

Towards 100% Reliability in Wireless Monitoring Networks

Lance Doherty
Dust Networks
30695 Huntwood Avenue
Hayward, CA 94544, USA
510-400-2953

ldoherty@dustnetworks.com

Dana A. Teasdale
Dust Networks
30695 Huntwood Avenue
Hayward, CA 94544, USA
510-400-2900

dteasdale@dustnetworks.com

ABSTRACT

High end-to-end reliability is a quality demanded by those with critical monitoring and actuation requirements. To date, Wireless Sensor Network (WSN) users have often accepted sub-optimal reliability as being intrinsic to wireless technologies. We describe a centralized monitoring TDMA network with policies chosen to maximize the number of received packets while maintaining low power characteristics. The methods for detecting and diagnosing packet loss are presented along with expected bounds on their relative impacts. This diagnosis allows for a cataloguing of all the known loss mechanisms and for the analysis of loss in a 50-node network running at 99.99% steady-state end-to-end reliability.

Categories and Subject Descriptors

C.3 [Special-Purpose and Application-Based Systems]:

Realtime and embedded systems, microprocessor/microcomputer applications, signal processing systems.

General Terms

Algorithms, Performance, Design, Reliability.

Keywords

Wireless mesh networks, Performance analysis, Experimental evaluation, Case studies.

1. INTRODUCTION

A *wireless monitoring network* is a WSN where data is periodically generated at all network nodes and collected through multi-hop transmissions at a single node called the manager. Knowing the number of nodes and the generation rate allows for the expected number of packets to be calculated: let *reliability* be a unit-less quantity describing the number of unique packets received at the manager divided by the number of unique packets that should be received. Academic research in WSNs is directed at providing optimal solutions: minimum energy cost, lowest latency, or highest bandwidth. While these elements are important to industrial and commercial WSN users, their main concern is often that packets generated in the network should be reliably received at a centralized data repository and available for analysis. The actual stress on the

network is typically not that high: an example would be a 50-node network with 10-byte data payloads generated at each node every minute with a 76.8 kbps radio transceiver. This network is using less than 0.1% of the raw available bandwidth to the manager. It is not surprising that many researchers shy away from studying such trivial requirements. It is, however, surprisingly difficult to meet strict reliability requirements even in generously provisioned scenarios. The time-synchronized protocol presented delivers a high-reliability wireless network solution that also enables low-power routing and high available bandwidth. The motivating goal is not necessarily to attain perfect reliability, but rather to understand the reason for every lost packet and be able to bound each loss mechanism.

An installation of a WSN in a historically wired domain [5] identifies low reliability as one of the main risks of adopting the technology. The metric of reliability is often discussed in systems papers dealing with WSN deployments, but is rarely the main focus. Projects such as ExScal [2] are continually extending the limits of WSN size and overall performance. This particular work cites 86% end-to-end reliability and hints at the lack of published data on the faults and variability that plague reliable operation. Another self-described “industrial sensor network” project [7] comments on the importance of reliability but indicates that most nodes in the experiment report more than 80% of their data suggesting that overall reliability is nowhere near the 100% range. It is clear that these authors recognize the importance of reliability in commercial applications, but equally clear that loss of individual packets is not of paramount concern. A revealing plot in [3] shows the span of published networks in the node number/lifetime space but makes no direct comment on the reliability of these networks. Up to now, it has been sufficient that a network remain operational for a given time span, not that the network necessarily is capable of reliable packet delivery for this time. The implication of the work in [4], which seeks to use a network to detect “rare, random and ephemeral events,” is that the network must be ready to sense over vanishingly small temporal and spatial regions. The continuation of this principle suggests that a single packet generated as the result of such a brief but important event should be held to the same high standards.

Recent new platforms and systems [6],[8] have focused on redesigning WSNs from first principles but have done so without the explicit goal of providing near-perfect reliability. A survey of the broader area of wireless mesh networks in general [1] mentions only briefly the subject of reliability, suggesting that even outside the realm of sensor networks, wireless multi-hop has not been focused on this goal.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PE-WASUN'06, October 6, 2006, Torremolinos, Malaga, Spain.

Copyright 2006 ACM 1-59593-487-1/06/0010...\$5.00.

2. THE WIRELESS NETWORK

Wireless channels are notoriously variable over time and space. To overcome this challenge, we use several techniques and policies to maximize reliability:

1. TDMA scheduling to ensure adequate bandwidth
2. Mesh topology to provide redundant routes
3. Required ACK prior to deleting a message
4. Frequency hopping to avoid blocked channels
5. Topology optimization to choose better routes
6. High-power mode to recover disconnected nodes
7. Consistent use of reliability diagnostics

The network consists of wireless nodes and repeating directional TDMA communication slots between select node pairs. Each such node pair defines a *path* and can consist of one or more slots in one TDMA superframe. Every transmission and reception is scheduled so no collisions occur from network traffic. This selection was made with low power in mind but it also allows for better reliability in congestion than contention-based protocols. Each node is assigned a sufficient number of slots in the TDMA schedule to satisfy the maximum bandwidth demand anticipated for itself and descendants. The “monitoring network” assumption is required here: it is difficult to adequately provision a network with unknown traffic requirements. A TDMA network requires tight time synchronization which poses its own set of challenges.

By constructing the network such that the manager has no parents and such that no loops are allowed in the digraph generated by the collection of paths, traffic flows upwards to the manager. With each node seeking two parents, the resulting mesh structure minimizes lost packets due to single path failures. Manager-node control communication is accomplished through broadcast flooding along the same paths in the opposite direction. Details of the TDMA protocol are not critical to this discussion; what is important is that sufficient bandwidth be scheduled to minimize the effects of unexpected congestion or interference from external sources. The high-reliability networks are over-provisioned in terms of bandwidth at the sole cost of additional idle listen slots. This allows queues to remain empty for most of the network operation. On the transmit side, the radio is activated only if there are packets waiting to be sent, so there is no energy cost for idle slots. However, the receiver must always listen for a brief period at the beginning of a slot in case of a transmission so the more bandwidth assigned, the more energy is spent. Mesh networks may pay upfront latency and energy penalties due to some traffic traveling along sub-optimal routes to the manager, but these costs ensure that few connectivity losses ensue from the inevitable path stability variations.

Our networking paradigm is that once a packet is generated it should never be discarded, only passed node-by-node topologically closer to the manager. The manager is a central data repository to which all packets are routed and is often connected to a user-accessible database. Multi-hop propagation is accomplished through a series of handshakes: packets are not deleted from a child until an ACK is received from the parent. This requires that a reasonably-sized packet queue is available on each node. When the queue fills up, the node can no longer store newly generated local data and begins to lose these packets. Full queues do not result in external packets being lost as the node instead NACKs messages from its children in this case.

To overcome local narrow-band interference that could adversely impact reliability, our networks spread the communication spatially and temporally over multiple channels. Transmissions hop over 50 channels in the 900MHz band and 16 channels in the 2.4GHz band. If there is interference on one or some of the channels, transmissions can still succeed on a subsequent attempt on a different channel. Again, the over-provisioning of bandwidth allows for networks to maintain high reliability even with blocked channels. Frequency hopping contributes to longer join times as new nodes must scan over several channels to locate their peers, but long-term reliability in dynamic radio environments favors multiple-channel strategies.

The network architecture lends itself to path optimization during network formation and operation. The manager is continuously searching for ways to redirect the network graph to ensure that reliability is maximized. Scoring for optimization is based on path quality (either empirically measured or RSSI), the number of hops to the manager and the lifetime of the parent. This allows us to continually ensure that the network is using the best available routes for data.

The network can operate in either low-power or high-power mode. During network formation and following loss of connectivity, the network automatically switches to high-power mode to facilitate and expedite node joining. Time spent by nodes outside the network is time that packets are lost; to minimize lost packets, more energy is spent to recover the missing nodes as quickly as possible.

Finally, as will be described in detail in the following section, we consistently employ several independent diagnostic methods to track lost packets. As reliability levels increase, lost packets become more difficult to track and the especially rare events cannot be tracked during short experiments. As such, using the diagnostic tools on *all* operating networks gives the most opportunities to discover all loss mechanisms.

3. IDENTIFYING LOST PACKETS

Lost packets are identified in our network by comparing two independent packet counts: periodic diagnostic packets from each node and the list of unique packets received by the manager. The diagnostic packet informs the manager both of the number of data packets originating at the node during the most recent collection interval and the number of locally-generated packets that were dropped due to full message buffers. The goal is to be able to consistently track the number of lost packets and the cause for each loss. In addition to these methods, we have also used direct wired and wireless node queries, network simulation, and node emulation to determine the origin of packet losses. Based on the data available during standard network operation, we can sort any lost packets into the following categories:

1. Congestion
2. Device resets
3. Device failures
4. CRC corruption
5. Accounting errors
6. Unknown

Congestion losses are reported directly by the nodes in the diagnostic packets and are hence easiest to track. However, continued congestion results in diagnostic packets being aborted to

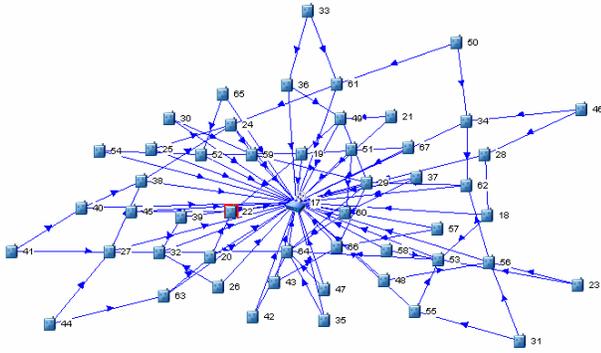


Figure 1. Network graph showing child-parent relationships in the test deployment.

discourage further losses, so packets reported as lost by the manager for node n during an interval when the diagnostic packet for node n does not arrive are assumed to be congestion losses. Node n can only lose its own messages due to congestion; nodes NACK children when their message buffers are full. This does not change the notion that every packet entering the network is kept safe, it rather shows how some expected packets are kept from ever entering the network.

A node resets its software when it loses connectivity to all parents. This results in the node losing all packets in its queue which could include packets from other nodes. A reset event is detected by the manager both through alarm packets generated by the node's parents and children and further through a request by the node to rejoin the network. Packets lost during the interval when a node is known to have been out of the network are assumed to be due to this resetting node regardless of their initial origin. Similarly, a device failure requiring some maintenance (such as battery replacement) can initially result in losses from several sources but only packets from the failed node are lost as time progresses.

A 16-bit CRC is appended to all messages to ensure link-level integrity. However, with longer packets, this CRC is not unique and certain error combinations can result in the error-ridden packet having the same CRC as the error-free packet. Another end-to-end encryption MIC is used, however, so the error is detected (but not correctable) at the manager upon decryption.

Several iterations of accounting bugs have been found and fixed and there is no guarantee that none remain. We assume that a packet is missing until it is reported by the manager, so all errors are false positive losses. Still, it is impossible to distinguish an accounting error from a packet otherwise lost so the worst must always be assumed. The most trouble comes from sequences of packets that are massively out of order and span across different fifteen minute intervals. For this study, an independent trace of all the incoming messages was kept to determine if any accounting errors were present in the automated analysis.

One year ago this list of loss mechanisms was much different. Some problems were solved while others were discovered. Pushing to progressively higher reliability, network sizes and traffic levels, more obscure and surreptitious losses will be revealed. Elements that factor at the 10^{-6} level are invisible when 10^{-3} level losses are prevalent. Again, this underscores the need to be able to account for every lost packet of known loss type. As will be seen in the next section, the current dominant form of packet loss in the test network

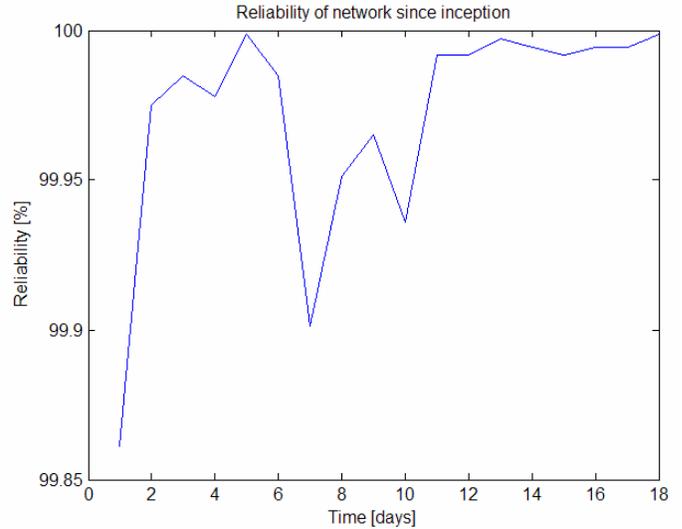


Figure 2. Daily reliability over the lifetime of the network.

is through unidentified mechanisms. We speculate that this is due to logical interference with other unsynchronized networks in the same radio space.

4. CASE STUDY: A 50-NODE NETWORK

As an example of the level of reliability that is attainable, a 50-node monitoring network is deployed in and around our office buildings. Nodes are placed on two vertical levels inside our workspace and on top of the roof. The network was designed to operate at very high reliability while having low node-to-node channel stability; a 1-minute reporting rate for each node was chosen to minimize the chance of packets lost to congestion. This network generates 72,000 packets per day which are wirelessly transmitted through an average of 2.0 hops each. The node furthest from the manager averages 3.6 hops per packet.

The network is running a TDMA superframe of 199 slots (6.2 seconds). The network digraph is depicted in Figure 1. During the first day of operation the reliability was 99.86%, but a full week of at least 99.99% reliability was observed. Reliability data for 17 days without human intervention is shown in Figure 2.

With nodes on different vertical planes, the average path stability is a rather low 60%, meaning that four out of ten transmission attempts fail and are retried. The low path stability was chosen to encourage all possible classes of error to reveal themselves in the network. All nodes are identical – none are distinguished as being “router” nodes – and all run on the same AA batteries. The shortest node lifetime on a pair of AA batteries for the network is 13 months. The mean packet latency over the lifetime of the network is 3.9s. The network has operated during periods of rain and surface temperatures measured by the rooftop nodes have exceeded 70 degrees Celsius.

Table 1 summarizes the nature of all the lost packets up to the date of writing. The total number of lost packets is 367 and 1.2 million unique packets have been received.

Table 1. Total number of lost packets by type for the lifetime of the network.

Loss Mechanism	Number of Packets	Loss Rate
Node Reset	265	2×10^{-4}
Unknown	71	6×10^{-5}
Congestion	26	2×10^{-5}
CRC Failure	5	4×10^{-6}
Failed Device	0	0
Accounting	0	0

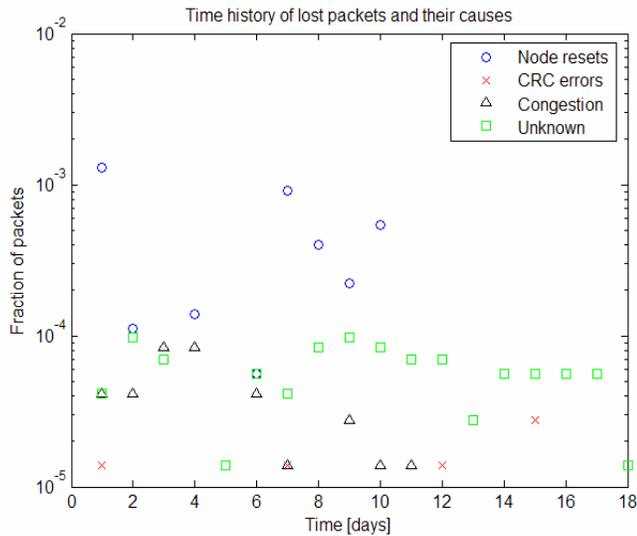


Figure 3. Mechanisms of daily packet loss over the lifetime of the network.

A further day-based analysis of the loss mechanisms is shown in Figure 3. In the most recent week of operation, no packets were lost either to node resets or to congestion. This is due to the optimization routines running in the network – the topology has been optimized to virtually eliminate these components of packet loss by ensuring that all the paths used are of high enough stability to ensure reliable operation. During all of this, individual paths are created and fail, but due to the mesh architecture, this does not necessarily contribute to packet loss.

Most of the node reset events occurred during network formation and optimization; the expectation is that these losses will be much less frequent during steady-state operation. The number of unknown and CRC failures should continue at their current rates and we expect the network to continue to report data at 99.99% reliability from this point to at least 13 months in the future.

5. CONCLUSIONS AND FUTURE WORK

The proposed TDMA-based WSN policies have allowed a low-power 50-node network to operate at a steady-state 99.99% level of

reliability without human intervention. This level of reliability was attained by a careful design of the entire system around the goal of delivering every packet to the centralized manager. Changes to the protocol could result in the network losing below 10^{-4} of its packets including the joining period by preserving queue contents during reset and increasing the queue length to eliminate congestion losses.

Going forward, the goal is to identify the sources and reduce the number of unknown losses. The magnitude of these errors is high enough that the CRC-type errors are relatively of little concern. Further discovery will be accomplished through broader use of wired connections to directly monitor individual node pairs. The network did not exhibit any node failures, but this level of reliability requires failures to either be predictable (through battery monitoring, for example) or extremely rare. At reliability levels of 99.9%, we can afford failure downtimes on the order of 1/10000. If each failure replacement takes a day, we need device failures to occur less frequently than once per 27 years.

While the test network consists only of 50 nodes, the protocol scales to larger networks and has similar per-node reliability. For monitoring networks, the reporting rate must be adjusted accordingly to ensure that the level of congestion does not increase, but there are no fundamental changes to the protocol. We have run networks of 250 nodes prior to a few bug fixes for periods in excess of one month with >99% reliability. Providing that nodes can stay connected at the same level, the rest of the operation is identical and should result in the same loss fraction.

6. REFERENCES

- [1] Akyildiz, I.F. and Wang, X. A Survey on Wireless Mesh Networks. *IEEE Communications Magazine*, vol. 43, no. 9, Sept. 2005.
- [2] Bapat, S., Kulathumani, V. and Arora, A. Analyzing the Yield of ExScal, a Large-Scale Wireless Sensor Network Experiment. *ICNP 05*, Nov. 2005.
- [3] Dutta, P., Grimmer, M., Arora, A., Bibyk, S. and Culler, D. Design of a Wireless Sensor Network Platform for Detecting Rare, Random and Ephemeral Events. *IPSN 05*, Apr. 2005
- [4] Dutta, P., Hui, J., Jeong, J., Kim, S., Sharp, C., Taneja, J., Tolle, G., Whitehouse, K. and Culler, D. Trio: Enabling Sustainable and Scalable Outdoor Wireless Sensor Network Deployments. *IPSN 06*, Apr. 2006.
- [5] Healy, W. M. Lessons Learned in Wireless Monitoring. *ASHRAE Journal*, Vol. 47 No. 10, October 2005, pp. 54-60.
- [6] Kling, R., Adler, R., Huang, J., Hummel, V., Nachman, L. Intel Mote: Sensor Network Technology for Industrial Applications, *IPSN 05 demo*, Apr. 2005.
- [7] Krishnamurthy, L., Adler, R., Buonadonna, P., Chhabra, J., Flanigan, M., Kushalnagar, N., Nachman, L., Yavis, M. Design and deployment of industrial sensor networks: Experiences from a semiconductor plant and the North Sea. *ACM SenSys 05*.
- [8] McIntire, D., Ho, K., Yip, B., Singh, A., Wu, W., Kaiser, W. J. The Low Power Energy Aware Processing (LEAP) Embedded Networked Sensor System. *IPSN 06*, Apr. 2006.