

Research Challenges in Wireless Networks of Biomedical Sensors*

Loren Schwiebert
Department of Computer
Science
Wayne State University
Detroit, MI 48202
loren@cs.wayne.edu

Sandeep K. S. Gupta
Department of Computer
Science and Engineering
Arizona State University
Tempe, AZ 85287
sandeep.gupta@asu.edu

Jennifer Weinmann
Department of Electrical and
Computer Engineering
Wayne State University
Detroit, MI 48202
aa5685@wayne.edu

ABSTRACT

Implanted biomedical devices have the potential to revolutionize medicine. *Smart sensors*, which are created by combining sensing materials with integrated circuitry, are being considered for several biomedical applications such as a glucose level monitor or a retina prosthesis. These devices require the capability to communicate with an external computer system (base station) via a wireless interface. The limited power and computational capabilities of smart sensor based biological implants present research challenges in several aspects of wireless networking due to the need for having a bio-compatible, fault-tolerant, energy-efficient, and scalable design. Further, embedding these sensors in humans add additional requirements. For example, the wireless networking solutions should be ultra-safe and reliable, work trouble-free in different geographical locations (although implants are typically not expected to move; they shouldn't restrict the movements of their human host), and require minimal maintenance. This necessitates application-specific solutions which are vastly different from traditional solutions.

In this paper, we describe the potential of biomedical smart sensors. We then explain the challenges for wireless networking of human-embedded smart sensor arrays and our preliminary approach for wireless networking of a retina prosthesis. Our aim is to motivate vigorous research in this area by illustrating the need for more application-specific and novel approaches toward developing wireless networking solutions for human-implanted smart sensors.

1. INTRODUCTION

The recent development of high-performance microproces-

*This research was supported in part by National Science Foundation Grants #DGE-9870720, #ANI-0086020 and #ANI-0196156.

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 publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 2001 ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

sors and novel sensing materials has stimulated great interest in the development of smart sensors – physical, chemical, or biological sensors combined with integrated circuits. It is not uncommon to place multiple sensors on a single chip, with the integrated circuitry of the chip controlling all these sensors. Since the context should be clear, we also use the term sensor when referring to a smart sensor, *i.e.*, the combination of the sensor and the integrated circuit. These smart sensors can be relatively inexpensive to build, allowing for the large-scale deployment of networks of smart sensors. Technical advances are expected to improve the capabilities and performance of these devices.

The applications of smart sensors are extensive and describing all the uses of these devices is probably not possible. However, a few example application domains will help explain the interest in these devices. For example, military interest in networks of smart sensors is motivated by many problems that can be safely and effectively solved with smart sensors. For example, networks of smart sensors could be deployed in combat scenarios to track troop movements. Sensors placed on small robots could conduct land mine detection. Smart sensors could detect the use of biological or chemical weapons and, via network communication, report their presence in time to protect troops.

Civilian applications of smart sensors are also many and varied. One example is pollution detection along beaches, with smart sensors distributed along the shoreline and using wireless communication to relay information to a base station for further processing. Smart sensors could also be distributed throughout the exhaust system of an automobile to detect levels of emissions and efficiently reduce pollution. Networks of smart sensors have also been proposed for climate control in large office buildings, where consistent temperatures have been difficult to achieve.

This list of applications only scratches the surface of proposed and potential applications of novel smart sensors. One thing in common with the aforementioned applications, however, is that using a wireless interface to these devices is superior to a wired connection, even in cases where a wired connection may be possible. For example, for some military applications, the distribution of sensors could be done by ejecting them from aircraft. For automobile exhaust monitoring, wires would require extensive protection to operate in this high-temperature environment.

In this paper, we consider another application area of smart sensors – biomedical applications. In particular, we

are interested in embedded smart sensors that operate within the human body to compensate for various diseases. Because these devices are placed inside the human body and also because of the anticipated uses of these devices, the research problems to be solved differ from those for networking generic smart sensors. In the following, we first describe a number of proposed biomedical applications of smart sensors, including a biomedical application we are currently working on, the artificial retina. Then we describe the problems with wireless networking of biomedical applications, followed by some of our ideas on solving these problems.

2. BIOMEDICAL SENSOR APPLICATIONS

Although the technology for biomedical smart sensors is still relatively new, a number of interesting proposed applications have emerged. These applications are briefly described to give an idea of the potential impact that biomedical sensor networks will have on the health and well-being of future medical patients.

Although only a small cross-section of possible applications are described, it should be obvious that there are many potential uses for wireless biomedical sensors. As this research progresses, many more will become apparent.

2.1 Artificial Retina

The example biomedical application we discuss in detail in this paper is the artificial retina. This should help highlight the different requirements of biomedical sensors. In this section, we first describe the retina prosthesis developed within the Smart Sensors and Integrated Microsystems (SSIM) project at Wayne State University and the Kresge Eye Institute. One of the goals of this project is to build a chronically implanted artificial retina with sufficient visual functionality to allow persons without vision or with limited vision to “see” at an acceptable level. In order to reach this goal, communication protocols that are significantly more efficient and scalable than current designs must be developed.

Currently, smart sensor chips, each with 100 microsensors, have been built for an ex-vivo testing system of a retina. Figure 1 depicts an illustration of the design. The smart sensor has two components: an integrated circuit and an array of sensors. The integrated circuit is a multiplexing chip, operating at 40KHz, with on-chip switches and pads to support a 10×10 grid of connections. The circuit has both transmit and receive capabilities. Each connection has an aluminum probe surface where the micromachined sensor is bonded. This is accomplished by using a technique called backside bonding, which places an adhesive on the chip and allows the sensors to be bonded to the chip, with each sensor located on a probe surface. Before the bonding is done, the entire IC, except the probe areas, is coated with a biologically inert substance. Several prototype chips have been fabricated through MOSIS and the sensor arrays have already been built. The sensors form a 10×10 electrode array; each sensor is a microbump that will rest on the retina. The distance between adjacent microbumps is approximately 70 microns. These microbumps start with a rectangular shape and near the end taper to a point. This point is placed on the retina tissue, allowing contact between the smart sensor array and the retina without perforating the retina.

An implantable version of the current ex-vivo microsensors




Figure 1: Illustration of the Microsensor Array

array, along with its location within the eye, is shown in Figure 2. The microbumps rest on the surface of the retina rather than embedding themselves into the retina. Unlike some other systems that have been proposed, these smart sensors are placed upon the retina and are small enough and light enough to be held in place with relatively little force. These sensors produce electrical signals which are converted by the underlying tissue into a chemical response, mimicking the normal operating behavior of the retina from light stimulation. The chemical response is digital (binary), essentially producing chemical serial communication. A similar design is being used for a cortical implant, although the spacing between microbumps is larger to match the increased spacing between ganglia in the visual cortex.

As shown in Figure 2, the front side of the retina is in contact with the microsensor array. Transmission into the eye works as follows. The back side of the retina is stimulated electrically (via an artificial retina prosthesis) by the sensors on the smart sensor chip. These electrical signals are converted into chemical signals by the ganglia and other underlying tissue structures and the response is carried via the optic nerve to the brain. Signal transmission from the smart sensors implanted in the eye works in a similar manner, only in the reverse direction. The resulting neurological signals from the ganglia are picked up by the microsensors and the signal and relative intensity can be transmitted out of the smart sensor. Eventually, the sensor array will be used for both reception and transmission in a feedback system and chronically implanted within the eye. Although the microsensor array and associated electronics have been developed, the challenge at this point is the wireless networking of these microsensors with an external processing unit in order to process the complex signals to be transmitted to the array.

The reasons for transmitting information *out of* the retina are perhaps less obvious. The purpose of this reverse flow of data is to determine the mapping between the input image and the resulting neurological signals which enable us to see that image. Currently, the actual mapping from input image to brain pattern is not well understood. The ultimate goals of this research will not be realized without a better understanding of these processes. It is not feasible to do the processing internally using the capabilities of only the sensor arrays. Thus, work on interconnecting these smart




Figure 2: Location of the Smart Sensor within the Eye

sensors with an external processing system is a fundamental aspect of realizing the potential of an artificial retina. On-going diagnostic and maintenance operations will also require transmission of data from the sensor array to an external host computer. These requirements are in addition to the normal functioning of the device, which uses wireless communication from a camera embedded in a pair of eyeglasses into the smart sensor arrays.

Figure 3 depicts the processing steps from external image reception to transmission to the retina prosthesis. A camera mounted on an eyeglass frame could direct its output to a real-time DSP for data reduction and processing (e.g., Sobel edge detection). The camera would be combined with a laser pointer for automatic focusing. The DSP then encodes the resultant image into a compact format for wireless transmission into (or adjacent to) the eye for subsequent decoding by the implanted chips. The setup in Figure 3 shows a wireless transceiver that is inside the body, but not within the retina.

Our ultimate research goal is to support an array of 1600 smart sensor chips, each with a 25×25 grid of electrodes. The investigations into understanding the visual processing of the brain will indicate whether or not the sensor arrays will be implanted with uniform distribution. Functionally, electrode arrays within the center of the macula (the central retina) will have to stimulate the retina differently than peripherally placed electrode arrays, since the functions of these various parts of the retina are very different. Centrally, in the macula, we perceive our high resolution detail vision, while in the periphery, the retina is better at detecting motion or illumination transients. (For example, most persons can perceive their computer monitor's vertical refresh when looking at the monitor using peripheral vision, since the peripheral retina has better temporal resolution,

but poorer spatial resolution than the macula.) Thus, a multi-electrode array visual prosthesis will have to encode the visual scene slightly differently, depending upon where on the retina each electrode array is placed. The peripherally placed electrodes need to generate signals based on lower spatial resolution with greater emphasis on temporal events, while centrally placed sensor arrays upon the macula need to encode more spatially-oriented information. Each array will have to transmit some common information such as the overall luminosity of the visual scene. So, each smart sensor will have to be coordinated with other smart sensors based on an image processing algorithm designed to control a set of smart sensor arrays, each separate, sending input to functionally different retinal areas.

2.1.1 Networking and Processing Requirements

In this section, the functional requirements of the system are described along with some design options. The proposed solutions are motivated by these requirements.

In order to achieve the envisioned functionality, two-way communication will be needed between an external computer and cortical implant so that we can provide input to the cortical implant and determine if the desired image is "seen". We also need two-way communication with the retinal implant so that we can determine that the sensors in the retina are operating as expected. Besides input from the camera, we also need the ability to provide direct input to the retinal implant to determine if the patient sees what is expected from that input pattern. This will validate our understanding of the signaling between the camera and the smart sensor array as well as the operation of the wireless communication protocols.

The plan is to scale up to a 25×25 smart sensor array on a single chip for human retinal implantation. As the research progresses, many of these chips will be inserted for a final system of a 1000×1000 sensor array (40×40 array of chips, where each chip has a 25×25 sensor array). The rods and cones fire at an interval of approximately 200 – 250ms. Therefore, the processing will be performed periodically in a 200 – 250ms processing loop. Hence, data will be transmitted four or five times per second. Although the actual rods and cones in the eye operate in an analog manner (variety of possible values), our initial system will operate in a strictly on/off mode. In other words, one bit of data per sensor every 200 – 250ms. We plan on eventually moving to multiple-level stimulation.

The power must be carefully controlled to avoid damage to the retina and surrounding tissue. Each sensor array operates with less than one microamp of current. The power can be provided in different ways. One option is to use wires to provide the power, although we would still require wireless data communication to limit the number of wires. The power could be provided by implanting a battery near the eye. A second option is to use inductance, provided by RF or IR signals. A third option is a photo-diode array, which converts light to power. It is important to note that even if the power source is wired, the data communication needs to be wireless in order to minimize the number of wires and improve the flexibility of the system.

After considering all factors, the decision has been made to use radio frequencies for both power inductance and data transmission. The human body is mostly water and thus has attenuation characteristics similar to water. This prevents

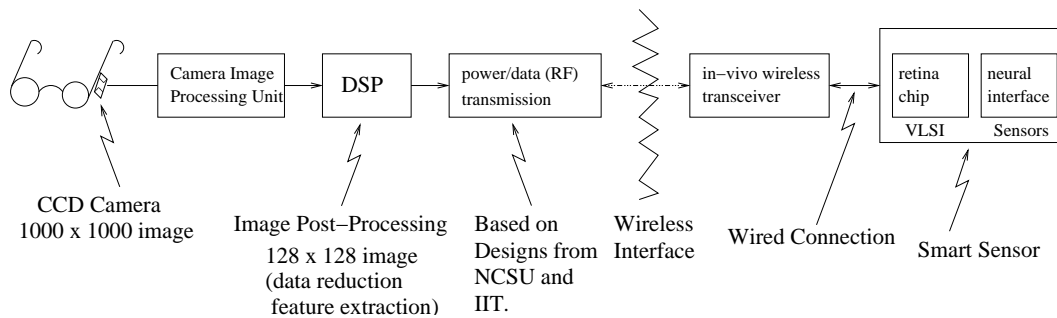


Figure 3: Processing Steps for the Retina Prosthesis Project

our design from using extremely high frequencies, although frequencies in the low GHz range can be used. Because of obvious restrictions on the size of the antenna that can be implanted in the eye, a high frequency, in the GHz range, will be used.

2.2 Glucose Level Monitors

The US National Institute of Health's (NIH) National Institute of Diabetes and Digestive and Kidney Disease reported 15.7 million people – 5.9% of the US population – had diabetes in 1999 [16]. Complications that arise from diabetes include heart disease, stroke, high blood pressure, blindness, kidney disease, and amputations. As the seventh leading cause of death in the United States, it is clear that this is a disease that needs serious attention [16]. Typical treatment for diabetes includes a strict regimen of diet, exercise, insulin injections, and home blood monitoring. Wireless biomedical sensors may afford a more effective way to treat diabetes, by providing a more consistent, accurate, and less invasive method for monitoring glucose levels. Currently, to monitor blood glucose levels, a lancet is used to prick a finger; a drop of blood is placed on a test strip, which is analyzed either manually or electronically. This constant pricking several times a day over a period of years can damage the tissue and blood vessels in that area. Wireless biomedical sensors could be implanted in the patient once. The sensor would monitor the glucose levels and transmit the results to a wristwatch display. This would allow for a less invasive way to monitor the glucose levels several times daily, as well as providing more accurate results. The glucose would also be monitored more regularly, alerting the patient sooner to fluctuations, allowing for corrective measures earlier. This system could even be taken one step further: insulin could automatically be injected when a certain threshold glucose level is reached. This would allow for a more accurate insulin amount to correct the problem.

2.3 Organ Monitors

In early 2001, there were over 74,500 people waiting for an organ donation in the US [17]. Nearly 5,000 patients on donor lists died while waiting for an organ in 1998. Unfortunately, more than one third of the hearts and livers available for transplant go unused, for there is a very short window of time, depending on the organ involved, in which to transplant the organ before it is no longer viable. For instance, preservation time is 4-6 hours for a heart, 12-24 hours for a liver, and 48-72 hours for a kidney [17]. It is

possible that having a better understanding of the organ's condition could lengthen these times. Some researchers are looking into developing a gas monitor with smart sensors to achieve this. The sensor would monitor levels of in vitro heart carbon dioxide, oxygen, and methane to determine heart viability during transplant. This could also be used for basic laboratory cardio-research experimentation.

2.4 Cancer Detectors

Cancer is the second leading cause of death in US, with 563,100 deaths reported in 1999. Currently 9 million living people have had a cancer diagnosis, with 1,221,800 new cases diagnosed in 1999 [21]. Although there is no conclusive evidence on how to prevent cancer, early detection is key to surviving this life-threatening disease. Wireless biomedical sensors may play a key role in early detection. Studies have shown that cancer cells exude nitric oxide, which affects the blood flow in the area surrounding a tumor [10]. A sensor with the ability to detect these changes in the blood flow can be placed in suspect locations. It is likely that any abnormalities could be detected much sooner with the sensors than without. This is especially useful for people who have a family history of certain types of cancer, are at high risk, or are in remission. Research is also being conducted on placing sensors on a needle, enabling physicians to diagnose tumors without having to do a biopsy [3]. The sensors used in this device have the ability to differentiate between different types of cells, identifying cancerous cells.

2.5 General Health Monitors

Wireless biomedical sensors have also been proposed for use as implanted or ingested general health monitors. NASA has been working on several different biotelemetry and bioinstrumentation sensor systems to support Life Sciences research during space flight and ground-based studies [1]. Each of these pill-sized packages would have a variety of sensors that would constantly monitor various substances or tissues in its environment. One proposal would help people with digestive problems. Swallowing one pill containing a smart sensor might enable wireless transmission information about intestinal acidity, pressure, or contractions of intestinal smooth muscle, allowing doctors to better diagnose gastrointestinal diseases in a non-invasive manner. Another such pill might measure ECG, EEG, EMG, and/or heart rate for people with heart disease. These monitors can be used not only in a doctor's office setting, but also to monitor the health and performance of soldiers or firefighters.

3. EXISTING CHALLENGES

In this section, we describe the many challenges for wireless communication with biomedical sensor networks that are unique or fundamentally different from other sensor network application domains. Although some subset of these requirements may be shared by other problem domains, the combination of these requirements make wireless networking of biomedical sensors significantly different from other types of sensor networks. As a result of these differences, different research problems arise and different solutions to these problems are required.

3.1 Low Power

Wireless sensor networks have power restrictions, whether they are biomedical or otherwise. This is due not only to the small physical size of the sensor, but also because of the absence of wires. Without wires to supply a constant source of power, there are only a few options. One is to have an integrated power supply, such as a battery. This supplies power for only a limited time, however, even with conservationist techniques. If the node is implanted in the body, it is not practical to replace the battery as often as would be required. Passive power sources, such as solar and vibration, provide insufficient power for continuous operation. This leads to an external, wireless source of power being the only feasible solution. Options for power supply are discussed in section 3.4.

Wireless sensor networks require some form of energy to operate. This energy is used for various functions in each node, including running the sensors, processing the information, and data communication. Ideally, this power will be evenly distributed and consumed among the nodes in the network. An even consumption of power would allow the nodes to be recharged simultaneously, thereby reducing the use of bandwidth for recharging.

Biomedical wireless sensors add additional constraints with regard to power, including the heat dissipated from using the power. Depending on where in the body the sensor is placed, the allowable amount of heat dissipated varies. Chronic implantation requires much lower dissipation, so as to not damage the tissue surrounding the sensor. In the case of the retina prosthesis, for instance, the temperature of the eye is lower than the core body temperature. Medical doctors are unsure of the effect of raising the temperature of the eye. It is possible that bacteria that are usually absent will thrive in the warmer temperature with the sensors as a growth place and the immune system won't be able to effectively combat these bacteria. It is uncertain whether any subtle effects might occur in the enzymatic reactions as well. Although we are still not certain of the effects of chronic implantation on the body, researchers are currently studying the effects.

3.2 Limited Computation

Computation is directly limited due to the limited amount of power. Although it is acknowledged that sensors are not expected to have the computing power of workstations or even mobile hand-held devices, the amount of computation that is possible with a biomedical sensor is significantly less than the capabilities of a generic sensor. This is true for a variety of reasons, one of which is the power limitation. Communication is very expensive in terms of power, as relatively more power is used to communicate than to compute.

Since we are dealing with a small, finite quantity of power and communication is vital, it is not surprising that little power will be remaining for computation. There are several possibilities to overcome the issues that arise because of this constraint. One method would be to employ a data compression mechanism.

Data compression allows the same amount of data to be transmitted in fewer bits. There are many different algorithms for compressing data, each with advantages that are appropriate for different applications. In many cases, such as with the insulin monitor, the sensor will perform periodic polling of the environment. The sensor may continually process this data and relay the cumulative readings once every 5 minutes. Although the sensor itself will likely perform some local processing of the data, large amounts of data might still be left to transmit. As energy efficiency is key, the less data to communicate, the less power the sensor will consume. In addition to signal compression, image compression might also be needed for the retina prosthesis. In this instance, the images coming from the external camera will likely be compressed before being transmitted into the internal sensors, reducing the amount of information that needs to be transmitted through the limited bandwidth. One method for compressing the image uses a form of segmentation. A small area of the image is examined and highlights are noted. This might be a point, line, or edge, depending on the method being employed. Once these highlights are discovered, adjacent areas are then compared to see if any of the trends follow. If they do, this might indicate a larger feature is present. Point, line, and edge detection are all very computationally intense, however. An external processor will need to assist this process.

On the one hand, good signal quality would dictate a high redundancy rate. This would insure that the correct information is getting from the sensor to the external device or vice versa. High redundancy, however, is very costly with the precious amount of power that is available. The same is true of a common technique called data fusion, which entails several nodes pooling their information together for increased computational power processing and accuracy. There are some ways to work with these issues. One might call for limited wired connections. For instance, a sensor node might be wired to a nearby communication node [4]. This would allow for the sensor node to sense and process some data while the communication node could then send out this processed information. The communication node could use its energy explicitly for sending or receiving data, yielding a longer time before the necessary recharge.

It is expected that for some applications, such as blood glucose monitoring, the ability to transmit data to an external device will be required for further data processing. This external device may be the result display unit, for instance. This is a very attractive idea, for it allows further computation in a place with greater processing capabilities and less inhibited power supply. Other applications might demand sensors with varying capabilities that communicate with each other and send out one collaborative data message. For instance with the artificial retina prosthesis, there might be one type of sensor for the peripheral vision and a different type for the fine detail vision, as different parts of the retina process information differently. This could give the most accurate replication of what is normally occurring in the eye.

3.3 Material Constraints

Because biomedical sensors will be implanted within the body, the shape, size, and materials will be restricted. The size and shape are determined in part by the application. For example, a smart sensor designed to support the retina prosthesis must be small enough to fit within an eye. Assuming multiple chips, each with many sensors, will be used, each covering a small portion of the retina, each chip must be small enough to accommodate this. A gastrointestinal monitor, on the other hand, need only be swallowed, so the size of a pill would be sufficient. The sensors may have to curve around the retina or accommodate a bone or cartilage as well, having an impact on the sensor design and material use.

Even if all of the computational issues previously discussed can be resolved, it is of no use if the sensor cannot safely and effectively function in the body. The human immune system is designed to combat foreign substances in the body. It is therefore of utmost importance to have the sensor nodes composed of biocompatible materials. These materials must act as a barrier between the fluids and the sensor components; however, they must not inhibit the operation of the sensor. The body may begin to reject the implant, which has an adverse effect on the sensor's ability to function properly. This might include the sensor node becoming encapsulated by fibrous tissue. Materials are also restricted to ones that are not adversely affected by bodily fluids. Implants are prone to a condition called "thrombosis," which occurs when the platelets in the blood begin to coagulate on the surface of the sensor. If the conditions inside the body cause a reaction with the sensors, normal operation over an extended period of time is not possible. Similarly, the normal operation of the smart sensor should not cause damage to the body. For example, leakage of toxic or hazardous substances, including those that occur as a result of biological interaction with the device, must be avoided.

There are options for packaging the sensor nodes in biocompatible materials. Common materials for biomedical devices include synthetic polymers or biological polymers, metals, or ceramics, although not all of these will be suitable for packaging sensors. While glass packaging is inert with respect to the body, it also has poor thermal properties, dissipating much less heat than other materials. The excess heat is then dissipated through parts of the sensor node that are not encased in glass, such as the sensor contact points. This localized heat dissipation is likely to cause permanent damage to the surrounding tissue. Ideally, the sensor package should be part of the smart sensor node itself. Some crystalline materials that can be grown on a silicon substrate include diamond and aluminum nitride [9]. These materials have high thermal conductivity properties, which makes them good choices for sensor packaging. Diamond, unlike aluminum nitride, does not have an oxidation layer, however. This means that the diamond will react with iron in the body. Aluminum nitride, on the other hand, is inert with respect to the body, making it a feasible choice for biomedical sensors.

3.4 Continuous Operation

Wireless sensor networks have the potential to enhance our lives in many ways, and researchers are thinking of new applications daily. Most of these sensor networks, however, are designed to operate on limited battery power. This

might be because the sensors are for military reconnaissance with no hope of retrieval or because of cost constraints. Because of these limitations, much research effort is currently being directed toward improving the battery life of sensors in a network, for the main purpose of prolonging the usefulness of the network. This usually translates into a need for low power requirements, for once the battery has been depleted, the sensor node is defunct.

This is not the case, however, for biomedical sensors. Ideally, once a sensor is implanted, it will remain in operation for several years, possibly decades. The medical procedures and risks associated with sensor implantation make frequent placement and adjustment of sensors impractical. Furthermore, based upon how the body reacts with the sensor, it may be difficult to swap one sensor for another. For example, in the case of the retina prosthesis composed of smart sensors, the specific placement of each sensor relative to the nerves and ganglia in the underlying tissue will determine the operation of the sensor network. Depending on the proximity between the ganglia and the sensor probes, different voltage levels may be required. Based on feedback from the patient, the behavior of the sensors will be adjusted. It is thought that the patient's brain will also adapt to this new form of stimuli, although more testing must be completed to validate this. Implanting new sensors could drastically alter this interaction, even if the placement of the sensors does not change radically. Hence, the implanted sensors are expected to be operational for extended periods of time. In the case of chronically implanted devices, years or even decades of operation are expected. Ideally, the sensors would remain implanted for the life of the patient. Although power management is still important, the working model requires that biomedical sensors remain functional, perhaps with the regular infusion of additional power. Since the extracorporeal (external to the body) power supply should be compact and portable, power restrictions still apply, but the sensors will require sufficient power for chronic functionality.

The long-term placement would require a renewable power source, as battery power would be insufficient for the length of time that is being proposed. There are several alternative power sources, however. Interesting research is being conducted on using motion from walking to generate power. A special insole made of a piezoelectric material is worn inside of the shoe. When strides are taken, stress is put on that material, which is then converted into electricity. Solar energy or mechanical vibrations can also be harnessed to power sensors. However, with any of these choices, the energy generated still needs to be transferred wirelessly into the body, as leads coming out of the body are not practical. It is also unclear that these passive techniques can reliably generate the amounts of energy needed. Other possibilities include using inductance, provided by radio frequencies (RF) or infrared (IR) signals. Since the energy travels in waves, there is no need for wires to carry the energy from outside to inside the body. RF inductance has the ability to communicate to many sensors simultaneously. IR, on the other hand, is a much more directed beam, only able to transmit to a single destination. IR is also bound by line of sight; that is, the receiver must be visible to the transmitter. This is likely to cause issues when dealing with components inside the body. The FCC has specified the 2.4-2.48 GHz range for unregulated Industrial-Scientific-Medical uses of RF. This range is not constrained by line of sight for short distances, such as

those that would be used for biomedical sensor networks. It should be noted, however, that there is more absorption in the body at higher frequencies than at lower ones. In special applications such as the retina prosthesis, a photodiode array, which converts light into electrical power, may also be considered.

3.5 Robustness and Fault Tolerance

Biomedical sensors are expected to last a long time, as it is not desirable to surgically adjust sensors every week, month, or even year. Ideally, a sensor would last forever, although a more realistic lifetime would be in the range of decades. Because of this, the sensor network should be extremely robust and fault tolerant. In particular, the failure of one node should not cause the entire network to cease operation. The best method of achieving this is making a distributed network, one where each sensor node can operate autonomously from its neighbors, though still cooperate when necessary. If a sensor does stop working, then the sensors in the surrounding area should still function as normal.

Indeed, this concept should work within individual sensor nodes as well: if a sensor can no longer process data, for instance, that data should still be received and passed on if communication is functioning. One way to achieve this would be to have several small components wired together that function as a group. This node might be comprised of a sensing block, a communication block, a scheduling block, and/or a data block. This would be a good way to isolate the malfunctioning block from the rest of the components in the node, as well as reducing power consumption among the various components.

In order to ensure that the proper data is being sent and received, there are a few options that can be used, whether singly or in combination, including checksums, parity check, and cyclic redundancy check [15]. One of these is the common idea of checksums, a simple error detection scheme. With checksums, each data message is sent along with a checksum value based upon the ones' complement or two's complement of the sum of the words in the message. When the message is received, a checksum is computed on the incoming message. The two values are then compared. If they are the same, then there is a high probability that the message was sent without data loss. If they are different, however, then the message has been altered in transmission and needs to be sent again.

Other forms of error correction include parity checking and cyclic redundancy check (CRC) [15]. Any of these schemes can be used to insure the integrity of the data being transmitted, although they all have trade-offs. While parity checking is rather simple and does not use any memory, it has limited error detection capabilities. Specifically, a one-bit parity check can detect a single bit error. This is an acceptable choice when errors are rare. CRC, on the other hand, can detect a greater variety of errors. The message is treated as a number, which is divided by a predetermined generator. The remainder of this computation is appended to the message. When the message is received, the remainder is computed and compared. Based upon the generator that is chosen, all single-bit errors, almost all double-bit errors, and any odd number of errors can be detected. In order to compute a CRC, additional hardware, such as a shift register and XOR gates, is required. Based upon the application, there may not be room for these additional com-

ponents on a biomedical smart sensor.

Diagnostics are another important feature that these systems will need to incorporate. Since these sensors are going to be implanted in and monitor the human body, there are likely to be varying responses from different people. In the case such as the glucose level monitor, different patients have different "normal" levels of glucose in their blood stream. The sensor will have to be fine-tuned to give the proper care to each patient. In terms of the retina prosthesis, there is no regular spacing between the ganglia. Because of this, a sensor might be placed in the area between two ganglia, rather than directly on top of this. Voltages will need to be adjusted at each node to insure proper vision, based on feedback from the patient. Diagnostics will be vital for testing prototypes of these devices as well. Without them, there would be no clear feedback on what is actually occurring inside the body.

Materials must also be carefully chosen to avoid unintended interaction with the body. This will also aid in reducing the failure rate. Designing simple sensors and implanting multiple sensor nodes will lead to more robustness than implanting one monolithic sensor node. Many small, low-power smart sensors distributed among a larger area will likely cause less damage to the tissue than one large, power-consuming sensor block.

3.6 Scalability

Previous constraints mentioned, such as the need for low power and confined available placement areas, forces the sensor to be relatively small with limited functionality. Yet there is still the need for substantial amounts of functionality. In order to ensure the sensor networks are still as powerful as necessary, many small sensors may need to be placed and work in conjunction with each other. With multiple sensors communicating amongst themselves, as well as to the output device, lots of communication will be taking place, demanding the efficient use of the wireless spectrum. Current wireless sensor protocols do not sufficiently support the number of low-power sensors needed, nor is it clear that these protocols will scale to the size necessary for biomedical needs.

It is also important to realize that it is unclear how many sensors will need to be placed to make the system functional. At this point, it is better to build in the ability to increase the number as the technology becomes more refined than to redesign the system each time the sensor count is increased.

3.7 Security and Interference

As with any medical information, the issue of confidentiality arises. Although it is important that the physician or the patient see the feedback from the sensors, it is not information that necessarily should be broadcast publicly. In addition, it is not desirable for an outsider to gain access to the sensors or the display. It could be dangerous, even fatal, for false readings to appear on a patient's glucose monitor output. The same is true of erroneous images being projected in the eye or for any other function of the biomedical sensors. Because of these requirements, it is advisable that strict security mechanisms be put in place that would prevent malicious interaction with these systems. Although it may seem attractive to encrypt all of the data, any meaningful, strong encryption would be too computationally intensive to be practical for these uses. This is clearly an open

issue that has no obvious solution at this time.

Between all of the proposed applications for biomedical wireless sensor networks and the great number of people afflicted by diseases that might be helped with the use of these networks, it is not unreasonable to expect hundreds of thousands of these networks in place in the next decade. This can lead to the problem of interference between wireless networks in people standing next to each other or even the possibility of conflicting signals within one person. It is clear that information from one person's network should not manifest itself on someone else's display system. Interference might also come from other types of wireless communication, such as microwave ovens or Bluetooth devices. Several options exist for dealing with this, though it is not clear at this time which one will be the most effective.

3.8 Regulatory Requirements

As these biomedical sensors are being used to treat disease, it is not surprising that the Food and Drug Administration (FDA) regulates their testing and use. There must be some evidence that these devices will not harm, and potentially help, the test subjects. Procedures for protecting patients have been developed for clinical trials. While animal testing is used for preliminary data, it is often insufficient for determining the effectiveness of the sensor networks. With the retina prosthesis, for instance, using the feedback and diagnostics, a researcher can stimulate the sensors and view the brain activity of a test cat. The cat, however, is unable to verbalize what it is seeing, if anything. Verbal feedback from human test subjects will be necessary for examining the limitations and usefulness of prototype devices, as well as guiding procedures of refinement before eventual chronic implantation in humans.

For this reason, even prototype devices will have to meet the strict standards of patient safety before any human testing can be done. The wireless transmission of data must not harm the surrounding tissues and the chronic functioning and power utilization of these devices must also be benign. Design for safety must be a fundamental feature of biomedical sensor development, even at the earliest stages. Reasonable evidence of design efficacy will be required even for prototype devices.

Biomedical sensors create additional, fundamental issues that are not inherent in most other types of sensor networks, namely, should we do this? Because of the debilitating effects of many of these diseases, sufferers of these diseases are likely to be willing volunteers to advance science in this area. Researchers in biomedical sensors must consider ethical and moral issues that do not arise in most other sensor applications. It is conceivable that some unscrupulous researchers could perform tests and trials with devices that are dangerous to the volunteers. Even with informed consent, it is possible that people that suffer from these ailments would overlook obvious dangers in the search for a cure. Therefore, it is imperative to have diligent oversight of these testing operations.

4. RESEARCH APPROACH

Biomedical wireless sensor networks have certain advantages over other types of wireless networks. In general, wireless nodes are considered to be mobile, such as with cellular phones [14, 19, 23]. Implanted biomedical sensor nodes, however, are relatively stationary. This fixed placement

can be exploited to create efficiencies that are impossible in mobile networks. There is no need for a self-organization protocol to allow individual sensors in the network to determine the proximity and number of neighboring sensors. Nor is there any requirement to develop a routing protocol among sensors to determine paths among sensors that need to communicate. Although such protocols are invaluable for networks of sensors that are placed by simply ejecting a large number of sensors from an airplane or manually strewn about a field, these issues do not arise in biomedical sensor applications, where the sensors are surgically implanted and the number and placement of the sensors are determined beforehand. By encoding this information into the sensors and exploiting this information in the design of the sensor application, optimized communication strategies can be defined for biomedical sensing tasks. The power expenditure for building routing tables and developing knowledge of the network topology can be avoided. Thus, energy efficiency can be obtained from an application-specific biomedical sensor network. Currently, researchers are investigating which topologies might be the most effective for the constraints set forth for a biomedical sensor network, particularly power management.

Although we can ideally exploit these fixed topologies, there are some exceptions. There is the issue of node failure as discussed above in section 3.5. There is also the case in which a sensor or sensors are placed in limbs. If the sensors are placed solely in the limb, this should not cause any problems, for the distance between nodes will not be changing. When nodes are also placed in other parts of the body, however, a non-stationary placement relative to the torso occurs. Ideally, the sensor reading shouldn't vary if an arm is out-stretched, extended above the head, or down by the patient's side, although this variance in distance could give alternative readings. Because of this possibility, the communication protocol must allow for this situation in some applications. On the other hand, if the clustering can be done without too much power consumption in the clusterhead, the clustering approach should lead to an overall reduced power requirement.

Wireless networks of biomedical sensors will require highly customized and application-specific protocols to be developed. We are currently working on developing protocols that have general applicability for sensor networks of these types as a starting point from which certain optimizations can be realized for specific biomedical applications. In other words, we do not intend to design a unique solution for every biomedical application. Instead, we posit that techniques can be developed for biomedical sensors that can be used as building blocks for an optimized solution of specific applications. As an example, efficient wireless bandwidth usage is required for biomedical applications. Some trade-offs exist between processing overheads and communication efficiencies. Depending on the sensor capabilities, the application may arrive at different points on this design spectrum, but a flexible design for optimizing bandwidth usage for different processing levels allows for an informed decision. Following are some specific issues that we have to consider when designing protocols for wireless communication in biosensor networks:

- The frequency range we select for communications plays an important role in the design and performance of a biosensor network. There is a direct relationship be-

tween frequency and tissue warming. The higher the frequency of the EM signal, the higher is its absorption by the tissue and more the tissue warming. Hence it is desirable to use lower frequencies for communications. However, the lower the frequency, the larger will the antenna dimensions have to be. We will therefore have to make a trade-off between antenna dimensions on the one hand and greater tissue warming on the other.

- The human body is composed mostly of water. Since water absorbs a lot of radiation, the percentage of water in the tissue has a bearing on the signal strength of the EM signal. However, the water content is not uniform throughout the body. Fat and bone are relatively dry and absorb less energy when compared to wetter tissues such as muscle [18]. This variation in energy absorption, coupled with the possibility of tissue damage due to heating by radiation, complicates the design of biosensors and biosensor networks.
- In order to avoid excessive radio interference as well as having to obtain FCC licenses, it may be advisable to use the designated, unlicensed ISM frequency bands. An FCC regulation, however, makes it mandatory for all communication on the ISM frequency band to use spread spectrum techniques. This increases the complexity of the transceivers on the biosensors.
- As was mentioned earlier, biosensor networks are power constrained. Therefore, the communications protocols we develop will have to consume a minimal amount of energy. If feasible, we should avoid contention-oriented protocols such as CSMA/CA and adopt a TDMA approach instead.

Similarly, we are currently analyzing the performance issues associated with different network topologies. Unlike previous studies of these issues, mobility is not an issue. The question we are seeking to answer is what is the best topology for a wireless network of sensors, assuming we can control the placement of these sensors and the sensor locations are fixed relative to each other. One factor in the choice of topology is the amount of contention for the wireless media. The level of contention will vary with the application, since the message pattern and overall message generation rate are functions of the biomedical application. However, our study should provide some insights which can be used along with knowledge of the biomedical application to select an appropriate topology. Again, the goal is not to find a single topology that is appropriate for all applications, but rather provide a structured analysis of the options and give guidance on the best choices so that a more informed decision is possible.

5. PRELIMINARY RESULTS

In this section, we briefly describe the results from our preliminary investigation on power-efficient topologies and wireless networking protocols. We will refer to each biosensor chip as a *node*.

5.1 Power-Efficient Topologies

Our initial studies have focused on analyzing the performance of wireless sensor networks when arranged in different topologies. Since we assume control over the placement of

these sensing nodes and do not require mobility of the sensors relative to each other, the research problem changes.

The summarized results presented next are for network topologies where each wireless sensor is placed within range of a specific number of neighbors. For example, the nodes are arranged into a 2D mesh with a maximum of four neighbors. The sensors on the edges of the network have fewer neighbors, but each sensor has the same transmission range.

In [20], we have analyzed the power dissipated with respect to the network topology with differing numbers of neighbors. We considered two-dimensional network topologies with three, four, six, and eight neighbors as well as a three-dimensional network with six neighbors. Some of these network topologies are shown in Figures 4 – 7.

The power used by a transmitter is a function of the distance and number of bits transmitted. Similarly, the receiver expends power to receive the message. A clear trade-off exists: having more neighbors implies a network with a smaller diameter, while having fewer neighbors implies reduced wireless channel contention and less power expended receiving a message. For our preliminary work, we have adopted the same energy model as Heinzelman *et al.* [13].

5.1.1 Two Dimensional Analysis

Figures 4 – 7 show the topologies with differing numbers of neighbors. It is clear that as the number of neighbors increases the number of alternative paths increases.




Figure 4: Topology with 3 Neighbors

In table 1 the power usage for a route across the diameter of the network through the interior is shown. Table 2 shows the power usage when the path instead travels along the edges of the network. In both cases, the network consists of 36 nodes. Edge routing consists of moving messages to the outer edges of the network where there are fewer neighbors. Interior routing keeps the messages in the middle of the network, where there is a consistent number of neighbors for each node. In some cases, longer paths were chosen to give a similar number of transmissions. From tables 1 and 2, edge routing dissipates less power than interior routing in all cases except for 3 neighbors. This is because the 3 neighbor




Figure 5: Topology with 4 Neighbors




Figure 6: Topology with 6 Neighbors

network makes edge routing difficult. With either routing strategy, as the number of neighbors increases the power dissipated increases for the same number of transmissions.

In table 3, we consider the power dissipated between the source and destination for a message spanning the diameter of the network for networks with 3 and 6 neighbors. See figures 4 and 6 for depictions of these two network topologies.

As we can see from table 3, increasing the number of neighbors decreases the number of transmissions and the total power dissipated in the system. This result can only be attributed to the availability of a shortest path between the source and destination. Similar analysis can be concluded from table 4, which shows a similar comparison for networks with 4 and 8 neighbors. There is a trade-off between the number of neighbors and the total power dissipated in the system. However, this trade-off breaks in special cases where the availability of alternative shortest paths can be used as an advantage for the power budget calculations.

Table 1: Interior Routing, 2D

Neighbors	Transmissions	Receptions	Joules Used
3	10	27	9.473×10^{-4}
4	10	36	11.777×10^{-4}
6	10	52	15.873×10^{-4}
8	10	69	20.225×10^{-4}

Table 2: Edge Routing, 2D

Neighbors	Transmissions	Receptions	Total Power
3	14	33	12.034×10^{-4}
4	10	28	9.729×10^{-4}
6	10	37	12.033×10^{-4}
8	10	46	14.337×10^{-4}

Table 3: Routing Freedom and Power Dissipation; 3 and 6 Neighbors

Neighbors	Transmissions	Receptions	Joules Used
3	10	27	9.473×10^{-4}
6	5	27	8.193×10^{-4}

Table 4: Routing Freedom and Power Dissipation; 4 and 8 Neighbors

Neighbors	Transmissions	Receptions	Joules Used
4	10	36	11.777×10^{-4}
8	5	38	11.009×10^{-4}




Figure 7: 3D Topology with 6 Neighbors

Table 5: Edge and Interior Routing Power Dissipation

Network	Path	T_x	R_x	Joules Used
2D	Interior	10	36	11.777×10^{-4}
	Edge	10	28	9.729×10^{-4}
3D	Interior	7	27/33	$8.705 - 10.241 \times 10^{-4}$
	Edge	7	25	8.193×10^{-4}

5.1.2 Three Dimensional Analysis

Figure 7 depicts a three-dimensional network topology with six neighbors, which has some advantages due to its inherent symmetry. In 3D, the routing paths between any given source and destination without misrouting would always result in the same number of transmissions but different number of receptions, just as in two dimensional case of 4 neighbors.

Table 5 compares 2D with 4 neighbors and 3D with 6 neighbors with a fixed number of transmission. Considering both the interior and edge routing, we conclude the following observations: (1) Edge routing in the case of the 3D network has lower power dissipation than interior routing. (2) The number of transmissions and receptions, and the total power dissipated in the three dimensional network is less than the two-dimensional network for both edge as well as interior routing.

Table 6 presents the results of transmitting 100,000 messages, each with a randomly chosen source and destination, through each of the different networks. Two different routing protocols were used, although only one is presented here. The first selects a shortest path, while the second uses the available power at each neighbor to select the neighbor closer to the destination that has the most remaining power. Details of these two routing protocols can be found in [20], where it is shown that power-aware routing reduces the power usage of the network. The results show that there is a trade-off between the number of neighbors and the total

Table 6: 100,000 Node Pairs

Neighbors		DSAP routing (Not Power Aware)		
		T_x	R_x	Joules Used
2D	4	388540	1369487	45.010465
	6	331801	1723883	52.629757
	8	274405	1908911	55.896402
3D	6	302160	1312998	41.35415

power dissipated in the system. Topologies with fewer number of neighbors dissipated less power, although the topologies with more neighbors require fewer hops.

Because the number of neighbors differs with different topologies, one expects different topologies to have different power usage rates. Even our simulations of the contention-free case show that different topologies have different levels of power efficiency. The preliminary results show that the total power consumption is reduced for topologies with fewer neighbors; although the topologies with more neighbors require fewer hops, the power expended by many nodes to receive these messages increases the power usage. Among the 2D topologies, the best power efficiency is achieved with the 2D topology which allows up to four neighbors. The 3D topology performs best, although a 3D topology may not be feasible for some applications.

5.2 Wireless Communications Protocols

Here we present two communication protocols that were designed to reduce energy consumption. For the purpose of designing general purpose communication protocols for wireless biomedical systems we assume that nodes in the system need bi-directional communication with an external processor via the external transceiver (henceforth called the *base station*). Since the placement of nodes within the body is predetermined and fixed, we assume that the base station is aware of the positioning. All communication is wireless. TDMA is used for media access and we assume that all nodes are perfectly synchronized and are aware of the beginning of each slot.

We observe that energy can be conserved by reducing or eliminating medium access contention. An interesting fact about our biosensor network is that the communications pattern is deterministic and periodic. Each node has to transmit its data once in 250ms. This leads us to adopting a fixed TDMA method at the MAC sub-layer. A TDMA scheme has the added advantage that nodes can sleep when they are not sending/receiving data. This leads to less power usage and extended battery life. In the following we present and compare two approaches, cluster-based and tree-based, from the perspective of biosensor wireless networks.

5.2.1 Cluster-based Approach

In our cluster-based protocol, we work on the idea that energy consumption can be reduced by stipulating that only a small fraction of the nodes are allowed to communicate with the base station. These nodes are called *leaders* and each leader is in charge of data from a cluster of nodes around it. Each leader collects data from nodes in its cluster, compresses the data to eliminate redundancy and transmits the compressed data to the base station. This method is similar to the LEACH protocol proposed by Heinzelman et al [13]. In their study, the authors assume that the system has a




Figure 8: Cluster Based Approach.

limited and non-renewable source of energy in which nodes die as the energy level falls. Their goal is to extend system lifetime and this is achieved by dynamic clustering of nodes. However, this is not an issue for biosensor networks as they will be provided with a continuous source of energy. Hence our sole concern is that of reducing energy consumption. Our static clustering approach has two major advantages:

1. The data is compressed at the source and hence only relevant data is transmitted to the external processor; and
2. Only a small subset of the nodes make a long distance transmission and hence energy is conserved.

There are several issues involved here. First, we must have some way of deciding which nodes should be leaders. Second, the remaining nodes should join a cluster. How does a node decide which cluster to join? Since the location of the sensors are predetermined and fixed, and only the base station has global knowledge about the system, it is best for the base station to nominate the leaders. If we elect the leaders probabilistically, there is a good chance that an elected leader may not be the best choice. For example, the leaders may not be evenly distributed throughout the network and several of them may be on an edge of the network. On the other hand, the base station has global knowledge of the system and is in a better position to nominate leaders. For example, while deciding leaders, the base station can take into consideration the node concentrations and allocate more leaders in denser areas and less in sparser areas. It can also take into account various other factors such as cost of communication etc., and improve overall system performance by choosing an optimal number of leaders with optimal positioning.

Once the leaders have been nominated, each node in the system has to choose a cluster to join. In our scheme, all leaders transmit a signal concurrently but at different frequencies. The frequencies can be allocated by the base station. Every non-leader node in the system scans all the frequencies and identifies the one frequency that offers it the best signal-to-noise ratio. The node then joins the cluster represented by that frequency.

To join a cluster, the node has to inform the leader of the cluster about its intent. But there is a possibility that several nodes send their join requests at about the same time. This will cause a collision. To avoid such collisions, we adapt a TDMA scheme that was originally used as a channel reservation protocol. In this scheme, nodes transmit a binary one in slots corresponding to their address. The




Figure 9: TDMA scheme for cluster-based protocol




Figure 10: Tree Based Model

added advantage of this scheme is that all nodes in that cluster will know which other nodes have joined and in what order.

After clusters have been formed, nodes in each cluster will have to send data to the leader of that cluster. Here again, there is the possibility of collisions if the media is not regulated. To avoid this problem, we propose to use TDMA. Each node in the cluster will be allocated a slot in a frame. The allocation need not be explicitly done as we can simply follow the order in which the nodes joined the cluster. As we stated earlier, each node is aware of who joins the cluster and when, and hence each node is aware of its position in the frame. Once the leader collects data from all the nodes in its cluster, it can compress the data and transmit it to the base station.

5.2.2 Tree-based Approach

In this approach the base station selects one or more nodes to be its children based on factors such as their proximity to itself and node density across the system. These selected nodes then make a low intensity transmission, each at a different frequency. Nodes that receive this transmission at a predefined minimum signal-to-noise ratio can then request the transmitting node to be its parent. To make this request, the prospective children will have to use the subscribing protocol we mentioned earlier. Once the children have selected their parents, it is their turn to solicit children. This continues until all nodes in the system are covered.

By following the above method, we will obtain a spanning tree for the network. When a node wants to send data, it will send it to its parent node by a low-energy transmission. The parent will collect data from all its children, compress the data if required, and in turn transmit it to its parent. Only the children of the root node (base station) will be required to make the high energy transmit to the base station. The TDMA frame structure is given is figure 11.

5.2.3 Performance Comparison

In this section we present the results of our analysis of the two protocols. The parameters we used during the com-




Figure 11: TDMA scheme for tree-base protocol

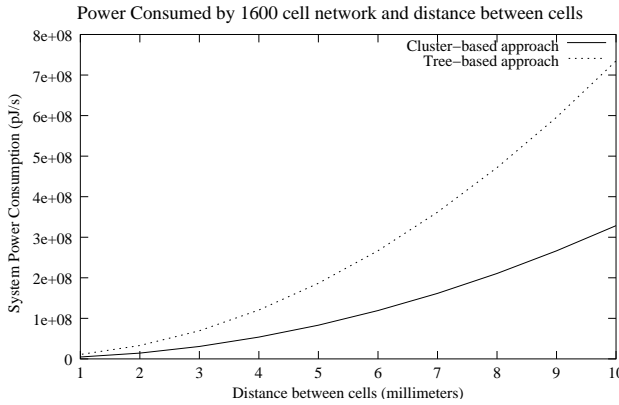


Figure 12: Power consumption vs. distance between nodes for Tree-based and Cluster-based approaches.

parison are given in Table 7. Our network consisted of a 2-dimensional array of nodes with the base station placed above the array. The distance of the base station from the center of the array was assumed to be twice that of the inter-node distance. So far in our analysis, we have concentrated on the power performance of the two protocols. We studied the variation of power consumption with respect to distance between nodes as well as the number of nodes in the network.

Table 7: Parameters and their values used in the performance analysis

Parameter	Value
Transmission cost	100 pJ/bit/m ²
Reception cost	50 pJ/bit
Circuitry cost	50 pJ/bit
Number of nodes	50 - 1600
Distance between nodes	1 - 10 mm

Figure 12 shows the relationship between power consumption and distance between nodes. As can be seen, the cluster-based approach shows better energy-efficiency. Figure 13 compares the power consumption of the two protocols as a function of the number of nodes in the network. Once again the cluster-based approach performs better.

6. RELATED WORK

Existing communication protocols are not necessarily sufficient for the needs of biomedical wireless sensor networks, due mainly to the constraints and requirements discussed thus far. In addition, there are networking requirements that also must be considered before designing a new protocol. When considering a protocol for biomedical sensor networks, it may be acceptable for it to have a high initial-

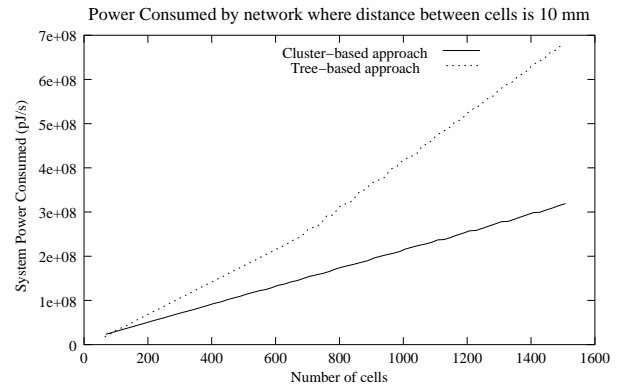


Figure 13: Power consumption vs. number of nodes in the network.

ization cost, for the network will usually be in continuous operation and not require regular restarting. This would mean that any addressing or look-up table initialization could take place in this initial phase. There is also a need for bi-directional communication. This might be needed for upgrading the protocol or software in the body. It will also give the flexibility of a standard for biomedical wireless sensors. Some sensors, such as for glucose monitoring, would receive information from the blood and transmit it to an external device (output display). The retina prosthesis, on the other hand, will be receiving its information from the external device (camera) and transmitting a signal to the smart sensors in the eye. One standardized protocol for both devices would be ideal, but perhaps not realistic.

There is also the issue of multiple sensors communicating within the network through limited bandwidth. Current protocols have been created to deal with this in the cellular realm. A short-range communication alternative may be Bluetooth [22]. Since its standardization, Bluetooth is becoming one of the more popular protocols for wireless connection between personal electronics such as mobile phones, PDAs, and pagers. Bluetooth is designed for wirelessly communicating with these devices at short distances, no more than approximately 100 meters. Keeping with its intended design, Bluetooth also offers lower data rates. This may prove to be insufficient for biomedical wireless sensors that are data intensive, such as the retina prosthesis, which will have hundreds of sensors needing updates four to five times each second.

Another protocol is the IEEE 802.11 standard for wireless LANs [12]. This gives a sufficiently high data rate for most foreseeable applications, between 2 Mbps and 11 Mbps. IEEE 802.11 is also attractive because of the built-in authentication and encryption mechanisms to prevent tampering of the data, which is an inherent concern when dealing with confidential medical data. Because this standard is intended for LANs, such as those found throughout buildings, IEEE 802.11 is capable of transmitting over long distances. The major drawback to this is that it consumes much more power than other alternatives.

One method that may be more applicable to biomedical wireless sensors is the Time Division Multiple Access (TDMA) protocol. TDMA allows multiple users to access a single radio frequency without interference. It achieves this by allocating unique time slots to each user of a chan-

nel. TDMA has a data rate capable of between 64 kbps and 120 Mbps, in 64 kbps increments. These data rates should be sufficient for all currently perceived biomedical applications. Another advantage of using the TDMA protocol is the energy-efficiency that can be gained by scheduling. Since each task is assigned a time and channel slot, it is possible for a node to "sleep" for a given period of time. By allowing the sensor to power down when not in use, great amounts of power can be conserved at each node.

Heidemann *et al.* [7] conducted a study comparing TDMA and IEEE 802.11 in terms of energy consumption. For varying numbers of nodes in the simulated network, IEEE 802.11 consistently dissipated 20 times more heat than TDMA for the data diffusion protocol. The reason cited for this large discrepancy in energy consumption is that TDMA is more energy-conserving. Thus, IEEE 802.11 is inherently inappropriate for sensor networks. Instead, an energy-conserving MAC such as TDMA is recommended for the sensors to insure their long life.

As mentioned earlier, security and reliability is a major requirement for the type of medical applications considered in this paper. However, current solutions to wireless networking, such as IEEE 802.11 and Bluetooth, have security and reliability problems as illustrated by the following examples:

- Security Problems: IEEE 802.11 provides limited support for confidentiality through the wired equivalent privacy (WEP) protocol which has significant flaws in the design as shown in [24]. Recently, Arbaugh *et al.* [2] have shown that all the security mechanisms provided in IEEE 802.11 are completely ineffective.
- Performance/Reliability Problems: The 2.4GHz Industrial, Scientific, and Medical (ISM) band is currently shared by wireless LAN systems (e.g. IEEE-802.11x excluding IEEE802.11a, HomeRF), other communication systems (e.g. Bluetooth, proprietary cordless phones), and non-communication systems (e.g. microwave ovens). There is much concern over the mutual interference between Bluetooth and IEEE 802.11. Studies done by companies thus far indicate that if these two technologies operate at least 2 meters apart, interference would not be significant. There is, however, severe degradation in performance if these two technologies are colocated [5].

A number of research groups have been studying sensor networks. Although these sensor networks do not match the requirements for biomedical sensor networks, a brief review of some of this work is included for completeness.

Byers and Nasser [4] have proposed dividing the sensors into different types to reduce power usage. Some sensors are used for communication and others for the sensor application. Heinzelman *et al.* [13] have proposed using clusters of sensors and rotating the cluster head to equalize power usage among sensors and extend sensor life. Bhagwat *et al.* [8] propose emphasizing battery lifetime and cost over maximizing wireless bandwidth efficiency. None of these models are directly applicable to biomedical sensors applications as described in this paper.

Some recent work [14, 19, 23] has also examined improving power efficiency through topology control. In these papers, however, the authors have focused on adjusting the transmission range, and hence the topology, for a mobile network.

Another interesting research effort on creating extremely small and low-power sensors is the Smart Dust project [6]. This work is not directly applicable to biomedical sensors, however, because optical transceivers are used for communication. Optical transceivers require a line-of-sight path between transmitter and receiver and hence are usable in implanted devices.

Estrin *et al.* [11] have proposed a communication model called directed diffusion for scalable co-ordination in sensor networks. This paper motivates the need for tightly integrating the routing functions with the application. Such approaches would also be useful for biomedical applications.

Some of the issues involved in designing communication protocols for wireless communication within human body are addressed in Personal Area Networks (PANs) research. A PAN enables data communication between electronic devices on and near human body by capacitively coupling pi-coamp currents through the body [25].

7. CONCLUSIONS

Smart sensors offer the promise of significant advances in medical treatment. Networking multiple smart sensors into an application-specific solution to combat disease is a promising approach, which will require research with a different perspective to resolve an array of novel and challenging problems.

As wireless networks of sensors are developed for biomedical applications, the knowledge gained from these implementations should be used to facilitate the development of sensor networks for new applications. In many cases, the sensing materials and low-power electronics already exist for solving these problems. The design of wireless networks for these sensors is the time-critical next step in realizing the benefits of these sensors. By increasing the community of researchers addressing these problems, the time until the human benefits are achieved will be dramatically reduced.

Expedient development of implanted smart sensors to remedy medical problems presents clear benefits to individuals as well as society as a whole. There is the obvious benefit to persons with debilitating diseases and their families as these patients gain an enhanced quality of life. Biomedical implants that monitor for cancer will help recovering patients maintain their health. Visual problems such as retinitis pigmentosa and age-related macular degeneration affect millions and improving sight for these individuals will contribute greatly to their well-being. Not only will these individuals personally benefit from their improved health and well-being, but society will also benefit from their increased productivity and societal contributions. Once the technology is refined, medical costs for correcting chronic medical conditions will be reduced. As the world population ages, the demand for such devices will only increase.

8. ACKNOWLEDGMENTS

The authors would like to thank the anonymous referees for their constructive comments which helped to improve the quality of this paper. We also thank Raymond Iezzi and Pat McAllister (Wayne State University, Medical School) and Pepe Siy (Wayne State University, Department of Electrical and Computer Engineering) for their input on many aspects of this work.

9. ADDITIONAL AUTHORS

Additional authors: Ayad Salhieh (Wayne State University, Department of Electrical and Computer Engineering, email: ai4874@wayne.edu), Vikram Shankar (Arizona State University, Department of Computer Science and Engineering, email: vikram.shankar@asu.edu), Valliappan Annamalai (Arizona State University, Department of Computer Science and Engineering, valliappan.annamalai@asu.edu), Manish Kochhal (Wayne State University, Department of Electrical and Computer Engineering, ah2471@wayne.edu), and Greg Auner (Wayne State University, Department of Electrical and Computer Engineering, gauner@eng.wayne.edu).

10. REFERENCES

- [1] National Aeronautics and Space Administration. Pill Shaped Implantable Biotelemeters, 2000. <http://ctoserver.arc.nasa.gov/techbreak/biotelemeters.html>.
- [2] W. A. Arbaugh, N. Shankar, and Y. C. Justin Wan. Your 802.11 Wireless Network has No Clothes. <http://www.cs.umd.edu/waa/wireless.pdf>.
- [3] M. Braukus, J. Bluck, and M. Goodkind. NASA Developing Computerized Breast Cancer Diagnostic Tool. Press Release 98-184, National Aeronautics and Space Administration, October 1998. Available at <ftp://ftp.hq.nasa.gov/pub/pao/pressrel/1998/98-184.txt>.
- [4] J. Byers and G. Nasser. Utility-Based Decision-Making in Wireless Sensor Networks (Extended Abstract). In *IEEE MobiHoc*, 2000.
- [5] Mobilian Corporation. Wi-Fi (802.11b) and Bluetooth Simultaneous Operation: Characterizing the Problem. <http://www.mobilian.com>.
- [6] B. Warneke *et al.* Smart Dust: Communicating with a Cubic-Millimeter Computer. *IEEE Computer*, 34(1):44–51, January 2001.
- [7] J. Heidemann *et al.* Effects of Detail in Wireless Network Simulation. In *SCS Multiconference on Distributed Simulation*, pages 3–11, January 2001.
- [8] P. Bhagwat *et al.* System Design Issues for Low-Power, Low-Cost Short Range Wireless Networking. In *IEEE International Conference on Personal Wireless Communications*, 1999.
- [9] P. Gluche *et al.* Diamond for Medical Applications. In *High Care 2000*, February 2000.
- [10] S. Reveneau *et al.* Nitric Oxide Synthase in Human Breast Cancer is Associated with Tumor Grade, Proliferation Rate, and Expression of Progesterone Receptors. *Laboratory Investigations*, 79(10):1215–1225, October 1999.
- [11] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. Next Century Challenges: Scalable Coordination in Sensor Networks. In *MobiCom*, pages 263–270, 1999.
- [12] The Working Group for WLAN Standards. IEEE 802.11 Wireless Local Area Networks. <http://www.manta.ieee.org/groups/802/11/>.
- [13] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy-Efficient Communication Protocols for Wireless Microsensor Networks. In *Hawaii International Conference on System Sciences*, 2000.
- [14] L. Hu. Topology Control for Multihop Packet Radio Networks. *IEEE Transactions on Communications*, 41(10):1474–1481, October 1993.
- [15] S. Keshav. *An Engineering Approach to Computer Networking*. Addison-Wesley Publishing Company, Reading, Massachusetts, 1997.
- [16] National Institute of Diabetes, Digestive, and Kidney Diseases. Diabetes Statistics. NIH Publication No. 99-3892, National Institute of Health, March 1999.
- [17] U.S. Department of Health and Human Services. 2000 Annual Report of U.S. Scientific Registry for Transplant Recipients and the Organ Procurement and Transplantation Network: Transplant Data: 1990–1999.
- [18] National Council on Radiation and Measurements. A Practical Guide to the Determination of Human Exposure to Radio Frequency Fields. Report no. 119, NCRP, Bethesda, MD, 1993.
- [19] R. Ramanathan and R. Rosales-Hain. Topology Control of Multihop Wireless Networks using Transmit Power Adjustment. In *INFOCOM*, pages 404–413, 2000.
- [20] A. Salhieh, J. Weinmann, M. Kochhal, and L. Schwiebert. Power Efficient Topologies for Wireless Sensor Networks. In *International Conference on Parallel Processing*, September 2001. to appear.
- [21] American Cancer Society. Cancer Facts and Figures. <http://www3.cancer.org/cancerinfo/>.
- [22] Bluetooth Open Specification. Bluetooth Home Page. <http://www.bluetooth.com/>.
- [23] Z. Tang and J. J. Garcia-Luna-Aceves. A Protocol for Topology-Dependent Transmission Scheduling in Wireless Networks. In *IEEE Wireless Communications and Networking Conference*, 1999.
- [24] J. Walker. Unsafe at any key size: an analysis of WEP encapsulation. <http://grouper-ieee.org/groups/802/11/Documents/DocumentHolder/0-362.zip>.
- [25] T. G. Zimmerman. Personal Area Networks: Near-field intrabody communication. *IBM Systems Journal*, 35(3-4):609–617, 1996.