

Transmission Power Control in Body Area Sensor Networks for Healthcare Monitoring

Shuo Xiao, Ashay Dhamdhere, Vijay Sivaraman, and Alison Burdett

Abstract—This paper investigates the opportunities and challenges in the use of dynamic radio transmit power control for prolonging the lifetime of body-wearable sensor devices used in continuous health monitoring. We first present extensive empirical evidence that the wireless link quality can change rapidly in body area networks, and a fixed transmit power results in either wasted energy (when the link is good) or low reliability (when the link is bad). We quantify the potential gains of dynamic power control in body-worn devices by benchmarking off-line the energy savings achievable for a given level of reliability. We then propose a class of schemes feasible for practical implementation that adapt transmit power in real-time based on feedback information from the receiver. We profile their performance against the off-line benchmark, and provide guidelines on how the parameters can be tuned to achieve the desired trade-off between energy savings and reliability within the chosen operating environment. Finally, we implement and profile our scheme on a MicaZ mote based platform, and also report preliminary results from the ultra-low-power integrated healthcare monitoring platform we are developing at Toumaz Technology.

Index Terms—Body Area Sensor Network, Healthcare Monitoring, Transmission Power Control.

I. INTRODUCTION

AN ageing population, combined with sedentary lifestyle and poor diet, is resulting in an increasing number of people living for years or even decades with chronic conditions that require ongoing clinical management. This is putting huge cost pressures on healthcare today. Wireless sensor network technologies have the potential to offer large-scale and cost-effective solutions to this problem. Outfitting patients with tiny, wearable, vital-signs sensors would allow continuous monitoring by caregivers in hospitals and aged-care facilities, and long-term monitoring by individuals in their own homes.

To successfully deploy body area networks that can perform long-term and continuous healthcare monitoring, it is critical that the wearable devices be small and lightweight, lest they be too intrusive on patient lifestyle. This places fundamental limitations on the battery energy available to the device over its lifetime. Typical prototype devices in use today, such as MicaZ motes [1] used in Harvard’s CodeBlue [2] project, operate on a pair of AA batteries that provide a few Watt-hours (a few tens of kilo-Joules) of energy. Truly wearable health monitoring devices are emerging that have orders of magnitude lower

S. Xiao, A. Dhamdhere and V. Sivaraman are with the School of Electrical Engineering and Telecommunications, University of New South Wales, Sydney, Australia. E-mail: shuo.xiao@student.unsw.edu.au, {ashay, vijay}@unsw.edu.au. A. Burdett is with Toumaz Technology Limited, Abingdon, Oxfordshire, UK. E-mail: alison.burdett@toumaz.com

Manuscript received 15 Jan 2008, revised 26 Sep 2008. This work was supported by the Australian Research Council (ARC) under the Linkage Grant LP0776902.

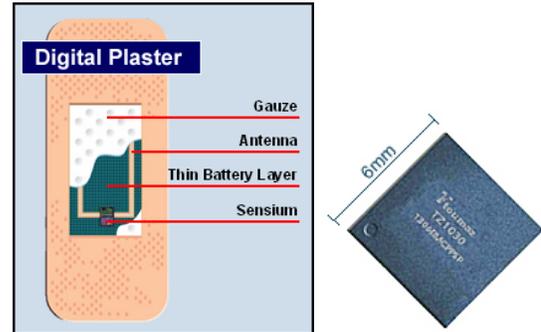


Fig. 1. Toumaz Sensium™ Digital Plaster

battery capacity – at Toumaz Technology we are building a new generation of single-chip low-cost disposable “digital plasters” (shown in Fig. 1) that provide non-intrusive ultra-low power monitoring of ECG, temperature, blood glucose and oxygen levels. Our Sensium™ chip operates on a flexible paper-thin printed battery [3] with a capacity of around 20 mWatt-hours (approximately 70 Joules). Such stringent energy constraints necessitate very careful energy management.

Communication is the most energy consuming operation that a sensor node performs [4], and can be optimised at multiple layers of the communication stack. At the physical layer, we at Toumaz have innovated an ultra-low-power radio [5] suited to body-area networks: our radio provides a proprietary 50 kbps wireless link over a distance of 2-10 metres, and consumes 2.7 mW at a transmit strength of -7 dBm (compare this to the CC2420 radio [6] in MicaZ motes that consumes 22.5 mW for -7 dBm output). At the data-link layer, energy can be saved by intelligent medium access control (MAC) protocols that duty-cycle the radio, i.e. by turning the radio off whenever packet transmission or receipt is not expected. Several such MAC protocols have been developed in the literature (see [7] for a survey). The B-MAC [8] protocol included in the TinyOS distribution provides versatility to the application in controlling the duty-cycling of the radio, while at Toumaz we have developed our proprietary MAC protocol [9] suited to body area networks. However, these MAC protocols only control *when* the radio is switched on, they do not determine the *output power* of the radio when it is on. The focus of this paper is to study the impact of transmission power (for any given transmission schedule) in trading off reliability of the time-varying wireless link for energy efficiency at the transmitting node. We note that the ability to control the transmission power is available on most platforms: the CC2420 radio in Crossbow’s MicaZ motes provides 32 transmission levels (ranging from -25 dBm to 0dBm output) selectable at run-time by configuring a register, while

our SensiumTM platform similarly supports 8 levels (ranging from -23dBm to -7dBm output).

Our **first** contribution presents extensive empirical trace data that profiles the temporal fluctuations in the body area wireless channel. The *large* variations show fixed transmit power to be sub-optimal: when link quality is poor, low transmit levels result in reduced reliability, whereas when link quality is good, high transmit levels waste energy. Furthermore, the *rapid* variations render existing schemes (discussed in detail in §II), that adapt transmit power over long time scales (hours and days), inappropriate for use in body area networks. Using trace data we compute off-line the “optimal” power control scheme, i.e. one that minimises energy usage subject to a given lower bound on reliability. Though infeasible to realise in practise, the optimal gives insight into the potential benefits and fundamental limitations of adaptive power control in body area networks, and also provides a benchmark against which practical schemes can be compared.

Our **second** contribution develops a class of practical on-line schemes that dynamically adapt transmission power based on receiver feedback. These schemes are easy to implement, and can be tuned for desired trade-off between energy savings and communication reliability. We show conservative, balanced, and aggressive adaptations of our scheme that progressively achieve higher energy savings (14-30%) in exchange for higher packet losses (up to 10%). We also provide guidelines on identifying algorithm parameter settings appropriate to application requirements and operating conditions.

As our **third** and final contribution, we present a real-time implementation of our power control scheme on a MicaZ mote based platform, demonstrating that energy savings are achievable even with imperfect feedback information. We also present preliminary observations of the efficacy of our scheme in the ultra-low power platform for continuous healthcare monitoring being developed at Toumaz Technology. Our work shows adaptive transmit power control as a low-cost way of extending the battery-life of severely energy constrained body wearable devices, and opens the doors to further optimisations customised for specific deployment scenarios.

The rest of this paper is organised as follows: §II briefly discusses prior work relevant to this paper. §III presents empirical observations on channel variability in body area networks and motivates dynamic power control. Optimal off-line power control is explored in §IV, while practical on-line schemes are proposed and analysed in §V. §VI describes our implementation and experiments on the MicaZ mote and Toumaz SensiumTM platforms, while conclusions and directions for future work are presented in §VII.

II. RELATED WORK

Transmit power control has been studied extensively in the literature in several contexts with different objectives. A large number of works e.g. [10]–[15] consider the IEEE 802.11 environment for wireless ad-hoc networks, and propose adjusting transmit power or transmit rate for data packets based on several factors such as wireless channel conditions (probed via RTS/CTS messages), payload length, etc. In spite of the useful insights from these works, there are significant

differences between the IEEE 802.11 WLAN environment and a body-area network (BAN) for healthcare monitoring such as: (a) WLAN devices send sporadic data while BAN devices typically send data periodically, (b) packet sizes are much smaller in BANs than in WLANs, and (c) BAN devices operate on or very near the human body which makes radio propagation (and hence channel conditions) in BANs markedly different from WLANs. These important differences merit study of power control in the specific context of body area networks, which to the best of our knowledge has not been undertaken to-date.

A large body of work in power control has also targeted multi-hop networks with the objective of enhancing throughput [16], increasing connectivity [17], [18] or reliability [19], and reducing delays [20]. Joint routing, scheduling, and power control schemes have also been proposed [21], [22]. Our work considers a single-hop body area network where such issues do not arise. Moreover, the environment presented by a body area network is much more dynamic than the pseudo-static scenarios considered by these works.

More relevant to our work in this paper are existing power control schemes that target energy savings in wireless sensor networks. The study in [23] proposes two algorithms that adapt transmission power by exchange of information among nodes and based on signal attenuation, while [24] proposes a linear prediction model for estimating the optimal transmission power based on measured link quality. However, these studies have targeted static deployments (such as for environmental or structural monitoring applications) wherein variability in wireless link quality has been shown empirically [25], [26] to be slow. In contrast, our work considers wearable mobile devices for which the wireless link quality can change significantly and rapidly since it is very susceptible to position and orientation of the human body [27]. To the best of our knowledge adaptive power control for body-wearable devices has not been studied before.

Two recent papers make interesting observations that are complementary to our work in this paper: the authors in [28] investigate the number of different power levels which can be effectively leveraged by power control algorithms in an indoor wireless LAN environment. They show that, due to multipath and fading effects, there is significant overlap between the RSSI distributions for nearby power levels, making them practically indistinguishable at the receiver, and claim that as few as 4 power levels may suffice to make power control attractive. We leverage this fact when we shift from the MicaZ (32 power levels) to the SensiumTM (8 power levels). Lastly, the study in [29] points out that power control by itself does not lead to significant energy savings unless complemented by a good MAC protocol that has low duty-cycling of the radio. We concur with this observation and note that the MAC protocol used in the SensiumTM platform has very low duty cycle.

III. THE CASE FOR TRANSMIT POWER CONTROL IN BODY AREA NETWORKS

We begin by empirically profiling the temporal variations in the quality of the wireless link between a body-worn device and a fixed base-station, as a patient wearing the device

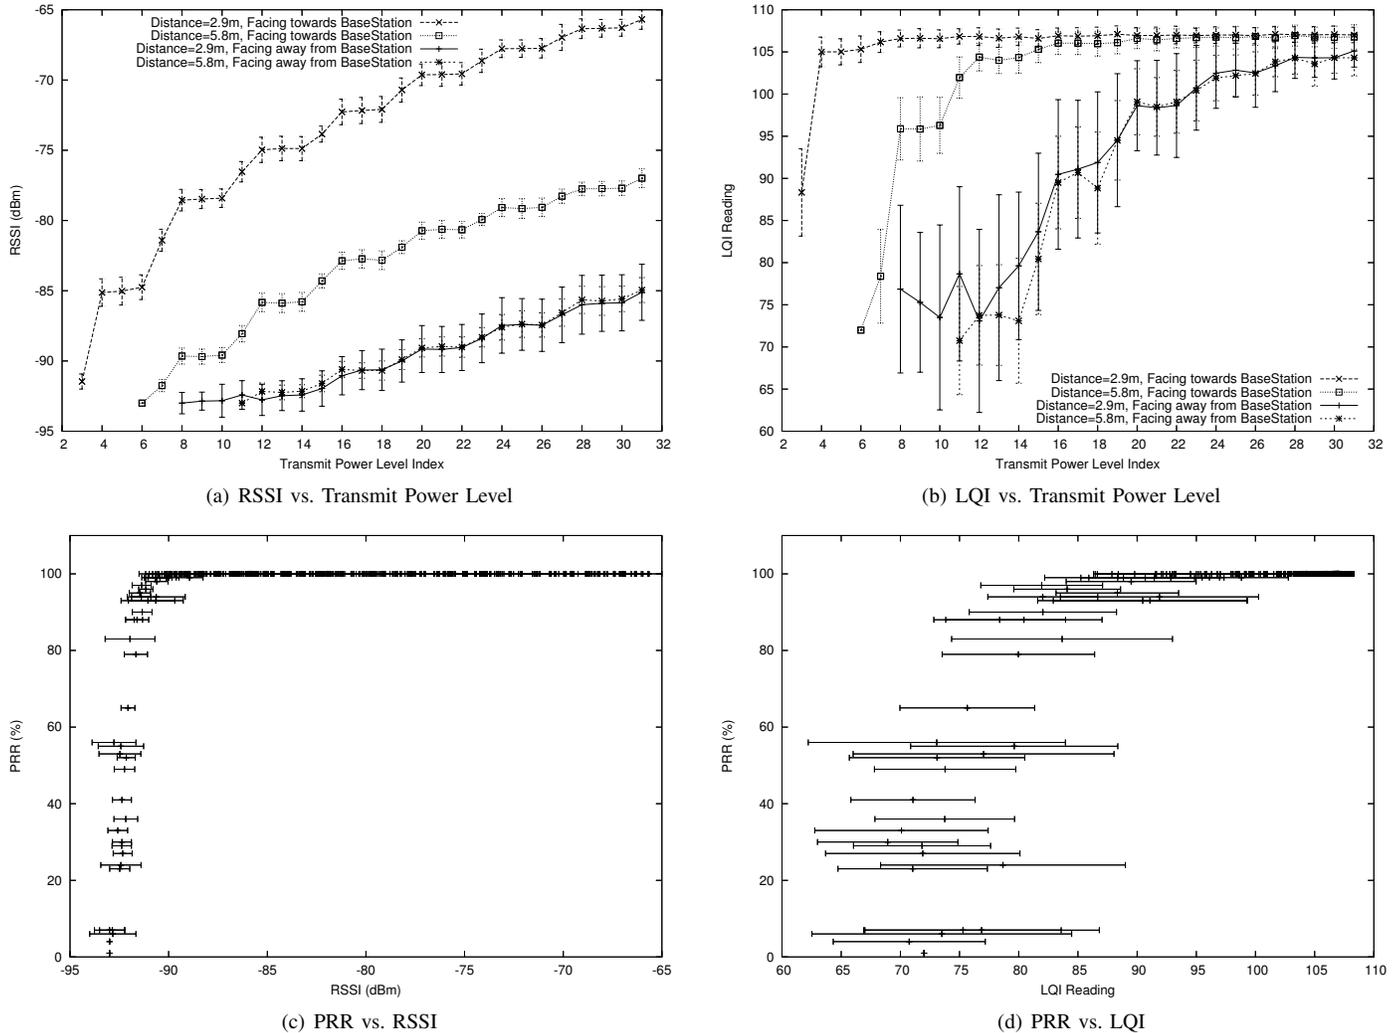


Fig. 2. Comparison of RSSI and LQI as indicators of channel quality in a body area network.

performs various activities. The patient was played by the first author. Our experiments in this section use the MicaZ motes from Crossbow Technologies [1], while some preliminary results that use the Toumaz SensiumTM platform are presented in §VI. In each experiment the device was strapped around the patient's chest, simulating continuous monitoring of heartbeat and ECG. We also conducted several experiments in which the device was strapped around the patient's arm (for monitoring blood pH and glucose); the results were qualitatively similar and are omitted here due to lack of space. The experiments were conducted indoors in an office space containing 10 cubicles. The base-station was placed close to one side of the room at an elevation (atop a shelf) to provide better line-of-sight coverage across the office space.

transmit level	output (dBm)	power (mW)
31	0	31.3
27	-1	29.7
23	-3	27.4
19	-5	25.0
15	-7	22.5
11	-10	20.2
7	-15	17.9
3	-25	15.3

TABLE I
CHARACTERISTICS OF THE MICAZ CC2420 RADIO

The MicaZ mote operates in the 2.4 GHz frequency band, and can support a 250 Kbps data rate. It supports 32 RF output power levels, controllable at run-time via a register; the output power (in dBm) and corresponding energy consumption rate (in mW) for various levels are shown in table I. Since our goal is to save energy at the body-worn device (which typically has lower energy resources than the base-station), our experiments involve emitting packets periodically from the body-worn device at various power levels, and measuring the link quality at the receiver (base-station). Two metrics of link quality are available on the mote platform: received signal strength indicator (RSSI), which is computed internally in the radio by averaging the signal power over eight symbol periods of the incoming packet, and link quality indicator (LQI) metric which measures the chip error rate for the first eight symbols of the incoming packet. For a static scenario, previous research [24], [30], [31] has observed LQI readings to suffer from early saturation and be relatively less stable. We conducted experiments to verify this in a body area network scenario. Fig. 2 (a)-(b) show the RSSI and LQI as a function of transmit power level for various scenarios of patient distance and orientation relative to the base-station. While the RSSI seems to show a smooth increase with transmit power, the

LQI saturates early on for several scenarios, and also seems to exhibit a large variance (shown as error-bars) which makes the readings less reliable. Fig. 2 (c)-(d) show the RSSI and LQI as a function of the packet receive ratio (PRR) observed across several experiments. While the RSSI seems to provide a good estimate of packet loss rates (e.g. RSSI of -90 dBm or larger always corresponds to PRR of 95% or more), the LQI is seen to be a much weaker indicator of PRR due to its high variance. Our subsequent study therefore uses RSSI as an indicator of channel quality.

We profile, at different radio transmit levels, the changes in link quality with time as patients perform their routine activities involving resting, moving, turning, etc. To compare link quality at different power levels, we should ideally take *simultaneous* measurements at all power levels, which is infeasible. As an approximation, we make the body-worn device transmit every packet multiple times in quick succession at sixteen different transmit levels 31, 29, 27, \dots 1. The receiver (base-station) can thus record, more-or-less simultaneously, the signal strength corresponding to each transmit level. Measurements for three scenarios are described next.

A. Normal Walk

This scenario has the patient walking back and forth in the room for a few minutes at a normal walking pace; the patient stays between 1 and 8 metres from the base-station at all times. The body device, strapped on to the patient's chest, generates a packet every second (which is not entirely unrealistic for a heartbeat/ECG monitor) and transmits it at 16 different output levels. The RSSI is recorded at the base-station for each packet at each transmit level, and plotted in Fig. 3(a) against time for four of the transmit levels. For any fixed transmit power, the received signal strength fluctuates widely: at fixed maximum transmit output (level 31 at 0dBm), the signal strength at the receiver changes from -64 dBm (at 34 sec) to -94 dBm (at 86 sec): a change of 30dBm under a minute. There are nevertheless some discernible trends: for example, in the interval 30-50 sec, the receive signal is consistently above -72 dBm (at the maximum transmit level) due to the clear line-of-sight presented by the patient walking towards the base-station, while the subsequent interval 50-70 sec exhibits RSSI below -75 dBm (again at the maximum transmit level) due to the patient turning and blocking the line-of-sight with his body. The question of whether these patterns present opportunities for energy savings by adapting transmit power will be tackled in §IV.

B. Slow Walk

In this scenario we consider a slowly moving person (such as an elderly or handicapped person with restricted mobility) who takes an exaggeratedly long time (over six minutes) to walk a distance of three metres and back. As before, packets are transmitted every second at several power levels, and the received signal strength at the base-station is recorded and depicted in Fig. 3(b). The trend in the plot is very evident: the RSSI is fairly low for the first half, when the patient's body blocks the line-of-sight between the body-worn device and the

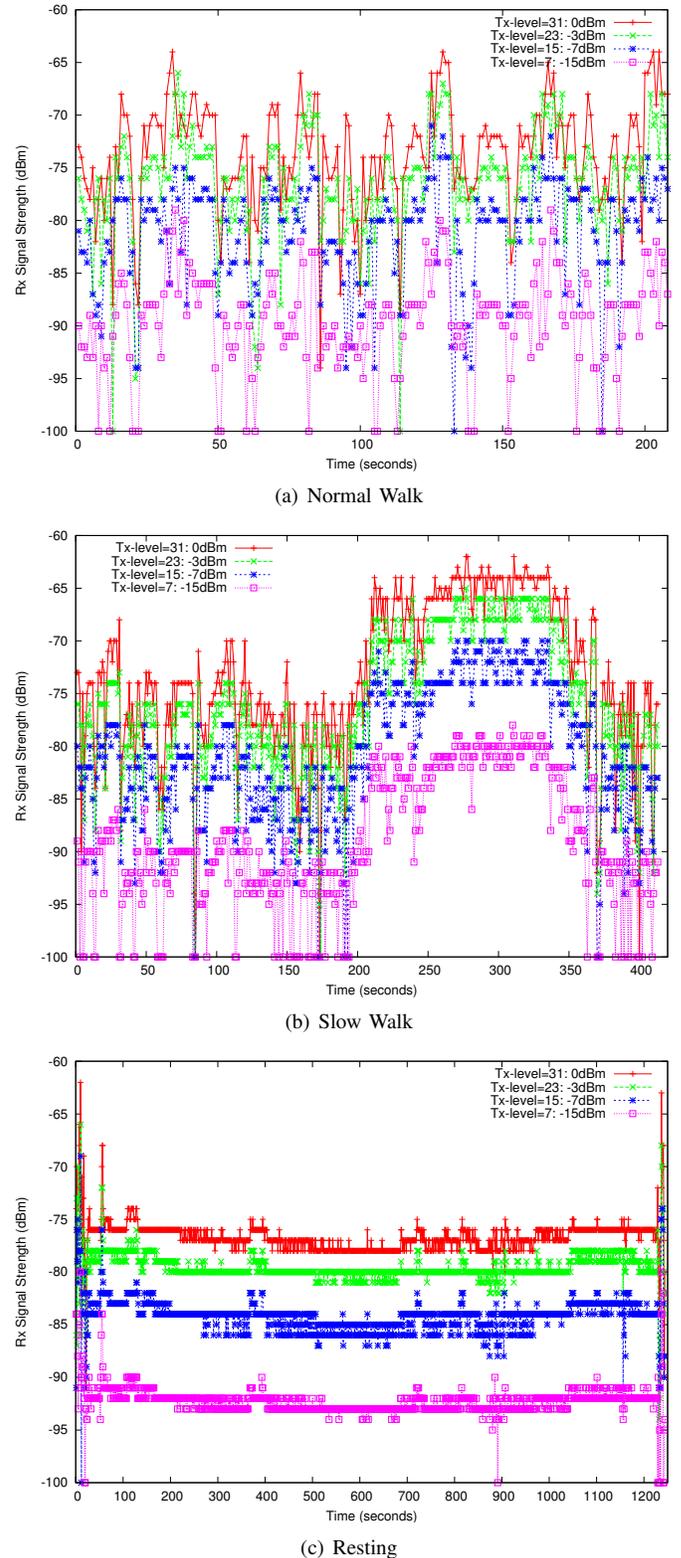


Fig. 3. RSSI vs. time for various patient scenarios.

base-station, and then rises to a perceptibly higher value in the second part of the experiment when the patient is walking facing the base-station (barring the last few seconds when the patient turns again). This scenario depicts the shortcomings of fixed transmit power: a low transmit level would result in weak signals (and packet loss) during the first half, while a high transmit level would unnecessarily waste energy in the

latter half. Such a scenario therefore presents opportunities for saving energy by reducing transmit power adaptively when the good channel persists.

C. Resting

In this scenario the patient sits down to rest for approximately 20 minutes on a chair at a distance of about 6 metres from the base-station. Fig. 3(c) plots the RSSI over the entire period, at several transmit levels. The wireless link is found to be fairly stable when the patient is at rest (in spite of a few other people moving around at several points in the experiment). This is in some sense an “ideal” environment with tremendous potential for energy savings, particularly with patients who are resting for a major part of the day. These energy savings would be unattainable if the transmit level were fixed, since a fixed setting would have to cater to the worst-case scenario of a poor channel.

Having gained an understanding of the wireless channel under various patient activity scenarios, the next section quantifies the potential benefits of adaptive transmit power control.

IV. OPTIMAL OFF-LINE TRANSMIT POWER CONTROL

To quantify the potential benefits of adaptive transmit power control, we compute what the “optimal” transmission level might be for each of the scenarios considered before. We define the “optimal” as the lowest required transmit power level (as a function of time) to achieve a minimum target RSSI. Based on our studies in §III relating packet loss with RSSI for the MicaZ motes in a body area network setting, we choose a conservative target RSSI of -85 dBm. The computation of the optimal transmission level defined thus is done off-line, i.e. using the traces shown in the previous section. For each scenario, at each time instant, we know the RSSI for each transmit power level, and we can therefore identify the lowest transmit power at which the signal strength at the receiver is no lower than the threshold of -85 dBm (if all received signal strengths are below the lower threshold we set the transmit level to be the maximum). We note that such a scheme is not implementable in practise, since it would require the transmitter to have instantaneous knowledge of the RSSI at the receiver for each choice of transmit power level, which is infeasible given that the channel varies with time.

The optimal transmit levels, and their associated RSSI values, for each of the three scenarios, are depicted as a function of time in Fig. 4. Sub-plot (a) shows, for the normal walk scenario, that the optimal changes rapidly to track the rapid fluctuations in channel quality, thereby maintaining a fairly stable RSSI as shown in sub-plot (b): for example, in the time interval 50-75 sec, the optimal transmit level fluctuates multiple times between a high of 29 and a low of 9. Based on the energy draw for each transmit power level (shown in Table I), we can compute the energy savings of optimal power control to be around 34% as compared to using the maximum transmit power. However, as the rapid fluctuations in the optimal level indicates, a practical scheme is unlikely to be able to predict the current optimal transmit level based on prior channel quality.

The optimal transmit power for a slow walk in Fig. 4(c) shows high sensitivity to body orientation, even when the motion is very slow. The rapid changes during the first 200 sec arise from minor variations in the patient’s body orientation while blocking the line-of-sight between the body-worn device and the base-station (indeed a few packets are lost even at the highest transmit power). But when the patient turns (at approximately 200 sec), there is a clear line of sight, and the wireless link is relatively stable permitting the optimal transmit power to remain low for a considerable length of time (more than 2 minutes). This indicates that if the body orientation is favourable, periods of slow activity could be capitalised by a transmit control scheme to save energy without compromising reliability.

When the patient is resting, the link is fairly stable and the optimal transmit power level is near-constant as shown in Fig. 4(e), which permits an energy savings of over 38% compared to maximum transmit power. It would seem the quiescent wireless channel in this case gives ample opportunity for practical schemes to reduce transmit power without sacrificing reliability. The design of such schemes is discussed next.

V. PRACTICAL ON-LINE TRANSMIT POWER CONTROL

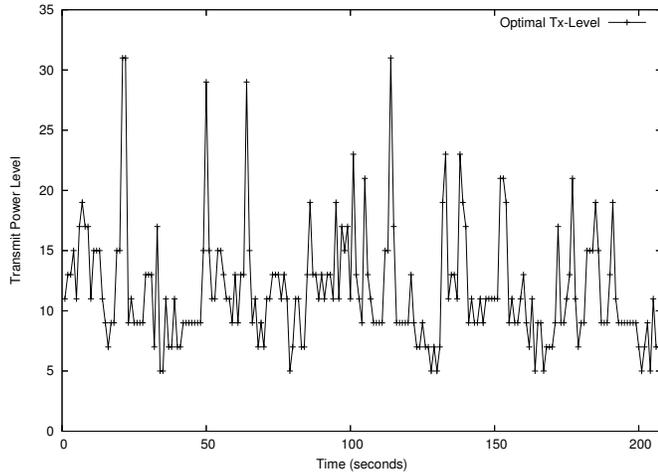
The optimal transmit power control scheme above was performed off-line and required the sender to have *a priori* knowledge of the link quality at the receiver, which is infeasible in reality. This section develops practical algorithms that are then benchmarked against the optimal.

A *predictive* approach to designing on-line schemes for power control would require a wireless channel model for body area networks. Modeling the propagation of electromagnetic waves around the human body, e.g. via “creeping waves” [32], is fairly complex as it needs to account for the permittivity and conductivity of the different layers of bone and tissue in the human anatomy. Additionally, the model has to contend with changes in orientation of the human body, mobility of the patient, and other spatio-temporal aspects (such as room layout, people in the vicinity, etc.). We believe such predictive models are too complex to implement on energy-constrained wearable devices, and do not pursue them in this paper.

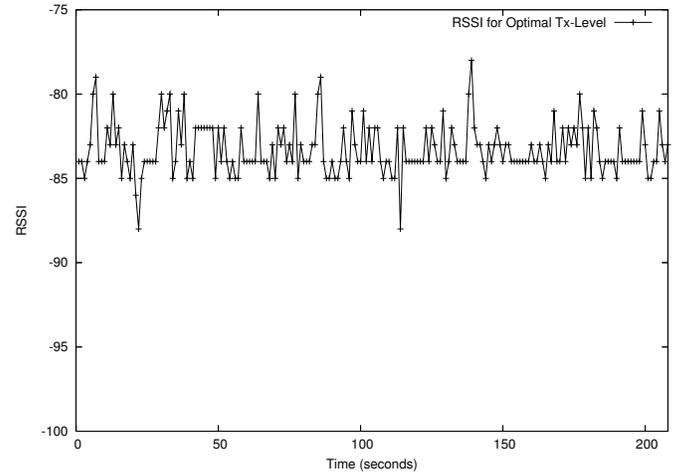
We focus instead on *reactive* schemes that adjust transmit power based on feedback from the receiver (inspired by the way TCP adjusts its transmission rate in reaction to congestion in the Internet). Specifically, the base station measures the RSSI for each received packet and feeds it back to the body-worn device in the acknowledgment packet. In this section we assume that the feedback information is perfect (i.e. acknowledgment packets are never lost); this assumption will be relaxed in §VI-A.

A. A Simple and Flexible Class of Schemes

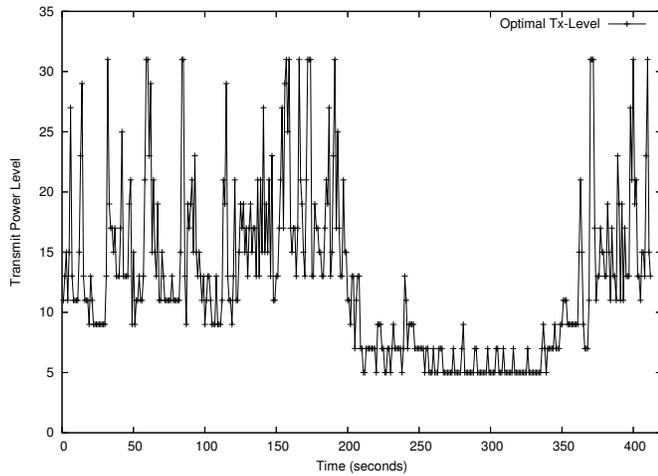
At its core, any reactive power control scheme must ramp up transmit power when the channel quality deteriorates (so as to avoid packet loss), and decrease transmit power when the channel quality improves (in order to save energy). We propose a general class of schemes that allow these operations to be tuned via appropriate parameter settings.



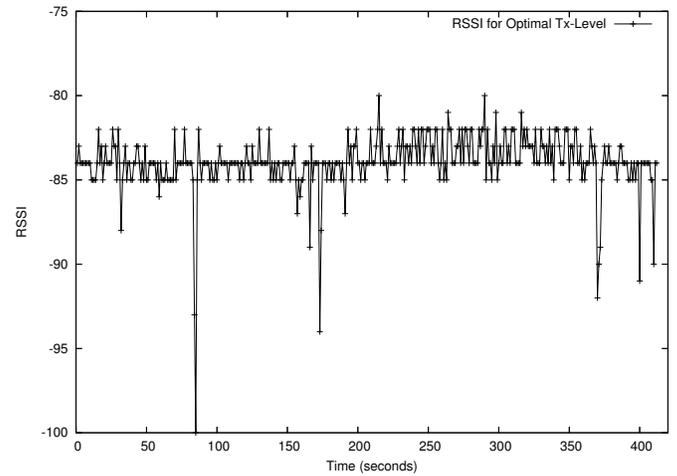
(a) Normal Walk: Optimal transmit level



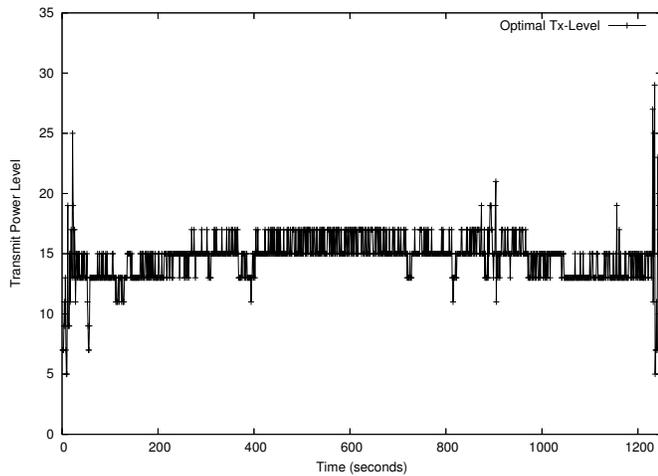
(b) Normal Walk: RSSI for optimal transmit level



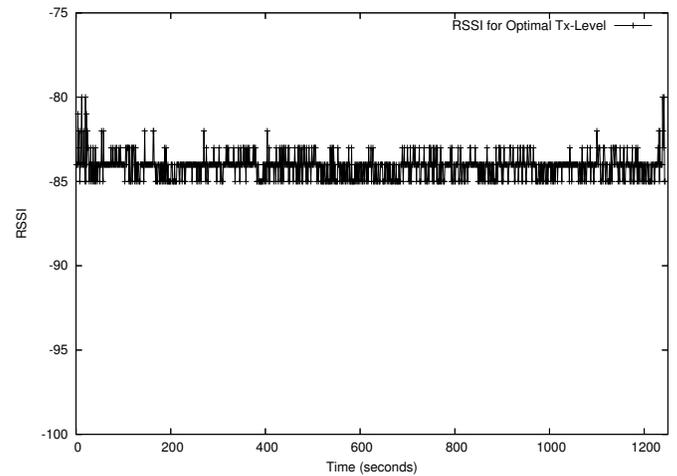
(c) Slow Walk: Optimal transmit level



(d) Slow Walk: RSSI for optimal transmit level



(e) Resting: Optimal transmit level



(f) Resting: RSSI for optimal transmit level

Fig. 4. Optimal transmit power and associated RSSI for a normal walk, slow walk, and resting position

Algorithm 1 depicts our class of schemes, and is characterised by four parameters: α_u , α_d , T_L and T_H . The scheme maintains a running average \bar{R} of the RSSI, computed by exponential weighted averaging (steps 2 and 4) of each newly obtained sample. The averaging weight α_u for a sample representing an improving channel can in general be different

from the weight α_d used for a sample representing a deteriorating channel; this gives flexibility to a scheme in reacting differently to a perceived increase or decrease in channel quality, thereby placing different emphasis on energy savings versus packet loss.

The scheme increases or reduces transmit power by compar-

Algorithm 1 A Class of Power Control Schemes

Require: R {RSSI from the current sample}
Require: \bar{R} {Average RSSI}

- 1: **if** $R \leq \bar{R}$ **then**
- 2: $\bar{R} \leftarrow \alpha_d R + (1 - \alpha_d) \bar{R}$
- 3: **else** $\{R > \bar{R}\}$
- 4: $\bar{R} \leftarrow \alpha_u R + (1 - \alpha_u) \bar{R}$
- 5: **end if**
- 6: **if** $\bar{R} < T_L$ **then**
- 7: Double the transmit power
- 8: **else if** $\bar{R} > T_H$ **then**
- 9: Reduce the transmit power by a constant
- 10: **else** $\{T_L \leq \bar{R} \leq T_H\}$
- 11: No action is required
- 12: **end if**

ing the running RSSI average \bar{R} to lower and upper thresholds T_L and T_H . Based on our previous analysis of RSSI and packet loss (see Fig. 2(c)), we believe a lower threshold $T_L = -85$ dbm is appropriate for the MicaZ platform. If \bar{R} falls below this threshold (step 6), the transmit power is immediately doubled (Step 7) to avoid imminent packet loss in the deteriorating channel. If, on the other hand, the average RSSI exceeds an upper threshold T_H (step 8), the transmit power is reduced by a small amount (step 9). Our experimental traces indicate that $T_H = -80$ dBm is appropriate for the MicaZ platform; lower values make the target RSSI range $[T_L, T_H]$ too narrow while larger values lead to higher transmit power (and hence higher energy usage) than required. Readers may note the similarity of our scheme to TCP in that the transmit power increase is multiplicative while the decrease is additive, and thus has a clear bias towards reacting faster to a deteriorating channel than an improving one.

The class of schemes outlined in algorithm 1 is fairly easy to implement in platforms with very limited CPU and memory resources. It is also very flexible, and the parameters α_u and α_d can in particular be tuned for appropriate trade-off between energy efficiency and reliability. The performance of the schemes for specific parameter settings is discussed next.

B. Example Adaptations of the General Scheme

A *conservative* approach to energy savings might be warranted in applications where data loss is critical. A relatively high value $\alpha_d = 0.8$ could be used so that a low RSSI value triggers a rapid ramp-up in transmit power, while keeping $\alpha_u = 0.2$ low so that transmit power is reduced cautiously when good channel conditions prevail. Applications in which energy is at a high premium and data loss is not as critical may adopt an *aggressive* strategy with a high $\alpha_u = 0.8$ that reacts quickly to improvements in channel conditions while setting $\alpha_d = 0.2$ so that a transient bad channel is ignored. A *balanced* approach that equally values energy savings and packet loss may place equal emphasis on whether the channel is getting better or worse by setting $\alpha_u = \alpha_d = \alpha = 0.8$.

We tested the efficacy of the conservative, aggressive and balanced schemes as described above on the trace data for the

Scheme	Normal Walk		Slow Walk		Resting	
	power (mW)	loss (%)	power (mW)	loss (%)	power (mW)	loss (%)
Maximum	31.32	0	31.32	1.2	31.32	0
Optimal	20.60	0	20.74	1.2	22.35	0
Conservative	26.90	1.4	25.52	2.18	24.36	0.08
Balanced	25.15	3.85	24.11	3.85	24.22	0.16
Aggressive	21.78	9.6	21.96	7.28	23.74	0.32

TABLE II
POWER DRAW AVERAGED ACROSS PACKETS, AND LOSS RATE, FOR VARIOUS POWER CONTROL SCHEMES

three scenarios described earlier. The transmitter is assumed to know the RSSI for each transmitted packet via feedback from the receiver; it performs the exponential averaging using the parameters listed for each scheme, and chooses an appropriate transmit power level for the subsequent packet transmission. Fig. 5 shows the transmit power level and the corresponding RSSI for each of the three schemes for each scenario. The average power draw per packet, as well as packet loss rates, for each of the schemes under the different scenarios, are summarised in Table II.

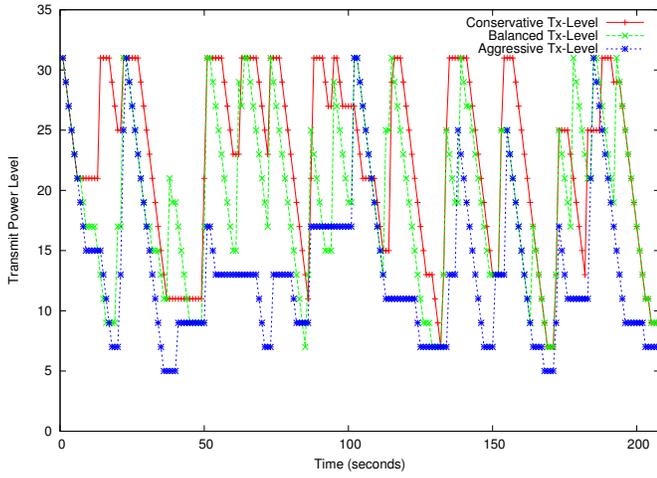
For the normal walk, Fig. 5(a) shows that all schemes exhibit considerable fluctuation in their transmit power, since the channel varies quite rapidly. The conservative scheme yields good link reliability (only 1.4% loss) but has high energy usage (30.58% above optimal), while the aggressive scheme yields good energy savings (5.73% within optimal) by sacrificing link quality (9.6% packet loss). As one might expect, the energy usage and packet loss under the balanced scheme are between those of the conservative and aggressive.

For the slow walk scenario, Fig. 5(c) shows the aggressive scheme to be fairly stable in the interval 210-350 sec when the patient is walking slowly facing the base station, whereas the conservative scheme shows some fluctuations in that region. Again, the conservative scheme uses 16.21% more energy than the aggressive scheme, but has a packet loss rate much lower than the aggressive scheme, showing that energy and reliability can be traded-off in our schemes by choosing parameters appropriately. As before the balanced scheme lies between the other two in both its metrics.

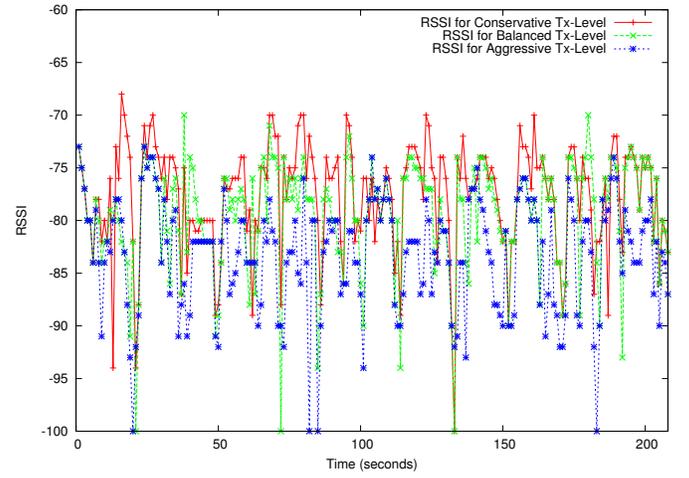
The transmit level under the three schemes for the resting scenario, shown in Fig. 5(e), is fairly stable, barring some rogue glitches in the link quality (e.g. at 905 sec) that the conservative scheme over-reacts to. The energy savings, as well as the packet loss ratios, are comparable under all three schemes, indicating that under quiescent channel conditions the parameter settings do not influence the algorithm performance significantly.

C. Tuning the Parameters

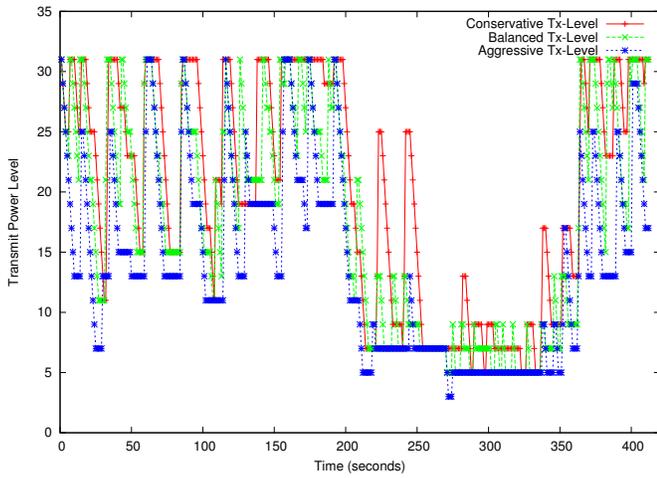
In this section we undertake a more detailed study of the impact of the algorithm parameters α_u and α_d on the performance of our class of schemes. We conducted several experiments in which the patient is fairly active, since the parameters have a larger impact on energy and loss under such scenarios. Fig. 6 plots the energy consumption (left column) and packet loss (right column) of our scheme for various parameters settings for a representative scenario. Fig. 6(a)-(b) show that for any fixed α_d , the energy consumption falls



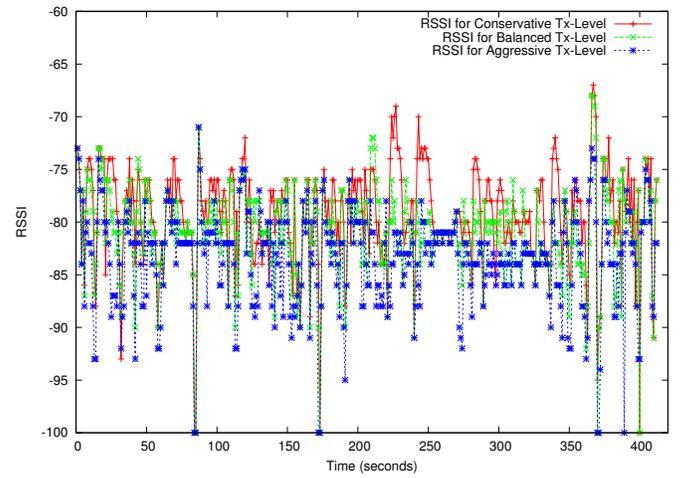
(a) Normal Walk: transmit level



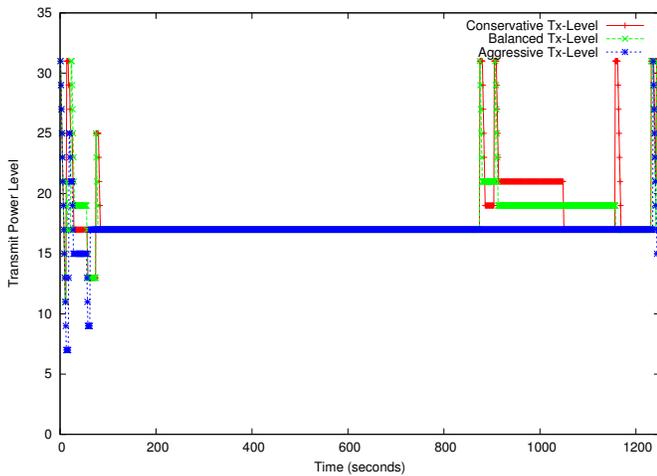
(b) Normal Walk: corresponding RSSI



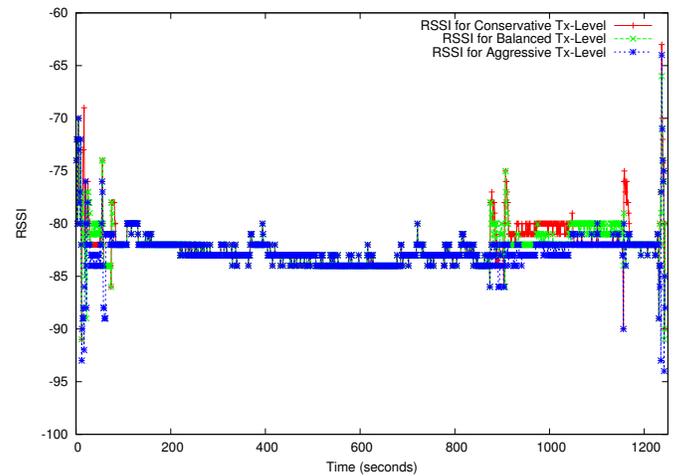
(c) Slow Walk: transmit level



(d) Slow Walk: corresponding RSSI



(e) Resting: transmit level



(f) Resting: corresponding RSSI

Fig. 5. Transmit power and RSSI under conservative, balanced, and aggressive schemes for various scenarios.

monotonically with α_u : this is because a larger α_u makes the scheme react faster to an improving channel, thereby reducing transmit power to save energy. When α_u is held constant and α_d increases, the scheme react more quickly to a degrading channel, and the packet loss rate diminishes (Fig. 6(d)) at the expense of increased energy consumption (Fig. 6(c)). The plots show clearly that increasing α_u pulls the algorithm in

the direction favouring energy savings, while increasing α_d pulls it towards reliability – these opposite trends justifying our decision to devise a scheme with two separate α values. For a balanced scheme that uses $\alpha_u = \alpha_d = \alpha$, Fig. 6 (e)-(f) show that a scheme that is very reactive to both good and bad channel conditions (i.e. high α) improves loss performance as expected, but incurs a penalty in terms of slightly increased

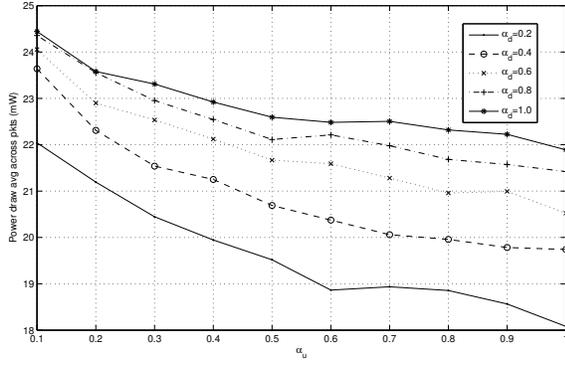
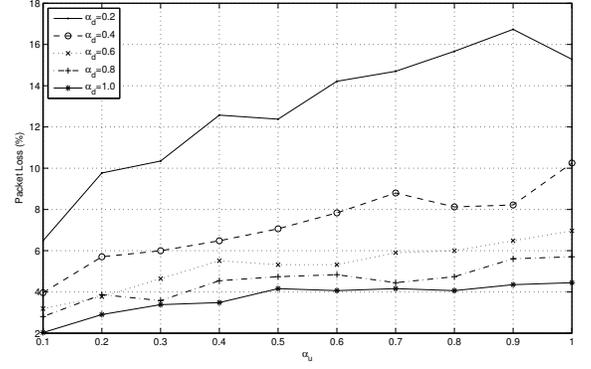
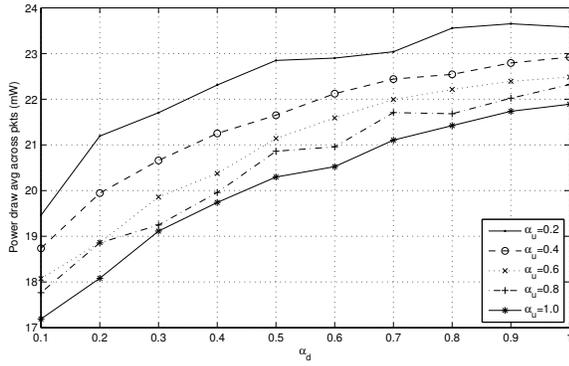
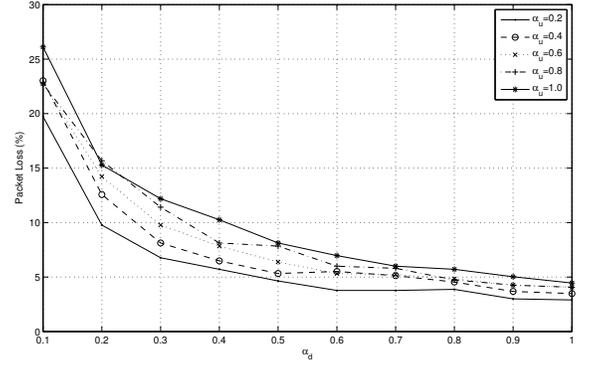
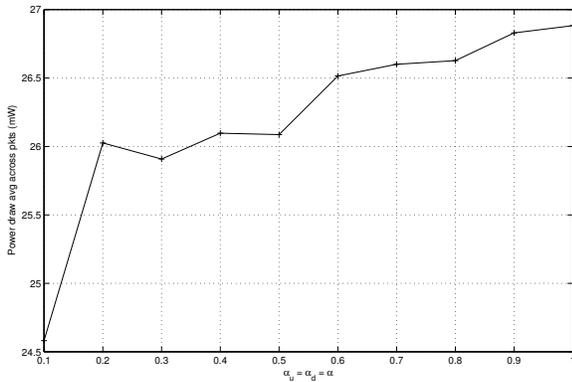
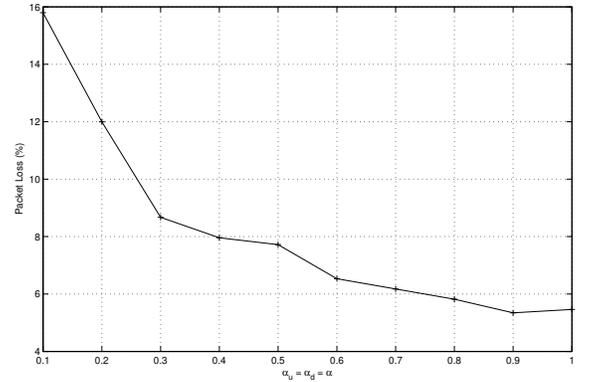
(a) Fixed α_d : Energy Consumption vs. α_u (b) Fixed α_d : Packet Loss vs. α_u (c) Fixed α_u : Energy Consumption vs. α_d (d) Fixed α_u : Packet Loss vs. α_d (e) $\alpha_u = \alpha_d = \alpha$: Energy Consumption vs. α (f) $\alpha_u = \alpha_d = \alpha$: Packet Loss vs. α

Fig. 6. Energy Consumption and Packet Loss for Various Parameter Settings

energy usage.

An interesting question concerns the optimisation of the parameters for a given operating scenario. For the trace considered in this section, we found that $(\alpha_u = 1.0, \alpha_d = 0.3)$ yields maximum energy savings when the loss budget is high at 15%, implying that the ramp down in transmit power should be immediate while a ramp up happens slowly over several poor RSSI samples. When the loss margin is stringent at 5%, energy savings are maximised for $(\alpha_u = 0.8, \alpha_d = 1.0)$, i.e. the scheme should react immediately to a degrading channel, and relatively slowly to an improving channel. Though the optimal settings obtained here do not transfer directly to other scenarios, they provide insight which could conceivably be

used by a system to adjust the parameters at run-time based on application requirements and operating conditions.

VI. PROTOTYPING AND EXPERIMENTATION

This section reports on a real-time implementation of our power control schemes on a MicaZ mote-based testbed, as well as preliminary experiments on the Toumaz SensiumTM platform.

A. MicaZ Mote Platform

The previous section evaluated the performance of the transmit power control schemes using trace data wherein the sender (body-worn device) is assumed to know the RSSI at

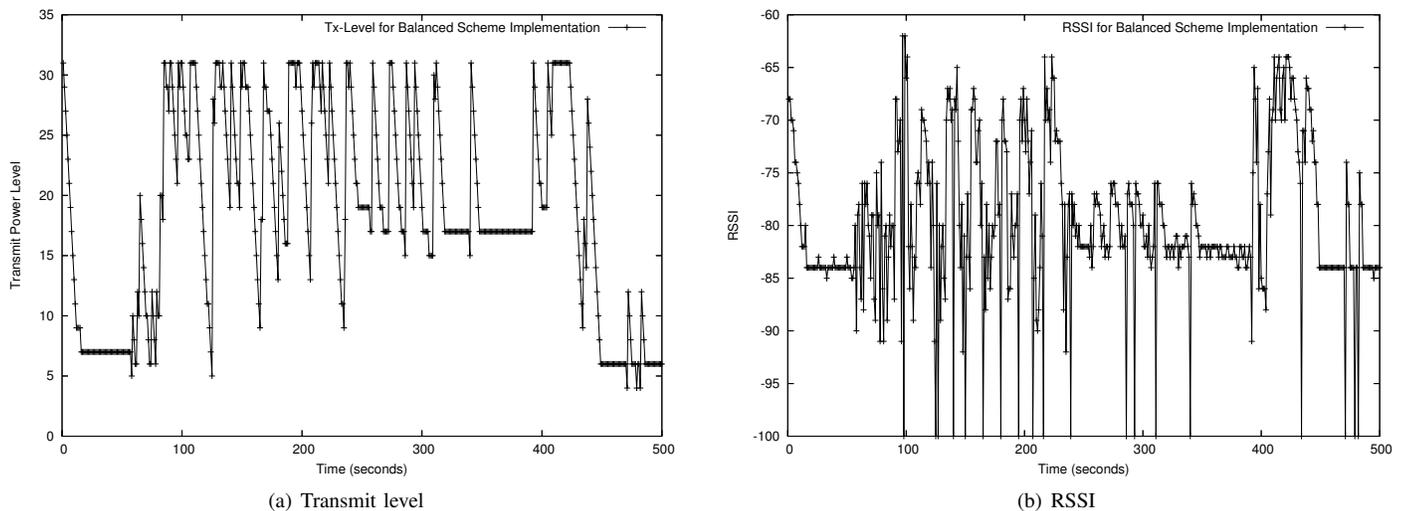


Fig. 7. Transmit power and associated RSSI for real-time implementation of Balanced scheme on the MicaZ platform

the receiver (base-station) for each of the packets it has thus far transmitted. In a real operational setting, such feedback information would be contained in acknowledgement packets (from the receiver to the sender) which may also experience loss. This section undertakes a real-time implementation of our scheme on body-worn MicaZ motes to evaluate the efficacy of our power control scheme under imperfect feedback. We note that the base-station is assumed to have abundant energy reserves and does not implement power control, always using the highest power level for transmitting acknowledgements.

The implementation performs exactly the steps shown in algorithm 1, with the additional step that if an acknowledgement is not received for a transmitted packet, the RSSI for that sample is taken to be -100 dBm (thereby signaling a bad channel to the algorithm). We present results with balanced parameter setting $\alpha_u = \alpha_d = \alpha = 0.8$ for a scenario where the patient undertakes a mix of walking and resting. Fig. 7(a) shows the transmit power level under our scheme and Fig. 7(b) the corresponding RSSI. Our scheme performs quite well, yielding average energy consumption rate per packet of 20.23 mW (a 35.4% savings in energy compared to using maximum transmit power), and packet loss rate of 3.8%. Unfortunately there is no way to benchmark this result (recall that computing the optimal requires a trace that includes the RSSI for all power levels at all times), but a careful look at the plots will show regions where the RSSI is high and yet the transmit power is not reduced, for example in the interval (410, 420) sec. It was found that several acknowledgement packets were lost in this interval (the overall acknowledgement loss rate for this scenario was 11.82%), and the absence of feedback information forced our scheme to use a higher transmit power level than would have been necessary with perfect feedback. High acknowledgement loss rates arising from asymmetric link qualities can have a negative impact on the efficacy of feedback-based power control schemes, and merit deeper study in future work.

B. Toumaz SensiumTM Platform

A major goal of this project is to evaluate the costs and benefits of adaptive power control in the real-world

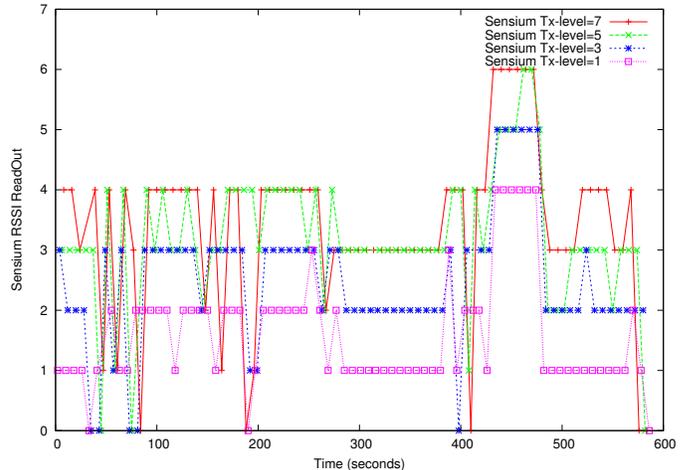


Fig. 8. RSSI vs. time using the SensiumTM platform

continuous health monitoring platform being developed by Toumaz Technology. Unfortunately the digital plaster (which includes the SensiumTM chip, printed battery, etched antenna, and water-protective covering) is still in the packaging process; so we present here preliminary results obtained from using a SensiumTM chip mounted on a development board strapped to the patient's chest, as preliminary indicators on the feasibility and benefits of adaptive transmit power control in the SensiumTM platform.

Tx level	output (dBm)	power (mW)	Rx level	input (dBm)
7	-6	2.8	7	> -35
6	-7	2.7	6	-46 to -35
5	-9	2.6	5	-52 to -46
4	-10	2.5	4	-58 to -52
3	-12	2.4	3	-64 to -58
2	-15	2.2	2	-70 to -64
1	-18	2.0	1	-76 to -70
0	-22	1.8	0	< -76 dBm

(a) Transmit Characteristics

(b) Receive Characteristics

TABLE III
SENSIUMTM RADIO TRANSMIT AND RECEIVE CHARACTERISTICS

The hardware-optimised design of the SensiumTM presented several constraints in our experimentation: (i) the radio has only 8 transmit power levels (unlike 32 in the MicaZ radio),

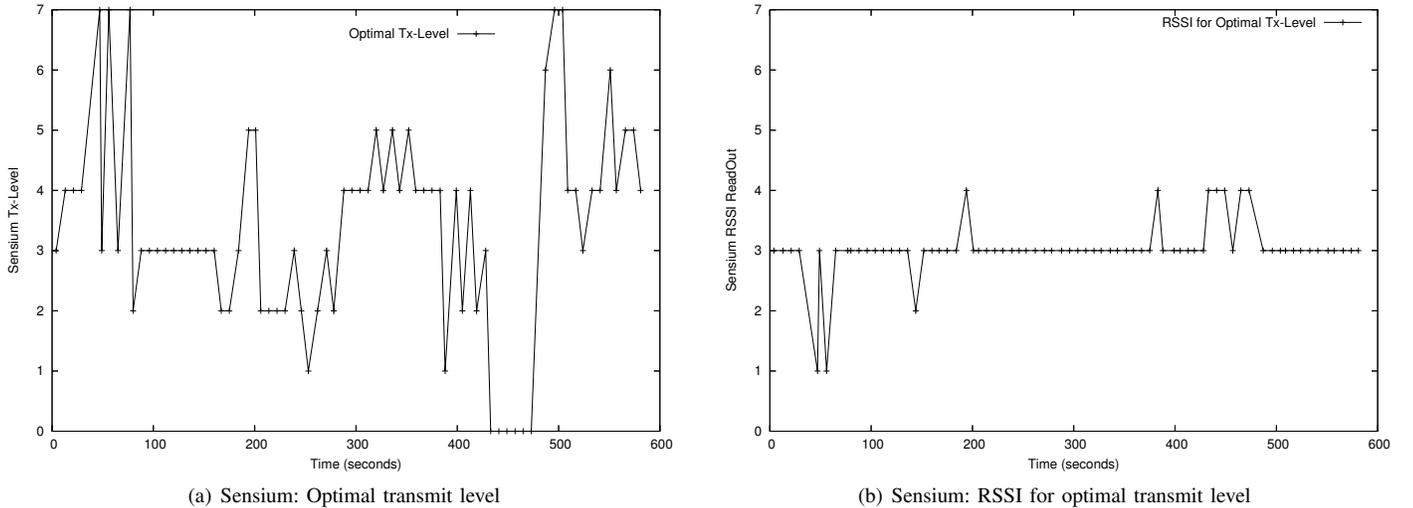


Fig. 9. Optimal transmit power and associated RSSI for the SensiumTM

and the energy savings are upper-bounded at 35% due to the limited transmission power range (see Table III(a)), (ii) the RSSI readout is only 3 bits (compared to 8 bits in the MicaZ), which gives only a coarse estimate of the RSSI (see Table III(b)), (iii) a sleep time of at least one second is imposed between successive packet transmissions (for energy efficiency), which restricts our ability to sample the channel at various transmit power levels without substantial change in patient position/orientation, (iv) the RSSI feedback from receiver to sender requires software modification of the acknowledgement packets which introduces a one second lag in the feedback. These limitations notwithstanding, we believe there is great value in experimenting with this real-world wearable device, and these limitations can be addressed in subsequent revisions of the SensiumTM hardware.

As with the motes, we strap the SensiumTM to the patient's chest, and make it transmit a packet every second, cycling through the the 8 available power levels, while the RSSI is recorded at the base-station. A ten-minute extract of the recorded RSSI at various transmit power levels is shown in Fig. 8: in this scenario, the patient is walking back-and-forth in interval 30-80 sec, during which period the RSSI fluctuates rapidly, while in interval 280-380 sec and 420-470 sec the patient is stationary (facing away and towards the base-station respectively) and the RSSI is stable. This confirms that the body area wireless channel presents significant temporal fluctuations in the 862-870 MHz frequency range (much like the 2.4 GHz range of the MicaZ radio), and fixed transmit power is sub-optimal.

For the above trace, we compute off-line the optimal transmission power schedule for desired RSSI level of 3 (corresponding to the range -64 to -58 dBm). Fig. 9 shows the optimal transmit power level (at the sender) and the corresponding RSSI level (at the receiver), and clearly depicts that the transmit power level can be reduced during quiescent periods such as 420-470 sec, while preserving reliability during periods when channel conditions are poor. For this scenario, optimal transmit power control uses 14.63% less energy than fixed maximum transmit power, which is a significant saving

for this severely energy-constrained device. Further results from our testing on the SensiumTM platform are reported in our subsequent work [34].

VII. CONCLUSIONS AND FUTURE WORK

In this paper we have investigated the potential benefits and limitations of adaptive radio transmit power control as a means of saving precious energy in body-wearable sensor devices used for medical monitoring. We experimentally profiled the radio channel quality under different scenarios of patient activity, and showed that fixed transmit power either wastes energy or sacrifices reliability. We then quantified the theoretical benefits of adaptive transmit power control, and showed that across different scenarios it can save nearly 35% energy without compromising reliability. We then developed a general class of practical power control schemes suitable for body area networks, and showed specific instances that save 14-30% energy (as compared to using maximum transmit power) in exchange for 1-10% packet losses. We demonstrated that the adjustment of parameters allow our schemes to achieve different trade-offs between energy savings and reliability, making them suitable across diverse applications in different operating conditions. Finally, we demonstrated that a real-time implementation of our scheme on the MicaZ mote based platform is effective even in the presence of imperfect feedback information, and presented preliminary experimental results indicating the potential for saving precious energy in Toumaz's real-world platform for continuous healthcare monitoring.

Our work represents a first step in understanding dynamic transmit power control in body area networks. There is much further study required in exploring its potential for specific health monitoring environments (e.g. critical care in hospitals, aged care, athlete monitoring, etc.) which have different characteristics in terms of patient mobility, periodicity, and criticality of collected data. Our work in [33] addresses the performance impact of sensor data periodicity on adaptive power control, and in our future work we intend to perform more extensive experimentation [34] with the Toumaz platform, possibly in deployments with real patients.

REFERENCES

- [1] Crossbow-Technologies, "Mica2 and MicaZ notes." [Online]. Available: <http://www.xbow.com>
- [2] V. Schnayder *et al.*, "Sensor Networks for Medical Care," Division of Engineering and Applied Science, Harvard University, Tech. Rep. TR-08-05, 2005.
- [3] Power-Paper, "Power Patch Platform." [Online]. Available: <http://www.powerpaper.com>
- [4] D. Culler, D. Estrin, and M. Srivastava, "Overview of Sensor Networks," *IEEE Computers*, vol. 37, no. 8, pp. 41–49, Aug 2004.
- [5] A. Wong, G. Kathiresan, T. Chan, O. Eljamaly, and A. Burdett, "A 1V Wireless Transceiver for an Ultra Low Power SoC for Biotelemetry Applications," in *ESSDERC/ESSCIRC*, Munich, Germany, Sep 2007.
- [6] Chipcon, "CC2420: 2.4 GHz IEEE 802.15.4 / ZigBee-ready RF Transceiver." [Online]. Available: <http://www.chipcon.com>
- [7] K. Langendoen and G. Halkes, *Embedded Systems Handbook*. CRC Press, 2005, ch. Energy-Efficient Medium Access Control.
- [8] J. Polastre, J. Hill, and D. Culler, "Versatile Low Power Media Access for Wireless Sensor Networks," in *ACM SenSys*, Baltimore, MD, Nov 2004, pp. 95–107.
- [9] O. Omeni, O. Eljamaly, and A. Burdett, "Energy Efficient Medium Access Protocol for Wireless Medical Body Area Sensor Networks," in *Proc. IEEE-EMBS Symposium on Medical Devices and Biosensors*, Cambridge, UK, Aug 2007, pp. 29–32.
- [10] S.-L. Wu, Y.-C. Tseng, and J.-P. Sheu, "Intelligent Medium Access for Mobile Ad Hoc Networks with Busytone and Power Control," *IEEE J. Sel. Areas Comm.*, vol. 18, no. 9, pp. 1647–1657, Sep 2000.
- [11] J.-P. Ebert, B. Stremmel, E. Wiederhold, and A. Wolisz, "An Energy-efficient Power Control Approach for WLANs," *J. Communications and Networks*, vol. 2, no. 3, pp. 197–206, Sep 2000.
- [12] J. Monks, V. Bharghavan, and W. Hwu, "A Power Controlled Multiple Access Protocol for Wireless Packet Networks," in *IEEE Infocom*, Anchorage, AK, Apr 2001, pp. 219–228.
- [13] J. Pavon and S. Choi, "Link Adaptation Strategy for IEEE 802.11 WLAN via Received Signal Strength Measurement," in *IEEE ICC*, Anchorage, AK, May 2003, pp. 1108–1113.
- [14] E.-S. Jung and N. Vaidya, "A Power Control MAC Protocol for Ad Hoc Networks," *Wireless Networks*, vol. 11, no. 1-2, pp. 55–66, Jan 2005.
- [15] D. Qiao, S. Choi, and K. G. Shin, "Interference Analysis and Transmit Power Control in IEEE 802.11a/h Wireless LANs," *IEEE/ACM Trans. Networking*, vol. 15, no. 5, pp. 1007–1020, Oct 2007.
- [16] T. ElBatt, S. Krishnamurthy, D. Connors, and S. Dao, "Power Management for Throughput Enhancement in Wireless Ad-Hoc Networks," in *IEEE ICC*, New Orleans, LA, Jun 2000, pp. 1506–1513.
- [17] R. Ramanathan and R. Hain, "Topology control of multihop wireless networks using transmit power adjustment," in *IEEE Infocom*, Tel-Aviv, Israel, Mar 2000, pp. 404–413.
- [18] M. Kubisch, H. Karl, A. Wolisz, *et al.*, "Distributed Algorithms for Transmission Power Control in Wireless Sensor Networks," in *IEEE WCNC*, New Orleans, LA, Mar 2003.
- [19] D. Son, B. Krishnamachari, and J. Heidemann, "Experimental Study of the Effects of Transmission Power Control and Blacklisting in Wireless Sensor Networks," in *IEEE SECON*, Santa Clara, CA, Oct 2004, pp. 289–298.
- [20] O. Chipara, Z. He, G. Xing, Q. Chen, X. Wang, C. Lu, J. Stankovic, and T. Abdelzaher, "Real-time Power-Aware Routing in Sensor Networks," in *IEEE IWQoS*, New Haven, CT, Jun 2006, pp. 83–92.
- [21] G. Xing, C. Lu, Y. Zhang, Q. Huang, and R. Pless, "Minimum Power Configuration in Wireless Sensor Networks," in *ACM MobiHoc*, Urbana-Champaign, IL, USA, May 2005, pp. 390–401.
- [22] I. C. Paschalidis, W. Lai, and D. Starobinski, "Asymptotically Optimal Transmission Policies for Large-Scale Low-Power Wireless Sensor Networks," *IEEE/ACM Trans. Networking*, vol. 15, no. 1, pp. 105–118, Feb 2007.
- [23] L. Correia *et al.*, "Transmission Power Control in MAC Protocols for Wireless Sensor Networks," in *ACM/IEEE MSWiM*, Montreal, Canada, Oct 2005, pp. 282–289.
- [24] S. Lin, J. Zhang, G. Zhou, L. Gu, T. He, and J. Stankovic, "ATPC: Adaptive Transmission Power Control for Wireless Sensor Networks," in *ACM SenSys*, Boulder, CO, Nov 2006, pp. 223–236.
- [25] G. Zhou, T. He, S. Krishnamurthy, and J. Stankovic, "Impact of Radio Irregularity on Wireless Sensor Networks," in *ACM MobiSys*, Boston, MA, Jun 2004, pp. 125–138.
- [26] J. Zhao and R. Govindan, "Understanding Packet Delivery Performance in Dense Wireless Sensor Networks," in *ACM SenSys*, Los Angeles, Nov 2003.
- [27] P. S. Hall and Y. Hao, *Antennas and Propagation for Body-Centric Wireless Communications*. Artech House, 2006.
- [28] V. Shrivastava, D. Agrawal, A. Mishra, S. Banerjee, and T. Nadeem, "Understanding the Limitations of Transmit Power Control for Indoor WLANs," in *ACM/Usenix Internet Measurement Conference (IMC)*, San Diego, CA, Oct 2007, pp. 351–364.
- [29] J. Jeong, D. Culler, and J.-H. Oh, "Empirical Analysis of Transmission Power Control Algorithms for Wireless Sensor Networks," in *Intl. Conf. Networked Sensing Systems (INSS'07)*, Kanazawa, Japan, Jun 2007, pp. 27–34.
- [30] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Lewis, "Understanding the Causes of Packet Delivery Success and Failure in Dense Wireless Sensor Networks," in *ACM SenSys*, Boulder, CO, Nov 2006, pp. 419–420.
- [31] K. Srinivasan and P. Lewis, "RSSI is Under Appreciated," in *Workshop on Embedded Networked Sensors (EmNets)*, Boston, MA, May 2006.
- [32] J. Ryckaert, P. D. Doncker, R. Meys, A. de Le Hoye, and S. Donnay, "Channel model for Wireless Communications around Human Body," *Electronics Letters*, vol. 40, no. 9, pp. 543–544, Apr 2004.
- [33] A. Dhamdhare, V. Sivaraman, V. Mathur, and S. Xiao, "Algorithms for Transmission Power Control in Biomedical Wireless Sensor Networks," in *IEEE Workshop on Wireless Network Algorithms (WiNA)*, Yilan, Taiwan, Dec 2008.
- [34] A. Dhamdhare, V. Sivaraman, and A. Burdett, "Experiments in Adaptive Power Control for Truly Wearable Biomedical Sensor Devices," in *Intl. Workshop on Adaptation in Wireless Sensor Networks (AWSN)*, Sydney, Australia, Dec 2008.



Shuo Xiao received his B.E. from Beijing University of Aeronautics and Astronautics, China, in 2005, and his M.E. from the University of New South Wales (UNSW), Sydney, Australia, in 2008. He is currently working as a part-time research assistant at UNSW. His research has focused on transmission power control algorithm design for wireless sensor networks.



Ashay Dhamdhare received his B.Tech. from the Indian Institute of Technology (IIT) Mumbai in 1998, and his MS and PhD from the University of California, San Diego in 2001 and 2004 respectively. He has worked at the Royal Institute of Technology (KTH), Stockholm, the Swedish Defence Research Agency, Linköping, and Airbus Industrie, Hamburg. He is currently a Senior Research Associate at the University of New South Wales, Sydney. His research interests include wireless sensor networks, cross-layer network design and resource allocation.



Vijay Sivaraman (M '94) received his B.Tech. from the Indian Institute of Technology in Delhi, India, in 1994, and his Ph.D. from the University of California at Los Angeles in 2000. He has worked at Bell-Labs and a silicon valley start-up manufacturing optical switch-routers. He is now a senior lecturer at the University of New South Wales in Sydney, Australia, and works part-time at the CSIRO. His research interests include sensor networking, optical switching, network design, and QoS.



Alison Burdett (SM) received her B.Eng. and Ph.D. from the Imperial College of Science, Technology and Medicine, London University, in 1989 and 1992 respectively. She was a Senior Lecturer at Imperial College before joining Toumaz Technology Ltd. as Director of Engineering in 2001. Her expertise is in the area of high frequency analogue and wireless integrated circuits. She has co-authored over 50 academic papers and articles, and designed several commercially successful silicon chips.