Impact of Wireless Channel Temporal Variation on MAC Design for Body Area Networks

ATHANASSIOS BOULIS and YURIY TSELISHCHEV, NICTA, University of Sydney
LAVY LIBMAN, University of New South Wales
DAVID SMITH and LEIF HANLEN, NICTA, Australian National University

We investigate the impact of wireless channel temporal variations on the design of medium access control (MAC) protocols for body area networks (BANs). Our measurements-based channel model captures large and small time-scale signal correlations, giving an accurate picture of the signal variation, specifically, the deep fades which are the features that mostly affect the behavior of the MAC. We test the effect of the channel model on the performance of the 802.15.4 MAC both in contention access mode and TDMA access mode. We show that there are considerable differences in the performance of the MAC compared to simulations that do not model channel temporal variation. Furthermore, explaining the behavior of the MAC under a temporal varying channel, we can suggest specific design choices for the emerging BAN MAC standard.

Categories and Subject Descriptors: I.6 [**Computing Methodologies**]: Simulation and Modeling; C.4 [**Computer Systems Organization**]: Performance of Systems

General Terms: Design, Performance

Additional Key Words and Phrases: Body area networks, channel modeling, medium access control (MAC) protocols

ACM Reference Format:
Boulis, A., Tselishchev, Y., Libman, L., Smith, D., and Hanlen, L. 2012. Impact of wireless channel temporal variation on MAC design for body area networks. ACM Trans. Embed. Comput. Syst. 11, S2, Article 51 (August 2012), 18 pages.
DOI = 10.1145/2331147.2331161 http://doi.acm.org/10.1145/2331147.2331161

1. INTRODUCTION

Body area networks (BANs) are one of the core technologies to deliver the next generation of healthcare and, more generally, implement specialized sensor networks for monitoring human performance (e.g., in sports) and health. The characteristics of the new wireless channel, namely in and around the moving human body, introduce a radically different environment for telecommunications and networking than the more traditional open-field environments. These new characteristics, along with the drive for ultra energy efficiency, that is, sensor nodes wirelessly connected to an on-body gateway should have a lifetime of a few years yet be very compact and unobtrusive, create great interest for the development of new radios (physical layer design) as well as new medium access control (MAC) solutions. The commercial and scientific

Authors' addresses: A. Boulis and Y. Tselishchev, Networked System ATP Laboratory, NICTA, 13 Garden St., Eveleigh, NSW 2015, Australia; email: athanassios.boulis@nicta.com.au; L. Libman, School of Computer Science and Engineering, University of New South Wales, Sydney, NSW 2052, Australia; D. Smith and L. Hanlen, Networked Systems CRL Laboratory, NICTA, 7 London Circuit, Canberra, ACT, Australia.

DOI 10.1145/2331147.2331161 http://doi.acm.org/10.1145/2331147.2331161

importance of such new solutions is recognized by the formation of the IEEE task group for the creation of the new BAN standard, namely, IEEE 802.15.6.[1] NICTA is involved in the efforts to develop the standard, focusing on channel modeling and MAC design.

We have designed and implemented specialized hardware to enable us to take multiple simultaneous signal strength measurements around the human body at a high rate (200Hz) while the person is conducting every day activities. Analysis of the acquired data allowed us to create a model of the wireless channel around the body which takes into account both small-scale and large-scale fading statistics. This measurements-based model captures the salient characteristics of the channel more accurately compared to generic large-scale statistic wireless channel models.

Our end goal is to study how this more accurate model affects high-layer behavior and, consequently, can help us design better network solutions. For this purpose, we have implemented our BAN channel model into the open-source Castalia network-level simulator[2]. In this article, we show how more accurate modeling affects MAC design. More specifically, we show the impact that temporal variations, particularly deep fades, have on MAC performance. An existing low-power IEEE standard MAC protocol, namely, the IEEE 802.15.4 MAC, is used to investigate MAC behavior under different simulation scenarios. Although 802.15.4 has been shown to be not the most suitable solution for BAN, our purpose is not to reinforce or validate this fact. Rather, we are using the 802.15.4 MAC as a vehicle to reveal new aspects of MAC performance that have not been studied before and, in some cases, are non-intuitive, thus making their explanation particularly important for efficient BAN MAC design. Furthermore, 802.15.4 uses many techniques that are candidates for efficient BAN protocols, and indeed, several standard proposals already are basing their design on 802.15.4's functionality and terminology. Lacking a clear BAN MAC candidate, 802.15.4 is currently the best base case to study and reveal the different behaviors resulting from new channel conditions. We are particularly interested in pointing out the MAC performance differences between channel models that use temporal variation modeling and models that do not. Our results show that, indeed, there are important differences in the MAC layer based on the choice of channel model, with differences of up to 38% found in metrics such as the rate of packets received correctly or delay constraints and differences of up to 16% in saturation point traffic. Moreover, our explanations of these differences point to specific MAC design directions.

By pointing out the importance of the channel's temporal variation on the MAC, we do not exhaust the MAC design issues. For instance, the issue of co-existence (i.e., the efficient operation of a BAN in the presence of other BANs and other sources of radio communication in the same spectrum that cannot be centrally managed) is of great significance and should be considered in combination with accurate channel modeling. However, the co-existence issue is already receiving some attention in the IEEE BAN task group, unlike the impact of the channel on the MAC design. We are hoping that our work will provide a solid scientific background to bring this issue onto the agenda and serve to produce a more robust and efficient standard.

The rest of the article is organized as follows. Section 2 discusses related work. Section 3 briefly describes our channel modeling. Section 4 presents our simulation setup and assumptions. Section 5 presents the results from the various simulation scenarios, explaining the behaviors that produce them and pointing to specific MAC design choices. Finally, Section 6 concludes the article.

[1]IEEE BAN Task Group. `http://www.ieee802.org/15/pub/TG6.html`.
[2]Castalia Simulator. `http://castalia.npc.nicta.com.au`.

2. RELATED WORK

There are several research efforts to characterize the BAN wireless channel. Some are trying to fit the standard log-normal shadowing model to signal strength measurements acquired [Fort et al. 2006; Reusens et al. 2007], while others have tried to capture some large-scale temporal statistics [Cotton and Scanlon 2007]. In general, existing efforts are based on the observation of ensemble averages (such as that of Cotton and Scanlon [2007] in which a Nakagami-m fading distribution was observed); on the other hand, the main focus of our model is on the temporal correlations. Thus, our model captures the correlation between the last observed signal strength value, the time past, and the current signal strength value, which leads to more accurate capturing of the distribution and duration of the fades. This property is essential when investigating the effects of the channel on the MAC.

There are also several proposals for BAN MAC protocols. Su and Zhang [2009] propose a BAN MAC for healthcare applications that takes into account battery characteristics as well as varying channel characteristics. Although offering valuable ideas for BAN MAC design, they do not seek to explain behavior resulting from channel variations. Moreover, they only consider a quasi-static channel model, which does not accurately capture the typical time variation of received signal strength. Timmons and Scanlon [2009] propose a BAN MAC, while at the same time arguing the non-suitability of the 802.15.4 MAC for BAN. They compare the two MAC protocols, but their results are based on ideal channels with no temporal variation. Omeni et al. [2008] propose a new BAN MAC that they implement on hardware along with their physical layer design proposal. They offer important ideas for BAN MAC techniques that have proven efficient in certain scenarios and evaluate their MAC through experimentation. However, the experiments are limited to the static demo boards and are not extended to live experiments on the body. Thus, it is difficult to properly evaluate the MAC proposal since the adverse channel conditions are not accurately captured. Within the IEEE BAN task group, MAC protocols have been proposed, such as those of Davenport et al. [2009b] and Ding [2009]. They have not been evaluated yet, so it is unknown on upon what channel assumptions they have been designed. Our work does not propose a new MAC but rather lays the foundation for MAC design. Using appropriate simulation tools and channel models while, at the same time, not being tied to specific preconceived MAC techniques, we explore MAC techniques using the 802.15.4 MAC as our base. To the best of our knowledge, this is the first design groundwork effort of its kind.

3. BAN CHANNEL MODELING

Mobile and fading channel models have largely neglected slow-speed channels. In general, works such as Cotton and Scanlon [2007] have measured channel statistics over various periods of time (sometimes a few minutes to an hour) and collated the ensemble statistics. In particular, the models have provided overall fading rates (level crossings), average outage duration (fade length), and overall probability densities of the received signal strength. These are then used to create simulated channels by sampling at a particular rate and assuming the channel is fixed between samples. These so-called quasi-static channel models ignore the partial correlation between a particular channel sample and another at an arbitrary later point.

We have developed small and wearable hardware which enables the recording of a large number of simultaneous signal strength measurements, at the 2.4GHz carrier frequency, around the human body [Hanlen et al. 2009] whilst the wearer performs every day activities, including office work, walking, running, and sleeping. The design allows for continuous 200Hz sampling for up to 12 hours ($200 \times 60 \times 60 \times$

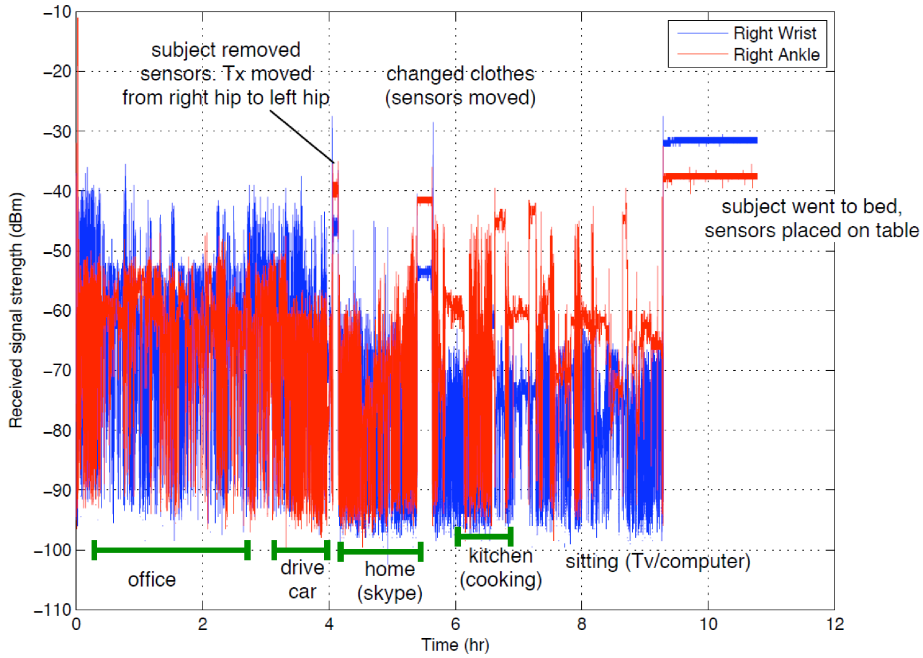

Fig. 1. Signal strength measurements in time.

12 = 8.6 Msamples per link). We have found that the signal strength varies continuously at up to 50Hz [Zhang et al. 2009]. The signal strength measurements have been used to develop time-variable- and every day movement-based models of the signal propagation. We have observed that the time-variation of the signal strength is the dominant component of the BAN wireless channel. Shadowing due to the human body, which absorbs approximately 60dB of signal power at this frequency, means that intuitive models (such as distance-vs-path loss models) of the physical layer are not appropriate.

The main components of the time-based signal are the following.

— Large fluctuations (+/– 30dB) of the channel gain around the median.
— Not dependent on distance (close to human body) [Hanlen et al. 2009].
— Log-normal or Weibull-distributed RSSI values [Zhang et al. 2009].
— For reasonable link margins (low-power transmissions), channel outages on the order of 50ms to 100ms [Miniutti et al. 2009].

Figure 1 shows a typical 2-link plot (the male subject wore a hip-mounted network hub, a transceiver on the right wrist, and another transceiver on the right ankle). It can be seen that activity type is not clearly distinguishable in the plot, although periods when the sensors are occluded by the subject's body or by furniture result in a large variation in the channel (e.g., when sitting at a desk and using a computer).

The temporal variation of the channel is used to develop statistical models of specific links (e.g., right wrist to left hip) as well as ensemble-average (generic) links. The analysis details may be found in Hanlen et al. [2009] and Zhang et al. [2009].

From the measurements acquired, we calculate the likelihood of a given sample value over the entire sample set. The overall probability density of the entire set of links, not surprisingly, turns out to be given by a log-normal (or nearly log-normal) distribution. This defines the long-term fading distribution. Given a particular sample

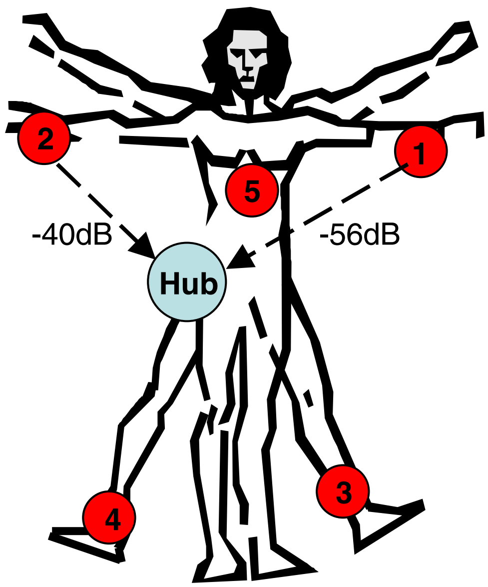

Fig. 2. The simulated network topology.

point, a second sample point is well modeled by the long-term distribution only when that second sample point is independent of the first. In the case of the body area channel, the time required to achieve true independence may be on the order of 100's to 1,000's of milliseconds (ms). Between correlated (within the coherence time of approx. 25–50ms) and independent (approx. 1–2 seconds) time scales, the channel exhibits partial correlation, that is, the probability of a delayed sample value given the current sample is dependent on the time delay.

$$P(s_{T>t} = X|s_t) = f(T - t, s_t).$$

Accordingly, we produce quantized conditional probabilities for a time-delayed sample (given the current sample) using various time intervals. This process allows us to transition from a continuous time (sample-driven) channel simulation to an event-based channel simulation, where the time steps are not regular. This kind of modeling is required due to the fact that Castalia, like most network simulators, is event-driven rather than time-driven. Thus, at every channel event the question asked is, ”Given that in the last event, a signal of value S was observed and that the last event happened T seconds ago, what will be the current signal value?” For large T, we expect the answer to be governed by the limits of the ensemble distributions. However, for shorter timeframes (e.g., in transmit-NACK-ACK-retransmit cycles) the conditional probability dominates.

4. SIMULATION SETUP

We are using the Castalia open source simulator; all simulations described in this article are released with Castalia 2.3, which assists our results in being reproducible and, furthermore, facilitates the evaluation of other BAN MAC proposals (such as the ones mentioned in the related work) using Castalia and our channel model.

Figure 2 shows the setup we have used throughout our simulations. One coordinator node is located on the right hip, and five sensor nodes (at the extremities

Table I. Radio Parameters

Data rate	1,024Kbps	512Kbps
Modulation	Diff QPSK	Diff BPSK
Rx sensitivity	−87dBm	−91dBm
Noise bandwidth	1MHz	
Noise floor	−104dBm	
Tx power	−10, −12, −15, −20, −25dBm	
CCA time	1ms	
Tx→Rx and Rx→Tx transition times	20μs	

of the four limbs and the chest) regularly send packets of 130 bytes (including overhead) to the coordinator. The average path loss between every pair is calculated from our experimental measurements. The figure shows the average path loss for a couple of links. The complete average path loss map can be found in the file Simulations/BANtest/pathLossMap.txt in Castalia's distribution. We have also implemented the time value correlated channel model described in Section 3 into Castalia. If the time difference of two events is less than 10ms, then the signal does not change. If the time difference is more than 5s, then the signal values are considered independent. For time differences in between, we have calculated several probability density functions that depend on the previously observed signal strength value.

Our physical radio model follows the model described in Davenport et al. [2009a]. Table I gives the various parameters of the two radios defined. Unless otherwise specified, the simulation results shown in the article use the fast radio.

We have also implemented most aspects of the 802.15.4 MAC standard, such as beacon mode, CAP (a contention-based access mode), and GTS (a TDMA-based access mode), which are the most relevant for the BAN case (i.e., single coordinator, multiple sensor nodes one hop away, sending data to the coordinator). We have chosen 120ms frames with 25% duty cycle (i.e., 30ms active time within one frame). Beacons are transmitted by the coordinator/hub at the start of each frame. Beacons are used for time synchronization between nodes and broadcasting control information, such as slot assignment in GTS mode. The standard defines 16 slots in each active period, 15 of which can be used for GTS (at least one slot has to remain available for contention-based traffic). When using GTS, we are statically assigning all 15 slots (three consecutive slots per node). This is an optimal assignment for the traffic rates used in the simulation scenarios (i.e., constant rates, identical for all nodes). Data packets are acknowledged. If an ACK does not arrive within a predetermined timeout (1.2ms), then a retransmission is attempted. By default, we use two transmission attempts (i.e., one retransmission), but in some simulations, we vary this parameter. Although the standard does not specify it, we are using a 32-packet incoming buffer for the MAC.

Given the data rates of the radios (1,024/512Kbps), the duty cycle (25%), and the number of transmitting nodes (five), we can calculate the theoretical maximum sustained traffic per node. This is 25.6Kbps for the slow radio and 51.2Kbps for the fast radio. Obviously, when control overhead, packet errors, and quantization (i.e., the fact that a time slot does not fit an integer amount of packets) are taken into account, the sustained maximum rate is considerably less. The application of each of the five sensor nodes produces traffic at a constant rate which is destined for the coordinator. In the simulations, the rate is expressed at a percentage of the theoretical maximum sustained traffic.

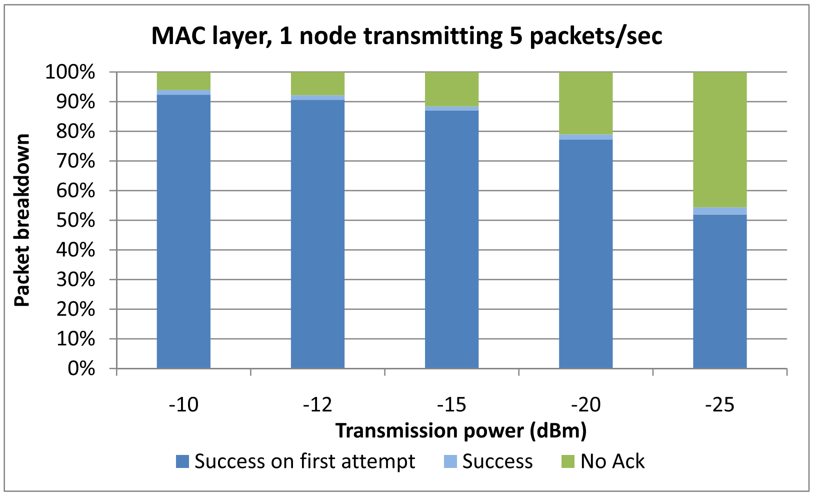

Fig. 3. Packet breakdown at the MAC layer.

5. PERFORMANCE OF 802.15.4 MAC UNDER BAN CHANNEL CONDITIONS

We have run a series of simulations using the 802.15.4 MAC, varying different parameters each time and observing different performance metrics, such as packet outcome breakdown and latency. The packet breakdown is observed at two different layers. One is the physical layer at the receiver (i.e., the coordinator node), where we can have the breakdown in terms of received packets, interfered packets, and packets lost to fades. Packets in this context refers to all packets coming out of the senders' radios, including control packets as well as retransmissions of previously failed data packets. The other observed packet breakdown is at the MAC layer of the senders; in this context, packets refers only to original packets from the application layer, excluding control packets and retransmissions. At the MAC layer, the packet categories are the following.

(1) First-time success (i.e., an ACK was received on the first transmission attempt).
(2) Success (i.e., an ACK was received on the second or later transmission attempt).
(3) NoACK (i.e., the packet received no ACK and was transmitted at least once on the radio).
(4) Channel busy (i.e., the packet failed because the CSMA mechanism never found the channel free (after several back-off efforts) in all transmission attempts).
(5) Overflown (i.e., the MAC buffer was full, so the packet was rejected).

The MAC layer breakdown is a straightforward metric for gauging the MAC performance. Furthermore, the physical layer breakdown can provide indications for the causes of some MAC behaviors, which is why it is used on some occasions. Another metric is the end-to-end delay of a packet which is described as *latency* and is given either as a histogram (quantized probability distribution) or as an average value over all packets.

5.1. Varying the Transmission Power

We begin our evaluation with a seemingly simple simulation. We allow only one node to transmit (node 3 with 59dB average path loss to the coordinator), and we vary the transmission power. Figure 3 shows the packet breakdown at the MAC layer of node 3 at low produced traffic (5 packets/sec = 5.2Kbps).

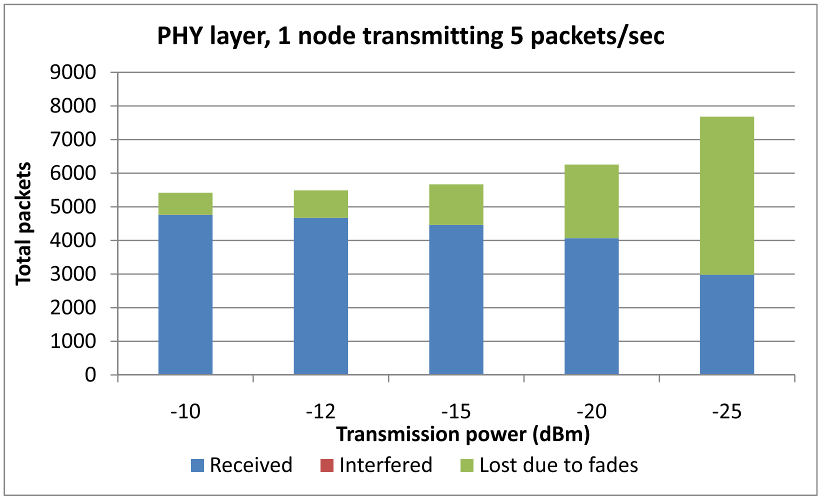

Fig. 4. Packet breakdown at the PHY layer.

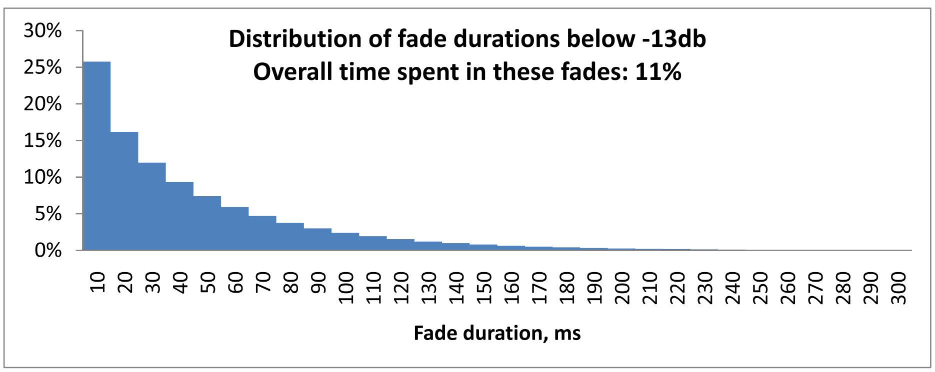

Fig. 5. Fade duration distribution for link margin −13dB [= −87dBm – (−59dB – 15dBm)].

We can see that the portion of failed (not acknowledged) packets increases exponentially as the transmission power is decreased. This is a rather expected result. We also see a small percentage of packets succeed on the second attempt. This portion is small compared to the failed packets, which means that retransmissions are largely not effective. To see this more clearly, Figure 4 shows the physical layer breakdown at the receiver.

Figure 4 shows the absolute number of packets to make our point more clear. As the transmission power is decreased, the portion of faded packets increases (corresponding to an increased number of NoACK packets in the MAC). The overall number of physical-layer packets increases as well, because more retransmission attempts are made. We note that these retransmissions are mostly failing: the graph shows a strong coupling between increased packets and increased failed packets. This is also expected, as the fades in our channel last for at least 10ms, which is well within the retransmission time. In some cases, if the first transmission happens towards the end of a fade, the retransmission can be successful; such events are not very frequent but explain the existence of some successful retransmissions. Figure 5 shows the distribution of

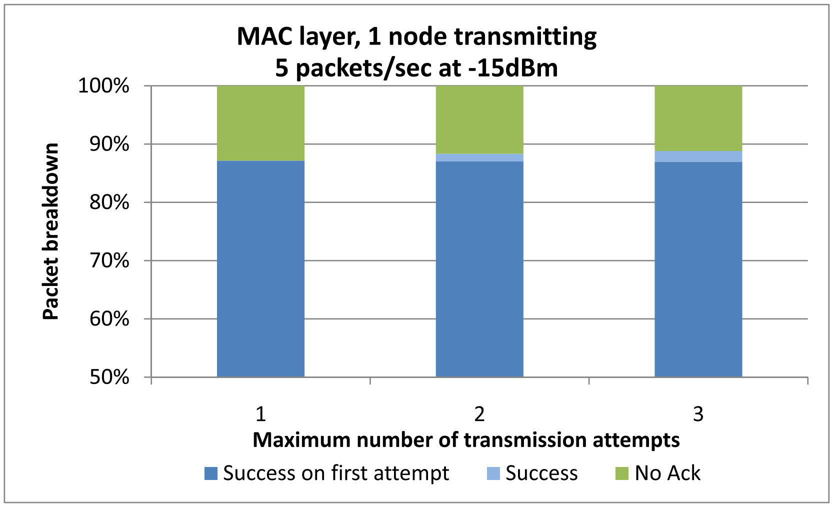

Fig. 6. Packet breakdown at the MAC layer, low traffic rate.

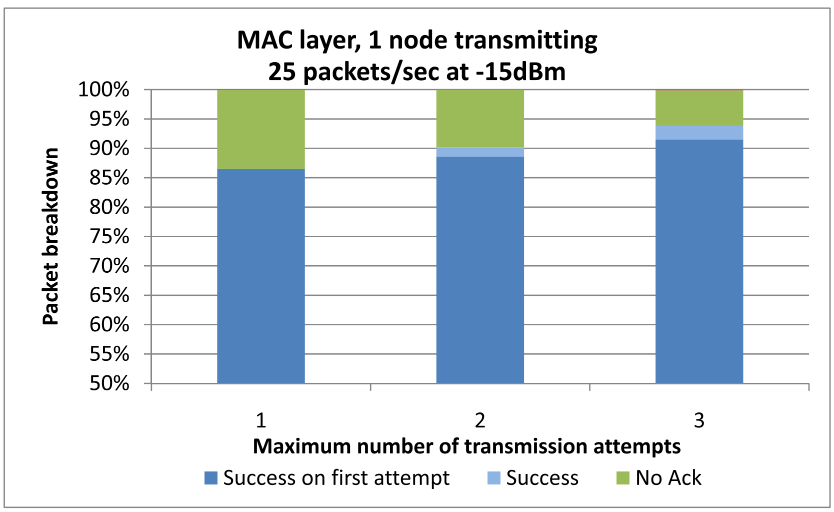

Fig. 7. Packet breakdown at the MAC layer, high traffic rate.

the duration of deep fades (i.e., fades that bring the signal below the radio sensitivity) for node 3 (59dB path loss) when −15dBm transmission power is used.

A finer point to note is that energy consumption may increase for the nodes if retransmissions are allowed, despite the lower power used. For the 802.15.4 MAC, the energy does not change much since the nodes have an active radio for the CAP duration time and their GTS assigned slots whether they have something to transmit or not. For this reason, we do not show energy results for the 802.15.4 MAC. However, a MAC designed specifically for BAN might only allow an active radio when the node is transmitting. In this case, and depending on the radio parameters, retransmissions can be costly in terms of energy, resulting in only a small percentage of additional successes.

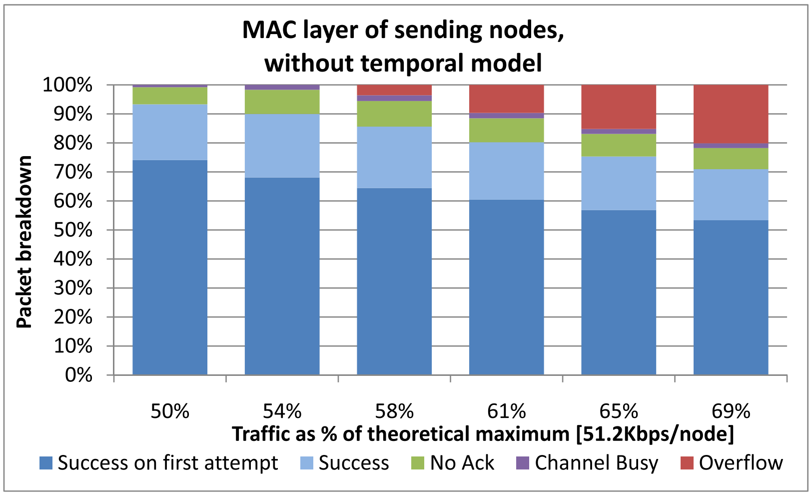

Fig. 8. Packet breakdown at the MAC layer, no temporal signal variation.

5.2. Varying the Number of Transmission Attempts

Investigating the issue of retransmissions, we conducted a series of simulations in which, again, only node 3 was sending packets to the coordinator, the transmission power was set to −15dBm, and the transmission attempts varied. Figure 6 shows the results at the low rate of 5 packets/sec = 5.2Kbps.

As expected, almost the same number of packets succeed on the first attempt, and we observe a slight increase in the number of packets successful on subsequent attempts.

The results become more interesting when we consider the same simulation but with higher traffic: 25 packets/sec = 26Kbps. With this rate of traffic and taking into account the duty cycle (30ms active, 90 ms sleeping), several packets are queued to be transmitted at the start of an active period. Figure 7 shows the results.

We notice that the percentage of first-time successful packets increases. Upon first inspection, this is counterintuitive since the channel statistics remain the same. However, this is explained by the correlation among subsequent packets. As mentioned previously, at the start of each active period, there are several (three, on average) packets waiting to be transmitted. Whereas in the low-rate case, the packets are transmitted sparsely without much correlation to each other, now most packets immediately follow the transmission of another packet. This implies that if the first packet fails and tries to retransmit, the retransmissions will take a considerable amount of time from the "bad" channel state. In turn, this will result in a better-than-average probability for the next packet to be transmitted over a "good" state channel and succeed on the first time. Notice that the key issue here is not the retransmissions but the total time they take. Alternatively, upon not receiving an acknowledgment, a BAN MAC protocol can simply wait for 10ms (or another duration based on the channel statistics) before either trying a retransmission or proceeding with the next packet. Delayed retransmissions in this manner can help improve overall success rates for low traffic rates; however, for higher rates, such retransmissions may not be the preferred solution, as they will increase the latency of every packet. In other words, for a high traffic-rate case, although achieving similar success percentages at the physical layer, each successful packet will

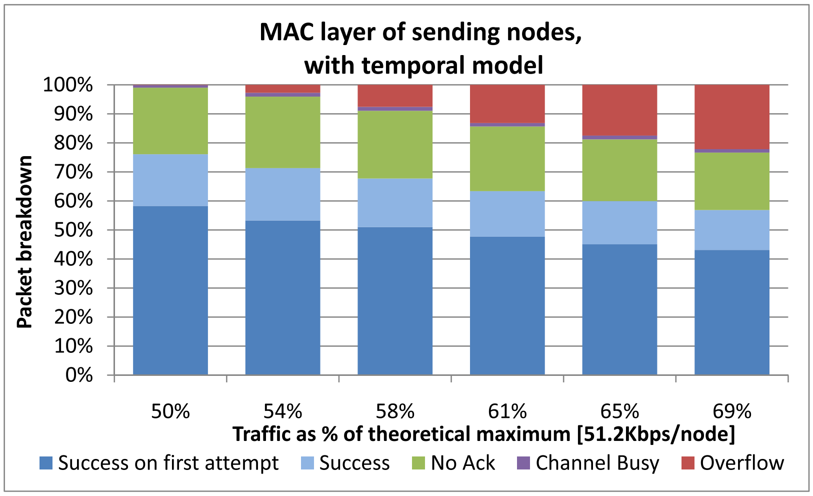

Fig. 9. Packet breakdown at the MAC layer, with temporal signal variation.

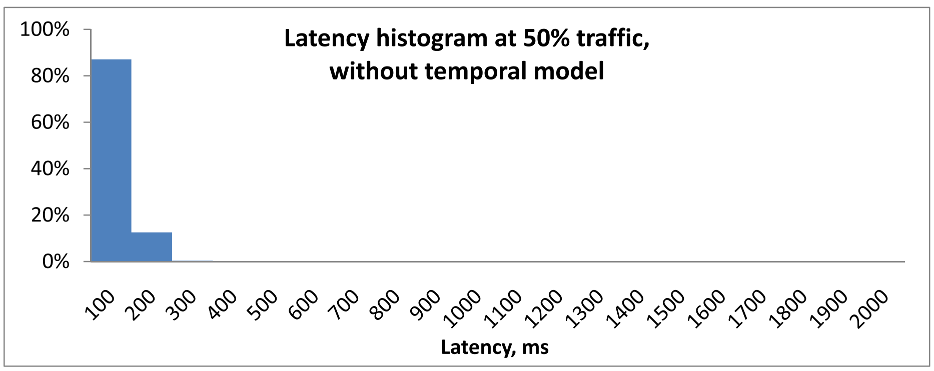

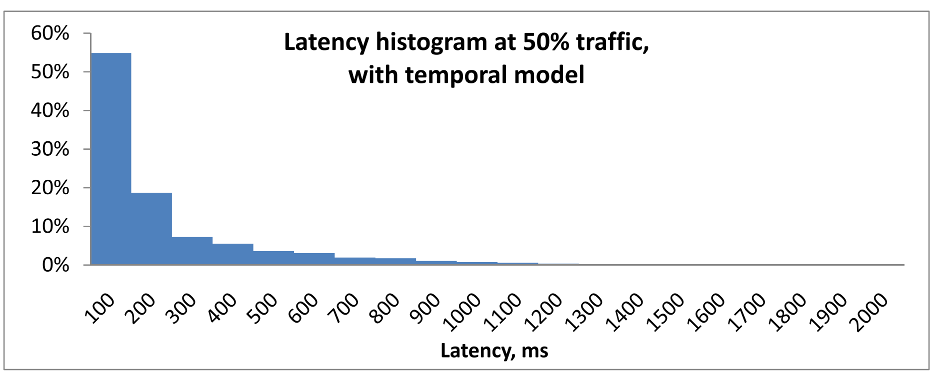

Fig. 10. Latency distributions for 50% traffic load.

be delayed for longer. We will delve more into latency issues in the following sections. This particular design trade-off falls outside the protocol specification of 802.15.4, so we do not test it in this paper. However, it is a design direction that BAN MACs should explore.

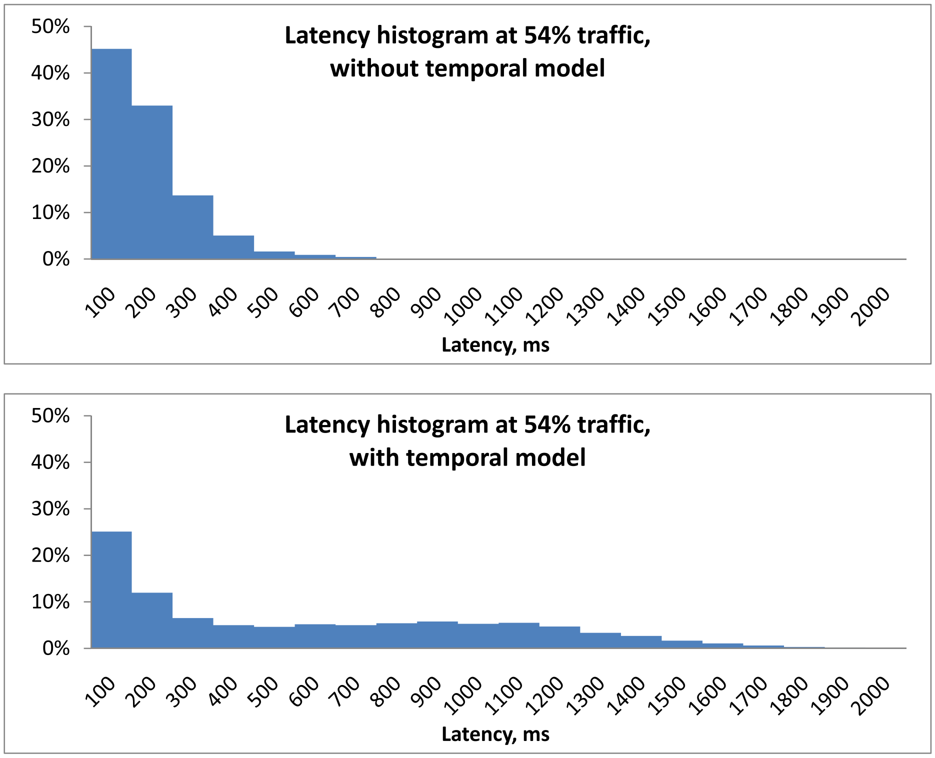

Fig. 11. Latency distributions for 54% traffic load.

5.3. CAP Performance Under Varying Traffic

So far, the cases we explored had only one node transmitting. Obviously, we would like to study the behavior of the MAC when all five nodes are sending data packets to the coordinator. We are interested in the performance of the MAC when our time-varying channel model is used and also when no such model is used. In the simpler cases explored previously, a comparison with the nontemporal model was not provided since the latter always produced perfect results. We first investigate the contention-based access mode of 802.15.4, called CAP. We would like to note here that it was not trivial to adjust the parameters of the 802.15.4 MAC in order to achieve acceptable collision levels with the given radio model we have. We had to explore the min and max back-off exponent parameters over a range to decide on the best choice. The chosen parameters result in about 10% collisions (as measured at the receiver's physical layer) across a wide range of traffic loads (from low traffic to saturation). Figure 8 shows the performance of the MAC for a channel with no temporal variation for various traffic loads around the saturation point (notice the overflown packets). Below 50% traffic load, the performance is similar to the 50% load for both the temporal and the nontemporal model. Above 69% of the portion of overflown packets increases at the expense of the other categories.

Around 55% of the theoretical maximum supported rate (51.2Kbps) the network gets saturated, with 90% of the packets getting through ($<$ 70% on the first attempt).

Figure 9 shows the corresponding results with the use of a temporal model.

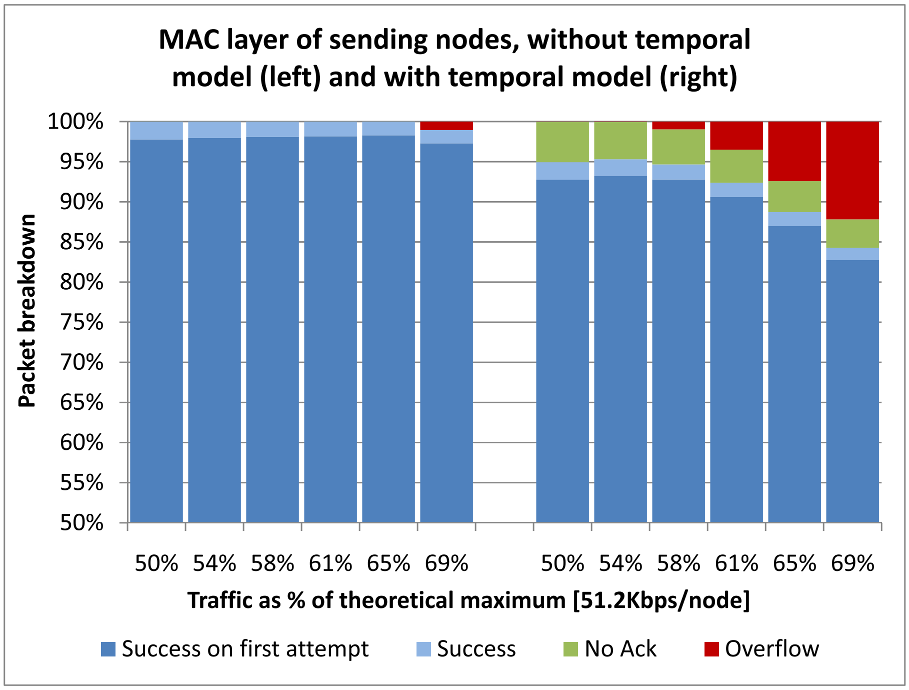

Fig. 12. Comparison of GTS performance with and without a temporal model.

Notice that (1) the saturation starts earlier than the nontemporal case, (2) that there is a higher portion of NoACK packets, and (3) that there are fewer successes and fewer first successes.

5.3.1. Latency Performance. When gauging the performance of an MAC, especially close to the saturation point, we should not overlook the latency, or end-to-end delay, of packets. The BAN requirements document from the IEEE BAN task group [Zhen et al. 2008] mentions latency bounds of 125ms for emergency traffic and 250ms for normal traffic. Looking at the latency, we discover important differences between temporal and nontemporal models. Figure 10 shows the latency distribution for the 50% traffic load with and without the temporal model. Notice that without a temporal model, all successful packets meet the latency requirements, whereas only 75% of the packets simulated with a temporal channel model meet the requirements. Taking this into account, we now see a much greater difference in successful packets that meet the latency requirements between temporal and nontemporal (approx. 93% vs. 57% of all packets). Figure 11 shows similar latency histograms for 54% traffic. There, the signs of saturation even without a temporal model start to show. At 58% (not shown in a graph), most packets experience large delays.

5.4. GTS Performance Under Varying Traffic

A contention-based mode for a BAN MAC is not really a viable candidate since it does not make good utilization of the channel and nodes can spend considerable time in collisions and channel sensing. A TDMA approach, on the other hand, seems much more suitable, especially since we have a natural coordinator and all nodes are within one-hop distance of each other (at least when average path losses are taken into

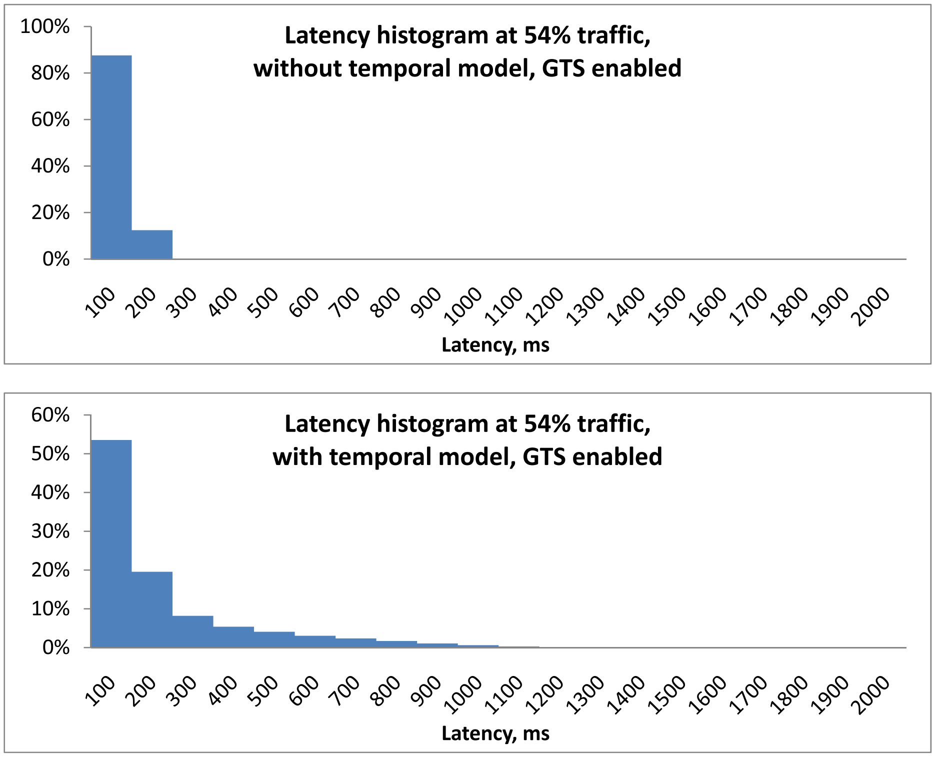

Fig. 13. Latency distributions for 54% traffic load.

account). For this reason, an investigation of the GTS performance of the 802.15.4 MAC with and without a temporal model can be very useful. When we simulate without a temporal channel model, then the results are perfect (100% success rate) up to the saturation point of the network. We get a small percentage (2%) of successes from retransmissions. These are due to collisions, which in turn happen in the only slot of the active frame that allows CAP. As explained in Section 4, when in GTS mode, 15 of the 16 slots are assigned to nodes, and one slot (the first) remains for contention-based access, resulting in possible collisions. The network is saturated at 69% (much less than the theoretical 100%) for the reasons explained in Section 4.

The performance of GTS is considerably different when a temporal channel model is used. Figure 12 shows these differences in a compact manner.

We notice that even before saturation, the percentage of packets that are not acknowledged (in all transmission attempts) is around 6%. This is an expected effect of the deep fades. Successful packets on second attempts are about the same as those of the nontemporal case, which implies that these successes come mainly from the few collided packets and not from faded packets. The other major point is that saturation arrives much earlier for the temporal case (58% traffic load instead of 69%).

5.4.1. Latency Performance. Figure 13 shows the latency distributions for the 54% traffic load with and without a temporal model. For the temporal model case, we see the start of saturation, although the packet breakdown graph in Figure 12 does not show any overflow packets at this particular load. Only 75% of the packets are within the latency requirements, making the performance with the temporal model

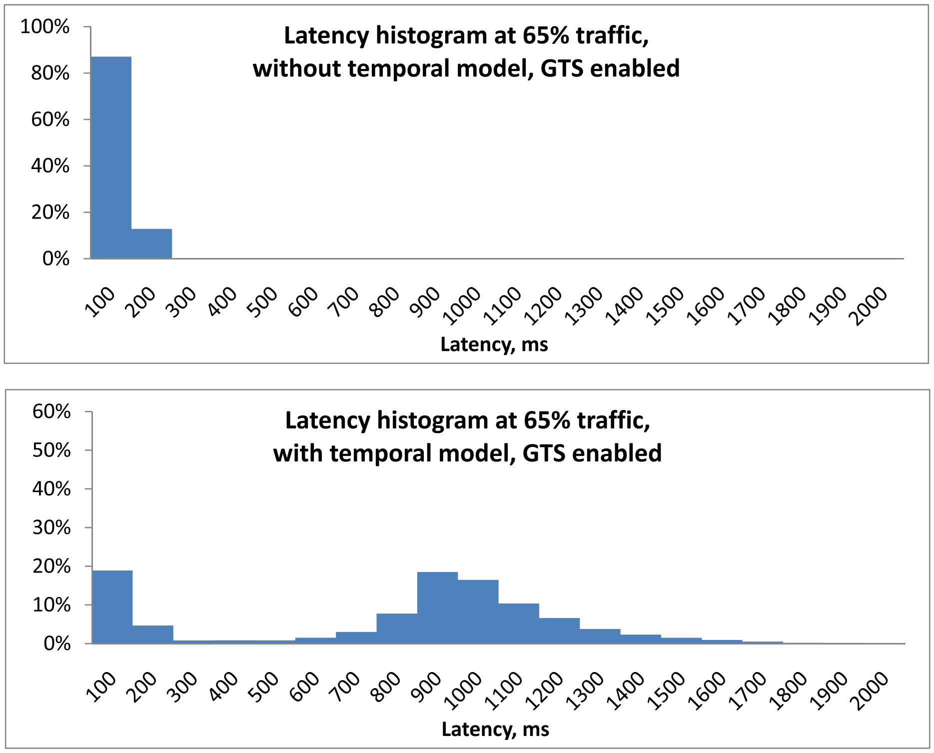

Fig. 14. Latency distributions for 65% traffic load.

considerably worse than without it. More specifically, instead of the 100% vs. 95% difference of successful packets, we now have [100% vs. 95%·75% = 71%] difference of successful packets that meet the latency requirements.

Performance without a temporal model continues to produce ideal latency results up to 65%. For the same load, performance with a temporal model results in most packets being over the latency limit. Figure 14 shows the respective latency histograms. At 69%, both cases result in a saturated network with most packets experiencing large delays.

The interesting thing about the lower part of Figure 14 is that there is a portion of packets (about 1/5 of the total) that experience very low delays compared to the rest of the packets. Looking deeper into debug traces, we discover that these packets come almost exclusively from node 2. Not surprisingly, node 2 is the best connected node to the coordinator with only −40dBm average path loss. This implies that it rarely experiences deep fades, which in turn implies that it does not miss many beacons from the coordinator nor does it try many retransmissions due to fades.

In order to further investigate the effect of lost beacons (i.e., the lost opportunity to transmit for a whole frame) and the effect of retransmissions in overall latency, we carried out the following simulations. Keeping all other parameters fixed, we increased the transmit power of the coordinator only, from −15dBm to −10dBm (the maximum allowed transmit power around the human body is still under discussion but many limit it to −10dBm). In this way, the sensor nodes lost only 3.2% of all beacons, as opposed to 7% when −15dBm was used. The effects of this change to the average packet latency for the load cases we considered before is shown in Figure 15.

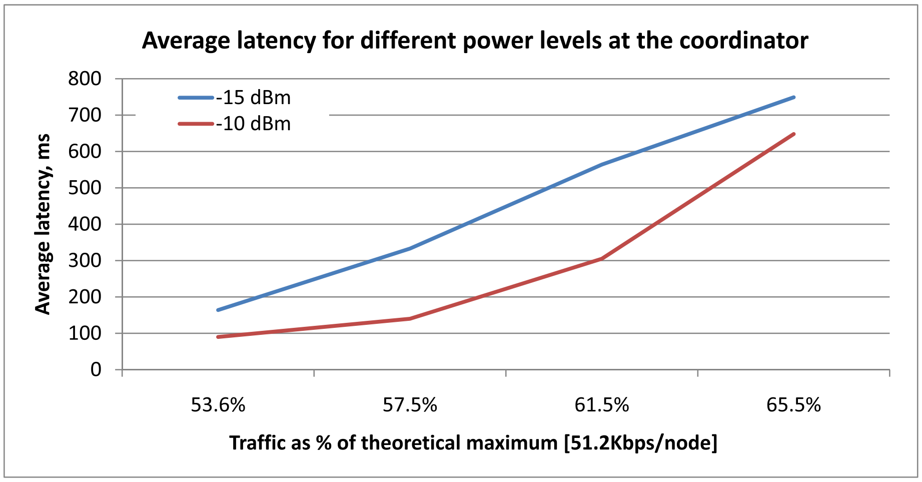

Fig. 15. Average packet latency for different coordinator Tx power.

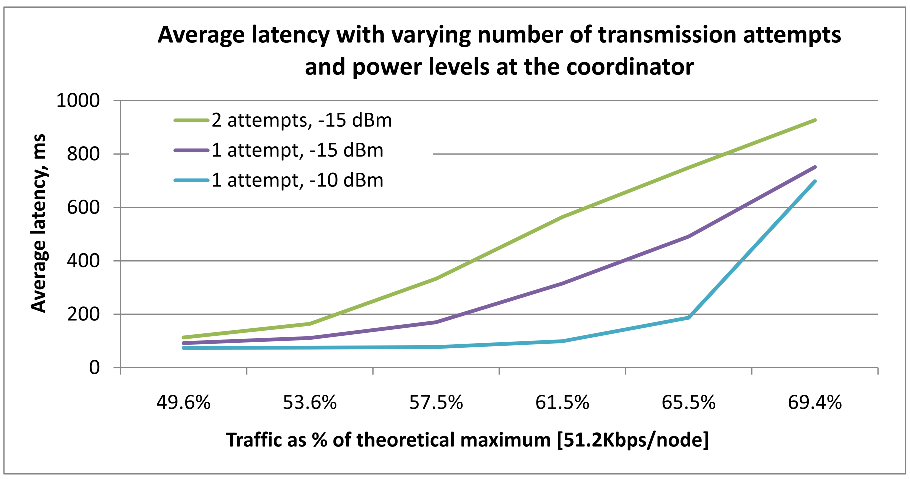

Fig. 16. Average packet latency for different schemes.

We notice that there are considerable differences, especially for rates below 61%. We also notice that this change alone does not equate to the results with the nontemporal model latency, so retransmission should affect latency. To explore this rationale, we simulated the scenarios with only one transmission attempt and, finally, with the combined effect of one transmission attempt and the coordinator transmitting at -10dBm. Figure 16 summarizes the results.

We observe that having only one retransmission improves the average latency in a significant way. However, the combination of fewer lost beacons and one transmission attempt makes the average latency comparable, even for traffic rates close to the saturation point (saturation for the nontemporal case). Of course, it should be noted that allowing only one attempt decreases the rate of successful packets, not only due to excluding packets successful on the second attempt but also by indirectly reducing

the percentage of the first successes, as explained in Section 5.2. Whether this is a beneficial method will depend on the application priorities (e.g., reliability, timeliness, throughput) and the specifics of the radio and traffic. On the other hand, the increase of the coordinator power is a positive move in most scenarios, since the increase in the energy consumption of the coordinator is generally not a problem (in most BAN scenarios, the coordinator can be considered rechargeable).

6. CONCLUSION

We have shown that the time-varying channel associated with a human body considerably influences the performance results of the 802.15.4 MAC, especially in the GTS mode. Whilst we have considered the 802.15.4 MAC in the context of BAN (802.15.6), our results are extensible to the design of BAN MAC protocols since they refer to generic MAC techniques (e.g., beacons, TDMA, retransmissions) as well as specific BAN channel conditions. Based on our results, we advocate an asymmetric power selection between the coordinator and the sensor nodes and the consideration of avoiding retransmissions. BAN MAC protocols should deter transmission when in a deep fade. In order to estimate the duration of these fades, an accurate channel model is needed. MAC designers should also account for the reduced capacity of the network because of the temporal variation. We believe that these findings along with further usage of accurate time-varying channel models on specific BAN MAC proposals will result in a more reliable and efficient BAN MAC standard.

REFERENCES

COTTON, S. L. AND SCANLON, W. G. 2007. Characterization and modeling of the indoor radio channel at 868 MHz for a mobile bodyworn wireless personal area network. *IEEE Antennas Wire. Propag. Lett. 6,*, 51–55.

DAVENPORT, D., SEIDL, N., MOSS, J., PATEL, M., BATRA, A., HO, J., HOSUR, S., ROH, J., SCHMIDL, T., OMENI, O., AND WONG, A. 2009a. MedWin physical layer proposal documentation. IEEE Document P802.15.09-0329-00-0006.

DAVENPORT, D., SEIDL, N., MOSS, J., PATEL, M., BATRA, A., HO, J., HOSUR, S., ROH, J., SCHMIDL, T., OMENI, O., AND WONG, A. 2009b. MedWin MAC and security proposal documentation. IEEE Document 802.15-09-0326-01-0006.

DING, G. 2009. Olympus MAC proposal presentation. IEEE Document P802.15-09-0311-00-0006.

FORT, A., RUCJAERT, J., DESSET, C., AND DE DONCKER, P. 2006. Ultra-wideband channel model for communication around the human body. *IEEE J. Select. Areas Commun. 24,* 4, 927–933.

HANLEN, L. W., MINIUTTI, D., RODDA, D., AND GILBERT, B. 2009. Interference in body area networks: Distance does not dominate. In *Proceedings of the International Symposium on Personal Indoor and Mobile Radio Communication*.

MINIUTTI, D., HANLEN, L. W., SMITH, D. B., ZHANG, J., RODDA, D., GILBERT, B., AND GENGFA, F. 2009. NICTA proposal. IEEE Document P802.15-09-0345-00-0006.

OMENI, O., CHI WAI WONG, A., BURDETT, A. J., AND TOMAZOU, C. 2008. Energy efficient medium access protocol for wireless medical body area sensor networks. *IEEE Trans. Biomed. Circuits Syst 2,* 4, 251–259.

REUSENS, E., JOSEPH, W., VERMEEREN, G., AND MARTENS, L. 2007. On-body measurements and characterization of wireless communication channel for arm and torso of human. In *Proceedings of the 4th International Workshop on Wearable and Implantable Body Sensor Networks*. 264–269.

SU, H. AND ZHANG, X. 2009. Battery-dynamics driven TDMA MAC protocols for wireless body-area monitoring networks in healthcare applications. *IEEE J. Select. Areas Commun. 27,* 4, 424–434.

TIMMONS, N. F. AND SCANLON, W. G. 2004. Analysis of the performance of IEEE 802.15.4 for medical sensor body area networking, In *Proceedings of the 1st Annual IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks*. 16–24.

TIMMONS, N. F. AND SCANLON, W. G. 2009. An adaptive energy efficient MAC protocol for the medical body area network. In *Proceedings of the 1st International Conference-Wireless Communication, Vehicular Technology, Information Theory and Aerospace & Electronic Systems Technology*. 587–593.

ZHANG, J., SMITH, D. B., HANLEN, L. W., MINIUTTI, D., RODDA, D., AND GILBERT, B. 2009. Stability of narrowband dynamic body area channel, IEEE Antennas and Wireless Propagation Letters, 2009. Stability of narrowband dynamic body area channel. *IEEE Antennas Wirel. Propag. Lett.*

ZHEN, B., PATEL, M., LEE, S., WON, E., AND ASTRIN, A. 2008. IEEE 802.15.6 Technical Requirements Document, IEEE Document 802.15-08-0644-09-0006.

Received October 2009; revised July 2010; accepted September 2010