

Methodological Evaluation of Architectural Alternatives for an Aeronautical Delay Tolerant Network

Rubén Martínez-Vidal^{a,*}, Ramon Martí^a, Cormac J. Sreenan^b, Joan Borrell^a

^a*Department of Information and Communications Engineering
Universitat Autònoma de Barcelona
08193 Bellaterra, Spain*

^b*Department of Computer Science
University College Cork
Cork, Ireland*

Abstract

In this paper, we use graph analysis to evaluate the network architecture of a large scale delay tolerant network (DTN) of transoceanic aircraft. At LCN (Local Computer Networks) 2014 we analyzed information propagation inside a pure opportunistic version of this network, a scenario constructed from more than 2,500 traces of transatlantic flights in which communications relied only on the sporadic contacts between airplanes. As only a small percentage of the nodes were capable of performing efficient air-to-ground communications we concluded the need to devise a more suitable network architecture by combining opportunistic and satellite communication systems. We propose a generic methodology based on graph analysis (both static and dynamic temporal) to evaluate the different ways to create this new architecture. We show the architectural combination that most improves the network delivery performance while minimizing its deployment costs.

Keywords: Delay Tolerant Networks, Transoceanic aircraft, Simulation, Graph Analysis

1. Introduction

Data communication in aeronautical scenarios is very limited [18]. Airplanes periodically send to ground stations their identification and position, and some company data (few tens of bytes per second). Any other control information is sent over traditional voice-based radio links. These limitations in data communication become more evident when looking for the causes of aerial accidents, as in the case of the ill-fated MH370's disappearance [24].

Regarding passenger data communication, some companies have deployed, for domestic flights, solutions based on external antennas directed to ground [2]. As to transoceanic flights, very few companies choose to use expensive satellite links [23, 7] while, in most of the cases, passengers must remain disconnected for hours. On the other hand, most of the research papers on aeronautic transoceanic communications propose the use of communication systems with large coverage (from 300 to 600 km) and thus consider airplanes connected to form an Aeronautical Ad-hoc Network (AANET) [19, 22, 18].

Looking for a cheaper solution than those based on radio links with high coverage or satellite links, we proposed in LCN (Local Computer Networks) 2013 [20] a new approach based on aircraft equipped with IEEE802.11b/g communication links of only 50 km. This limited range led to a network with a reduced node degree, fully in line with the characteristics of Delay Tolerant Networks (DTNs) [8, 3]. To build our network, instead of using synthetic data, we used the actual traces of 2,550 transatlantic flights to produce a characterization of a large-scale real-world aeronautical DTN. We analyzed the distributions of active nodes, node degree, number of clusters, any contact time (ACT), and inter-any-contact time (IACT) of that network.

*Corresponding author

Email addresses: rmartinez@deic.uab.cat (Rubén Martínez-Vidal), rmarti@deic.uab.cat (Ramon Martí), cjs@cs.ucc.ie (Cormac J. Sreenan), jborrell@deic.uab.cat (Joan Borrell)

In LCN 2014 [21] we analyzed how data propagate inside our previous transoceanic aircraft delay tolerant network, in order to evaluate the feasibility of using this network as a reliable way to deliver on-flight generated data (either from passengers or company) to the ground. First, we provided an evaluation of two different routing protocols for this scenario: A basic routing protocol using an epidemic dissemination approach, and a version of the protocol which uses a decision function based on an Expected Landing Time (ELT) metric to reduce the data overhead. Then, we showed how only a small percentage of the nodes of the network (7%) was capable of performing efficient air-to-ground communications, by delivering on-flight generated data to the ground at least 5,000 seconds before the landing of the source aircraft, i.e at least before one third of the average time in which airplanes are disconnected.

Our final conclusion in [21] was that the architecture of our network had to be improved in order to increase the percentage of nodes capable of performing successful air-to-ground communications. To this end, we planned as a future work to devise a more suitable network architecture by combining nodes with just opportunistic communications with nodes equipped with default aircraft satellite communications.

The main objective of this paper, as an extension of [21], is to evaluate the different ways to create this new architecture in order to find the combination that improves the network performance while minimizing its deployment costs. We conduct this evaluation through the analysis of the graph representation of the network. Graph analysis is an established methodology applied to social networks that has been recently used as a powerful and general tool to forward data in DTNs [5, 14, 12]. We consider several centrality graph metrics [11], such as degree or betweenness, to determine the smallest possible set of nodes to be equipped with satellite links. Not only static graph analysis is used but also dynamic temporal analysis [33] to achieve a further reduction of the set of nodes and thus of the deployment costs.

Our contributions are as follows:

- We propose a novel ad-hoc network architecture for the aeronautical network scenario we defined in earlier papers, by combining opportunistic and satellite communication systems.
- We propose a generic methodology based on graph analysis techniques and centrality metrics to choose which nodes of the scenario have to be equipped with satellite communication systems. This graph analysis allows us to conduct a complete evaluation of the architecture's performance for use in air-to-ground communications using a custom delay metric (*delivery time reduction*).
- We empirically determine the most efficient centrality metric (*betweenness*) and the optimal parameters that maximize performance while minimizing deployment costs. Finally we show how a dynamic temporal analysis of the associated graph allows further reduction of those costs.

The rest of the paper is organized as follows. Section 2 presents the background and previous work on Aeronautical Ad-hoc Networks, on performance evaluation of DTN over realistic scenarios, and on the use of graph analysis on such networks. In Sections 3 and 4 we provide a comprehensive description of our previous papers [20] and [21]. In Section 5 we present the new architecture, the two types of nodes which compose it, and the characteristics of the new scenario. In Section 6 we construct the associated contact graph to our network and we analyze it to infer which nodes are the most suitable to be equipped with satellite links, according to several graph metrics. In Section 7 we further refine our analysis taking into account temporal restrictions of the dynamic graph. In Section 8 we discuss our results regarding the costs of deployment. Finally, in Section 9 we present the overall conclusions of the paper and our intended future work.

2. Related Work

In this section we review the previous work in design of routing protocols in AANETs, performance evaluation of Delay Tolerant Networks under realistic scenarios, and graph analysis of such networks.

AANETs have the distinguishing feature that they are composed by nodes moving at high speeds that results in highly dynamic network topologies. This has a great impact on the performance of routing protocols, and as a result, providing reliable communications in such environments is a great challenge. Most papers agree with the necessity of a multi-hop based approach for the successful delivery of data within the network. It is also commonly remarked the inability of common Mobile Ad-hoc Networks (MANETs) routing protocols to efficiently deliver data packets when

used in AANET environments. There have been a few studies on the field of routing in aeronautical networks, each one presenting their own efficient way to select the optimal route.

One of such studies is the geolocation based routing protocol AeroRP [15]. This routing protocol takes decisions based on a TTI (Time To Intercept) metric which gives an estimation of how soon a neighbor node will be within transmission range of the destination. Some studies regarding this routing protocol offer a performance evaluation ([29, 28, 4]) under different traffic conditions and compare it against common MANET routing protocols.

Another example of such a protocol is DASR (Delay Aware Routing Protocol) [13], a reactive routing protocol which takes decisions based on a expected node delay metric. The authors use performance metrics related to traffic demand and network stability to compare their approach against other routing protocols. Lastly, it is worth to mention MUDOR (Multipath Doppler Routing Protocol) [31], a reactive routing protocol which uses a decision function based on the relative speed of nodes. This protocol aims to select stable paths within the network, and to guarantee links duration long enough to successfully forward data.

An important aspect of previous papers is that they treat this kind of networks as highly mobile MANETs, with common link failures produced by reduced contact times and highly changing topologies. Some of them, like [30], are aware of this problem, and offer their own custom protocols to guarantee reliable services over this kind of networks.

In our case, we prefer to directly consider this network as a DTN and apply a methodology specifically designed for this kind of networks. For the validation of our results we will use the realistic scenario proposed on our previous paper [20], fully described in section 3.

DTNs, as opposed to MANETs, are disconnected most of the time and thus routing performance is directly dependent on the frequency and duration of contact opportunities. Consequently, when studying the performance of routing protocols in a DTN, it is important to use a model that accurately represents these opportunities. We can find several examples of routing performance analysis for DTNs based on realistic characterizations of scenario mobility, such as [16], [37] which base their model on actual mobility traces from buses inside the cities of San Francisco and Amherst (Massachusetts) respectively. Other examples that actually study large-scale networks using real world data can be found in [1], where the authors base their study on roughly 1,200 buses from the public transportation network of Seattle, and [36] which is based on mobility traces of approximately 12,000 taxis in Beijing.

Also regarding DTNs, graph analysis has been used by many studies as an established methodology for prediction of future contacts. Common methods involve aggregating past observed contacts to define a social graph, study the network structure and forward data to nodes in important network locations. Some examples of this are Simbet [5] and BubleRap [14], whose authors use the concepts of community and centrality. It is assumed that the mobility of nodes forms communities of related nodes, and those nodes located in central locations within the community are chosen as relays of data. Another interesting approach that makes use of graph analysis in DTN networks is [12], that is based on a public bus transport network, as opposed to previous works which create their graph based on contacts between mobile nodes. In this paper the nodes of the graph are the bus stops, and the edges appear when a bus moves between two stops. The authors use temporal centrality metrics to locate which stops act as key mediators of information over time, and they define a hierarchical network architecture where those nodes occupying a higher position act as a network backbone.

3. Network Model

In this section we recall our scenario from [20, 21] and proceed to describe the methodology used to define the network model and its characteristics. This model has been implemented using the ns-3 simulator [25] and currently represents the different features of the network in great detail, from link layer to application layer, including a custom mobility model and the routing protocols.

This model defines an aeronautical ad-hoc network in which nodes have limited communication range and rely on sporadic encounters to communicate in an opportunistic manner. The area of deployment is based on the North Atlantic ocean (which is the ocean with the highest concentration of flights). We define a geographically bounded oceanic area where the planes are disconnected (see Figure 1), this zone will be referred as **disconnected zone**. Our study focuses on the time span where the aircraft traverses this disconnected area. This time interval will be referred as the aircraft's **disconnected flight duration**. And the nodes operating within this zone will be called **active nodes**. When the nodes leave the disconnected area, we assume that they are capable of delivering their data using the infrastructural Internet, and thus their opportunistic capabilities are inactive.




Figure 1: Bounds of the disconnected zone over the north Atlantic oceanic area.

Our model includes all the flights departing on the same concrete day (Thursday, 03/14/2013) for all time-zones involved. All departure times are converted to UTC time format in order to provide a standardized time for the whole network. Simulation starting time is synchronized with the instant 0:00 AM - UTC time zone.

3.1. Mobility Model

In this section, we describe the design and characteristics of the mobility model. A high-level pseudocode representation of the node mobility set-up procedure is provided in Algorithm 1. The basic assumptions that constitute this mobility model are the following:

Trajectory: nodes follow geodesic patterns constructed using the datasets from [26], which is a user maintained public domain database based on the latest DAFIF¹ public release. This database contains a total of 59,036 routes between 3,209 airports from 531 airlines. We use the subset of routes that cross the North Atlantic ocean (up to 239 routes).

Scheduling: This defines the number and frequency to which planes traverse the previously defined routes. Data on aircraft schedules was obtained from several flight status providers like [10] and [9] using web scraping techniques. We include all the flights with scheduled departure in a single day and use a unique simulation node for every departing flight (up to 2,878 nodes).

Node speed: Mobile nodes move at a constant speed of Mach 0.84 (285.852 m/s), which is (according to specifications) the standard cruise velocity for most commercial aircraft models.

Altitude: Node altitude only changes during take-off or landing. All nodes flight at the same altitude, and we do not consider multiple flight levels. The flight level has little relevance due to the use of omnidirectional antennas with a range of several kilometers (details given in the next section).

Separation regulations: Our model, uses a simplistic separation ATC (Air Traffic Control) model which only takes into account same course longitudinal separation guidelines. Upon scheduling a new flight the simulator checks it against previously recorded flights, in case that two aircrafts come within a 10 min flying time of each other, the new flight departure is postponed. Also, note that the unique set of guidelines OTS (Organized track system) used in the North Atlantic ocean is not being applied.

3.2. Network Configuration

In this section, we describe the lower layers of the communication stack and the traffic generation model.

Physical/MAC/IP: As a communication hardware we propose a cost-effective system using IEEE 802.11b/g with high gain and high sensitivity (low noise) transceivers, directly connected to omnidirectional antennas. As an example, Bullet 2HP transceivers [34], that for 1/54 Mbps transfer rates provide 20/15 dBm transmission power with a

¹Digital Aeronautical Flight Information File, that is a complete database of up-to-date aeronautical data managed by the National Geospatial-Intelligence Agency.

Algorithm 1 Mobility setup procedure: Each node is assigned a mobility model defined as a list of waypoints. A waypoint is a structure with two components: a position vector (x,y,z) and a time variable.

```
nodes ← []
routes ← LoadRoutes()
for each route in routes do
  flights ← LoadFlights(route)
  for each flight in flights do
    node ← []
    # Creates the list of waypoints required to traverse a route and initializes their position component.
    waypoints ← ComputeWaypoints(route)
    # Enforces longitudinal separation, postponing the new flight departure if necessary.
    flight ← SeparationGuidelines(flight)
    # Assigns the waypoints' time component which is computed assuming constant speed.
    waypoints ← ComputeTimes(flight, waypoints)
    node.waypoints ← waypoints
    nodes ← nodes + node
  end for
end for
```

-95/-72 dBm sensitivity, could be used, together with 3 dBi omnidirectional antennas in order to have a low directivity in the elevation (vertical) plane. According to hardware specifications, they can achieve a maximum effective communication range of 50 km.

We have modeled this communication channel using a range propagation loss model, which depends only on the distance between transmitter and receiver. The bandwidth of the link has been set to 1Mbps using Direct-Sequence Spread Spectrum (DSSS) modulation. We use IEEE 802.11g as Datalink layer and MAC type set to ad-hoc mode. Lastly, we use IPv4 with private range addresses.

Traffic generation: Each aircraft is equipped with a CBR (constant bit rate) traffic generator that follows an on-off pattern alternating at periodic intervals between two states: traffic generation and idling. The characteristics of the CBR are defined by the size of each packet, which is set to 512 bytes, and the bit rate, which is set to 256 bytes/second. The traffic generators start as soon as the airplane enters the disconnected zone and stop when the data generated attains a maximum of 10 KBytes or once the airplane leaves the disconnected area, whichever comes first.

Our traffic generation model does not represent any specific network application as it is currently intended as a tool to model information flow. This information flow is formally shown in Section 6 using a graph structure. Our selected value seeks a compromise between the information flow modeling (for which a single packet would be enough) and the evaluation of protocol performance.

3.3. Network Characterization

There are several characteristics that need to be taken into account for defining proper routing strategies in this scenario. In this section we highlight the most relevant:

Active node distribution: This measure describes the variation in number of active nodes. Nodes become active when they enter the disconnected the area and are disabled upon leaving it.

The distribution over time can be seen in Figure 2a. The nodes show two main peaks of activation which is consistent with the usual trend of North Atlantic flights: airplanes start flying west (from Europe to America) during daytime (~ 06:00 UTC), and east (from America to Europe) during the night (~ 22:00 UTC). This periodic pattern is repeated daily, and thus only two peaks of activation are shown in Figure 2a. Another interesting point is that the intervals of activation overlap (some stragglers from Europe are still active at the time flights start departing from America) resulting in active nodes during the whole period.

This distribution offers insight on the characteristics of mobility and contact opportunities that should be taken into account during the design and deployment of network applications.

Average node degree: this is a measure of the number of neighbors a node has at any time. Since it is applied to a mobile network the changes of topology during the whole studied time interval are considered. This metric offers us a quick way to quantitatively measure the degree of network connectivity.

The node degree is influenced by the number of active nodes and the contacts between them. In Figure 2a we can see the evolution of average node degree over time in relation to the number of active contacts. There is a direct relationship between the number of active nodes and the node degree. On the other hand, the average node degree is defined by the number of contacts and thus influenced by the radio range used. Since a higher radio range implies more contact opportunities and thus higher degree, there is a linear relationship between the node degree and the radio range as shown in Figure 2b.

Any contact time (ACT): this metric defines the expected duration of a contact between any two nodes of the network. It can be used to estimate the amount of data that can be exchanged in every encounter. We compute it as the elapsed time between the first and the last consecutive messages received during a contact situation. The contact time is directly influenced by the radio range, since a higher radio range ensures a longer contact. The linear dependence between them is shown in Figure 2b.

Inter-any-contact time (IACT): This metric measures the elapsed time between two consecutive meetings of any pair of nodes. As such, it offers insight on the frequency of communication opportunities, and it is useful for the design and configuration of routing protocol parameters. As can be seen in Figure 2b, there is an inversely proportional dependency between radio range and inter-contact-time, in the same way that the duration of contacts increases the time between contacts decreases.

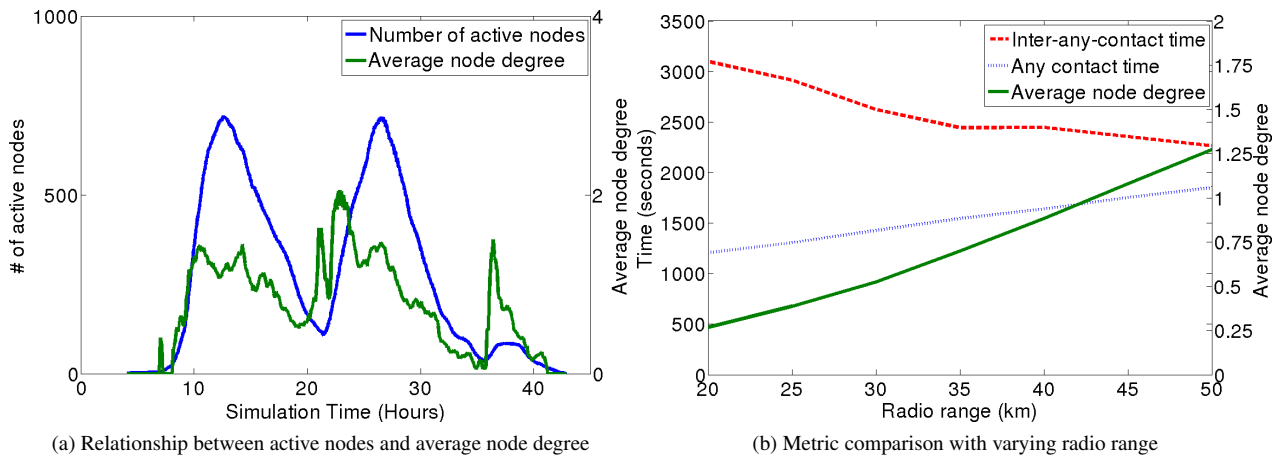


Figure 2: Network topology metrics

A much more detailed analysis and characterization of the network topology may be found in our previous work [20]. In the following sections we focus our study on a network configuration based on 50 km radio links, which displays higher connectivity values but remains within the parameters commonly seen in DTN network scenarios.

3.4. Departure Delay Impact

Daily air traffic commonly suffers deviations from scheduled departure times. In this section we consider those changes and analyze the impact they have in our network model. To perform this analysis, we compared scheduled departure times (gathered one day before departure) against the actual departure times. This data was gathered from all departing flights during a time lapse of 5 days. The histogram of the gathered samples is shown in Figure 3a, which displays a bell-shaped distribution centered around a positive delay of 980 seconds. As there is no good fit using commonly known probability distributions we use kernel density estimation [27] to describe the probability density function of the gathered samples, which is provided by the expression:

$$\hat{f}_h(x) = \frac{1}{nh} \sum_{i=1}^n K\left(\frac{x - x_i}{h}\right) \text{ for } h = 174.491 \text{ and a Gaussian kernel } K(u) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}u^2}$$

We compute the cumulative distribution function under the range of observed delays and use the inverse-transform sampling method [6] to generate random numbers according to this distribution. These values are added to the simulation as a random parameter that describes the deviation of each plane from its initially scheduled flight plan. We evaluate the impact of including those random delays to the simulation using the metrics described in Section 3.3. In Figure 3b we show the comparison between a static network model and a model including random delays. As we can see, the values of each metric change but the overall growing trend and line slope is preserved. A paired-sample t-test determines that the two data-sets are not significantly different at the 5% significance level with a p-value of 0.5201, and thus we keep our model simple, with just the data from the day mentioned before.

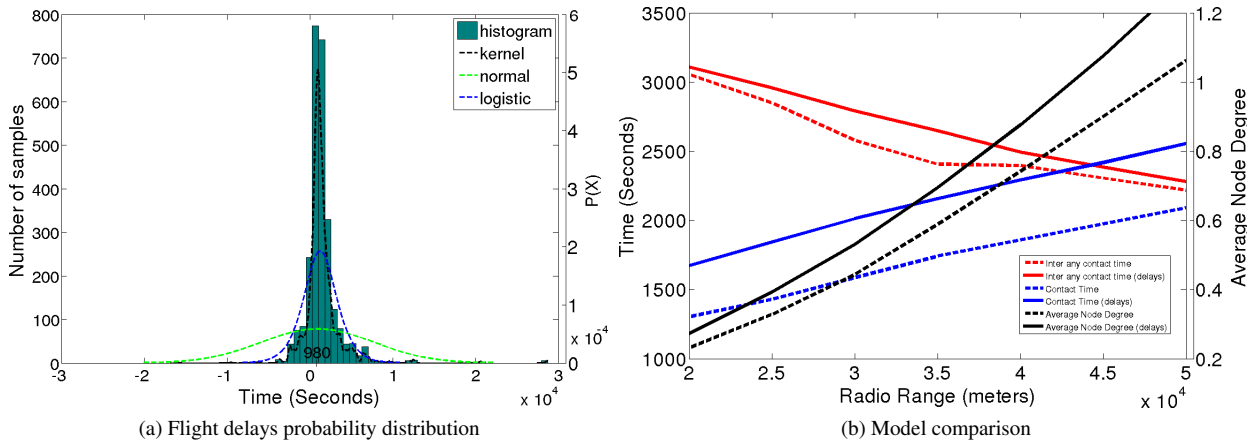


Figure 3: Departure Delay Impact

3.5. Communication Models

In our scenario, we can define three types of communication-based on its purpose, the origin and the destination of the data:

- **Air-To-Ground:** This kind of communication seeks to deliver data generated in-flight to the ground, such as passenger data or aircraft status notifications.
- **Air-To-Air:** This kind of communication seeks to deliver data generated in-flight from one node to another. One example of this may be the exchange of flight route information between planes (such as the meteorological conditions of a previously traversed area).
- **Ground-To-Air** This kind of communication tries to deliver data generated on-ground to a node that is currently on-flight. An example of such application may be key distribution or any two-way communication (e.g. web browsing).

We started focusing on achieving reliable cost effective communications for the first case: air-to-ground communications. In this case scenario, the routing priority is to reduce the time required for data delivery. This purpose is achieved exploiting the periodic encounters with other nodes, and as such we start addressing this problem using an opportunistic network approach that is detailed in the next section.

4. Opportunistic Network Architecture

In [21], we analyzed the opportunistic network approach for information propagation in the large scale delay tolerant network (DTN) of transoceanic aircraft presented in the previous section. In this kind of approach, nodes use ad-hoc wireless communications and exploit sporadic encounters to forward their data using a multi-hop scheme.

In this section, we summarize the experiments performed, the metrics used, and the results obtained. We consider three main well known performance metrics to evaluate the performance of the protocols, and one scenario specific metric (delivery time reduction) to assess the degree of accomplishment achieved towards reducing communication time.

4.1. Routing Protocols

From the previous network characterization we can infer that this is a sparsely connected network, and as such it is our priority to take advantage of every contact opportunity. Therefore, we have chosen an epidemic based flooding scheme [35]. This routing scheme allows to maximize the delivery ratio and improves delivery time by exploring all possible paths available.

In these schemes, nodes maintain a fixed length buffer that stores entries until the data can be transmitted or upon expiration after a fixed period of time. Possible transmission opportunities are found by periodically broadcasting beacons, when a contact is produced; both nodes exchange a copy of every packet stored in their respective buffers. Unfortunately, this sort of protocol offers no control over data forwarding, produces data overhead and generates a great volume of network traffic which eventually leads to a reduced delivery ratio due to network congestion.

We also implement a new version of the protocol in order to reduce this congestion. This new protocol uses a decision function based on a custom metric called Expected Landing Time (ELT). This protocol uses a limited flooding approach, and its inner processes are analogous to those of epidemic routing protocols. Specifically, this protocol has a custom discovery phase where the nodes exchange their expected time of arrival, perform a comparison based on the ELT metric, and proceed to make a copy of the packet. This last step is only performed when the encountered node has a lower ELT than the one currently holding the data. In this way, we reduce the number of transmissions and overhead through the network.

4.1.1. Routing configuration

The performance of these routing protocols is directly influenced by four parameters: **length of buffer** for packet storage, **expiration time** for entries in the buffer, maximum **number of hops** per packet, and **beacon interval**.

From our previous analysis [20] we found that the average IACT is around $\sim 2,200$ seconds (Figure 2b) and, additionally, that roughly 800 seconds is the smallest inter-any-contact time (IACT) displayed by nodes. This is the time lapse between successive contacts between any two nodes of the scenario. Based on this result, we set the **beacon interval** to half of that value (400 seconds), that ensures that the chance of missing contact opportunities is kept to a minimum while reducing network load.

In disconnected scenarios such as this one, the delivery delays can span to several hours. We need to guarantee that the buffer can store packets long enough to be delivered, at least by some of the nodes carrying a copy of them. As such, our next step is to study how long does a node remain in the disconnected area, in order to define the maximum time required to deliver a packet. The most extreme case is when the source aircraft leaves the disconnected area and delivers the packet directly. In Figure 4a, we show the complementary cumulative distribution function (CCDF) of the disconnected flight duration for all the nodes of the scenario. This figure shows the observed disconnected flight duration for each aircraft as a point of the plot, the ordinate represents the percentage of samples above a particular time given by the abscissa. We use the average of all samples, approximately 14,400 seconds, as the packet **expiration time**. This ensures an appropriate packet expiration drop ratio inferred from the specific characteristics of the scenario.

As shown in Figure 2b the average node degree is slightly above one unit, this implies that nodes have barely any neighbors most of the time. From this result, we can infer that packet replication will be low. Now, we study the most extreme case for which we use buffers with limitless storage, this us a higher bound that can be observed in Figure 4b). Since our objective is to maximize the spread of information, by taking advantage of every possible path. We set the maximum **number of hops** to 40, in accordance with the highest achievable replication value.

Finally, accounting for the low number of contact opportunities in the network and the amount of data generated by each node, the **length of buffer** is set to 100 packets. Each aircraft is capable of holding a traffic from up to 5 other nodes which is enough for most of the cases. This value establishes an environment constrained enough to display a proper amount of packet drops for the purpose of our experiment.

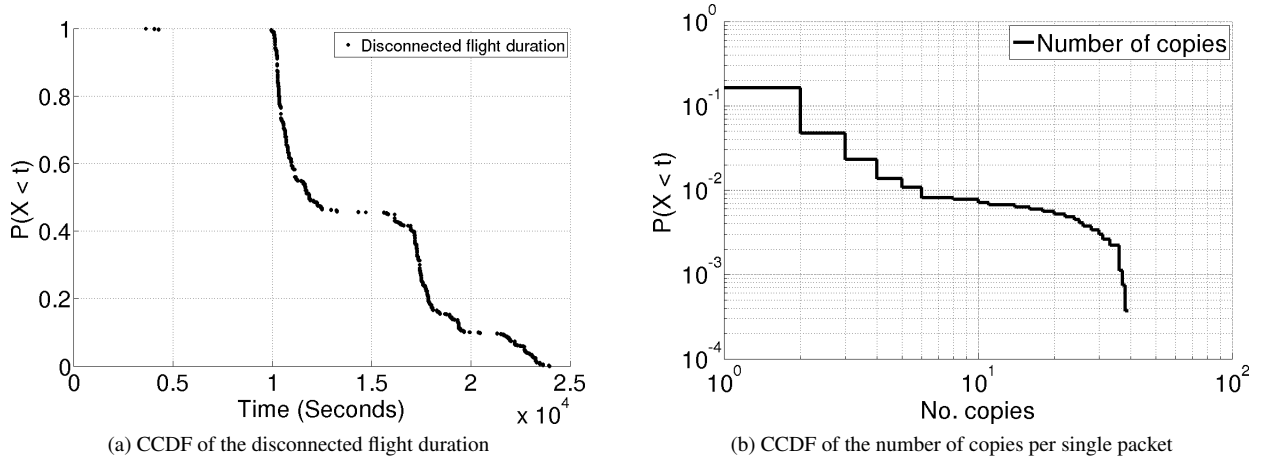


Figure 4: Complementary cumulative distribution functions of network node characteristics

4.2. Configuration Parameters

For convenience we provide a summary of all the relevant parameters defined through the section along with their default values, this can be seen in Table 1.

Beacon Interval	400 s.	Max. N. of hops	40 hops
Traffic per node	10 KB	Radio Range	50 km
Buffer Capacity	100 packets	Expiration Time	14,400 s.

Table 1: Default configuration parameters

4.3. Protocol Performance

In this section, we describe the evaluation of the performance of the routing protocols. For this purpose, we considered the following three metrics:

- **Delivery ratio:** This is the ratio of successfully delivered packets to the total number of packets generated in the network at the end of the simulation. The delivery ratio may be lower than one because of packet drops caused by either insufficient buffer capacity, packet expiration or by exceeding the allowed number of hops.
- **End-to-end delay:** This is defined as the time required for a packet to reach its destination. It is calculated by accumulating the delay of each hop. We perform this calculation at Network layer and only for packets successfully delivered.
- **Network traffic per packet:** This refers to the average network traffic per packet, and it is calculated by dividing the number of forwarded packets by the number of generated packets.

In the following sections, we describe the evaluation of the the performance of the routing protocols under changing conditions we performed. Concretely, we modified specific configuration parameters (such as: radio range, traffic demand, and buffer size) and observed the changes of behavior.

4.3.1. Varying radio range

In this experiment, we studied how the variation of the node radio range impacts the performance of the epidemic and ELT protocols. In Figure 5 we can see that there is a direct dependence between radio range and performance. Increased radio ranges imply higher values of **network traffic**, this is explained by the fact that longer radio ranges result in higher frequency of encounters and longer contact durations. Both **end-to-end delay** and **delivery ratio** remain mostly stable despite the variation on radio range.

Our conclusion is that the considered variations in radio range are not enough to generate encounters with new routes. As a result, there are no new paths from which the data can be delivered faster to the destination, and, therefore, there are no significant communication improvements. This fact is consistent with results observed in [20]. This paper analyzed two aeronautical models and offered analytical expressions showing that the increases in connectivity were not relevant until radio ranges of several hundred kilometers were used, in the domain of connected AANETs.

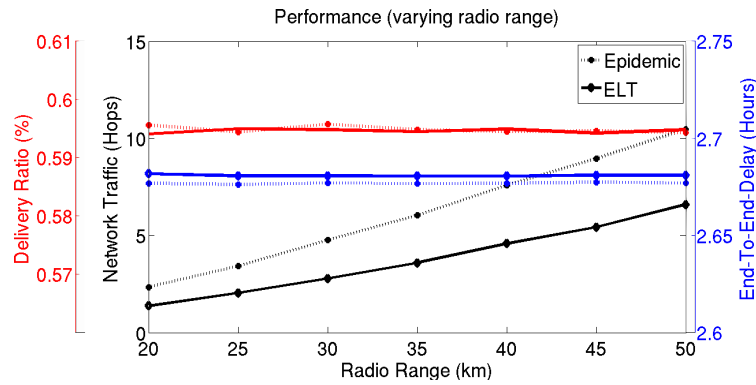


Figure 5: Varying radio ranges: impact on Delivery Ratio (red lines), End-to-End Delay (blue lines), Network Traffic (black lines)

4.3.2. Varying traffic demand

In this experiment, we studied how the increase of traffic demand influences the network, by changing the maximum amount of data generated by the traffic generators. As shown in Figure 6, a limited buffer size and a high expiration time implies that increases in traffic demand result in a high number of filled buffers. As a result, packets are not being forwarded, and thus packet drops increase due to congestion. This situation has a negative impact in both **delivery ratio** and **network traffic** per packet. Regarding **end-to-end-delay** little deviation can be observed. Since this metric is mostly influenced by node mobility (in our scenario), packet drops have limited impact.

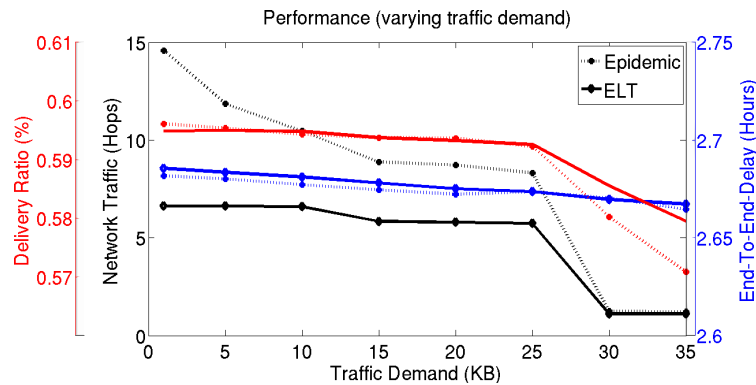


Figure 6: Varying traffic demand: impact on Delivery Ratio (red lines), End-to-End Delay (blue lines), Network Traffic (black lines)

4.3.3. Varying buffer size

In this experiment, we studied the impact of the buffer size, by changing the maximum amount of packets the data buffer in each node can hold. As shown in Figure 7, the **network traffic** per packet increases along with the buffer size until reaching a maximum value (at size 40), and then it becomes stable. For small buffer sizes, the buffer fills faster causing the forwarded packets to be dropped, and this affects the **delivery-ratio** that increases along with the size of the buffer. Finally, it can also be seen that packet drops have barely any impact on the **end-to-end-delay**.

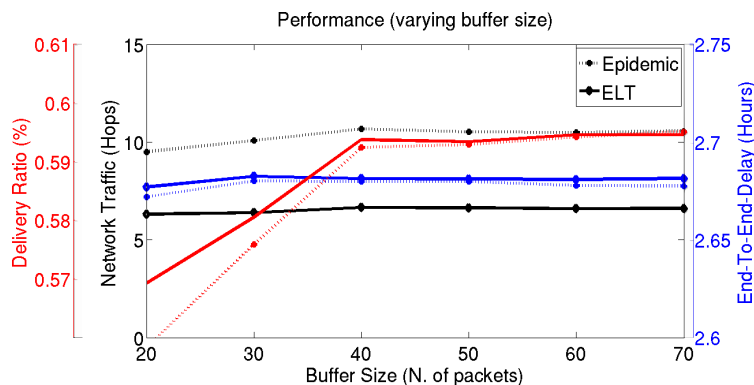


Figure 7: Varying buffer size: impact on Delivery Ratio (red lines), End-to-End Delay (blue lines), Network Traffic (black lines)

4.3.4. Protocol comparison

Table 2 shows a comparison of the two protocols using default parameters. Under these constraints, we get a relatively low delivery ratio of $\sim 60\%$, mostly due to expiration drops. The network traffic displayed by the epidemic routing is ~ 10.5 on average (taking into account the number of nodes of the network, 2,887, this represents a very small percentage). Also notice that the configuration parameters have little influence on the end-to-end delay, as it usually spans several hours and is mostly influenced by mobility patterns.

In regards to which protocol is preferable, we can say that the epidemic routing protocol offers slightly better performance under low traffic or bigger buffer sizes. Otherwise, the ELT based is preferable because it achieves similar results with a great reduction in network traffic (see Table 2). It is also interesting to note that the limited spread of packets makes the ELT a good decision function for this scenario. In most cases, the packets replicate very few times and remain in nodes that traverse similar routes, which makes choosing the one with smallest time of arrival better than a pure epidemic approach, see [21] for a full comparison.

Protocol	Epidemic	ELT
Min. End-to-End Delay (h)	2.6770778	2.6810722
Max. Delivery Ratio (%)	0.594838	0.594317
Network Traffic (hop)	10.4784	6.59046

Table 2: Performance comparison of the schemes

4.4. Network Usability

In order to evaluate the possible use of this network for air-to-ground communications, we define a custom metric based on delivery times.

Delivery time reduction: Is the gain in seconds obtained by forwarding the packet through the network instead of performing a direct delivery upon leaving the disconnected area.

This metric is based on the following concepts. Assume that an aircraft gets disconnected at time t_d , and regains connectivity at time t_c . During this interval a packet is generated, say at time t_g , then the time required for its delivery would be $t_c - t_g$. Now the aircraft, instead of waiting for t_c , decides to take advantage of the possible encounters with

other aircraft, and forwards the packet. When the packet finally gets delivered by a different aircraft, at time t_f , then the time for delivery will be $t_f - t_g$. In the case that t_f is smaller than t_c there is a reduction in delivery time ($t_c - t_f$) and we call this improvement *delivery time reduction*.

Figure 8 shows the tail distribution of the delivery time reduction in linear scale. As it can be seen, the network displays moderate improvement against direct delivery. First, roughly half of the samples offer an improvement of less than 185 seconds, and a significantly better delay (1,000 seconds) is only displayed in 23% of the cases. Finally, less than 10% of the samples show improvements of 5,000 seconds or more. Taking into account the average disconnected flight duration (14,400 seconds) this improvement represents one-third of the flight duration. Despite that, the percentage of flights in which we can find this improvement is small.

These results show that almost half of the samples have almost no improvement (less than 185 seconds). We can explain this result with two facts. First, the contacts between planes are not homogeneously distributed: some aircraft display high number of encounters while others remain isolated for most of the flight. In Figure 4b, we show the complementary cumulative distribution function (CCDF) of the number of replicated packets. In this figure, we can see that less than 20% ($2 \cdot 10^{-1}$) of the samples achieve a degree of replication higher than 2. The second factor is related to the geographic location of the encounters. In [20, 21] we showed that the contact locations are closely related to the aircraft trajectory and that encounters are produced between aircraft following similar routes (Figure 1). As a result, planes that encounter each other will leave the disconnected area at similar times. Therefore, packets do not have the chance to spread much further from its source node.

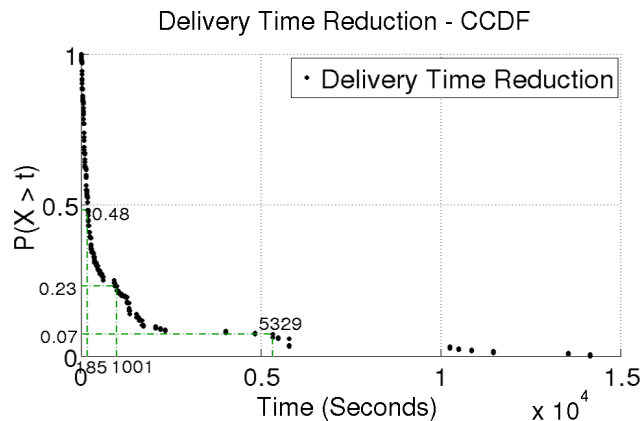


Figure 8: CCDF of the delivery time reduction

From these results, it can be concluded that achieving reliable air-to-ground communications using a purely opportunistic approach is challenging. There are several reasons for this: First, the delivery ratio is relatively low although this can be improved using sufficiently large buffers without expiration timers. Nevertheless, the main drawback is the limited dissemination of information and the apparition of node communities isolated from each other. A direct consequence of this is that more than half of the nodes show no delivery time reduction when forwarding data. Nevertheless, we would like to note that there is a small percentage ($\sim 7\%$) that offer fairly good results. This result shows that this network can achieve the initial objective at least in some cases.

5. Network Architecture with Satellite Nodes

The previously defined opportunistic network architecture offered very limited results. Therefore, in this section we propose a new approach that expands the network architecture to a hierarchical one with two differentiated node categories. We propose to combine opportunistic communications with the usage of satellite links (current state of the art aircraft air-to-ground communication system for disconnected environments). With this approach, we try to achieve a balanced, cost-effective solution that minimizes the requirement for expensive satellite links with the use of almost cost-negligible opportunistic communications. In this new network architecture, we define two differentiated node types:

- **Basic nodes:** Those nodes have the same characteristics as those used in the previous section. They are only able to use short range radio communications and will rely on opportunistic encounters to forward their data.
- **Satellite forwarder nodes:** These nodes have the basic characteristics of a basic node plus the capability of using satellite communications. These nodes will provide fast delivery of packets during flight-time allowing an alternative way for data delivery; instead of waiting for the data to be delivered when aircraft land they could be forwarded to these nodes during flight-time.

The inclusion of this new kind of node will allow faster air-to-ground communications resulting in improved delivery time reduction if compared to a purely opportunistic approach. The extent of this improvement will be directly determined by the number of satellite forwarder nodes that can be deployed. The main priority is to minimize this number (and thus reduce the economical cost) while maximizing delivery time reduction. This requires an optimal selection of the nodes most suited to perform this task. Basic nodes in the network still rely on opportunistic communications to transmit their data, and thus the selection of nodes must be based on factors such as the number, frequency and timing of contacts between those nodes. We achieve this selection by defining a graph representation of those contacts and then applying several graph analysis techniques to determine appropriate nodes capable of efficiently performing as satellite forwarders.

5.1. New Scenario Setup

The inclusion of the new communication model introduces several changes on our network scenario. Node configuration is analogous to the one used in previous experiments as shown in section 4.2. The requirements for routing protocols change; now data can be delivered both by nodes exiting the disconnected area and by nodes capable of satellite communications. Therefore, the use of the expected landing time (ELT) of airplanes is no longer justified and we only consider **epidemic routing**. In regards to nodes capable of satellite communications, we treat them as data sinks and mark any packet that reaches them as successfully delivered. We do not account for possible delays in the use of satellite links nor include a packet loss model for them as we consider these negligible in comparison to the ones displayed by the nodes using opportunistic communications.

6. Graph Analysis

In order to select the most appropriate nodes while taking the mobility and opportunistic factors into account, we propose a method based on defining the contact graph of the network and applying several graph analysis techniques on it. This graph displays the relationships between nodes and allows us to obtain detailed insight on the roles of each node within the network, and, therefore, a criteria to select which nodes would be the most appropriate.

6.1. Constructing the Contact Graph

We establish an undirected labeled dynamic graph $G(N, E)$. Each of the aircraft i is a node of the graph, $n_i \in N$, and the edge $e_{i,j}$ represents that nodes i and j had a contact between them. Contacts between the same pair of nodes happen only once, therefore, $e_{i,j}$ has no assigned weight. Each $n_i \in N$ has a limited lifetime represented by a label $n_i^{(start,end)}$ depicting the time interval in which this node is part of the graph. Each $e_{i,j}$ is also labeled with the time offset of the contact. Additional labels are included to provide geographical and trajectory information.

6.2. Metrics and Concepts

We study the properties of the contact graph and examine the relationship between the nodes. For this study, we consider the following concepts and metrics [11]:

- **Connected Component:** of an undirected graph is a subgraph in which any two nodes are connected by a path, and has no connection to nodes in the supergraph.
- **Network Diameter and Average Path Length:** represent the longest and average (respectively) of all the shortest paths between nodes in a graph.

- **Centrality metrics:** are a set of metrics that measure the structural importance of a node in the network.
 - **Degree:** is measured as the number of direct edges that reach a given node.
 - **Betweenness:** is equal to the number of shortest paths linking other nodes that pass through the node being measured.
 - **Closeness:** is based on the length of the shortest path linking to all other nodes.

6.3. Graph Structure

We start by analyzing some structural characteristics of the contact graph. For this structural analysis, we use the aggregated static version of the contact graph, constructed by aggregating all nodes and edges belonging to the graph at any interval of time.

The graph has 2,866 nodes and 9,530 edges representing the contacts between nodes. There is a total of 416 connected components, and 345 of those correspond to isolated nodes without any contact opportunity. In Figure 9a we show the distribution of nodes into connected components. There are two main connected components: one composed of 2,249 nodes (Component α) and another of 78 nodes (Component β). In Figure 9b, we can observe that the connected components of size smaller than 10 represent a 7.7 % of the network, but only a 1.43 % of the edges (contacts). Therefore, the structural location of isolated nodes and small connected components provide little advantage towards improving network connectivity and thus we disregard them.

We now study both connected components and measure the degree of interconnection between nodes. For component α , we observe a network diameter of 18 and average path length of 6.024. For Component β , a diameter of 9 and average path length of 3.531. These values mark the number of hops required for communication between any two nodes of the network.

In the following sections we offer a generic methodology that can be applied to all the connected components independently. Anyway, from this point onward we will be using component α , since it is the most heavily interconnected and represents the majority of the network. The other components are of little interest due to the scarcity of nodes and the inability to communicate with nodes in different components.

6.4. Inferring Node Notoriety

In this section, we locate the most influential nodes within the network, and for this purpose we use the centrality metrics defined in section 6.2. After computing the value of the metric for each node, we create a list of nodes ordered by it in decreasing order. We define the subsets D , B , and C , each subset contains a configurable percentage of the total number of nodes. The nodes in each subset are picked from the top of the list in decreasing order.

- $D \subset N$ (based on degree) is comprised of popular nodes with a high number of contacts with others, their structural position favors the role of conduits of information.
- $B \subset N$ (based on betweenness) is comprised of nodes whose structural position allows control over the information flowing through the network.
- $C \subset N$ (based on closeness) is comprised of nodes whose structural position allows a fast spread of information.

Nodes on these sets display interesting characteristics that make them good candidates to perform special roles within the network and, therefore, have a high chance of improving delivery times. Taking this into account we try to achieve the following goals: First, determine the extent of the performance improvement achieved by a hierarchical architecture over a purely opportunistic one. Secondly, evaluate the efficiency of using centrality metrics as the criteria to choose the adequate nodes, and determine which is the most effective centrality metric to perform this task.

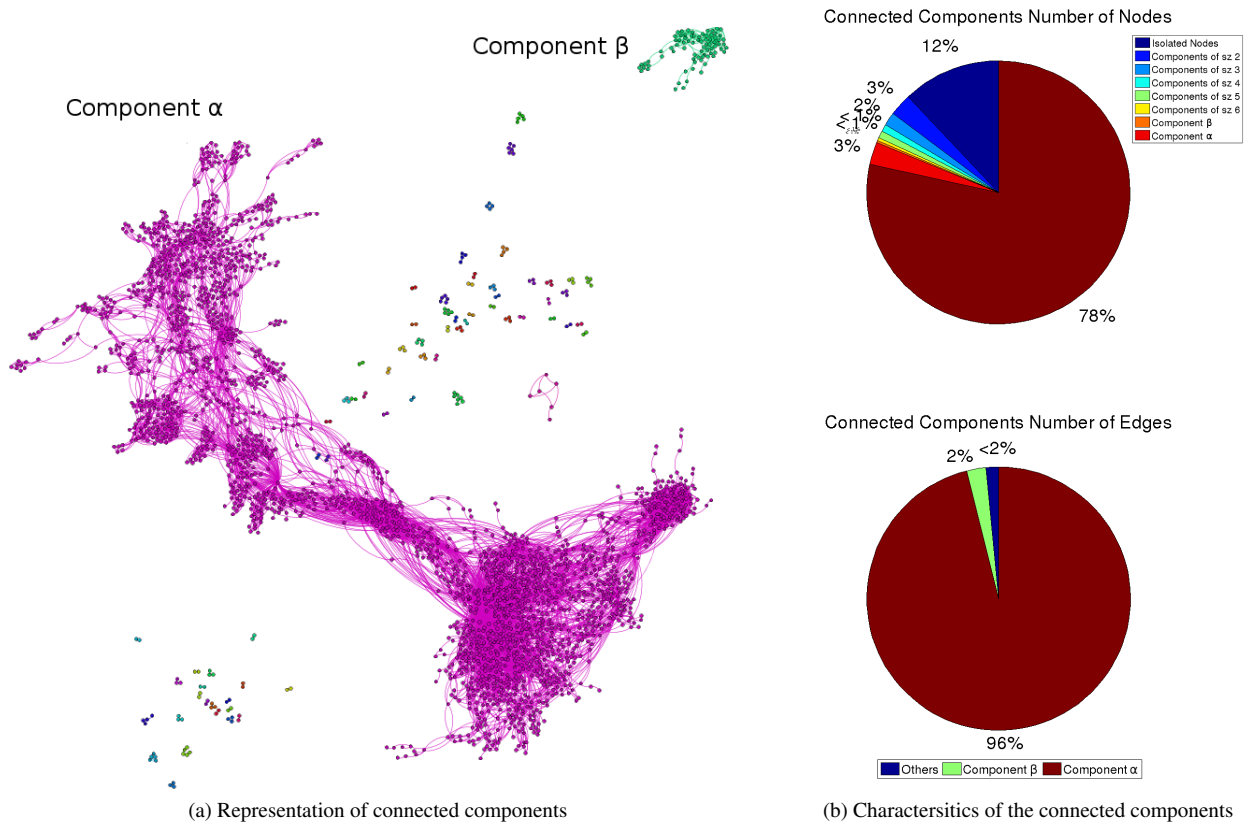


Figure 9: Contact graph

6.5. Experimental Results

To evaluate our results we use the network metric **delivery time reduction** as defined in section 4.4. Additionally, we define a fourth set U that is comprised of nodes chosen uniformly at random and will be used as a baseline for comparison purposes.

We evaluate delivery time reduction after the inclusion of satellite nodes, using all four sets B, C, D , and U with a varying percentage of selected nodes (from 5% to 90%). In Figure 13 we show the average delivery time reduction corresponding to the four sets for every percentage of nodes. In Figures 10, 11, 12 we show the tail distribution in linear scale corresponding to the configurations 5%, 15% and 40% of the nodes, which are relevant points of the exponential growth displayed by the values of Figure 13.

First, we focus on Figure 10. We can see that all four selection criteria provide a fair improvement when compared with a pure opportunistic approach (Figure 8). In this case, more than half of the samples offer an improvement higher than 5,000 seconds (one-third of the average disconnected flight duration). The next interesting fact is that using closeness offers very poor results, even lower than selecting nodes at random. Conversely, using betweenness and degree proves to be a good selection criteria as they offer far better results than using a random selection.

From these results, we conclude that delivery time reduction is mostly influenced by nodes acting as hubs for information flow (betweenness) or those with a high quantity of contacts themselves (degree). As opposed to closeness, which evaluates how fast the information may spread from a concrete node but it does not imply that the node itself handles a big amount of information.

In Figure 11, we can observe a higher percentage of samples with increased values for delivery time reduction than those of the previous figure. In the case of betweenness, we can see that more than 88% of the samples achieve a delivery time reduction of 5,000 seconds or higher. In the case of degree, this value is about 82%. In Figure 12

these values are further increased. The threshold of 5000 seconds is achieved in 95% of the cases by the samples corresponding to betweenness and 90% by the samples corresponding to degree.

As the number of satellites nears the 100%, the sample distribution tends to stabilize. All four distributions converge, and the shape of the distribution is analogous to that of the disconnected flight duration seen in Figure 4a. We can see that the disconnected flight duration defines the optimal upper limit. This is a logical result because the delivery time reduction can not increase any further than the duration of the flight itself.

Finally, in Figure 13, we can observe that the delivery time reduction increases exponentially based on the number of nodes capable of satellite communications. First, we can see that both betweenness and degree increase exponentially reaching a maximum and stabilizing at around 30% of the nodes. The random set does not follow a clear distribution, and several oscillations can be observed. The samples based on closeness have a much more limited growth and are significantly lower than the rest. Finally, we can see that all four curves eventually tend to similar values. A theoretical limit close to the average disconnected flight duration as observed before. We can also derive the actual delivery time for each configuration by subtracting the average time reduction to the average delivery time of the network with no satellites (~ 14900 seconds).

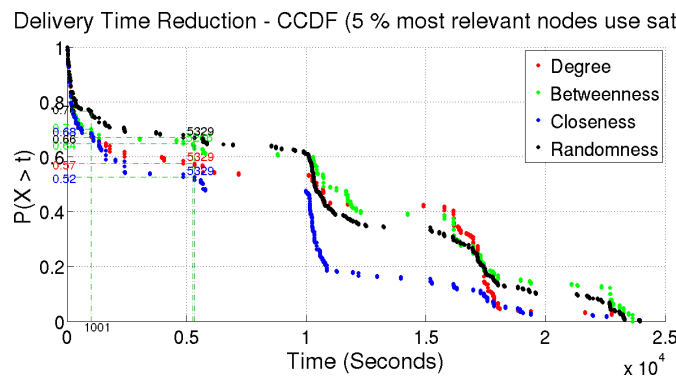


Figure 10: CCDF of the delivery time reduction using 5% of total nodes selected from all sets




Figure 11: CCDF of the delivery time reduction using 15% of total nodes selected from all sets

7. Ensuring Temporal Continuity

In this section, we consider the dynamic graph in order to analyze the temporal relationships between nodes. Once we consider that contacts happen at specific time intervals, there are several facts that should be taken into account. First, if an edge connects two nodes, it is ensured that they have a contact at some point in time and also that they will be able to exchange information. Despite that, if the path linking two nodes has a distance greater than one hop, there is no guarantee that the information can be exchanged between these two nodes. If we assume that nodes i and j are

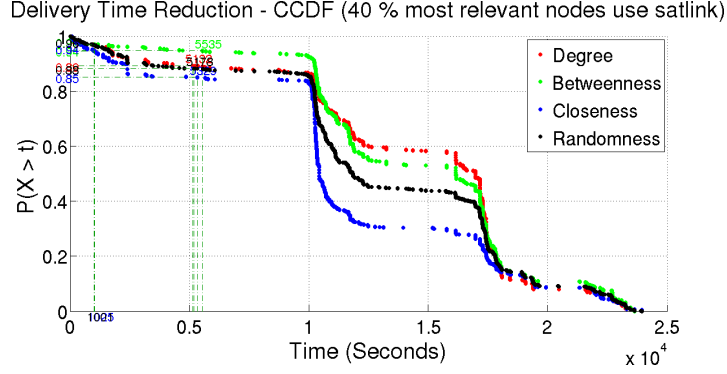


Figure 12: CCDF of the delivery time reduction using 40% of total nodes selected from all sets




Figure 13: Average delivery time reduction for all thresholds

linked by a n hop path, node i is only capable of delivering data to node j if each contact along the path is produced at a later time than the previous one. Our objective in this section is to reveal which paths can be successfully exploited to communicate any given two nodes of the network. Specifically our primary interest is which paths can connect basic nodes with nodes inside one of the notorious subsets.

To this end, we define the two following subsets:

- $R \subset N$: stands for relevant nodes and may be equal to either B , D or C .
- $I \subset N$: stands for irrelevant nodes and is defined as the relative complement of R in N ($I = N \setminus R$).

We consider a directed dynamic graph in which edges are directed towards nodes in R , and we study to what extent a node in R can be reached by nodes in I . We consider the *temporal path* [33] between each node in I towards each node in R . The temporal path between two nodes i and j is defined as a sequence of n hops through a distinct node n_k^t at time t_k : $p_{ij} = (i, n_1^t, \dots, n_k^t, j)$ where $t_{k-1} < t_k$. If a *temporal path* from i to j does not exist, (i, j) is a *temporarily disconnected node pair*.

We define the numeric value *available temporal paths* which is assigned to each node in R according to the following procedure:

- Compute all possible combinations of node pairs (i, j) where $i \in I$ and $j \in R$.
- For each node pair, find the shortest path between nodes i, j .
- Evaluate if this path is a *temporal path*.
- If the conditions are meet, then the value of *available temporal path* for node j is increased.

Once all the paths have been explored, we use the aggregated values of all explorations to define a new subset of nodes:

- $R_{tmp} \subseteq R$: is an ordered set comprised of nodes from R ordered decreasingly according to the *available temporal paths* values.

We compute the new R subset for each of the notorious subsets (D, B, C) and analyze their similarities. We apply Spearman rank correlation to compare R and R_{tmp} , the average of the correlation coefficient for the sets D ($\rho = 0.0067$ with p-value of 0.44661), B ($\rho = 0.0012$ with p-value of 0.4583), and C ($\rho = 0.0158$ with p-value of 0.5158), and it shows that there is little correlation between both set orders. This value indicates that occupying a significant position within the graph does not necessarily ensure good communication with all the other nodes when a temporal factor is taken into account.

We estimate that notorious nodes with a reduced value of *available temporal paths* should have a low impact on delivery time reduction because other nodes are not able to deliver data to them. Based on this concept, we strive to further optimize satellite link usage by discarding nodes located on the lower part of the R_{tmp} subset.

7.1. Experimental Results

To evaluate our results we use the network metric **delivery time reduction** as defined in section 4.4. We compare the performance of satellite links based on both R and R_{tmp} . As R we use the subset B , since it is the one that offers the best performance in the static analysis of the contact graph. After constructing B_{tmp} we discard a 10% of the lower ranked nodes in such a way that $B_{tmp} \subset B$.

In Figure 14 we show the results obtained using three sets of nodes: First, the base set B using a 20% of network nodes. Then, the corresponding set B_{tmp} discarding a 10% of the nodes. Lastly, we define a reduced version of B for comparison purposes. We call this set B_r and it has the same size as B_{tmp} . We construct this set from B by directly discarding the last 10% of nodes (instead of performing the temporal path exploration). As we can see, the samples from B_{tmp} display slightly lower performance as those of B , but a noticeable improvement if compared to the sample B_r of its same size.

In Figure 15 we show the average delivery time reduction for every percentage of nodes using the previously described three node selections. As we can see, B and B_{tmp} display very similar results but always slightly better than B_r . All three samples show an exponential curve. The relationship between the subsets can be observed in the first stages of the exponential growth, at higher percentages of nodes all three curves stabilize and converge to similar values.

These results show that B_{tmp} is an optimized version of B that requires fewer satellite links to achieve slightly lower performance but with definite better results than using B with fewer satellites. From these results, we can see that the new order imposed by the temporal restrictions allows us to discard the nodes that despite being relevant will have a poor contribution to network performance.




Figure 14: Comparison between basic and temporal ordering performance

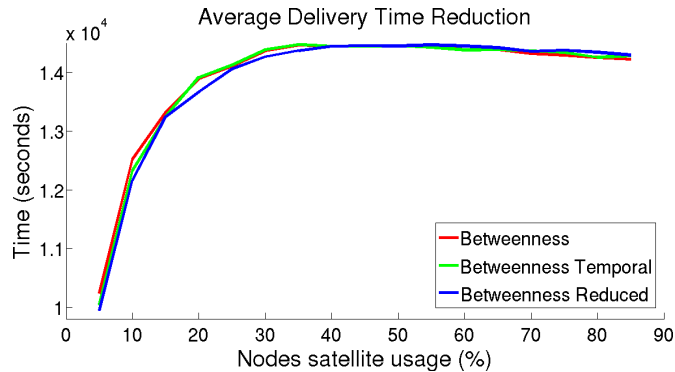


Figure 15: Comparison between basic and temporal ordering performance

8. Discussion

Nowadays, airlines operate in rather thin profit margins, and one of the main concerns is optimization and reduction of costs. Our paper offers a new approach that takes the currently deployed technology (satellite links) and offers an alternate extension at negligible cost (opportunistic communications) that if used in tandem can greatly reduce the expenditure in satellite links.

According to [32], equipping a single aircraft with satellite based equipment for air-to-ground Wi-Fi capabilities has a minimum cost of \$100,000. Retrofitting unused equipment from one plane to another is also a very costly and time consuming task. It is also stated that some airlines are currently offering on-board Wi-Fi at a loss in part due to the high investment required for equipment. In regards of usage costs, satellite providers operating in the KU-Band are the most economic offers available on the market. According to [17], the estimated cost of those services is 20 cents per megabyte, which is about \$200 per gigabyte. If we compare it to the rates of similar services such as 4G LTE (usually around \$10 per gigabyte) it is obviously a very costly service.

Our approach provides a way to install satellite links in just a reduced percentage of nodes and yet still obtain delivery-times fairly close to those that would be obtained by installing a satellite link per aircraft. From the results obtained we conclude that the optimal selection criteria is to use betweenness, as seen in Fig. 13, which is able to reach optimal values with approximately just using a 35% of total nodes. It is straightforward to affirm that this approach reduces the cost of fitting an airline’s fleet to almost one third of its price and this base cost may be further reduced by applying optimization based on temporal relations.

9. Conclusions and future work

In this paper we provided a detailed evaluation of a novel aeronautical ad-hoc network architecture based on opportunistic communications and satellite communication systems. We provided an off-line generic methodology based on analyzing the graph representation of the contacts between nodes. The analysis of its structural characteristics allowed us to discern isolated nodes apart from heavily connected nodes. Afterwards, while focusing on the most relevant connected components, we applied some centrality metrics to identify the most relevant nodes. We deployed a network infrastructure based on opportunistic communications but with some of the nodes, selected with the insight acquired from the centrality metrics, also capable of using satellite communications. We analyzed the performance of such network for its use in air-to-ground communications. Our results show that equipping just a percentage of nodes (~35%) with satellite communication systems, according to a selection based on betweenness centrality, it is possible to obtain performance on par with those architectures that equip satellite systems on all nodes. We further reduced the number of nodes required by proposing a methodology that considers the temporal relations between nodes and discards those that have little impact on network performance. Those contributions can be used to greatly reduce the deployment cost of satellite based architectures in these kind of scenarios and are not specifically dependent on the model used.

This paper offered an effective network architecture for data delivery in air-to-ground communications. The immediate future work revolves around the study and deployment of real applications in such a network, focusing on applications which do not require data round-trip, and considering two main types based on the final user:

- Passenger data: examples of such applications are delivery of mail or twitter messages.
- Airline data: an example of such application is early notification of faulty components prior to landing for compliance with MEL (Minimum Equipment Lists), which are safety regulations lists that define which items (instrument, equipment or systems) are required for the aircraft to be allowed to depart after an intermediate stop. These lists can include non-critical components such as cabin lighting items, and a departure may suffer delays due to the time required for finding a spare and replacing such component. For this reason, early notification of any faulty equipment during flight can allow for faster response after landing.

Another topic for future work is the evaluation and study of the architecture assuming other communication scenarios such as air-to-air or air-to-ground communications.

Acknowledgment

This work was partly supported by the Catalan AGAUR 2014SGR-691 project, Spanish MINECO TIN2014-55243-P project, and PhD grant UAB PIF 472-01-1/E2010. We would like to thank Adrián Antúnez for his collaboration with the early implementation of the graph modelization.

10. References

- [1] Ahmed, S., Salil, S. K., 2010. Characterization of a large-scale delay tolerant network. In: Proceedings of the 2010 IEEE 35th Conference on Local Computer Networks. LCN '10. IEEE Computer Society, Washington, DC, USA, pp. 56–63.
- [2] Aircell Business Aviation Services, 2014. <http://www.aircell.com>.
- [3] Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., Weiss, H., April 2007. Delay tolerant network architecture. IETF RFC 4838.
- [4] Cheng, Y., Çetinkaya, E. K., Sterbenz, J. P., October 2011. Performance comparison of routing protocols for transactional traffic over aeronautical networks. In: Proceedings of the International Telemetering Conference (ITC). Las Vegas, NV.
- [5] Daly, E. M., Haahr, M., 2007. Social network analysis for routing in disconnected delay-tolerant manets. In: Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing. MobiHoc '07. ACM, New York, NY, USA, pp. 32–40.
- [6] Devroye, L., 1986. Non-Uniform Random Variate Generation. Springer-Verlag.
- [7] Ebsco Host, 2013. Intelsat to Provide Satellite Connectivity Capacity for Gogo's International In-flight Wi-Fi Service. <http://connection.ebscohost.com/c/articles/82832211/intelsat-provide-satellite-connectivity-capacity-gogos-international-in-flight-wi-fi-service>.
- [8] Farrell, S., Cahill, V., 2006. Delay and Disruption Tolerant Networking. Artech House.
- [9] Flightaware, 2013. Flightaware, flight tracker , flight status , flight tracking. <http://es.flightaware.com/>.
- [10] Flightview, 2013. Real time flight tracker & airport delays. <http://www.flightview.com/TravelTools/>.
- [11] Freeman, L. C., Mar. 1977. A Set of Measures of Centrality Based on Betweenness. *Sociometry* 40 (1), 35–41.
- [12] Galati, A., Vukadinovic, V., Olivares, M., Mangold, S., 2013. Analyzing temporal metrics of public transportation for designing scalable delay-tolerant networks. In: Proceedings of the 8th ACM Workshop PM2HW2N. ACM, New York, NY, USA, pp. 37–44.
- [13] Gu, W., Li, J., He, F., Cai, F., Yang, F., 2012. Delay-aware stable routing protocol for aeronautical ad hoc networks. *Journal on Information & Computational Science*.
- [14] Hui, P., Crowcroft, J., Yoneki, E., Nov 2011. Bubble rap: Social-based forwarding in delay-tolerant networks. *Mobile Computing, IEEE Transactions on* 10 (11), 1576–1589.
- [15] Jabbar, A., Sterbenz, J. P., October 2009. AeroRP: A geolocation assisted aeronautical routing protocol for highly dynamic telemetry environments. In: International Telemetering Conference (ITC) 2009. Las Vegas, NV.
- [16] Jain, S., Fall, K., Patra, R., Aug. 2004. Routing in a delay tolerant network. *SIGCOMM Comput. Commun. Rev.* 34 (4), 145–158.
- [17] Karp, G., 2012. Sky-high wifi nearly ready to widely fly. <http://articles.chicagotribune.com/>.
- [18] Kingsbury, R., 2009. Mobile ad hoc networks for oceanic aircraft communications. Master's thesis, Massachusetts Institute of Technology. Dept. of Aeronautics and Astronautics.
- [19] Li, H., Yang, B., Chen, C., Guan, X., 2010. Connectivity of aeronautical ad hoc networks. In: GLOBECOM Workshops (GC Wkshps), 2010 IEEE. pp. 1788–1792.
- [20] Martínez-Vidal, R., Martí, R., Borrell, J., 2013. Characterization of a transoceanic aircraft delay tolerant network. In: Proceedings of the 38th Annual IEEE Conference on Local Computer Networks. IEEE Computer Society, pp. 590–597.
- [21] Martínez-Vidal, R., Martí, R., Borrell, J., 2014. Analyzing information propagation in a transoceanic aircraft delay tolerant network. In: Proceedings of the 39th Annual IEEE Conference on Local Computer Networks. IEEE Computer Society, pp. 116–123.

- [22] Medina, D., Hoffmann, F., Ayaz, S., Rokitansky, C.-H., 2008. Feasibility of an aeronautical mobile ad hoc network over the north atlantic corridor. In: *Sensor, Mesh and Ad Hoc Communications and Networks*, 2008. SECON '08. 5th Annual IEEE Communications Society Conference on. pp. 109–116.
- [23] Mitchell, J., Mar. 4 2003. Aircraft satellite communications system for distributing internet service from direct broadcast satellites. US Patent 6,529,706.
- [24] New Straits Times, Malaysia, 2014. Missing MH370: Plane's ACARS cut off time irrelevant to its location. <http://www.nst.com.my/latest/font-color-red-missing-mh370-font-plane-s-acars-cut-off-time-irrelevant-to-its-location-1.519898>.
- [25] ns-3, Network simulator 3, 2014. <http://www.nsnam.org>.
- [26] OpenFlights.org, 2012. Airport, airline and route data. <http://openflights.org/data.html>.
- [27] Parzen, E., 1962. On estimation of a probability density function and mode. *Ann. Math. Statist.* 33 (3), 1065–1076.
- [28] Peters, K., Çetinkaya, E. K., Jabbar, A., Sterbenz, J. P. G., October 2010. Analysis of a geolocation-assisted routing protocol for airborne telemetry networks. In: *Proceedings of the International Telemetering Conference (ITC)*. San Diego, CA.
- [29] Peters, K., Jabbar, A., Çetinkaya, E. K., Sterbenz, J. P., March 2011. A geographical routing protocol for highly-dynamic aeronautical networks. In: *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*. Cancun, Mexico, pp. 492–497.
- [30] Rohrer, J. P., Sterbenz, J. P., October 2009. Performance and disruption tolerance of transport protocols for airborne telemetry networks. In: *International Telemetering Conference (ITC) 2009*. Las Vegas, NV.
- [31] Sakhaee, E., Jamalipour, A., Kato, N., June 2006. Multipath doppler routing with QoS support in pseudo-linear highly mobile ad hoc networks. In: *Communications, 2006. ICC '06. IEEE International Conference on*. Vol. 8. pp. 3566–3571.
- [32] Schawalder, J., 2014. The future of inflight entertainment in europe, according to passenger expectations: Why airlines should embrace consumer technology. Master's thesis, Grenoble Ecole de Management.
- [33] Tang, J., Musolesi, M., Mascolo, C., Latora, V., Nicosia, V., 2010. Analysing information flows and key mediators through temporal centrality metrics. In: *Proceedings of the 3rd Workshop on Social Network Systems. SNS '10*. ACM, New York, NY, USA, pp. 3:1–3:6.
- [34] Ubiquiti networks, 2013. Bullet: Revolutionary outdoor radio device. <http://www.ubnt.com/bullet>.
- [35] Vahdat, A., Becker, D., 2000. Epidemic routing for partially-connected ad hoc networks. Tech. rep.
- [36] Xia, C., Liang, D., Wang, H., Luo, M., Lv, W., 2012. Characterization and modeling in large-scale urban dtns. In: *Local Computer Networks (LCN), 2012 IEEE 37th Conference on*. pp. 352–359.
- [37] Zhang, X., Kurose, J., Levine, B. N., Towsley, D., Zhang, H., 2007. Study of a bus-based disruption-tolerant network: Mobility modeling and impact on routing. In: *Proceedings of the 13th Annual ACM International Conference on Mobile Computing and Networking. MobiCom '07*. ACM, New York, NY, USA, pp. 195–206.