Seven Nines in 3 ms: Diversity Metrics in Wireless Factory Automation

Brian Kilberg, Craig Schindler, Arvind Sundararajan, Alex Yang, Kristofer Pister
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley
Berkeley, California, USA
Email: bkilberg@berkeley.edu

***Abstract*—Wireless sensor networks for factory automation and control will require strict latency and reliability requirements on the order of 1 millisecond end-to-end latency and 10^{-8} packet error rate. In order to meet these requirements, wireless sensor networks will require effective modes of network diversity that are also compatible with low-latency time scales. We developed a metric for quantifying diversity efficacy, and we evaluated the efficacy of temporal, hardware, spatial, and frequency diversity in cooperative IEEE 802.15.4 wireless networks. We demonstrated a wireless sensor network topology that achieved 99.99999% of reliability bounded by a worst case end-to-end latency of 3 milliseconds.**

I. Introduction

Ultra low power wireless sensor networks are deployed world-wide in industrial process automation, for both monitoring and relatively high latency control, using standards such as WirelessHART (IEC62591), and ISA100.11a (IEC62734), which use 2.4 GHz 802.15.4 radios running variants of the Time Synchronized Mesh Protocol TSMP [1]. These systems enable sensor data to be gathered reliably from industrial systems without the need for expensive and physically cumbersome wiring. The natural extension of these networks is their use for low latency wireless control in factory automation.

While ultra low-power wireless sensor networks have enabled high-reliability wireless control in process automation, current high-reliability systems do not deliver tight enough latency bounds to be suitable for wireless control in factory auatomation [2]. Previous work has demonstrated the use of a multi-hop 802.15.4e 6TiSCH network to wirelessly stabilize an inverted pendulum, but found the latency introduced by the wireless system was a significant challenge [3]. Wired industrial control systems generally have a round trip latency of 1 ms with reliability of 10^8 [2]. Wireless industrial control systems must eventually mirror these specifications to become viable.

Significant work has gone into creating reliable (>99.999% [4]) implementations of IEEE802.15.4 and many additional standards [5], [6], [7], [8] with end-to-end packet latency on the order of tens of milliseconds. Yang et al. demonstrated better than four nines reliability at under 40 ms latency using an unmodified OpenWSN stack [9]. Given that the default time slot length in these standards-based protocols is 10 ms, improvement below 10 ms is unlikely. In order to explore strategies for end-to-end packet latency of less than 10 ms we jettison the networking stack and build up our test setup starting with only the OpenWSN BSP interface, the 802.15.4 PHY layer and the lower MAC layer.

II. Network Diversity and Reliability

Redundant packet transmissions and receptions that are ideally independent produce network diversity, which prevents network performance degradation due to interference and multipath fading [10]. Network diversity is key for high-reliability wireless networks [2][11]. Spatial, frequency, temporal, and hardware diversity can all be used as diversity modes in a network system. Figure 1 illustrates these types of diversity. In low-latency networks, temporal diversity is not very useful due to coherence times being on the same time scale as the latency requirements [11].

The primary goal of network diversity is to create redundant data paths that fail independently from each other. For example, consider a wireless sensor network with one transmitting mote and one receiving mote that has a packet delivery ratio (PDR) of 0.9999, and thus a packet error rate (PER) of 10^{-4}. Then consider a second receiving mote with a PER of 10^{-4} added to the network. This network now has diversity because it has a redundant packet delivery path. However, the extent that this diversity improves the total packet error rate (PER_t) depends on the efficacy of the specific diversity mode used. The PER_t isn't simply 10^-8 because packet error events aren't necessarily independent of each other. The correlation between each packet delivery channel dictates the efficacy of the network's diversity mode, with higher correlation yielding lower diversity efficacy. Let us model the packet reception and error events for both motes as the following Bernoulli random variables:

$$X_1 = \begin{cases} 1 & \text{if Mote 1 Receive Packet} \\ 0 & \text{if Mote 1 Misses Packet} \end{cases}$$
$$X_2 = \begin{cases} 1 & \text{if Mote 2 Receive Packet} \\ 0 & \text{if Mote 2 Misses Packet} \end{cases} \quad (1)$$

These variables have the following probability mass functions:

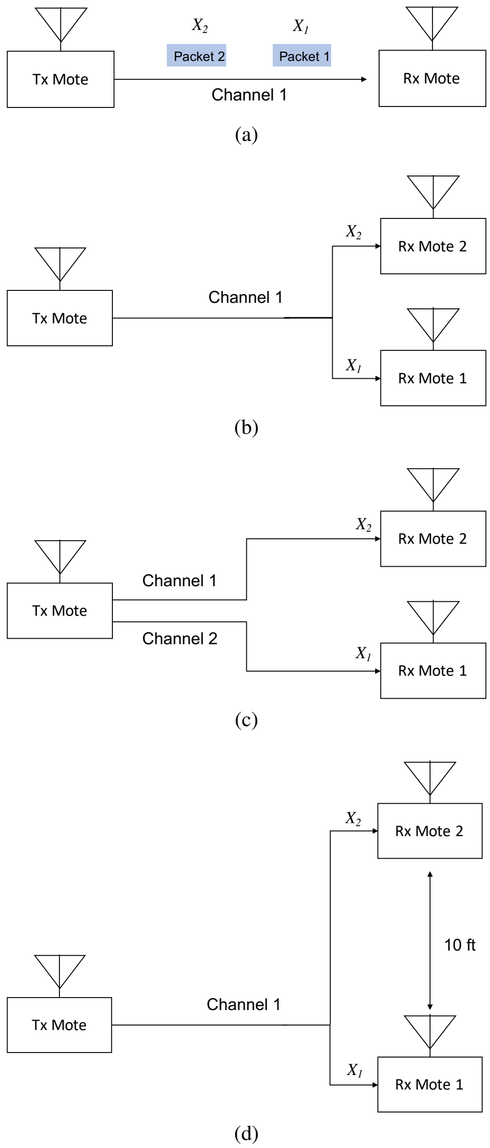

Fig. 1. Network diversity in wireless sensor networks. (a) temporal diversity where each packet is retransmitted multiple times in order to decrease the likelihood of a missed packet. (b) Hardware diversity where two receive motes are in the same location. (c) Frequency diversity where the transmitting mote broadcasts on two channels and each receiving mote listens on a different channel. (d) Spatial diversity where the two receiver motes are in different locations.

$$f_{X_1}(x_1) = \begin{cases} PDR_1 & \text{if } x_1 = 1 \\ PER_1 & \text{if } x_1 = 0 \end{cases}$$
$$f_{X_2}(x_2) = \begin{cases} PDR_2 & \text{if } x_2 = 1 \\ PER_2 & \text{if } x_2 = 0 \end{cases} \quad (2)$$

The total PER is defined below:

$$\begin{aligned} PER_t &= P(X_1 = 0 \ \cap \ X_2 = 0) \\ &= P(X_1 = 0|X_2 = 0)P(X_2 = 0) \\ &= P(X_2 = 0|X_1 = 0)P(X_1 = 0) \end{aligned} \quad (3)$$

The conditional probabilities $P(X_1 = 0|X_2 = 0)$ and $P(X_2 = 0|X_1 = 0)$ describe the probability of one mote missing a packet given that the other mote misses a packet and are useful metrics for evaluating the efficacy of network diversity. If the network's mode of diversity is highly effective, the random variables X_1 and X_2 would be independent, so $P(X_1 = 0|X_2 = 0) = P(X_1 = 0)$, yielding $P(X_1 = 0 \cap X_2 = 0) = 10^{-8}$. Conversely, if the network's diversity is completely ineffective, X_1 and X_2 will be perfectly correlated and $P(X_1 = 0|X_2 = 0) = 1$, yielding $P(X_1 = 0 \ \cap \ X_2 = 0) = 10^{-4}$. In this case, the second mote would provide no benefit to the reliability of the network. This example shows the importance of a diversity mode's efficacy for a network's reliability. With sufficiently effective diversity modes, only a few extra packet delivery paths are necessary to obtain high-reliability networks. We define our diversity efficacy metric as the average ratio of an independent packet failure to a conditional packet failure; this is also equal to the ratio of PER_t to the multiplication of the conditional packet failure probabilities:

$$\begin{aligned} D &= \left(\frac{1}{2}\right)\left(\frac{P(X_1 = 0)}{P(X_1 = 0|X_2 = 0)} + \frac{P(X_2 = 0)}{P(X_2 = 0|X_1 = 0)}\right) \\ &= \frac{P(X_1 = 0 \ \cap \ X_2 = 0)}{P(X_1 = 0|X_2 = 0) \ P(X_2 = 0|X_1 = 0)} \end{aligned} \quad (4)$$

A diversity score of 1 implies X_1 and X_2 are completely independent. A diversity score of less than 1 implies X_1 and X_2 are correlated, and a diversity score of more than 1 implies that X_1 and X_2 are anti-correlated.

$$PER_t = \frac{PER_1 PER_2}{D} \quad (5)$$

III. Efficacy of Network Diversity Modes

We evaluated the efficacy of temporal, spatial, frequency, and hardware diversity on an 802.15.4-based sensor network. The topology of the network had n-to-k transmitters to receivers with various forms of diversity implemented. OpenMotes [12] based on the Texas Instruments CC2538 wireless SoC were used for the network hardware. Multiple network topologies were tested by sending 10^6-10^7 packets through each network and capturing the arrival and latency of each packet. The test harness used for data collection was modified from the system used in [9], which consisted of a computer-controlled logic analyzer that triggers the Tx mote to transmit a packet and waits for the Rx motes to toggle a GPIO pin when they receive the packet. The transmitting mote was programmed to send packets whenever the test harness toggled its GPIO pin. Each packet contained a 8-byte payload in addition to preamble and CRC. Inside the packet was a 32 bit passphrase unique to the network under test. Receiving motes only toggled their packet-received GPIO once the packet's CRC and passphrase was verified. Histograms of packet latencies were generate from the data. The receiving motes were located approximately 4 m from the transmitting mote. Currently, there are a small number of packets that measured as arriving earlier than they are physically able to. We count these packets as either missed, or at their next reception time, if they were received again. Future work will focus on isolating and eliminating these spurious packets, which will be necessary to demonstrate 10^{-8} and 10^{-9} reliable networks.

As stated previously, temporal diversity at the millisecond time-scale that low-latency networks require was expected to be ineffective. Experiments showed that temporal diversity only produced a diversity metric of $D = 1.5419 * 10^{-4}$ (Figure 2). However, retransmitting packets three times was sufficient to decrease the intrinsic PER by an order of magnitude from $1.3 * 10^{-4}$ to $1.22 * 10^{-5}$. temporal diversity can still be useful for improving reliability, but its fundamental nature requires a temporal-latency trade-off, and is not a very efficient form of diversity.

The hardware diversity network topology consisted of one mote transmitting to two receiving motes in the same location, on the same channel. Each receiving mote should experience approximately the same network conditions. Hardware diversity is intended to alleviate any hardware-originated packet loss or high-latency events that could be caused by non-deterministic interrupt-handling or hardware failure. For this diversity mode, $D = 5.521 * 10^{-3}$, a slight improvement over temporal diversity (Figure 3).

The efficacy of frequency diversity depends on the magnitude of frequency difference between each channel used for communication. Frequency-dependent interference from sources like WiFi can span multiple 802.15.4 channels, which would increase the correlation of packet errors on adjacent channels. The first experiment evaluated frequency diversity where the communication channels only differed by one channel, which is 5 MHz for 802.15.4 channels. The transmitting mote would send each packet twice, once on channel 12, and then once on channel 13. 10^6 packets were sent from one transmitting mote to two listening motes, mote 1 was listening on channel 12, and mote 2 was listening on channel 18. The diversity efficacy of 5 MHz channel hop frequency diversity was $D = 8.926 * 10^{-3}$. Increasing the channel hop to six channels, or 30 MHz, improved the diversity metric to $D = 6.329$ (Figure 4). Since $D > 1$ in this case, it is actually less likely for a for a packet to be missed if the packet on the other mote is missed. This suggests that in this case, packet error events are actually anti-correlated, where a packet error on one channel makes it less likely for a packet to be missed on the other channel.

Spatial diversity should help eliminate the effect of local interference and multipath effects. In the spatial diversity experiment, each receiving mote was placed at a distance of 4m from the transmitting mote, with an angle of 90 degrees between them, with reference to the transmitting mote. Only one frequency channel was used for this experiment. The resulting diversity efficacy of this mode was 0.0287 (Figure 5).

A. *Combining Multiple Network Diversity Modes*

Combining multiple modes of networking diversity could yield high reliability networks without sacrificing latency or increasing network complexity greatly. In this experiment, a combination of temporal, mote, spatial, and frequency diversity were used to achieve reliability of $< 10^{-7}$ with a latency bound of 3 ms. 10^7 packets were sent during this experiment.

B. *Evaluation of Network Diversity Modes*

Table I shows a summary of the efficacy of different flavors of network diversity. Frequency diversity, when the channels are 30 MHz apart, was the most effective form of network diversity. Frequency and temporal diversity both require packet retransmissions when using a single transmit mote, which leads to an increase in latency bound. Spatial and hardware diversity only require more receiving motes, and do not affect the potential upper bound of latency. It is interesting to note that mote and temporal diversity have about the same effect on network reliability. This suggests that at the short time intervals low latency networks operate within, resending packets over time doesn't provide any benefit to increasing the number of motes listening on the channel. While frequency diversity was the most effective form of diversity, the other modes were still useful when used simultaneously. The 3 ms latency-bounded 10^-7 reliability network topology is a testament to the utility of combining multiple modes of diversity.

IV. Scalability and Resistance to Interference

Industrial automation and control settings will require large scale networks with numerous transmitting and receiving motes. Therefore, low-latency high-reliability industrial wireless networks must be able to scale effectively as the number of active motes increases. A network topology such as a one-to-two mote network utilizing two-channel frequency diversity could easily scale to eight independent networks, provided that the active frequencies for each network were prudently chosen. However, this topology could encounter scalability issues once all sixteen 802.15.4 channels are used. At this point, the channel occupancy of each channel will increase and the likelihood of packet collisions and interference will increase, leading to degraded network reliability. We evaluated the performance of a two channel one-to-two mote network as the adversarial channel occupancy of its channels increased. A second transmitting mote was used to randomly transmit 10-byte payload packets on random channels in the vicinity of the one-to-two mote system. Adversarial channel occupancy was defined as the percentage of time the interfering mote spent transmitting on each channel, in expectation. The one-to-two mote frequency diversity network occupied approximately 4% of each of the two channels it was using. Figure 7 shows the effect that channel occupancy of the interfering mote had on the diversity efficacy and total PER on the one-to-two mote network. Diversity efficacy (D) remained on the order of 1 as adversarial channel occupancy increased from 0% to 3.3%, while total PER increased from 1 ppm to

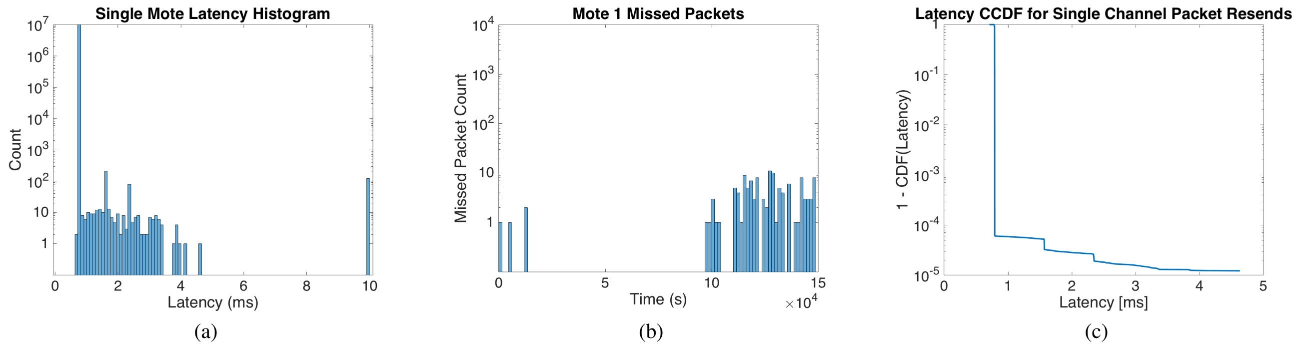

Fig. 2. Experimental data from the temporal diversity experiment. 10^7 packets were sent, were each packet was retransmitted three times. (a.) This latency histogram shows the distribution of latencies observed during the experiment. Missed packets appear as latencies at 10 ms. (b.) Time series histogram of missed packets. Most packets were missed in the last third of the experiment. (c) Complimentary cumulative distribution function shows the fraction of packets that arrive after a specific latency on the graph.

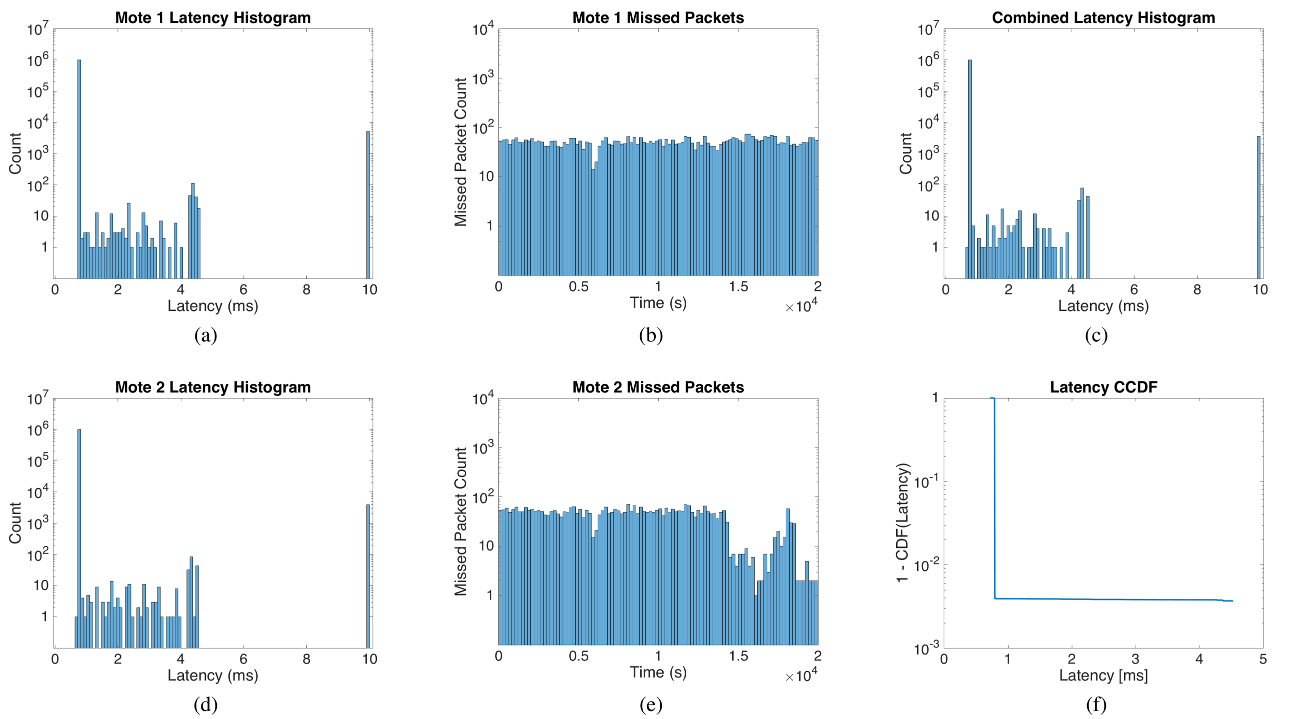

Fig. 3. Experimental data from hardware diversity experiment. 10^6 packets were sent (a,d) Histograms illustrating the latency distributions for each receiving mote. Missed packets are denoted by the 10 ms bin. (b,e) Time series histograms of PER over time. (c) Histogram showing the combined (the latency at which the packet was first received by mote 1 or mote 2) latency distribution from the two motes.

TABLE I
SUMMARY OF NETWORK DIVERSITY EFFICACY

Network Diversity Mode	PER_1	PER_2	PER_t	$P(X_2 = 0 \mid X_1 = 0)$	D
Temporal	$1.3 * 10^{-4}$	$1.3 * 10^{4}$	$1.22 * 10^{-5}$	0.8431	$1.5419 * 10^{-4}$
Hardware	$5.2 * 10^{-3}$	$3.92 * 10^{-3}$	$3.66 * 10^{-3}$	0.81	$5.521 * 10^{-3}$
Frequency(Channel 12 and 13)	$9.05 * 10^{-4}$	$1.657 * 10^{-3}$	$1.68 * 10^{-4}$	0.1856	$8.926 * 10^{-3}$
Spatial	$2.887 * 10^{-3}$	$5.087 * 10^{-3}$	$1.1 * 10^{-3}$	0.38205	0.0287
Frequency (Channel 12 and 18)	$2.549e * 10^{-3}$	$2.483 * 10^{-3}$	10^{-6}	$3.974 * 10^{-3}$	6.329
Temporal, Spatial, and Frequency	$1.22 * 10^{-5}$	$1.008 * 10^{-4}$	$< 10^{-7}$	$< 4.5944 * 10^{-3}$	N/A

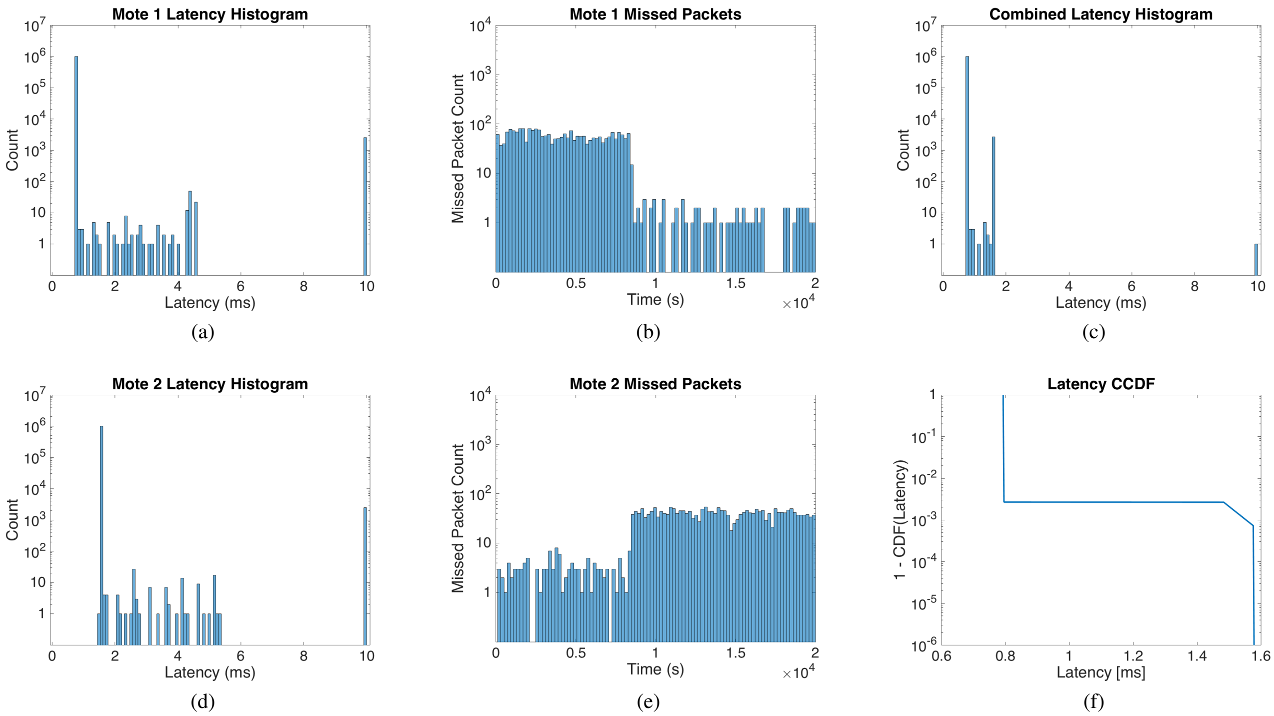

Fig. 4. Experimental data from frequency diversity experiment with a 30 MHz channel hop. (a,d) Histograms illustrating the latency distributions for each receiving mote this experiment. Missed packets are denoted by the 10 ms bin (b,e) Time series histogram. Approximately halfway through the experiment, an unknown source of interference switched channels, which can be seen in the PER of mote 2 increase while the PER of mote 1 decreasing. (c) Histogram showing the combined (the latency at which the packet was first received by mote 1 or mote 2) latency distribution from the two motes. (f) CCDF of latency

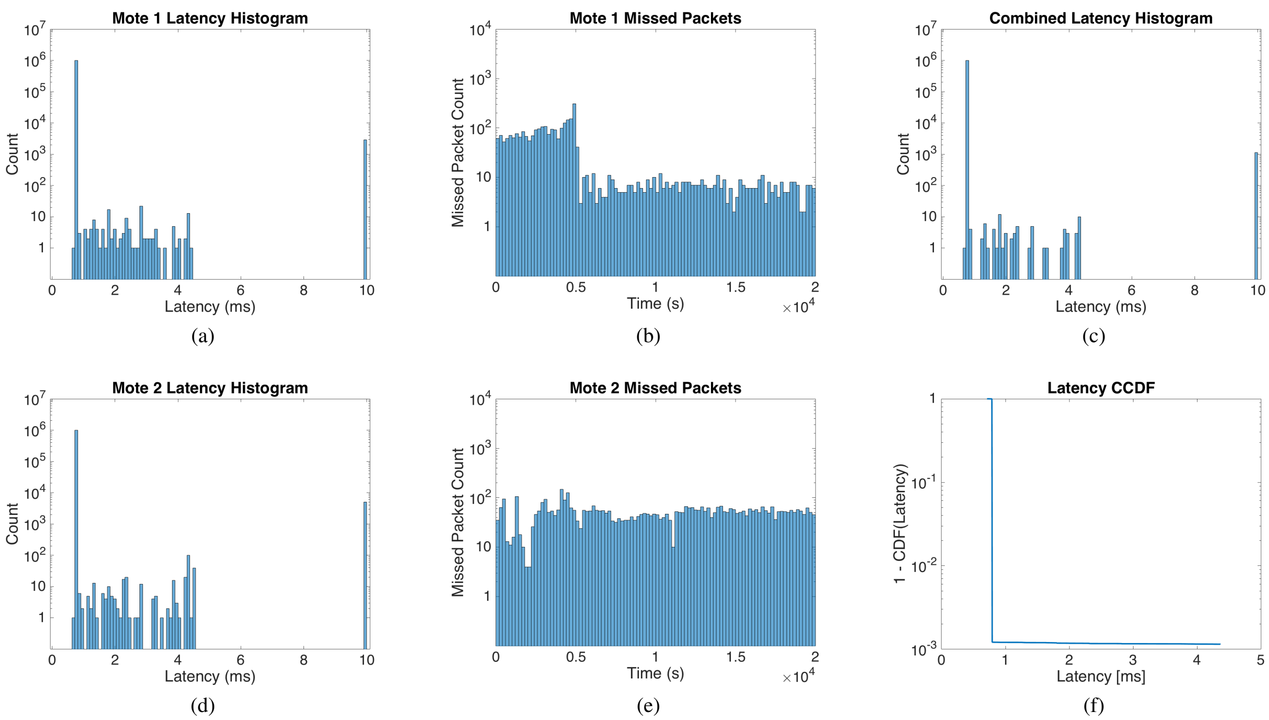

Fig. 5. Experimental data from spatial diversity experiment. 10^6 packets were sent from one transmitting mote to two listening motes, with each mote listening on the same channel. (a,d) Histograms illustrating the latency distributions for each receiving mote this experiment. (b,e) Time series histograms showing missed packets over time. (c) Histogram showing the combined latency distribution from the two motes. (f) CCDF of latency

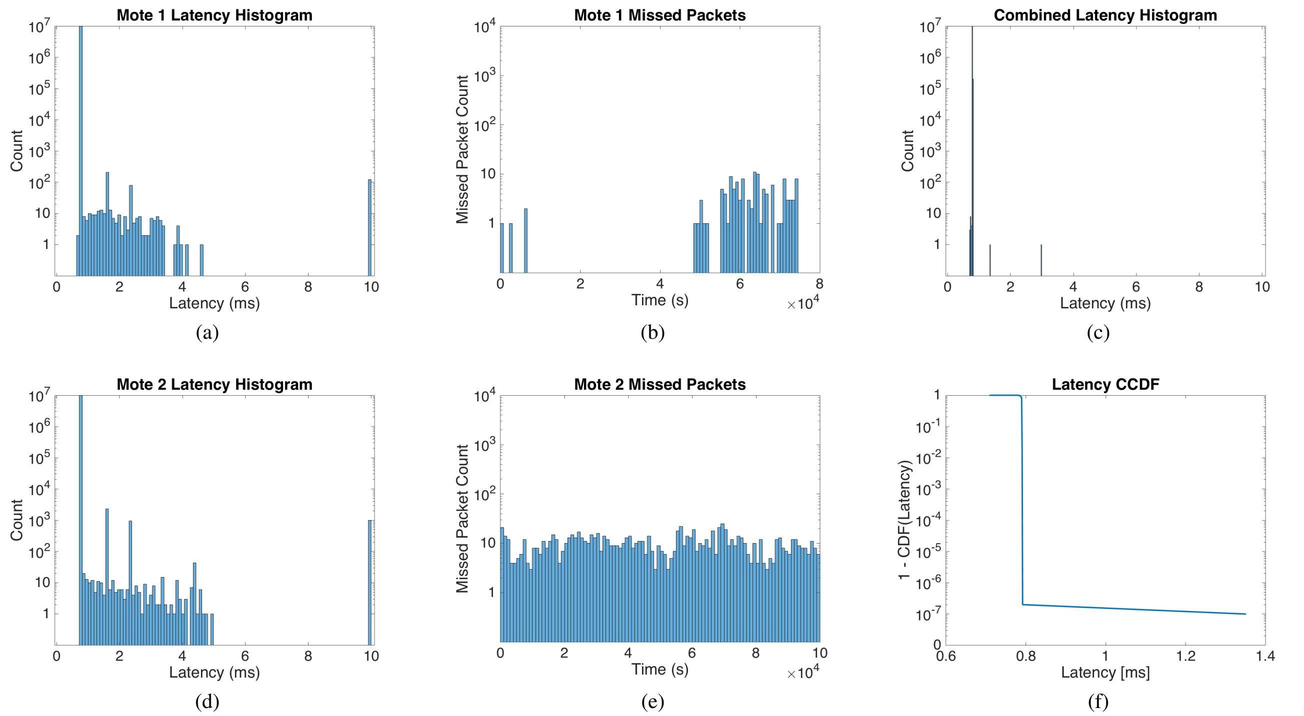

Fig. 6. Experimental data from multi-mode diversity experiment. In this experiment, two transmitting motes, each on different channels, were sending packets to two receiving motes, each listening on its own channel. 10^7 packets were sent, and each packet was retransmitted three times. (a,d) Histograms illustrating the latency distributions for each receiving mote this experiment. Missed packets are denoted by the 10 ms bin (b,e) Time series histograms showing how the packet error rate of each mote evolves through time. (c) Histogram showing the combined PER from the two motes. (f) CCDF of latency

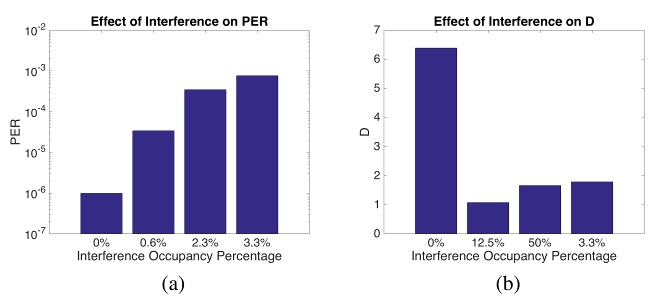

Fig. 7. The effect of interference on total PER and diversity efficacy of a one-to-two mote network with frequency diversity. The transmitting mote first transmits on channel 14 and the channel 20. The interference transmitter is sending packets with 10-byte payloads at random times on random channels. The occupancy percentage is the amount of time the interfering transmitter is using a given channel.

782 ppm. Since the diversity efficacy remained approximately equivalent, the total PER increase was caused by an increase in the individual PERs for each mote, rather than an increase in correlation between packet error events. While frequency diversity remains effective in this situation, care must be taken to avoid degraded network performance due to crowded channels. Decreasing channel occupancy and thus throughput is one method of accommodating multiple networks sharing channels.

V. Conclusion

Latency-bounded, high-reliability wireless sensor networks depend on effective modes of network diversity. Diversity modes with high efficacy enable increased reliability without sacrificing latency performance and unduly increasing network complexity. In 802.15.4 networks running on commercial-off-the-shelf hardware, frequency diversity was the most effective mode of diversity. A one-to-two mote network with frequency diversity demonstrated 10^{-6} packet error rate with a 2 ms latency bound. While showing less efficacy, other forms of diversity such as temporal, hardware, and spatial diversity could still be effective if used in conjunction with frequency diversity. By combining temporal, hardware, frequency, and spatial diversity, we demonstrated a 3 ms latency-bounded network with 10^{-7} reliability. These network topologies can scale to more than 8 simultaneous networks if channel occupancy and throughput are adjusted accordingly. Future work will focus on using latency-bounded high-reliability networks to demonstrate reliable wireless control.

Acknowledgment

The authors thank Azbil Corporation and the Berkeley Sensor & Actuator Center for supporting this research.

References

[1] K. Pister and L. Doherty, "Tsmp: Time synchronized mesh protocol," *IASTED Distributed Sensor Networks*, pp. 391–398, 2008.

[2] V. N. Swamy, S. Suri, P. Rigge, M. Weiner, G. Ranade, A. Sahai, and B. Nikolic, "Cooperative communication for high-reliability low-latency wireless control," *IEEE International Conference on Communications*, vol. 2015-September, pp. 4380–4386, 2015.

[3] C. B. Schindler, T. Watteyne, X. Vilajosana, and K. S. Pister, "Implementation and characterization of a multi-hop 6TiSCH network for experimental feedback control of an inverted pendulum," *2017 15th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks, WiOpt 2017*, 2017.

[4] T. Watteyne, J. Weiss, L. Doherty, and J. Simon, "Industrial ieee802. 15.4 e networks: Performance and trade-offs," in *Communications (ICC), 2015 IEEE International Conference on*. IEEE, 2015, pp. 604–609.

[5] T. Watteyne, X. Vilajosana, B. Kerkez, F. Chraim, K. Weekly, Q. Wang, S. Glaser, and K. Pister, "Openwsn: a standards-based low-power wireless development environment," *Transactions on Emerging Telecommunications Technologies*, vol. 23, no. 5, pp. 480–493, 2012.

[6] D. Dujovne, T. Watteyne, X. Vilajosana, and P. Thubert, "6tisch: deterministic ip-enabled industrial internet (of things)," *IEEE Communications Magazine*, vol. 52, no. 12, pp. 36–41, 2014.

[7] T. Winter, P. Thubert, and A. Brandt, "Rpl: Ipv6 routing protocol for low power and lossy networks. rfc 6650," Internet Draft, March, Tech. Rep., 2011.

[8] N. Kushalnagar, G. Montenegro, and C. Schumacher, "Ipv6 over low-power wireless personal area networks (6lowpans): overview, assumptions, problem statement, and goals," Internet Draft, March, Tech. Rep., 2007.

[9] A. Yang, A. Sundararajan, C. B. Schindler, and K. S. J. Pister, "Analysis of low latency tsch networks for physical event detection," in *2018 IEEE Wireless Communications and Networking Conference Workshops (WCNCW)*, 2018.

[10] J. N. Laneman and G. W. Wornell, "Distributed Space Time-Coded Protocols for Exploiting Cooperative Diversity in Wireless Networks," *IEEE Transactions on Information Theory*, vol. 49, no. 10, pp. 2415–2425, 2003.

[11] M. Weiner, M. Jorgovanovic, A. Sahai, and B. Nikolie, "Design of a low-latency, high-reliability wireless communication system for control applications," *2014 IEEE International Conference on Communications, ICC 2014*, pp. 3829–3835, 2014.

[12] X. Vilajosana, P. Tuset, T. Watteyne, and K. Pister, "Openmote: open-source prototyping platform for the industrial iot," in *International Conference on Ad Hoc Networks*. Springer, 2015, pp. 211–222.