

Experiments in Adaptive Power Control for Truly Wearable Biomedical Sensor Devices

Ashay Dhamdhere, Vijay Sivaraman
School of Electrical Engineering and Telecommunications
University of New South Wales
Sydney, NSW 2052, Australia
Emails: {ashay, vijay}@unsw.edu.au

Alison Burdett
Toumaz Technology Limited
Building 3, 115 Milton Park,
Abingdon, UK
Email: alison.burdett@toumaz.com

Abstract

Emerging body-wearable devices for continuous health monitoring are severely energy constrained and yet required to offer high communication reliability under fluctuating channel conditions. Such devices require very careful management of their energy resources in order to prolong their lifetime. In our earlier work we had proposed dynamic power control as a means of saving precious energy in off-the-shelf sensor devices. In this work we experiment with a real body-wearable device to assess the power savings possible in a realistic setting. We quantify the power consumption against the packet loss and establish the feasibility of dynamic power control for saving energy in a truly-body-wearable setting.

1. Introduction

Healthcare today is experiencing enormous cost pressures from the necessity for monitoring patients with chronic medical conditions. The global demographic trend towards an ageing population, coupled with a sedentary lifestyle and poor diet is leading to an increasing number of people living for years (or even decades) with chronic conditions requiring ongoing clinical management. The healthcare information systems in use today are mainly designed to manage acute illness, and are ill-equipped to cope with demands for pervasive monitoring of chronic conditions. Wireless body-worn sensor devices have the potential to provide a large scale and cost-effective solution to this challenge. At Toumaz Technology, we are developing the wireless infrastructure for intelligent, non intrusive, continuous patient monitoring.

For a body-wearable patient monitoring solution to be successful the devices must be sufficiently small and lightweight so as to not impede patient activity. At Toumaz

we are developing the ultra low-power Sensium™ “digital plasters” (shown in figure 1) which are designed to be truly body-wearable. This is a significant reduction in size over existing prototype devices (such as the MicaZ motes [4] used in Harvard’s CodeBlue [11] project) while allowing for monitoring of ECG, temperature, blood glucose and oxygen levels. The Sensium™ can also interface to 3-axis accelerometers and pressure sensors; and multiple digital plasters can be attached to a patient, allowing physiological data to be relayed back to a basestation Sensium™ plugged into a PDA or Smartphone.

A small form factor inherently constrains the battery size. The Sensium™ operates on a flexible paper-thin printed battery [10] with a capacity of around 20 mWatt-hours (in comparison, the MicaZ motes operate on a pair of AA batteries which provide a few Watt-hours of capacity). Such stringent energy constraints require careful energy management.

Communication is the most energy consuming operation that a sensor node performs [5], and can be optimised at multiple layers of the communication stack. At the physical layer, we have innovated an ultra-low-power radio [13] 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 [2] 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 [6] for a survey). The B-MAC [9] protocol included in the TinyOS distribution and its enhancements [1] provide versatility to the application in controlling the duty-cycling of the radio, while at Toumaz we have developed our proprietary MAC protocol [8] 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. Further, the ability to vary the transmission power is available on most platforms (the Sensium™ supports 8 levels, ranging from -23dBm to -7dBm output, while the CC2420 radio in the MicaZ motes provides 32 levels ranging from -25dBm to 0dBm output).

In our earlier work [14] we had proposed dynamic power control as a means of improving the energy performance of body area sensor devices, and validated our proposed algorithms using MicaZ motes. In this work we experiment with the Toumaz Sensium platform and establish the efficacy of dynamic power control for truly-wearable devices which are specifically designed for health monitoring. The Sensium™ operates in the 900MHz frequency spectrum (as opposed to 2.4GHz for the MicaZ motes), and has a lower number of output power levels (8 as opposed to 32 in the MicaZ motes). Our focus is therefore on assessing the impact of power control in trading off reliability of the wireless link for energy efficiency in a realistic body-worn setting. Note that the variable nature of wireless links in sensor networks has been recorded in several empirical studies [16, 15], and the idea of dynamically adapting transmit power has been explored before [7, 3, 12]. 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 first contribution is to profile the wireless channel using the Sensium™ platform by emulating patient activity. We show that the wireless channel can exhibit large variations in quality, depending on the patient activity and orientation relative to the base station, thus making fixed transmit power sub-optimal; setting a low transmit level results in excessive losses when the link quality is poor, while setting a high transmit level wastes energy when the chan-




Figure 1. Toumaz Sensium™ Digital Plaster

nel is good. Next we assess the potential savings afforded by power control, by evaluating offline, the optimal power to be used at each time instant, so as to expend minimum energy while maintaining reliable communication. Though infeasible to implement in practice, this evaluation allows an understanding of the potential benefits of power control, and serves as a benchmark to compare practical schemes. Finally, we adapt our previously proposed power control scheme to the Sensium™ platform, and quantify the energy savings and reliability figures achieved by our power control scheme. We show that the proposed scheme achieves energy savings varying between 14% and 32% with an acceptable decrease in reliability, highlighting the efficacy of dynamic power control.

2. Channel Profiling

We begin by empirically profiling the temporal variations in the quality of the wireless link between a body-worn Sensium™ device and a fixed Sensium™ base-station, as a patient wearing the device performs various activities. The patient was played by the first author. In each experiment the device was strapped around the patient’s chest, simulating continuous monitoring of heartbeat and ECG. The experiments were conducted indoors in an office space containing 8 cubicles. The base-station was placed at an elevation on one side of the room to provide better line-of-sight coverage.

transmit level	output power (dBm)	power draw (mW)
7	-6	2.8
6	-7	2.7
5	-9	2.6
4	-10	2.5
3	-12	2.4
2	-15	2.2
1	-18	2.0
0	-22	1.8

Table 1. Output power setting and typical battery power draw for the Toumaz Sensium™ radio

The Sensium™ supports 8 RF output power levels which can be controlled at runtime, and can provide an RSSI value (varying between 0 and 7, 0 being the lowest) for each transaction. The output power (in dBm) and corresponding energy consumption rate (in mW) for these 8 levels are shown in table 1. To profile the wireless channel, we send packets from the body-worn device to the base station at 125 millisecond intervals. Each packet is sent at a successively higher power level, reverting to level 0 after all 8 levels have been transmitted. This effectively allows us to sample the




Figure 2. Fast walk: RSSI vs. time

channel at all 8 power levels every second. At each power level, the average RSSI at the base station is recorded; this allows us to map the transmitted power level to the received RSSI, across time (when a packet is lost, we record the RSSI for that packet as -1). While we would like to ideally take simultaneous measurements at all 8 power levels, using a small inter-packet spacing allows more-or-less simultaneous measurement. We simulate three types of patient behaviour, detailed below.

2.1. Fast Walk

This scenario has the patient walking back and forth in the room for a few minutes at a reasonably active pace; the patient stays between 1 and 5 meters from the base-station at all times. The body device, strapped on to the patient's chest, is assumed to generate a packet every second (simulating a heartbeat/ECG monitor) and transmits it at all 8 output levels.

Figure 2 plots the RSSI recorded at the base station for power levels 0, 4 and 7 as a function of time (for ease of visibility we only plot the first 200 seconds). Note that for any power level, the brisk movements of the patient result in strong fluctuations in the RSSI, with the RSSI routinely fluctuating by 2 – 3 levels in a matter of seconds. Note also certain trends with respect to time; as the orientation of the patient changes (as the patient walks towards and away from the base station), the channel quality shows a corresponding increase or decrease.

2.2. Slow Walk

This scenario considers a slowly moving person, such as an elderly or handicapped person with restricted mobility, who walk slowly (one step every six seconds) towards and away from a base station (simulating movement around a




Figure 3. Slow walk: RSSI vs. time

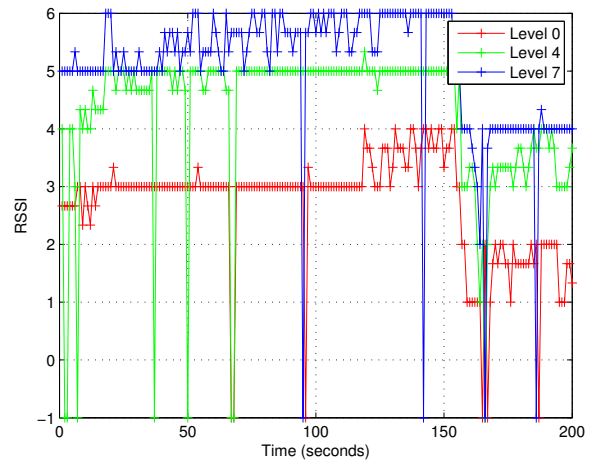


Figure 4. Resting: RSSI vs. time

room, or to a toilet). The RSSI recorded at power levels 0, 4 and 7 are plotted in figure 3 (once again, we plot the first 200 seconds of the trace). Note that the variation in RSSI is much less rapid as compared to the fast walk scenario. The channel quality shows definite trends as the patient moves towards and away from the base station, and maintains its value when the patient is stationary. Again, it is clear that a fixed transmit power would perform sub-optimally, wasting power when the channel is good and causing packet losses when the channel is bad. Clearly, the RSSI for lower power levels is below that for higher levels, as we would expect.

2.3. Resting

In this scenario the patient sits down to rest on a chair a few feet away from the base-station. Figure 4 plots the RSSI over the first 200 seconds at transmit levels 0, 4 and 7. Note that in this case the wireless link remains stable

for extended periods, changing occasionally to reflect the patient’s orientation. Power control could potentially have major benefits in this scenario; at the same time, a fixed power scheme would once again be wasteful, since it would have to cater to the worst case channel scenario.

Having gained a better understanding of the wireless channel under various patient activity scenarios, we quantitatively assess the potential performance impact of adaptive transmit power control in the next section.

3. Optimal Transmit Power

We quantify the potential benefits of adaptive transmit power control by computing what the “optimal” transmission level might be for each of the three scenarios considered before. We define the “optimal” as the lowest required transmit power level (as a function of time) to achieve a RSSI value greater than 0. Further, we treat an RSSI of 0 as a packet loss; while this is conservative, it ensures that the optimal scheme tries to keep the RSSI away from the lowest possible level. Since we know the RSSI level corresponding to each transmit power level, the optimal simply picks the lowest power level which has an RSSI above 0. Though this scheme is not achievable in practice, it provides an upper bound on potential savings.

We plot the chosen transmit level, as well as the received RSSI for the optimal scheme in figure 5 (for ease of comparison the results of our practical power control scheme are superimposed on this plot and will be discussed in the next section). When the patient mobility is high and the channel is varying rapidly the optimal scheme must vary the transmit power frequently to maintain an RSSI of 1, as seen in subplots (a) and (b). Using the energy consumption values in table 1) we compute the energy consumption of the offline optimal scheme. Comparing with the maximum power scheme (where every packet is transmitted at power level 7) the offline optimal results in an energy savings of 29%. Note however, that the channel varies rapidly in this case, and a practical scheme (which would have to predict future channel quality based on past samples) may not be able to track the variations accurately. We evaluate the performance of a practical scheme in the following section.

In contrast the channel variations during slow walk are significantly more gradual, as seen in figure 3. The channel shows variations depending on the patient’s orientation (facing towards or away from the base station) as well as the distance to the base; note the gradual increase in channel quality from 30 to 65 seconds, followed by a gradual decrease from 65 to 80 seconds). However the channel quality remains good most of the time (with power level 0 returning an RSSI value 1 for significant periods). As a result the optimal scheme can maintain a transmit level of 0 for significant periods as seen in figures 5(c) and 5(d),

occasionally increasing the power to account for variations (e.g. at 9 seconds power level 0 is not received, and power level 3 is chosen to maintain acceptable RSSI). The optimal saves 34% power compared with maximum power; further, this scenario affords a practical scheme a better chance of achieving savings, since past channel quality is a better indication of future channel behaviour.


Finally when the patient is resting, the link is very stable and the optimal scheme can almost always use power level 0. The scheme has energy savings of 35% as compared with maximum power. The stability of the channel in this case affords a practical scheme ample opportunity to conserve energy without sacrificing reliability. In the following section we discuss such a scheme and evaluate its performance.

4. Practical Transmit Power Control Scheme


The optimal scheme described above was performed off-line and is infeasible to realise in practice since it requires the sender to have *a priori* knowledge of the link quality at the receiver. On the other hand, practical on-line power control schemes rely on channel quality feedback from the receiver to adjust the transmit power. There is thus an inherent delay between sampling the channel and adjusting the transmit power, which limits the performance of practical schemes. In our earlier work [14] we had proposed practical power control schemes for body area devices. In this section we analyze and evaluate the performance of the “aggressive” scheme in the three scenarios under consideration.

Our scheme was inspired by the TCP congestion control mechanism, wherein the receiver maintains a running average of the RSSI values. Using a weighted average immunizes the scheme against rogue samples. We use exponential averaging to update the running average; thus for every new RSSI sample R the running average \bar{R} is updated as $\bar{R} \leftarrow (1 - \alpha)\bar{R} + \alpha R$. Here α is a weighting factor which decides how heavily the current RSSI sample is weighted. We use $\alpha = 0.8$ to make the scheme reactive to the current sample, (and hence the name “aggressive” for this scheme). The scheme attempts to keep the running average between thresholds T_L (lower) and T_H (upper). We adapt our original scheme to account for the 8 power levels on the SensiumTM (as compared to 32 for the MicaZ motes); thus when the running average falls below the lower threshold T_L we increase the transmit power by 2 levels, while if it rises above the upper threshold T_H we reduce the transmit power by 1 level. As long as the running average remains between T_L and T_H no action is taken. We outline the algorithm in figure 6.


We choose $T_L = 1$ and $T_H = 2$; thus our scheme attempts to keep the RSSI at level 1 or above (since RSSI level 0 is the lowest receivable RSSI, our scheme attempts



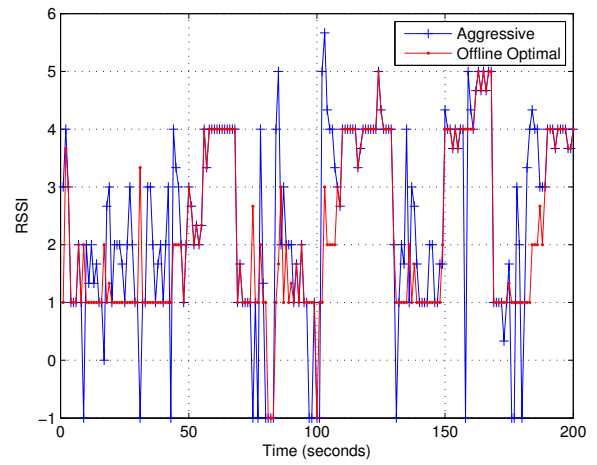
(a) Fast Walk: Chosen transmit levels



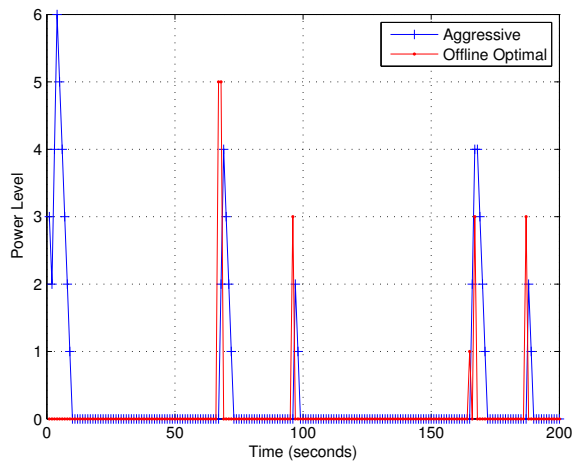
(b) Fast Walk: RSSI at Base Station




(c) Slow Walk: Chosen transmit levels



(d) Slow Walk: RSSI at Base Station



(e) Resting: Chosen transmit levels



(f) Resting: RSSI at Base Station

Figure 5. Transmit power and associated RSSI for fast walk, slow walk, and resting scenarios under Offline Optimal and Aggressive power control schemes

1. $\bar{R} \leftarrow (1 - \alpha)\bar{R} + \alpha R$
2. if $\bar{R} < T_L$ increase the transmit power by 2 levels (capped at 7)
3. if $\bar{R} > T_H$ reduce the transmit power by 1 level (capped at 0)
4. if $T_L \leq \bar{R} \leq T_H$ no action is required

Figure 6. Practical Power Control Scheme

to keep the RSSI above this level to increase reliability). As in the case of the offline optimal, we consider a packet lost if it is received at an RSSI of 0 or lower.

We evaluate the efficacy of our practical scheme on the three scenarios outlined before. At each time instant the algorithm makes a decision on the transmit power based on the weighted average. This transmit power is then used for the next transaction. From our trace data we know the received RSSI for every power level at every time instant, allowing us to implement the power control scheme. The energy usage and reliability of our scheme is summarized in table 2, while the transmit power and corresponding RSSI are plotted in figure 5.

For the fast walk scenario the transmit power varies appreciably to track the channel variations. Note that our practical scheme is slower to react to a poor channel than the offline optimal; this is due to the inherent delay in the feedback loop. As a result, the practical scheme is liable to lose some packets, as seen in RSSI plot in figure 5(b). Further, after experiencing a bad channel, the practical scheme is also correspondingly slower in reducing power compared to the optimal. This is due to the effect of thresholds; as long as the weighted average remains between T_L and T_H , the scheme maintains the same power level. This is reflected in the RSSI; the practical scheme maintains RSSI at a higher level than the optimal (in essence using more energy). In spite of this, our practical scheme saves 14% energy as compared to maximum power (using 21% more energy than the optimal). The price we pay is a decrease in reliability; the practical scheme results in 17% packet loss. This again illustrates the energy vs. reliability tradeoff inherent in practical schemes.

The practical scheme performs better for slow walk, saving 25% energy compared to maximum power, and uses only 13% more energy than the optimal. The packet loss is also correspondingly smaller at 11%. Once again the scheme exhibits the same inherent characteristics; there is a delay in reacting to a bad channel (due to the feedback nature of the scheme), and it maintains a higher power level than the optimal for some time in a good channel (due to our specific choice of T_H and T_L). For example, in the range 139 to 149 seconds the RSSI at the receivers remains

between T_L and T_H ; thus the scheme maintains transmit power at level 1 (while the optimal level is 0). Note however that this behaviour guards against a rogue sample in a predominantly bad channel, which may otherwise cause packet loss.

The practical scheme performs even better in the resting scenario, saving 33% energy compared to the maximum, and is within 2.5% of the optimal. Reliability is also highest among the three scenarios, with a packet loss of 5%. In this case the wireless channel remains relatively stable for significant periods, leading to large energy savings with respect to the maximum.

4.1. Discussion

This work establishes the feasibility of dynamic power control in a truly-wearable body area device. We have used the three scenarios as a guideline to measure the performance of power control under different operating conditions. Clearly, the performance of our scheme would depend heavily on the specific patient characteristics and operating conditions. For long term monitoring, we make the assumption (as in our previous work) that the patient spends an equal amount of time in the three scenarios. This results in an average power savings of 24% as compared to maximum power. Note that our algorithm is simple to implement at the base station, and these energy savings are therefore easy to achieve. Focusing again on body-worn disposable “digital plasters”, our scheme could result in almost 25% increase in the life of these devices, a handy increase.

At the same time our work allows us to bound the potential improvements possible using power control in body-worn devices. Note that the maximum savings achieved by the optimal scheme are 34% (compared to maximum power). Comparing this with innovations such as the ultra low power radio developed at Toumaz (which uses an order of magnitude less energy than the CC2420 radio in MicaZ motes) we see that dynamic power control, would be secondary to other energy saving techniques, but can be used in conjunction with other techniques to improve energy efficiency.

While we restricted our study so far to three illustrative scenarios, much further work is required in exploring the potential of dynamic power control 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, etc.

5. Conclusions and Future Work

In this work we build upon our earlier results on power control in body-worn devices. We profile wireless chan-

Scenario	Max (fixed)		Optimal		Aggressive	
	Energy per Packet (mW)	Packet Loss (%)	Energy per Packet (mW)	Packet Loss (%)	Energy per Packet (mW)	Packet Loss (%)
Fast walk	2.8	0.67	1.99	0.67	2.41	16.96
<i>better than "max" by:</i>	—		29%		14%	
<i>worse than "optimal" by:</i>	40.95%		—		21.20%	
Slow Walk	2.8	2.90	1.85	2.90	2.09	10.99
<i>better than "max" by:</i>	—		33.99%		25.40%	
<i>worse than "optimal" by:</i>	51.5%		—		13.01%	
Resting	2.8	1.26	1.84	1.26	1.88	5.37
<i>better than "max" by:</i>	—		34.46%		32.79%	
<i>worse than "optimal" by:</i>	52.57%		—		2.55%	

Table 2. Summary of energy and loss performance for various power control strategies

nel in a body-worn healthcare monitoring system using the Toumaz Sensium™, and assess the theoretical potential for saving energy without compromising reliability. Next we adapted our practical power control algorithm for use on the Sensium™, and showed that energy savings of 24% are possible on average without undue loss of reliability. We show that even with the reduced number of power levels and different operating frequency of the Sensium™ (as compared with MicaZ motes), power control remains as effective in saving energy. Since our power control algorithm is easy to implement, these savings are easily achievable. While our current work focused on continuous long-term patient monitoring, in future we plan to study the efficacy of power control in more dynamic settings such as athlete monitoring.

References

- [1] M. Avvenuti, P. Corsini, P. Masci, and A. Vechio. Energy-efficient reception of large preambles in mac protocols for wireless sensor networks. *Electronics Letters*, 43(5), Mar 2007.
- [2] Chipcon. CC2420: 2.4 GHz IEEE 802.15.4 / ZigBee-ready RF Transceiver.
- [3] L. Correia et al. Transmission Power Control in MAC Protocols for Wireless Sensor Networks. In *ACM/IEEE MSWiM*, pages 282–289, Montreal, Canada, Oct 2005.
- [4] Crossbow-Technologies. Mica2 and MicaZ motes.
- [5] D. Culler, D. Estrin, and M. Srivastava. Overview of Sensor Networks. *IEEE Computers*, 37(8):41–49, Aug 2004.
- [6] K. Langendoen and G. Halkes. *Embedded Systems Handbook*, chapter Energy-Efficient Medium Access Control. CRC Press, 2005.
- [7] 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.
- [8] 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*, pages 29–32, Cambridge, UK, Aug 2007.
- [9] J. Polastre, J. Hill, and D. Culler. Versatile Low Power Media Access for Wireless Sensor Networks. In *ACM SenSys*, pages 95–107, Baltimore, MD, Nov 2004.
- [10] Power-Paper. Power Patch Platform.
- [11] V. Schnayder et al. Sensor Networks for Medical Care. Technical Report TR-08-05, Division of Engineering and Applied Science, Harvard University, 2005.
- [12] 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*, pages 289–298, Santa Clara, CA, Oct 2004.
- [13] 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.
- [14] 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.
- [15] J. Zhao and R. Govindan. Understanding Packet Delivery Performance in Dense Wireless Sensor Networks. In *ACM SenSys*, Los Angeles, Nov 2003.
- [16] G. Zhou, T. He, S. Krishnamurthy, and J. Stankovic. Impact of Radio Irregularity on Wireless Sensor Networks. In *ACM MobiSys*, pages 125–138, Boston, MA, Jun 2004.