

Algorithms for Transmission Power Control in Biomedical Wireless Sensor Networks

Ashay Dhamdhere, Vijay Sivaraman, Vidit Mathur and Shuo Xiao

School of Electrical Engineering and Telecommunications

University of New South Wales

Sydney, NSW 2052, Australia

Emails: {ashay, vijay}@unsw.edu.au, {z3131001, shuo.xiao}@student.unsw.edu.au

Abstract—Wireless sensor networks are increasingly being used for continuous monitoring of patients with chronic health conditions such as diabetes and heart problems. As biomedical sensor nodes become more wearable, their battery sizes diminish, necessitating very careful energy management. This paper proposes feedback-based closed-loop algorithms for dynamically adjusting radio transmit power in body-worn devices, and evaluates their performance in terms of energy savings and reliability as the data periodicity and feedback time-scales vary. Using experimental trace data from body worn devices, we first show that the performance of dynamic power control is adversely affected at long data periods. Next for a given data period we show that modifying the transmit power at too long time-scales (around a minute) reduces the efficacy of dynamic power control, while too short a time-scale (few seconds or less) incurs a high feedback signaling overhead. We therefore advocate an intermediate range of time-scales (when permitted by the data periodicity), typically in the few tens of seconds, at which the control algorithms should adapt transmit power in order to achieve maximal energy savings in body-worn sensor devices used for medical monitoring.

I. INTRODUCTION

Wireless sensor network technologies have the potential to offer large-scale and cost-effective solutions to the challenge of monitoring patients with chronic medical conditions. 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 tens of kilo-Joules of energy. In contrast, emerging truly-wearable health monitoring devices such as the “digital plasters” being developed at Toumaz Technologies have orders of magnitude lower battery capacity (the Toumaz SensiumTM chip shown in figure 1 operates on a flexible paper-thin printed battery [3] with a capacity of approximately 70 Joules). Such stringent energy constraints necessitate very careful energy management.

Communication is the most energy consuming operation performed by a sensor node [4], and can be optimized at multiple layers of the communication stack. At the data-link layer, several intelligent medium access control (MAC) protocols [5], [6], [7] have been designed which duty-cycle the radio, i.e. turn the radio off whenever packet transmission or receipt is not expected, thereby saving energy. 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. Moreover, most sensor platforms have the ability to

control the radio transmit power dynamically: the CC2420 radio in Crossbow’s MicaZ motes provides 32 transmission levels ranging from -25dBm to 0dBm output, while the SensiumTM platform supports 8 levels ranging from -23dBm to -7dBm output.

Our earlier paper [8] showed that transmit power control can help body-worn devices save energy with minimal impact on reliability. Our algorithm embedded power control information in acknowledgment packets sent in response to each transmission. However, this inherently limits the frequency of power updates to the periodicity of the data. Our first contribution is to show that in certain scenarios the efficacy of power control is degraded at long data periods.

Next for a given periodicity of data, transmit power adaptation requires the sensor node to listen to signaling packets from the base – either in the form of link-quality feedback information (if the algorithm operates at the sensor node), or control messages telling the sensor node what transmit power to use (if the algorithm runs on the base). However, listening to signaling packets consumes energy at the sensor node.

Our second contribution is to investigate whether further energy savings can be obtained by reducing the signaling overhead. Lesser signaling reduces listening power consumption, but also implies that the sensor node will adjust its transmit power less often, potentially increasing energy wastage and/or packet loss. In this paper we empirically quantify this trade-off and show that there is an optimal region of signaling frequency that maximizes energy conservation at the sensor node.

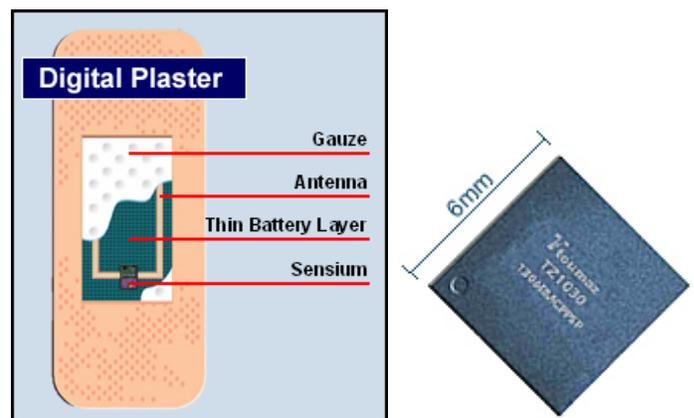


Fig. 1. Toumaz SensiumTM Digital Plaster

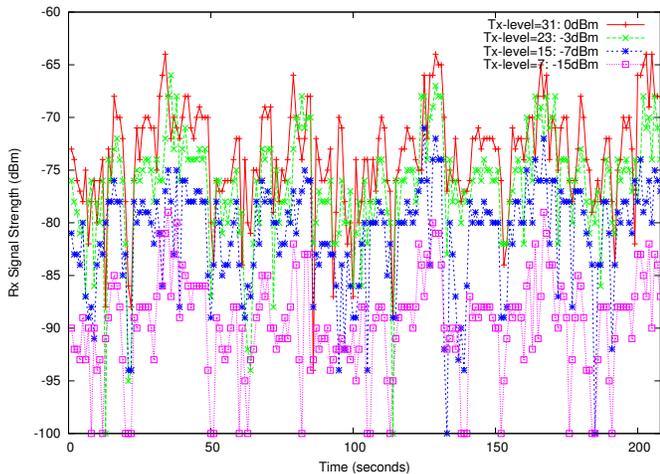


Fig. 2. Fast walk: RSSI vs. time

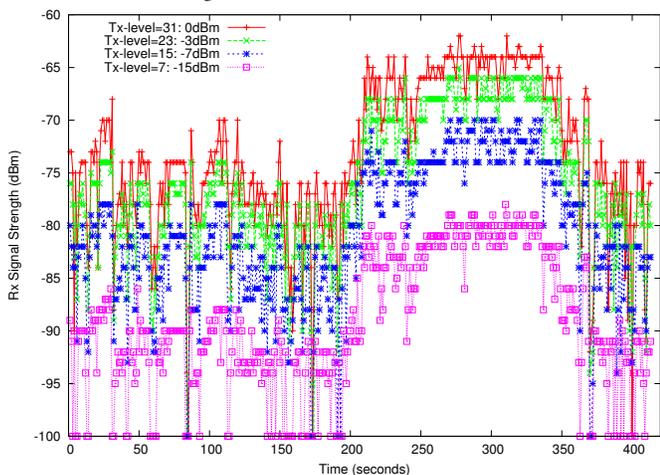


Fig. 3. Slow walk: RSSI vs. time

II. BACKGROUND: ADAPTIVE POWER CONTROL FOR BODY DEVICES

Several empirical studies have recorded the variable nature of wireless links in sensor networks [9], [10], and the idea of dynamically adapting transmit power has been explored before [11], [12], [13]. However, these earlier studies have targeted static deployments, such as for environmental or structural monitoring applications, wherein variability in wireless link quality over time is lower and slower. In contrast, our work considers wearable mobile devices where the wireless link quality can change significantly and rapidly since it is very susceptible to position and orientation of the human body.

Our earlier work [8] motivated the need for dynamic power control by recording extensive trace data, and proposed a power control algorithm for the sensor node. The traces record the RSSI at the base station when packets are sent by the target at 16 different power levels every second. Three scenarios were considered: in *fast walk* the patient walks back and forth in a room for a few minutes at a reasonably active pace, while *slow walk* considers a slowly moving person (such as an elderly or handicapped person), who takes over six minutes to walk a distance of three meters. The *resting* scenario has the patient resting for approximately twenty minutes on a chair, six meters

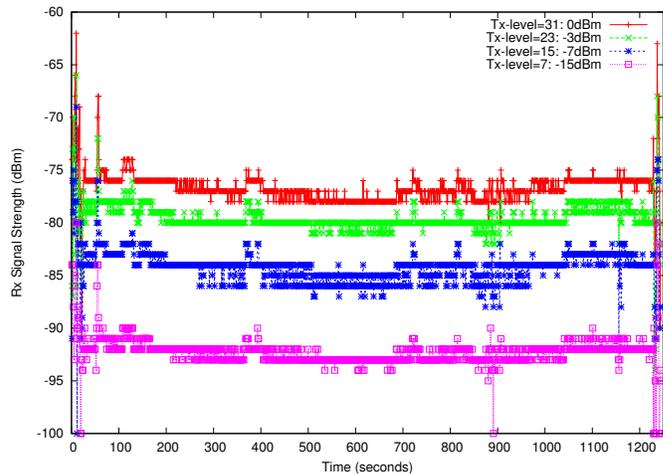


Fig. 4. Resting: RSSI vs. time

away from a base station. The RSSIs for 4 power levels are plotted against time in figures 2, 3 and 4.

Note that the RSSI for a given power level varies significantly with time, particularly for the fast walk and slow walk scenarios. Thus transmitting at a fixed power level can be wasteful (if the chosen level is high) or cause unreliability in packet reception (if the chosen level is low), motivating the need for dynamic power control.

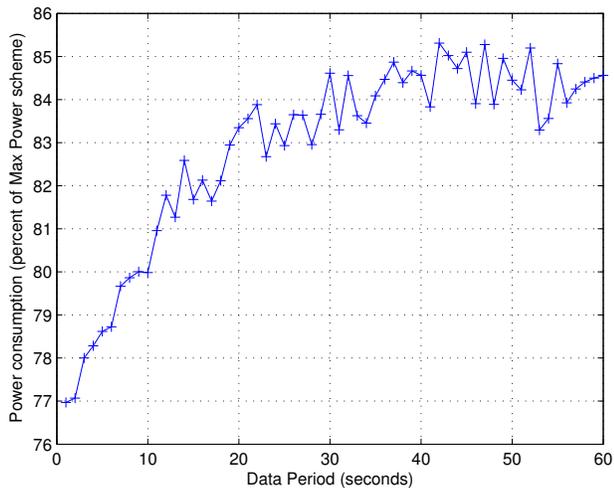
- 1) $\bar{R} \leftarrow (1 - \alpha)\bar{R} + \alpha R$
- 2) if $\bar{R} < T_L$ double the transmit power level
- 3) if $\bar{R} > T_H$ reduce the transmit power level by 2
- 4) if $T_L \leq \bar{R} \leq T_H$ no action is required

Fig. 5. Power control algorithm with per-packet feedback

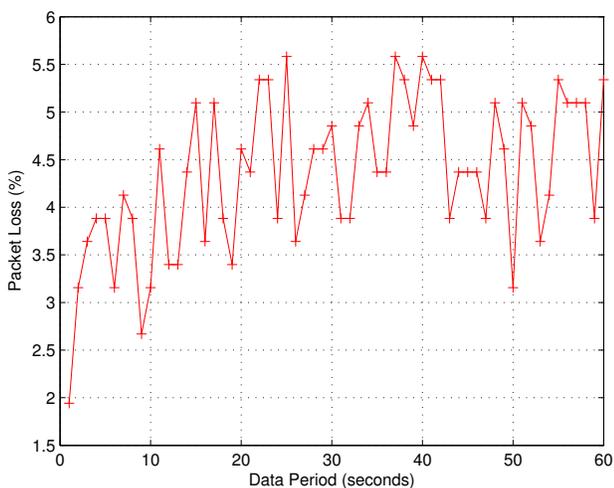
We reproduce our original power control algorithm in figure 5. R denotes the RSSI at the base station for the current packet which is conveyed to the target in an acknowledgment packet, while \bar{R} is the average over time. The algorithm attempts to keep \bar{R} between the thresholds T_H and T_L by suitably adjusting the transmit power level (among the 32 available in the MicaZ mote). We chose the value $\alpha = 0.8$ to make the algorithm reactive to the current sample.

In our earlier work we had evaluated the algorithm for an ECG monitoring scenario, where data was generated every second. However different biometric data types have different periodicities: for instance the data period for blood glucose and temperature monitoring may be significantly longer. Further, our algorithm assumed that every data transmission was acknowledged. However this may not be required in many situations (e.g. in the case of temperature monitoring, an acknowledgment may be required only when the temperature is in an abnormal range). On the other hand, listening for acknowledgments consumes energy at the body worn device.

We therefore examine two aspects in this work. We first explore the effect of data periodicity on the efficacy of dynamic power control in §III, and show that long data periods can degrade power control performance in some scenarios. Next we consider the effect of update period on power consumption, for applications that do not require per-packet acknowledgments. In §IV we present a power control algorithm, operating at



(a) Slow Walk Power Consumption



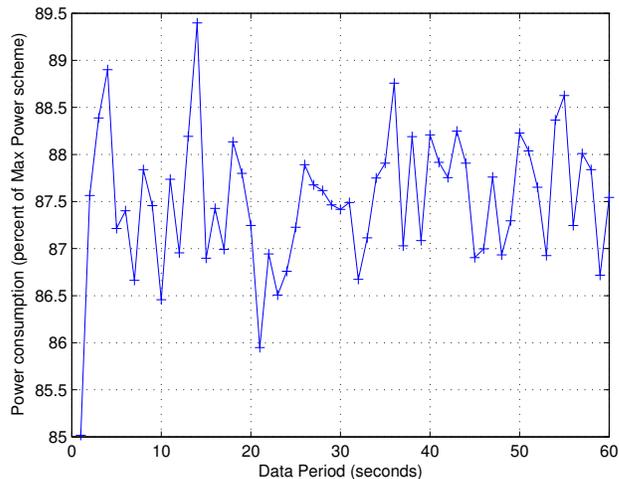
(b) Slow Walk Packet Loss

Fig. 6. Power consumption and packet loss for slow walk

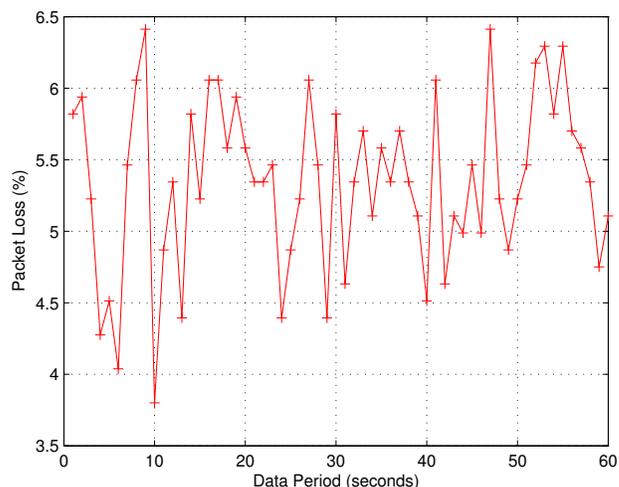
the base station, that sends a power control message to the target every n packets. Our intent is to reduce the receive power consumption at the body node, by potentially sacrificing some transmit power due to a coarser time-scale of power adaptation. Through evaluation on trace data in §V we show that there is an optimal range for the feedback frequency n which minimizes total power consumption. Finally we prototype our algorithm on MicaZ motes and verify our results through experiments in §VI.

III. EFFECT OF DATA PERIODICITY ON TRANSMIT POWER CONTROL

The data period inherently limits the frequency of power updates for the algorithm in figure 5. As data is sent less frequently, the channel may change by the time the updated power is used, adversely affecting energy usage and packet loss. For the fast walk, slow walk and resting scenarios we evaluate the transmit power consumption and packet loss as the data period varies. Results are shown in figures 6, 7 and 8, with the transmit power consumption compared against a



(a) Fast Walk Power Consumption



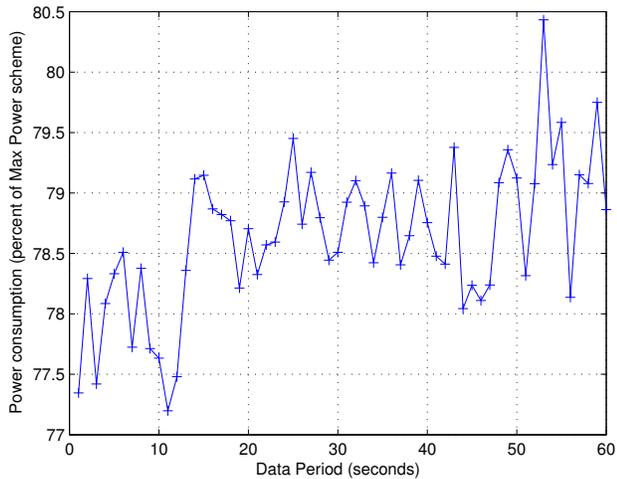
(b) Fast Walk Packet Loss

Fig. 7. Power consumption and packet loss for fast walk

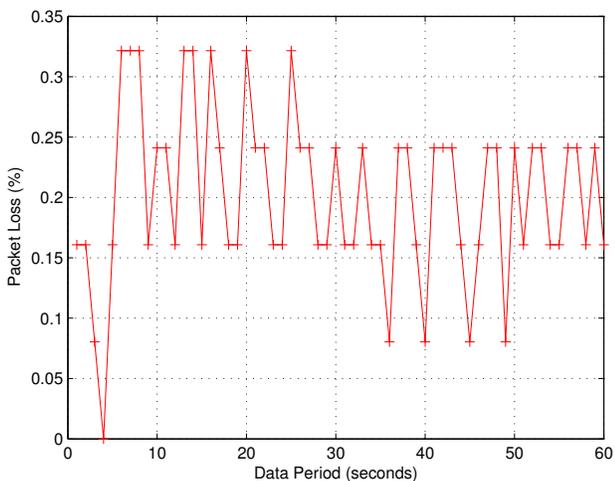
constant transmit power scheme where the maximum power level is used for all transmissions.

From figure 6(a) it is clear that data periodicity has a pronounced effect on energy consumption in the slow walk scenario. Referring to figure 3, we note that the change in channel quality is relatively gradual (compared to fast walk), and the channel quality is more correlated across time. When data is sent frequently, the power control loop is able to follow the channel variations closely, thus saving power. As data period increases, power control becomes less effective and the scheme consumes more power (note the 7.5% difference in energy consumption between data periods of 1 and 60 seconds). Similarly the packet loss is lowest when frequency of data transmission is highest, as seen in figure 6(b).

In contrast, the performance for fast walk is relatively insensitive to the data period. Figure 2 shows that the channel varies rapidly in fast walk, with the channel quality often oscillating every second. Since the correlation in channel quality across time is inherently low, sampling the channel finely does not yield additional benefits. Thus, power control performance is essentially identical as the data period changes:



(a) Resting Power Consumption



(b) Resting Packet Loss

Fig. 8. Power consumption and packet loss for resting

note that power consumption lies between 86.5% and 88.5% of the maximum power scheme, and the packet loss lies between 4.5% and 6%, for most data periods we consider.

Similarly for the resting scenario we find that performance is relatively unaffected by data period. From figure 4 we note that the resting scenario is characterized by long periods of stable channel behaviour, interspersed with some oscillations in channel quality. Since the channel is relatively stable in this scenario, it is not necessary to sample it very closely; thus even when data is transmitted infrequently, the power control algorithm performs reasonably well. Note that power consumption lies between 77.5% and 79.5% of the maximum power scheme for most data periods considered.

It is clearly seen that transmit power consumption depends greatly on both, the rate of channel variations and the periodicity of data transmissions. When permitted by the periodicity of the data, fine-grained power control can yield benefits in some scenarios; however this benefit has the associated cost of listening to power update commands. The following section evaluates the energy cost of signaling for power control.

IV. POWER CONTROL ALGORITHMS WITH TUNABLE FEEDBACK COST

Listening for feedback packets from the base station consumes energy. To assess the impact of different feedback frequencies, we present a modified power control algorithm which is implemented at the base station. Rather than sending feedback information for every transmitted packet, the base station maintains a running average R_n of received RSSI. A power control decision is made only every n th sample, by updating \bar{R} using R_n and comparing with the thresholds as in the original algorithm. As before, the RSSI of the current sample is denoted by R . The complete algorithm is presented in figure 9.

- 1) for sample k : $R_n = R_n + R/n$
- 2) if $(k \bmod n) = 0$,
 - a) $\bar{R} \leftarrow (1 - \alpha)\bar{R} + \alpha R_n$
 - b) if $\bar{R} < T_L$ double the transmit power level
 - c) if $\bar{R} > T_H$ reduce the transmit power level by 2
 - d) if $T_L \leq \bar{R} \leq T_H$ no action is required
 - e) Set $R_n = 0$

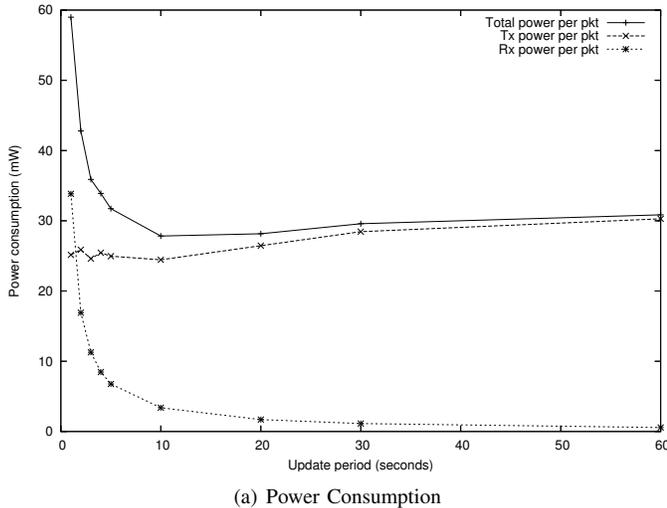
Fig. 9. Power control algorithm with tunable feedback frequency

At small values of n the target node must expend energy to listen frequently to power update commands; however it is responsive to channel variations. At large values of n the target is less responsive to channel variations (and may expend more energy in transmissions), but must spend less energy in receiving update commands. The parameter n allows us to trade off between transmit and receive power consumptions. In the following section we evaluate our algorithm on trace data to identify an optimal frequency for update information.

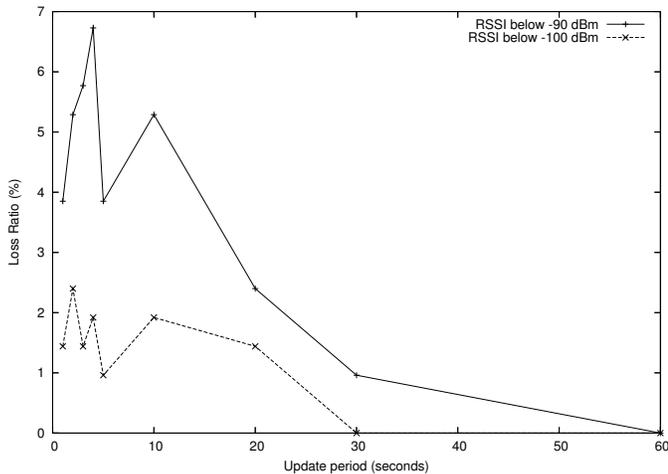
V. ANALYSIS USING TRACE DATA

The transmit power consumption rises with an increase in update period for both the slow and fast walk scenarios, as seen in figures 10(a) and 11(a). At high update periods the algorithm is slow in responding to a good channel, and consumes more energy. On the other hand the receiving power consumption falls as the update period increases, as the receiver must listen less often to feedback information. As a result, the total power consumption (shown by the solid curve) starts at a high value (dominated by the listening cost), and climbs again for large values of the update period (due to degradation in power control performance), passing through a minimum value which occurs at an update period of 10 seconds for fast walk and 20 seconds for slow walk. *There is thus an optimal signaling interval which effectively balances the energy cost of listening to feedback with the need to update transmit power.* This interval, while depending on the particular scenario, lies in the range of a few tens of seconds.

Further, packets are lost if the RSSI is below a communication threshold dictated by the receiver sensitivity of the sensor node. We plot the packet losses in figures 10(b) and 11(b) for two thresholds, -90dBm and -100dBm . As we would expect, the price we pay for the lowered power consumption



(a) Power Consumption



(b) Packet Loss

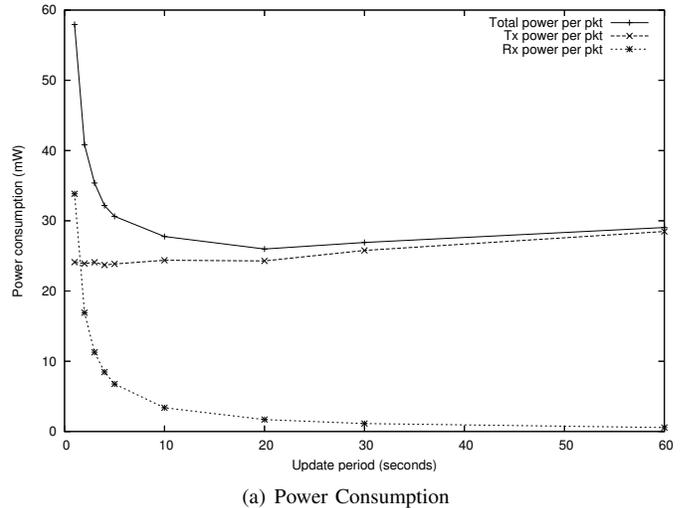
Fig. 10. Trace data results for a fast walking patient: Power consumption and packet loss as a function of update periodicity

is a relative increase in packet loss, clearly seen in figure 11(b) for slow walk, with a peak in packet loss coinciding with the trough in power consumption.

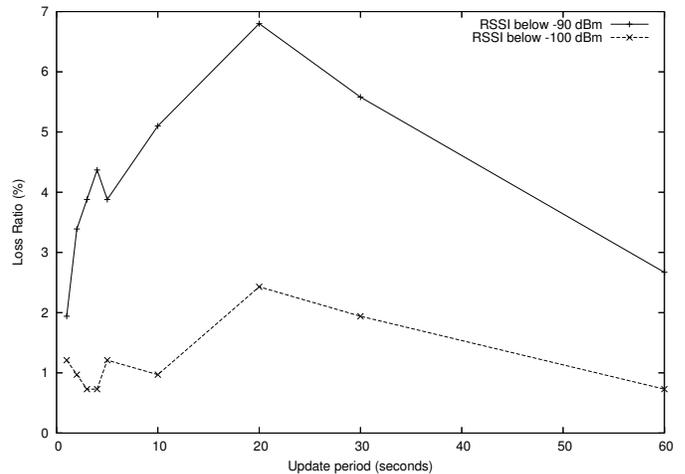
VI. PROTOTYPING AND EXPERIMENTATION

We corroborate our earlier analysis by prototyping our algorithm on the MicaZ motes, and evaluating them under two scenarios. In the first, the patient performs a fast walk in a straight line, towards and away from the base station. In the second, the patient walks 2 meters and proceeds to sit in a chair 2.5 meters from a base station. The experiments are performed with 8 different values of the update period, $n = 1, 2, 3, 4, 5, 10, 20, 50$. For each value of n the patient performs the same movements, though some variability in the results is to be expected due to the nature of the wireless channel. The chosen transmit power and received RSSI for three update periods ($n = 1, 10$, and 50) are plotted in figures 12(a) and 12(b) for a walking patient.

With a higher update period ($n = 50$) the transmit power is slow to react to a good channel, and is therefore held at a higher value compared to a low update period ($n = 1, 10$);



(a) Power Consumption



(b) Packet Loss

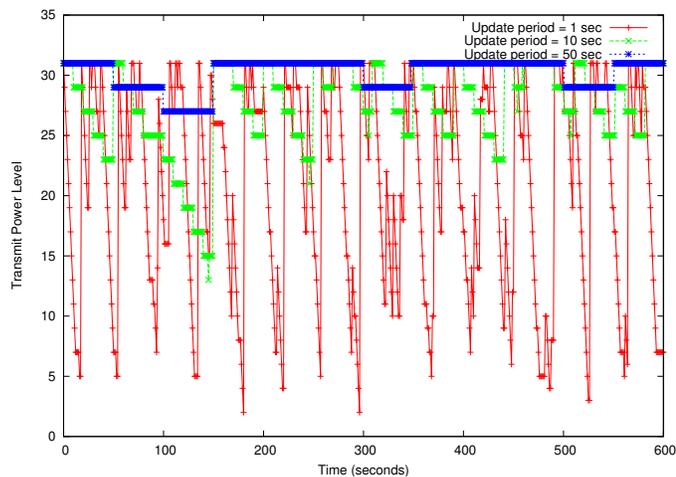
Fig. 11. Trace data results for a slow walking patient: Power consumption and packet loss as a function of update periodicity

thus a lower update period results in transmit power savings as seen in figure 13(a). Note again, that the power savings comes at a price; the RSSI falls below -90 dBm more often for $n = 1, 10$ than for $n = 50$, resulting in higher packet losses at lower update periods. These observations concur precisely with results from trace analysis.

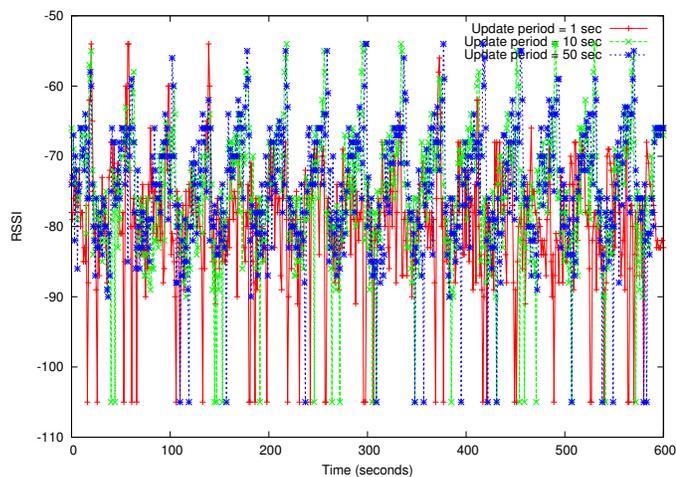
Finally, our experiments corroborate our observation regarding an optimal signaling interval, with the minimum total power consumption occurring at an update period of 20 seconds for the walking patient, and more markedly, at 10 seconds for the resting patient.

VII. CONCLUSION

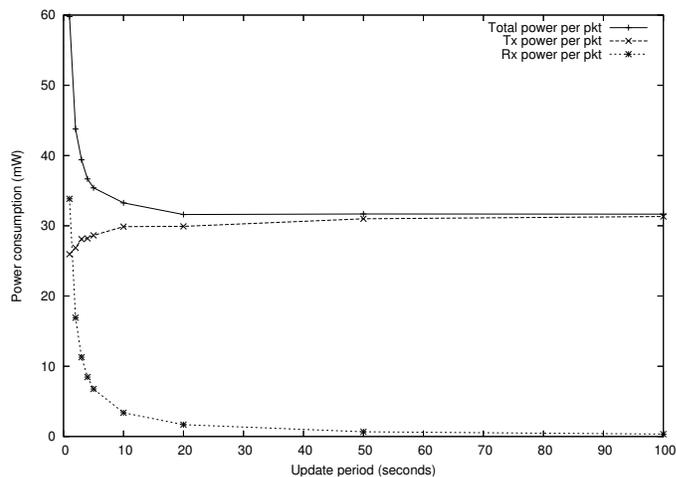
In this work we address two aspects which affect dynamic power control, namely the periodicity of data transmissions and the frequency of transmit power updates. We bring out the dependence between data periodicity and the patient's mobility, and show that under certain scenarios the performance of dynamic power control is degraded as the data period increases. Next we explore the impact of the transmit power update frequency on the total power consumption. We



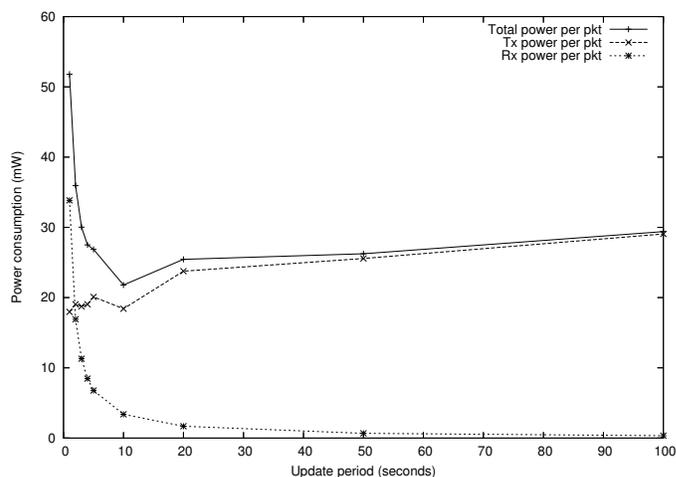
(a) Power Level



(b) RSSI



(a) Walking Patient



(b) Resting Patient

Fig. 12. Experimental results for a walking patient: Transmit power level and RSSI variation with time

show that using a very small or very large update interval is inefficient, and further point out that an update interval in the few tens of seconds minimizes power consumption in the body-worn sensor node.

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] K. Langendoen and G. Halkes, *Embedded Systems Handbook*. CRC Press, 2005, ch. Energy-Efficient Medium Access Control.
- [6] 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.
- [7] 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.

Fig. 13. Experimental results for walking and resting patients: Power consumption as a function of update periodicity

- [8] S. Xiao, V. Sivaraman, and A. Burdett, "Adapting Radio Transmit Power in Wireless Body Area Sensor Networks," in *Proc. ACM BodyNets*, Tempe, Arizona, USA, Mar 2008.
- [9] 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.
- [10] J. Zhao and R. Govindan, "Understanding Packet Delivery Performance in Dense Wireless Sensor Networks," in *ACM SenSys*, Los Angeles, Nov 2003.
- [11] 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.
- [12] 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.
- [13] 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.