

Analyzing the Use of E-textiles to Improve Application Performance

Mark Jones, Tom Martin, Zahi Nakad, Ravi Shenoy, Tanwir Sheikh, David Lehn, Joshua Edmison,
Madhup Chandra
Virginia Tech
Bradley Department of Electrical and Computer Engineering
Blacksburg, VA 24061

E-mail: {tlmartin, mtj}@vt.edu

Phone: (540) 231 1739

Fax: (540) 231 8292

Abstract: E-textiles are an alternative to radio-based personal area networks, with potentially significant advantages in hardware cost and energy consumption, through a reduction in the number of communication modules and the energy consumed by data transmission. This paper presents a methodology based on simulation and physical experimentation for quantifying this savings on a large e-textile that determines the bearing/location of one or more distant audible targets.

1 Introduction

Electronic textiles (e-textiles) are fabrics that have electronics and interconnections woven into them. Interest in e-textiles has arisen due to the recognition of the manufacturing capabilities and economies of the textile industry as well as the increasingly sophisticated fiber technologies being produced by materials science research. E-textiles offer the promise of computing elements, sensors, and actuators seamlessly integrated into familiar form factors such as shirts, hats, parachutes, and blankets. New fibers are being produced that function as batteries, durable wires, and speakers [7]; new packaging techniques for digital components allow for manufacturable integration in a durable, washable textile [4]. Current e-textile research is examining the applications and computing infrastructure for e-textiles [5][6].

E-textiles provide an alternative to radio-based personal area networks, with potentially significant advantages in terms of hardware cost and energy consumption. This savings arises from a reduction in the number of communication modules required as well as from the lower energy requirements for transmission over the physical media. This paper presents and carries out a methodology for quantifying this savings.

To provide a meaningful analysis, it is necessary to select benchmark applications. In a more mature research area, such benchmarks would be readily available. However, because the field is still in its infancy, meaningful benchmarks have yet to be developed. The benchmark used in this paper is a large e-textile (approx. 10 meters) that can determine the bearing and/or location of one or more distant audible targets. Energy consumption must be minimized because the device should operate untethered for days or weeks in the field without human intervention. This application was selected because it is

documented in the literature, we have prototypes with which to work, and, we argue in the following section, it is representative of a range of potential e-textile applications.

The methodology presented in this paper uses a combination of simulation and physical experiments to perform the analysis. The design state space to be explored is quite large, necessitating the use of simulation to effectively explore it. Given the immature state of e-textile technology, it is necessary to ground this simulation in reality through experiments with physical prototypes.

In Section 2, the application is presented and its suitability as a benchmark is discussed. In Section 3, physical experiments carried out on prototypes are presented. Three alternative communication schemes are described and their energy consumption is measured: (1) a pure wireless solution based on Bluetooth (2) a “wires-only” implementation in e-textiles, and (3) a hybrid of the two methods. In addition to these measurements, the potential range of component costs, capabilities, and energy consumption is delineated based on the literature as well as analytical analysis. In Section 4, a simulation environment based on Ptolemy II [9] is described that allows for the exploration and evaluation of a wide range of designs using each of the three schemes. Results from the simulation experiments are presented to quantify the potential savings. A summary and conclusions are given in Section 5.

2 Benchmark Application

For this investigation, we examine a large-scale acoustic beamformer. Because the application can employ differing numbers of sensors in many different configurations, it allows for several physical architectures for the fabric to be explored. In addition, the application can be solved with schemes that vary in their balance between computation and communication.

The acoustic beamformer uses an array of microphones to determine the angle from which an acoustic signal is arriving. This is accomplished via beamforming where the differences in signal phase at each of the microphones are compared as shown in Figure 1. Fundamental to this acoustic application is the dependence of the accuracy of the computed direction on the sampling rate as well as on the position and number of microphones. If the microphones are too far apart, then aliasing will occur and the phase difference will be predicted

incorrectly; this distance is an inverse function of the frequency. If the microphones are too close together, then the difference detected in phase will be small relative to the error in determining phase. Similarly, if the sampling rate is low, then the number of samples per wave is small, leading to

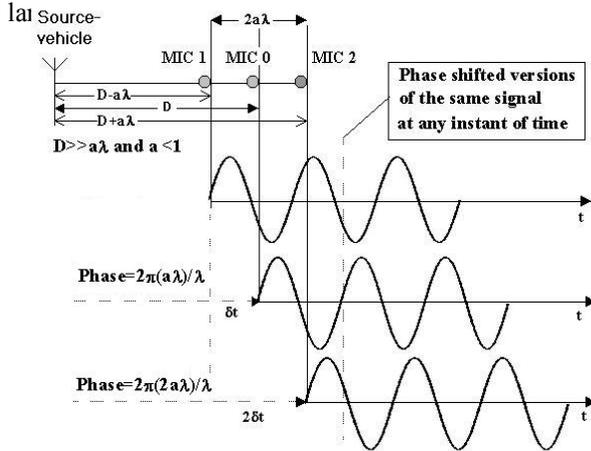


Figure 1: An illustration of the difference in phase in the same signal arriving at different microphones.

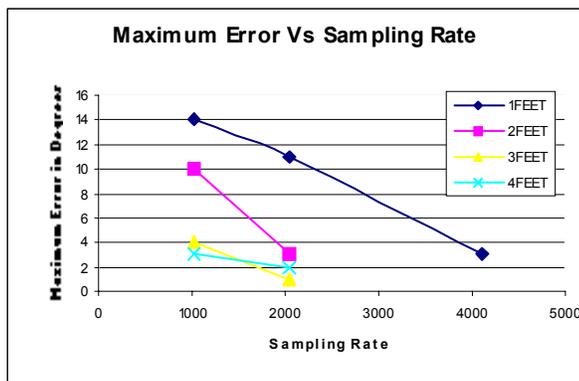


Figure 2: The error in computed beam direction for a 100 Hz signal as a function of the sampling rate for different microphone spacings.

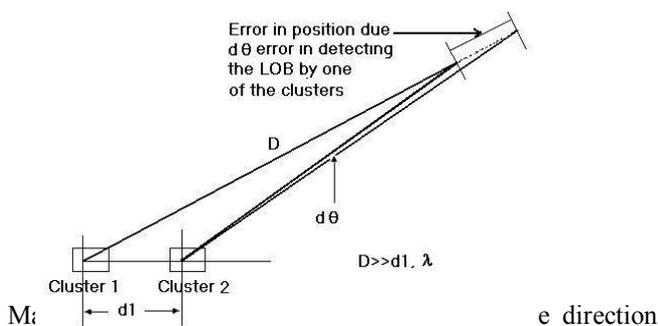


Figure 3: A small distance between clusters can lead to a large error in computed position.

requires that the position of the microphones relative to one another be known. The flexibility of an e-textile poses a challenge as the microphones may change position with respect to one another. The beamformer may be completely free to move in any direction (depending on the deployment method) and small errors in microphone position information may lead to large errors in computing direction. A solution to this problem is discussed in [5].

Effectively managing energy consumption while maintaining acceptable accuracy involves selecting a set of microphones to be sampled that allows for the proper spacing for the expected frequency, while compensating for any problems with microphone spacing by (a) increasing the sampling rate and/or (b) filtering out frequencies high enough to result in aliasing. The graph in Figure 2 shows the effect of microphone spacing and sampling rate on the accuracy of the computed direction (results collected from the simulation in Section 4). The set of microphones from which to select may not be a choice free of other costs as it may involve using multiple processing units as well as communicating time series between these units.

This type of management is representative of what is required in other e-textile applications in which a large set of sensors/actuators exists with which to accomplish a task. In most applications, using all of the sensors will require a large amount of computation when results of the same quality could be computed at much lower cost from the correct subset of sensors.

The acoustic beamformer determines the location of a constantly emitting acoustic source; for example, a large vehicle. It computes the position by combining the directions computed by multiple clusters within the textiles; these results are combined via triangulation as shown in Figure 3. Each cluster of microphones can compute the direction of an acoustic source. Theoretically, directions from only two clusters are required to triangulate the location of a source, but by combining multiple results, the answer can be improved upon and faults tolerated. If the clusters are spaced closely together, then small errors in the directions computed by the clusters will lead to large errors in the calculated position (see Figure 3); to combat this, the e-textile is constructed to be long. Note that direction can be computed accurately with a single cluster; it is only position that requires the very large fabric.

3 Experimental Measurements

This section describes the prototypes and apparatus for experimentally measuring power consumption for the benchmark applications. A set of potential textile architectures is described followed by measurements and analysis of the power consumption and cost of these components.

3.1 Textile Architectures

This paper will examine three basic architectures: a pure e-textile, wireless, and hybrid. The basic building block of each architecture is a microphone cluster, which consists of a set of microphones and a processing unit capable of communicating

with other processing units. Design choices to be made include the type of communication between processing units, the type of communication between processing unit and microphone, the number of clusters, and the number of microphones per cluster.

A pure e-textile architecture uses the fabric as the media for all communication. At the other extreme is a pure wireless architecture for all communication; each microphone would have an associated processor and wireless unit. Combining the two results in a hybrid architecture in which microphone-processor communication occurs via the fabric and processor-processor communication is wireless. Of these architectures, only the pure wireless requires significant communication to perform meaningful work. Each processor must acquire data from at least three microphones to form a beam, and a processor must receive an acoustic time series from two other processors to compute an angle.

Once angles are formed within clusters, the communication requirement between clusters is limited to the results of the beamformer (angle, signal level, and cluster location). The clusters exchange beamforming results to determine the location of the target and to improve the quality of the results through the use of redundant information.

There are several components to the system which each behave differently in different modes of operation. These components are the analog filter/amplifier, the A/D unit, and the DSP. For the system, we define four basic state of operation: sleep, data collection, beamforming, and communication. The table below depicts which units are actively consuming significant power in each state.

State	Analog	A/D	DSP
Sleep	Yes	No	No
Data collection	Yes	Yes	Minimal
Compute Angle	Yes	No	Yes
Send Results	Yes	No	Minimal

3.2 Current Prototype

We have constructed multiple acoustic beamforming prototype textiles, allowing us to measure power of each subsystem. These prototypes are based around a DSP-based data collection/processing unit that attaches to the textile. This unit is capable of sampling up to seven microphones and communicating with up to six other data collection/processing units. The primary components of the unit are an Analog Devices ADSP-2188M, the AD7888 A/D converter, and analog circuits for interfacing to each of the microphones (see the diagram in Figure 4). The operating voltage for this unit can be substantially reduced in size and repackaged for a more acceptable form factor.

This unit can be in one of two basic states: (1) processing and data collection are asleep and only the analog interface is active, and (2) the entire system is actively collecting and analyzing data. The DSP and A/D are activated by the analog interface whenever a sound signal of sufficient intensity is received or when another unit sends a signal to the sleeping unit.

In addition to the data collection/processing unit, the textiles

have multiple microphones and batteries attached to fiber wires woven alongside cotton fibers. The fiber wires are strong, flexible structures constructed from many 100% stainless steel fibers with a resistivity of 10 Ohms per meter. These fibers, supplied by Bekintex, come in a variety of other metals resulting in differing strengths and resistivity [3].

A photograph of a single cluster textile is given in Figure 5. This prototype allows for the power characteristics of the system in operation to be characterized, including the cost of analog interface circuits, data collection, data processing, and communication between units. In the following subsections, we experimentally characterize the power consumption of these prototypes.

The experimental set up for making the power measurements described in the following sections consisted of an Agilent 3458A digital multimeter capable of 100K samples/second connected to a PC via a GPIB card [1]. The following sections describe the power measurements of the subsystems of the acoustic beamformer.

3.3 Analog Interface Circuits

These circuits are always on to allow for one or more microphones to awaken the processor out of its sleep state. A low-pass filter and op-amp circuit is applied to each microphone input that is then passed to the A/D converter and a simple digital wakeup circuit. The A/D converter is only active/awake when the DSP is awake. The power consumption of the analog circuit is constant over the operational modes of the system and was measured to be 9 mW.

3.4 Data Collection

The acoustic beamformer collects data when one or more microphones are detecting signals of a sufficiently high level; otherwise the DSP goes into a sleep state. During data collection, the DSP queries the A/D converter to sample all or a subset of the microphones. This requires interrupt handling and minor processing to fill up the appropriate buffer. The power consumption of data collection is dependent on both the sampling rate and the number of microphones, as shown in Table 1.

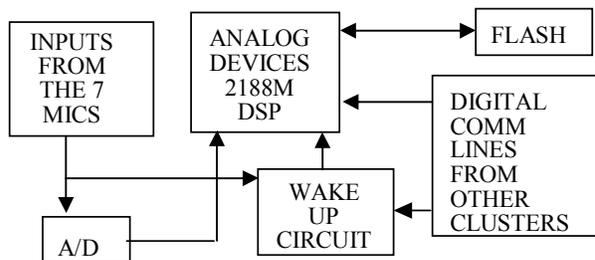


Figure 4: Block diagram of the signal collection and processing unit.

3.5 Data Processing

The data processing consists primarily of running a beamforming algorithm on the given data samples. The beamforming algorithm estimates the direction of an acoustic source given a set of samples from three or more microphones.

The beamforming algorithm that we have chosen is derived from an Army Research Laboratory algorithm modified to achieve low power consumption [10]. This algorithm assumes that the microphones are in the same plane as the acoustic source. The beamforming algorithm is run on the most recently collected set of samples from the microphones; typical sample lengths are in the 512-1024 range per microphone. The energy consumption of the beamforming algorithm is 200 mW, and the runtime is linearly proportional to the number of microphones and the length of the samples. Table 2 gives the energy consumption of the board during data processing for typical sample lengths and numbers of microphones. Note that data collection is also occurring in the background, but is not a significant contributor to the load (it results in a slight lengthening of runtime).

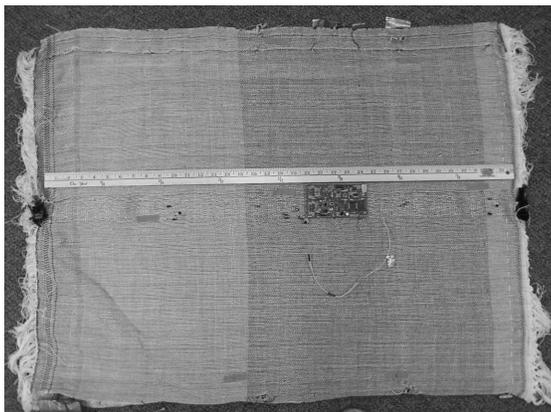


Figure 5: Photograph of single cluster e-textile beamformer prototype, approx. 1 meter square.
Table 1: Power consumption during data collection vs. sample rate and # of microphones.

Number of microphones	Sample per Second	Power Consumption(mW)
7	1000	61
7	2000	70
7	4000	88
7	8000	124
5	1000	59
5	2000	66
5	4000	79
5	8000	104
3	1000	57
3	2000	62
3	4000	70
3	8000	89
1	1000	55
1	2000	57
1	4000	62
1	8000	70

Table 2: Energy vs. number of microphones and sample length.

Number of Microphones	Sample Length	Runtime (msec)	Energy Consumed (mJ)
3	512	61	12
3	1024	76	15
3	2048	108	20
5	512	85	17
5	1024	98	19
5	2048	130	25
7	512	112	22
7	1024	127	25
7	2048	160	31

3.6 Communication on the Textile

In this section we will detail the power costs of sending and receiving on the processing node of the beamformer. For these measurements the data sampling and the beamforming code was disabled; only the communication was active.

The maximum length for any data communication is estimated to be less than 8 meters. The maximum baud rate of data transmission is dependant on the signal degradation. We found experimentally that 256 kbps was the maximum communication rate for which the signal integrity is preserved. Thus 256 kbps was used as the maximum communication rate for the power measurements of wired communication.

Power consumption measurements were carried out on different wire lengths, namely 2, 4, and 8 meters and different transmission speeds. An increase in power consumption would be expected when increasing the length of the communication channel (the wire in our case). Table 3 reports the values recorded with communication over the longest wire.

Sending and receiving on this processing node are controlled by a timer interrupt. Processing of this interrupt determines whether data is sent or received. "Sending" and "Receiving" columns indicate the values for transmitting and receiving respectively. The measurements were performed on independent transmitter and receiver boards, with disabled reception on the transmitter and vice versa.

Power consumption decreases as the communication rate decreases. Less processing in a specific time period is required for slower communication (lower timer interrupt rate). The extra power consumed on the transmitter is attributed to driving the communication channel.

Table 3: Power vs. mode and communication rates.

Communication Rate	Wait (mW)	Send (mW)	Recv (mW)
256 Kbps	0.135	0.179	0.159
128 Kbps	0.094	0.117	0.102
64 Kbps	0.074	0.082	0.078
32 Kbps	0.063	0.071	0.066
16 Kbps	0.058	0.061	0.059
8 Kbps	0.056	0.058	0.056

3.7 Wireless Communication

A wireless communication medium permits nodes to communicate in the event of tears in the fabric. In our current prototype, we are considering the use of Bluetooth for short-range wireless communication.

The Bluetooth Application Development Kit by Teleca ComTec AB [11] has been used for the initial test and development purposes. This employs the Ericsson ROK101008, a Class 2 (0 dBm) module, compliant with Bluetooth version 1.1. Two such devices were used for the transmitter and receiver sections. The maximum data rate over the UART is 460 kbps; our test results are for 57.6kbps.

Table 4: Power consumed in Bluetooth modes.

Operation Mode	Power Consumption(mW)
Standby	100
Inquiry:Scan	204
Inquiry:Send	261
Paging:Send	261
Paging:Scan	100

This module was tested for the basic three modes of operation: standby, inquiry and paging mode. These values represent only the communication cost without accounting for the necessary nodal processing power. Bluetooth can overcome fabric tears but at a higher power cost as can be seen from the results in Table 4. The same fault tolerance is achievable with the wire-based solution by redundancy.

4 Simulation Results

To compare the three architectures on the metrics of energy consumption and accuracy, a simulation program was constructed in Ptolemy II [9] that calculated power consumption and execution time, as well as ran the beamforming code. In this simulation, two versions of the application were constructed: (1) a beamforming textile with three clusters spaced at forty meter intervals, and (2) a beamforming textile with nine clusters spaced at ten meter intervals. For each version, we modeled an acoustic source whose position changed during the course of the simulation, i.e. a moving vehicle.

The accuracy of the beamforming application was computed as the average difference in the computed position and the actual position, as a percentage of the actual distance. The number of clusters in both beamforming textiles led to an overdetermined linear system which was solved using linear least squares; the intention was to improve the quality of the solution by using more clusters with a wider maximum spacing (three versus nine).

The energy consumption was compared across the three basic architectures: wired, wireless, and hybrid. For the beamforming textile, the energy consumption in the wired architecture was modeled as four states: (1) collecting data samples, (2) beamforming computation, (3) send/receive beamforming results to/from the collection point, (4) idle (not sleeping). In the hybrid architecture, the energy consumption was broken into two parts, the processor and the BlueTooth

module. The processor was in states (1), (2), or (4) above, while the BlueTooth module was either sending/receiving beamforming results as in (3) above, or in an idle state. The wireless model was similar to the hybrid architecture with the added complication that each of the microphones had a processor/BlueTooth module paired with it. To accomplish this, each cluster had a master node (which was connected to one of the microphones) that received data from each of the other BlueTooth modules in the cluster; this required the additional power consumed by these modules as well as the additional communication to be modeled. Collecting beamforming results in all of the architectures was modeled as point-to-point communication. Each cluster communicates with one of its neighbors in a binary tree pattern to collect results.

To assess the systems, two basic parameters were varied, the number of microphones (5 or 7) and the sampling rate (1024 or 2048 per second). The sample length and wired communication were kept fixed at 512 and 256,000, respectively, as these were found to be consistently the best in simulation results. Comprehensive results for accuracy and energy consumption are given in Table 5. Note that these points reflect the superior position of the wired architecture in terms of energy consumption. Alternatively, one could suppose a given energy budget and compute the most accurate answer possible.

An important aspect of energy consumption for each of these applications has not been included in this analysis. These devices will not have an acoustic source during the majority of their operational lifetime, making their sleep power consumption very important. All of the processing boards can be awakened only upon receiving an acoustic signal of appropriate strength, but only the wired communication architecture allows for the board to be awakened upon receipt of a communication signal. It is, therefore, simple in the wired architecture to operate in the 9mW sleep mode awaiting either a communication or acoustic signal; in the wireless case, this is not possible, forcing the processor to either keep the BlueTooth module awake (~100mW) or only periodically awaken itself to check for incoming communications at the cost of longer latencies in reporting results. Note that it is not sufficient to simply have every node wait to be awoken by an acoustic signal as differences in microphone gain and positioning (e.g., behind a rock) may lead to some nodes not being awake when needed.

5 Summary and Future Work

Both simulation and prototype measurements support the conclusion that architectures based on e-textiles result in power savings over wireless architectures, not a surprising conclusion given the relative costs of wired versus wireless communication. Several other interesting conclusions, however, can also be drawn from the results.

Table 5: Simulated energy and average location error for 9-cluster and 3-cluster beamformer vs. architecture, sample rate, number of microphones

Arch.	Mics	Sample rate	Energy(J) 9/3 cluster	Error % 9/3 cluster
Wired	5	1024	0.45/0.15	29.1/27.2
Hybrid	5	1024	1.34/0.45	29.1/27.2
Wireless	5	1024	6.41/2.14	29.1/27.2
Wired	5	2048	0.35/0.12	3.6/3.8
Hybrid	5	2048	1.25/0.42	3.6/3.8
Wireless	5	2048	5.94/1.99	3.6/3.8
Wired	7	1024	0.51/0.17	11.9/10.6
Hybrid	7	1024	1.40/0.47	11.9/10.6
Wireless	7	1024	9.00/3.00	11.9/10.6
Wired	7	2048	0.41/0.14	2.6/2.8
Hybrid	7	2048	1.31/0.44	2.6/2.8
Wireless	7	2048	8.34/2.78	2.6/2.8

The flexibility to select the best of sensors for a given task allows for both solution quality and power consumption to be improved; when the set of sensors is non-optimal, larger data sets must be collected and processed, trading off power consumption for solution quality. The flexibility to select a set of sensors comes at the cost of supporting a communication network capable of moving the right data to the right processors. There is, therefore, an interaction between solution quality, power consumption, and the communication network. In addition, the type of communication network and the number of sensors and processors needed to support a given solution quality over a range of conditions has implications for the cost of the system. The simulation results support the conclusion that the wireless system requires more processing/communication modules to achieve the same solution quality as the e-textile architecture.

While they are the focus of this paper, solution quality and power consumption are not the only criteria to be used when making architectural decisions. Other factors include cost, fault tolerance, electromagnetic emissions, and the ease with which modifications can be made. While the last two factors are outside the scope of the current simulation environment, fault tolerance is an excellent and important candidate for investigation.

Causes of faults can include manufacturing defects, tears in the fabric during use, failed components, and jammed radios; not all of the flaws affect each of the architectures in the same way. Determining reasonable fault rates, however, is a challenge during a time when e-textile assembly methods are a research topic and there are results from widespread textile usage.

References

- [1] Agilent Technologies, Agilent 3458a Multimeter Data Sheet, <http://cp.literature.agilent.com/litweb/pdf/5965-4971E.pdf>
- [2] Bluetooth Specification, <http://www.bluetooth.com/dev/specifications.asp>
- [3] Bekintex, <http://www.bekintex.com>.
- [4] Jung, S. Lauterbach, C., and Weber, W. "Integrated Microelectronics for Smart Textiles," Workshop on Modeling, Analysis, and Middleware Support for Electronic Textiles, October 2002.
- [5] Jones, M., Martin, T., and Nakad, Z. "A Service Backplane for E-Textile Applications," Workshop on Modeling, Analysis, and Middleware Support for Electronic Textiles, October 2002.
- [6] Marculescu, D., Marculescu, R., and Khosla, P. "Challenges and Opportunities in Electronic Textiles Modeling and Optimization," Proceedings of the Design Automation Conference, June 2001, pp. 175-180.
- [7] Measurement Specialties. Piezo Film Sensors Technical Manual. http://www.msiosa.com/piezo_documentation.htm.
- [8] Post, E.R., Orth, M., Russo, P., and Gershenfeld, N. "E-broidery design and fabrication of textile-based computing," IBM Systems Journal, Vol. 39, Nos. 3&4, 2000.
- [9] Ptolemy Project, Heterogeneous Modeling and Design. UC Berkeley, <http://ptolemy.eecs.berkeley.edu>.
- [10] Riley, R. Thakkar, S., Czarnaski, J., and Schott, B. "Power-Aware Acoustic Beamforming," Proceedings of the Fifth International Military Sensing Symposium, December 2002.
- [11] Teleca Comtec AB, *Bluetooth Development Kit information sheet*, http://www.comtec.teleca.se/info_development.asp