Cumulus Humilis: Wireless Mesh Networking for Gliders

Roel Baardman and Nirvana Meratnia
Pervasive Systems (PS), University of Twente
PO Box 217, 7500 AE Enschede, The Netherlands
rbaardman@zonnet.nl, n.meratnia@ewi.utwente.nl

Presented at the XXX OSTIV Congress, Szeged, Hungary, 28 July - 4 August 2010

Abstract

Digital communication between gliders would be beneficial because exchange of information has the potential to provide insight not only to the pilot but also other pilots and people on the ground. However, such a communication does not exist. In this paper we present a solution for the exchange of information using opportunistic dissemination – pro-actively spreading information - through a wireless mesh network - a network in which each participant (node) can act as a router. A challenge faced in designing our opportunistic dissemination is how to handle different types of information, as information regarding an emergency has a higher priority than information regarding the weather, for instance. We evaluate our solution via simulation.

Introduction

During cross-country flights, receiving information about the weather can be valuable for the pilot. Delivery of large-scale weather developments to the pilot enables them to adapt their strategy early and not to be surprised by the changing weather conditions. For example, when high clouds approach blocking incoming sunlight and, thus, reducing the formation of thermals, the pilot, then, could be informed to head home earlier to prevent an outlanding.

Being able to follow gliders can be valuable in two ways. Firstly, this enables friends and family of the pilot to see where the glider is and, thus, share part of the excitement of cross-country gliding. They can see how the flight progresses. Secondly, if an accident takes place, this allows for quickly determining the location of the pilot. Knowing the last location of the pilot can reduce the search area for the ground crew, and increase the chance for survival of the pilot. During measurement flights, for example in the Mountain Wave Project[1], information exchange enables feedback to the measuring glider. Analysis on the ground can determine if the region which is currently being measured is relevant to the goal of the measurement, and give feedback to the pilots performing the measurements. If multiple gliders are performing measurements simultaneously, they can collaborate and scan an area for interesting phenomena. When an interesting phenomenon is found, all gliders, then, can fly towards the area and perform detailed measurements of the phenomenon.

During these flights, however, no digital communication exists between the glider and other gliders, or the glider and individuals on the ground. This communication, however, would be beneficial for cross-country flights and measurement flights.

Problem statement

In this paper we present a solution for the exchange of information between gliders using a wireless mesh network, a network in which each participant (node) can act as a router. This section will describe problems associated with wireless transmission and wireless mesh networking.

In wireless transmission, the hidden terminal problem is a problem related to medium access control, and is visualized in Fig. 1. One node, in this case the "Hub", is within the transmission range of both node A and node B. When both node A and node B are transmitting simultaneously, the Hub will receive none of the two transmissions, since these will interfere with each other. However, node A and node B are not in each other's transmission range. This makes node A and node B unable to coordinate their respective transmissions such that the Hub node can receive those.

Another common problem in wireless mesh networking is the changing network topology. Due to the fact that nodes in the network move around, connections can break and new connections can become available. To allow two nodes to continue sending packets to each other when a connection breaks, these changes should be known and packets should be routed appropriately considering the new network topology. When topology is such that no direct route from the source of a packet to the destination of a packet is readily available, a store-and-forward approach can be utilized to deliver the packet at the destination node.

Providing Quality of Service (QoS) in mesh networking is another challenge. Since both network topology and signal strength change, prioritizing one type of data over another type of data is not straightforward. Packet loss and channel characteristics are not predictable, thus sending high-priority data before low-priority data can result in packet loss of the high-priority data. The solution we present in this paper has to address these issues, in order to function properly.

Analysis

Before we look at related work, we analyze logfiles for real simultaneous flights. This provides insight on the circumstances we can expect, in terms of how fast the network topology changes.

We analyze logfiles of real simultaneous flights to gain insight on possible connection times, the distance between two connected nodes, the number of connections a single node has, and the effect of the speed and direction to which two nodes are traveling during the time they can be connected. Downloaded from Online Contest[2], we group logfiles of the same days. The reason for doing so is to have groups containing both good and poor cross-country flying conditions. This enables us to analyze both conditions with many and few gliders air-borne.

To analyze the files, we create a script written in the PHP scripting language. This script reads all the logfiles and produces a table for each logfile. Each entry in this table contains a timestamp and the location of the glider at that time. However, not all logfiles contain Global Positioning System (GPS) locations at precisely the same timestamps. Therefore, for fixed timestamps the location of the glider is calculated. If the location for a specific timestamp is not known, it is calculated via interpolation.

Optionally we use also another file, which contains the locations of all Dutch gliding clubs. This file enables simulation with stationary nodes at airfields, which can improve connectivity and allow individuals on the ground to access the network.

After reading all these files, we check connectivity. At every time instance, connectivity between every pair of two nodes - computers in a glider cockpit or a computer at an airfield – is checked. Two nodes are assumed to have a connection when their relative distance is less than 40 kilometers, which is the maximum range in The Netherlands of the XBee PRO 868 transceivers[3] used in this paper. In the case of available connectivity, a record for each link will be created containing:

- The nodes involved.
- The first moment the connection is available.
- The last moment the connection is available.
- The duration of the connection.
- The average distance between the involved nodes during the available connectivity.
- The vector dot-product between the average speeds of the two nodes involved.

Next, the fan-out of each node at each point in time is analyzed. At each point in time, for each node, a set with all the node's direct neighbors and the node's neighbor’s neighbors is created. The overall maximum fan-out is stored, as an indication to how sparse or dense the network is for certain conditions. It provides insight on, for example, the effect of nodes at airfields.

We write also a Keyhole Markup Language (KML) file - which can be viewed in Google Earth - containing records of all GPS locations at all available timestamps, and records of all links. This allows a visual check on the correctness of the script and to get insight on the circumstances. Links are colored depending on the distance between the two involved nodes: a distance of 20 to 40 kilometers is colored red, a distance of 10 to 20 kilometers is colored yellow, and a distance smaller than 10 kilometers is colored green. Figs. 2 and 3 illustrate the generated KML files.

Related work

This section examines existing work in the field of mesh networking. Analysis of logfiles downloaded from Online Contest[2] presents circumstances where connectivity is good enough to not have disconnected segments in network coverage, as seen in Fig. 2. For these circumstances, we look at existing approaches for unicast routing and multicast routing in mesh networks, which find routes to specific hosts in the network. For circumstances in which connectivity is not good enough and the network sometimes has disconnected segments of network coverage, which can be seen in Fig. 3, we look at existing opportunistic dissemination approaches. Dissemination protocols focus on spreading data in the network where no route is available between hosts.

Routing in mesh networks can be categorized into pro-active protocols and reactive protocols. Pro-active protocols, like Destination-Sequenced Distance Vector (DSDV)[4] and Optimized Link Source Routing (OLSR)[5], exchange routing information continuously. Each node in the network broadcasts its current routing table to its neighbors, who do the same. When a node wants to send a packet, it can simply look up the destination of the packet in its routing table to determine the next hop for the packet. Thus, the time it takes for a node to send the packet is relatively short. However, since the nodes are continuously exchanging routing information, packets are sent constantly between nodes. Even when no data packet is sent over the network, the nodes are busy exchanging routing information. Thus, control overhead is high in this situation. When connectivity in the network is short, there is also the risk that nodes route information based on outdated routing information.

In reactive protocols, like Ad-Hoc On-demand Distance Vector (AODV)[6, 7, 8, 9], Dynamic Source Routing (DSR)[10] and On-Demand Multicast Routing Protocol (ODMRP)[11, 9, 12], the nodes in a network do not continuously exchange routing information, but only when a node wants to send a data packet. This node then initiates the process of path discovery. During path discovery, the node discovers only the current route to the destination it wants to send the packets. All other available routes in the network should be discovered separately. The latency before sending a packet is higher, since the node originating the packet has to first discover the path to the destination of the packet.

Dissemination can be categorized into reactive approaches - like the typical client/server model in which consumers request for specific information - and pro-active approaches - like broadcasting, publish/subscribe mechanisms and multicast - in which information is pro-actively disseminated whenever such information is available in the network.

Reactive approaches for dissemination are hard to implement in wireless mesh networks with mobility[13]. Therefore, we will not look into these approaches.

Opportunistic dissemination[13, 14, 15, 16, 17] is a type of proactive approach and transport the information on the network, without the information specifically being requested. This means that the node does not need to send requests to the network, instead the network delivers information to the node.

Cumulus Humilis

Our Cumulus Humilis (CuHum) protocol aims to disseminate data among gliders in both sparse and dense glider networks. More specifically, it aims to achieve a large delivery rate with a small propagation delay while honoring relative packet priorities. For this purpose, we take the following approaches:

- Desynchronizing hosts using random transmission moments to prevent the hidden terminal problem.
- A store-and-forward communication model to on nodes physically carry packets to new locations whenever a multi-hop connection is not available between gliders.
- Assigning each type of traffic a priority, allowing nodes to prioritize packets with a high priority over packets with a low priority.

We make the following assumptions for the gliders:

- A mobile computer, for example an iPaq - is available in the cockpit of gliders. This is the mobile computer which already runs existing in-flight navigation software.
- This mobile computer is connected to a GPS receiver.
- Enough batteries are available in the glider to power all electronic equipment for the duration of all flights during the day. Thus, power consumption is not a factor in the operation of our protocol. The X-Bee PRO 868 modules used in this paper have a maximum operating current of 800 mA at 3.3V.
- The mobile computers have limited storage available for storing packets.

We also make the following assumptions on antenna radiation and signal propagation:

- The transceivers available at the gliders have a maximum transmission range of 40 km.
- Signal strength decays over distance in the order of distance to the forth (d^4).
- The antennas used spread all transmission power equally in all directions. Thus, the received signal strength of a radio transmission does not depend on the relative orientation of the sender and receiver of the transmission. Only the relative distance affects the received signal strength.

Protocol operation

The core of the protocol we propose is basic dissemination. Basic dissemination in our protocol regularly broadcasts received packets locally. When a node receives packets it has not encountered before, it stores this packet in memory to disseminate the packet at a later time. A packet is regularly broadcasted by all nodes receiving it, until nodes run out of memory.

To prevent packets from residing in the network for extended periods, each packet carries a deadline, i.e. maximum age it can reach. When a packet's age has reached the deadline, the packet is removed from memory. Since nodes have only limited storage available to store packets, sometimes packets will have to be deleted from storage. In this case, CuHum deletes the packet with the lowest priority. This causes packets with high priority to remain in storage whereas packets with low priority will be deleted when storage becomes scarce.

To maximize the chance different broadcasts of a packet from one node are received by as many nodes as possible, we spread the different broadcasts evenly over time. However, when multiple packets have to be broadcasted at the same time by the node, scheduling is needed to optimally spread all different packet broadcasts over time. We utilize Earliest Deadline First (EDF)[18] scheduling for this purpose. Each packet broadcast is assigned a release date, i.e. the time at which the packet broadcast is available, and a due date, i.e. the time at which the packet should have been sent. We, then, sort all packet broadcasts based on their deadline and broadcast the packets in this order.

Performance evaluation

To evaluate Cumulus Humilis, we write a discrete-time simulator in the C++ programming language. The simulator uses GPS logfiles downloaded from Online Contest[2] to model glider movement. The locations of gliding airfields are also used. The simulator uses discrete time steps of 1 millisecond. Each millisecond the movement of every node, and the protocol behavior – for instance scheduling a packet for transmission or processing a received packet - of every node is determined. The radio model used in the simulator assumes that signal strength decays over distance following d^4. A transmission is readable for a node, when, during the entire transmission, the signal-to-noise ratio is higher than 10 (a common value for the signal-to-noise ratio[19]). For all other transmissions, the noise is calculated again through d^4. If a transmission exists which has a ratio of 10 in signal strength over all other transmissions, we assume this transmission is readable by the node.

Metrics used in simulation

We use the following metrics for our evaluation:

- Successful Transmission. We define a transmission to be successful when all the bytes in the transmission have been sent by the transceiver at the sending node.
- Successful Reception. We define a reception of a packet to be successful when the receiving node has successfully received all the bytes of a transmission and has decoded the transmission.
- Number of Successful Receivers of a packet. We define the number of receivers of transmissions to be number of Receptions divided by the number of Transmissions.
- Latency. We define the latency of a packet as the age of the packet when it is received for the first time by a node.
- Packetloss. We define packetloss as the ratio of all packets not successfully received divided by all packets successfully sent. Since all traffic in the network is broadcast traffic, one packet transmission has multiple receivers. Thus, the amount of packets not successfully received can be larger than the amount of packets successfully sent.
- Delivery ratio. We define the delivery ratio of an information type as the number of nodes which on average receives packets for the type of information, divided by the number of nodes. Thus, it illustrates the probability any node will receive a packet from this type of information.

Simulation results

During simulation we look at two specific scenarios:

- Scenario 1: A scenario with good connectivity and airfields turned on to evaluate the performance of the protocol in circumstances in which a lot of traffic is sent.
- Scenario 2: A scenario with bad connectivity and airfields turned off to evaluate the performance of the protocol in circumstances in which the opportunistic dissemination features of the protocol have to be used to deliver packets.

First we evaluate a scenario with good connectivity for which we use GPS logfiles from a day of the Dutch National Championships. To increase connectivity we enable also airfields. One of the airfields disseminates information about the weather. One of the gliders is performing a measurement flight during this day and disseminates these data. Figure 5, Fig. 6 and Fig. 7 show, respectively, the latency of packets containing positional information, measurement information and weather information. We see that positional information and measurement information deliver most packets with a low latency. Positional information has priority over measurement information, causing the latter to quickly be erased from the network. However, it spreads better than weather information. This information spreads much less, which can be seen by the much lower number of packets overall delivered. Figure 10 shows the delivery ratio during this scenario. We see that positional information has the best delivery ratio, followed by weather information and measurement information. Although measurement information has a higher priority in the network, and, therefore, would have a higher delivery ratio, it has an earlier deadline compared to weather information. This prevents the measurement information from reaching more nodes.

The other scenario we evaluate is a scenario with sparse connectivity for which we use GPS logfiles from a day during October 2009. To further decrease connectivity, we disable airfields. During the day network segments meet and separate again. The connections between these segments are through one route with several hops. Figure 8 and Fig. 9 show, respectively, the latency of packets containing positional information and measurement information. In these figures the effect of the regular broadcast can be clearly seen. Since two large networks of gliders are connected through a few hops, both positional information and measurement information propagates through these hops. The result of this is the increase of the number of packets for the latency between 20 and 45 seconds which illustrates the time it approximately takes for the packets to travel from one large network of gliders to another. The second increase of packets for the latency above 50 seconds shows opportunistic dissemination. Packets which are stored in memory are eventually delivered when a connection is possible between two segments of the network. Figure 11 shows the delivery ratio during this scenario. We see that the limited number of nodes in the network causes both types of information to be delivered equally well. Since the network is often disconnected during the scenario, not all nodes can be reached.

Figure 12 shows successful transmissions, successful receivers and the calculated packet loss for both scenarios.

Conclusion and future work

In this paper we have proposed CuHum (Cumulus Humilis), an opportunistic dissemination protocol for the exchange of information between airborne gliders using a wireless mesh network. It applies a store-and-forward communication model to physically carry packets to new locations and assigns information a priority. Simulations show the protocol is able to handle situations with rich connectivity and situations with sparse connectivity.

The work in this paper is part of a master thesis[20] carried out at the University of Twente, The Netherlands.

Future work could focus on reducing packetloss for the broadcast transmissions used in CuHum, as this would increase the amount of traffic the network can handle. Also, adding a distributed publish/subscribe layer might improve efficiency, allowing the network to scale further.

References

[1] R. Heise, K. Ohlmann, and J. Hacker, “Mountain wave project," Apr 2010. http://www.mountain-wave-project.com/index-1.html

[2] Segelugszene GmbH, OLC – Gliding - Overview, Jun 2010. http://www.onlinecontest.org/olc-2.0/gliding/

[3] Digi International Inc., XBee-PRO 868 OEM RF Modules, Nov 2008. http://ftp1.digi.com/support/documentation/90001020_a.pdf

[4] C. E. Perkins and P. Bhagwat, “Highly dynamic destination-sequenced distance-vector routing (dsdv) for mobile computers," In *SIGCOMM '94 Proceedings of the Conference on Communications Architectures, Protocols and Applications*, pp. 234-244, 1994.

[5] T. Clausen and P. Jacqueet, “Rfc3626: Optimized link state routing protocol (olsr)," Oct 2003. http://tools.ietf.org/html/rfc3626

[6] C. E. Perkins and E. M. Royer, “Ad-hoc on-demand distance vector routing," In *IEEE Workshop on Mobile Computing Systems and Applications*, pp. 90-100, 1999.

[7] C. Perkins, E. Belding-Royer, and S. Das, “RFC 3561: Ad hoc on-demand distance vector (AODV) routing," Jul 2003. http://tools.ietf.org/html/rfc3561.

[8] C. E. Perkins, *Ad Hoc Networking*. Addison Wesley, 2001.

[9] T. Kunz and E. Cheng, “On-demand multicasting in ad-hoc networks: Comparing aodv and odmrp," *International Conference on Distributed Computing Systems*, 2002.

[10] D. Johnson, Y. Hu, and D. Maltz, “Rfc4728: The dynamic source routing protocol (dsr)," Feb 2007. http://tools.ietf.org/html/rfc4728

[11] S.-J. Lee, M. Gerla, and C.-C. Chiang, “On-demand multicast routing protocol," In *Wireless Communications and Networking Conference 1999. WCNC. 1999 IEEE*, pp. 1298-1302, 2002.

[12] S.-J. Lee, “Routing and Multicasting Strategies in Wireless Mobile Ad hoc Networks.” PhD thesis, University of California, 2000.

[13] A. Datta, “Autonomous gossiping in ad-hoc networks," 2003. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.58.3986&rep=rep1&type=pdf

[14] S. Chachulski, M. Jennings, S. Katti, and D. Katabi, “Trading structure for random ness in wireless opportunistic routing," In *Proc. of ACM SIGCOMM 2007*, Aug 2007.

[15] W. R. Heinzelman, J. Kulik, and H. Balakrishnan, “Adaptive protocols for information dissemination in wireless sensor networks," In *MobiCom '99 Proceedings of the 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking*, pp. 174-185, 1999.

[16] A. P. Sistla, O. Wolfson, and B. Xu, “Opportunistic data dissemination in mobile peer-to-peer networks," in *Advances in Spatial and Temporal Databases*, pp. 346-363, 2005.

[17] Y. Yu, R. Govindan, and D. Estrin, “Geographical and energy aware routing: A recursive data dissemination protocol for wireless sensor networks," 2001. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.21.8533&rep=rep1&type=pdf

[18] W. Horn, “Some simple scheduling algorithms," *Naval Research Logistics Quarterly*, **21**, pp 177-185, 1974.

[19] K. Xu, M. Gerla, and S. Bae, “How effective is the IEEE 802.11 rts/cts handshake in ad hoc networks?," *Global Telecommunications Conference, 2002, GLOBECOM '02. IEEE*, Vol. 1, pp. 72-76, 2002.

[20] R. Baardman, “Cumulus humilis: Wireless mesh networking for gliders," Master's thesis, University of Twente, Enschede, Overijssel, The Netherlands, 2010.

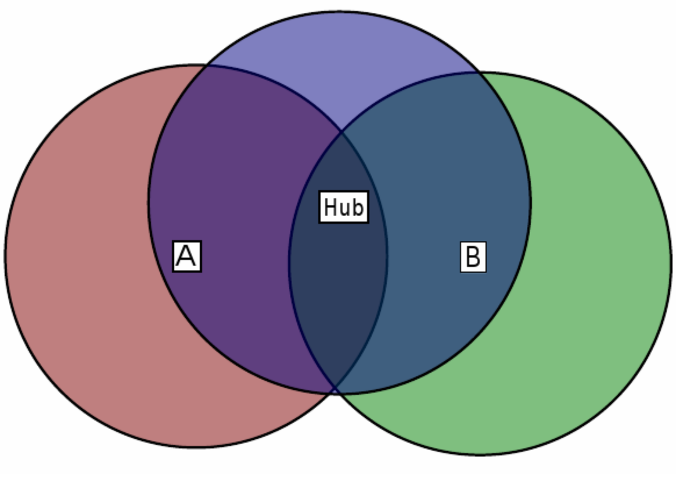

Figure 1 The hidden terminal problem.

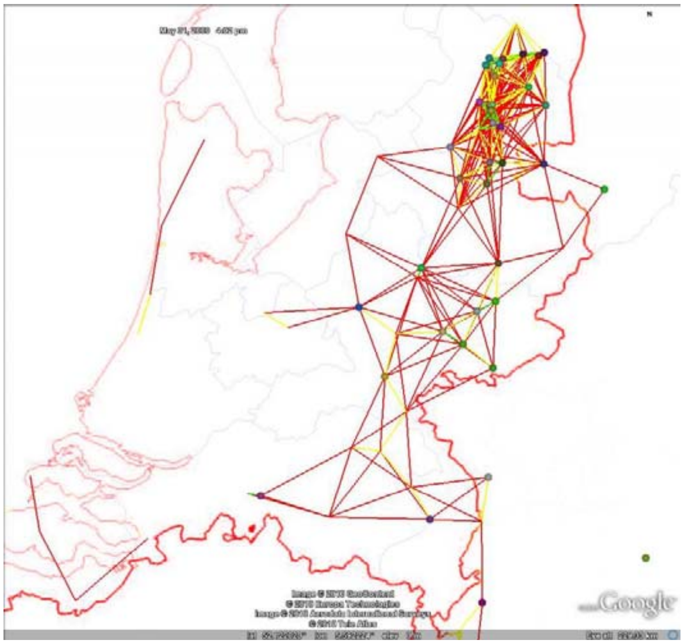

Figure 2 Flight analysis visualized in Google Earth. The 31th of May 2009, a day during the Dutch National Championships, is shown. Airfields are turned on. Communication is possible from the north to the south of The Netherlands.

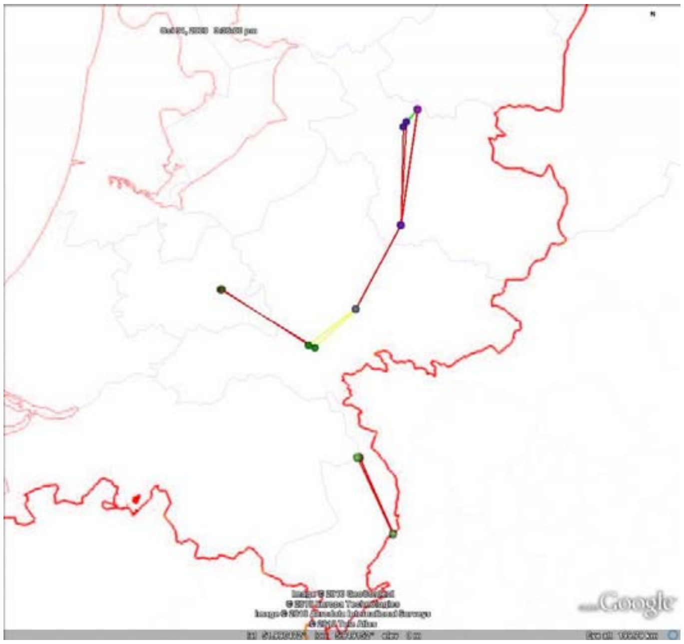

Figure 3 Flight analysis visualized in Google Earth. The 31th of October 2009, a day with not so good cross-country conditions, is shown. The network during this day is almost always partitioned.

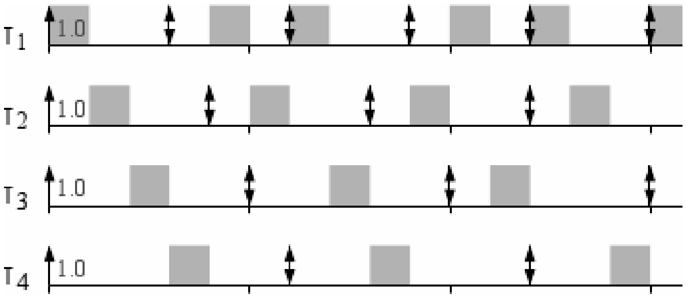

Figure 4 EDF scheduling visualized for 4 packets. Arrows pointing upwards indicate release times, arrows pointing downwards indicate deadlines. Grey blocks indicate the transmission of the packet.

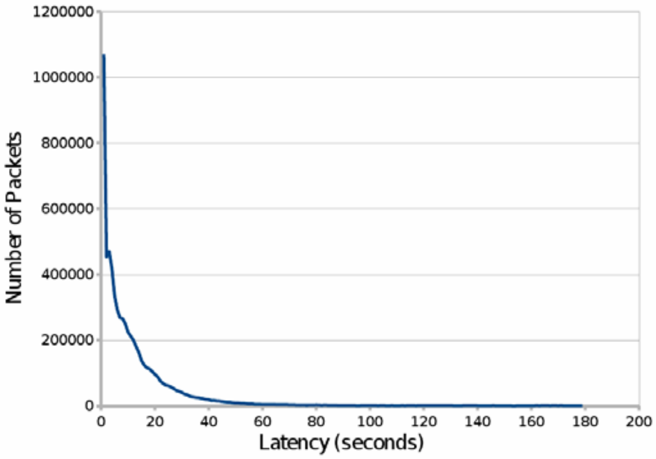

Figure 5 Latency for positional information during a day with good connectivity.

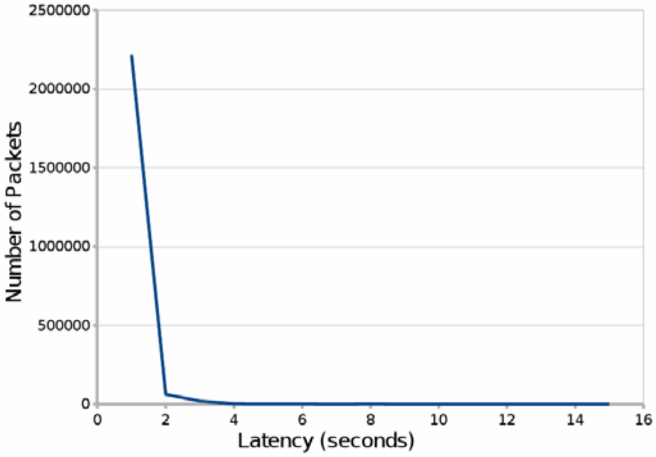

Figure 6 Latency for measurement information during a day with good connectivity.

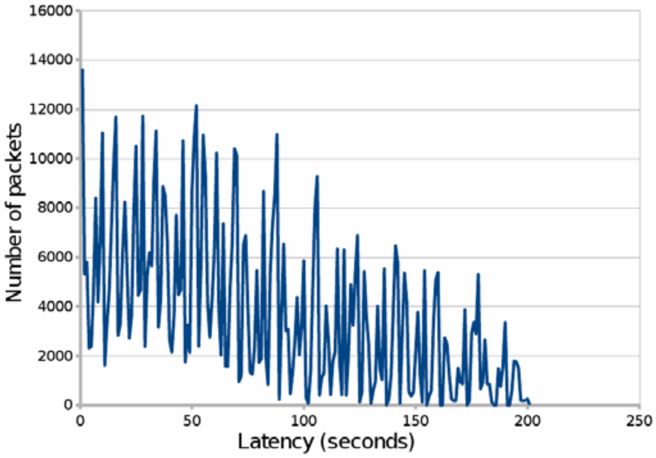

Figure 7 Latency for weather information during a day with good connectivity.

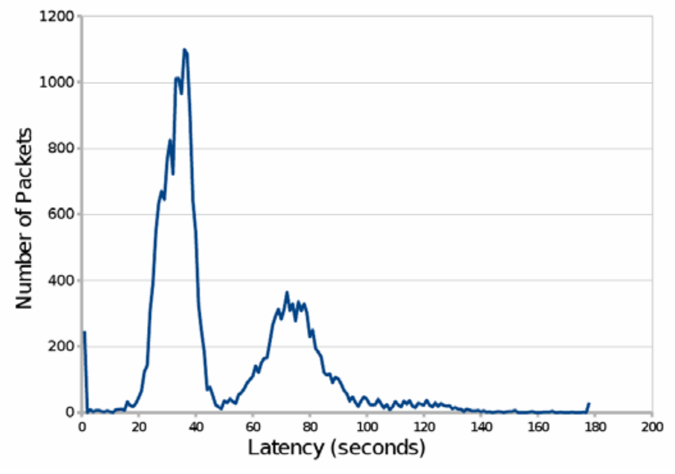

Figure 8 Latency for positional information during a day with sparse connectivity.

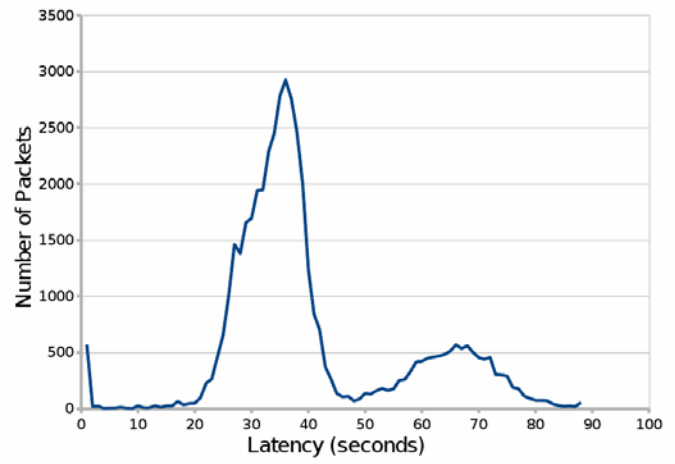

Figure 9 Latency for measurement information during a day with sparse connectivity.

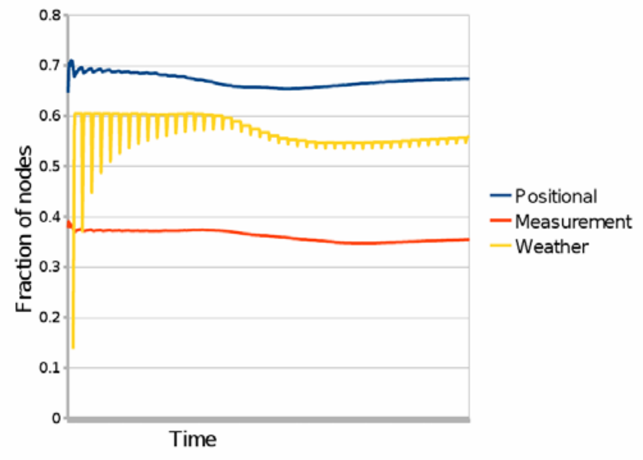

Figure 10 Packet arrival rate plotted over time for scenario 1, a day with good connectivity.

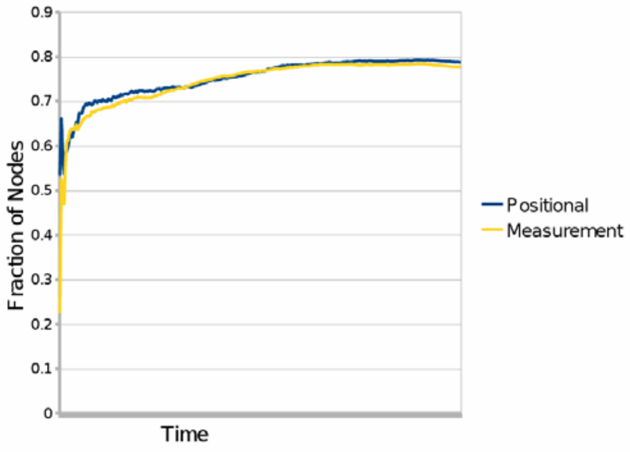

Figure 11 Packet arrival rate plotted over time for scenario 2, a day with sparse connectivity.

Scenario	Successful transmissions	Successful receivers	Packet-loss
Scenario 1	**21227286**	**24,6301128**	**36,38363%**
Scenario 2	**297314**	**6,6310205**	**13,24223%**

Figure 12 A summary of results from the two scenarios simulated.