Challenges in Body Area Networks for Healthcare: The MAC

Athanassios Boulis, NICTA

David Smith and Dino Miniutti, NICTA and the Australian National University

Lavy Libman and Yuriy Tselishchev, NICTA and the University of Sydney

ABSTRACT

Body area wireless sensor networks (BANs) are a key component to the ubiquitous healthcare revolution and perhaps one of its most challenging elements from a communications standpoint. The unique characteristics of the wireless channel, coupled with the need for extreme energy efficiency in many healthcare applications, require novel solutions in medium access control protocols. We present the main characteristics and challenges associated with BANs from a healthcare perspective, and present some MAC techniques based on studies of the BAN channel that could be used to address these challenges.

INTRODUCTION

Ubiquitous healthcare promises continuous gathering and analysis of physiological, behavioral and other such health-related information and either acting on it, providing feedback, or delivering it to health professionals. The aim is to provide better healthcare to people via continuous monitoring of small autonomous wireless body-worn sensors, often outside of typical healthcare settings such as hospitals [1]. If it is to become commonplace, the gathering of such health-related information needs to be unobtrusive and cheap, yet its delivery must fulfill certain reliability and timeliness guarantees for the data to remain useful. Body area wireless sensor networks, or body area networks (BANs), are an emerging technology that aims to satisfy these requirements [2–6]. A BAN consists of one or more small sensors (e.g., temperature, blood pressure) placed on or inside the human body that communicate their readings back to a hub device on or near the individual (e.g., BAN-enabled smart phone), which may in turn communicate to others via another longer-range network (e.g., WiFi, 3G).

The focus of this article is the short-range wireless communication network that is formed between the sensors and the hub in a healthcare BAN, particularly at the medium access control (MAC) layer. There are many challenges in designing such a network, with reliability and energy efficiency arguably the most important considerations. The importance of reliability is obvious when one considers that BANs may be employed to monitor the vital signs of people in poor health; one missed emergency signal could be the difference between life and death. Energy efficiency is also important for BAN sensors as their small size limits the size of their batteries.

In this article we first comment on the important challenges for the BAN MAC and then show how the proposed IEEE 802.15.6 standard for BANs improves on existing standards to help meet these challenges. We then present some results from our studies of the BAN channel, including a summary of its statistical properties and behavior. We argue that this knowledge is key to the design of reliable and efficient BANs, and conclude by demonstrating how it can be used by presenting four MAC techniques for improving reliability and efficiency.

IMPORTANT CHALLENGES FACED BY BAN

Reliability for BAN healthcare applications is vitally important, and it needs to be met in the face of these four particular challenges for BANs.

Extreme energy efficiency: To deliver the levels of comfort and unobtrusiveness required for widespread adoption, BAN sensor nodes must be small and have batteries that last on the order of days to years, depending on the application. The size requirement obviously limits the size of the batteries that will power the nodes (energy scavenging is another option, but the amount of power available from such techniques is relatively small), so BAN nodes must be extremely frugal in their energy usage.

Unique characteristics of the wireless channel: The behavior of the wireless channel around the human body poses a unique set of challenges to reliable communication. The first of these challenges is severe attenuation of the wireless signal between the sensor and the hub that can occur (attenuation of over 100 dB has been observed) and may push the received signal power below the level required for reliable com-

munication. This level is referred to as the receiver's *sensitivity*, and is typically limited in BAN nodes due to their relatively small antennas and simple energy-efficient designs. Unlike other longer-range networks where the distance between the transmitter and receiver dominates signal attenuation, the strength of a BAN signal is most affected by the physical location and orientation of the nodes in relation to each other as well as the human body, which can "shadow" or attenuate the signal. Moreover, people move about, which constantly changes the attenuation at a rate that depends on the type of physical activity.

Managing interference: The nodes in a BAN can be centrally coordinated by the hub, thus allowing a large number of devices to coexist in a single network without having them interfere with each other. Things become more complicated when multiple people wearing BANs come into range of each other. In this case coordination may become impossible. The difficulty comes from the peoples' actions, which are unpredictable from a network's viewpoint, and can result in networks moving into and out of range of each other. In such a situation there is no natural way of choosing a network coordinator. As a corollary, any interference mitigation scheme will need to adapt faster than the rate at which the network topology changes if it is to be successful in minimizing interference

Application requirements: BANs also need to support a wide range of throughput rates (1 kb/s to 10 Mb/s) to accommodate higher throughput applications such as video, while still delivering the high reliability and low-latency required in many medical applications.

THE PROPOSED IEEE 802.15.6 STANDARD

There are a number of existing standards, such as Bluetooth and the IEEE 802.15.4 standard for wireless personal area networks (WPANs) [7, 8], that at first glance seem appropriate solutions for meeting the important challenges faced by BANs outlined in the previous section. This may indeed be true for some BAN applications, but these existing standards are not flexible enough to cover the wide range of BAN applications; nor can they meet BANs' tough reliability and energy efficiency requirements. Hence, the IEEE 802.15.6 working group was created in November 2007 to draft physical and MAC layer standards for wireless communication to address the particular challenges found in BANs [9].

One of the key features of the IEEE 802.15.6 BAN standard is its flexibility, which allows it to support diverse applications, ranging from healthcare devices, such as medical implants, to consumer entertainment devices, such as video goggles. Each of these applications has its own set of requirements; most notably data rate, reliability, and energy efficiency, which cannot be simultaneously maximized. The IEEE 802.15.6 standard gives device manufacturers the flexibility to design their devices in such a way as to trade off each of these features against each other, depending on the application. Other competing standards offer some such flexibility, but not to the extent of the IEEE 802.15.6 standard.

In terms of flexibility, the most important feature of the IEEE 802.15.6 standard is that it offers multiple physical layer modes of operation: narrowband, ultra-wideband (UWB), and human body communication (HBC). Furthermore, depending on the mode, a number of carrier frequency bands are supported.[1] This is important since some applications perform better at certain frequencies; for example, lower radio frequencies offer less attenuation and shadowing from the human body, while ultra-wideband frequencies can offer higher data rates. Moreover, specific applications are regulated to only operate within specific frequency bands (e.g., implants may only operate at 402–405 MHz worldwide).

The IEEE 802.15.6 MAC also offers a great deal of flexibility by offering a number of different access modes. As do other low-power standards, 802.15.6 employs a network coordinator, which sends out beacons to organize time into superframes (i.e., intervals between beacons) and slots (i.e., small intervals within a superframe allocated using a multiple access mode). Using IEEE 802.15.4 as a baseline for comparison, IEEE 802.15.6 adds polling/posting, also known as "improvised access," whereby the hub/coordinator can inform sensor nodes that they have been granted one-off exclusive time slots to transmit or receive information; this is unlike regular time-division multiple access (TDMA) which follows a regular schedule. Scheduled TDMA traffic is still supported, and the standard allows the nodes to be scheduled to periods greater than the superframe period. Carrier sense multiple access (CSMA) is also supported, and the standard offers many more ways to combine the different access modes, controlling their length and their place in the superframe, and allowing them to be woven together in the same superframe, whose length can also be changed with a more fine-grained specification. Energy efficiency is also increased via mechanisms that allow sensor nodes to enter a low-power sleep mode for a long time (several beacon periods) before transmitting/receiving. Furthermore, each device need not implement all of the access mechanisms, choosing only those that are suitable for its operation, giving device designers more flexibility to choose the mode that will meet the cost, reliability, energy, and other requirements that are needed for the device's specific application.

Reliability is also a key requirement in many BAN applications, especially in healthcare, and the IEEE 802.15.6 standard attempts to maximize reliability in a number of ways. As wireless spectrum becomes increasingly crowded, the interference from other networks (BAN and non-BAN) is also increasing, reducing potential throughput and reliability. The IEEE 802.15.6 standard attempts to minimize the impact of interference by shifting the beacon transmission by a known offset each beacon period. This is important as the beacon is crucial to the network, coordinating all of the traffic and providing a timing reference for each of the nodes. Moreover, the standard provides a dynamic

One of the key features of the IEEE 802.15.6 BAN standard is its flexibility, which allows it to support diverse applications, ranging from healthcare devices, such as medical implants, to consumer entertainment devices, such as video goggles.

[1] *402–405 MHz, 420–450 MHz, 863–870 MHz, 902–928 MHz, 950–956 MHz, 2360–2400 MHz and 2400–2483.5 MHz for narrowband; 3245–4742 MHz and 6240–10234 MHz for UWB; 14–18 MHz and 25–29 MHz for HBC.*
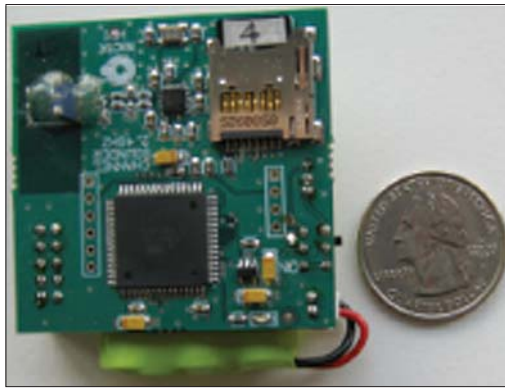
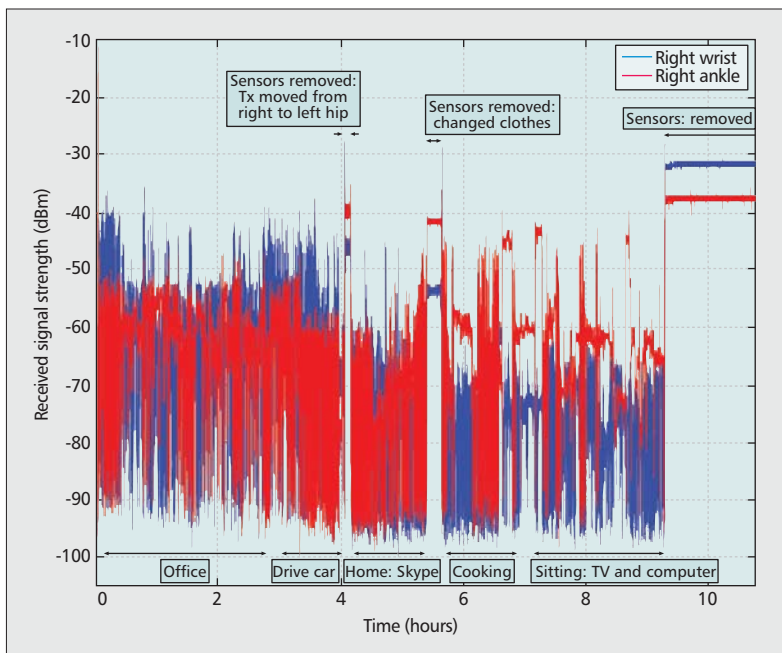**Figure 1.** *Our "channel sounder" — a small wearable radio.*

Figure 2. *Example of received signal strength we have observed over several hours of "everyday" activity.*

channel hopping mechanism that helps the network avoid other narrowband interferers. Finally, the standard also allows a single relay (i.e., dual-hop) to be used in cases where the typical single-hop star topology cannot maintain the required levels of reliability. This is especially useful in cases where the person is not moving, such as when they are sleeping, as some links may be attenuated below the receiver sensitivity for tens of minutes at a time [10].

Despite the suitability of IEEE 802.15.6 for BAN, successful systems are highly dependent on the characteristics of the BAN wireless channel and need to be designed with an understanding of its behavior, which the next section outlines briefly.

THE BAN WIRELESS CHANNEL

Compared to ultrawideband and human body conduction systems, narrowband wireless communication is arguably best suited to the greatest

number of healthcare applications and is thus the focus of this section. At channel bandwidths typical of narrowband BAN systems, the radio channel has been shown to be essentially slow and flat-fading, with an insignificant amount of intersymbol interference from multipath. Consequently, the received signal strength is a good measure of the channel at any point in time. That said, the movement of the human body has a dramatic effect on the strength of the received signal; hence, static measurements of the BAN channel at a single point in time provide limited useful information to those designing BAN systems; long-term measurements, which are characterized statistically and capture a wide variety of "everyday activities," are far more relevant.

The authors of this article have performed a wide variety of dynamic channel measurements of "everyday" mixed activities near the 900 and 2400 MHz industrial, scientific, and medical (ISM) bands totaling many hundreds of hours, mainly using custom-built "channel sounders," a collection of small body-worn radios than can transmit, measure, and log received signal strength at 2.4 GHz [11]. Activities and environments have included walking, running, driving, at the office, at home, outdoors, and sleeping. One of our channel sounders is depicted in Fig. 1, and Fig. 2 illustrates one of our everyday channel sounder measurements, plotting the signal strength between two body-worn devices and a hub worn on the hip.

The results of our "everyday" channel sounder measurements have shown that attenuation, predominantly due to shadowing, can be significantly large, with median channel attenuation greater than 70 dB. Furthermore, we have confirmed significant temporal correlation for samples of the channel gain in periods of 500 ms or more, with the samples only becoming independent for periods greater than 1 s. Speaking informally to illustrate the point, one can think of the channel gain moving around between stable levels for short periods of time. If one looks at any such stable period, the channel gain will be highly correlated, but the aggregate of all such periods results in an independent signal with an apropos first-order distribution. Furthermore, and important, a slow changing channel means that we can make predictions at the MAC level about the channel's future state and use this knowledge to improve the reliability.

Given that the BAN received signal strength can remain relatively stable for significantly long periods of time (for a wireless system), this suggests that the signal strength could remain poor for long periods too. These periods of large attenuation are vitally important as they can limit throughput or even make reception impossible; we call a continuous period of attenuation below sensitivity an "outage." The red curves in Fig. 3 plot the mean duration of such outages against the maximum attenuation a system can support before going into outage, ranging from 20 ms to 10 s. The blue curve shows the probability of being in an outage at any point in time. Both curves are plotted for "everyday" scenarios, as well as for someone sleeping in a bed.² Note that more sensitive receivers can reduce the likelihood and duration of outages, but options are

limited in the physical layer due to size, cost, and complexity constraints.

MAC Techniques for a Reliable and Efficient BAN

It is important to understand that the mechanisms provided in the draft IEEE 8021.15.6 standard do not constitute a complete MAC protocol. Rather, they outline the basic elements that must be used to ensure interoperability among IEEE 802.15.6 compliant devices, such as packet formats and message exchange protocols to achieve simple tasks (e.g., assign an allocation interval, or acknowledge the reception of a packet) [10]. The draft standard leaves open higher level questions such as: Should we use contention-based, scheduled, or improvised access, and when? In what order do we schedule allocation intervals? When should relays be used? When should one use retransmissions to cope with failed packet reception?

The understanding of the issues arising from the above questions can lead to appropriate control decisions at the MAC layer that can make a significant difference in terms of MAC efficiency, reliability, and throughput; all of which could be vitally important for healthcare. Knowledge of the channel characteristics can help us make smart decisions. For example, immediately attempting to retransmit a failed packet should be avoided since the channel is likely to still be in an outage (the most probable cause of failure).

Here we present four MAC techniques that can be used to increase IEEE 802.15.6 system reliability while achieving high energy efficiency. First, we consider dynamic allocation of slots in an effort to avoid outages. This technique increases reliability without increasing energy consumption. The second technique further increases reliability by allowing retransmissions, thus spending additional energy, but builds on the dynamic allocation idea by achieving the best improvement in reliability for any given limit on the additional energy that may be spent. The third technique concerns outages that last too long to be remedied with retransmissions by employing relay nodes. Finally, the fourth technique explores the potential of transmission power control and can be applied concurrently with the previously mentioned techniques. The performance of these techniques have been verified through simulations that make use of our real-world "everyday" channel sounder data to model channel behavior.

Dynamic Slot Allocation

As mentioned earlier, the outages experienced in BANs are a major source of packet losses. Indeed, when using contention-free access, packets are lost when the link between a sensor and hub goes into an outage. Not only are packets lost; energy is wasted, and a time slot that could have been used productively by a different sensor node is also squandered.

Ideally, sensors would only be allocated time slots when their links are good. This might be impossible to achieve practically, but one can do
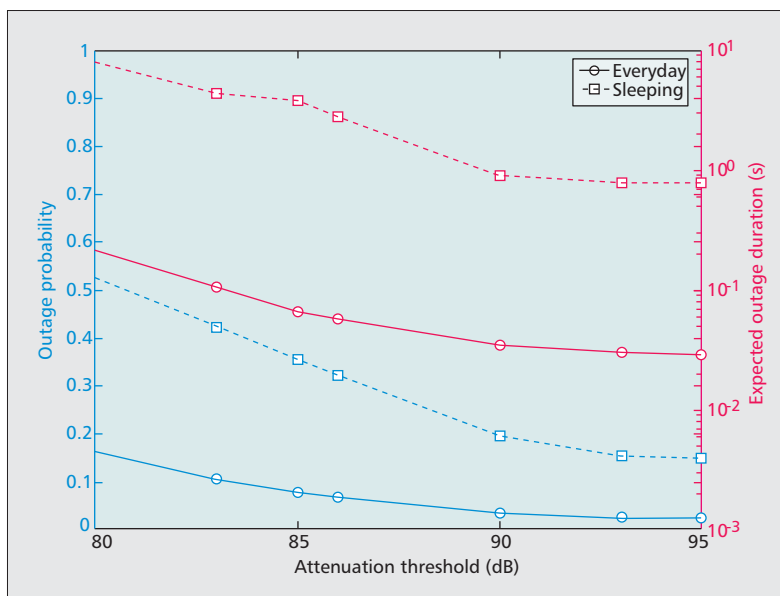

Figure 3. *Probability and expected duration of outages given the channel attenuation threshold that a wireless system can support.*

better than either statically or randomly assigning slots by better understanding the characteristics of the BAN channel. By understanding that the channel is slow-moving, and that outages last in the order of tens of milliseconds to seconds, this implies that whatever information one has about the state of the link will likely hold true in the near future. For example, knowing that a transmission attempt failed, or barely succeeded, implies that transmitting a packet straight away may not be worthwhile.

Knowing something about the future state of the link gives rise to the idea of using a variable slot allocation scheme to increase the reliability of the network. This idea is similar to opportunistic scheduling in cellular networks, where the base station dynamically selects a user to transmit in the next available time slot. However, traditional opportunistic scheduling approaches are not compatible with BANs as they require that the slave nodes are *continuously* available for communication. This contradicts the energy-saving mechanisms of BAN devices, which are in sleep mode most of the time, only waking up to listen for beacons and transmit data. However, since sensor nodes that want to send data must wake up to listen for the beacon in the superframe they want to transmit in, variable slot allocation can work at the superframe level of granularity. The hub can use link information gathered from previous superframes to schedule slot allocation for the next superframe, conveying this information to the sensor nodes at the start of the superframe.

The authors of this article have investigated such a variable scheduling scheme in BANs, assuming a two-state Markov model for the link state (good/bad) [13]. Each link, or channel between a sensor node and the hub, is characterized by two parameters, the "stability" S, the long-term average probability of the link to be in the good state; and the "volatility" V, a measure of how frequently the link transitions between

good and bad states. The values of these two parameters have been extracted from a large volume of our experimental channel sounder data [11].

While there are many scheduling techniques that could be applied to such a model, we have shown [13] that a simple "Flipping" algorithm can provide close to optimal performance, while not requiring knowledge of S and V for each link. The Flipping algorithm schedules all bad links of the previous superframe last, preserving the order in time in which they were observed, while scheduling all good links first, but reversing the order in which they were observed in the previous superframe. The rationale behind Flipping is that all bad links are given the longest time to recover (i.e., get out of the outage), while the flipped ordering of the good links takes advantage of the most recently observed good links to ensure high probability of success in the next round of transmissions.
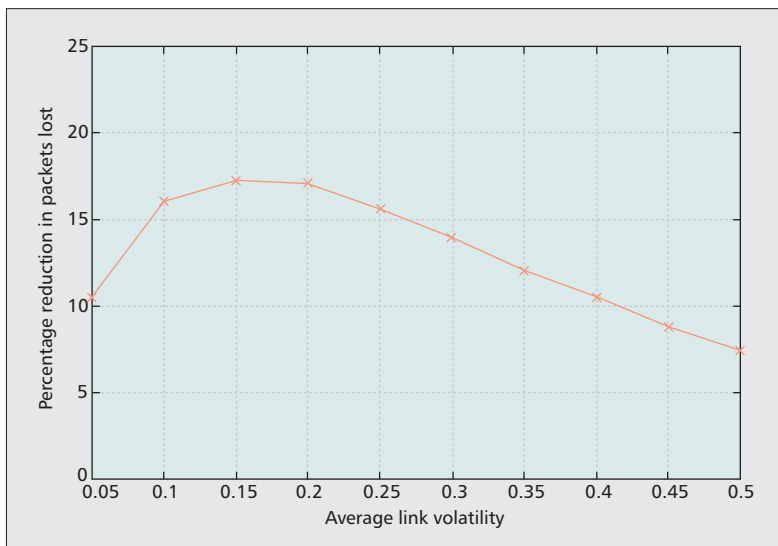

Figure 4. *Reduction in packets lost from Flipping, compared to static scheduling.*
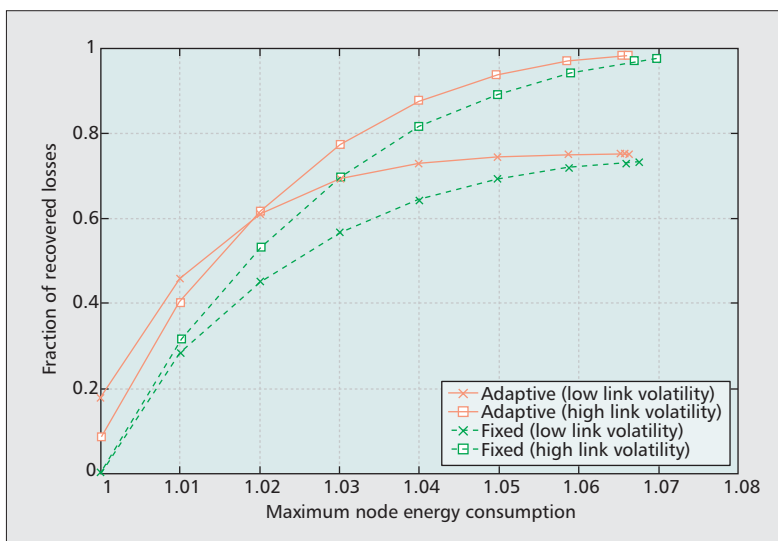
Figure 4 presents simulation results for an 8-node network; details on the setup and assumptions can be found in [13]. The figure shows the percentage reduction in packets lost when comparing the Flipping algorithm to static scheduling against average link volatility. The Flipping scheme can reduce packet loss by up to 17 percent compared to static allocation, and does not expend any additional energy.

SCHEDULING OF RETRANSMISSIONS

Assume an "everyday" activity BAN link with an average successful packet reception rate of 96 percent, as is typical for systems with an attenuation threshold of 90 dB (Fig. 3). A variable scheduling technique, such as Flipping, could improve the successful receptions to 96.7 percent by reducing packet losses by 17 percent. While this is a modest gain, it is achieved without using more energy. However, a success rate close to 100 percent, required by many healthcare applications, cannot be achieved without retransmission of failed packets. Retransmissions have to be considered together with the extra energy they consume, and the delivery deadlines imposed on the data (e.g., 125–250 ms is typical for healthcare applications). The issue of additional energy consumption is usually handled by imposing an energy budget that limits the extra energy that nodes may spend on retransmissions.

The BAN channel greatly influences the effectiveness of retransmissions; as explained earlier, retransmissions should be delayed until the channel is more likely to have returned from a state of outage. One method of delaying retransmissions is to allocate retransmission slots at the end of each superframe. This kind of fixed allocation scheme yields improvements in reliability, but more efficient schemes are possible. We have devised an adaptive retransmission technique [14] that takes into account the energy budget and adaptively allocates retransmission slots to anywhere in the superframe, according to channel conditions.

Figure 5 presents simulation results for an 8-node network; details on the setup and assumptions can be found in [14]. The figure shows the fraction of recovered packets that the adaptive and fixed retransmission schemes achieve, plotted against a range of energy budgets. The x-axis represents the maximum energy consumed by a single node in the network, a quantity limited by the energy budget. The baseline of energy consumption is 1, and represents the energy used by "regular" traffic. Numbers greater than 1 represent the additional energy used by retransmissions. For example, 1.04 means that an additional 4 percent of the baseline energy was used to retransmit failed packets. As expected, more packet losses can be recovered by spending more energy in retransmissions, albeit with diminishing returns. The adaptive retransmission scheduling outperforms the fixed one, especially when the energy budget is small. In other words, it pays to use smart scheduling especially when we do not have a lot of extra energy to spare. Also notice that retransmissions can recover many more packets in highly volatile channels as there is a greater chance of the links transitioning from bad to good state within the limited delivery deadline.

Figure 5. *Fraction of packets recovered by adaptive and fixed retransmission schemes.*

USING RELAYS

Retransmission schemes are only effective when the duration of the outages is significantly shorter than the packet delivery deadline, but this is not always the case. For example, because people move less while sleeping in bed, we have found that such "sleeping channels" can experience long outages, often minutes or tens of minutes at a time. The only alternative here is to route around the bad links via a relay node, which is an option supported by the IEEE 802.15.6 standard.[3]

There is a cost, however, to the energy consumption of a relay node as it must remain awake, listening for packets to relay to the hub from the sensor that cannot reach the hub directly. Minimizing energy consumption is critical in healthcare BANs, so relays should be used as efficiently as possible. Such a pursuit raises many interesting questions, for example: How does the hub decide that a node needs a relay (as opposed to attempting retransmissions)? Which node should act as a relay?

Our experiments are beginning to answer some of these questions. For example, we have considered sleeping subjects wearing on-body sensors and being monitored by an off-body device that is acting as the BAN hub [10]. This would be a common healthcare scenario where a monitoring device is placed beside the head of the bed. If we consider channel attenuation of more than 90 dB resulting in outage, there is a node that could act as a successful relay 85 percent of the time the sensor node (Tx) to hub link (Rx) is in outage. If, instead, the hub is placed on the subject's left hip, there is a viable relay 80 percent of the time that the direct link is in outage considering an attenuation threshold of 90 dB, as illustrated in Fig. 6. Considering that our measurements show that long outages occur for about 15 percent of the time that subjects are sleeping, it is evident that relays can play a significant role in improving BAN reliability.

CONTROLLING THE TRANSMIT POWER

It is possible to both improve transmission reliability and energy efficiency and reduce total energy consumption by controlling the transmit power level for each node in a network. Such power control can be applied concurrently to the three previous MAC techniques. This is especially effective in BANs as the slowly changing channel exhibits temporal correlation for up to 1 s into the future, making it relatively easy to predict over short periods of time. If one assumes that a sensor transmits a single packet every superframe, which typically lasts on the order of tens to hundreds of milliseconds, a simple sample-and-hold transmit power scheme can be effective. In such a scheme, the transmitting device estimates channel attenuation from the last packet that it has received and then adjusts its transmit power accordingly. Given the reciprocity of BAN wireless channels [11], then the device could use this predicted transmit power level for up to 1 s into the future.

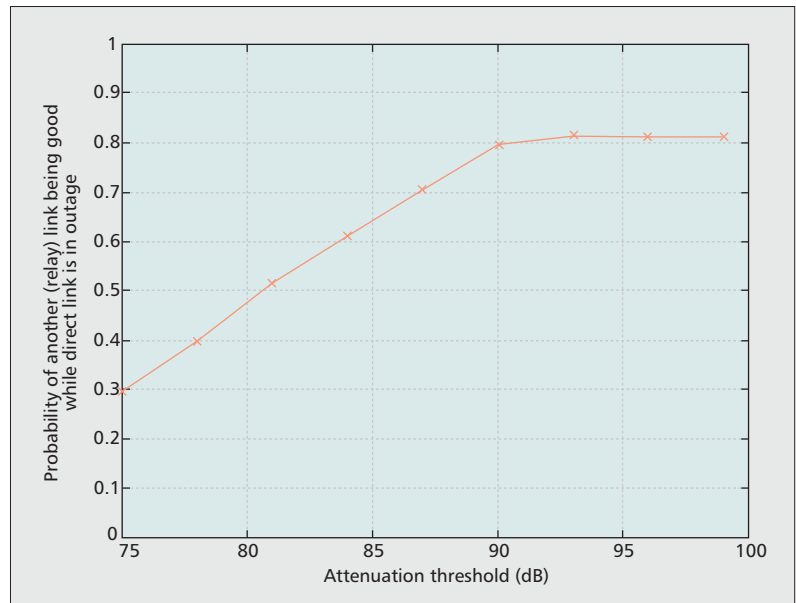Compared to using a fixed transmission power of 0.1 mW, our simulations over "every-

Figure 6. *Empirical probability of there being another good (relay) link while the direct link is in outage for a subject sleeping and the hub is situated at the front of the left hip [10].*

day" channel sounder measurements have shown a two- to fivefold improvement in packet loss by using this simple sample-and-hold scheme [15]. In addition to the improvement in reliability, the transmitter also consumes 15 percent less energy, more if one were to factor in the energy saved from not needing to retransmit as many failed packets. Of course, constantly transmitting at the maximum allowed power of 1 mW in IEEE 802.15.6 will increase reliability, but it will be at the cost of greatly increased energy consumption. Transmit power control can achieve similar reliability results to a 1 mW transmitter, while still reducing power consumption compared to a moderately powerful 0.1 mW transmitter.

CONCLUDING REMARKS

A number of challenges for the operation of the MAC in the context of healthcare and the current IEEE 802.15.6 draft standard for BANs have been described. We have described particular means to make best use of the provisions for the MAC in the current standard, including intelligently dealing with outages; retransmitting wisely; using relays in particularly difficult scenarios, such as people sleeping; and transmit power control. Of course, there are many other techniques that could enhance the standard, some of which will be driven by application-specific demands. Some techniques we are investigating include cross-layer implementation techniques, whereby information is shared between the physical and MAC layers to improve performance; and also methods of interference mitigation in the MAC for collocated networks.

It is clear that there are many exciting prospects for the further development of BAN technologies, particularly in the MAC layer, that will help fully exploit the potential of ubiquitous healthcare.

---

[3] We note that "sleeping channels" are not the only example of channels for which relays are useful; these may be useful in various "everyday" channels. However, sleeping channels are just an extreme vital use case for relays.

> It is clear that there are many exciting prospects for the further development of BAN technologies, particularly in the MAC layer, that will help fully exploit the potential of ubiquitous healthcare.

ACKNOWLEDGMENT

REFERENCES

[1] J. Penders *et al.*, "Human++: from Technology to Emerging Health Monitoring Concepts," *Int'l. Wksp. Wearable and Implantable Body Sensor Networks*, Hong Kong, China, June 2008,pp. 94–98.

[2] P. Khan, Md.A. Hussain, and K. S. Kwak, "Medical Applications of Wireless Body Area Networks," *Int'l. J. Digital Content Technology and its Applications*, vol. 3, no. 3, Sept. 2009.

[3] J. Penders *et al.*, "Potential and Challenges of Body Area Networks for Personal Health," *Engineering in Medicine and Biology Society*, *EMBC 2009, Annual Int'l. Conf. IEEE*, 2009, pp. 6569–72.

[4] S. Ullah *et al.*, "A Comprehensive Survey of Wireless Body Area Networks: On PHY, MAC, and Network Layers Solutions," *J. Medical Systems* (Springer), Aug. 2010.

[5] H. Cao *et al.*, "Enabling Technologies for Wireless Body Area Networks: A Survey and Outlook," *IEEE Commun. Mag.*, vol. 47, no. 12, Dec. 2009, pp. 84–93.

[6] P. S. Hall *et al.*, "Antennas and Propagation for On-Body Communication Systems," *IEEE Antennas and Propagation Mag.*, vol. 49, no. 3, June 2007, pp. 41–58.

[7] R. C. Shah and M. D. Yarvis, "Characteristics of On-Body 802.15.4 Networks," *WiMesh*, Reston, VA, Sept. 2006, pp. 138–39.

[8] A. Koubâa *et al.*, "TDBS: A Time Division Beacon Scheduling Mechanism for ZigBee Cluster-Tree Wireless Sensor Networks," Springer, Real-Time Systems, vol. 40, no. 3, Dec. 2008, pp. 321–54.

[9] A. Astrin *et al.*, "Tg6 draft," IEEE802.15.6 Technical Contribution, document ID:15-10-0245-040245-06-0006-tg6-draft, May 2010.

[10] D. Smith, D. Miniutti, and L. Hanlen, "Characterization of the Body-Area Propagation Channel for Monitoring a Subject Sleeping," *IEEE Trans. Antennas and Propagation*, vol. 59, no. 11, Nov. 2011, pp. 4388–92.

[11] L. Hanlen *et al.*, "Open-Source Testbed for Body Area Networks: 200 Sample/Sec, 12 hrs Continuous Measurement," *Pers. Indoor and Mobile Radio Commun.*, Istanbul Turkey, Sept. 2010.

[12] D. Smith *et al.*, "First and Second-Order Statistical Characterizations of the Dynamic Body-Area Propagation Channel of Various Bandwidths," *Annals of Telecommunications*, vol. 66, nos. 3–4, Dec. 2010, pp. 187–203.

[13] Y. Tselishchev, L. Libman, and A. Boulis, "Reducing Transmission Losses in Body Area Networks using Variable TDMA Scheduling," *12th IEEE Int'l. Symp. A World of Wireless Mobile and Multimedia Networks (WoWMoM 2011)*, Lucca, Italy, June 2011.

[14] Y. Tselishchev, L. Libman, and A. Boulis, "Energy-Efficient Retransmission Strategies Under Variable TDMA Scheduling in Body Area Networks," *36th IEEE Conf. Local Computer Networks (LCN 2011)*, Bonn, Germany, Oct. 2011.

[15] D. Smith *et al.*, "Simple Prediction-Based Power Control for the On-Body Area Communications Channel," *2011 IEEE ICC*, 5–9 June 2011, pp. 1–5.

BIOGRAPHIES

ATHANASSIOS BOULIS (athanassios.boulis@nicta.com.au) is a senior researcher at National ICT Australia (NICTA) since 2004. He received his B.E. degree in electronic and computer engineering from the Technical University of Crete, Greece, in 1997, and his M.Sc. and Ph.D. degrees in electrical engineering from the University of California, Los Angeles, in 1999 and 2003, respectively. His main research interests are in low-power wireless systems, MAC protocol design and implementation, and sensor network applications. He has several publications on body area networks, and has contributed to the IEEE 802.15.6 body area network standardization activities. He is also involved with developing sensor networks to monitor the structural health of bridges. He is the creator of the popular open-source Castalia simulator for low-power wireless networks.

DAVID SMITH is a senior researcher at NICTA and an adjunct fellow with the Australian National University (ANU) since 2004. He received his B.E. degree in electrical engineering from the University of N.S.W. Australia in 1997. He obtained an M.E. (research) degree in 2001 and a Ph.D. in 2004 both from the University of Technology, Sydney (UTS), both in the area of wireless communications engineering. His research interests are mainly in wireless communications including technology and systems for wireless body area networks; game theory for distributed networks; radio propagation and electromagnetic modeling; MIMO systems; coherent and non-coherent space-time coding; and antenna design, including the design of smart antennas. He has also had a variety of industry experience in telecommunications planning; radio frequency, optoelectronic and electronic communications design and integration. He has published numerous technical articles, has received a best paper award for research on body area networks at an international conference, and has contributed to IEEE 802.15.6 body area network standardization activities.

DINO MINIUTTI is a senior researcher at NICTA and an adjunct fellow with ANU. He received his Bachelor of Economics and Bachelor of Engineering from ANU in 1998 and 1999, respectively. He received a Ph.D. from ANU in 2005, with his thesis on the topic of low-complexity timing recovery in time-dispersive channels. His research interests are in digital signal processing, with wireless communication the main application. His current research is involved in body area networking, disaster recovery, mesh networking, and other forms of wireless communication where reliability is paramount. He has had a variety of experience, including starting a video-on-demand company to being a voting member for the IEEE 802.15.6 and IEEE 802.15.3c standards.

LAVY LIBMAN [S'97, M'04, SM'08] received his B.Sc. degrees in electrical engineering and computer engineering, and his M.Sc. and Ph.D. degrees in electrical engineering from the Technion — Israel Institute of Technology in 1992, 1997, and 2003, respectively. In 2003, he joined the Networks research group at NICTA, and between 2009–2011 he was a senior lecturer in the School of Information Technologies, University of Sydney. He is an Associate Editor of *IEEE Transactions on Wireless Communications*. His main research interests are in cooperative communication, network coding, and game-theoretic modeling and protocol design in wireless networks.

YURIY TSELISHCHEV is a Ph.D. student at the University of Sydney, School of IT, in partnership with NICTA, Australia's Information and Communications Technology Research Centre of Excellence. He received his B.S. degree in information technologies from the University of Sydney, Australia, in 2007. His research interests are focused on low-power sensor networks, and specifically on body area networks, targeting reliability and energy efficiency issues.