

Modeling and Simulating Electronic Textile Applications

Thomas Martin, Mark Jones, Joshua Edmison, Tanwir Sheikh, Zahi Nakad

Bradley Dept. of Electrical and Computer Engineering

Virginia Tech

Blacksburg VA 24061

{tlmartin, mtj, jedmison, tanwir, znakad}@vt.edu

ABSTRACT

This paper describes our design of a simulation environment for electronic textiles (e-textiles) and our experiences with that environment. This simulation environment, based upon Ptolemy II, enables us to model a diverse range of areas related to the design of electronic textiles, including the physical environment they will be used in, the behavior of the sensors incorporated into the fabric, the on-fabric network, the power consumption of the system, and the execution of the application and system software. This paper focuses on two aspects of the system, modeling the motion of a person wearing the e-textile and modeling the effect of faults in the e-textile system. To partially validate this environment, we compare simulation results against results from two different physical prototypes, a large-scale acoustic beamformer and a pair of shape-sensing pants.

Categories and Subject Descriptors

C.4 [Performance of Systems]: Design studies

General Terms

Design, performance, experimentation, measurement.

Keywords

Electronic textiles, wearable computing, smart fabrics, context awareness.

1. INTRODUCTION

Electronic textiles (e-textiles) are fabrics with interconnections and electronics woven into them. The electronics consist of both processing and sensing elements, distributed throughout the fabric. Potential applications for e-textiles include wearable computing (e.g., context-awareness [1][6], medical monitoring [7][13][18], and military uniforms) and sensor networks. E-textiles are embedded systems with a unique form factor, and are an extreme form of distributed

computing: physically spread over a relatively smaller space, but having a greater dependence on physical locality of computation, lower bandwidth for communication, less available energy, and requiring knowledge of the dynamic shape of the fabric.

The design space for e-textiles is large, with many choices for yarns, weaves, components, system software, interconnection networks, and sensor placement. While the area is still in its infancy, it is clear that prototyping alone cannot adequately explore the design space for most applications. Consequently, a simulation environment is necessary, one that allows a designer to adequately model a range of areas, including the physical environment in which the e-textile will be operating, the behavior of sensing and processing elements, and software execution. At the same time, simulation results are not useful unless the simulation environment has been validated against experimentally measured data. Consequently, we have built prototypes and compared their actual behavior to the behavior predicted by simulation.

In this paper, we describe the current status of our simulation environment for e-textiles and present results generated by both the environment and associated prototypes for two applications, a large-scale acoustic beamforming fabric for locating vehicles and a pair of pants for classifying and analyzing wearer motions. The remainder of the paper is organized as follows. Section 2 presents background material on e-textiles. Section 3 describes the elements of the simulation environment. Results from the environment for both applications are presented in Section 4. Finally, Section 5 summarizes the paper, presents conclusions, and discusses future work.

2. BACKGROUND

In order to help the reader understand the elements in the e-textiles simulation environment, this section provides background material on current applications of e-textiles and the design issues arising from those applications. The first part of this section reviews extant research in e-textiles and their applications. The remainder of the section examines the issues involved in the design of e-textiles.

2.1 E-textile Applications

The Wearable Motherboard project and related work at Georgia Tech has led to the creation of a system for monitoring a user's health, including heartbeat and respiration as well as the location of a bullet wound [11, 19]. Applications include monitoring infants for Sudden Infant Death Syndrome as well as monitoring the status of soldiers on the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

LCTES'04, June 11–13, 2004, Washington, DC, USA.

Copyright 2004 ACM 1-58113-806-7/04/0006 ...\$5.00.

battlefield. In these projects, wires were woven into the fabric for communication of data along with optical fibers to detect the location of bullet holes. Discrete sensors were attached and computing analysis was performed outside of the garment. Further work at Georgia Tech has investigated the use of FPGAs as self-configuring, fault-tolerant switches.

In industry, a range of products is either on the market or under development. ElekSen has developed a fabric keyboard that serves a dual-purpose in that it can be folded into a carrying case for a PDA device. This company has also developed fabrics for health-care applications and user interaction with vehicle interiors as well as health-care applications [5]. Durability tests show these fabrics are as robust as normal textiles and that the sensing capability does not degrade even over millions of user cycles. Infineon has produced a wearable MP3 player [9]. The primary contribution of Infineon's work is a method for packaging and attaching the digital and analog components in a washable, durable form factor with pins at a suitable pitch for fabric.

At Virginia Tech, the authors have constructed a number of e-textile prototypes, including large-scale (up to thirty feet long) acoustic-beamforming textiles that can determine the location of vehicles [8] and garments for context awareness. The acoustic-beamforming textile, a small prototype of which is shown in Figure 1, detects the bearing and location of a moving acoustic source using clusters of microphones and a digital signal processing board [8]. The context-awareness pants, shown in Figure 2, use an array of button-like eTAGs which function as sensors and communication devices to monitor and classify the wearer's motion [10]. In each case, these prototypes route all power and communication through the fabric and are capable of running for significant periods of time on a standard nine-volt battery [8]. We have also investigated the inclusion of novel materials, including thin piezoelectric films, into e-textiles [4]. Our long-term vision for e-textiles includes weaving new fibers into the fabric to provide capabilities such as chemical sensing, power storage, and dynamic color changing [25].



Figure 1: A single cluster prototype of the acoustic beamforming array. The cluster consists of seven microphones, a digital signal processing board, and communication/power lines woven into the fabric.



Figure 2: A prototype shape-sensing e-textile garment with button-like e-TAGs attached for sensing, acquiring, and analyzing data.

2.2 E-textile Design Issues

In a previous paper, we described the design issues associated with e-textiles [12]. This section summarizes these design issues here to make clear the relationship between the modules of the simulation environment described in Section 3. A successful simulation environment for e-textiles must encompass the full range of these areas, which includes

- physical environment
- sensor behavior,
- human body and motion,
- motion/draping of clothing,
- manufacturability (weaving & piecework),
- networking,
- power consumption, and
- software execution.

Applications of e-textiles will have to sense the physical environment, both for wearable computing applications and non-wearable computing applications. Non-wearable applications will typically manifest as sensor networks, which by definition are monitoring the physical environment. A particularly direct and appealing wearable computing application for e-textiles is context-awareness. Context-awareness can be summarized simply as knowing what the user is doing, where the user is, and what is going on around the user, so that the system can adjust its behavior to the user's current situation [1, 26]. Consequently, any simulation tool for e-textiles must be able to model the physics of the environment, and the behavior of the sensors used to monitor that environment. For wearable computing applications, an e-textiles simulation tool must also be able to capture aspects of the human body, e.g., body size and limb motion, in order to model the dynamic location of the sensors. Because most garments do not fit tightly on the human frame, the draping and motion of the cloth in response to the underlying movement of the wearer may also have to be modeled to fully account for the movement of the sensors.

With respect to manufacturability, one of our goals for e-textiles is to remain within the confines of existing manufacturing techniques in the textile and garment industries. This will enable e-textiles to take full advantage of the efficient and cost effective mass production techniques employed by those industries. Because garments are constructed of pieces cut from a large rectangle of fabric and then sewn together, the fabric must be designed such that the pieces can be easily interconnected electrically, allowing sensing and processing elements on different pieces of the garment to communicate.

An e-textile simulation environment must also account for the computing aspects of the application: networking, power consumption, and software execution. Sensing and processing elements must be interconnected through an on-fabric network [16]. The network topology is constrained by the properties of the weaving process. In particular, like the threads in the fabric, electronic interconnections must be run in two perpendicular directions. The network must also be fault tolerant, because the fabric will be subject to manufacturing defects as well as wear and tear during use.

Like other mobile and wearable computing devices, e-textiles will typically be battery-powered, and thus power consumption must be modeled to allow for energy-efficient design. E-textile systems differ from other low power systems in that the power sources will be modular and distributed in order to maintain flexibility. Whereas other low power systems must optimize energy use from a single power source, fabric substrates will likely have to optimize energy use from multiple power sources. Consequently, the optimization problem is more difficult for fabrics. In order to achieve true system level fault tolerance in an e-textile, power fault tolerance in addition to network fault tolerance, must be implemented. Paths from power sources to sensor and computation nodes must be dynamic, allowing power to be routed around damaged sections of the fabric. Because of the physical locality of computation, e.g. a node will search for other nodes within a given physical region for beamforming, power consumption may be non-uniform across the fabric. As a result, the number and distribution of power sources on the fabric will have a great impact on lifetime and application performance [23]. This issue will be elaborated upon later in the paper.

3. SIMULATION ENVIRONMENT

To construct an appropriate simulation environment, we have selected Ptolemy II as a means of integrating a wide range of simulation capabilities. Ptolemy II provides the type of diverse environment required to handle very different simulation domains and an open architecture that can accommodate interfacing to other environments [21]. For example, Ptolemy II has provisions for simulations in the continuous domain, which is useful for simulating the physics of the environment, while simultaneously providing a discrete event domain that is appropriate for computation.

Figure 3 presents the general framework of our simulation environment. Modules within the dashed line are implemented within the Ptolemy II environment, while those on the outside represent extant systems. For example, our models of the physical environment are implemented in the Ptolemy continuous simulation domain while our data from human motion is read from processed data files created by a video motion capture system [2]. Not all of the modules in the figure are present in every simulation.

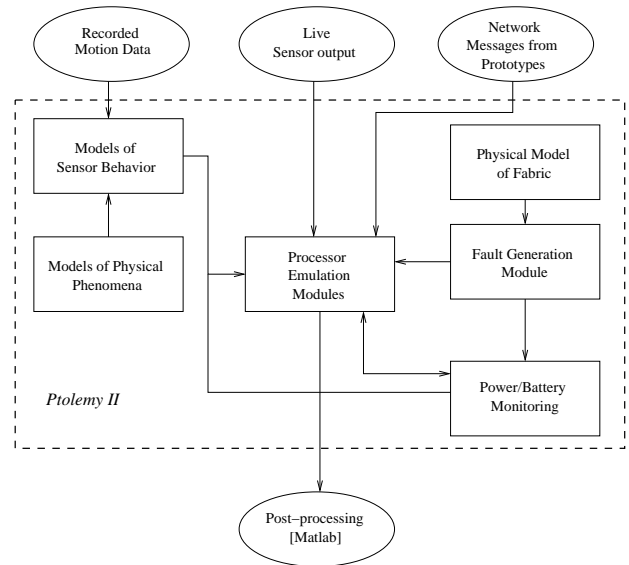


Figure 3: A block diagram of the structure of the simulation environment. Modules within the dashed line are implemented within the Ptolemy II simulation environment.

Some aspects of the system are application dependent, some are fabric configuration dependent, and others are constant across applications. The models for the sensors and the associated models of the physical environment are application dependent. For example, the microphone models and the propagation of vehicle acoustics are specific to the beamforming application and are described in the next section. The simulation environment has a representation of the number and location of sensors, processors, communication wires, and power lines. The specific configuration is clearly dependent on the characteristics of the e-textile being modeled, but this configuration is not dependent on the application. To allow for exploration of a range of architectures and sensor configurations, we have created an automatic generation tool that allows for the Ptolemy model of the textile to be generated from a configuration file.

To describe the fabric configuration, the physical model of the fabric considers the placement of the components on the fabric and the interconnections between pieces of fabric. It can be made to automatically generate locations of power and ground lines through each piece based upon an input wiring density per unit length. The model then automatically connects each component on the fabric to the nearest power/ground line pair. The power/battery model assumes that the battery has a fixed amount of energy and keeps track of how much energy has been drawn from it and how much remains. To model the energy consumption, the processor emulation modules must provide a table of power management states and the amount of power consumed in each state to the power/battery model, which we fashioned after the ACPI power management interface [3]. The power/battery model can then track how long each component is in a given state, and uses the timing information and the power for that state to calculate how much energy had been consumed. A more elaborate battery model that accounts for changes in battery capacity with load power is

possible [14], but this first-order model is sufficient for our purposes.

Finally, there are aspects of the system that are constant across application and textile configuration. Most important of these is the processor emulation module, with one module representing each processor in the e-textile. Each module accepts input from sensors as well as raw network data from other processor modules. The processor module implements an interrupt-driven model of computation, with several types of interrupts managed, including sensor data arrival, network data arrival, and clock data. The processor module allows for the user to insert their own C code to handle these interrupts as well as initialize the system. This C code has been ported intact to actual processors, allowing for emulation of e-textile prototypes in a simulation mode. The processor model is capable of outputting results to a network port (both to a debug port and to other processor modules) and providing results on power consumption and processing time. It is important to note that this is not a cycle-by-cycle processor simulator; the data on power consumption and processing time is based on user estimates for a given processor and interrupt handler to be modeled.

Because e-textiles are at an early stage of development and a range of applications are being investigated, it is not possible to confine all exploration of the design space to a pure simulation mode. Further, there is a need for extensive debugging in this novel hardware/software environment, but there is a dearth of such tools. To partially address these needs, we have augmented the environment with the capability to accept live sensor data as well as a hybrid mode of operation. In the live sensor data mode, we can replace select sensor models with sensor data provided over a TCP/IP link. In cases where the data processing cannot keep up with the sensor input rate, prerecorded data is fed from a file. The hybrid mode of operation, inspired by [20], allows for the simultaneous integration of the simulation environment with physical e-textile prototypes. For example, we have constructed a physical acoustic beamforming e-textile with four processing elements [8], but would like to investigate the performance of this prototype in an architecture with tens of processors. The hybrid mode bridges the simulated and physical e-textile networks over a TCP/IP link, allowing for extensive investigation and emulation of network and application behavior [17]. Finally, this same bridging technique can be used to decompose simulations of large e-textiles across multiple processors to reduce simulation time.

4. APPLICATION CASE STUDIES

In this section we present results from two different applications, a large-scale acoustic beamforming fabric for locating vehicles and a pair of pants for context awareness. These results illustrate the utility of the simulation environment in the exploration of the e-textiles design space. For the first set of results, the simulator was used to explore the trade-offs between fault tolerance and accuracy involved in the configuration of the batteries, utilizing the physical model of the fabric, the fault generation module, and power/battery model blocks of Figure 3. For the second set of results, the simulator was used to study the choice of sensors and sensor placement for a garment that monitors the wearer’s motions, using recorded motion data and models of two types of sensors.

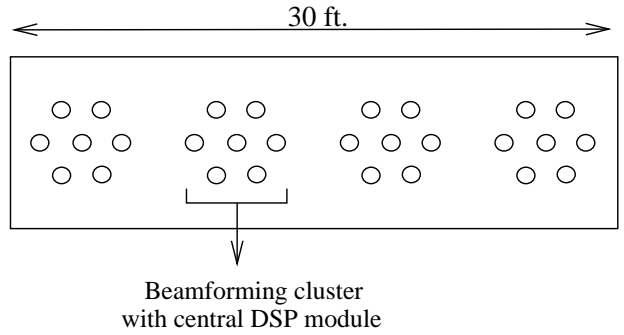


Figure 4: A diagram of the four-cluster acoustic beamforming e-textile. Each cluster has seven microphones, all of which are connected to a single digital signal processing board. Each cluster computes the line-of-bearing to the acoustic signal. Clusters then communicate this information amongst themselves to find the location via triangulation.

4.1 Large-scale acoustic beamformer

The purpose of the first application, the large-scale acoustic beamformer, is to locate passing vehicles based on their acoustic emissions. The beamformer, depicted in Figure 4, consists of four clusters. Each cluster has seven microphones and a DSP processing module. Using the difference in the time of arrival of a sound to each microphone, a cluster can calculate a bearing to a sound source. The clusters then communicate their angles to each other and triangulate to find the location of the sound source. A picture of one cluster of the prototype is shown in Figure 1.

We have previously described the use of the simulation environment to analyze the impact of communication schemes and sampling rates on the power consumption and accuracy of the beamformer. This analysis demonstrated that a trade-off exists between accuracy and power consumption based upon the sampling rate and the set of microphones chosen within each cluster [8]. In particular, if an e-textile has the ability to select an optimal set of sensors, then less data must be collected and processed, reducing the energy required for a given level of accuracy. However, being able to select an optimal set of sensors comes at the cost of an on-fabric network that can transmit data from the best set of sensors to the processing elements.

Another interesting trade-off arises when fault tolerance is considered, as is discussed in the following paragraphs. Because of manufacturing defects and normal wear and tear on the fabric, it is likely that there will be broken and shorted wires in the fabric over the course of its lifetime. If the system is powered by a single battery, then a short circuit in one portion of the fabric may cause the entire fabric to lose power. One solution to this problem is to have multiple batteries, with each battery powering an isolated portion of the fabric. However, this is less than optimal if no faults occur, because if a portion of the fabric is active more often, its battery may be exhausted while other, less active portions of the fabric still have energy remaining. The system may still be able to operate using only these portions, but its accuracy may be reduced.

We modeled this problem using our simulation environment by creating a physical fault model that introduces




Figure 5: The state of the acoustic beamformer after the introduction of faults on the left two clusters. The components disabled by the faults are denoted with an X; the remaining functioning components are denoted with a diamond.

short- and open- circuits into the fabric. Because e-textiles are not yet widely available, there is no information on the types and frequency of the faults that occur. As a first step, we assumed that faults could either be short- or open-circuits, and we generated a random set of faults of various types, shapes, and densities that we think are appropriate given the intended location for the deployment of the beam-former (a roadside in a hostile military environment). For example, one type of shape was a rectangular strip at some angle across the fabric such as might occur if a vehicle were to drive over the fabric. Another type was fashioned after a shotgun blast, a collection of random point faults over a constrained area.

In addition to faults, it was also necessary to model how many batteries are present and how they are configured. We investigated three different configurations of the batteries for the beamformer. Configuration 1 had one battery for the entire beamformer, configuration 2 had one battery for each pair of clusters (two batteries total), and configuration 3 had one battery for each cluster (four batteries total). The overall energy capacity was equal for each of the three configurations.

Figure 5 shows the acoustic beamformer in configuration 3. Each cluster has its own battery, as indicated by the small square near the right of each cluster. Two faults have been introduced, a vehicle track fault on the leftmost cluster and a shotgun fault on the cluster next to it. The faults are all short circuits, and the placement model has determined that the components on those two clusters have been disabled by the short circuits, as indicated by the X's. Components on the right two clusters are still functional, as indicated by the diamonds.

After a fault is introduced into a cluster, the battery begins to rapidly discharge, as shown in Figure 6. When the vehicle track fault occurs on the leftmost cluster, battery 1 begins to discharge rapidly, and then when the shotgun fault occurs on the cluster next to it, battery 2 also begins to discharge rapidly. The short-circuit faults are modeled with a finite resistance, so that the batteries do not discharge instantaneously. On the other hand, the two clusters on the right do not have any faults, and so their batteries (batteries 3 and 4) discharge at a normal rate.




Figure 6: Battery charge versus time for clusters with faults in the fabric (batteries 1 and 2) and batteries on clusters without faults (batteries 3 and 4).

The model can also introduce open-circuit faults, in which case the battery discharges more slowly than normal, as illustrated by Figure 7, which shows the battery charge versus time for the beamformer in configuration 1 (one battery for the entire beamformer) both under normal operation with no faults and with an open-circuit fault present in one cluster. The open-circuit fault causes some of the electronics to be disconnected, reducing the power drain on the battery, and thus reducing the rate at which the battery discharges.

In addition to the battery life, the accuracy of the beamformer is also a major concern. Only two clusters are required to triangulate the location of a sound source, and generally the farther apart those two clusters are, the more accurate the location estimate will be [24]. Consequently, using the leftmost and rightmost clusters gives the most accurate estimate, while using two adjacent clusters gives the least accurate estimate. The accuracy also varies with the location of the sound source, with higher accuracy for sources near a line drawn perpendicular to the midpoint of a line connecting the two clusters.




Figure 7: Battery charge versus time for an open-circuit fault introduced into one cluster and for no faults, with a single battery powering all four clusters (configuration 1).

To test the accuracy of the beamformer as faults are introduced, we modeled a sound source moving along an arc 60 feet from the cluster at a velocity of 7.85 ft/s, so that the arc is traversed in 12 seconds. For each of the three battery configurations, one fault was introduced each second.

Figure 8 shows the accuracy of the three configurations after the introduction of a series of faults. Configuration 1, as expected, fails after the first fault, as its single battery is discharged. Configuration 2 is able to withstand a few additional faults before failing, while configuration 3, with its one battery per cluster is able to withstand nearly twice as many faults before it fails entirely. The trade-off is that, in the event that there are no faults, even though all three systems have the same battery life, for configuration 3 the outermost two clusters run out of battery charge halfway through the battery life, leaving the inner two clusters to triangulate. The accuracy is then reduced by approximately 5%. So in the no-fault case, configuration 3 has a lower average accuracy over the entire life of the system. Thus, fault-tolerance must be balanced against performance.

4.2 Shape-sensing pants

Our second example of using the simulation environment involves the design of a pair of shape-sensing pants. The intended function of the pants is context-awareness, in particular, the classification of the current activity of the user, such as running, walking, standing. A number of other groups have investigated using sets of sensors to determine user activity [1, 6, 26, 22]. All of these groups built hardware prototypes of their sensor systems and then tested their functionality on a few individuals, usually other people working in their laboratory. In contrast, before building the hardware prototype, we simulated the garment using recorded motion data [2] and modeling the physical behavior of two types of sensors.

Simulating the system before building a prototype has several advantages over prototyping alone. First, by using motion data collected from a wider set of individuals, there is a greater confidence that our system will function for a broad segment of the population. The motion database has




Figure 8: This graph gives the accuracy of the location of the beamforming cluster as the cluster experiences an increasing number of faults. Configuration 1 fails after a single fault, configuration 2 experiences a rapid degradation of functionality after a few faults, while configuration 3 is able to run for up to a dozen faults.

individuals covering a range of body sizes and shapes. If simulation were not used, then each iteration of the hardware prototype must be tested on individuals covering the same range of body sizes. A second advantage is that we are able to study issues of the type of sensors used and where those sensors are placed on the body without having to first build a hardware prototype. Thus we can more rapidly perform feasibility studies and decrease the number of iterations that a prototype must go through before it has the desired functionality. Finally, simulation using recorded human motions makes our experiments more repeatable. When testing a hardware prototype, the subject will not make exactly the same motion each time the experiment is run, introducing variability in the results.

We used this environment to design a first iteration garment that incorporates two types of sensors, accelerometers and piezoelectric film strips. We created physical models for each of these sensor types, and for input to those models, we used data from a publicly available human motion database [2]. The motion database consists of three-dimensional position data for various locations on a subject's body for each frame of video. The data was available for the subjects making several different motions, such as walking, running, and jumping. Due to high frequency noise in the data, the data was smoothed using the Matlab spline toolbox. From this smoothed position data, the acceleration of each monitored location on the body was calculated from a sliding window of three position samples and the sampling frequency of the video system (120 Hz), i.e., the difference between each pair of consecutive positions yielded the velocity, and the difference between each pair of consecutive velocities yielded the acceleration. The resulting accelerations were then used to drive models of an accelerometer and a piezoelectric film strip.

The accelerometer we modeled was the dual-axis ADXL311 from Analog Devices. The accelerations calculated from the motion database were converted from mm/s^2 to g 's, a multiple of Earth's gravitational constant $9.81m/s^2$. The specifications for the ADXL311, specifically the number of millivolts per g , were used to calculate the final ac-




Figure 9: The angles used to compute the correction to simulated acceleration. The angle between the shin and the ground is Θ and the angle between the shin and thigh is α .

celerometer output voltage. This, however, did not completely model the behavior of the sensor. Because the sensor and the body are closely coupled, the accelerometer may undergo a change in orientation, particularly if the sensor is used on a part of the body that exhibits a high degree of freedom such as the wrist or ankle. To account for the change in orientation, a rotation matrix derived from Figure 9 was used as shown in the following equations

$$X_{corrected} = X_{uncorrected} \cos \Theta + Y_{uncorrected} \sin \Theta, \quad (1)$$

and,

$$Y_{corrected} = Y_{uncorrected} \cos \Theta + X_{uncorrected} \sin \Theta. \quad (2)$$

The accelerometer model included inputs for position data, such as the knee and heel for the ankle, which are used to calculate the angle necessary for computation of the rotation matrix. The output of the accelerometer model for a step using $X_{corrected}$ and $Y_{corrected}$ corresponds very closely to actual sensor output, as shown in Figure 10. The top graph in Figure 10 is from simulation of the accelerometer model using the recorded motion data, while the bottom graph is from an actual accelerometer being worn while walking. While the curves do not match exactly because they are not from the same individual, they have the same general features and ranges of acceleration.

The other sensor that was modeled for this application was a piezoelectric film strip. One property of piezoelectric film strips is that they generate a voltage in response to a physical stimulus, such as being bent or having a force applied. Thus they can be used to sense the angle of a joint or to sense the forces resulting from a motion, such as the force acting on the heel when it contacts the ground. However, for the simulation model for the piezoelectric film we only modeled sensing the angle of a joint due to the lack of force data in the motion database. The mathematical model of a piezoelectric film is an ideal voltage source in




Figure 10: The top figure gives the output of the simulated acceleration sensor while the bottom figure gives the actual data from a physical accelerometer. The simulation was generated from video data of a subject walking while the acceleration data was taken from a different subject wearing an accelerometer while walking.

series with a capacitor (typically tens of nano-farads) across an input impedance, forming a high-pass RC circuit with a cutoff frequency equal to $\frac{1}{2\pi RC}$ [15]. Even for large values of R , the resulting cutoff frequency is typically hundreds of hertz. Above this cutoff frequency, the film produces voltage proportional to the physical stimulus; below this cutoff frequency, the film produces a voltage proportional to the rate of change of the input. Because the frequency of most human motions (typically 1 Hz or less) is well below this frequency, particularly joint angles, we can assume that the piezoelectric film will measure the angular rate of change if placed on the joints.

The simulation model consists of inputs for three points that form a joint angle, such as the hip, knee, and heel, which form the knee joint. The angle formed by these three points as shown in Figure 9 is calculated and the discrete derivative of the resulting angles is the final piezoelectric sensor output. The model was verified using the analytical results from a pendulum in [12]. For this particular appli-




Figure 11: The top figure gives the output of the simulated piezoelectric sensor while the bottom figure gives the actual data from a physical piezoelectric sensor. The simulation was generated from video data of a subject walking while the acceleration data was taken from the same subject walking.

cation, the sensor was placed at the knee. Figure 11 gives a comparison of the output of the simulation model and the output of an actual piezoelectric film strip. The top graph in Figure 11 shows the simulation output, while the bottom graph shows the output of the actual sensor. These results are for the same set of steps from a single individual. As would be expected, the actual results have more noise than the simulation results, but other than that, the simulation results match the actual results very well.

Given the models for the accelerometer and the piezoelectric film strip, the simulation environment was used to train a neural network for classifying user motion into walking, running, jumping, and standing, using a method similar to that described in [26]. Ideally, one would like a context-aware garment that works for the entire population without requiring any training by the individual, or at least, a set of garments of that work for various ranges of body sizes. An open question is whether this goal is feasible, and if so, what

is the subset of the population should be used to train the garment to achieve this goal. We have previously reported on using the environment to explore the the dynamic range required of the sensors when placed on different locations on the body and the accuracy of the application when trained on different subsets of the population [12]. The results show that the accuracy is highly dependent on which range of individuals is used to train the neural network. The best results were obtained when using individuals who were near the mid-range of the population in terms of maximum heel acceleration while walking, while the worst results were obtained for individuals who were outliers at either end of the population. Somewhat surprisingly, using the average of two outliers, one from each end of the population, worked nearly as well as using an individual who was in the middle of the population. The important point is that these findings would have been difficult to obtain without simulation. Using a hardware prototype to study the accuracy of the application versus the segment of the population used for training would have been prohibitively time-consuming. Simulation permits a greater exploration of the design space.

5. CONCLUSIONS

Design considerations for e-textile applications encompass a range of areas, including the physical environment, sensor behavior, human body size and motion, the shape of the cloth, garment and fabric manufacturability, and standard computer engineering areas such as networking, power consumption, and software execution. A simulation environment for e-textiles must be able to readily model aspects of each of these areas. This paper has described a simulation environment designed to aid in the exploration of the e-textile design space. In addition, features allowing for the inclusion of live sensor data and the integration of physical and simulated elements of an e-textile make it a useful tool for debugging e-textile prototypes.

Aspects of this environment were demonstrated in two case study applications. In the case of fault and power analysis of the acoustic beamformer, we were able to explore the effects of fault introductions that would be prohibitively expensive to study using actual e-textile prototypes. In particular, we showed that the choice of battery configuration will affect both the fault tolerance and the accuracy of the e-textile, and that a trade-off exists between them. For the context-aware pants, the results of the simulation environment were verified against ground truth data, showing a close correspondence between the simulation and actual data.

For our future work, we would like to extend the model to include the manufacturability of the e-textile. One of our goals is to stay as close as practically possible to existing manufacturing techniques for textiles and garments. By doing so, e-textiles will be able to take advantage of the cost-effectiveness of textile and garment manufacturing techniques. We envision that e-textiles are an enabling technology for pervasive computing, and as such, must be inexpensive if they are to be widely deployed. Adding manufacturability to the environment would provide feedback to the e-textile designer about the cost of points in the design space.

Another aspect that we believe must be added to the environment is the motion and draping of the cloth. At the present time, we assume that the sensors for wearable

e-textiles are at fixed position on the body, which would be true for tight-fitting garments (e.g., spandex). However, most garments move with respect to the body and change shape as the user moves about. Consequently, sensors will not be at fixed locations on the body. For some applications the change in position of the sensors relative to joints and other important features of the body must be accounted for. If they are not accounted for, then the simulation environment will not provide accurate information about the characteristics of the e-textile design space.

6. ACKNOWLEDGMENTS

This material is based upon work partially supported by the National Science Foundation under grant no. CCR-0219809 and by DARPA under contract F30602-00-2-0548. The authors gratefully acknowledge Dana Reynolds and Eloise Coupey for their efforts in the weaving of the fabrics and construction of the pants.

7. REFERENCES

- [1] B. Clarkson, K. Mase, and A. Pentland. Recognizing user context via wearable sensors. In *Proceedings of the Fourth International Symposium on Wearable Computing*, pages 69–75. ISWC 2000, October 2000.
- [2] CMU Graphics Laboratory. CMU Graphics Lab Motion Capture Database, 2003. Available from: <http://mocap.cs.cmu.edu/search.html>.
- [3] Compaq Computer Corp., Intel Corp., Microsoft Corp., Phoenix Technologies Ltd., and Toshiba Corp. Advanced configuration and power interface specification, revision 2.b, 2002. Available from: www.acpi.info/spec.htm.
- [4] J. Edmison, M. Jones, Z. Nakad, and T. Martin. Using piezoelectric materials for wearable electronic textiles. In *Proceedings of the Sixth International Symposium on Wearable Computing*, pages 41–48. ISWC 2002, 2002.
- [5] Eleksen, Ltd., 2004. Available from: <http://www.eleksen.com>.
- [6] J. Farrington, A. Moore, N. Tilbury, J. J. Church, and P. Biemond. Wearable sensor badge and sensor jacket for context awareness. In *Proceedings of the Third International Symposium on Wearable Computing*, pages 107–113. ISWC 1999, October 1999.
- [7] P. Grossman. The lifeshirt: A multifunctional ambulatory system that monitors health, disease, and medical intervention in the real world. In *International Workshop on New Generation of Wearable Systems for eHealth*, pages 73–80, December 2003.
- [8] M. Jones, T. Martin, Z. Nakad, R. Shenoy, T. Sheikh, D. Lehn, and J. Edmison. Analyzing the use of e-textiles to improve application performance. In *Proceedings of the IEEE Vehicular Technology Conference 2003, Symposium on Wireless Ad hoc, Sensor, and Wearable Networks*. VTC 2003, October 2003.
- [9] S. Jung, C. Lauterbach, and W. Weber. Integrated microelectronics for smart textiles. In *Workshop on Modeling, Analysis, and Middleware Support for Electronic Textiles*, pages 3–8. MAMSET 2002, October 2002.
- [10] D. Lehn, C. Neely, K. Schoonover, T. Martin, , and M. Jones. e-TAGs: e-Textile Attached Gadgets. In *Proceedings of Communication Networks and Distributed Systems: Modeling and Simulation*, January 2004.
- [11] E. Lind, R. Eisler, G. Burghart, S. Jayaraman, R. Rajamanickam, and T. McKee. A sensate liner for personnel monitoring applications. In *Digest of Papers of the First International Symposium on Wearable Computing*, pages 98–105. ISWC 1997, 1997.
- [12] T. Martin, M. Jones, J. Edmison, and R. Shenoy. Towards a design framework for wearable electronic textiles. In *Proceedings of the Seventh International Symposium on Wearable Computing*, pages 190–199. ISWC 2003, October 2003.
- [13] T. Martin, E. Jovanov, and D. Raskovic. Issues in wearable computing for medical monitoring applications: A case study of a wearable ecg monitoring device. In *International Symposium on Wearable Computers*, pages 43–50, October 2000.
- [14] T. Martin and D. Siewiorek. A power metric for mobile systems. In *Proceedings of the 1996 International Symposium on Low Power Electronics and Design*, pages 37–42, August 1996.
- [15] Measurement Specialties, Inc. Piezo technical manual, piezoelectric film properties, 2002. Available from: <http://www.msiousa.com/PART1-INT.pdf>.
- [16] Z. Nakad, M. Jones, and T. Martin. Communications in electronic textile systems. In *Proceedings of the 2003 International Conference on Communications in Computing*, pages 37–43. CIC 2003, June 2003.
- [17] Z. S. Nakad. *Architectures for e-Textiles*. PhD thesis, Bradley Department of Electrical and Computing Engineering, Virginia Tech, 2003.
- [18] R. Paradiso, G. Loriga, N. Taccini, M. Pacelli, and R. Orselli. Wearable system for vital signs monitoring. In *International Workshop on New Generation of Wearable Systems for eHealth*, pages 161–168, December 2003.
- [19] S. Park, C. Gopalsamy, R. Rajamanickam, and S. Jayaraman. The wearable motherboard: An information infrastructure or sensate liner for medical applications. *Studies in Health Technology and Informatics*, IOS Press, 62:252–258, 1999.
- [20] S. Park, A. Savvides, and M. B. Srivastava. Sensorsim: a simulation framework for sensor networks. In *Proceedings of the 3rd ACM international workshop on Modeling, analysis and simulation of wireless and mobile systems*, pages 104–111. ACM Press, 2000.
- [21] Ptolemy Project. Heterogeneous modeling and design, 2003. Available from: <http://ptolemy.eecs.berkeley.edu>.
- [22] C. Randell and H. Muller. Context awareness by analysing accelerometer data. In *Proceedings of the Fourth International Symposium on Wearable Computing*, pages 175–176. ISWC 2000, October 2000.
- [23] T. Sheikh. Modeling of power consumption and fault tolerance for electronic textiles. Master's thesis, Bradley Department of Electrical and Computing Engineering, Virginia Tech, 2003.
- [24] R. Shenoy. Design of e-textiles for acoustic applications. Master's thesis, Bradley Department of

- Electrical and Computing Engineering, Virginia Tech, 2003.
- [25] G. Spinks, B. Kim, G. Wallace, and R. John. Electroformation of conductive polymers in a hydrogel support matrix. *Polymer*, 41:1783–1790, 2000.
- [26] K. Van Laerhoven and O. Cakmakci. What shall we teach our pants? In *Proceedings of the Fourth International Symposium on Wearable Computing*, pages 77–83. ISWC 2000, October 2000.