, J.-C. Zufferey, D. Floreano, and B. Rimoldi, "Implementation of a wireless mesh network of ultra light MAVs with dynamic routing," in *IEEE Globecom Workshops 2012*, pp. 1591–1596, 2012.
- [14] C.-M. Cheng, P.-H. Hsiao, H. Kung, and D. Vlah, "Performance measurement of 802.11a wireless links from UAV to ground nodes with various antenna orientations," in *IEEE ICCCN 2006*, pp. 303–308, 2006.
- [15] E. Yanmaz, R. Kuschnig, and C. Bettstetter, "Achieving air-ground communication in 802.11 networks with three-dimensional aerial mobility," in *IEEE INFOCOM 2013*, pp. 120–124, 2013.
- [16] E. Chittimalla, A. Venkateswaran, V. Sarangan, and R. Acharya, "On the use of nodes with controllable mobility for conserving power in MANETs," in *IEEE ICDCS 2006*, pp. 88–88, 2006.
- [17] D. Goldenberg, J. Lin, A. Morse, B. Rosen, and Y. Yang, "Towards mobility as a network control primitive," in *ACM MobiHoc '04*, pp. 163–174, 2004.
- [18] A. So and B. Liang, "Enhancing WLAN capacity by strategic placement of tetherless relay points," *IEEE Transactions on Mobile Computing*, vol. 6, no. 5, pp. 522–535, 2007.
- [19] A. Srinivas and E. Modiano, "Joint node placement and assignment for throughput optimization in mobile backbone networks," in *IEEE INFOCOM 2008*, pp. 1130–1138, 2008.
- [20] W. Zhao, M. Ammar, and E. Zegura, "Controlling the mobility of multiple data transport ferries in a delay-tolerant network," in *IEEE INFOCOM 2005*, vol. 2, pp. 1407–1418, 2005.
- [21] C. Dixon and E. Frew, "Maintaining optimal communication chains in robotic sensor networks using mobility control," *Mobile Networks and Applications*, vol. 14, no. 3, pp. 281–291, 2009.
- [22] M. Grossglauser and D. Tse, "Mobility increases the capacity of ad-hoc wireless networks," in *IEEE INFOCOM 2001*, vol. 3, pp. 1360–1369, 2001.