

Semantic Multimodal Compression for Wearable Sensing Systems

Saro Meguerdichian, Hyduke Noshadi, Foad Dabiri, and Miodrag Potkonjak
Computer Science Department, University of California, Los Angeles
{saro, hyduke, dabiri, miodrag}@cs.ucla.edu

Abstract—Wearable sensing systems (WSS’s) are emerging as an important class of distributed embedded systems in application domains ranging from medical to military. Such systems can be expensive and power hungry due to their multi-sensor implementations that require constant use, yet by nature they demand low-cost and low-power implementations. Semantic multimodal compression (SMC) mitigates these metrics in terms of data size by leveraging the natural tendency of signals in many types of embedded sensing systems to be composed of phases. In our driving example of a medical shoe with an insole lined with pressure sensors, we find that the natural airborne, landing, and take-off segments have sharply different and repetitive properties. SMC models and compresses each segment independently, selecting the best compression scheme for each segment and thus reducing total transmission energy.

I. INTRODUCTION

Wireless embedded networked sensing has been a realm of great technical interest over the past decade. Sensor networks have enabled real-world, practical applications of remote environmental monitoring and control in the home or in the field for personal, business, and military use. More recently, the research effort has shifted to wearable sensing systems (WSS’s), where the sensing “environment” in question is no longer remote but is instead very personal: the human body. WSS’s borrow the same core concepts from traditional wireless sensor networks, but they introduce a whole new set of concerns and challenges which demand innovation, including comfort, mobility, safety, and privacy.

Next-generation WSS’s can range from full-body military “smart suits” for infantrymen to small disposable patches that monitor physiological parameters while a patient sleeps to detect sleeping disorders. In fact, a majority of the recent research in WSS’s has been for medical devices, resulting in the creation of the new booming field of wireless health. Much of the effort in the wireless health domain has been geared toward aiding medical professionals in detecting and diagnosing ailments in patients by removing the patient from a static monitoring environment (e.g. the hospital) and injecting constant non-invasive monitoring into the patient’s daily life. This can result in more reliable, comprehensive, and long-term data which is more likely to be indicative of a patient’s true habits, since the patient is not necessarily always aware of the direct monitoring.

Although wireless health and WSS’s in general have been the recent “hot spot” for technical innovation, industrial realization of the field has nevertheless been slow in compar-

ison. This is largely due to the complexity in designing a truly deployable WSS, which must meet the aforementioned challenges and constraints while keeping reliability high and cost, energy, and operational complexity low. In this paper, we address the energy problem in the data transmission stage by leveraging unique properties of WSS’s to compress signals in a novel way. We use as our driving example the Hermes platform, which is a medical shoe with an insole lined with pressure sensors [1].

An analysis of the data from users of the shoe reveals a crucial observation that can generally apply to many embedded sensing systems and in particular to WSS’s: these systems produce signals that are subject to natural phases. For example, a heart rate sensor on someone who is exercising exhibits warm-up, workout, and cool-down phases. In the Hermes platform, pressure sensors in the shoe follow natural human motion, in airborne, landing, and take-off phases.

The key contribution of this paper is to demonstrate that we can take advantage of the existence of these phases to design more energy efficient WSS’s, by segmenting a signal into its corresponding phases and processing each segment independently. Essentially, by segmenting the signal we can reduce its entropy, since the entropy within one segment is likely to be significantly smaller than that of all the segments together. This observation is key; information entropy is a commonly used metric for estimating the upper bound for compressibility of data, and is defined in Section IV. Generally, lower entropy in the data allows for a higher compression ratio.

In this paper, we develop the notion of semantic multimodal compression (SMC), where we segment the signal into its natural phases and then compress each segment independently. In Section III-A, we highlight the relevant details of the Hermes platform, our driving example, then describe its phases in Section III-B. In Section III-C, we discuss some important real-time system parameters and define our assumptions and system constraints. We then develop our compression algorithm in Section IV and demonstrate experimental results in Section V. Finally, we propose some ideas to relax our assumptions in Section VI and conclude in Section VII.

II. RELATED WORK

Numerous scientific and engineering topics have been explored within the framework of embedded sensor networking and computational sensing [2] [3] [4] [5] [6] [7] [8] [9] [10] [11]. In particular, much of the attention has been on reducing

energy in wireless sensor networks, as covered by Anastasi, et al. in the survey paper [12].

WSS's, medical and otherwise, are a sub-domain of sensor networking that has been attracting a great deal of recent research [13] [14] [15]. There has also been prior research driven toward gait analysis via WSS's (e.g. sensor-equipped shoes) [16] [17] [18]. Surprisingly, although there has been much research in energy reduction in WSS's, even for medical shoe implementations, there has been little effort to leverage the unique properties of signal behavior in WSS's to enhance the compression ratio in particular.

Compression is one of the most popular information theory topics. While initially it was mainly used to reduce storage requirements, recent emphasis has been on bandwidth and energy savings. There are many general and special purpose (e.g. text, video, and image) textbooks and surveys [19] [20] [21] [22]. Several techniques that address compression in distributed embedded sensing platforms have also been proposed [23] [24] [25] [26] [27] [28] [29] [30] [31] [32]. To the best of our knowledge, this is the first effort to use semantic signal segmentation to improve compression rate and reduce energy requirements.

III. PRELIMINARIES

A. The Hermes Platform

Hermes is a lightweight smart shoe designed to assess mobile balance and instability in human patients [14]. It consists of low-cost off-the-shelf sensors integrated with a common embedded computing platform, and can interface with a smartphone with bluetooth and GPRS capabilities for data processing, storage, and transmission. Specifically, the shoe has an insole lined with either passive resistive sensors from Tekscan [33] or piezoresistive fabric from Eeonyx [34], along with gyroscopes and accelerometers, communicating with a MicroLEAP mote [35] that samples data at 60 Hz. The MicroLEAP processor consumes 2.7 mW of power while the radio, with a data rate of 250 kbps, consumes 57.5 mW. For sensor placement, we use the Pedar plantar pressure mapping system [36], which uses a grid of 99 pressure sensors distributed across the foot. Figure 1 depicts the Hermes shoe, its insole, and the Pedar sensor map.

This comprises a very powerful system for human gait, balance, and stability analysis, which can be used for a number of crucial applications ranging from fall prevention to shoe insole design for flat-footed patients. However, high data sampling and transmission rates mean that this system will have a short lifetime and high cost of operation, and thus becomes another example of technological innovation being slow to achieve industrial realization. Therefore, our goal in this paper is to take the next step in this system, to significantly reduce its energy consumption rate in the data storage and transmission stage of operation. To facilitate this, we collected extensive data from 8 people performing repeated sets of various tasks, such as running, walking, and limping.



Fig. 1. Hermes shoe platform (right) and Pedar sensor mapping (left) [1].

B. Signal Phases and Predictability

An example signal from one sensor on one Hermes shoe is shown in Figure 2, with applied pressure on the y-axis and time on the x-axis. Here, the subject is walking normally. We can clearly see that each “blip” in the signal corresponds to a single step, which is comprised of the following phases:

- (i) *Airborne (blue)*. In the airborne phase, the foot is applying essentially zero pressure onto the sensor; this corresponds to the period of time when the foot is completely off the ground.
- (ii) *Landing (red)*. This is the time between when the heel first touches the ground to when the toe touches the ground. We see this as a near-quadratic incline (with decreasing slope) in pressure applied to the sensor.
- (iii) *Take-off (green)*. The take-off phase corresponds to the foot leaving the ground, from when the heel begins to rise to when the toe is completely off the ground. The signal in this phase is an essentially linear decline in reported pressure.

We can also see from Figure 2 that the sensor signals are highly predictable. Phase and amplitude on the same sensor are relatively constant between consecutive steps. Furthermore, there is high correlation between some sets of sensors. The two sensors shown in Figure 3 are sensors 9 (blue) and 86 (red) in the sensor map in Figure 1. Since both sensors are located on the same longitudinal axis in the middle of the foot, one the heel and one in front of it on the toe, pressure is applied to one with a relatively constant time offset from the other.

Both the predictability of the sensor signals and the distinctness of their phases are attributed to the cyclic, rhythmic nature of human behavior. Thus, these phenomena are found across different types of WSS's, and are one example of why WSS's can be an area of tremendous innovation. In the following sections, we use the signal segmentation and predictability described here to develop SMC and reduce transmission energy.

C. Assumptions and Constraints

For our implementation of SMC, we make a number of preliminary assumptions and set some usage constraints. First, we observe that the MicroLEAP's radio consumes energy

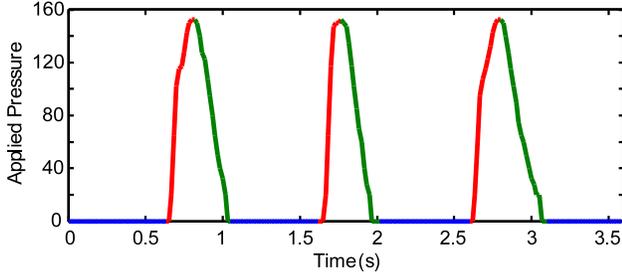


Fig. 2. Example signal from a single sensor over 3 steps. The three phases are clearly visible: airborne (blue); landing (red); and take-off (green). The signal maintains rough phase and amplitude between steps.

at a much higher rate (over 20x) than its processor; hence, we limit our algorithm and evaluation to transmission energy reduction under the assumption that the benefit here will far outweigh the small energy cost in processing and storage for our compression algorithm.

Furthermore, for this discussion, we assume that the real-time latency and storage constraints are such that we can store, process, and transmit entire steps (all 3 phases) at a time. This is a reasonable assumption in practice because much of the useful gait characteristic information is contained in the entire step, and steps are quick enough that storage is not a problem. Nonetheless, in Section VI we propose at a higher level some modifications to our SMC algorithm that will relax these assumptions and be effective under stricter constraints.

IV. SEMANTIC MULTIMODAL COMPRESSION (SMC)

In our driving example of the medical shoe, we have three segments—airborne, landing, and take-off—as described in Section III-B. A typical step of a signal obtained from a single sensor is depicted in red in Figure 4, and calculated values for its entropy are in the first column of Table I. In this case, information entropy is defined as:

$$h(S) = - \sum_{i=1}^n p_i \log p_i, \sum_{i=1}^n p_i = 1 \quad (1)$$

for a string S with n symbols, each occurring with probability p_i [37]. In general, compressibility of a signal is inversely proportional to its entropy; a signal with high entropy has less data redundancy and is therefore more difficult to compress. As expected, entropy for each segment is indeed lower than the entropy of the full signal. Note that the airborne segment has 0 entropy, which suggests the possibility for high compressibility via segmentation.

Furthermore, we can leverage the predictability of the data to reduce entropy of the landing and take-off segments even more without losing data. We describe this process in Subsection IV-A. Then, in Section V, we present typical results for our compression method on data from the medical shoe, showing decreased transmission energy when compared to a common basic compression scheme.

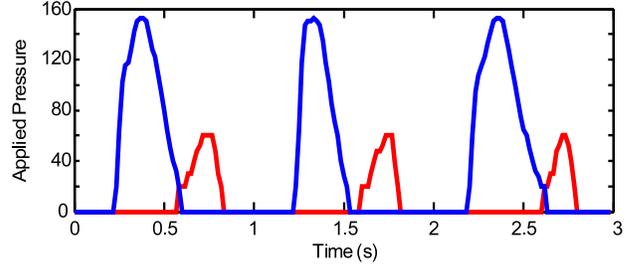


Fig. 3. Example signal from two sensors over 3 steps, with one sensor on the heel (blue) and one on the toe (red). The heel sensor touches down first, then the toe sensor, with a relatively constant phase shift.

A. Entropy Reduction via Signal Approximation

From Figure 2 we can see that a signal will maintain its rough shape, phase, and amplitude over a series of steps. Therefore, if we have some knowledge about the expected behavior of the signal, we can potentially reduce its entropy by an order of magnitude in the common case. Figure 4 demonstrates this; the red curve is the original signal, the blue curve is a simple piecewise linear model, and the green curve is the difference between them. The second column of Table I enumerates the new entropy values for the approximated (green) signal, which are significantly lower than those of the original.

To obtain the piecewise linear model, we only need the following 4 values: (1) the number of airborne samples; (2) the number of landing samples; (3) the number of take-off samples; and (4) the value of the sample with peak amplitude. With these 4 metrics, the original signal can be reconstructed without loss from the compensated signal. We show that it is beneficial for transmission energy consumption to encode and compress these 4 metrics and the compensated signal, as opposed to just the original signal, due to the large extent that entropy is thus reduced.

There is, however, further room for improvement. The 4 metrics described above are themselves predictable; over time, we can develop and store an expected value for each of the 4 metrics. Then, again, instead of encoding the values of the metrics themselves, we can encode the difference between each value and the expected value of the metric. Doing so will drive the encoded values closer to 0 and is more likely to create repetition and reduce entropy.

Storing a separate set of expected values for the 4 metrics for each individual will reduce the entropy of the signal and thus improve the quality of compression and decrease transmission energy more than storing a general set of expected values. This is intuitive, since averaging over the entire population will introduce noise into the expected metrics for a single individual. For example, we are likely to expect significantly different metrics from someone with flat feet compared to someone without flat feet, and in turn compared to someone with a limp.

Doing so, however, will require more storage and will likely be intractable for systems that support multiple individuals; on

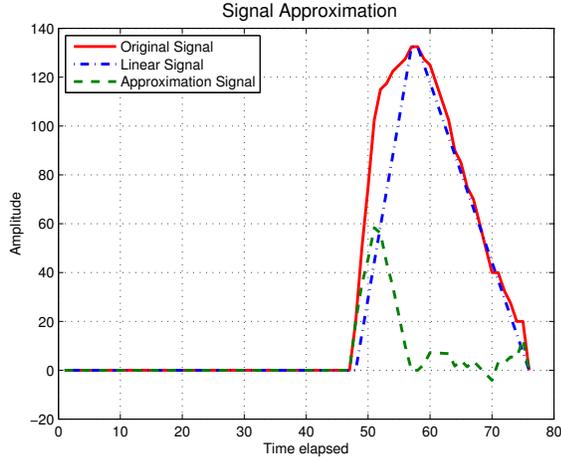


Fig. 4. We can approximate a single step to reduce entropy. The resulting approximated signal (green) is the difference between the original signal (red) and a simple piecewise linear model (blue).

TABLE I
ENTROPY PER SEGMENT FOR A SINGLE STEP FOR ORIGINAL AND APPROXIMATED SIGNALS.

	Original	Approximated
Airborne Segment	0	0
Landing Segment	$6.7 * 10^3$	$1.7 * 10^3$
Take-off Segment	$8.5 * 10^3$	$1.8 * 10^2$
Full Signal	$1.5 * 10^4$	$1.9 * 10^3$

the other hand, it is not unreasonable to assume that wearable sensor systems such as the medical shoe will be custom tailored for one individual. In our experiments, the benefit gained from storing individual metrics were negligible when compared to total transmission energy, so we limit further discussion to global metrics for this example.

Regardless, now that we have reduced entropy as much as we can without losing data or requiring extensive computation, we are ready to compress the new data, which consists of the following: a series of 4 values, as described above, and the approximated landing and take-off segments of the signal. To preserve our ability to decompress, we assume that the decoding party also has the established expected values for the 4 metrics. For compression, we can use any effective compression algorithm; for the medical shoe data, we use arithmetic coding (AC) [38].

V. TRANSMISSION ENERGY SAVINGS FROM SMC

To demonstrate the effectiveness of our entropy reduction techniques, we compare results from data compressed using SMC to the same data compressed using just AC. The effectiveness of compression will depend on the size of the chunk of data being compressed via AC. Typically, increasing the chunk size will to an extent increase redundancy, thereby reducing entropy and increasing the quality of compression.

Figure 5 depicts these results for a typical signal with 6000 samples. The red and blue curves show the transmission

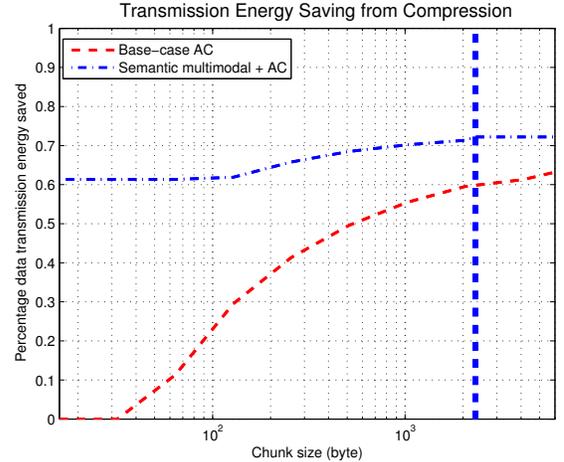


Fig. 5. SMC with AC (blue) reduces transmission energy significantly for all sizes, compared to vanilla AC (red). The vertical blue dotted line represents the pre-compression data size reduction from SMC before AC is applied.

energy savings from base-case AC compression and SMC, respectively, for various chunk sizes. Note that the base-case compression is already quite effective, since the natural redundancy and predictability (and thus low entropy) of the signal make it ideal for compression. Base-case AC compression shows up to a 63.2% savings in transmission energy if all 6000 segments are compressed at once, decreasing to 0 as the chunk size decreases. Nonetheless, using SMC, we reduce transmission energy by up to 9% over the base-case compression at its best, and 72.2% in total.

It is interesting to note that even when AC is ineffective at smaller chunk sizes, SMC provides great reduction in transmission energy. Without using AC at all, SMC still reduces transmission energy by 61.4%, which is almost as good as the base-case AC compression scheme under optimal conditions (maximum chunk size). This is due to the fact that for a given AC chunk size, SMC actually encodes more samples than the base-case compression, since it pre-compresses the data to reduce entropy for AC on-the-fly. The blue vertical line represents the chunk size (2319, 61.4% less than 6000) at which SMC has already encoded the entire 6000-sample signal. Therefore, to reduce computation costs, SMC can be applied with chunk size 0 (no AC) and still save transmission energy significantly.

VI. FUTURE WORK

In presenting our algorithm, we have made a series of assumptions, outlined in Section III-C, that are reasonable in practice for this particular system, but may not scale to the general case. These assumptions exist in our implementation of SMC because we must have the entire step in order to create the piecewise linear model. In other words, we must have the values of all 4 metrics before we begin compression. However, this may be relaxed by using a best-estimate value for metrics which we could not obtain within the allotted time. We can do this based on the recent history of the signal; if the

patient is walking at a brisk, steady pace, for example, we can estimate that the metrics will be constant for the next step and use the previous ones to create our piecewise linear model. With this method, we will likely sacrifice some compression ratio to meet the real-time constraints.

Alternatively, we could use real-time data from some sensors to estimate the metrics of another. We describe an example in Section III-B where applied pressure triggers one sensor at a constant time offset before another over a series of steps. Therefore, we can base our model off of the predictor sensor, with a phase shift, to create our approximation signal. Here, we meet the real-time constraints, and may even improve the compression ratio, depending on how strongly the two sensors are correlated. For example, if one sensor predicts another perfectly and reliably, the approximation signal for the predicted signal will essentially be very close to 0, with ultra-low entropy and thus high compressibility.

VII. CONCLUSION

We have presented a method for semantic multimodal compression that leverages the notion that signals in many embedded and wireless sensor networks have natural phases and predictable characteristics. We have applied SMC to the domain of wearable sensing systems, where the phases and predictability are derived from natural human behavior. In this paper, we have implemented this algorithm for the specific case of the Hermes medical shoe, which has an insole lined with pressure sensors that can be used to essentially map how a patient applies pressure to the ground during daily activities such as walking and running. We have also described how the SMC paradigm can be extended to other WSS's with different constraints or to sensor networks in general which display these characteristics. Our main focus has been to improve the compression ratio and thus reduce transmission energy, and have shown that SMC in our case study provides a 72.2% reduction in total transmission energy, up to 9% over the base-case compression benchmark at its best.

REFERENCES

- [1] H. Noshadi, S. Ahmadian, H. Hagopian, J. Woodbridge, N. Amini, F. Dabiri, and M. Sarrafzadeh, "HERMES: Mobile balance and instability assessment system," *BioSignals*, 2010.
- [2] S. Meguerdichian, F. Koushanfar, M. Potkonjak, and M. Srivastava, "Coverage problems in wireless ad-hoc sensor networks," *Infocom*, vol. 3, pp. 1380-1387, 2001.
- [3] S. Megerian, F. Koushanfar, M. Potkonjak, and M. B. Srivastava, "Worst and best-case coverage in sensor networks," *TMC*, vol.4, no. 1, pp. 84-92, 2005.
- [4] S. Meguerdichian, S. Slijepcevic, V. Karayan, and M. Potkonjak, "Localized algorithms in wireless ad-hoc networks: location discovery and sensor exposure," *MobiHoc*, pp. 106-116, 2001.
- [5] S. Slijepcevic and M. Potkonjak, "Power efficient organization of wireless sensor networks," *ICC*, pp. 472-476, 2001.
- [6] G. Veltri, Q. Huang, G. Qu, and M. Potkonjak, "Minimal and maximal exposure path algorithms for wireless embedded sensor networks," *Sensys*, pp. 40-50, 2003.
- [7] J. L. Wong, R. Jafari, and M. Potkonjak, "Gateway placement for latency and energy efficient data aggregation," *LCN*, pp. 490-497, 2004.
- [8] F. Koushanfar, M. Potkonjak, and A. Sangiovanni-Vincentelli, "Fault tolerance techniques in wireless ad-hoc sensor networks," *Sensors*, pp. 1491-1496, 2002.
- [9] J. L. Wong and M. Potkonjak, "Search in sensor networks: challenges, techniques, and applications," *ICASSP*, vol. 4, pp. 3752-3755, 2002.
- [10] J. Feng, F. Koushanfar, and M. Potkonjak, "System-architectures for sensor networks: issues, alternatives, and directions," *ICCAD*, pp. 112-121, 2002.
- [11] J. Feng, S. Megerian, and M. Potkonjak, "Model-based calibration for sensor networks," *Sensors*, pp. 737-742, 2003.
- [12] G. Anastasi, M. Conti, M. Di Francesco, and A. Passarella, "Energy conservation in wireless sensor networks: A survey," *Ad Hoc Networks*, vol. 7., no. 3, pp. 537-568, 2009.
- [13] A. Krause, D. P. Siewiorek, A. Smailagic, and J. Farringdon, "Unsupervised, dynamic identification of physiological and activity context in wearable computing," *ISWC*, pp. 88-97, 2003.
- [14] H. Noshadi, E. Giordano, H. Hagopian, G. Pau, M. Gerla, and M. Sarrafzadeh, "Remote medical monitoring through vehicular ad hoc network," *VTC*, pp. 1-5, 2008.
- [15] K. Lorincz, B. Chen, G. W. Challen, A. R. Chowdhury, S. Patel, P. Bonato, and M. Welsh, "Mercury: a wearable sensor network platform for high-fidelity motion analysis," *SenSys*, pp. 183-196, 2009.
- [16] S. Popovic, M. R. Dietz, V. Morari, M. Pappas, I. Keller, and T. Mangel, "A reliable gyroscope-based gait-phase detection sensor embedded in a shoe insole," *Sensors*, vol. 4, no. 2, pp. 268-274, 2004.
- [17] K. Oshima, Y. Ishida, S. Konomi, N. Thepvilajanapong, and Y. Tobe, "A human probe for measuring walkability," *SenSys*, pp. 353-354, 2009.
- [18] V. Erickson, A. U. Kamthe, and A. E. Cerpa, "Measuring foot pronation using RFID sensor networks," *SenSys*, pp. 325-326, 2009.
- [19] J. Storer, *Data Compression Methods and Theory*. Computer Science Press, Rockville, MD, 1988.
- [20] T. Bell, I. H. Witten, and J. G. Cleary, "Modeling for text compression," *CSUR*, vol. 21, no. 4, pp. 557-591, 1989.
- [21] T. C. Bell, J. G. Cleary, and I. H. Witten, *Text Compression*. Prentice-Hall, Upper Saddle River, NJ, 1990.
- [22] A. Gersho, and R. M. Gray, *Vector Quantization and Signal Compression*. Springer, New York, NY, 1991.
- [23] J. Kusuma, L. Doherty, and K. Ramchandran, "Distributed compression for sensor networks," *ICIP*, vol. 1, pp. 82-85, 2001.
- [24] A. Scaglione and S.D. Servetto, "On the interdependence of routing and data compression in multi-hop sensor networks," *MobiCom*, pp. 140-147, 2002.
- [25] K. Barr and K. Asanovic, "Energy aware lossless data compression," *MobiSys*, pp. 231-244, 2003.
- [26] D. Petrovic, R. Shah, K. Ramchandran, and J. Rabaey, "Data funneling: Routing with aggregation and compression for wireless sensor networks," *SNPA*, pp. 156-162, 2003.
- [27] S. Pattem, B. Krishnamachari, and R. Govindan, "The impact of spatial correlation on routing with compression in wireless sensor networks," *IPSN*, pp. 28-35, 2004.
- [28] N. Kimura and S. Latifi, "A survey on data compression in wireless sensor networks," *ITCC*, vol. 2, pp. 8-13, 2005.
- [29] C. M. Sadler and M. Martonosi, "Data compression algorithms for energy-constrained devices in delay tolerant networks," *SenSys*, pp. 265-278, 2006.
- [30] J. Xiao, A. Ribeiro, Z. Luo, and G. B. Giannakis, "Distributed compression-estimation using wireless sensor networks," *IEEE Signal Processing Magazine*, vol. 23, no. 4, pp. 27-41, 2006.
- [31] S. Gandhi, S. Nath, S. Suri, and J. Liu, "GAMPS: compressing multi sensor data by grouping and amplitude scaling," *SIGMOD*, pp. 771-784, 2009.
- [32] G. Reeves, J. Liu, S. Nath, and F. Zhao, "Managing massive time series streams with multi-scale compressed trickles," *VLDB Endowment*, vol. 2 no. 1, pp. 97-108, August 2009.
- [33] Tekscan pressure sensor, [Online document], Available HTTP: <http://www.tekscan.com>.
- [34] Eeonyx resistive fabric, [Online document], Available HTTP: <http://www.eeonyx.com>.
- [35] L. K. Au, W. H. Wu, M. A. Batalin, D. H. McIntire, and W. J. Kaiser, "MicroLEAP: Energy-aware wireless sensor platform for biomedical sensing applications," *BIOCAS*, pp. 158-162, 2007.
- [36] Pedar: plantar pressure system, [Online document], Available HTTP: <http://www.novel.de>.
- [37] G. Hansel, D. Perrin, and I. Simon, "Compression and entropy," *STACS*, pp. 515-528, 1992.
- [38] I. H. Witten, R. M. Neal, and J. G. Cleary, "Arithmetic coding for data compression," *CACM*, vol. 30 no. 6. pp.520-540, 1987.