Exploring the Feasibility of Wireless Sensor Networks Within the MRI Environment

By: Justin Enders

Supervised by: Dr. Hassan Naser, Dr. Laura Curiel, and Dr. Samuel Pichardo

August 2011

A thesis submitted in partial fulfillment of the requirements of the M.Sc.Eng degree in

Electrical and Computing Engineering

Faculty of Engineering

Lakehead University

Thunder Bay, Ontario

Abstract

Wireless communications have long since been utilized to connect a multitude of devices quickly and easily to a common network. However, such methods have not yet been implemented inside a Magnetic Resonance Imaging (MRI) environment, leading to a difficulty of implementing simple devices easily and quickly inside the environment, especially in significant numbers. The feasibility of having a network of MRI-compatible sensors is of great interest for image-guided interventions such as MRI-guided Focused Ultrasound Surgery (MRgFUS). MRgFUS will benefit from the collection of multiple data such as temperature, backscattered acoustic waves, among others, for the monitoring of the procedure.

The MRI environment is generally considered to be noise filled, with a large magnitude pulse periodically occurring at about 127MHz (for a 3 Tesla MRI) while the MRI is in use. Additionally, in order to contain the large magnitude pulses, the MRI environment is contained within a Faraday cage. Wireless sensor networks (WSNs) are an emerging technology that allows for easily deployable, repairable, and robust networks that can be noise resistant, span large areas, and do not require a large number of wires to construct. The primary problem with implementing a WSN within the MRI environment is that the wireless signals are generated within the MRI environment, whereas the signal-processing endpoints will likely exist outside the MRI environment, on the other side of the Faraday cage. Fibre-optic communi-

cations are proposed as a solution for propagating the wireless signals from the nodes within the environment through the Faraday cage to their sink node. Another possible problem is that the WSN must not interfere with the operation of the MRI in any significant way.

We propose to solve the problems involved with implementing a Wireless Sensor Network inside a Magnetic Resonance Imaging device's environment to help facilitate MRI guided medical treatments. This thesis will cover background information such as the purpose of the research, which was to design a WSN that operates within and overcomes all problems associated with the MRI environment. Also covered is the background information of the key supporting technology of the proposed system, which are WSNs and Fibre-optic Communications. Background information in Magnetic Resonance Imaging is introduced to provide information about the environment that the proposed system is to operate in. A targeted application, Magnetic Resonance guided Focused Ultrasound Surgery (MRgFUS), is introduced to show a possible application for the proposed systems, as well as other possible applications.

The proposed system is described in detail, which is divided into three separate blocks. Explanation on how the blocks of the system operate, their purpose, and how they communicate with the other blocks to further the overall purpose of the proposed system is provided. The system proposes to utilize wireless communications to overcome the problem of having a large number of wires within the MRI environment and passing through the Fara-

day cage. This thesis also proposes to utilize fibre-optics to propagate the wireless information through the Faraday cage, which may attenuate the wireless signal. This thesis details a prototype developed as a proof of concept, designed to prove that the wireless communications do not interfere beyond tolerances with the MRI scanning processes.

Results from experimentation show that it is possible for wireless communications to exist within the MRI environment. The results also show that while there is interference to the MRI scanner from the wireless communications, that interference is relatively small (>3dB when compared to images with the network offline). Interference to the proposed system from the MRI scan process was also observed, causing errors in 5-10% of sensor readings gathered during the scan. The thesis also discusses the sources of the interference, the steps taken to minimize that interference, and the next steps to eliminating it. The discussion includes the next steps for the proposed thesis, as well as suggests devices to build and experiments to run in the future.

Acknowledgment

To Dr. Naser, Dr. Curiel, and Dr. Picardo for allowing me to work on this research.

To my girlfriend, brother, mother, grandfather and the rest of my family for their support.

To Thunder Bay Regional Hospital for use of equipment and the MRI.

For Megan and Chris.

Contents

List of Figures

Chapter 1

Introduction

1.1 Motivation

Magnetic Resonance Imaging (MRI) devices are now being used for more applications than just imaging. New MRI guided interventions are being tested. The capability of using multiple types of sensors to augment the imaging capabilities are of great interest. One of the most efficient means of introducing multiple sensors of different kinds is through a Wireless Sensor Network (WSN). MRI guided interventions are now being used for various treatments. In particular, MR-guided Focused Ultrasound Surgery (or MRgFUS) is being studied as a means of non-invasively excising malignant growths.

At the moment, it is difficult for MRgFUS to treat areas that may be located behind dense structures, like bones. By utilizing sensor arrays and sending a test ultrasonic pulse, it is possible to predict and compensate for the effect those structures have on the high intensity ultrasound. However, the nature of the MRI environment makes it difficult for a large number of sensor devices to be implemented. Should multiple sensors be utilized, MRgFUS could be used for a greater variety of interventions. In addition, the system that implements the sensors could be easily applied to other applications. For example, the system could be used for a simultaneous Electroencephalogram (EEG, a test commonly used to determine brainwave patterns) and MRI test. WSNs are an efficient way of quickly implementing many sensors of varying types.

Wirelessly implementing sensors within the environment solves the limitation of devices, but has its own problems within the environment. The MRI environment is contained by a Faraday cage, and this aspect of the environment may make it difficult to implement a pure wireless solution. In order to have a reliable system, a method of propagating signals through the Faraday cage should be found. Additionally, the proposed system must be reliable without interfering with the MRI operation. The purpose of this thesis is to propose a system that can communicate wirelessly within the Faraday cage, communicate with systems outside the environment, and do this with minimal impact to the operation to the MRI scanner.

1.2 Thesis Contributions

This thesis examines a methodology to permit the use of WSN inside the MRI environment. It details the findings of several experiments designed to show that the system can operate with the MRI without causing interference. The thesis contributions can be summarized as follows:

1. Proposes a system that both operates as intended, causes minimal interference with the MRI processes, and is compatible with all types of MRIs. This system has been specifically designed to create a WSN within the MRI environment, and yet is able to transmit sensor data and receive new information from outside the environment. This system will be able to be quickly and easily setup and dismantled as the need arises. The system will also be able to add new devices to the network automatically.

2. Develops a proof of concept system to demonstrate the viability of the system. The proof of concept system will be a scaled down version of the full system, and will be able to demonstrate all the concepts used in the larger network.

3. Details experiments whose results demonstrate that the proof of concept system has the potential of being utilized for the purpose of improving MRI based interventions. The results from these experiments will also be able to be used as a benchmark for future experiments done for future developments for the proposed system.

1.3 Outline

Chapter 2 outlines background information of the technologies that were involved in the research of this thesis, as well as their various applications. It includes knowledge on: WSNs, Fibre-optic Communications, MRIs, MRgFUS, as well as other minor topics. It also covers the identified problems that this thesis attempts to address.

Chapter 3 covers the proposed system, and divides the system into three separate blocks to ease the explanation of the entire system. Each block is explained fully, and how each block interacts with each other is also explained. The proposed system's basic concepts of operation is also explained in this chapter.

Initial experimentation and more involved experimentation is examined in Chapter 4. In particular, four separate experiments are covered, each of them focusing on a different aspect of impact on the MRI processes, or designed to demonstrate that the system will operate as intended and quantifying the interference to the MRI.

The communication block and Future work is detailed in Chapter 5. This section covers specific short term goals (including recommendations on how to complete the communication block) and experiments that should be run to build upon the work and concepts introduced in this thesis. Eventual long term goals are also covered in this chapter.

The thesis will conclude in Chapter 6. The thesis is summarized in this chapter.

Chapter 2

Background

2.1 Technologies

2.1.1 Wireless Sensor Networks (WSNs)

Sensors are devices that capture instances that occur in the real world and convert them to electrical signals. Following that, sensor networks are groups of sensors that are all linked together by sending their information to a common source. The purpose of doing this is to increase sensor coverage area, increase redundancy, and confidence. Sensor networks can be comprised of single or multiple types of sensors, as well as actuators. There are two critical components: sensor node and sink node. The sensor node contains the sensor package, which is the sensor(s) and associated circuitry to support the sensor. The sensor node also relays information to the sink node. The sink node receives information and either acts upon it, stores it, or relays it.

Both types of nodes contain a micro-controller for control, an antenna for communication, and a power supply. Sensors are used to gather information in the real world for electrical devices. WSNs are implemented when it is impractical to have wired sensors in place. There are many different types of sensor nodes. There are many different considerations when designing WSNs. Also, there are a variety of security considerations for WSNs that must be addressed. WSNs have a variety of applications, including environmental monitoring, security, factory monitoring, and traffic monitoring.

Wireless sensors are simply sensors that have a power source and a capability to wirelessly transmit the gathered data. These networks are typically implemented when it is impractical or impossible to provide communication lines or power lines to sensors. This usually occurs when the sensor is very small, and considered disposable. They are also used when the possibility of sensor death is high, as it is easier to replace a wireless sensor than it is to replace a wired sensor. There are many different types of sensors that can be implemented, varying between how many types of phenomena the sensor can measure, the quality and range of the information the sensor is able to detect, and the expected life of the sensor [1].

There are two general types of sensors: micro and macro. There are also many different sensor packages that fall between the two extremes. However for the purposes of an introduction to this topic the two types are sufficient. A macro sensor package is considered to be a "high-end" sensor. Macro sensor packages, in general, have more than one type of sensor, superior sensors to be able to obtain a greater quality and range of information, and have high quality batteries or are connected to a power supply. Macro sensors also tend to have larger processors, and as such are more capable of implementing routing protocols faster as well as other algorithms. Macro sensors are unlikely to die a natural death. As a result, they are used in locations where long term sensor coverage is required, and in locations where unnatural death, such as theft or destruction, is unlikely [2].

Micro sensors packages are considered to be the "low-end" sensors. Micro sensor packages usually only have one type of sensor, and the sensors used are generally low quality. As a result, the information micro sensors gather are of lower range and quality than macro sensors. The only advantages micro sensors have are that they generally cost far less than a macro sensor, and they can be far smaller, and are able to fit into areas where a macro sensor cannot. Micro sensors also tend to have less processing power and memory than macro sensors. Micro sensor networks make up for their inferior range and quality with numbers. Because they cost less, micro sensors are implemented where it is impractical or costly to have macro sensors; they are also implemented where unnatural sensor death is a high possibility, or where only short term or a short range of sensor coverage is required [2].

There are four different types of nodes inside a WSN. First, the sink node is a location where all of the sensor information from the network is routed to. There can be more than one sink node inside the network for a variety of purposes, such as to increase redundancy or to reduce the amount of power required to transmit to a sink node. Second, the full-service node contains a sensor package to gather sensor information, implement data fusion, receive other nodes information and transmit sensor information to the next node on the path to a sink node. Third, the limited-service node also has a sensor package to gather sensor information and has the ability transmit it to another full-service node or a sink node. The limited-service node cannot route other nodes information or implement data fusion. This is usually because a limited-service Node is more "micro" than a full-service node and as such does not have the processing power, memory, or battery power required to do all that is required by a full service node. A limited-service node could also be a full-service node that has entered a critical battery state and has entered a conservation mode. The final component is the database, where the sensor information is eventually stored and acted upon. For example, weather maps can be made or decisions can be made based of off the gathered sensor information [3].

There are many considerations to be addressed when designing a WSN. The positions of the sensors, as well as the sensor types must be considered. Network topology, fault tolerance, and routing protocol are other important considerations. Among other conditions are information fusion, network lifetime, security, and the most important considerations, the purpose of the network. Information fusion is the action of combining information from multiple sources into one piece of information that is superior to the first two pieces of information. Sensor node position is one of the last considerations to actually be addressed in the design of a WSN, but due to the ease of the process it will be covered first. Proper sensor node positions can extend the overall life of the sensor nodes by minimizing the amount of power needed for transmission, and by reducing the load on the nodes closer to the sink node. By carefully positioning sensor nodes it is possible to obtain maximum sensor coverage. This limits the effect of sensors obtaining redundant information and maintains network functionality in the event of sensor node death, creating a fail-safe or fault tolerant network [1].

Figure 2.1 shows an example of poor sensor node positioning. The centralized sensor node is under great load, as it is required to route the other sensor nodes data to the sink node as well as transmit its own data. This causes the central sensor node to be active far more than any of the other sensor nodes. Should the central sensor node's power be sourced from a battery, the central node would be much more likely to fail when compared to the other sensor nodes. When the central sensor nodes fails, the other sensor nodes would be unable to communicate with the sink node. This would be a total network failure, even with most of the sensor nodes and the sink node operational. In addition, many of the sensor nodes are placed relatively close to one another. This increases the possibility of redundant data being transmitted by multiple sensor nodes.

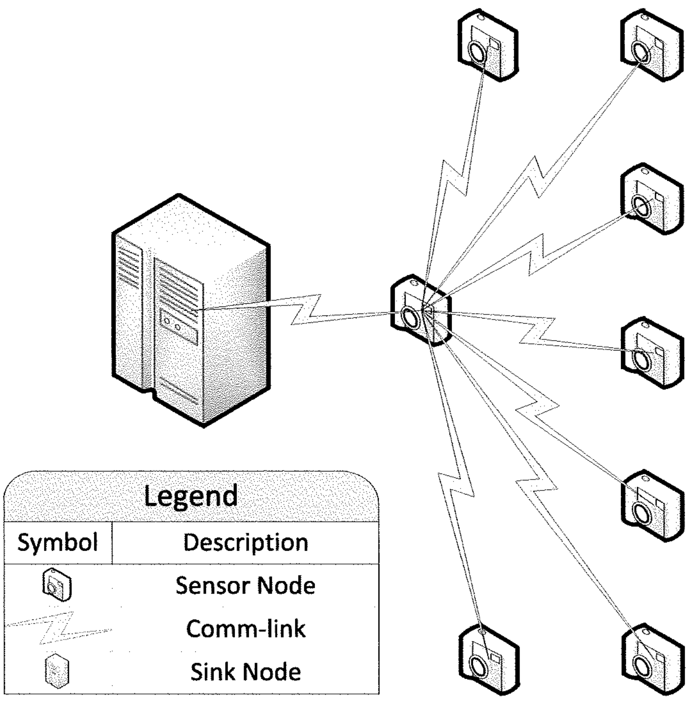

Figure 2.1: An example of faulty sensor node positioning.

Figure 2.2 shows an example of proper sensor node positioning. There is more than one path to the sink node, reducing the load of all the sensor nodes within transmission range of the sink node. In addition, each sensor node has more than one path to the sink node. This means that in case of a failure event, such as a sensor node failure or a communication link failure, the remaining sensor nodes should have another path to the sink node to communicate their data. This makes this network more fault-tolerant than the network shown in Figure 2.1. The sensor nodes are also positioned such that they cover a greater area without using more sensor nodes to do so.

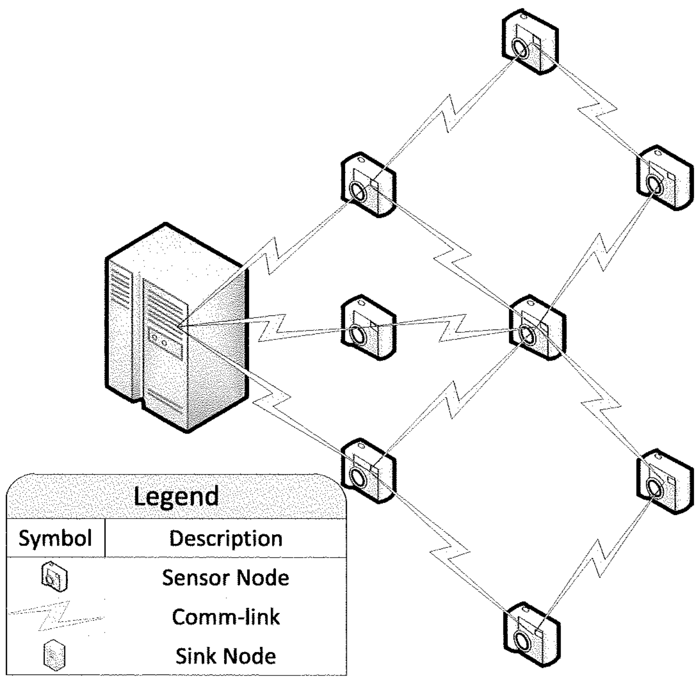

Figure 2.2: An example of proper sensor node positioning.

There are many different examples of network topologies. However, the more commonly used Cluster-Tree and Mesh topologies will be used for this introduction. Cluster-Tree topologies, shown in Figure 2.3 resemble a tree, in which the sink node would be the trunk, as all the information flows to it. Full-service nodes would be the branches as they can receive from the limited-service nodes and transmit to other full-service nodes and the sink node. The limited-service nodes would be the leaves, as they can only transmit to the sink node or full-service nodes. Cluster-Tree topologies usually have their paths generated for lowest power transmission. This is done to maximize the life of the network, as sensor nodes generally operate on battery power. Much like where a tree would have many leaves, Cluster-Tree topologies are most effective where the network contains many limited-service nodes. Given that Cluster Tree topologies can have a great number of limited service nodes it is a very economical topology to implement. The drawback being that the full-service nodes would be placed under greater load, and the network would be more fault intolerant, as the full-service "branches" would be vulnerable. If one of the full-service nodes fails, then every sensor node that was connected to the sink node through that path would be unable to transmit their information to the sink node [3].

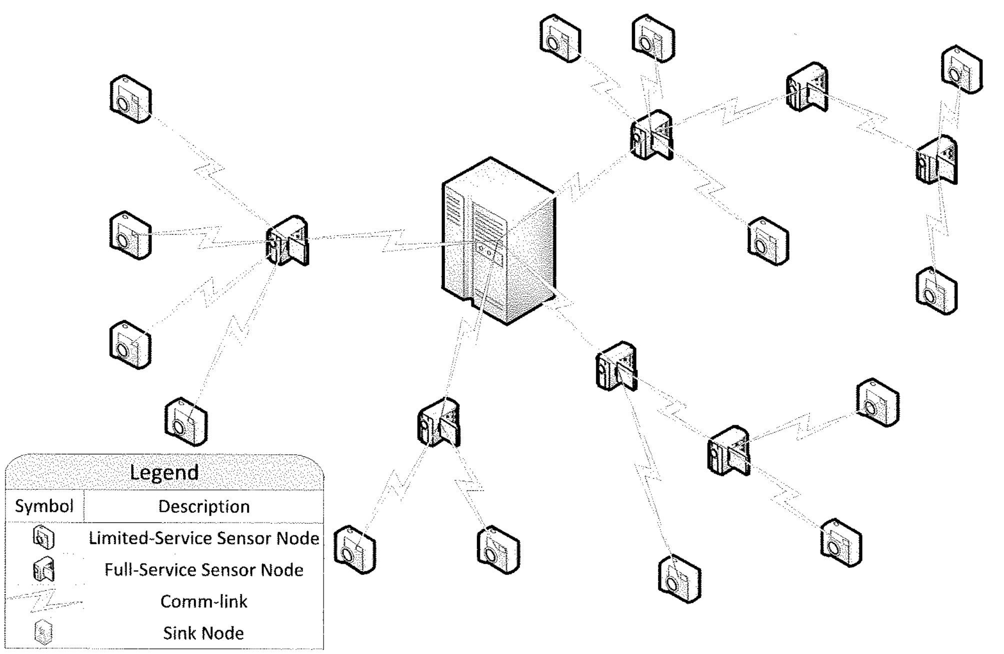

Figure 2.3: An example of a Cluster Tree Topology.

Mesh topology, shown in Figure 2.4, resembles a grid of nodes. Each sensor node is connected or is able to connect with each of its closest neighbours. Each of these nodes is able to find the best possible route to transmit its information across to the sink node. Most, if not all, of the nodes on a Mesh network should be full-service nodes. In Mesh topology, if a node dies its neighbours are likely able to find another route to the sink node. This makes this topology more fault tolerant. However, the drawback to this is that it would likely increase the load on those other nodes and make them more likely to run out of power [3].

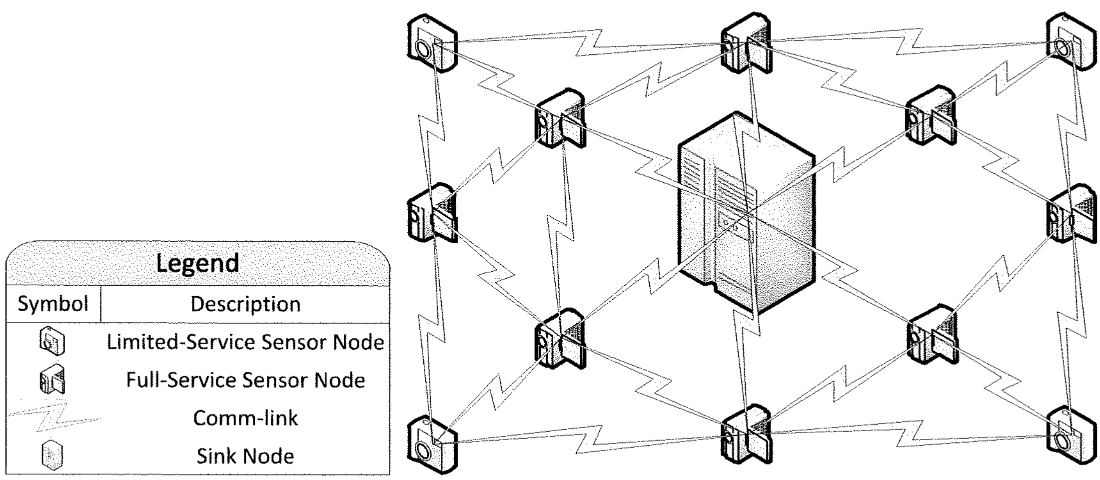

Figure 2.4: An example of Mesh Topology.

There are many different communication protocols that can be implemented in a WSN. In this introduction, two protocols will be compared. There are: ZigBee and LEACH. ZigBee [4], owned by the ZigBee Alliance, is a commercial protocol that was designed for nodes that must be able to last for long periods of time. The main feature that enables this is the ability for the nodes to go into a "sleep" mode, or to turn off all functions only activating on a received signal or on a periodic basis. ZigBee is a low cost, low power protocol that can operate in virtually any network topology, but one of its downsides is that it has a low data rate. The side effect of this is that nodes closer to the sink nodes would likely find themselves unable to enter "sleep" mode in larger networks due to the volume of data those nodes have to process and retransmit [3].

LEACH [2], created by Heinzelman et al., is also a low power, low cost protocol that is based off of a Cluster-Tree topology. Each full-service sensor node on the network has the ability to become a "Cluster Head". All other sensor nodes then find the closest Cluster Head and begin transmitting to it, forming "Clusters". The Cluster Head is responsible for transmitting to the Sink Node or to another Cluster Head to relay to the sink node. Super Clusters, which are essentially a larger cluster on top of smaller clusters, can also be implemented in sufficiently large networks. The main feature of the LEACH protocol is that it periodically reorganizes the topology of the network, so that the nodes that are Cluster Heads change. Given such, there is no typical position for a Cluster Head, as each full-service sensor node is capable of becoming a Cluster Head. This feature can extend the lifetime of the network to be eight times greater than other protocols, as the re-organization of the Cluster Heads can greatly extend the battery life of the sensor nodes. This is because it is unlikely that the same sensor node would repeatedly be chosen to be a Cluster Head [2].

Information fusion, introduced earlier, is used for three purposes: to reduce the amount of information that is transmitted on the network, to reduce the effect of redundancy, and to reduce spatial and temporal overlap of sensors. There are three basic types of information fusion: co-operation, redundancy, and complementary [5].

Co-operative information fusion is best implemented when two or more sensors transmit different information. At the receiving node, the information is combined to have a more complete data set, however the two pieces of information cannot be redundant, or identical, for this method to work. Co-operative information fusion is best used in situations where most nodes would transmit different data values of the same data type. For example, two temperature sensor nodes that sense two different temperatures that transmit to the same receiving node. Redundant information fusion is implemented when two or more sensors contain the same data. The information is fused so that the resulting data has better confidence or a better range, thus only one copy of the information is relayed to the sink node rather than two or more copies. Redundant information fusion is best implemented in situations where there are a significant number of nodes in the same area, and as such are likely to transmit identical information. Complementary information fusion is implemented when more than one type of sensor may exist in the same area. For example, visual, acoustic, thermal, wind, humidity sensors may exist in the same area. All the different types of information are fused to form a more complicated data set than the original information sets. Figure 2.5 is a diagram that visually shows how each of the information fusion

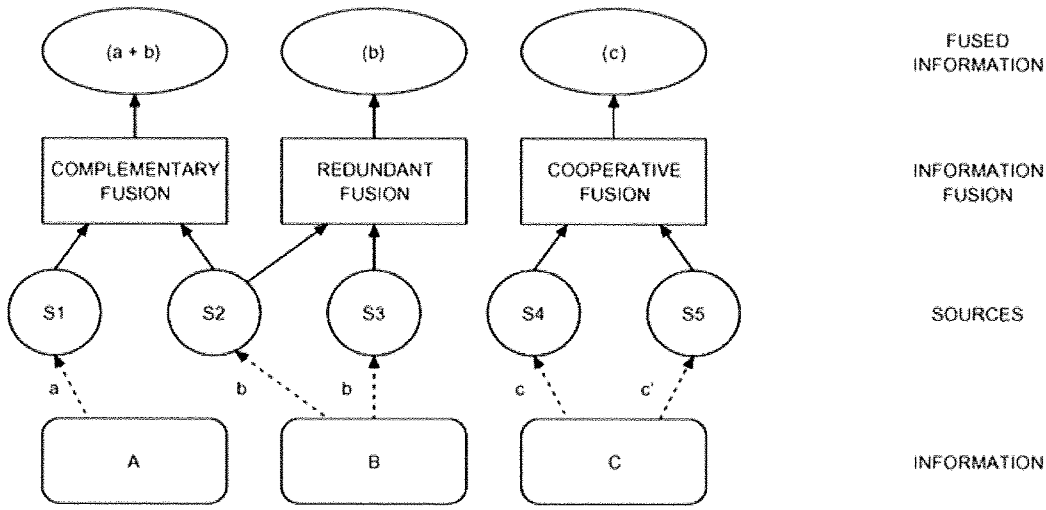

Figure 2.5: Basic Types of Information Fusion, from [5].

methods work [5].

Fault tolerance is a key consideration for WSNs. Wireless sensor nodes have a greater tendency to die when compared to wired sensors. This can be because of hardware failure, battery life, theft, destruction, etc. WSNs must be able to withstand node failures, maintaining a maximum amount of coverage as further nodes die. How a WSN reacts to sensor failures can determine how fault tolerant it is. As discussed before, topologies, protocols, placement, and hardware, among other factors, can influence fault tolerance [2]. After the sensor data has been gathered, it must then be acted upon. If the data is part of a long term study, then it would be compiled in a database for long term storage. The information can be arranged graphically for human viewing, showing a trend over time or space. WSNs can also facilitate automation, by having systems able to interpret the information that the sensor network provide and automatically take action upon it. This can be useful in situations where a person may not be easily or frequently able to access the data, or where the reaction time of a person may be insufficient [6].

WSNs have a variety of vulnerabilities which should be considered in their design. Nodes may be tampered with causing a malfunction, destruction of the node, or the node may be reconfigured to transmit faulty or incorrect data. Nodes may encounter message corruption, where the original message that the node transmits may get destroyed with a jamming signal or overwritten with a more powerful signal. Nodes may encounter a Denial of Service attack, where an unauthorized node may repeatedly be trying to transmit on the network. Even if the unauthorized node is constantly denied, the true purpose of the attack is to drain the battery of the node under attack, or to shut out real messages from nodes that are actually part of the network. A malicious node may be added to the network to either gain sensor information from real nodes on the network, or to pass along incorrect information. Sensor node messages are vulnerable to interception, and if they are decoded the attacker may be able to gain the nodes sensor information. Using this information may allow threats to pinpoint vulnerable nodes to target that allow them to do greater damage to the network [6].

To counteract these vulnerabilities, a WSN must be able to implement security methods. Transmission encryption is implemented to counter message interception and to prevent malicious nodes from entering the network. To counter malicious nodes, the network must be able to identify and isolate them from the network, and inform the administrators of their general location for removal so that the malicious nodes cannot launch further attacks. Nodes on a WSN should ideally be capable of self-healing should a successful attack on the network occur; meaning that the network should be capable of restoring service to as many nodes as possible following an attack. Should a node be stolen, the node should ideally be capable of self destruction. Hardware destruction is not necessary, but the node should wipe its memory if it detects that it has been removed from the network without authorization [6].

WSNs have a variety of purposes, such as environmental monitoring, traffic monitoring, factory condition monitoring, as well as security. In environmental monitoring, a variety of environmental conditions can be monitored. For example, temperature, humidity, wind, UV index can be measured at a variety of locations to create weather reports and maps. In factory applications, sensors can detect a variety of conditions, such as machine position, temperature, and relay the information to an operator, or even have systems act on the information independently. In security applications, businesses and security companies can employ cameras, as well as motion detectors, pressure sensors, and microphones to detect unauthorized entry [2].

WSNs have a variety of capabilities and purposes. The design of WSNs have a variety of challenges and considerations, such as the type of sensor being used, the purpose of the sensor, and how to keep the network secure. Other considerations for a WSN include information fusion, network protocol, network topography, and the overall purpose of the Sensor Network. WSNs become more viable and more common as technology reduces the size and cost of sensor packages.

2.1.2 Fibre-optic Communications

Fibreoptic communications is the concept of using light to transmit data by converting light from electricity, and then piping that light to its destination using glass fibres. This method of communication takes advantage of the speed of light, allowing for bit rates and capacity that is simply not available using copper lines. It is also distance intensive, allowing for data to travel kilometers without significant signal degradation. This makes fibreoptics ideal for both long distance and local data communications. There are a variety of methods that fibre-optic communications can implement, both active and passive.

In the "access" side, or the side where the end user can access the network, there are two basic types of optical networks: Passive and Active. These types of networks are also called "Access Networks", as this type of network is generally used to provide access for end users to the network. Passive optical networks (PON) have no electrical components, except at the endpoints of the network. This type of network is ideal for networks that are meant to last long periods of time, as generally the only type of failure occurs from

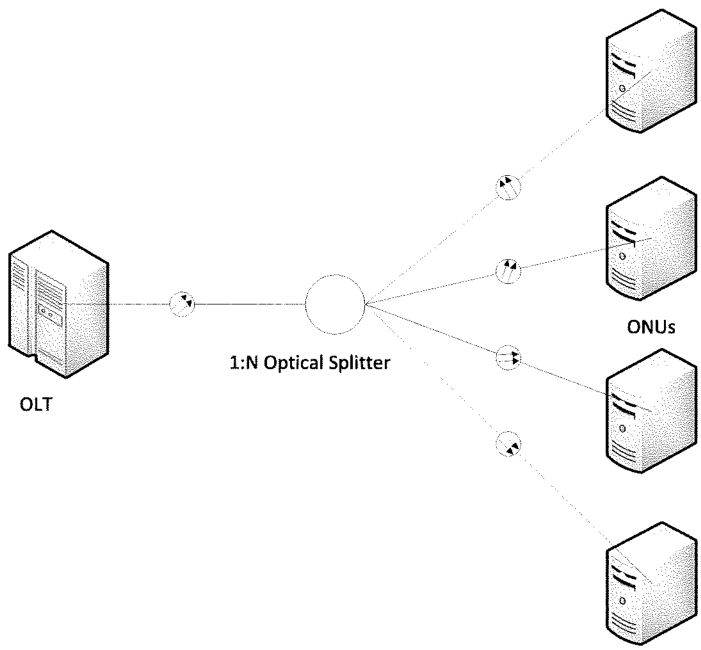

Figure 2.6: An Example of a Passive Optical Network

destruction. This type of network is also typically very power efficient, as only the endpoints of the network consume power. Typical examples of this type of network are ones implemented in the access networks, or as is becoming increasingly common, inside buildings. An example of a PON is shown in Figure 2.6 [10].

Active optical networks have electrical components between the endpoints of the networks, for signal boosting and regeneration. While this causes additional possible points of failure at the signal boosters and regenerators, it also greatly extends the distance and area that the network can cover. The additional points of failure are the added signal boosters and regenerators in the network, which can also suffer failures independent of the fibre or the

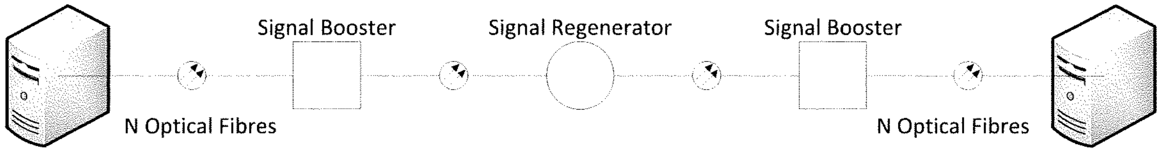

Figure 2.7: An Example of an Active Optical Network

end-points. Active optical networks also typically utilize more power than a PON, as the signal boosters and regenerators consume additional power. This type of network is ideal for long-haul networks and connections, such as connections between cities, countries, or continents. This type of network is also known as a long-haul network. An example of an active optical network is shown in Figure 2.7.

One of the main purposes of research for optical networks is to deliver fibre-optics as close to the end-user as possible. There are three general methods of accomplishing this: Fibre-to-the-Curb (FTTC), Fibre-to-the-Building (FTTB), and Fibre-to-the-Home (FTTH). These three basic methods, as well as some others, are denoted as Fibre-to-the-x (FTTx), where x is the location that the fibre is being delivered to. FTTC delivers the fibre to a city block, then continues on to the end-user using copper wire. This method is favorable when servicing a large number of end-users with a lower bandwidth and bitrate requirement. FTTB is much like FTTC, in that it is meant to serve a large number of end-users in a building such as an apartment building or a business complex. FTTH is when fibre-optics is installed in the end-users home or workplace, and is ideal when the end-user needs a high bitrate and bandwidth [7, 10].

FTTx is being heavily researched as it is a possibility of meeting the bandwidth demand of end-users; a demand that has been rapidly increasing for years. As there is a limit for how much capacity that copper lines can handle, this is creating a bottleneck at the last-mile. Without a method of either greatly expanding the amount of copper lines to deliver more bandwidth or developing a method of optimizing the amount of bandwidth available, copper will soon be unable to keep up with the demand. As the copper lines are replaced, there are more fibre-optic stations being installed to service the remaining copper lines, so that there is less copper line between an end-user and the fibre-optic stations. This is done as fibre-optics have much greater bandwidth available, and are more energy efficient. Eventually, it is possible that there may be a case where fibre-optics reach into the end-users home [7].

In PONs there are two primary types of equipment in place. The first equipment is the Optical Line Terminal (OLT), and the other is the Optical Network Unit (ONU). The ONU acts as a source of traffic demands on the network, and the OLT acts as the gateway to the internet. There are typically several ONUs and only one OLT. There can be as many as 128 or even 256+ ONUs when compared to the OLT, generally stated as 1:N, where N is the number of ONUs. The optical fibres are shared between the OLT and every ONU by an optical splitter. This splitter divides the downlink signal from the OLT to every ONU. However, when the ONUs attempt to communicate with the OLT, there may be traffic collisions as the ONUs are unable to detect when other ONUs are communicating. This will be discussed when the algorithms that are implemented to regulate the networks are discussed [7, 10].

The purpose of the OLT is to regulate the network by implementing rules for transmission for each ONU. The OLT acts as the gateway for the optical network, allowing transmissions from the ONUs to reach other ONUs, or to direct them out to the internet. The OLT also regulates received traffic, relaying the information to the correct ONU from other ONUs or the internet. What type of transmission rules are implemented is dependent on what type of multiplexing is used. There are three basic types of multiplexing, Time Division Multiplexing (TDM), Wave Division Multiplexing (WDM), and Space Division Multiplexing (SDM)[10].

Ethernet Passive Optical Network (EPON) is being considered for FTTx applications. Since it utilizes low cost ethernet technology and low cost optical components, making it a cost effective method. EPON has no active components in its communications path, making it a passive optical network. EPON can be implemented in almost all existing network topologies: such as star, ring, and tree topologies; however, the most common topology used is a 1:N point to multipoint tree topology. The fibre length can be as long as 20km, allowing EPON to be established city-wide. The communication method for EPON uses a common downlink wavelength and a shared uplink wavelength. The downlink is sent from the OLT to every ONU, and it is the ONUs responsibility to accept only the messages meant for it and reject all other messages. The uplink is shared between every ONU, and each message from an ONU to the OLT does not reach other ONUs. The normally used Carrier Sense Multiple Access (CSMA) collision detection method would not work well in this case, as the delay-bandwidth product of the network is large. This situation requires the OLT to schedule when the ONUs will transmit [8].

There are several ways of allowing multiple users to share a fibre-optic connection, usually through different methods of multiplexing. TDM involves separating time into blocks, and then assigning those blocks to different users to communicate with. Depending on the type of algorithm, these blocks may be of equal size or tailored to the needs of each individual ONU. There are many algorithms that have been developed to fairly distribute the size and amount of transmission blocks to each user. Examples of the algorithms that implement TDM are: COPS [9], PLSR [11], and AD-DBA [12].

Another method of sharing a fibre-optic connection is using WDM. This method is very similar to Frequency Division Multiplexing (FDM); as information is separated onto separate wavelengths or frequencies of light. The difference being that the wavelength used in WDM is a function of both frequency and velocity, while FDM uses frequency directly. This method has every user transmitting on a different wavelength of light, which allows each user to transmit simultaneously. This allows the end-users to have a large amount of bandwidth and bit rate without having to share with other users. The drawback to this method is that the OLT and the ONUs have to be far more complex when compared to other types of multiplexing methods. The OLT has to have receivers for every wavelength used, and the ONUs have to have either transmitters with a frequency unique to the network or a transmitter that is capable of changing its frequency[10].

There are also methods of combining WDM and TDM to have a hybrid communication method. This method, commonly known as Hybrid Division Multiplexing (HDM), applies a TDM algorithm to an optical network that has already implemented WDM. HDM applies the optimization gained from TDM to a network expanded by WDM. This type of multiplexing was designed for networks undergoing only a moderate expansion using WDM, in the case that the network administrators do not expect a rapid growth of network traffic very suddenly. HDM allows for a gradual expansion of network resources over a number of years, reducing the immediate cost and spreading the overall cost over a number of years; rather than a sudden expansion where the entire cost must be paid out immediately. There are several examples of algorithms that have been designed for this configuration: SUCCESS[13], SWDT and DWDT[14].

The final basic method of expanding fibre-optic capacity is to add more fibres to a fibre-optic trunk. This method is generally known as Space Division Multiplexing (SDM). This method does not alter the bitrate of the network; however it improves the network capacity by creating more links for network traffic to travel. This method has the advantage of being very inexpensive for shorter distances, and can be used in conjunction with TDM and WDM methods as well. With longer distances, such as part of a long haul network or an active optical network, this method becomes almost prohibitively expensive, as each new fibre added requires its own signal regenerator or amplifier [10].

A key advantage of fibre-optic networks when used in this application is that light is largely unaffected by electromagnetic (EM) radiation or other common sources of interference. This advantage makes utilizing fibre-optics the logical choice in areas where interference is a possibility. Instead of properly shielding a copper based communication line, it is possible to simply use fibre-optics enabled devices. Generally, the only sources of interference that fibre-optics receives are thermal based interference, losses due to bends or curves in the line, or when the line has been damaged. As fibre-optic cables are generally neutral in terms of ferromagnetivity and as a conductor for electricity, it is very well suited for communications with the MRI environment. The reason that fibre-optics are insensitive to external interference is that there are no electrical components in the fibre, and that the fibre is made of non-magnetic materials, so there is nothing that the MRI's magnetic field can affect [15].

A development caused by research into fibre-optic communications is fibre-optic sensors. These sensors share many characteristics of the more commonly seen types of sensors, such as being small, portable, and inexpensive. They are also immune to EM interference and they have large bandwidth capabilities; these sensors can go into environments with sources of large amounts of EM radiation, and are capable for having many sensors on the same fibre. However, many still prefer the more established electrical based sensors except in certain situations. Fibre-optic sensors are preferred when having an electrical current at the sensor may largely affect the reading of the sensor, such as a current sensor that is measuring a particularly small current. Fibre-optic gyroscopes are also considered superior due to the resolution that they are capable of and the fact that, unlike other gyroscopes, fibre-optic gyroscopes do not contain moving parts [16].

2.1.3 Magnetic Resonance Imaging (MRI)

The MRI scanner is a powerful, non-invasive imaging device that relies on powerful magnetic fields to operate. The medical community has come to rely on the MRI as a valuable diagnostic tool, as it is able to give a detailed view of organs in the human body without the need for opening or ionizing energy. The non-invasive and non-radioactive method that the MRI is capable of makes MRI technology able to be used safely on a greater variety of people in a greater amount of circumstances. [19]

MR imaging began with the theory that the angular momentum, or spin, of a hydrogen proton occurs at the "Larmor frequency" when inside a magnetic field. By altering the magnitude of the magnetic field at a linear rate, it was seen that the Larmor frequency also shifted at the same rate. This shows that the Larmor Frequency is proportional to the strength of the magnetic field. By spatially varying, or altering the magnitude of the magnetic field by distance, it is possible to spatially encode the locations of the hydrogen atoms. At this point, each hydrogen atom should have a slightly different Larmor Frequency. This allows the varying Larmor Frequencies to describe the location of the hydrogen atoms relative to the magnetic field. This discovery is what allows for Magnetic Resonance Imaging [17].

The Larmor Frequency can be found by the Larmor Equation $\omega_0 = \gamma_0 \times B_0$, where ω_0 is the Larmor Frequency (MHz), γ_0 is the gyromagnetic ratio, and B_0 is the strength of the magnetic field (in Tesla). A hydrogen proton has a Larmor Frequency of 63.9 MHz in a 1.5 Tesla magnetic field (or about 127 MHz in a 3 Tesla magnetic field). Generally, with no outside interference, slightly more than half of the total protons aligns parallel to the main magnetic field, or on the Z-axis. The other portion of the protons align in the opposite direction to the other protons. This generates the net magnetization, or longitudinal magnetization, denoted M_z. By introducing an electromagnetic (EM) wave to the system of the same frequency of the Larmor Frequency of the protons, it is possible to "excite" the spin of the protons and shift the magnetization from the z-axis towards the xy-plane, which generates the transversal magnetization M_{xy} [18].

When the magnetization is shifted by 90^o, the spin of the proton acts in the same manner as a generator, creating an alternating voltage at its Larmor Frequency. The receiver coils of the MRI can detect those signals and pass them to amplifiers and computers to be processed. The reason that the MR signal is weak enough to warrant amplification is that much of the signal is canceled out by the protons themselves. (The lesser portion of protons is aligned opposite of the larger portion of protons, thus canceling most of the signal generated.) The second type of coil in the MRI, called the gradient coil, spatially encodes the magnetic field as discussed earlier. This allows the computers to decode the received MR signals and generate the MR image with the magnitude and frequencies of the incoming signals [18].

After the EM wave that causes the spin of the proton stops, the MR signal gradually reduces due to two relaxation processes, spin-lattice (T1) and spin-spin (T2 and T2*). T1 relaxation is the tenancy of the protons spin to realign with the magnetic field, going from magnetization on the xy-plane to magnetization on the z-axis. This can take up to several seconds depending on the strength of the magnetic field and the nature of the magnetized target. As this happens, the MR signal reduces until it stops (when it has returned to equilibrium within the magnetic field). T2 relaxation refers to the phase of the spin between protons. The protons after being excited enter phase coherence, or spin with each other precisely. After the excitation stops, the spins loses coherence as the spins fall out of phase, and as this happens the MR signal deteriorates again. There are two main reasons that coherence is lost: by energy transfer where the spins give and lose energy to each other and generates the T2 time. The second T2 time, T2*, dictates the dephasing that occurs due to time based irregularities in the magnetic field. These irregularities are caused by the scanner itself, as well as any materials within the scan area. These two types of relaxations occur independently and simultaneously, where T1 takes seconds to complete, T2 happens faster, taking less than half a second. MR images then can be generated by the detected signal at a specific time point. The presence of the signal is influenced by how close the detection time is to the T1 and T2 relaxation times. This is known as T1 or T2 weighting of the images [18].

In order to decode the spatially encoded MR signal, Fourier analysis is employed. As the receiver coils cannot differentiate between the location of the different MR signals that it receives, its output appears to be the sum of all MR signals at any given point in time. To decode this information, it is possible to take the Fourier transform of the output of the receiver coils to obtain the Fourier coefficients of the signal, making it possible to reconstruct all the original MR signals. By only emitting a specific EM wave at a specific frequency while the field is varied in a single direction, a single "slice" in the target can be selected. However, in any given slice it is possible to exist two proton positions with the same Larmor Frequency, due to the gradient of the magnetic field being the same for the entire slice. The solution to spatially locate the different protons in the slice is to take into account the phase angle of all the MR signals as well. The frequency be along one axis, such as the x-axis(k_x), and the phase be along another axis, such as the y-axis(k_y), where k indicates spatial frequencies, then k_x and k_y form the axes of the coordinate system known as k space [19].

In order for a computer to perform the Fourier transform on the received signal from the receiver coils, that signal must be first converted to a digital signal. The received analog signal is passed through an analog-to-digital converter (ADC) to get the digital samples in time. The frequency at which the ADC must sample is at minimum twice the maximum received signal in order to accurately measure the signal. This frequency is known as the Nyquist frequency. A discrete Fourier transform is then applied to the digital signal to obtain the magnitudes of the k_x and k_y components of the signal. As the original signal was spatially encoded, by using this method it is possible to know the magnitudes of every MR signal that originally existed as a function of its spatial location. This information can then be graphically presented with the magnitude of every pixel (x, y) of the image corresponding to the magnitude of every signal (k_x, k_y) that existed in k-space [19].

Magnetic Resonance Imaging does have a few minor side effects. The most prominent of which is that of energy deposition. Energy deposition is when a material absorbs the RF energy that is sent during a MRI process, and converts some of that energy to heat. There are specific limits that have been set to the MRI to ensure patient safety during procedures, placing a strict limit on how much the temperatures in various parts of the body can be raised. However, it is unlikely that the MRI processes would ever endanger a patient. Even a patient whose ability to regulate temperature has been compromised would not be in danger. This is because in worst case scenarios the raise in temperature is less than a degree [19].

MRIs with high-field magnets are now being seen in clinical use, most at around 3 Tesla. The benefits of the higher strength magnetic field is primarily the improved Signal to Noise Ratio (SNR) that is generated. With this it is possible to reduce scan times, which would allow a greater patient throughput with the MRI, or improve the resolution of the resultant image, allowing for better, more accurate imaging. There are drawbacks to increasing the field strength. There would be more energy deposition in the target materials as the exciting EM signal would have to be applied for much longer. This may lead to greater required cool-down times for the patient. There is also the possibility of higher susceptibility to magnetic or partially magnetic materials. Should additional materials be utilized with a high-field MRI, those materials may contribute to a greater than normal image degradation [18].

2.2 Applications

2.2.1 MRI-guided Focused Ultrasound Surgery (MRgFUS)

Focused ultrasound surgery is a relatively new non-invasive method of treatment. Ultrasonic waves, which are mechanical waves (generally from 200 kHz to 4 MHz), can propagate through tissue and is usually used in Ultrasonic imaging or echography. Other applications have arisen from the side effects that are produced by the propagating wave. Some of the energy from the waves is lost in the form of heat. This heating can be used for different applications, from controlled hyperthermia to thermal ablation [32]. This ultrasound thermal surgery can also be combined with radiation and/or chemotherapy [20].

If living tissue is heated beyond a certain threshold of around $57 - 60\,^oC$ for more than a few seconds, then those cells will suffer irreversible damage, generally known as thermal ablations. However, this damage is not selective between healthy tissue and malignant tissue. It is therefore used as a tissue elimination technique compared to a surgical excision. Because of the non-selectivity and the risks this implies, this method of treatment is not feasible without a method that allows for precisely targeting of malignant tissues while minimizing damage to the healthy tissue. Early methods for ultrasound thermal ablation placed a probe in a target area, then heated or cooled the probe to damage the tissues. The problem with this method is that the probe did not accurately direct the heat, and the leads to the probe also conducted heat to non-target areas. This generally led to underheating, where target tissues were largely left intact, or overheating where non-target tissues were damaged [21].

Focused ultrasound was found to be a method that can very quickly and accurately bring a focal point within tissues to a temperature that can cause thermal ablations. Since this method causes little damage outside its focal point, and has little side effects from heat diffusion or circulation due to how quickly it can heat the focal point, as well as being a non-invasive treatment method, it is very attractive for clinical use. The treatment method can be lengthy, as the focal volume usually cannot cover larger volumes of tissues. In those situations multiple areas must be targeted sequentially with the focal point of the treatment. In order to minimize the effects of perfusion or circulation there must be a cool-down period between every application of ultrasound to allow the tissues and the skin to regulate their temperature; this can lead to long surgery times. To address this drawback, new ways of expanding and customizing the focal region are being studied [21].

Current evaluations of MRgFUS indicate that it is most effective on tissues that do not have any intervening bone or air. This is due to the intrinsic characteristics of ultrasound waves coming from reflection and attenuation of the wave when large differences in the speed of sound are found (such as between solids and liquids). This makes MRgFUS the treatment method for abdominal areas, and certain areas such as eyes. Other areas such as the heart, lungs, and brain are not recommended for this intervention (although using this treatment with the brain still continues to be explored). The reflection and attenuation of the ultrasound waves alter the focal volume of the treatment, causing the target volume to not be treated or causing untargeted areas to be ablated [20, 22].

A method of both expanding the focal region and gaining greater control of the ultrasonic energy during surgery is to use phased-array transducers. Normally, a transducer is only made up of a single element, and the only available methods to control the ultrasonic field is to either modulate power to the transducer or to physically move the transducer. In a phased-array transducer, there are hundreds or even thousands of elements. Using a phased-array transducer grants an additional method of controlling the acoustic field. By modulating the power and phase of each element of the array, it is possible to modulate the acoustic field. This allows for a more dynamic control of the ultrasonic field. Phased-array transducers are most commonly used to compensate for the distortions caused when an ultrasound hits a dense structure such as bones. Normally, Focused Ultrasound Surgery (FUS) would not be able to treat organs such as the liver or brain as they are behind bone structures which would distort the focal region of the treatment. However, as phased-array transducers are able to dynamically alter the focal region, it is possible to compensate for the distortions, provided that the information required to modulate the array correctly is available [21].

In order to achieve a non-invasive treatment method, there must be some type of monitoring to ensure that the treatment can be not only targeted but that it is progressing as planned. An example of failure that should be detected is a faulty element in an array during treatment that can cause the focal volume to change size and position, potentially covering a non-targeted area. An accurate and up to date image of the treatment area is required to be able to define the target regions correctly [21].

There are several candidates for guidance for ultrasound surgery: ultrasonic imaging and MRI imaging are among the most viable for clinical use. Ultrasonic imaging is the least desirable imaging method, as ultrasonic imaging is not sensitive enough to detect exact tumor regions and other important details, such as nerves. It is also unable to detect temperature changes, making it impossible for a focal region to be accurately followed during treatment. Without this ability, any error or change in the tissues between the imaging process and the treatment may cause undesirable results, such as an incomplete treatment or damage to healthy tissues [21].

MRIs are able to accurately detect the tumor regions and track minute temperature changes. These capabilities make MRIs able to perform a test immediately before treatment to define the focal volume correctly, allowing for real-time corrections to the focal regions in case of changes due to minor interferences [23]. The ability to accurately detect the boundary regions of the targeted area also allows for real-time corrections to the focal area to better cover the entire volume of the target region. MRI also allows for monitoring during the treatment, to ensure that the procedure is progressing as planned and allows for corrections as the treatment is occurring [21].

The ability to monitor the treatment as it is progressing is important for another reason other than for error correction. Living tissues usually have methods of dealing with an increased amount of heat to a given region, in order to prevent ablation. These cooling effects generally circulate the heat in the target regions, as well as the surrounding non-targeted regions that may be receiving a reduced thermal dose. A constant application of the thermal dose may cause thermal ablation effects in areas other than the target regions. Real-time monitoring helps ensure that only the target region suffers thermal ablation, and the ultrasonic treatments are only used in short bursts of less than a minute to minimize the risk of affecting regions outside the target zone [21].

Currently, there are many medical applications of MRgFUS. It is used for treating breast cancer [24], uterine fibrosis (benign growths in the uterus) [25], liver cancer [26], prostate cancer [27], as well as used in several other applications. It also has the potential to be used for brain surgery, as well as other areas of the body protected by bone structures. Further developments of the technology that supports MRgFUS may result in more accurate surgeries which would reduce the need for follow up action for the patient, as well as result in the ability to treat new areas of the body accurately and safely.

There are several devices that have been developed to perform MRgFUS. Among them include the Philips Sonalleve MR-HIFU [28] and the GE ExAblate series of systems designed for treating uterine fibrosis [29]. Other systems have been designed to treat other areas of the body, such as the Ablatherm [30] and the Sonablate for prostate cancer. Further systems are being designed and tested for various other applications, such as brain treatments, drug delivery, and other applications.

2.2.2 Multiple Scan Technologies

The ability to combine compatible scanning technologies is an interesting concept. The ability to perform two or more scans at once can be a time saving process that may reduce patient waiting times and allow doctors to come to a diagnosis faster. The ability to combine scan technologies may also provide additional insight that either scans on their own may not have provided. Additionally, the information can be registered and merged together, as the patient and organs are in the same position. A potential area for development would be to combine the MRI with an electrocardiogram (ECG) or an EEG. This may allow for a better understanding of the brain or other vital functions. As there are many wires normally associated with an ECG/EEG, it is normally difficult to implement both scans simultaneously.

2.2.3 Use of Wireless Sensor Networks for MRI Applications

There are several potential applications of a sensor network within the MRI environment. These include: sensor arrays for FUS and simultaneous imaging.

Sensor arrays for FUS allow for correction of defocusing due to the effects of solid tissue structures in the body with real-time information. However, there should be a sensor for every element in the phased-array transducer. This could potentially mean that there could be thousands of sensors employed in FUS. These sensors are currently in development at the moment. A potential problem with these sensors is that they are all currently wired. These wires have to pass either through the waveguide or a partition panel of the Faraday cage. The distance involved may attenuate the signal. The second complication of where to put the wires on the patient arises due to this as well [34].

With a WSN it is possible to remove most of the length of the wires that are attached to the sensors (At most a single meter as compared to the several meters required to reach the Faraday cage). This would solve the attenuation problem that a length of wire would create. It may also be unnecessary to build pre-amplifiers into the sensors, as they would not have to drive the signal out of the Faraday cage [33]. The sensors would presumably connect from their positions on the patient to the closest point on the MRI table, which would help alleviate the complication of where to position the sensor wires on the patient.

WSNs make it possible to wirelessly control compatible devices or elements that have been placed within the MRI environment. This allows for control of those devices without manually manipulating the devices or passing additional cables through the Faraday cage. They also make it possible to implement sensor arrays for other imaging methods. Arrays with a large number of elements can be controlled wirelessly rather than pass thousands of individual cables through the Faraday cage.

It is also possible to modify the MRI device itself using wireless sensors. New antenna coils are being developed with the ADC within the MR receiver coil itself, leading to the ability to treat the coil as a sensor package for a sensor node. This could lead to a increase of antenna coils used in MRIs, leading to an increase of accuracy in the images created by the MRI scan processes.

2.3 Problems

There are several inherent problems with implementing a wireless system within the MRI environment, due to intense magnetic fields as well as sensitive equipment. The following details the specific identified problems that this thesis attempts to resolve.

2.3.1 Interference With the MRI

As the wireless sensor system is meant to operate alongside the MRI at all times, it would not be viable for this system to cause significant interference in the resulting images the MRI generates. Thus, it must be shown that the system does not cause an unreasonable amount of interference. This is particularly important when the system may be used in conjunction with medical interventions, such as MRgFUS, as degradation in the resulting images may cause an incomplete treatment or cause the treatment to target healthy tissue. Interference may cause an entire scan sequence to be degraded; this would lead to a repeat of the scan (ideally after locating and removing the source of the interference), wasting both time and money. Additionally, artifacts could be generated in the image which could alter the interpretation of the images during FUS treatment, possibly causing safety issues.

Materials introduced to the MRI environment would cause interference to the images. The interference that can be caused by the materials can be observed in two parts of the obtained images: degradation of the magnitude and phase image. These interferences depend on the types and amounts of materials that is brought into the MRI environment. Magnitude image interference occurs when the materials cause a loss of the SNR in the magnitude images that the MRI generates. This is a large concern, as the quality of this image can determine how accurate the overall imaging process is, and affect the outcome of any diagnoses or procedures relying on this information. [18]

Phase interference can come from movements and inhomogeneities within the magnetic field, which can be aggravated by materials within the MRI environment. The phase images are representative of the magnetic flux traversing the scan area, and so any ferromagnetic materials causes significant amounts of interference. [18]

Should there be artifacts in the MRI image due to interference, then it may be possible to target healthy tissue with FUS, or to ignore malignant tissues. This can cause a safety concern, as such treatments and diagnoses require very accurate images to satisfy reliability concerns. Should a patient undergo a treatment with that inaccurate information, then it is possible that the treatment can fail due to incomplete targeting information or damaging unacceptable amounts of healthy tissue.

2.3.2 Signal Propagation Through the Faraday Cage

The Faraday cage is in place to attenuate the large amplitude pulses that the MRI generates when it performs scans; as well as attenuate any signal from external sources to keep them from causing noise in the images. This

Faraday cage may also attenuate the wireless signals that the system may generate, corrupting messages or completely blocking wireless communication. At the same time, passing wires for every sensor or device through the filter panel of the MRI is not practical. Therefore, the system must be able to propagate signals through the Faraday cage easily without causing additional interference in the MRI. It must be able to do this while keeping the wires that are being passed through the waveguide to a minimum.

2.3.3 High Network Density

One of the more interesting problems identified is the fact the WSN has a high nodal degree. Nodal degree is an average of how many other nodes would be in communication range of a single given node. The proposed network could potentially have a nodal degree of hundreds. Generally, most WSNs have an average nodal degree of around 4-6. The reason WSNs are generally designed with a significantly lower nodal degree is to take advantage of the fading channel, as well as minimize the energy used during transmission by communicating with a closer node. The fading channel effect states that on a given channel, particularly wireless channels, signal power decreases exponentially with distance. WSNs take advantage of this by having multiple nodes, that are not in communication range of each other, transmit on the same wireless channel. This allows a much higher data rate and node communication within the network. An example of a high density network is given in [35], which states in their testing of a dense network, 1000 nodes were placed in a 100 m x 100 m area. They were able to keep the nodal degree to an average of 4 or less by modifying the power each node used for transmission. In the MRI environment, it may be possible to have 1000 nodes in a 1 m x 1 m area, making it very unlikely that the nodes will be able to decrease nodal degree by lowering the amount of power used for transmissions.

The proposed network cannot take advantage of the fading channel effect in the same way, as every node within the network would be in communication range of each other node. Either only one node at a time would be able to transmit, or multiple wireless channels would be employed to allow more nodes to communicate at a time, or another method would have to be found to allow many nodes to simultaneously communicate. Only having one node at a time communicate would not be ideal, as it has been predicted that sensor events should occur at multiple areas at the same time. Conversely, having nodes communicate on multiple channels would cause a general increase of noise in the environment, causing a greater chance for interference with the MRI scanning process. Balancing these two factors is the key for overcoming this particular problem.

2.3.4 MRI Compatibility

MRI compatibility is a major concern for any research that involves introducing new materials to the MRI environment. As stated in 2.1.3, there is an intense magnetic field generated by the MRI. If there should be ferromagnetic materials within this field, magnetic attraction could cause those materials to be moved towards the MRI. This situation would be very problematic, as it is likely that the materials would not be able to easily be moved by normal means. It would take several man-hours and resources to remove the material, during which time the MRI would not be able to be used. Worse, should this occur with a patient in the MRI, this poses a significant safety risk as the material may injure the patient. Even if the amount of material is not significant enough to be attracted by the MRI, it still may cause inhomogeneities in the magnetic field causing image degradation. It is hoped that the system developed would be comprised of at least a Zone 2 device, if not Zone 1, as described in [36].

Chapter 3

Proposed System

The system that this thesis proposes to overcome the stated problems consists of the following major components: sensors (or sensor packages), sensor nodes, sink nodes, optical transmitters/receivers, and a computer. Optionally, actuators may be included to increase the functionality of the system. These components are divided into three separate blocks: the control block, the sensor block, and the communication block as shown in Figure 3.1. The control block consists of the controlling computer, its software, and hardware interfaces. The sensor block consists of the sensor packages and actuators that may be in place, and the sensor nodes that they are connected to. The communications block details how the sensor block and the control block communicates with each other. Figure 3.1 shows the simplified system as it would be used for MRgFUS.

In this example, there are sensor packages (hydrophones, with supporting

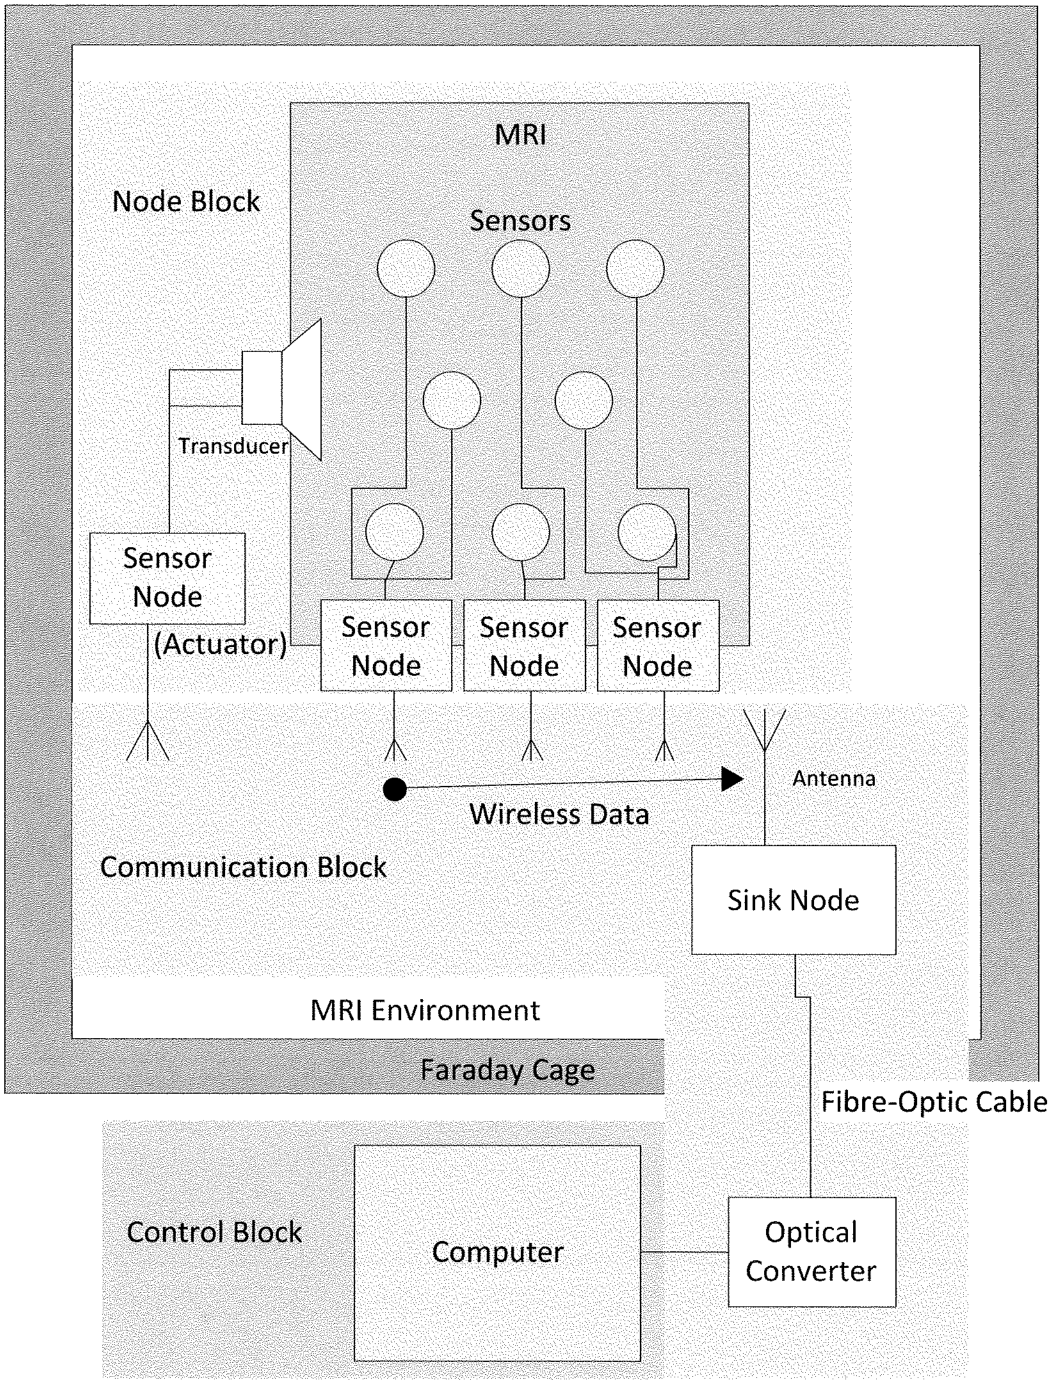

Figure 3.1: A block diagram of the proposed system.

amplifiers and ADCs) connected to sensor nodes. These sensors are close to the scan region of the MRI.

3.1 Control Block

The control block of the proposed system (shown in Figure 3.1) is where signal processing and system control decisions occur. The control block is primarily a computer with additional programming to handle the sensor outputs for signal processing, as well as decision making for the network. This block controls if the network is active or on standby, as well as other things such as controlling any actuators that may be connected to the network. The control block also acts as the sink node for the network, storing the information sent by the network and acting upon it.

The control block is responsible for creating and maintaining the wireless network. This involves registering each individual node that will be connected to the network, as well as what channel that node should communicate on should the network be established across multiple wireless channels. The control block will also be responsible for tracking failures across the network. There are two types of sensor node failures:

- Temporary failures, where interference from the MRI may cause the sensor data to be corrupted before it is transmitted by the sensor node (interference to the sensor package), or after it is transmitted (interference to the wireless network). These types of failures are instantaneous, and do not affect the long-term operation of the components.

- Permanent failures, where damage to the components has caused all data from that sensor package or sensor node to be corrupt or not transmitted at all. These types of failures require repairs to be done to the sensor node or package in question.

The control block must be able to identify and track these failures. The temporary failures would be compared to previous data and determined if it is a failure based off of an error threshold. This implies that the control block should be aware of what the general allowed data thresholds are in order to correctly identify the failures. Simply put, the system should be able to predict what the data should be within a given error threshold. The system should also be able to detect and manage permanent failures. It should identify the errors using the same method used with temporary failures or if the node simply stopped reporting. If the sensor node is continually sending erronous data, it should be identified as permanently failed. The control block should then ignore all further data from that sensor or node, and alert the user of the device failure.

The control block also has the user interface; the user interface outputs the information gathered by the network, as well as allows for the user to input changes or to review the information gathered. The user interface should be on the computer, and can be coded in a variety of methods tailored to the purpose of the network. The user interface should also be integrated with the other functions or user interfaces that may be utilized during its use. For example, if there is a transducer control user interface operating at the same time as the network interface, then an attempt should be made to integrate the two interfaces for ease of use.

Generally, the operating computer as well as any software that the network or user interface requires to operate is considered to be a part of the control block.

3.2 Node Block

The node block (shown in Figure 3.1) is where the nodes gather sensor information about the current environment. The nodes within this block can perform many tasks, such as being connected to sensors to gather environment or event information, being connected to actuators to be able to manipulate the environment or generate events, or a combination of the two. The nodes are responsible for transmitting and receiving data over the communication block to/from the control block.

There can be many different types of nodes on the same network, each node having different types of traffic patterns. Sensor nodes typically report their collected sensor data at regular intervals. They can also be programmed to report in data when the control block requests information. Different applications would require different types of updates. For example, should sensors be tracking a patients vitals, those nodes would have a regular traffic pattern. If there are sensors that are waiting for an event, then those nodes can wait in standby mode until the control block notifies the network that the event is about to occur.

Actuator Nodes would typically not output any traffic at all, but would receive traffic from the control block. Depending on the types of actuators deployed in the environment, the traffic frequency could range from intermittent to regular. The amount of data that the node receives would also be minimal, the payload consisting typically of one or two bytes per actuator on the node. Examples of actuators that may be used in this environment would be relays, motors, and surgical devices (should the system prove reliable).

There may also be nodes that are connected to both actuators and sensors. These nodes may be responsible for simple standalone devices that could also be added to the network. For example, an IV distributor may be utilized to provide medicine to a patient. Instead of interrupting the imaging process or the surgery to change or stop the dosage, the network can transmit the changes to the distributor. The distributor would also have to sense how much medicine is remaining, and can transmit that information to the control block. As the distributor usually stays with the patient, it would be beneficial if the distributors can come and go with the patient.

The node block may also be responsible for some simple signal processing. Signal processing is executed at the nodes in order to alleviate network traffic. If all signal processing were to occur at the control block, all the nodes would have to transmit raw sensor data, which has large data size. This amount of data that is transmitted over the network may tax the network's capacity, especially with a large amount of nodes on the network. Should there be more network traffic than what the network can handle, sensor data may be out of date when it reaches the control block or lost as new sensor events overwrite the old data as the node attempts to transmit. Even a basic amount of signal processing would decrease the amount of data that the nodes would be transmitting, allowing for more nodes to be added to the network.

For example, take a system utilizing the ZigBee protocol that has 1000 sensor nodes (N = 1000). Each sensor node has an analog sensor package that must transmit the received analog data within one minute (t = 60s) of time zero. In this example, we will assume that the system has implemented signal processing (Fourier transform) at the sensor nodes. The sensor nodes obtain 50 Fourier coefficients, each four bytes, resulting in a payload of 200 bytes (P = 200 bytes) per node. The time frame and payload represents what MRgFUS would need from the proposed system to reliably control the focal region of the ultrasonic field during the procedure. ZigBee has a maximum payload size of 104 bytes in any given frame, and each frame is up to 133 bytes long [4]. This leaves a overhead of 29 bytes per frame (H = 29) for network control, synchronization, addressing, and error checking. As the payload is 200 bytes, there must be 2 frames per node (F = 2).

$$BitRate_{min} = \frac{N \times 8\frac{bits}{byte} \times (P + F \times H)}{t}$$

$$BitRate_{min} = \frac{1000 \times 8\frac{bits}{byte} \times (200 + 2 \times 29)}{60}$$

$$BitRate_{min} = 34400\, bit/s\; or\; 34.4\, kbit/s$$

This bit rate of 34.4 kbit/s is entirely reasonable for any currently developed WSN protocol to be able to handle (ZigBee's maximum bit rate is 250 kbit/s [4]). This bit rate even leaves ample room for expansion, allowing for the ability to increase the number of sensor nodes, resolution of the obtained coefficient, or decrease the time that the nodes can transmit their information. The result is far different if no signal processing occurs at the sensor nodes.

For this example, assume the same number of sensor nodes, time frame, and network protocol. For the payload, assume that there must be 400 kilosamples of an analog signal with a resolution of two bytes. With the minimum frequency of an ultrasound being 200KHz in the best case scenario, the minimum sampling rate of the sensor package should be at least 400 kilosamples/s in order to reliably reconstruct the analog signal according to the Nyquist criterion. Assume that there is one second of received analog signal for each sensor for ease of calculation. This results in a payload of

800,000 bytes (P = 800000) per sensor node. Since the maximum data size per frame is 104 bytes, the payload must be transferred over 7693 frames (F = 7693) per node.

$$BitRate_{min} = \frac{N \times 8\frac{bits}{byte} \times (P + F \times H)}{t}$$

$$BitRate_{min} = \frac{1000 \times 8\frac{bits}{byte} \times (800000 + 7693 \times 29)}{60}$$

$$BitRate_{min} = 1.36 \times 10^{8}\, bit/s\, or\, 1.36 \times 10^{5}\, kbit/s\, or\, 136\, Mbit/s$$

The bit rate of 136 Mbit/s is a very high requirement for a wireless channel to fulfill. Adding to that is the fact that ultrasound can be as high as 4MHz, and the required bit rate increases correspondingly. This shows that signal processing is a great asset in keeping the amount of data that is to be transmitted to a minimum. This can be of great benefit if expanding the amount of nodes in the network, increasing the number of sensor packages per node, or decreasing the time frame in which the sensor nodes must transmit their data.

The node block generally consists of the sensors and actuators that would be in the MRI environment, as well as the microprocessors that control them.

3.3 Communications Block

The communications block (shown in Figure 3.1) consists of the wireless channel and the fibre-optic channel. The wireless channel that the nodes use is maintained by a sink node. This node is primarily responsible for establishing the wireless network, identifying all nodes on the network, and relaying information between the control block and the node block. The sink node should also have an optical i/o to be able to communicate effectively with the control block. On the other side of the Faraday cage, there should be a matching optical i/o, that converts the electrical signals back to their original format or to a separate format. Fibre-optics should be used in order to propagate the sensor data from the sensor block and commands from the control block across the Faraday cage. This is because if any power supply line or communication line were to be passed through the waveguide, that line would then act as an antenna, negating the signal attenuation effect of the Faraday cage. For example, a single sink node can usually communicate its data effectively using RS232 or USB. However, should there be more than one sink node in place to support the node block, converting the incoming serial data to a different format, such as USB, would be an effective method of multiplexing several sink nodes onto a single connection for the control block.

The communications block is created by utilizing a wireless communication protocol within the MRI environment and a fibre-optic relay through the Faraday cage. There are many existing wireless communication protocols, such as IEEE 802.11B/G/N (or more commonly known as Wi-Fi internet). Wi-Fi is not likely to be a candidate for use in this system, as it is geared for general purpose use and has many drawbacks for this situation, such as frame size and power usage. A better candidate for this situation would be a protocol designed for use for a WSN, such as LEACH or ZigBee. With comparatively smaller transmission frame sizes and lower power transmissions, these are better designed to handle situations where minimal interference is required. However, as stated in Chapters 2.1.1 and 2.3.3, these protocols are designed to take advantage of the fading wireless channel. When the system has a lower amount of nodes this is not particularly an issue, because the wireless channel is unlikely to be in high demand by the sensor nodes. Individually, or in small groups, sensor nodes do not have a large amount of data to transmit and so the wireless channel is mostly unused. However, when there are hundreds or even thousands of nodes, a method of regulating the wireless channel becomes required. Much like with fibre-optic networking, there are two general methods of regulation: Time Division Multiplexing (TDM) and Frequency Division Multiplexing (FDM).

TDM in this instance is slightly different from what is implemented in fibre-optics. The computer in the control block would have to create a scheduling scheme where the systems need for updating different types of information is weighted against how much of the channel is available to do so. This scheme would then have to be relayed to the sink node from the computer, which would then update all nodes for when they are able to transmit. The node would then transmit data using its assigned transmission block. The control block should be aware of what type of data is expected, and assign priority accordingly; as opposed to in fibre-optics where the OLT relies on the ONUs to prioritize their own traffic, and relies on updates from the ONUs to make decisions on how to distribute transmission blocks. The downside to TDM in this case is that it does not actually increase the bitrate of the network, it just optimizes the channel. If there would be more nodes added after a TDM channel is already at maximum capacity, FDM would be the next method to be used.

FDM in this case would be almost exactly the same as fibre-optics implements WDM. The control block would assign a frequency channel to each of the nodes to transmit on with each channel requiring its own sink node. This implies that the communication block would become proportionally more complicated and expensive as more channels are required by the system. As stated before, each sink node would require its own connection to the control block. There are several methods that can be implemented to achieve this. The sink nodes and control block can be optically networked in the same manner as utilized in PON, where a tree topology is implemented as in Figure 3.2, or each node can have a separate up/down connection as in Figure 3.3. The disadvantage of utilizing FDM in this situation is that utilizing more wireless frequencies would increase the interference to the MRI.

The final purpose of the communication block is to propagate data through

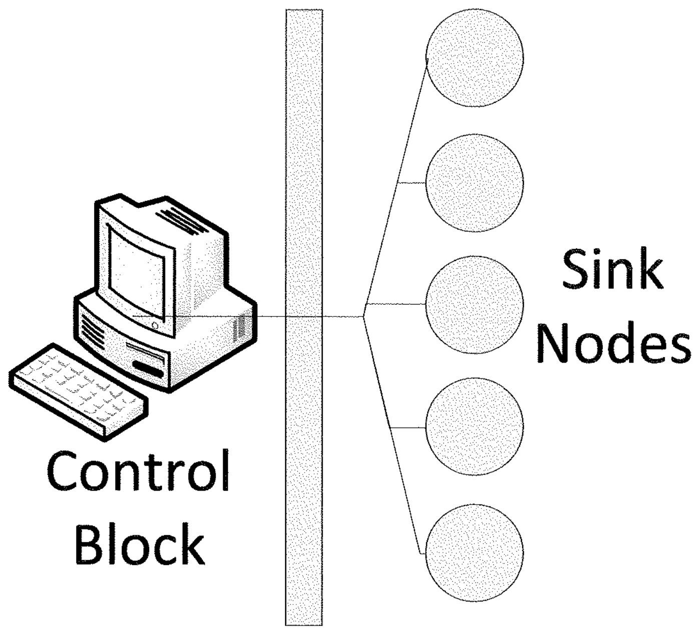

Figure 3.2: A Tree Topology implemented for each sink node

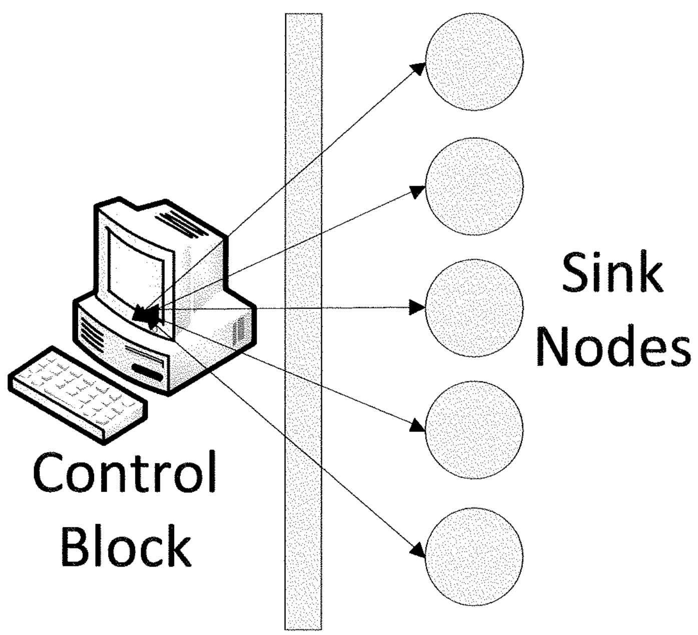

Figure 3.3: Separate fibre-optic connections for each sink node

the Faraday cage. This is done because the control block of the system would survive better and interfere less outside of the MRI environment. The Faraday cage would attenuate the wireless channel, and so to propagate the data fibre-optics are employed. The sink node(s) are equipped with fibre-optic transmitters, receivers, or transceivers. They transmit received data through the fibre-optics to a receiver that converts the received data to a form that the computer can interpret. There would be a receiver for every sink node. The receiver is part of a optical to electrical converter that multiplexes all data from all receivers onto one or more DB-9 or USB cables. These cables transmit the information at a higher bit rate than any one sink node to the computer. The computer then stores, displays, or acts upon the received information, and sends commands back to the sink node though the same manner. The sink node then relays those commands to the rest of the network through the wireless channel.

Utilizing fibre-optics to enhance WSN capabilities has been studied for various applications already. They have been studied to extend the range of WSNs to areas where wireless communications are limited, such as in mining conditions [37]. The use of fibre-optics extends the range of WSNs from hundreds of meters to hundreds of kilometers without dramatically increasing the number of nodes required to maintain network communications [38]. Research has also been done to minimize the amount of power required by the network as a whole [38, 39]. Utilizing WSNs in conjunction with fibre-optics with the applications for medical use in mind could potentially allow for more reliable medical imaging or allow for new and cutting edge interventions to be conceived.

In general, any component that is meant for data communications for the system or assists the system with that is considered to be part of the communications block.

Chapter 4

Experiments

4.1 Preliminary Explorations

A preliminary experiment was performed in order to determine if the wireless technology would function within the MRI environment. For the preliminary experiment, a sink node and a sensor node was implemented to establish wireless communications within the Faraday cage. The sensor node was programmed to regularly transmit a previously defined block of data to the sink node, which would send an acknowledgment to the sensor node. This was done as a sensor package was not yet connected to the sensor node.

4.1.1 Prototype

The ZigBee wireless communication protocol was used. ZigBee was chosen because it is representative of many WSN protocols and had relative ease of setup. Both the sensor node and the sink node were implemented using PICDEM Z motherboards, developed by Microchip Technology Inc. (shown in 4.3). The motherboards are preconstructed prototype boards that are designed to demonstrate the ZigBee protocol. They are constructed using a PIC18LF6450 low voltage micro-controller, a ZigBee transceiver board (a PICDEM Z 2.4GHz RF card, developed by Microchip Technology Inc.), an DB-9 connector for RS-232 communications, a temperature sensor, and had basic demonstration programs in C provided. [40] contains a schematic of the PICDEM Z motherboards. No additional sensor package was added to the prototype as of yet, as it was deemed unnecessary to test the wireless communications within the MRI environment.

4.1.2 Test Methodology

The sensor node and the sink node were introduced to the MRI environment. The sink node was placed away from the bore (the center of the magnetic field in the MRI scanner), where the sink nodes would normally be expected to be placed. The sensor node was placed on the examination bed approximately 200 cm away from the bore. A phantom was then placed on the bed and located in the bore for scanning. This setup is shown in Figure 4.2. The nodes were then activated and the sensor node communicated with the sink node. Network communications were not hampered by the MRI when it was in passive mode. The MRI was placed into scan mode to test the effect of the materials and network on the quality of the scan. The network communica-

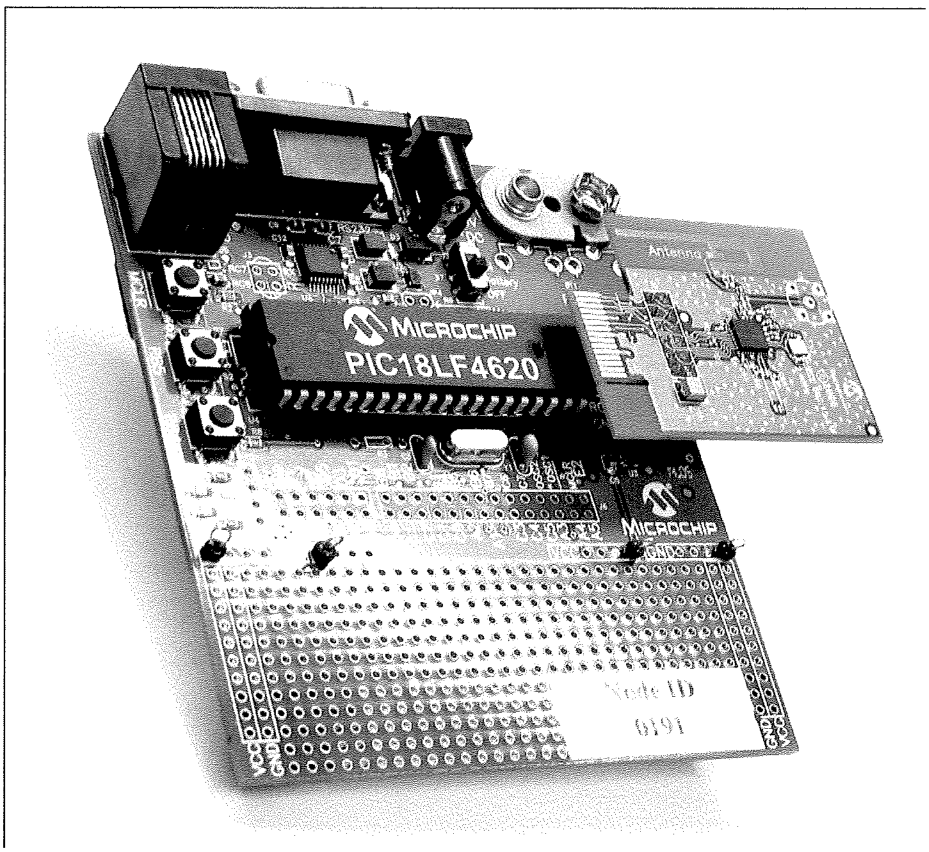

Figure 4.1: The PICDEM Z Motherboard, from [40]

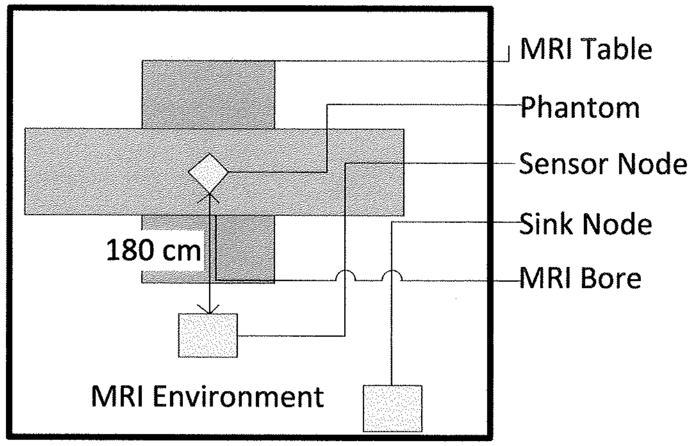

Figure 4.2: Diagram of Experiment Setup

tions were not greatly affected by the imaging process. However, the SNR of the MRI images were significantly impacted by the components on the sensor node and the sink node. This result was expected, as it was found that there were ferromagnetic materials used in the PICDEM Z motherboards.

4.1.3 Results

In general, initial experimental results were favorable to continue exploring this system. The wireless network was established and was able to maintain communication despite possible interference from the MRI device. The impact of the materials of the nodes onto the MRI scan, while not a favorable result, was expected under the circumstances. These results show that wireless communications have potential in being implemented in the MRI environment to assist in implementing MRI guided interventions. If the used motherboards were to be modified to be MRI compatible, the impact on the SNR and phase of the images would be significantly less than what was observed during this experiment.

4.2 Signal to Noise and Phase Impact Analysis

4.2.1 Setup

After the success of the preliminary experiment, work was done to make the PICDEM Z motherboards MRI compliant (shown in Figure 4.3). There were a variety of ferromagnetic materials that were on the boards that were not crucial to the function of the devices. The battery clips, the metal shield for the DB-9 connector, and the metal shields for the buttons were removed from the boards. The oscillator was also ferromagnetic, however, at the time an alternative was not available and it was not removed for the next set of experiments. By just reducing the materials, we expected a great reduction on the SNR loss and deformation of the phase image caused by the materials.

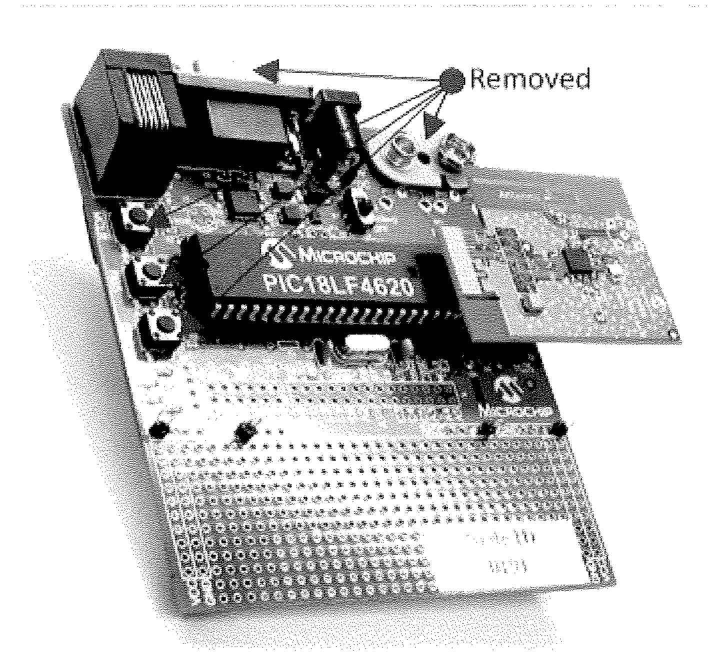

Figure 4.3: The PICDEM Z Motherboard with Removed Components, from [40]

The sensor node was also extensively modified. In order to get actual sensor data, eight Type-T thermocouples were added, along with a multiplexer (ADG507A, developed by Analog Devices), amplifier (AD595, developed by Analog Devices), and an analog to digital converter (ADC) (AD7813, developed by Analog Devices). The sensor nodes program was altered to control the multiplexer to poll each of the thermocouples, then store the output of the ADC to memory. After all the thermocouples outputs are stored, the node would then place the sensor information into a ZigBee frame and transmit to the sink node. This was programmed to happen approximately two times a second.

The preliminary experiment was then repeated with several modifications. The first modification was that as the battery clips were removed, a power supply for both nodes was employed. The 9V power cables were passed into the Faraday cage through the waveguide opening. The multiplexer and amplifier also required a separate power supply of 15 V, and those lines were also passed through the waveguide opening. Then the sensor node was placed in different positions increasingly closer to the phantom and the effects on the images at all these positions was studied; these positions are 1.9m, 1.4m, 1m, 0.6m, 0.4m, 0.2m, and 0m. The MRI would then capture three images at each position for the sensor node, both when the network was established and when unpowered. For reference, six images were taken with the phantom in the MRI and without the presence of any node in the chamber.

4.2.2 Results and Observations

The improved nodes, in terms of MRI compatibility, performed noticeably better than the preliminary experiment. A MATLAB program, (shown in Appendix A.1) was developed to accurately analyze the SNR of the images. The program takes the reference images, and averages them into a single image. The program then repeats this process separately for the images with inactive sensor and sink nodes (henceforth referred to as material images) and images with active sensor and sink nodes (henceforth referred to as network images) for given distances. The program creates a signal mask from comparing the magnitudes of the pixel points of the averaged reference image

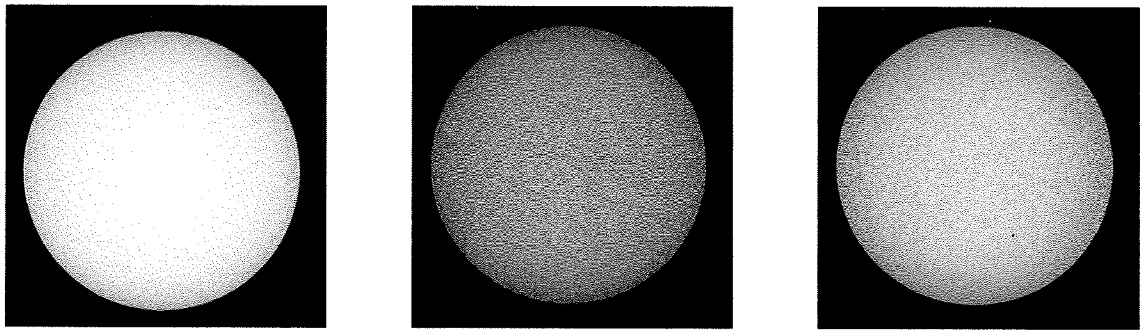

Figure 4.4: Examples of Reference (left), Material (center), and Network (right) Images

and compares it to a noise threshold. If the point's magnitude is above the noise threshold it is considered part of the signal mask. The program then calculates the SNR by using the calculation $SNR = 10log(\frac{Signal}{Noise})$, where signal is the average of all the points in the signal mask and noise is the average of all the points outside the signal mask. This program was executed for every position of the sensor node and the results are displayed in Figure 4.5.

The results show that even with ferromagnetic materials still on the nodes, the SNR of both the materials image and the network image are within 3dB of the reference image when the distance between the sensor node and the phantom was greater than 40 cm. However, as the sensor node was placed closer to the phantom, the SNR of both the materials and network images began to decay exponentially. Further, it is seen that the network images

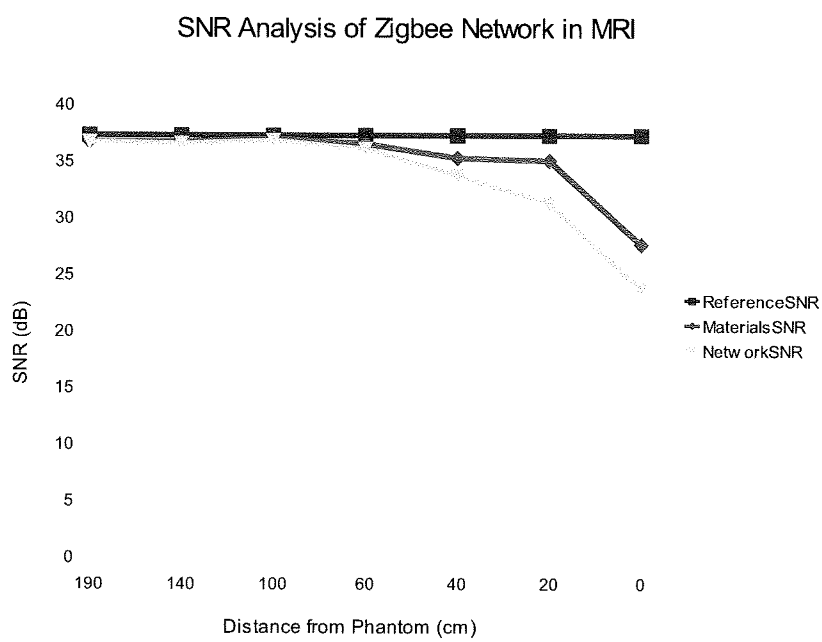

Figure 4.5: Experimentation Results for the First Run

begin to decay faster than the materials images, showing that the network communications do impact the image quality if the node is close enough to the target scan area. It can also be inferred at this point that even if the materials were perfectly MRI compatible and cause minimal impact to the MRI image quality, the wireless network communications would cause degradation in the quality of the MRI images.

In order to verify these results, the experiment was repeated, going in reverse. The sensor node was placed against the phantom and pulled back to the previously used distances. The results are shown in Figure 4.6. Again, the material and network image SNRs stay within 3dB of the reference images until about 50 cm away from the phantom. At this point the SNR of the material and network images begin decaying exponentially, with the network image quality decaying faster than the material image quality. These results are nearly identical to the first experiment's results, showing good repeatability of the experiment.

A second MATLAB program, shown in Appendix A.2, was also developed to analyze the impact of the inactive and active sensor and sink nodes on the phase of the images. Like the previous program, this program takes the reference images and averages them into a single image, and repeats for the material and network images. The program then creates a new image by subtracting the average material image from the reference image, and a second image by repeating this with the average network image. The program then requires user input to accurately determine a line of which has

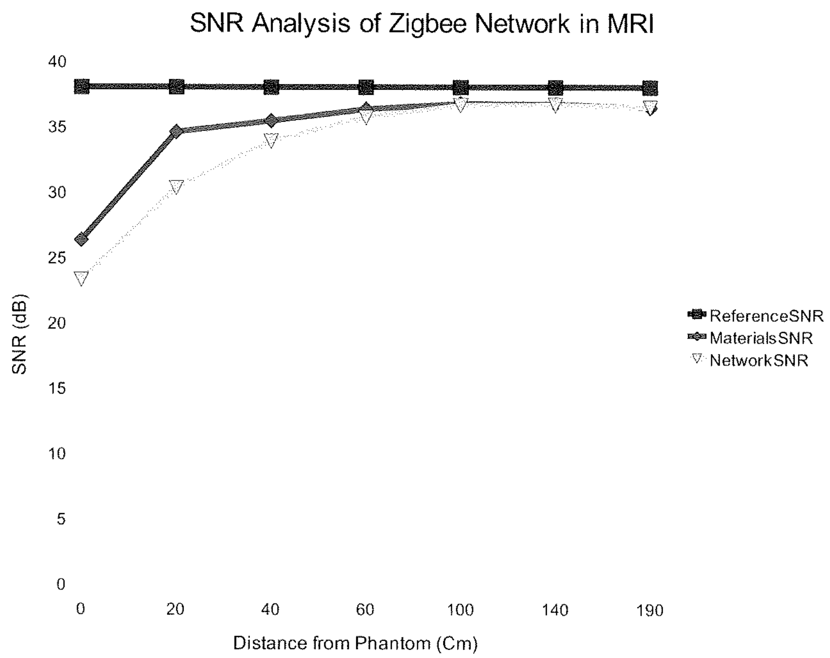

Figure 4.6: Experimentation Results for the Second Run

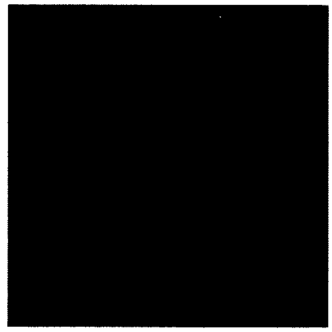

Figure 4.7: Reference Phase Image

the greatest variations and amount of variations on the images. Examples of these images and the user input required can be seen in Figures 4.8, 4.9, and 4.10. The line through each of the figures shows the user input to show the worst case of each of the images. The program then takes the values of all points on that line and maps the results on a chart. The program repeats this process for the second picture and overlays the result on the same chart as the first for comparison purposes. These results are shown in Figures 4.11, 4.12, 4.13, and 4.14.

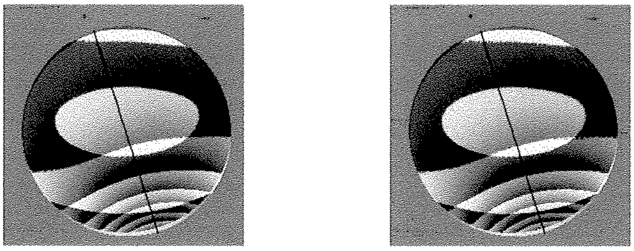

Figure 4.8: Examples of the Phase Images at 0 cm for material (top left) and network (top right).

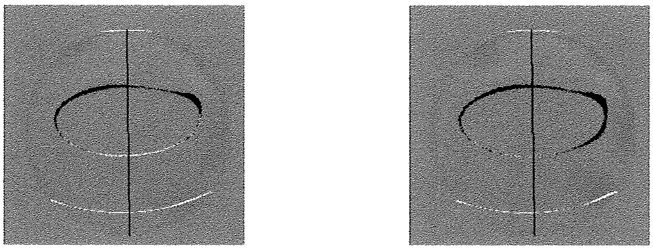

Figure 4.9: Examples of the Phase Images at 60 cm for material (top left) and network (top right).

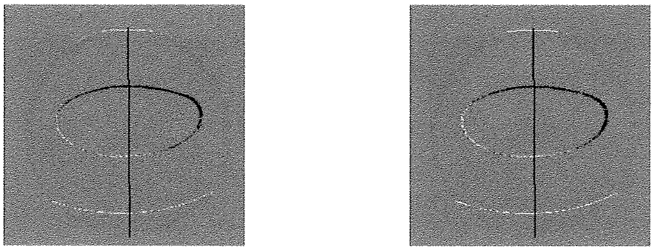

Figure 4.10: Examples of the Phase Images at 140 cm for material (top left) and network (top right).

These results show that there is a significant impact on the phase imaging of the MRI, particularly when the sensor node is close to the phantom. This indicates that there is ferromagnetic components in the sensor node, as the phase imaging is greatly susceptible to magnetic interference. Examination between the differences between the material analysis and the network analysis at each position shows that there is only extremely minor variations between the magnitudes of the two lines. The line representing the network image for most of the graphs fluctuates slightly more than its corresponding materials image.

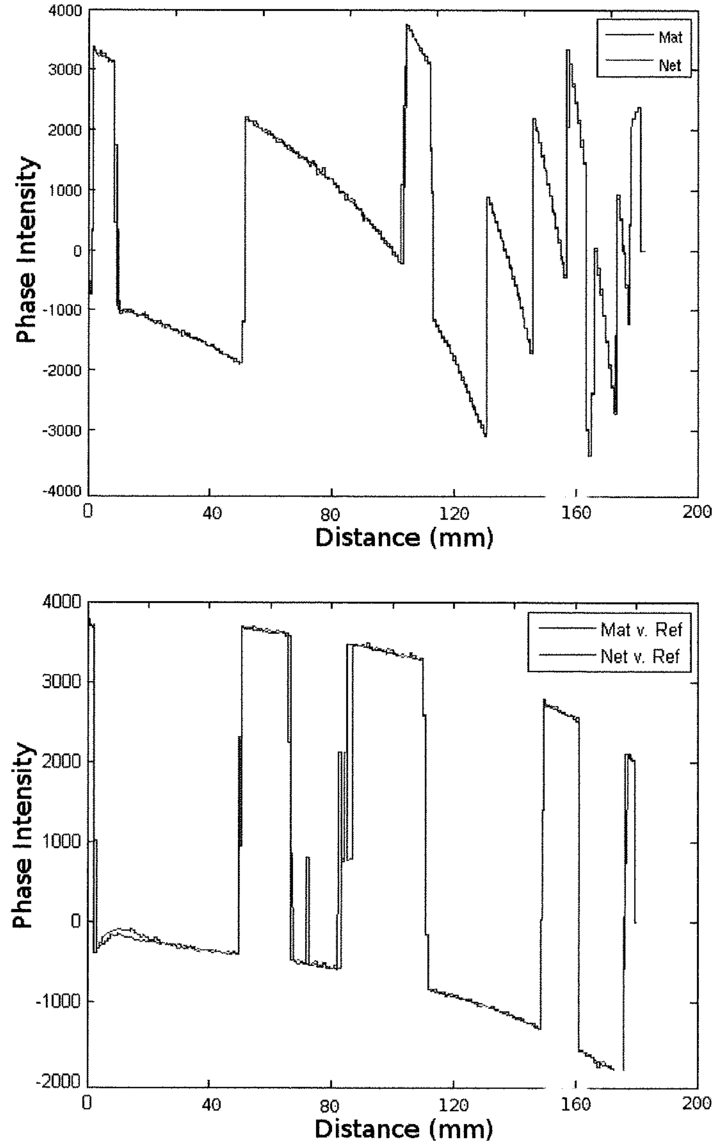

Figure 4.11: Phase Results at 0 cm (top) and 20 cm (bottom)

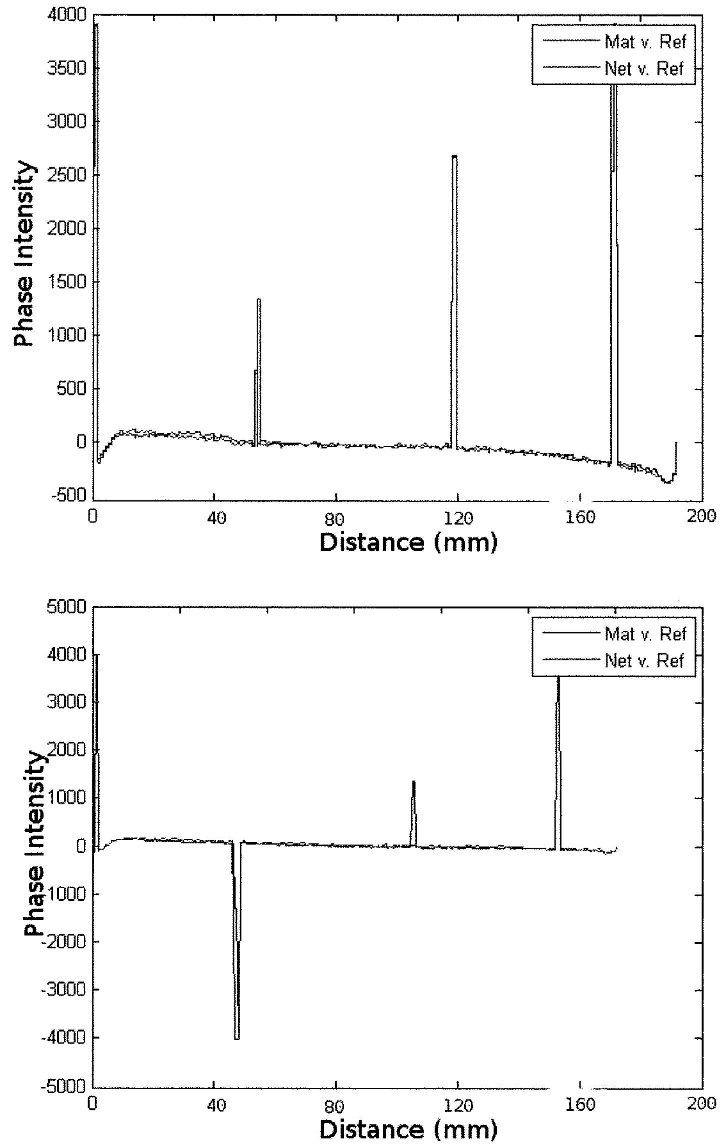

Figure 4.12: Phase Results at 40 cm (top) and 60 cm (bottom)

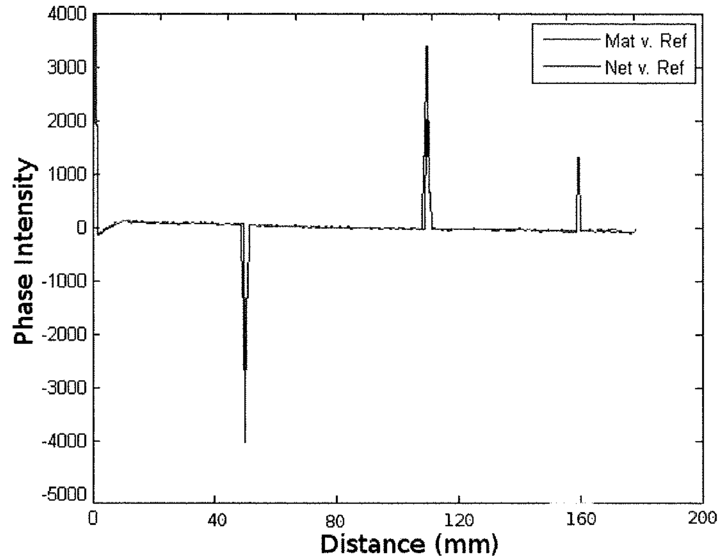

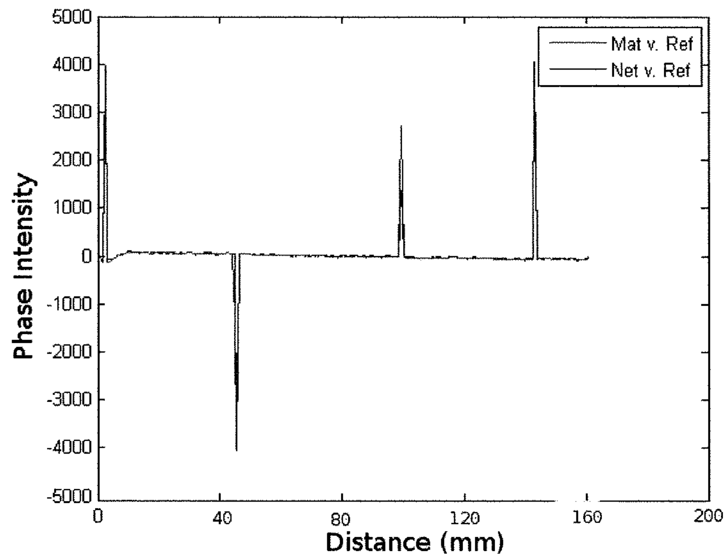

Figure 4.13: Phase Results at 100 cm (top) and 140 cm (bottom)

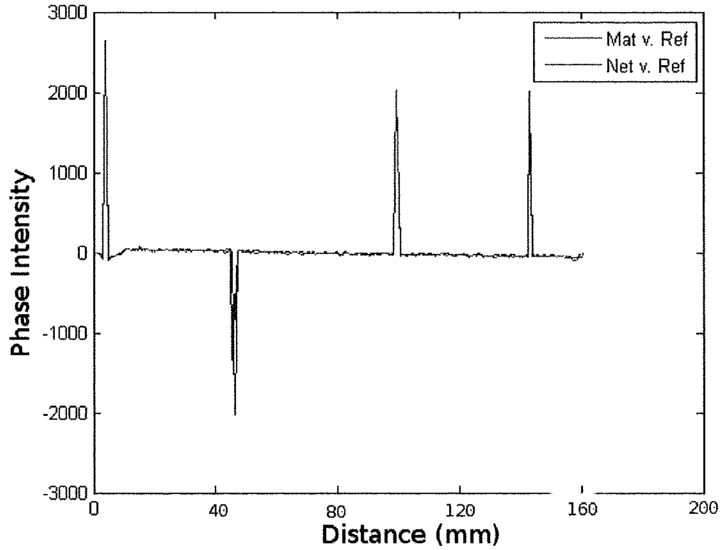

Figure 4.14: Phase Results at 190 cm

These results are understandable, as one would intuitively come to the conclusion that the network communications would not generate large amounts of additional magnetic flux that its materials would not already have. From these results, it can be concluded that the wireless network communications would not have a significant impact on the phase images of the MRI outside of the interference caused by its materials. If the materials would be MRI compatible, the interference caused would be diminished greatly.

4.3 Spectrum Analysis

4.3.1 Setup

After the results of the SNR and phase analysis experiments, there was a reason to examine the frequency spectrum of what the MRI's antennas were receiving. The reason for this was to verify conclusively that the wireless network system was interfering with the MRI. It was also suspected that the interference was at the signal detection (the coils) and not at the magnetic field. An spectrum analyzer (HP E4411B 9kHz - 1.5GHz, developed by Hewlett Packard) was connected to one of the output channels of the MRI amplifiers, and the frequency of the spectrum analyzer was set to 127MHz. This is the same frequency that the MRI pulses to align the atoms and molecules of the scan region. The span of the spectrum analyzer was set to 50MHz, in order to detect any possible sources of interference to the MRI.

Before the experiment was run, an analysis of wireless network communications was done outside of the MRI environment. The spectrum analyzer was set to the same settings as stated earlier, except with a span of 100MHz. The system was powered and the sensor node was communicating with the sink node. There were several minor signals noted on several frequencies close to what the MRI uses; mostly within 20MHz of 127MHz. This is why the span was set to 25MHz of either side of the MRI's primary frequency. Many of the possible signals at the various frequencies were near instantaneous, if frequent. To show this, the spectrum analyzer was set to max mode, to show the existence of the signals when compared to the signals normally generated by the MRI.

The sensor and sink nodes were then placed next to the phantom inside the MRI environment. A MRI scan was then executed and the results were captured. This process was repeated with wireless network communications enabled. This process was repeated for the previously used distances of 1.9 m, 1.4 m, 1 m, 0.6 m, 0.4 m, and 0.2 m.

4.3.2 Results and Observations

Before the nodes were put into the MRI environment, a reference was taken and examined with the spectrum analyzer, as seen in Figure 4.15. As was expected, there was one primary spike at 127MHz, with several minor spikes around that frequency. Figure 4.16 shows the results of the spectrum analysis at 0 cm. The materials image shows some interference on the closest side bands, showing that the materials themselves have some impact on the frequency spectrum around 127MHz. The network image shows that wireless network communications have a significant impact on the frequency spectrum that the MRI receives, so much so that the primary frequency of 127MHz is overtaken by the sidebands. It is also noticeable that when the network is engaged the power received at 127MHz is noticeably lower than the materials images and the reference image. There are also peaks at approximately 109MHz, 116MHz, 123MHz, 132MHz, 140MHz, and 148MHz. These peaks are consistent across all the images where the sensor node and the sink node are active and communicating.

As expected from the results from the previous experiments, the effect of both the materials and the network communications are reduced by distance, as shown in Figures 4.17, 4.18, 4.19, and 4.20.

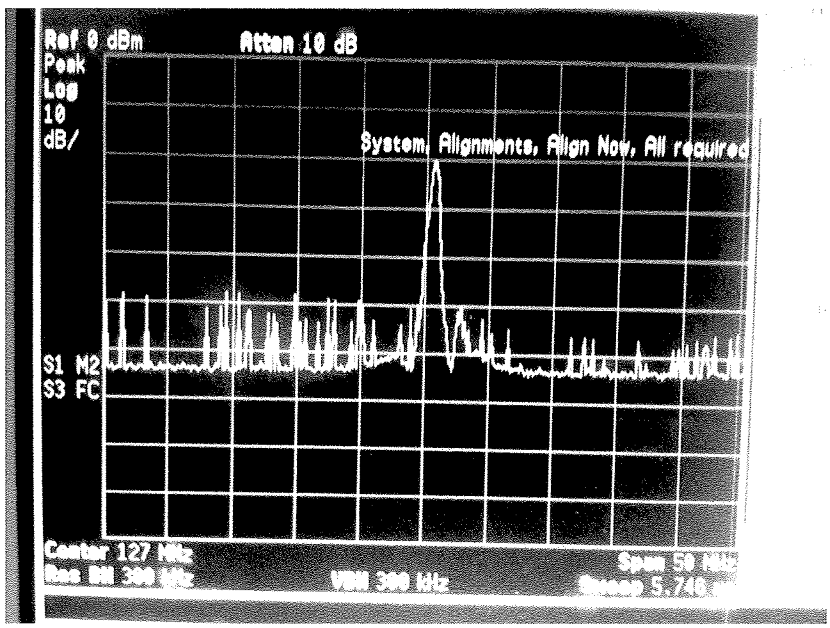

Figure 4.15: Reference Image with the Spectrum Analyser

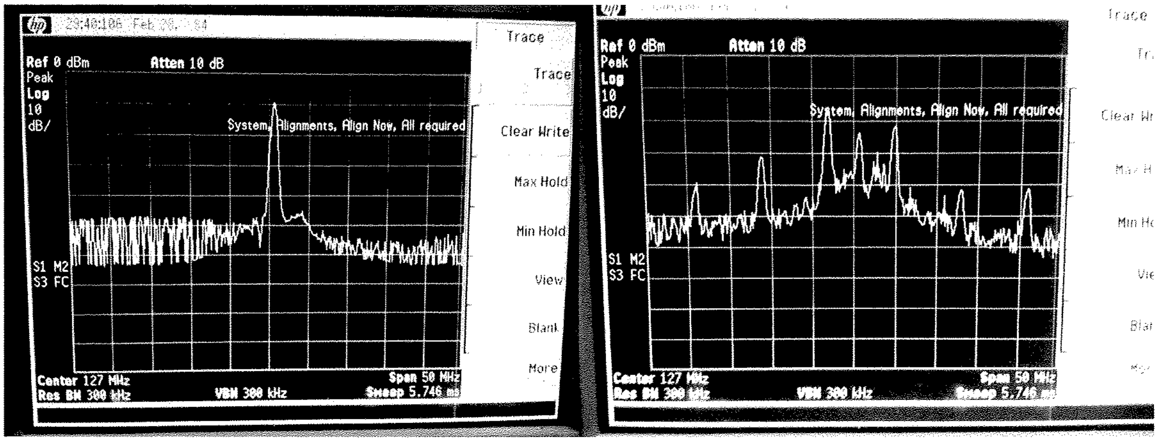

Figure 4.16: Spectrum Analysis of Material (left) and Network images at 0 cm (right)

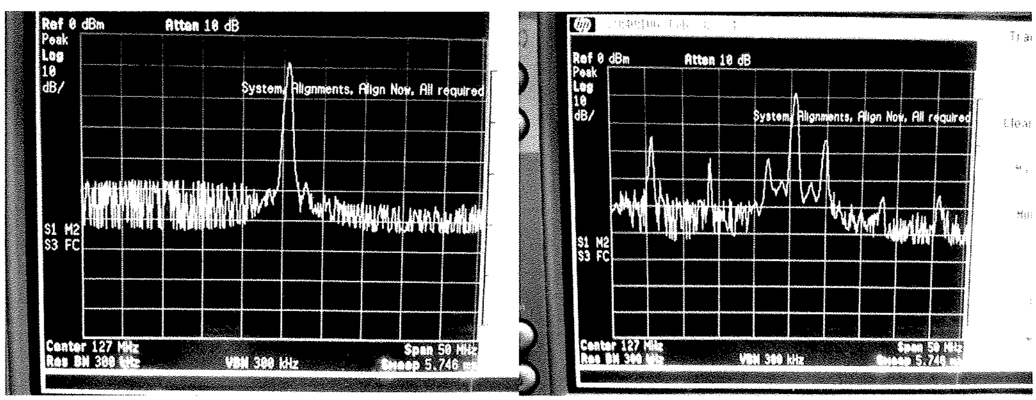

Figure 4.17: Spectrum Analysis of Material (left) and Network (right) images at 20 cm

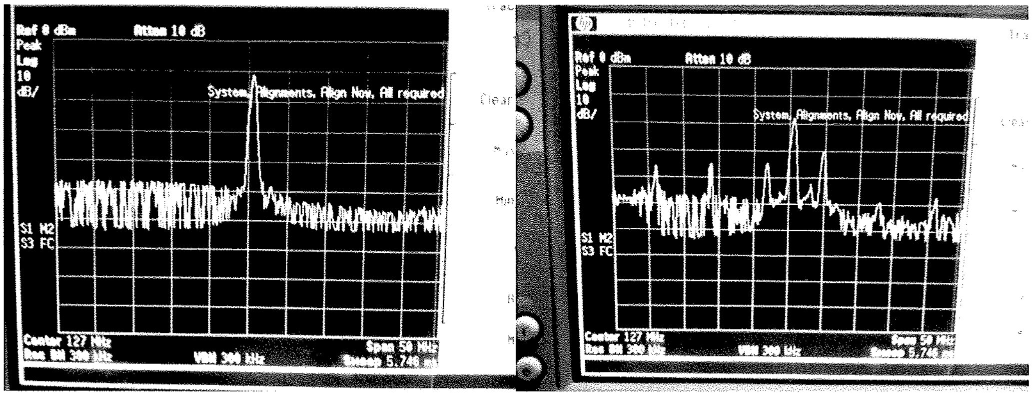

Figure 4.18: Spectrum Analysis of Material (left) and Network (right) images at 40 cm

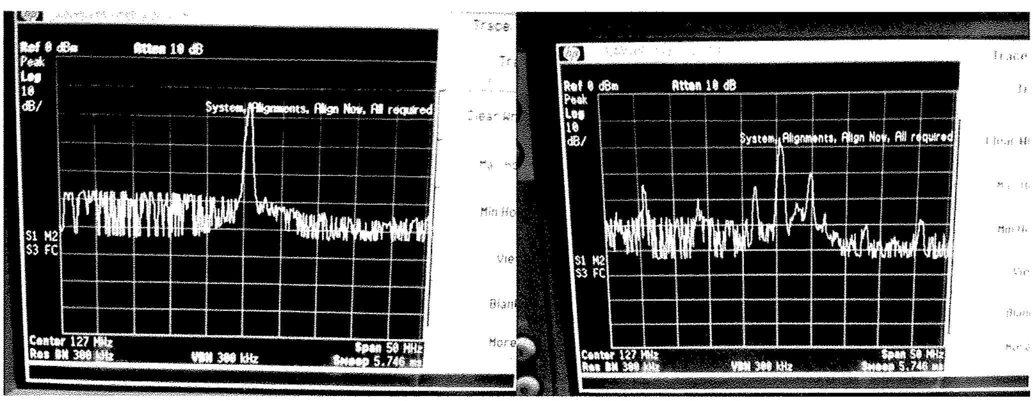

Figure 4.19: Spectrum Analysis of Material (left) and Network (right) images at 60 cm

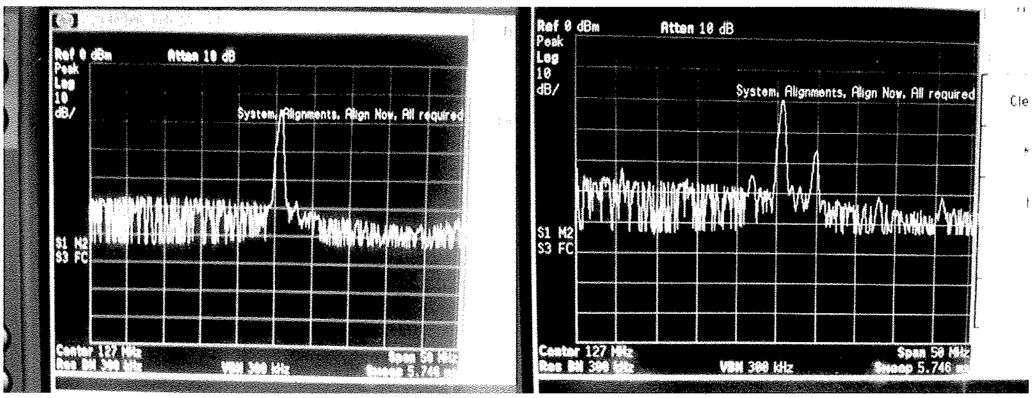

Figure 4.20: Spectrum Analysis of Material (left) and Network (right) images at 100 cm

Intuitively, one would not expect the 2.4GHz network communications to be interfering with the 127MHz signals generated by the MRI scan processes. To explain this, it is theorized that the MRI signal pre-amplifiers are being saturated by the network communications, and thus is the source of the SNR degradation.

4.4 Exploring the Effects of a Non-Ferromagnetic Oscillator

4.4.1 Setup

When the experiment performed in Section 4.2 was performed, the oscillator was the last ferromagnetic component on the motherboards. As MEMS oscillators are non-ferromagnetic, due to the fact that they are made of silicon, they were used to replace the current oscillators. It was thought that the removal of this ferromagnetic component would improve the phase images of the MRI. After the components were shown to have been compatible with the current motherboards, the experiment in Section 4.2 was repeated.

At the same time, a user interface (shown in Appendix A.3) was developed to visually inspect the output of the thermocouples. As there should be no heating effects occurring due to any of the experimental materials, any sudden spike in the readings would likely have been caused by interference to the sensor circuit from the MRI. Should these spikes exist, it would indicate that there is another possible source of interference that should later be accounted for. In order to connect the interface to the network, a DB-9 cable was passed through the waveguide of the Faraday cage and connected to the sink node. The sink node was programmed to output the received sensor data that the sensor node was transmitting wirelessly to it using RS-232. An error threshold of decimal 40 was applied to the results, which was equivalent to about 50 degrees Celsius. As the MRI environment is kept at 16-18 degrees Celsius, it was thought to be a reasonable threshold.

4.4.2 Results and Observations

As with the previous experiment, the SNR analysis program in Appendix A.1 was utilized to examine the impact of the non-ferromagnetic oscillator on the SNR of the MRI images. these results can be observed in Figure 4.21.

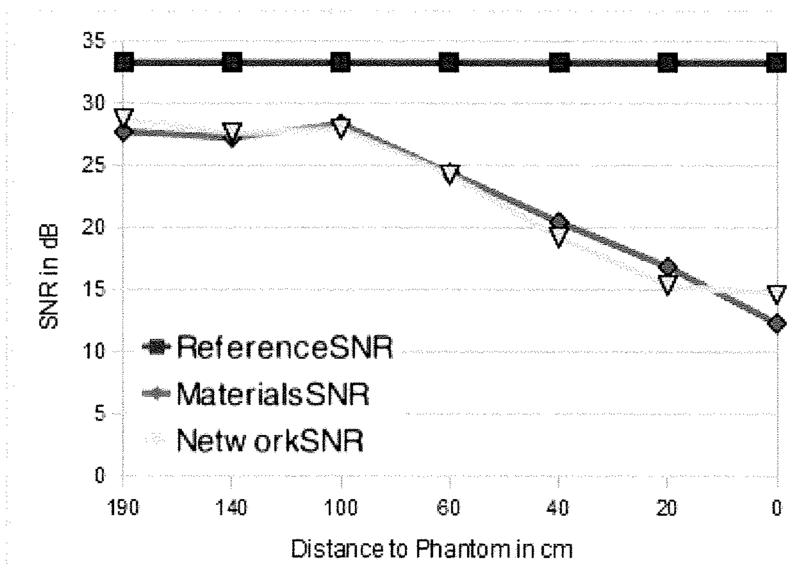

Figure 4.21: SNR Analysis Results

As these results show, there is a significant loss in SNR when compared to the results in Section 4.2. This is contrary to the expected results that the SNR would either improve or remain the same due to the removal of the ferromagnetic materials. It is likely that the inclusion of the DB-9 serial cable for testing the interface is the cause of the dramatic loss of SNR. Otherwise, the only other explanation, although very unlikely, could be that the MEMS oscillator is not MRI compatible as previously thought.

To analyze the phase images, the phase analysis program in Appendix A.2 was again employed. These results can be observed in Figures 4.22, 4.23, 4.24, 4.25.

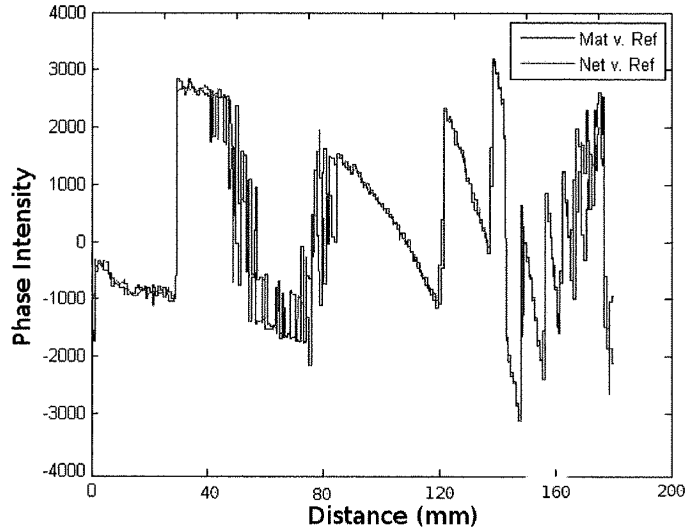

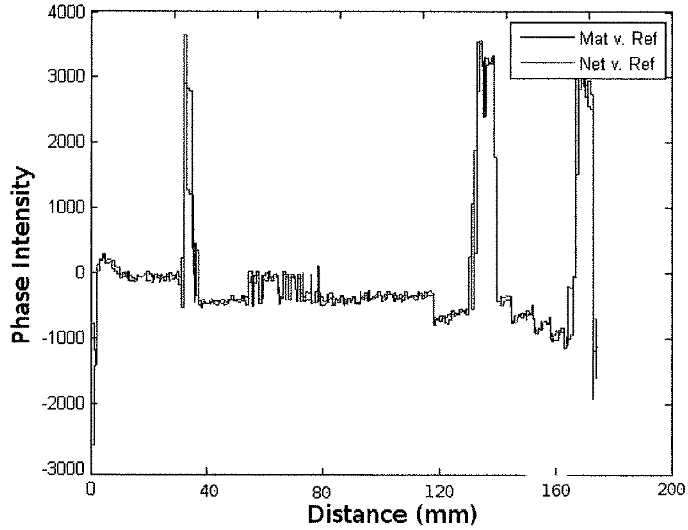

Figure 4.22: Phase Results at 0 cm (top) and 20 cm (bottom)

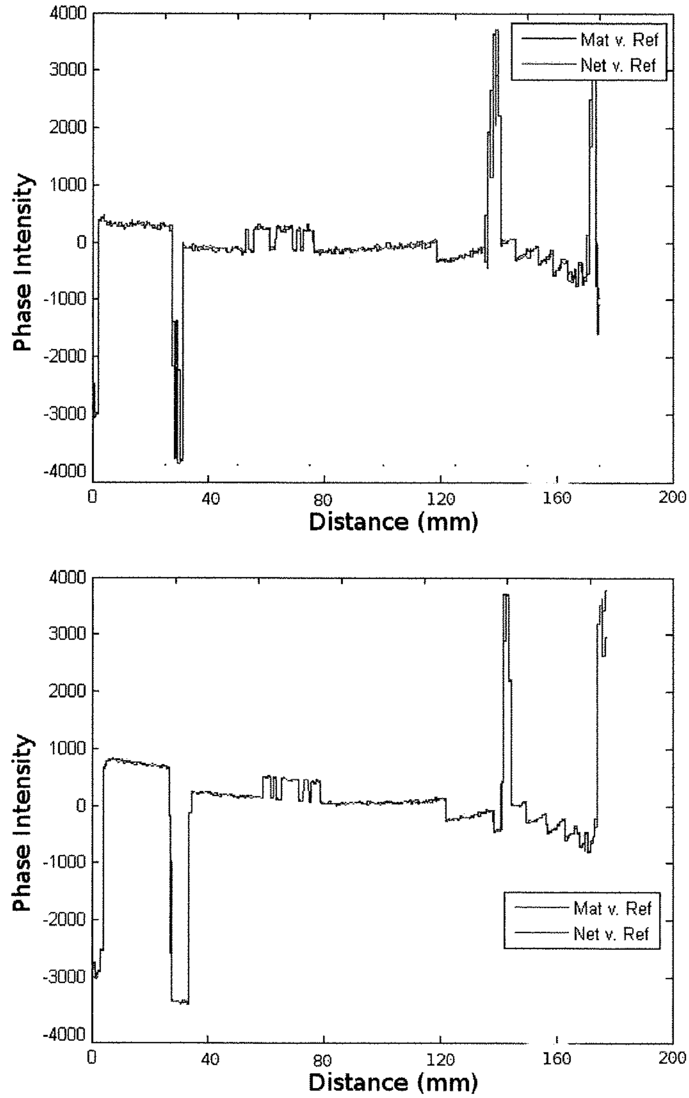

Figure 4.23: Phase Results at 40 cm (top) and 60 cm (bottom)

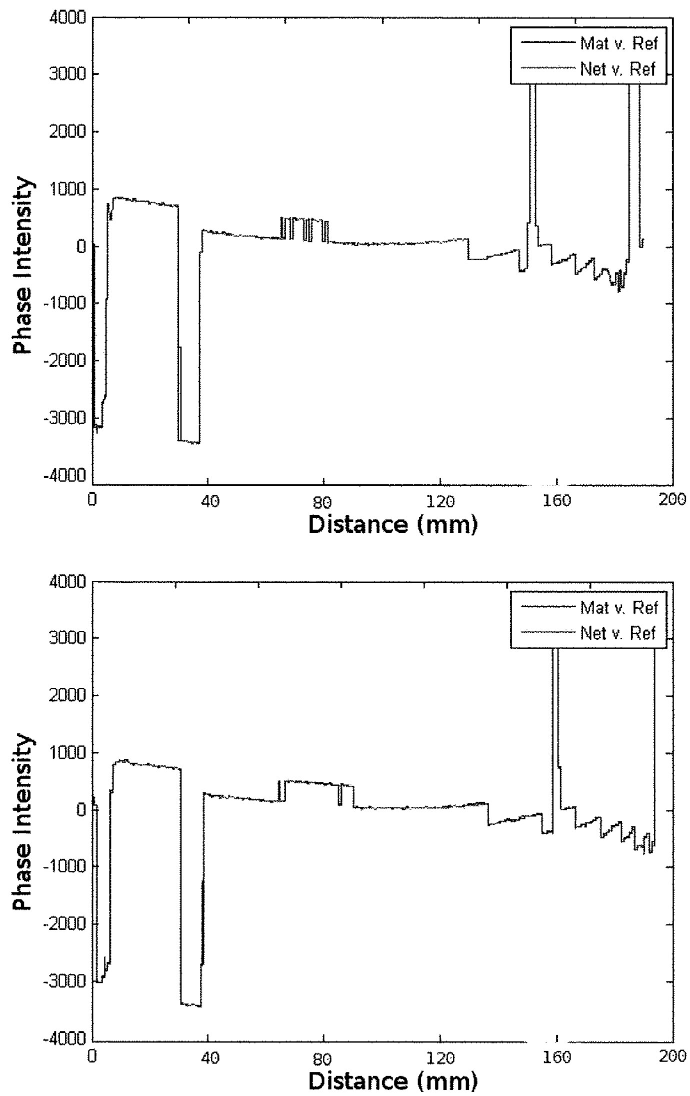

Figure 4.24: Phase Results at 100 cm (top) and 140 cm (bottom)

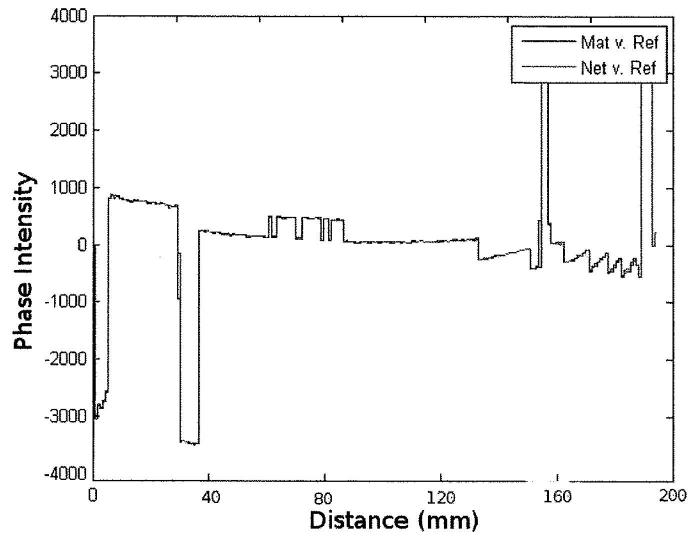

Figure 4.25: Phase Results at 190cm

These results show only minor improvement at short range, significant improvement at mid-range, and minor degradation when compared to previous results past 60 cm. This shows that the removal of the ferromagnetic oscillator does lessen the impact of the materials on the phase images of the MRI. The degradation at longer ranges was likely caused by the introduction of the DB-9 cable, which had some ferromagnetic components.

The interface also worked as expected, displaying the received sensor readings visually for inspection. While the MRI was not scanning, the sensor readings that were received by the sensor node and transmitted to the sink node were within the expected range. However, once the MRI began a scan sequence, a significant amount of errors (approximately 5-10% of total readings) began to be recorded by the interface. These errors are analog in nature, caused by voltage spikes in the physical sensors. This indicates that there is an additional source of interference to consider, the effect of the MRI scan processes on the physical sensors themselves. As these voltage spikes are likely caused by the 127MHz signal generated when the MRI is scanning, and the thermocouples generate a DC voltage, placing a low pass filter with a corner frequency lower than 127MHz on the output of the multiplexer before the amplifier may resolve this issue. Thermocouples may not be the only type of sensor used, however. The hydrophones used with MRgFUS, for example, have a much higher frequency on the output, and so the corner frequency of the low pass filter should be above the 4MHz that is typically used by ultrasound.

These results show that care must be taken when selecting the components for the sensor packages, as they can be affected by the MRI scanner. Given that the thermocouples are simply strands of wire made of different alloys, it is possible that the thermocouple acted as an antenna for the MR signal. Minimizing the length of the thermocouple (or the connecting wires for other sensors) may assist in reducing the amount of interference picked up by the sensors.

The DB-9 ports were used simply because the sink node already had the port built in. While this solution was sufficient for this experiment, it also had some drawbacks as discussed earlier. It would also not be sufficient in cases where multiple sink nodes are present. By implementing optical communications, it would eliminate the need for a DB-9 cable to be passed through the Faraday cage. As optical cables are generally not ferromagnetic and smaller than DB-9 cables, it reduces the amount of material that could potentially interfere with the imaging process.

RS-232 was selected as it is a simple to design and build transmission protocol. With a single sensor node and sink node, it is suitable for the amount of data that is being transmitted. As the amount of sensor nodes increase, RS-232 becomes more and more unsuitable as well, due to the baud rate limitation. A new transmission method, similar to USB or Ethernet, would be more suitable at these higher data rates.

Chapter 5

The Communication Block and Future Work

As was previously covered in Chapter 4, the basic operation of the sensor block and the node block was tested. Work was begun on constructing the communications block, such that to demonstrate the complete operational system. When these devices were being built, the only easily available optical transmitter/receiver pairs that were easily available were the TOTX147PL and the TORX147PL (Toshiba, Japan). They were designed for digital audio transmissions, and it was thought that they would be sufficient for this purpose. These components were added to the sink node, connecting to the micro-controller's UART transmitter. For the other side of the communications line, an optical to RS-232 converter was built.

It was noted that these optical components had a minimum required data

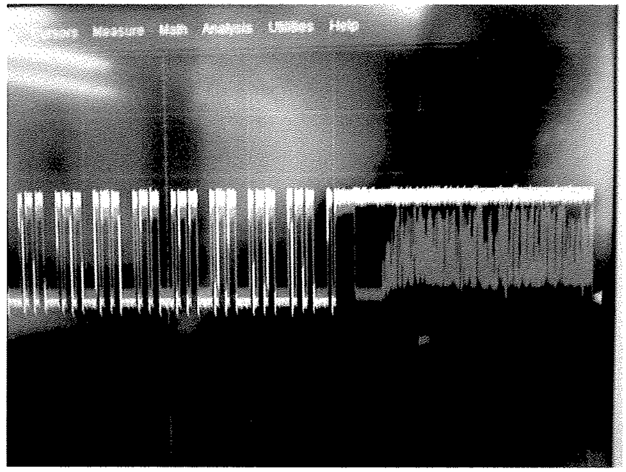

Figure 5.1: Unregulated Optical Receiver Output

rate in order for proper operation. To meet this requirement the sink node's serial transmitter's transmission baud was set to approximately 200 kbit/sec, twice the minimum recommended data rate of the optical components. In order to compensate for the instability of the receiver when the serial channel was idle, low pass filters were added to just after the optical receivers. This greatly reduced the amount of noise that the receivers output when there was no serial communication on the line. For an image of the unregulated output of the optical receiver, refer to Figure 5.1.

As seen in the image, the pink received optical signal matches the outputted signal from the sink node within tolerable levels. When the channel goes idle however, the receiver at first attempts to match the idle, drops low, then begins outputting a noise pattern. The filters removed the noise pattern as expected, however when the receiver drops low it is still seen as valid data by the computer. This causes the serial buffer to be forced multiple bits out of phase of what is actually transmitted. A solution to this problem while using the current optical receivers was not found, however, further research found that general purpose receivers such as the TORX150L (Toshiba, Japan) does not suffer from a minimum data rate, and was likely to be compatible with the current receivers and optical cables. Unfortunately, such receivers are obsolete and difficult to obtain. That series of receivers had been replaced with new receivers such as the TORX1950L (Toshiba, Japan). The TORX1950L is not yet available in North America as of yet, and so the completion of the Communication Block is pending the arrival of the new optical receiver.

On top of the immediate goal of finishing a proof of concept of the communication block, there are several short-term goals that should be looked into. Further research should be put into finding an optimal WSN protocol in this situation. A comparison of protocols should include the performance metrics of bit/s, amount of serviceable nodes, power used, and impact to the MRI. Power usage is a metric that should be considered, as the system may have to operate on a limited amount of power. This is due to the fact that as the power density of the system increases, the risk of greater interference with the MRI also increases. New sensor nodes should also be developed, specifically with hydrophones to begin to integrate testing of this technology with MRgFUS. In order to correct for the interference generated by wireless communications, as well as possible harmonics generated by the microprocessors itself, low pass filters should be implemented. These low pass filters should have a corner frequency above 127MHz and well below 2.4GHz should be implemented before the pre-amplifiers. Should this be done, the impact of the network communications on the SNR of the MRI images will be greatly reduced.

An intermediate term goal is to consider how the entire network should handle the signal processing of the sensor data. There are several methods that can be used to implement the required signal processing: processing at the nodes, processing at the receiving computer, or a combination of the two. Processing at the nodes have the advantage of greatly reducing the amount of data that has to be transmitted through the network to the sink node; its downsides are that this greatly increases the complexity of each of the nodes, and signal processing requires a significant amount of power, increasing the amount of interference and heat generated by the system as a whole. These downsides make it unlikely that this method of signal processing will be used with this system. The second method of signal processing, implementing it at the receiving computer, has the advantages of being able to keep the complexity of the sensor nodes to a minimum; at the downside of greatly increased network traffic. The third method of signal processing is a compromise between the two extremes. The nodes perform basic signal processing to reduce the amount of network traffic that has to be transmitted, and the receiving computer finishes that process. In order to implement this however, the balance between node complexity and network traffic should be explored in detail first.

Long term goals as presented by this thesis is to implement this system to assist with MRgFUS, or with other applications that may be developed between now and when that occurs. A full scale system may include hundreds or thousands of sensor and actuator nodes, multiple sink nodes and corresponding optical converters, and user interfaces that correspond to the applications that the system has been developed for.

Chapter 6

Conclusion

This thesis has proposed and begun testing of a system of implementing Wireless Sensor Networks (WSNs) to the Magnetic Resonance Imaging (MRI) environment. Previous to this system, any devices or sensors had to pass communication or power lines either through the waveguide of the Faraday cage or through a specially prepared partition in the Faraday cage. There was also a practical limitation on how many wires could be passed through the Faraday cage, limiting how many devices or sensors could be operating within the MRI environment at any one time.

Required background information is briefly covered by this thesis. The introductions of the supporting technologies for the proposed system such as WSNs, Fibre-optic Communications, and MRIs are provided. Information on possible applications for the system is also covered in this thesis. From this information, possible problems that this thesis addresses are identified and described.

With the introduction of WSNs to the MRI environment, further development of other interventions (such as MRgFUS) is possible. Without the limitation on how many sensors could be passed through the Faraday cage, techniques to treat areas behind solid masses (such as ribs or the skull) can become viable. The increased amount of sensor information increases the ability to compensate for the defocusing and attenuation caused by rigid structures in the path of the ultrasound field. The increased amount of sensor devices could also open avenues of research into simultaneous scanning methods.

Experimental results introduced by this thesis have shown that the proposed system is viable in the MRI environment. Those results have identified that the proposed system does generate additional interference to the MRI images beyond what the base material introduces. This interference to the SNR from the network communications is >3 dB less when compared to the images with only the materials. Interference to the proposed system from the MRI was also observed experimentally. This interference causes errors in 5-10% of sensor readings when the MRI is scanning. The thesis suggests methods of reducing and/or eliminating all observed sources of interference to both the proposed system and the MRI imaging process.

The thesis suggests methods of completing the system for experimental use for future work. It also suggests possible future devices to be developed (such as new sensor packages and nodes), and possible future topics of research (such as research into WSN protocols and signal processing). Long term goals of integrating the proposed system with existing developed technologies such as MRgFUS are stated.

Appendix A

MATLAB Programs

A.1 SNR Analysis Program

%%PARTIto each

%%Create find reference images, average them, create mask, and generate SNR

%%for the reference function Analysis

SNRresults=CompareMatVsRef({'IM_0534','IM_0535','IM_0536','IM_0537'},

{'IM_0562','IM_0563'},

{'IM_0564','IM_0565'});

function AvgImage=GetImages(NameofFiles)

numofref = length(NameofFiles);

reftotal=zeros(240);

for k = 1:numofref

```
refget = dicomread(NameofFiles{k});
reftotal= reftotal + double(refget);
end
AvgImage = reftotal/numofref;
function [refSNR, matSNR, NetSNR] = CompareMatVsRef(namefilesref,namefilesMaterial,nar
%% Average the images
refavg = GetImages(namefilesref);
%% Generate Signal Mask, assuming a 240x240 matrix
maskRefThres = refavg >30;
signaltotal = refavg(maskRefThres(:));
noisetotal = refavg(~maskRefThres(:));
signalavg = mean(signaltotal);
noiseavg = mean(noisetotal);
refSNR = 10*log((signalavg / noiseavg))
%%%%%%%%%%%%%%%
matavg = GetImages(namefilesMaterial);
signaltotal = matavg(maskRefThres(:));
noisetotal = matavg(~maskRefThres(:));
signalavg = mean(signaltotal);
noiseavg = mean(noisetotal);
matSNR = 10*log((signalavg / noiseavg))
%%%%%%%%%%%%%%%
Netavg = GetImages(namefilesNetwork);
```

```
signaltotal = Netavg(maskRefThres(:));
noisetotal = Netavg(~maskRefThres(:));
signalavg = mean(signaltotal);
noiseavg = mean(noisetotal);
NetSNR = 10*log((signalavg / noiseavg))
```

A.2 Phase Analysis Program

```
function PhaseAnalysis
    clear;
    close all;
    PhaseCompare ({'IM_0692', 'IM_0693', 'IM_0694', 'IM_0695', 'IM_0696',
'IM_0697'},
    {'IM_0689', 'IM_0690', 'IM_0691'},
    {'IM_0686', 'IM_0687', 'IM_0688'});
    function AvgImage=GetImages(NameofFiles)
    numofref = length(NameofFiles);
    reftotal=zeros(240);
    for k = 1:numofref
    refget = dicomread(NameofFiles{k});
    reftotal= reftotal + double(refget);
    end
    AvgImage = reftotal/numofref;
    function PhaseMagnitude = GetPhaseMagnitude(x,y, matcompare)
    prevy = 0;
    if (x(1) == x(2))
    xi = x(1);
    yi = y(1):1:y(2);
    for l = 1:1:length(yi);
```

```
PhaseMagnitude(l) = matcompare (yi(l),xi);
end
else
if (x(2) < x(1))
t = x(2);
x(2) = x(1);
x(1) = t;
t = y(2);
y(2) = y(1);
y(1) = t;
end
m = (y(2) - y(1))/ (x(2)-x(1));
b = y(1) - m*x(1);
xi = x(1):0.01:x(2);
for l = 1:1:length(xi)
yi(l) = m*xi(l) + b;
yi(l) = round(yi(l));
if(yi(l) ~= prevy)
xi(l) = round(xi(l));
PhaseMagnitude(l) = matcompare(yi(l), xi(l));
end
yi(l) = prevy;
end
```

```
pause();
end
function PhaseCompare (namefilesref,namefilesMaterial,namefilesNetwork)
%Average the reference images
refavg = GetImages(namefilesref);
%Average the Material images
matavg = GetImages(namefilesMaterial);
hold off;
%Average the Network images
netavg = GetImages(namefilesNetwork);
%Begin comparing the Materialv.Reference
%Subtract the material image from the referance image
matcompare = refavg-matavg;
figure(1);imshow(matcompare,[]);
hold on;
%Get the user to input a line on the resultant image
[x,y] = ginput(2);
PhaseLine = line(x,y);
plot(PhaseLine);
y = round(y);
x = round(x);
MatPhaseMagnitude = GetPhaseMagnitude(x,y, matcompare);
hold off;
```

```
figure(2);
plot (MatPhaseMagnitude);
%Begin comparing the Materialv.Reference
%Subtract the material image from the referance image
netcompare = refavg-netavg;
figure(3);imshow(netcompare,[]);
hold on;
PhaseLine = line(x,y);
plot(PhaseLine);
NetPhaseMagnitude = GetPhaseMagnitude(x,y, netcompare);
hold off;
figure(4);
plot (NetPhaseMagnitude);
figure(5);
plot (MatPhaseMagnitude, 'b');
hold on;
plot (NetPhaseMagnitude, 'r');
legend1 = legend('Mat v. Ref', 'Net v. Ref');
hold off;
```

A.3 System Interface

```
function varargout = interface(varargin)
gui_Singleton = 1;
gui_State = struct('gui_Name', mfilename, ...
'gui_Singleton', gui_Singleton, ...
'gui_OpeningFcn', @interface_OpeningFcn, ...
'gui_OutputFcn', @interface_OutputFcn, ...
'gui_LayoutFcn', [], ...
'gui_Callback', []);
if nargin && ischar(varargin{1})
gui_State.gui_Callback = str2func(varargin{1});
end
if nargout
[varargout{1:nargout}] = gui_mainfcn(gui_State, varargin{:});
else
gui_mainfcn(gui_State, varargin{:});
end
function interface_OpeningFcn(hObject, eventdata, handles, varargin)
handles.output = hObject;
% Update handles structure
guidata(hObject, handles);
function varargout = interface_OutputFcn(hObject, eventdata, handles)
```

```
varargout{1} = handles.output;
function pushbutton9_Callback(hObject, eventdata, handles)
%%Initialize
handles.var.errors = 0;
handles.var.readings = 0;
handles.var.loopcount = 0;
handles.var.loop = 1;
set (handles.loop, 'String', 1);
%%Set and open serial port
serialPort = serial('COM4', 'BaudRate', 19200);
serialPort.terminator = 'CR';
fclose(serialPort);
fopen(serialPort);
%%Loop until stop is clicked
while(handles.var.loop == 1)
errors = handles.var.errors;
readings = handles.var.readings;
%%Grab serial data
if (serialPort.BytesAvailable > 0)
go = 0;
while ((go == 0)&&(handles.var.loop==1));
lfstart = fscanf(serialPort, '%s', 2);
if (hex2num(lfstart) == hex2dec('AA'))
```

```
go = 1;
else
handles.var.loop = get(handles.loop, 'String');
end
end
for i = 1:1:8
textinput = fscanf(serialPort, '%s', 2);
input(i) = hex2dec(textinput);
if (input(i) < 100)
temp(i) = input(i);
readings = readings + 1;
else
errors = errors + 1;
end
end
handles.var.temp1 = temp(1);
handles.var.temp2 = temp(2);
handles.var.temp3 = temp(3);
handles.var.temp4 = temp(4);
handles.var.temp5 = temp(5);
handles.var.temp6 = temp(6);
handles.var.temp7 = temp(7);
handles.var.temp8 = temp(8);
```

```
set (handles.temp1, 'String', handles.var.temp1);
set (handles.temp2, 'String', handles.var.temp2);
set (handles.temp3, 'String', handles.var.temp3);
set (handles.temp4, 'String', handles.var.temp4);
set (handles.temp5, 'String', handles.var.temp5);
set (handles.temp6, 'String', handles.var.temp6);
set (handles.temp7, 'String', handles.var.temp7);
set (handles.temp8, 'String', handles.var.temp8);
handles.var.readings = readings;
handles.var.errors = errors;
set (handles.readings, 'String', handles.var.readings);
set (handles.errors, 'String', handles.var.errors);
else
end
loopcount = handles.var.loopcount;
handles.var.loopcount = loopcount + 1;
set (handles.loopcount, 'String', handles.var.loopcount);
drawnow();
handles.var.loop = str2num(get(handles.loop, 'String'));
end
fclose(serialPort);
function pushbutton10_Callback(hObject, eventdata, handles)
handles.var.readings = 0;
```

```
handles.var.errors = 0;
handles.var.loopcount = 0;
set (handles.readings, 'String', handles.var.readings);
set (handles.errors, 'String', handles.var.errors);
set (handles.loopcount, 'String', handles.var.loopcount);
function Stop_Callback(hObject, eventdata, handles)
set(handles.loop, 'String', num2str(0));
```

Bibliography

[1] Raghavendra, C. S., Krishna M. Sivalingam and Taieb Znati. Wireless Sensor Networks. New York: Springer Science+Business Media, 2004.Results and Observations

[2] Heinzelman, Wendi Rabiner, Anantha Chandrakasan and Hari Balakrishnan. "Energy-Efficient Communication Protocol for Wireless Sensor Networks." 3rd Hawaii International Conference on System Sciences. 2000. 1-10.

[3] Kuorilehto, Mauri, et al. Ultra-Low Energy Wireless Sensor Networks in Practice. Chichester: John Wiley & Sons Ltd., 2007.

[4] The ZigBee Alliance. "ZigBee Specification". ZigBee Standards Organization. January 2008.

[5] Nakamura, Eduardo F., Antonio A. F. Loureiro and Alejandro C. Frery. "Information Fusion For Wireless Sensor Networks:

Methods, Models, and Classifications." ACM Computing Surveys. August 2007: Article 9, 1-55.

[6] Li, Yingshu, My T. Thai and Weili Wu. "Wireless Sensor Networks and Applications." New York: Springer Science+Business Media, 2008.

[7] Stern, Thomas E., Georgios, Ellinas and Bala, Krishna. Multiwavelength Optical Networks: Architectures, Design, and Control. Second Edition. New York: Cambridge University Press. 2009.

[8] Zheng, Jun and Mouftah, Hussein T. "A Survey of Dynamic Bandwidth Allocation Algorithms for Ethernet Passive Optical Networks" Optical Switching and Networking. 2009: Volume 6, 151-162.

[9] Naser, Hassan, and Mouftah, Hussein T. "A Joint-ONU Interval-Based Dynamic Scheduling Algorithm for Ethernet Passive Optical Networks" IEEE/ACM Transactions on Networking. 2006: Volume 14, 889-899.

[10] Ramaswami, Rajiv, Sivarajan, Kumar N. and Sasaki, Galen H. Optical Networks: A Practical Perspective. Third Edition. Burlington: Elsevier Inc., 2010.

[11] Pereira, Flavio de Melo, Fonscea, Nelson L. S., and Arantes, Dalton S. "A Fair Scheduling Discipline For Ethernet Passive Optical Networks" Computer Networks. 2009: Issue 53, 1859-1878.

[12] Nishitani, Takashi et al. "Adaptive Dynamic Bandwidth Allocation Scheme for Multiple-Services in 10G-EPON System" IEEE Communications Society. 2011.

[13] An, Fu-Tai et al. "SUCCESS: A Next-Generation Hybrid WDM Optical Access Network Architecture" Journal of Lightwave Technology. 2004: Volume 22, 2557-2569.

[14] Dhaini, Amhad R. et al. "Dynamic Wavelength and Bandwidth Allocation in Hybrid TDM/WDM EPON Networks" Journal of Lightwave Technology. 2007: Volume 25, 277-286.

[15] Pal, Bishnu P. Fundamentals Of Fibre Optics In Telecommunication And Sensor Systems. Noida: New Age International Ltd. 1992.

[16] Lee, Byoungho. "Review of the Present Status of Optical Fibre Sensors." Optical Fibre Technology. 2003: Issue 9, 57-79.

[17] Haacke, Mark E. et al. Magnetic Resonance Imaging: Physical Principles and Sequence Design. New York: John Wiley & Sons Inc. 1999.

[18] Weishaupt, Dominik , Kochli, Victor D., Marincek, Borut. How Does MRI Work: An Introduction to the Physics and Function of Magnetic Resonance Imaging. New York: Springer. 2006.

[19] Smith, Robert C. and Lange, Robert C. Understanding Magnetic Resonance Imaging. Boca Raton: CRC Press LLC. 1999.

[20] Hynynen, Kullervo. "Review of Ultrasound Therapy." IEEE Ultrasonics Symposium. 1997. 1305-1313.

[21] Jolesz, Ferenc A. "MRI-Guided Focused Ultrasound Sugery." The Annual Review of Medicine. 2009. 417-430.

[22] Hudson, Susan B. A. and Stewart, Elisabeth A. "Magnetic Resonance-guided Focused Ultrasound Surgery" Clinical Obstetrics And Gynecology. 2008: Volume 51, Number 1, 195-166.

[23] McDannold N, Hynynen K, Jolesz F. "MRI Monitoring Of The Thermal Ablation Of Tissue: Effects Of Long Exposure Times". Journal of Magnetic Resonance Imaging 2001: Volume 13:421–27

[24] Furasawa, Hidemi et al. "Magnetic Resonance–Guided Focused Ultrasound Surgery of Breast Cancer: Reliability and

Effectiveness." Journal of American College of Surgeons. 2006: Issue 203, 54-63.

[25] Tempany, Claire M. C. et al. "MR Imaging– guided Focused Ultrasound Surgery of Uterine Leiomyomas: A Feasibility Study." Radiology. March 2003: Issue 226, 897-905.

[26] Kennedy, J. E. et al. "High-Intensity Focused Ultrasound for the Treatment of Liver Tumors." Proceedings of Ultrasonics International. April 2004: Volume 42, 931-935.

[27] Gelet, A. et al. "Transrectal High-Intensity Focused Ultrasound: Minimally Invasive Therapy of Localized Prostate Cancer" Journal of Endourology. August 2000: Volume 14, 519-528.

[28] Philips. Sonalleve MR-HIFU. June 2011. (http://www.healthcare.philips.com/main/products/mri/systems/sonalleve/).

[29] InSightec. Focused Ultrasound Application | Treatment of Uterine Fibroids. June 2011. (http://www.insightec.com/).

[30] Maple Leaf HIFU Co. About Ablatherm HIFU. June 2011. (http://www.hifu.ca/patient/about_ablatherm.php).

[31] Focus Surgery Inc. Sonablate 500. June 2011. (http://www.focus-surgery.com/Sonablate500.htm).

[32] ter Haar, Gail. "Review Therapeutic Ultrasound." European Journal of Ultrasound. 1999: Volume 9, 3-9.

[33] O'Reilly, Meaghan A. and Hynynen, Kullervo. "A PVDF Receiver for Ultrasound Monitoring of Transcranial Focused Ultrasound Therapy." IEEE Trans Biomed Eng. May 2010: Volume 57, Issue 9, 2286-2294.

[34] O'Reilly, Meaghan A. et al. "Design and Construction of a Passive Receiver Array for Monitoring Transcranial Focused Ultrasound Therapy." Ultrasonics Symposium (IUS). October 2010: 890-892.

[35] Yu, Jaewook, Noel, Eric, and Tang, K. Wendy. "Degree Constrained Topology Control for Very Dense Wireless Sensor Networks" IEEE Globecom. 2010: 1-6.

[36] Chinzei, Kiyoyuki et al. "MR Compatibility of Mechatronic Devices: Design Criteria." Medical Image Computing and Computer-Assisted Intervention. 1999: 1020-1030.

[37] Tang, Jun et al. "A Hybrid Radio over Fiber Wireless Sensor Network Architecture" IEEE Explore. 2007.

[38] Sodre, Arimar Cerqueira Jr. et al. "Performance Analysis Of A Radio Over Fiber System Based On IEEE 802.15.4 Stan-

dard In A Real Optical Network." Microwave and Optical Technology Letters. 2009: Volume 51, 1876-1879.

[39] Al-Fares, Mohammad S., Sun, Zhili, and Cruickshank, Haitham. "High Survivable Routing Protocol in Self Organizing Wireless Sensor Network." IAENG International Journal of Compuer Science. 2009: Volume 36.

[40] Microchip Technology Inc. "PICDEM Z Demonstration Kit User's Guide". 2007. <http://ww1.microchip.com/downloads/en/DeviceDoc/51524b.pdf>.