WIRELESS DATA ACQUISITION SYSTEM
FOR MULTI-PHASE ELECTRIC POWER EQUIPMENT

A Thesis

by

DOUGLAS ANDREAS GOODSELL

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

May 2008

Major Subject: Mechanical Engineering

WIRELESS DATA ACQUISITION SYSTEM
FOR MULTI-PHASE ELECTRIC POWER EQUIPMENT

A Thesis

by

DOUGLAS ANDREAS GOODSELL

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Approved by:

Chair of Committee,	Alexander Parlos
Committee Members,	Darbha Swaroop
	Jose Silva-Martinez
Head of Department,	Dennis L. ONeal

May 2008

Major Subject: Mechanical Engineering

ABSTRACT

Wireless Data Acquisition System
for Multi-Phase Electric Power Equipment. (May 2008)
Douglas Andreas Goodsell, B.S., Texas A&M University
Chair of Advisory Committee: Dr. Alexander G. Parlos

Industrial facilities that plan the shutdown of equipment for service have large financial savings compared to those managing unplanned shutdowns. To this end, a variety of algorithms have been developed and published in the literature that can monitor a machine's health and indicate when the machine starts to develop a fault. In order for such algorithms to be effective, they require raw data collected from machines. Often this involves the placement of accelerometers and other sensory devices for measurements of mechanical behavior. It is possible to extract much of the required information from the electrical signals of the equipment. This is normally a less expensive installation since one only needs access to the lines supplying electric power to the equipment. If these data acqusistion modules are accessible wirelessly, then one can monitor all the interfaced equipment from a central location. To successfully monitor such electrical equipment, a data acquisition unit is required that can sample on five or six channels simultaneously, depending on the switch gear configuration.

This thesis details the development of an "endpoint" device that samples the required number of channels to monitor the electrical signals of industrial equipment, and interfaces to a wireless network. The hardware and software design of the "endpoint" is discussed in detail. Also, the software design of the server that receives the data from the "endpoint" is presented.

The designed "endpoint" samples up to six channels simultaneously, at a rate of at least 8 kHz per channel, and a data resolution of 16 bits. The data are then transmitted wirelessly to a central server for processing. The system has been tested both in a laboratory environment and at an industrial environment. The desired specifications of the "endpoint" have been verified in both environments. Several "endpoints" have been assembled to form a network and have been tested in a laboratory setting.

This work has resulted in the demonstration that an "endpoint" can be constructed using off the shelf components that is suitable for the continuous health monitoring of multi-phase electric machines through a wireless network.

To My Wife

ACKNOWLEDGMENTS

I would like to thank my committee chair and advisor, Dr. Alexander G. Parlos. Without his guidance and patience, this work would not have been possible. I would also like to thank Dr. Darbha Swaroop and Dr. Jose Silva-Martinez for their support in this work.

Aninda and Parasuram desrve my thanks and appreciation for their contributions and assistance.

To my parents, I would like to convey a sincere thank you for their continued support in all that I do.

To my wife, thank you for your unconditional support and continued encouragement.

TABLE OF CONTENTS

CHAPTER Page

LIST OF TABLES

LIST OF FIGURES

FIGURE Page

CHAPTER I

INTRODUCTION

A. Motivation

If an industrial facility could plan the shutdown of equipment for service, then there would be large financial savings compared to unplanned shutdowns. To this end, algorithms have been developed that can monitor a machine's health and indicate when the machine starts to deteriorate, eventually leading to an unplanned failure [1]. In order for these algorithms to be executed, they require continuous data streams from the machine to be acquired or sampled. Often this involves the placement of accelerometers and other sensory devices for measurement of mechanical behavior. It is possible to acquire much of the required condition information from the electrical measurements of the equipment [1]. This is normally a less expensive installation since one only needs access to the low voltage side of the power cables supplying the equipment. Currently, to monitor the equipment a module is placed at the unit that collects and processes the data. This requires someone to visit each module to determine the machine's status. If these modules are accessible wirelessly, then one can monitor all the equipment from a central location. To successfully monitor such equipment, a unit is required that can continuously collect data on five or six channels simultaneously (three currents and at least two voltages) for a three phase installation.

The journal model is *IEEE Transactions on Automatic Control.*

B. Problem Definition

Currently, a cost effective, scalable solution for the condition monitoring of three phase power equipment does not exist. A scalable solution can be considered cost effective if the cost of adding a unit only involves the cost of the unit. For example, a cost effective solution would not require additional employees to cover the day-to-day operation of the unit. The algorithms for health monitoring, using sampled data, do exist, but have not been implemented in a cost effective, scalable solution. The problem to be addressed is the creation of data acquisition based on a standard wireless interface, sampling five or six channels at a rate of at least 4 kHz and a resolution of at least 14 bits, with the ability to adjust the analog input in order to capture the startup signal.

C. Proposed Approach

To design a cost effective, scalable solution, it is proposed to develop a wireless data acquisition specific for three phase power equipment. By using a wireless solution, the requirement for cost effective scalability is met as the data would be transfered to a central point, requiring minimal additional costs. To further reduce the cost, employing algorithms that use the electrical inputs to the equipment will reduce installation costs as the power input is much more readily available than the actual equipment for the placement of mechanical sensors.

In addition to reducing the cost of the installation, leveraging off the shelf components will reduce development and production costs. These components have already been verified for the environment they will be placed in. When available, already developed sub-systems further reduce associated development time.

The algorithms to be employed require floating point processing. This could be done at the equipment, using a floating point DSP, or at a central server. A central server approach was decided upon as the floating point DSPs are relatively expensive to deploy in large quantities. Further, to reduce operating costs, the health status needed to be available at a central point such as the server that could be used to process the data. Also, by using the server, it could be possible to make the data available by a remote interface such as a network interface.

D. Literature Review

The available literature documents many wireless sensor designs in addition to many units being commercially available. In his text, Fred Eady detailed the development of a wireless unit based on an 8-bit microcontroller and a compact flash 802.11 card [2]. Ferrari, et al., developed an 802.11 based unit using an 8 bit microcontroller, but with a PCMCIA 802.11b card [3]. For the monitoring of structures, Townsend, et al., developed an 802.11b based network using an 8 bit microcontroller with a low update rate [4].

Moving to a much greater complexity, Raygan, et al., contemplated the design of a wireless sensor system using Linux and 802.11b [5]. Also using a PC based system, Kohvakka, et al., used a PC/104 stack with Windows NT and a PCMCIA 802.11b adapted card [6].

Of the commercially available systems, Crossbow Technology Inc., has a system based on radio boards referred to as Motes. This is a proprietary network system intended for low power applications. These Motes are interfaced to many of the company's sensors in addition to external sensors [7]. Honeywell has a wireless sensor development kit based on the 802.15.4 standard with a data rate of 250 kbps [8].

Of those systems documented, none of them meet all of the requirements in the problem definition.

E. Research Contribution

The research conducted contributes a complete wireless solution for the data acquisition for three phase power equipment. This system provides a scalable wireless network solution to collect data at a sufficient bit rate to execute the existing algorithms to monitor the health of three phase power equipment.

F. Thesis Organization

This thesis starts with the development of the endpoint. This includes the details of the hardware selection and the software development. Following the development of the endpoint, the system verification procedures are described. Next, the system performance is presented. The conclusions of the development are then discussed with an outline of possible future work.

CHAPTER II

HARDWARE DESIGN

A. System Overview

The data acquisition system is composed of several blocks, as seen below in figure 1.

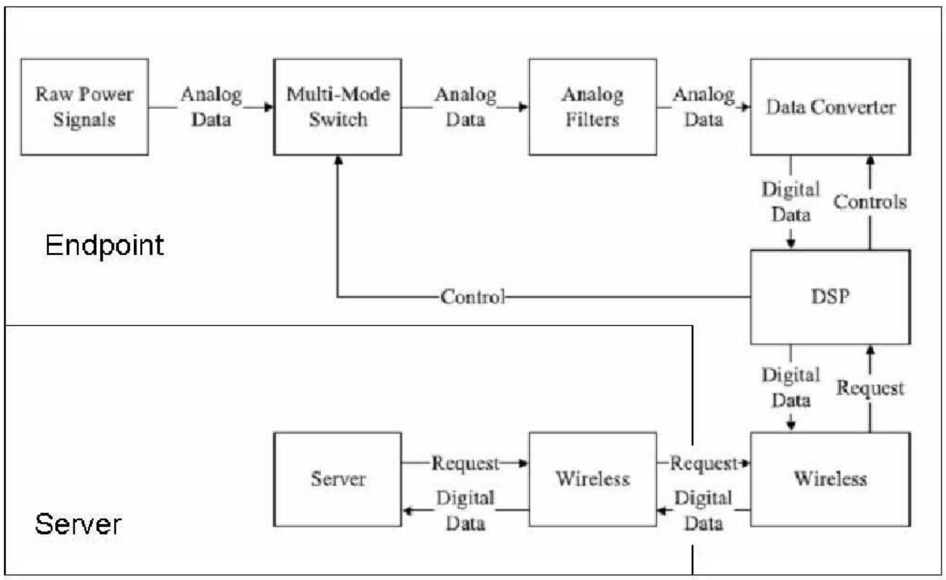

Fig. 1. System Data Flow Diagram

The raw power signals are collected via current transformers and brought into the multi-mode switch. This switch allows the steady-state signals and the transient signals to both be fitted to the dynamic range of the data converter. After the switch, the signals pass through analog filters that were designed previously to this work. These signals are then sampled by an analog to digital converter who passes the data on to a digital signal processor. This processor then needs to wirelessly transmit the data to a central server. This central server will then receive, store, and process the data, yielding the needed information regarding the power equipment's health.

B. Data Converter Selection

The AD73360 A/D converter from Analog Devices was selected for three reasons. The first reason is that the minimum channel requirement for monitoring three-phase power requirement is five channels. This results in monitoring two voltages and three currents in the delta configuration. Another possible configuration is the Y configuration, which requires measuring three voltages and three currents. This means the selected device needed to be capable of measuring six channels, which the AD73360 is capable of doing. In order to reduce the digital signal processing, it is advantageous to monitor the six channels simultaneously. If the channels are sampled sequentially, a phase shift would be introduced that would need to be accounted for. This can be accomplished using six separate A/D converters or a device like the AD73360. Six separate A/D converters create many complications with the interface for the selected processor and increase the complexity of the printed circuit board layout. The AD73360 provides the required simultaneous sampling of six channels in a single device[9].

According to the Nyquist criterion, to accurately digitize a signal, one must sample at a frequency of at least twice the highest frequency of interest in the signal. In practice, this is often much greater than a factor of two. For this application, others have determined that the signal information will be accurately captured if the signal is sampled at 3840 Hz[1]. The AD73360 has a minimum sampling rate of 8 kHz, which was more than adequate for this application. Due to the sampling rate being a non-integer multiple of the base frequency of three-phase power in the U.S., the digitized signal required some processing to down sample the signal to the 3840 Hz. This down sampling is beyond the scope of this thesis.

The third requirement of the digitizer was resolution. Since the signal information required for the assessment of the equipment is contained in the harmonics of the signal, significant digital resolution was required. This resolution has been determined to be at least 14 bits. The AD73360 contains six 16 bit converters with a minimum 12 bit signal to noise ratio. While this minimum is below what is desired, with sufficient signal conditioning, it should be possible to meet the requirement of at least 14 bits of effective resolution.

C. Multi-Mode Data Acquisition

Each of the endpoints is composed of several modules. The list of these is as follows:

- DSP
- Networking
- AD73360
- Power
- Analog signal conditioning

The first three of these have been discussed in the proceeding sections. The power module is comprised of transformers to step down and convert the available power source to those voltage levels needed to run the components on the endpoint. This was based on a pre-existing design.

The last module to be discussed is the analog signal conditioning. Each of the three voltage channels contain previously designed filters. In addition to the filters, the three current channels have a configurable voltage divider that is controlled by the DSP. The currents can be sensed by using either shunted or unshunted current transformers (CTs). The two versions of the CTs provide either a current or a voltage signal output. Through the use of jumpers on the endpoint, a current signal can be converted to a voltage signal by including a resistor in the circuit.

When the equipment is first turned on, there is a large current draw, many times the current draw of the steady state current[10]. This current draw can be correlated to the rated current of the equipment, thus providing needed configuration data for the server. In order to accurately capture this, the voltage divider needed to be designed to fit this signal into the available input range of the AD73360. If the steady state signal were to be sampled through this voltage divider, it would appear as nothing more than noise. By creating a dynamic voltage divider, the DSP can switch the amplification of the signal as it determines if the equipment is starting up or running in steady state.

To accommodate both types of current transformers, the start up, and steady state currents, several designs were evaluated. The first design used a series of muxes and demuxes. The outer set of muxes and demuxes were manually switched to select between a shunted or unshunted CT setup. The inner muxes and demuxes were controlled by the DSP depending on the state of the currents. In addition to being very complicated to implement, this design was found to introduce an impedance mismatch. This destroyed any useful information in the analog signals.

To protect the digital hardware from the analog signals, there is widespread use of isolation circuits. This complicates the design of the voltage bridge since it is possible that the common of the analog side may be at a different level than the common of the digital side. This difference in levels means that the DSP on the digital side of the isolators can not directly switch the analog circuit through standard transistors. With this in mind, opto-isolators were selected for the control mechanism in the divider circuit. The current flow through the isolators is regulated by an internal LED that is turned on and off by the DSP. This separates the analog and digital signals but still provides a means for the DSP to control the bridges. The schematic in figure 2 shows one channel of the dynamic bridge circuit. The DSP control is through PF9 in the figure.

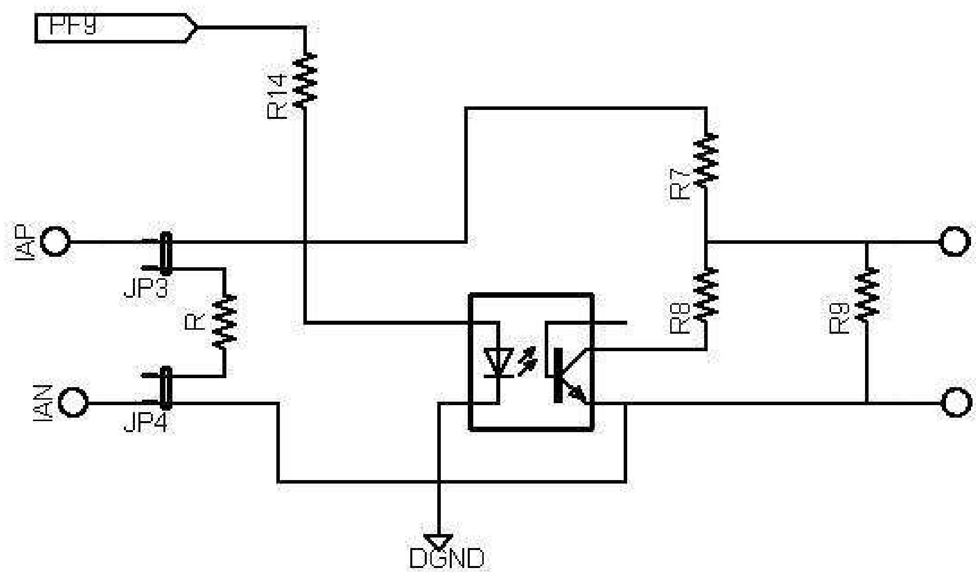

Fig. 2. Bridge Circuit Schematic

When the design went to the prototyper, it was recommended to switch from the opto-isolator used in the circuit shown in figure 2 to a solid state relay. This part operates in the same manor as the opto-isolator used in the bread-boarded design but does not introduce noise in the extremes of the analog signal. The breadboard implementation can be seen in figure 3.

The final piece to the dynamic bridges is deciding the thresholds that determine when the equipment has started up, is running in steady state, or has shutdown.

There are a variety of variables to consider in this. The first variable comes from the variation of electrical components. While these components meet their specifications very accurately, there is still some variation between components of the same specification. This variability results in variations in filter gains and signal biasing. The second variable comes from the noise inherent to analog signals. The amplitude of this noise can vary between channels on an endpoint and would not be consistent among the endpoints. Through processing of the digital signals, such as normalizing and filtering, these variables can be accounted for.

The variable that can not be accounted for is the equipment the endpoint is monitoring. The equipment will dictate the amplitude of the steady state signals and the amplitude of the startup signals. By using the processed digital signals, digital levels can be identified that clearly indicate when a signal has gone from the noise level to a startup, or from the steady state signal to the noise signal when the equipment shuts down. These levels will be unique to the equipment that is being monitored. The process by which the signals will be unbiased is discussed in the following chapter in addition to how the startup and shutdowns are detected.

D. Processor Selection

Prior to the selection of the DPAC radio, the possibility of interfacing a more generic DSP was considered. The evaluation board for the AD73360 made use of a PC's parallel port to download code to the RAM of the DSP and transfer data back to the PC. The evaluation board did not have enough memory to store the continuous 30 seconds of memory and would require streaming the data to the PC. In the preliminary stages of this design, it was found that the data loss with this streaming approach was unacceptable. As a result, higher speed interfaces were sought after.

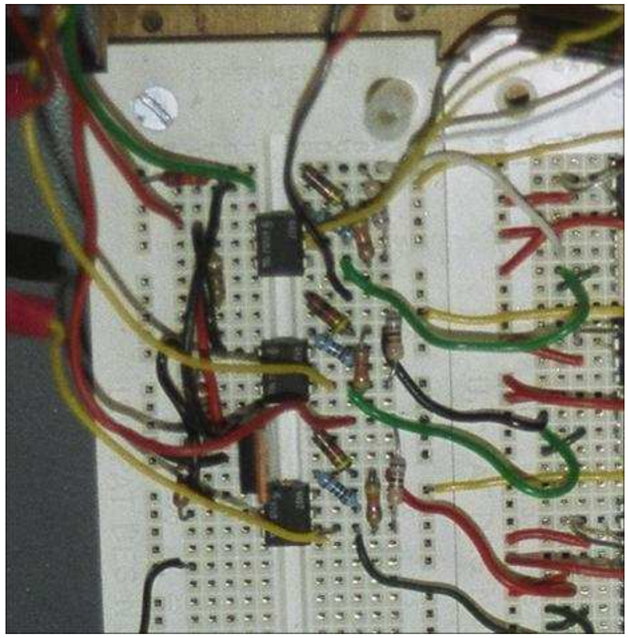

Fig. 3. Breadboard Implementation of the Multi-Mode Switch

With the selection of the AD73360 and the DPAC radio, the choices for a processor were severely limited. The AD73360 requires a full-duplex, synchronous serial interface that is not common on microcontrollers. The datasheet for the AD73360 outlined connections to Texas Instruments' Digital Signal Processors (DSPs) and Analog Devices' DSPs[9].

To accommodate the DPAC radio, an Ethernet interface was required. There were two options that could be pursued for this: interface Ethernet hardware to a port on a DSP and create the necessary software to support it or find a processor that already had Ethernet and use the manufacture's provided software. Analog Devices recently introduced their family of Blackfin DSPs among which are two versions, BF536 and BF537, that include a 10/100 Ethernet MAC. This greatly simplified the development process by only requiring the addition of an Ethernet physical layer chip to the DSP. Since the DSP had the MAC built in, Analog Devices also provided the necessary software tools for developing network enabled applications. These DSPs are fixed point processors with clock speeds of 300 to 600 MHz. This clock speed is more than adequate for handling the software the DSP will be running. The Blackin DSPs also included the needed serial interface for the AD73360. The development system for the Blackfin DSP can be seen in figure 4.

A single board computer (SBC) implementation was also considered and partially developed. This used an AMD Alchemy based SBC with available PCI slots and Ethernet port. With the use of an 802.11g PCI card, the wireless connection was achieved. The onboard Ethernet allowed the Blackfin DSP interface. This solution was deemed too expensive when compared to the DPAC radio as it still required an intermediate DSP and DSP resources.

The other block to this platform was the required development of an embedded operating system that ran on the MIPS architecture Alchemy processor. Such an operating system, based on the Linux kernel was attempted but abandoned due to development challenges that are beyond the scope of this thesis.

E. Wireless Selection

The decision to create a wireless endpoint was driven by the desire to reduce the total cost of the system. In the deployment of a diagnostic system, there are three costs to consider:

- Installation
- Operation
- Monitoring

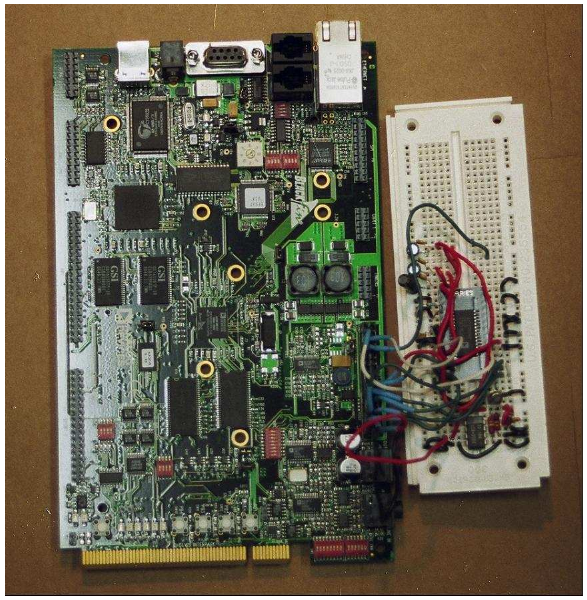

Fig. 4. Development Implementation of Processor and Data Converter

By going to the networked system, there is an increase in the installation costs of the system. A network requires additional hardware and connection cables. By choosing a wireless network, the connection cables are eliminated. The operation costs address how much money is required to have the system running. This cost may increase slightly as the wireless radios will increase the power drawn by the system, but this cost is more than offset by the savings in the other categories.

By switching to a networked diagnostic system, the monitoring cost is reduced. The networked system transfers all the information to a central point, eliminating the need for an individual to physically visit each endpoint. To choose a wireless network solution, several factors need to be considered. The first factor was transmission speed. It has been determined that the data would be collected according to the following formula (2.1):

$$\frac{Bits}{Second} = 6 Channels * 8000 \frac{Samples}{Channels * Second} * 16 \frac{Bits}{Sample} \tag{2.1}$$

It was determined that the data is required to be continuous for 30 seconds. This results in 23.04 megabits per data set. In order to allow for significant scaling of the system, this data needed to be transferred in 10 seconds or less. This meant a raw data rate of 2.304 megabits per second was required.

The second factor to be considered in choosing the wireless solution was scaling. The smallest setup includes one endpoint and one server but this would not be of much use in an industrial setting. A more practical setup includes a large number of endpoints and a single server. Thus, a wireless solution that requires a dedicated connection between each endpoint and the server is impractical as this would require the server to support as many connections as there are endpoints.

In the industrial environment, there were two more factors to be considered. The first was radio interference. It was advised that radio frequencies used by 802.11a, 5.8 GHz, often interfere with the equipment running in these environments. The second consideration was the expanse of these facilities. It is not impractical to require a radio range over 30 meters to allow signals to be received from remote equipment.

The first wireless option considered was a basic UART replacement. These systems generally operate in the 900 MHz spectrum and have transfer rates on the order of a standard serial port[11]. This was ruled out on the basis that the transfer rate is much to slow, and multiples of these radios do not easily interface to PCs since each radio link requires a unique connection on the PC. A second option considered was ZigBee (IEEE 802.15.4). This operates in the 2.4 GHz, 915 MHz and 868 MHz spectrums and has transfer rates on the order of 250 kilobits per second[12]. This specification is also too slow. Bluetooth was considered as well, but with transfer rates on the order of 1 megabit per second and ranges typically around 10 meters, this was inadequate[13].

The other option considered was the popular 802.11b/g standard. This standard allows for rates as high as 11 megabits per second for the b standard and 54 megabits for the g standard with the g standard being backwards compatible with the b standard. As with Bluetooth and ZigBee, 802.11b/g operates in the 2.4 GHz spectrum[14].

This specification has several advantages. The first is the already widespread adoption of the standard making the hardware very common. The second advantage is that it can use the TCP/IP protocol that is already commonly used on networks. This allows the system to easily integrate into an existing TCP/IP network, only requiring the addition of an 802.11b/g router if one does not exist. If the router exists, it is only a matter of configuring the endpoints and server to use the existing network, thus eliminating much of the installation costs described earlier. A third advantage available with the use of 802.11b/g systems is data encryption. With hardware support, the data transmitted can be encrypted with either Wired Equivalent Privacy (WEP) in 64 or 128 bits or Wireless Application Protocol (WAP). One last advantage is data integrity. The TCP/IP protocol will resend a packet until the endpoint receives the data. This is a big advantage over a protocol such as UDP.

The 802.11b/g specification definitely has the bandwidth required for the application, but the range needs to be considered. With low gain antennas, the range of the selected hardware was verified experimentally over 40 meters. This is more than adequate for the specified range and with the use of commercially available high gain or directional antennas, this range can be increased to much longer distances. As seen in Table I, 802.11b/g allows for the greatest range and data rates, thus making it the best solution for this application.

In an effort to reduce development costs, a commercially available, embeddable 802.11b/g radio was desired. Several options were identified from manufactures such

Table I. Wireless Specifications Comparison

Wireless Type	Data Rate	Range	Scalability	Cost
UART Replacement	115 kBits	Varies	Low	Low
ZigBee	250 kBits	20+ m	High	Low
Bluetooth	1 Mbits	10 m	High	Medium
802.11 b/g	11/54Mbits	40+ m	High	High

as DPAC Technologies and Lantronix. These companies provide a small module intended to be interfaced to embedded processors. Since these are embedded units, they are based on the 802.11b standard to help with power management.

The Lantronix units only interface by serial connections, thus limiting the available bandwidth to the speed of the serial connection[15]. These connections do not normally allow sufficient bandwidth for this application. DPAC Technologies has two different units available, the first unit also uses serial connections, but the second unit interfaces over a 10BaseT network connection[16]. This allows data transmission of 10 megabits per second, not a significant bottle neck to the 802.11b radio that is also contained in the unit. This DPAC radio was selected on the basis of providing the bandwidth required, over a standard that is in widespread use and with range acceptable for the application. The DPAC radio can be seen in figure 5.

F. Endpoint Arrangement

There were two possible options for the configuration of the system. The first option was to provide each endpoint with its own radio system and have them communicate through the radio to a central router and server. This arrangement is suitable for either a limited number of endpoints or endpoints that are spread out over a large

Fig. 5. Selected DPAC Wireless Module

area. A second configuration utilized a group of endpoints communicating with a hub that is equipped with a radio that then communicates with the central router and server. This option was suitable when several endpoints are located close enough together to make running the network cables from the endpoints to the hub affordable. Both of these configurations are diagrammed in figures 6 and 7.

G. Prototype Hardware

With the completion of the development of the breadboard hardware, the system was implemented in a prototype. In figure 8 one can see the resulting system.

In figure 9 one can see the processing platform and analog to digital conversion circuitry. Figure 10 shows the multi-mode switch hardware.

The last figure, figure 11, shows the implemented power circuitry and wireless module.

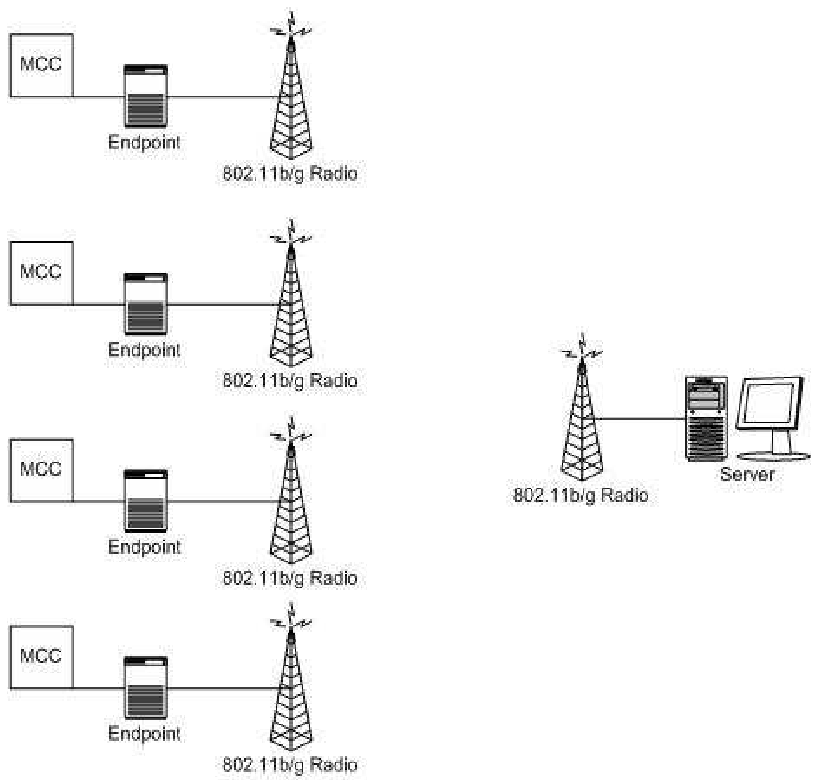

Fig. 6. Distributed Endpoint Arrangement

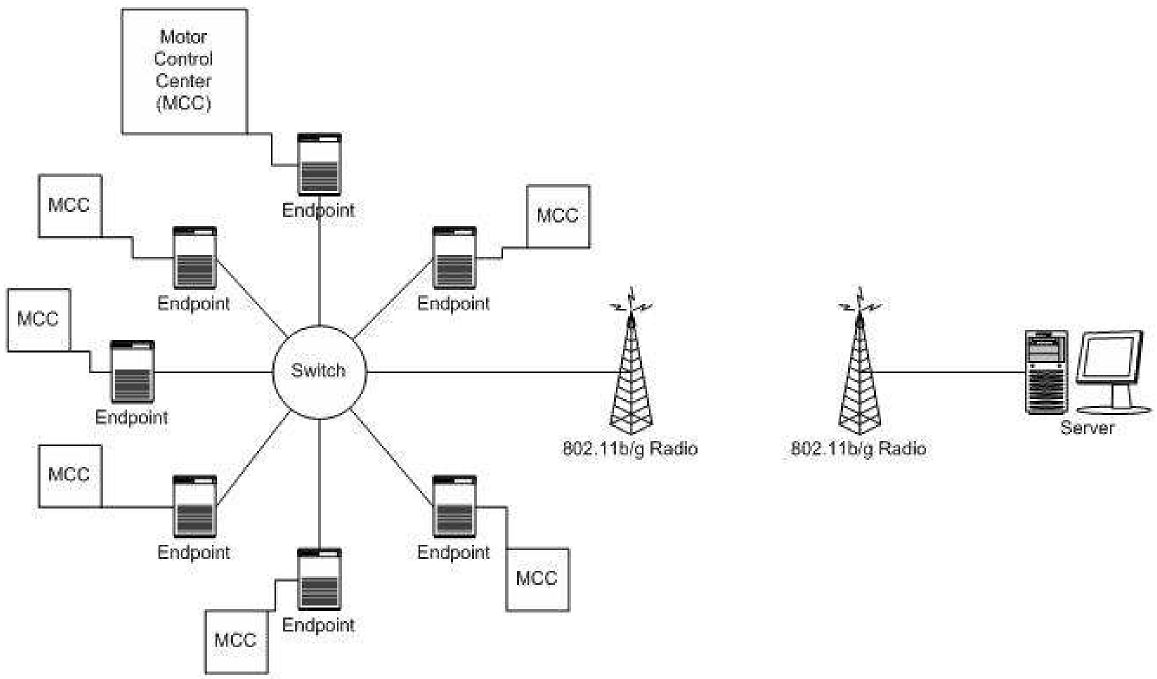

Fig. 7. Clustered Endpoint Arrangement

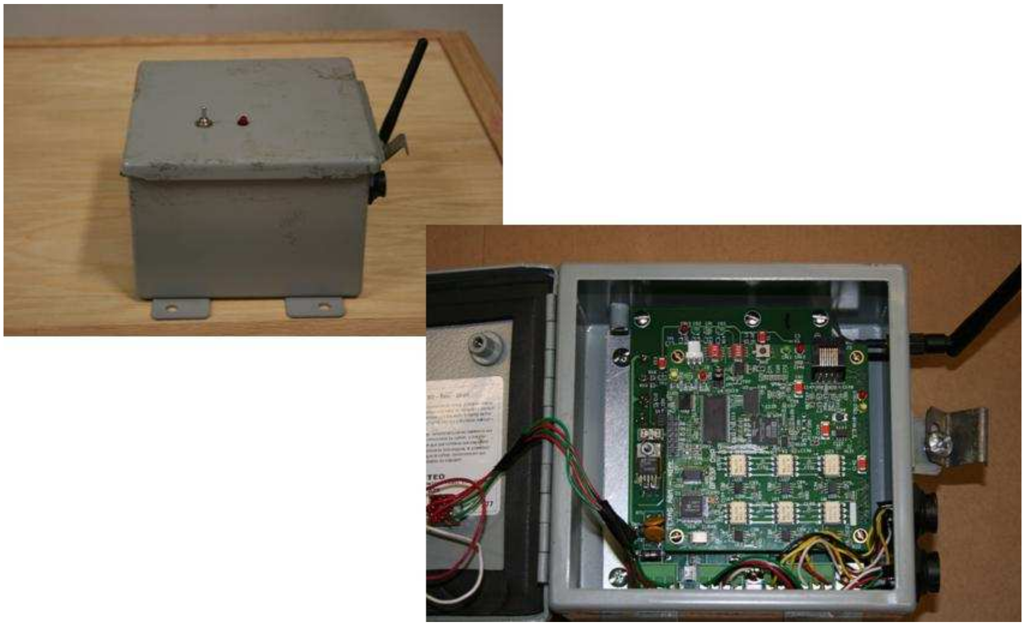

Fig. 8. Prototype Hardware in Enclosure

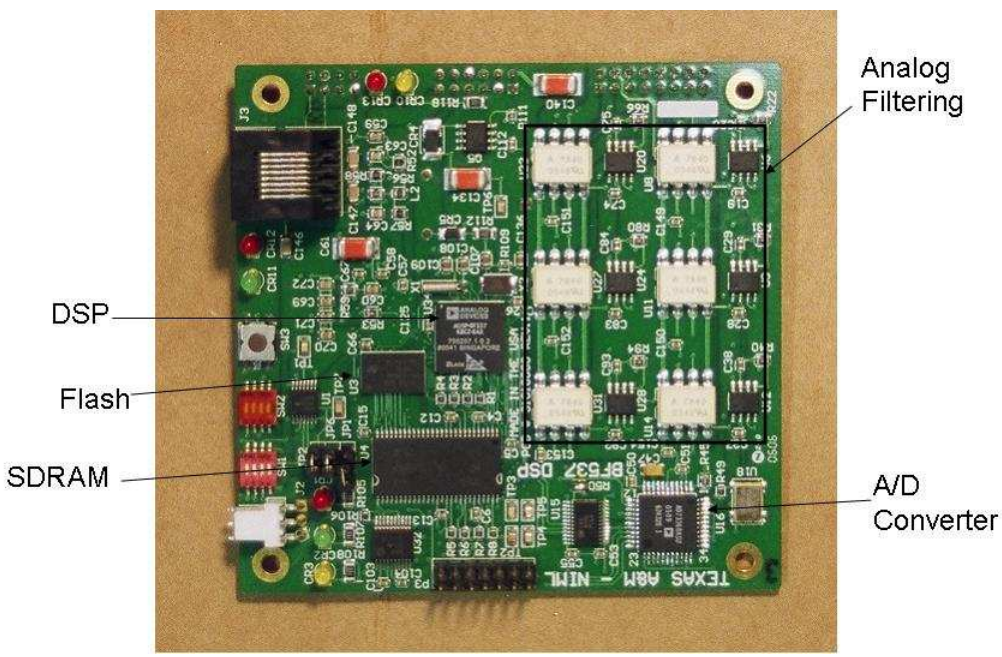

Fig. 9. Prototype Processor Board

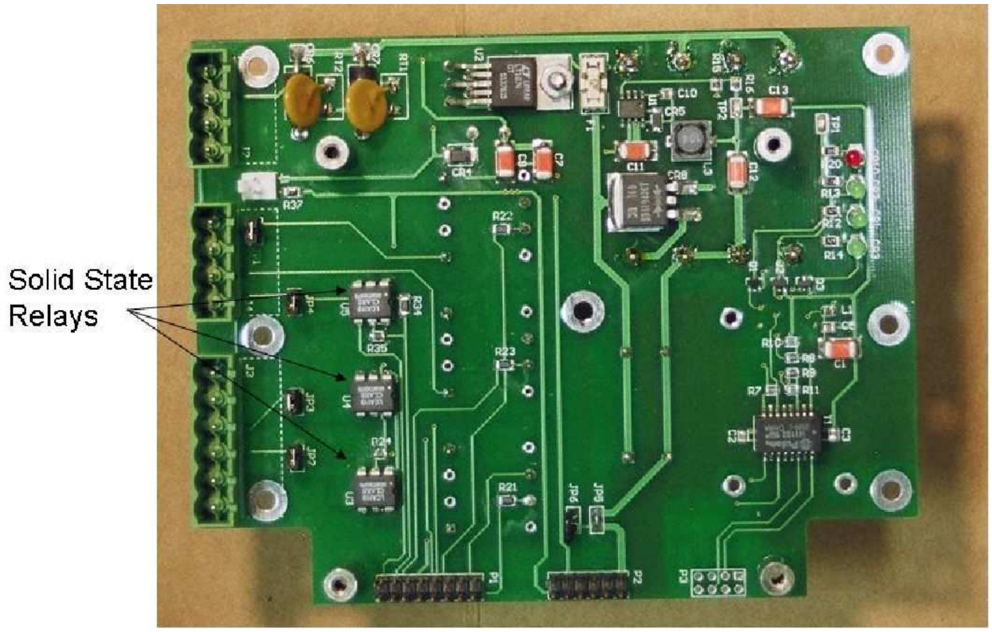

Fig. 10. Prototype Hardware Multi-Mode Switch

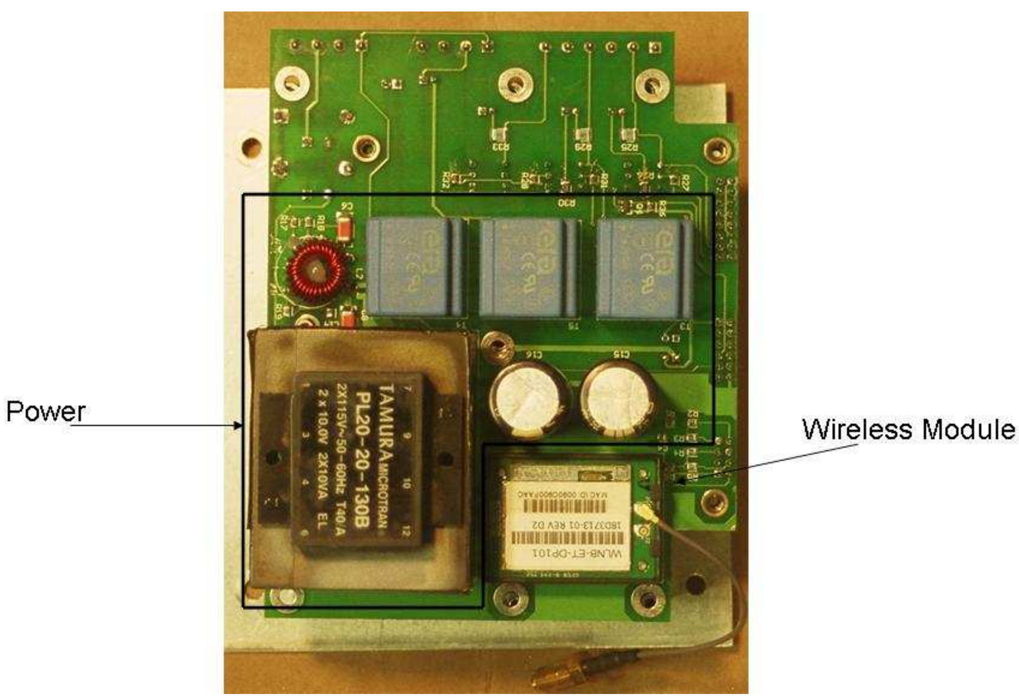

Fig. 11. Prototype Hardware Power and Wireless Board

CHAPTER III

SOFTWARE DESIGN

A. Endpoint Software

The endpoint software is built around the Analog Devices VDK kernel. This is a multi-threaded kernel that provides the necessary support for the Ethernet interface. This support comes in the form of driver software and a TCP/IP stack. Upon powering up, the kernel begins executing a system boot thread which handles the initialization of the system. This requires that the Blackfin, upon powering up, configures the AD73360 and the network interface. The configuration of the AD73360 requires setting the sampling rate to 8 kHz, turning on all six channels, and setting the programmable gain to zero. The configuration of the network interface requires setting the Blackfins IP address and establishing the link to the rest of the network. If the DPAC radio is to be used, it is pre-configured through a web browser on a PC and is not setup by the DSP.

Prior to the initialization of the AD73360, several data buffers are setup to store the collected data. This includes a buffer for the sampled noise, one for the startup data, and two large buffers for the steady state data. At the conclusion of the configuration of the AD73360, the interrupt controller is configured to interrupt when the AD73360 has completed the transfer of one set of samples. Each set of samples contains one sample from each of the six channels. At the conclusion of the interrupt, execution returns to the primary thread. After the AD73360 and interrupt controller are configured, the network is configured. This sets the Blackfins IP address, stores the MAC address for later use, and establishes the link to the rest of the network. Upon the conclusion of the network initialization, the boot thread starts a thread that will handle requests for data from the main server. After this thread is successfully started, the boot thread terminates. This process is outlined in figure 12.

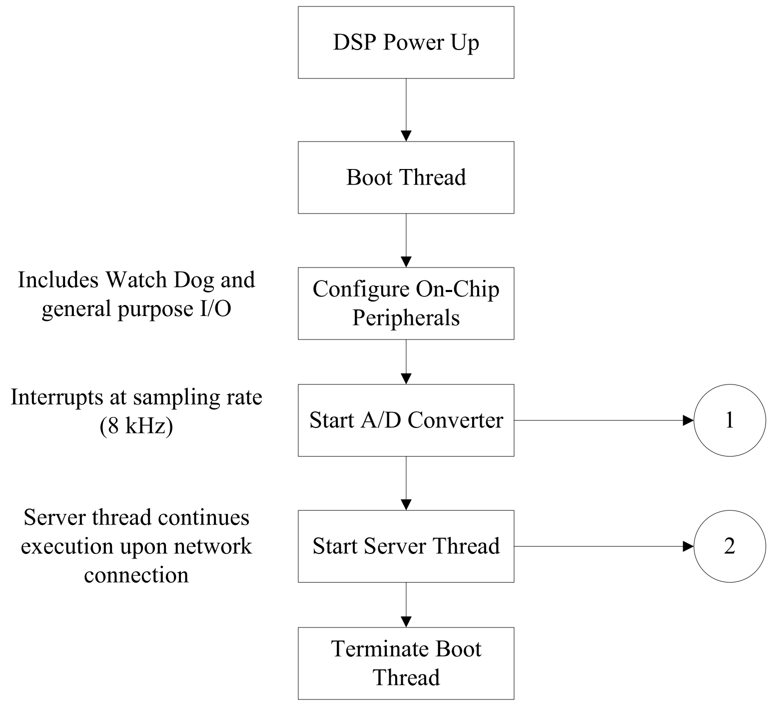

Fig. 12. Bootup Flowchart

The interrupt routine for the DSP handles receiving data from the AD73360 and determining the state of the equipment being monitored. When the AD73360 has completed the conversion of one sample for each of the six channels, the data is transferred over the serial interface to a DMA buffer on the Blackfin. When this DMA buffer is filled with the six samples, the interrupt is triggered.

When the interrupt routine executes, the routine first indicates that the interrupt has been handled. This allows the kernel to execute the routine the next time the buffer is full. The next step is to update several pseudo RMS values. These values represent a metric of the strength of the incoming signals for the three current channels and are calculated with the following formula (3.1):

$$RMS_X[n] = |Y[n]| + RMS_X[n-1] - \frac{RMS_X[n-1]}{N} \tag{3.1}$$

X in the expression is A, B, or C depending on the current channel being updated. The value of N determines how many points the formula uses to approximate the RMS of the signal. The input to the formula is the data point, Y, at sample n. The final term subtracts off 1/N of the previously known RMS value. This results in a smoothing of the incoming data, reducing the effects of noise.

If a limited number of samples have been collected, the interrupt routine will calculate a bias value for the signal. Through the use of the following formula (3.2), a value is determined to correct the above formula with. This reduces the effect of the bias variation present in the signals and allows the code to be used independent of calibration.

$$BIAS_X = \frac{RMS_X[n] + BIAS_X - \frac{BIAS_x}{N}}{N} \tag{3.2}$$

If a large number of samples have been collected, the above RMS value is compared with the pre-determined thresholds to determine the equipment state. If the equipment is in the noise state, the data is copied to the noise buffer. In this state, the DSP has signaled to the relays to configure the bridge for the startup signals. Upon detection of the startup, data is stored in the startup buffer. With the use of this two buffer system, the union of the buffers results in a continuous, complete view of the startup of the equipment. When the startup occurs, the data is stored in the startup buffer for a duration of five seconds. At the end of the five seconds, the DSP signals the relays to switch into the bridge for the steady-state signals. Upon conclusion of the five seconds, the equipment should have reached steady state and the data is stored in a steady state buffer until a shutdown is detected and the process repeats itself. The branches of logic used in the interrupt routine are diagrammed in figure 13.

The DSP's server thread, after successfully starting, will wait for a connection on the Ethernet interface. Upon this connection a hand shake with the server will occur to verify that the endpoint is a valid choice to poll. When the hand shake is complete, the DSP will perform a check of the equipments current state. In the unlikely event that a startup occurs and a polling request is made before there is a complete steady state buffer, the DSP will wait for the buffer to be filled before handling any of the data transfers. This ensures that the data collected is continuous. If the equipment is shutdown, this waiting is bypassed.

Once this check is complete, the DSP will send a single packet of data to the server. This packet indicates that the information contained is concerning the startup data and how many startups have occurred since the last polling. If no startups have occurred, the DSP will wait for an acknowledgment from the server that the packet was received. If startups have occurred, the DSP will immediately begin transferring the most recent noise and startup data. The server will see this as one continuous buffer as opposed to the two buffers that the DSP is storing. When the DSP finishes this transfer, it will await acknowledgment from the server. After the acknowledgment is received, the DSP will transfer another packet indicating if it has any steady state data to transfer. If there is no data available, the network connection will be closed and both the DSP and server will return to normal operation. If there is data to transfer, the data will be sent, after which, the connection will be closed. The DSP data transfer is diagrammed in figure 14.

With this logic, if there is no data to send at all, two packets will be sent. The first says there are no startups; the second says there is no steady state data. If there is no startup data, the first packet will be sent indicating there are no startups, the

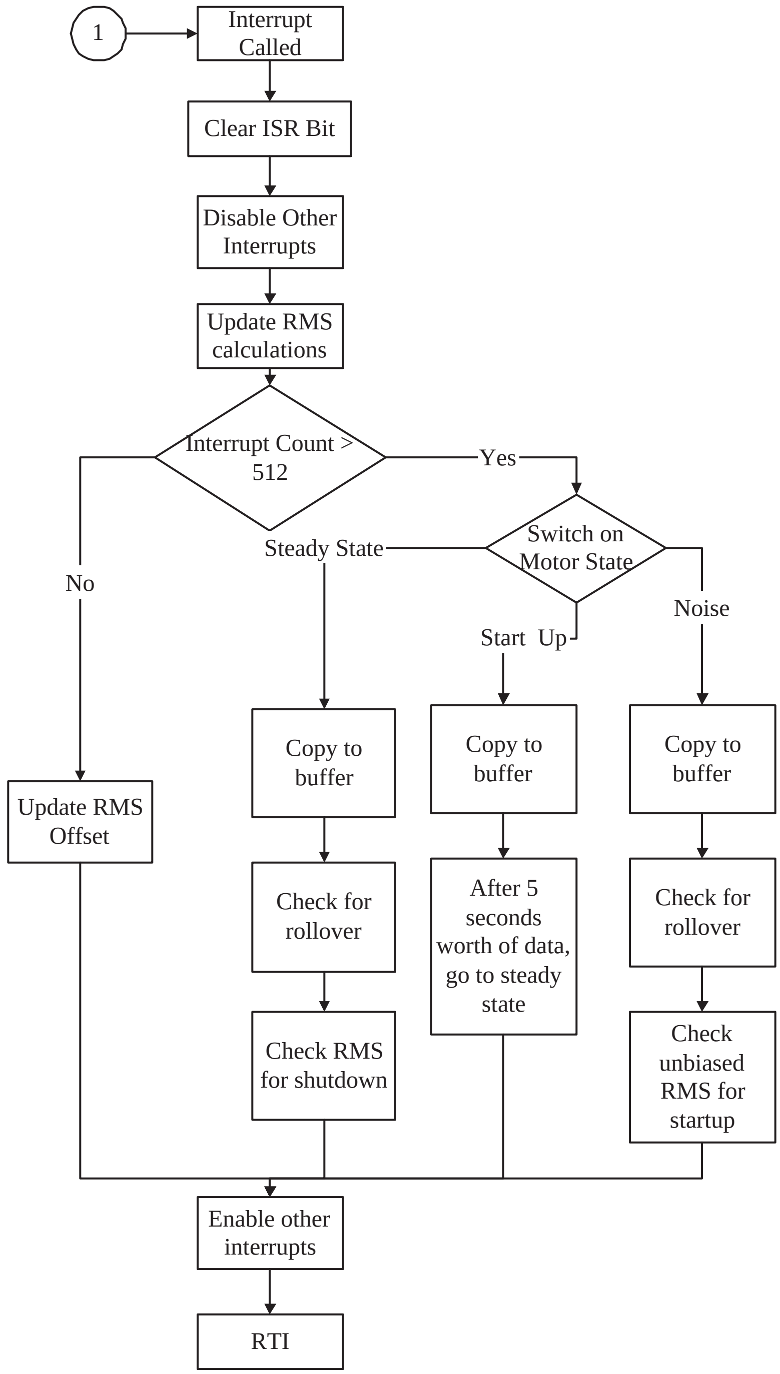

Fig. 13. Flowchart of the Interrupt Routine

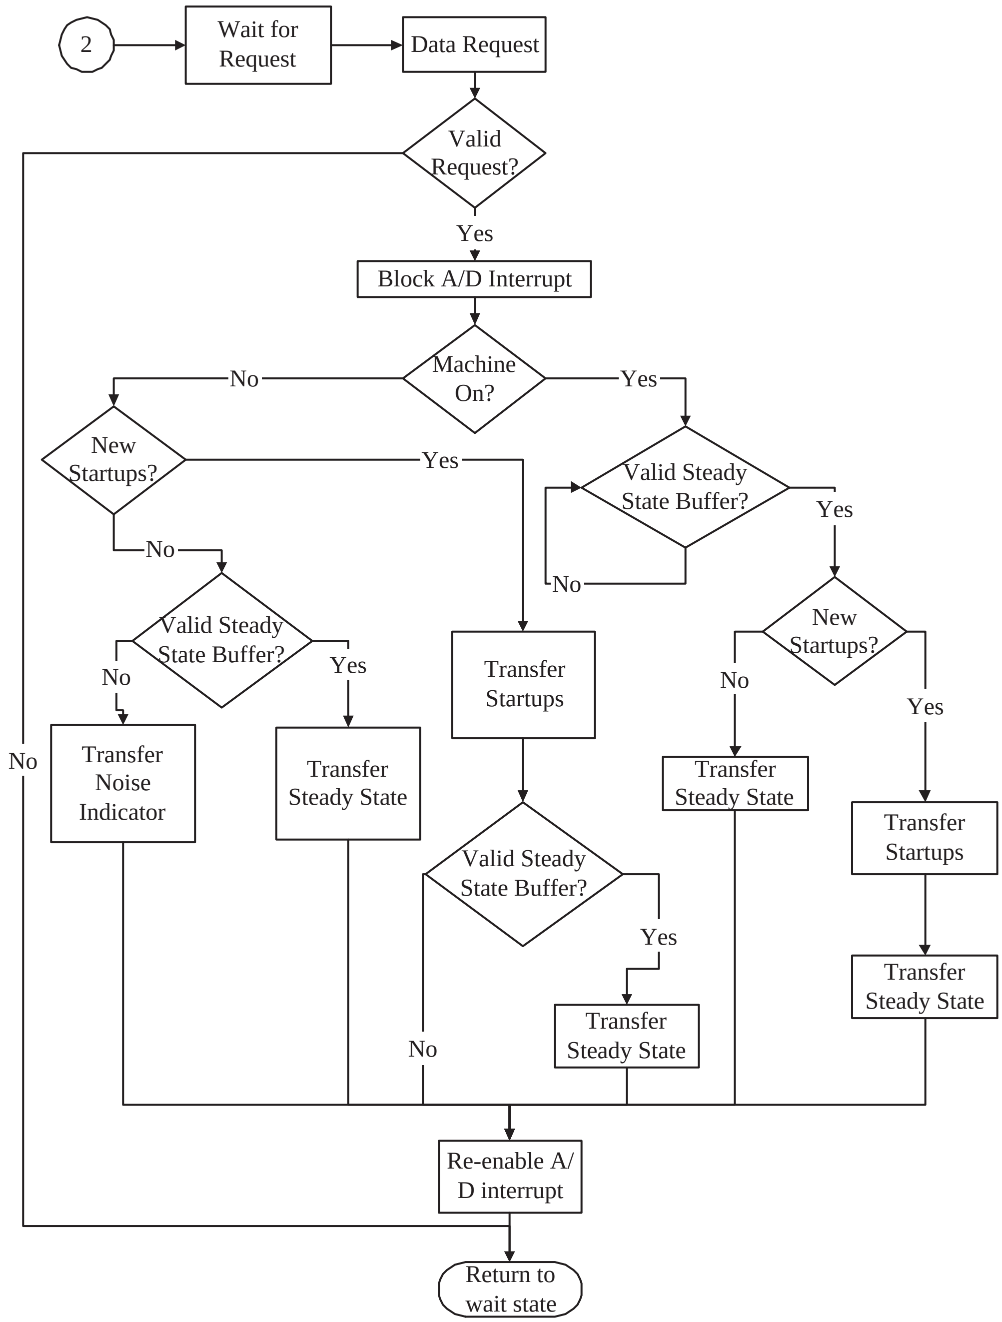

Fig. 14. Flowchart of the DSP Data Transfer

server will send the acknowledgment, and the DSP will transfer the steady state data. This order of operations will occur most often.

B. Server Software

The server is comprised of several modules. The first module handles the polling of the endpoints for data. A second module performs the actual processing of the data to determine the condition of the equipment being monitored. A third module that resides either in the processing server or another PC makes the equipment condition available through a web interface. The development of concern is the module that poles the endpoints and how this module interacts with the rest of the server modules.

Upon booting, the server code parses a configuration file. This file contains the needed information to uniquely identify all of the endpoints that are to be polled. The endpoints are then polled sequentially starting with the first node on the list. The server will first establish a connection with the endpoint at which time the endpoint will reciprocate with the hand shake packet. When the handshake packet is acknowledged by the server, the server then waits to receive the first packet from the DSP which indicates the availability of transient data. If transient data is available, the number of transients available from the endpoint is recorded. The sampled data is then received with the necessary byte ordering accommodated. A file is then created with the IP address of the endpoint and a transient indicator for the file name. The number of transients is then written to the file, followed by the transient data. If no transient data is available, a zero is written to the file to indicate to the other server threads that the transient data for the endpoint was polled but no data existed.

At the conclusion of receiving the transient data, the server sends another acknowledgment to the endpoint. At this point, the endpoint sends a single packet that indicates if there was steady state data available. If there is data available, the endpoint will send the data without waiting for any signals from the server. If there is no data available, the endpoint will close the network connection. At the conclusion of receiving the steady state data, a file is created with the IP address of the endpoint and no additional markers. The steady state data is written to the file and the loop repeats itself. If there was no data received, a zero is written to the file to indicate there was no available data. When the loop repeats, the configuration file is indexed to the next node in the list. The number of nodes scales easily by adding or deleting an entry from the configuration file.

CHAPTER IV

SYSTEM VERIFICATION PROCEDURES

A. Wireless Performance Determination

1. Transfer Times

Evaluating the transfer times occurred in two locations. The first location was again the lab facility in the WERC building. This provided a best performance baseline as this room was relatively clean from interference and the range was very short. The second testing location was in a power plant on the Texas A&M campus. This allowed the system to be tested in an industrial environment similar to the targeted environment of the system. This experiment collected two kinds of data; transfer time and signal quality. In order to perform this experiment, the system was loaded with a program that only sampled for steady state data. This allowed the experiment to be performed without connecting to any mechanical system. Since the quality of the sampled data was not relevant to the results, this was not a problem and allowed the experiment to be performed in the power plant without shutting down anything in order to connect the system.

2. Wireless Range

The evaluation of the wireless ranges often occurred at the same time as the transfer times were being tested. Therefore, this occurred in the WERC building and the physical plant. To evaluate the wireless range, the first test was done in a hallway of the WERC building and required separating the endpoint from the router and measuring the distance of separation. By conducting this test in the WERC building first, there was some confidence in having a low interference environment and thus measuring the maximum capabilities of the device.

The second set of tests for the wireless range were done in the physical plant and again required separating the endpoint and the router and measuring the distance. By doing this test in the physical plant, the performance was representative of what could be expected in an industrial environment. In figure 15 the layout of the physical plant experiments are shown. Each point labeled is the location of the laptop that was running the server software.

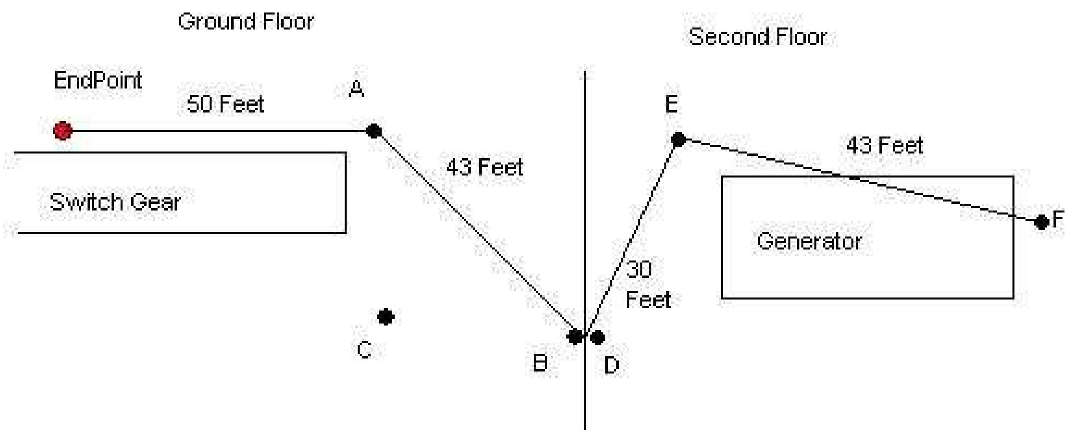

Fig. 15. Layout of Physical Plant Experiment (not to scale)

B. Data Conversion Verification

1. Channels Sampled Verification

One of the first things to be verified was that the minimum number of channels were sampled. For this design the minimum number is five channels, with six being desired. This requires a data set to be sampled and insure that all channels are sampled with no phase offsets occuring.

2. Verification of Sampling Rate

The second quantity to be confirmed is the sampling rate. This is to insure that the data converter is performing as expected. By assuming a sampling rate and calculating the power spectrum density, one can observe which frequency the carrier signal is operating at. In the United States, this frequency should be 60 Hz.

3. Dynamic Range

This procedure was required to determine and verify the effective dynamic range of the A/D converters for both the steady state and transient bridges. This was conducted on two test beds. The first test bed consisted of an unloaded one horsepower motor. This was powered from the utility grid, with no conditioning. The CTs used were shunted and already in place on the lines. The signals were recorded on a desktop PC running Windows XP and the server code described in the preceding chapter. After a thirty second sample had been collected on the PC, Mathworks' MATLAB was used to calculate the power spectrum of the signal.

To calculate a signal to noise ratio, the values at the carrier frequency and the noise basin of the power spectrum are of interest. The output of the MATLAB function is in dB/Hz which is not a viable unit for the calculation. These values can be converted back to a unit/Hz value by dividing by twenty and then taking ten to the power of the result. The signal to noise ratio can then be calculated as the carrier component in unit/Hz divided by the noise basin, also in unit/Hz. The units cancel out and the resulting value can be converted to dB using the standard calculation (20*log10(x)). This value can then be used to determine the effective resolution in bits.

This procedure was also done for the transient bridge. It is important to remember that the transient bridge was designed to capture a non-periodic event that decays to a very low amplitude signal, very quickly. This procedure was also performed on a three horsepower motor driving a centrifugal pump. Both of these motors are located in the lab facility in the Wisenbaker Engineering Research Center (WERC).

4. Startup Capture

The second procedure verified the accuracy of the startup capture routine. This experiment was again performed on the one and three horsepower motor test beds in the WERC lab space. Within the software there were several calls to general purpose I/O pins that were used for debugging purposes. Several of these pins were connected to light emitting diodes (LEDs) on the development board that provided a visual indicator of an event. One of these LEDs was used to indicate the state of the motor signal. If the LED was illuminated, the motor was running, if the LED was off, the motor was shut down. This was used as the indicator for the success of the startup capture. If the LED came on when the motor was turned on, the startup had been successfully detected.

5. Shutdown Trigger

The last piece of the data acquisition to verify was the detection of the motor shutting down. These experiments were again performed on the one and three horsepower motors in the WERC facility. This experiment was often conducted in conjunction with the motor startup experiment described above. When the motor was started, the indicator LED would come on, indicating a startup. When the motor was shut down, the indicator LED would shut off, indicating the motor had shut down.

CHAPTER V

SYSTEM PERFORMANCE

A. Benchmarks

The following sections pertain to the system performance as measured in the laboratory setting. These tests were run in the Wisenbaker Engineering Research Center (WERC) at Texas A&M University.

1. Wireless Performance

The evaluation of the wireless performance for a benchmark was rather informal. The main focus was to determine the best possible transfer time under the best circumstances. With easy access to the physical planet, there was less focus on these benchmarks. Having the end point near the wireless router, about 10 feet of separation, this transfer time was determined to be nine seconds. To determine a maximum range estimate, the system was setup in the hallway of the WERC bulding and a connection was attempted. The maximum distance for achieving a connection was about 270 feet which is the length of the hallway. None of the signal quality metrics were recorded.

2. Data Conversion Performance

a. Number of Channels Sampled

Figure 16 shows the data collected on six channels in the WERC lab space. One can observe that all of the channels have the correct phase offset. The sixth channel is not connected to a source so it does not display a meaningful signal, but it has been collected.

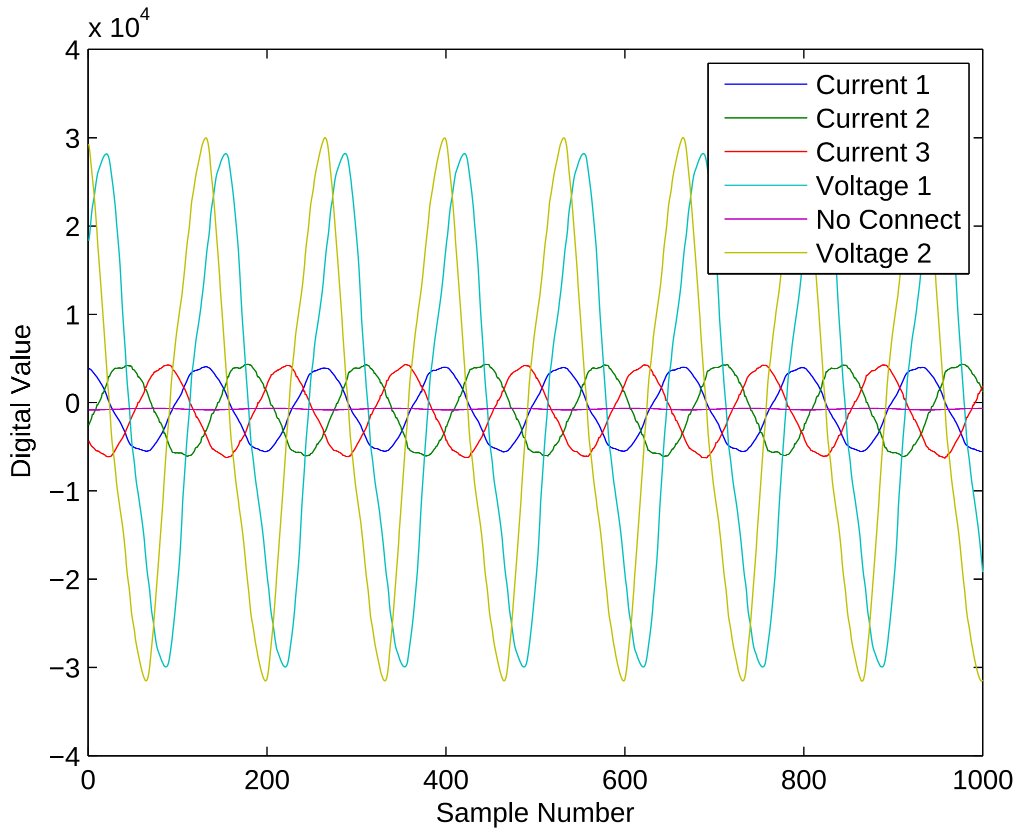

Fig. 16. Channels Sampled in the WERC Laboratory

b. Actual Sampling Rate

The power spectrum density (PSD) for a sample set configured for an 8kHz sampling frequency is shown in figure 17. One can observe that the principle component does occur at 60 Hz. This confirms that the sampling rate is in fact at 8 kHz.

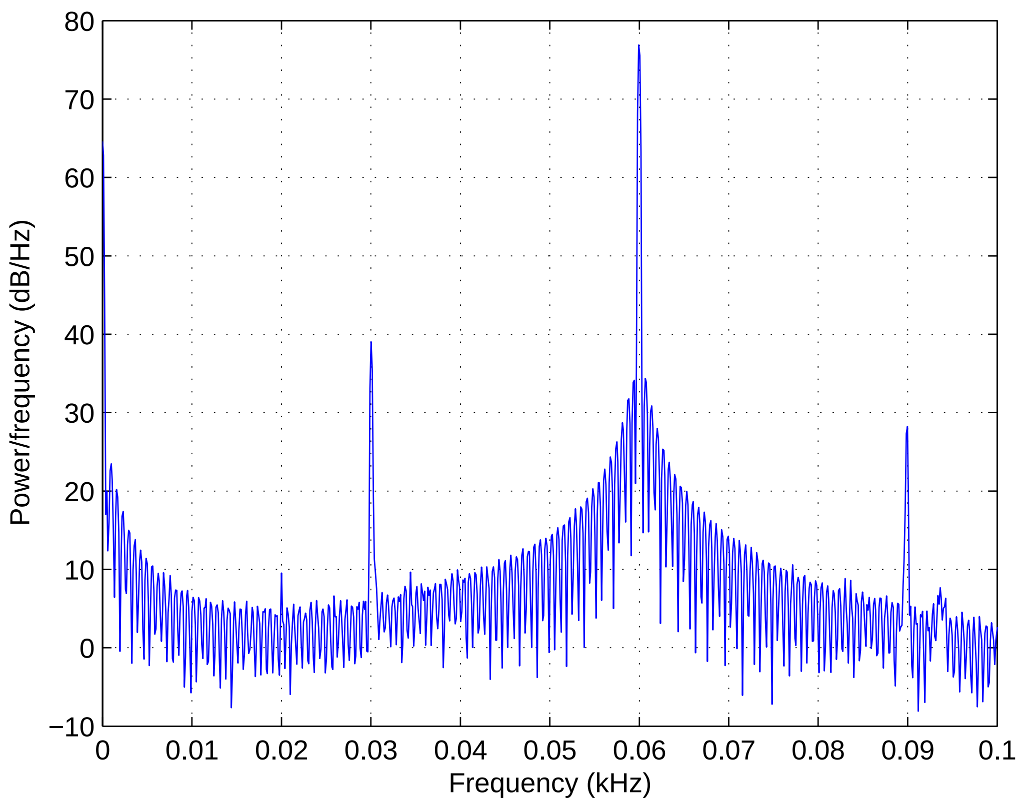

Fig. 17. Confirmation of Sampling Frequency

c. Steady State Sampled Signal Quality

In this section, the evaluation of the sampled signal quality in the system through the steady state voltage bridge is discussed. In equation 5.1 a relationship between the signal to noise ratio in dB and the effective resolution in bits(N) is expressed [17].

$$SNR(db) = 6.02N + 1.76 \tag{5.1}$$

The 60 Hz component of the signal can be seen as about 92 dB/Hz and the noise component is approximately at -15 dB/Hz in figure 18. Using the procedure outlined in chapter IV, this yields a signal to noise ratio of 107 dB. Using equation 5.1 the effective resolution of the steady state voltage signal in the lab is determined to be approximately 17 bits from the data in figure 18. This is better performance than the 16 bits that the A/D converter is designed for and meets the requirement of 14 bits.

In figure 19, the 60 Hz component has a strength of 78 dB/Hz and a noise level of -10 dB/Hz. The signal to noise ratio for the current signal is 88 dB. Using equation 5.1 the effective resolution of the steady state current signal in the lab is determined to be approximately 14 bits from the data in figure 19. This is also better than the 12 bits claimed as the minimum resolution of the A/D converters and meets the requirement.

d. Transient Signal Acquisition

This section evaluates the effectiveness of the system to detect and sample the in rush signal during the powered equipment startup. Also discussed is the system's ability to detect the shutdown of the powered equipment. Since the voltage applied to the system remains constant, only the current channels are handled differently for the startup event.

The data displayed in figure 20 shows the original pseudo RMS values for the startup data in addition to the unbiased pseudo RMS values for the same data set. It clearly shows the effect of the unbiasing algorithm. This also indicates how the unbiasing can allow a much tighter startup criteria to be set. Also clearly shown is

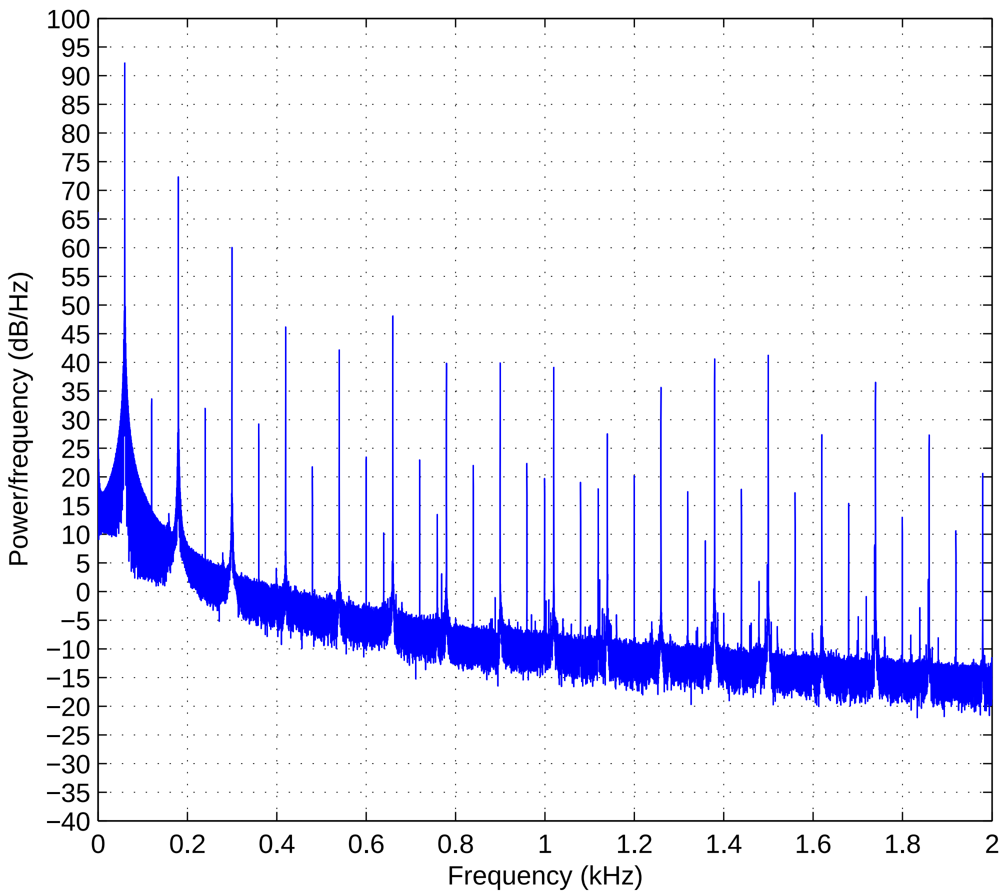

Fig. 18. Voltage Power Spectrum

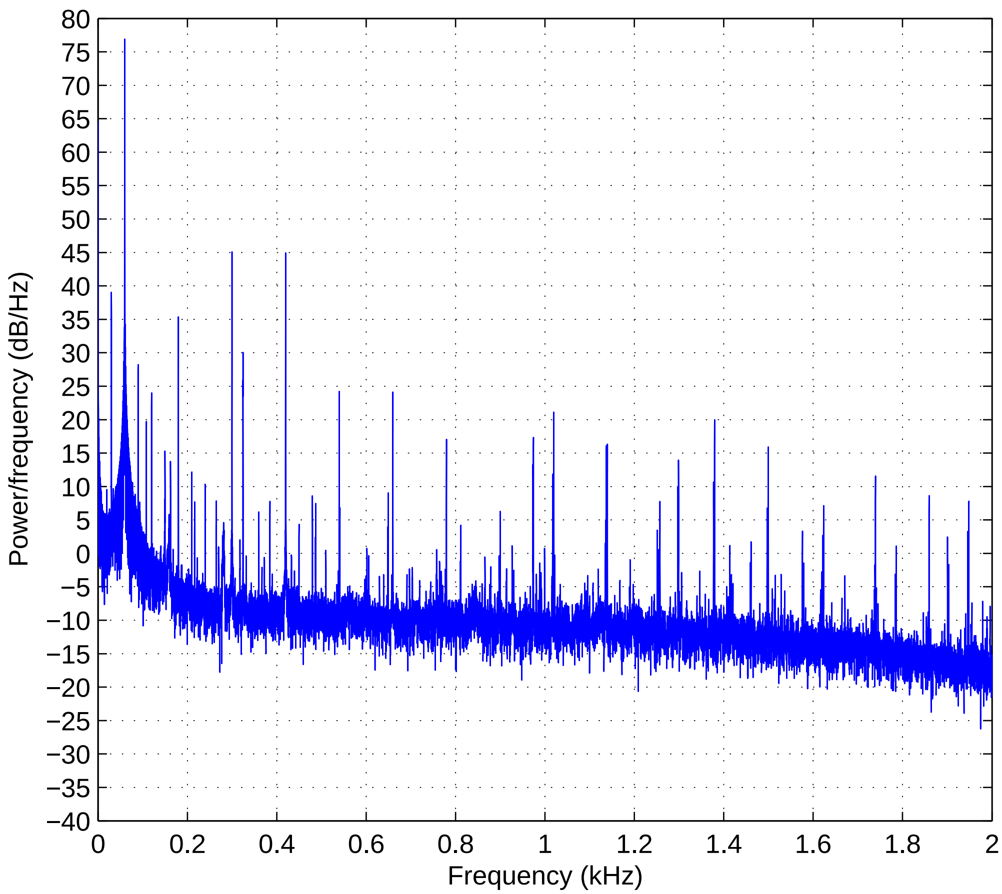

Fig. 19. Current Power Spectrum

the importance of removing the bias in each of the three current signals. Without removing this bias, the three channels vary between 50000 and 100000. With this much variation, it would be very difficult to set one level that one would be confident in working for all three channels.

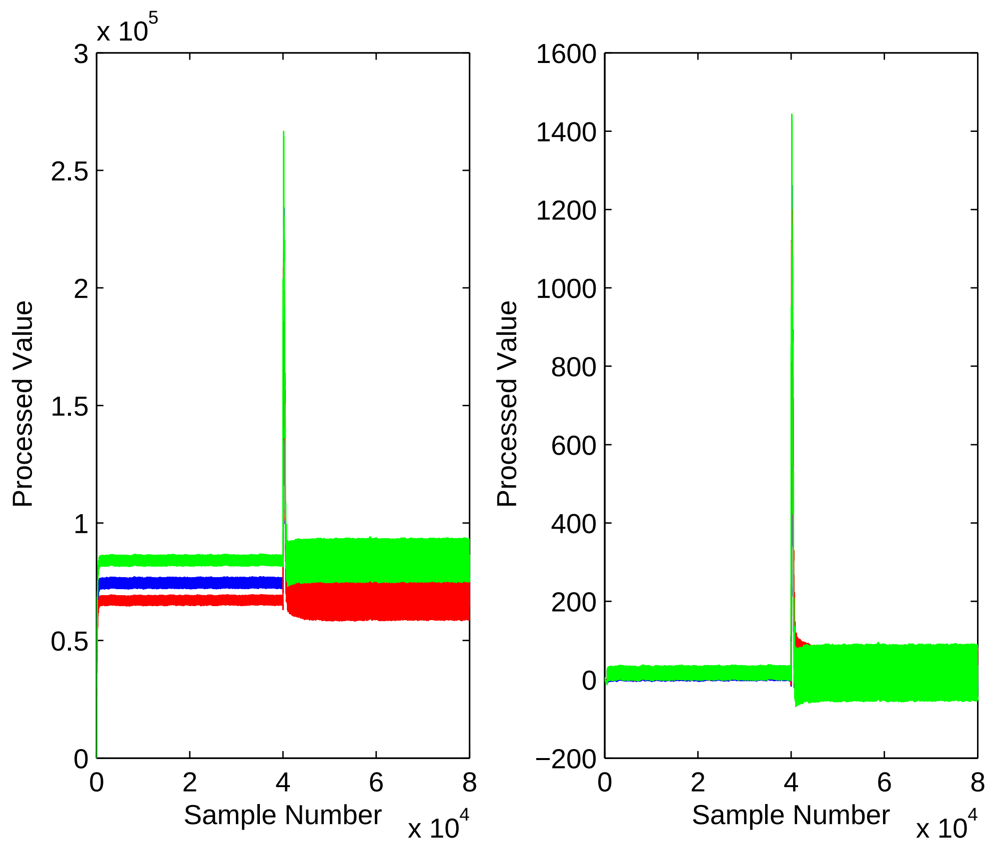

Fig. 20. Biased and Unbiased Current Signal

With the bias removed, one is now able to detect the startup. The data displayed in figure 21 shows the systems ability to capture the entire startup event. As designed, the actual event is located half way through the data set. This provides a clear picture of when the event has occured.

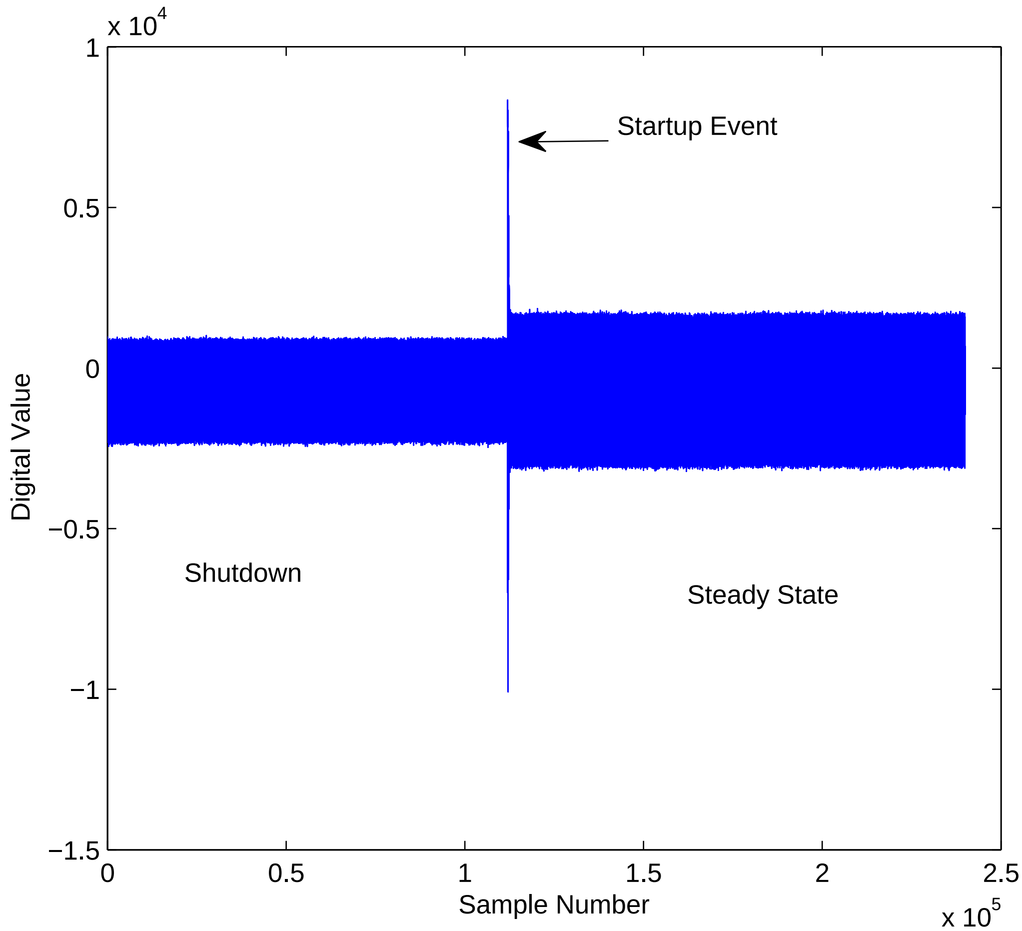

Fig. 21. Startup Signal on One Current Channel

Applying equation 5.1 to the startup data, effective resolutions can be determined. Since the startup signal is captured through a bridge designed to fit up to twenty times the steady state amplitude to the A/D converter, this signal has the potential to appear to be of much smaller amplitude. This will correspondingly decrease the effective signal to noise ratio.

There are two regions of interest in the startup signal to evaluate the effective resolution. The first region is the startup itself. This consists of the first group of oscillations where the amplitude is decaying to steady state. The second group is the steady state signal as seen through the transient bridge. Figure 22 shows the power spectrum of the entire startup data set and figure 23 shows the power spectrum of the steady state signal through the transient bridge.

By again applying equation 5.1, one can determine for these cases an effective resolution of 11 bits for the entire data set and 10 bits for the portion of the data that is at steady state. While these values are lower than the 12 bits promised by the AD73360, this is expected since the data is fit rather poorly to the device's input range.

B. Actual Operation

1. Wireless Performance

The actual performance of the wireless connection was evaluated in the physical plant. The initial evaluation of the system was done using the bread-boarded prototype. This system utilized a bread board attached to the Blackfin development board and the DPAC evaluation module connected to the development boards Ethernet port. This system demonstrated significant variations in transfer times and signal-to-noise ratios. When the prototypes were completed, they were also tested in the physical

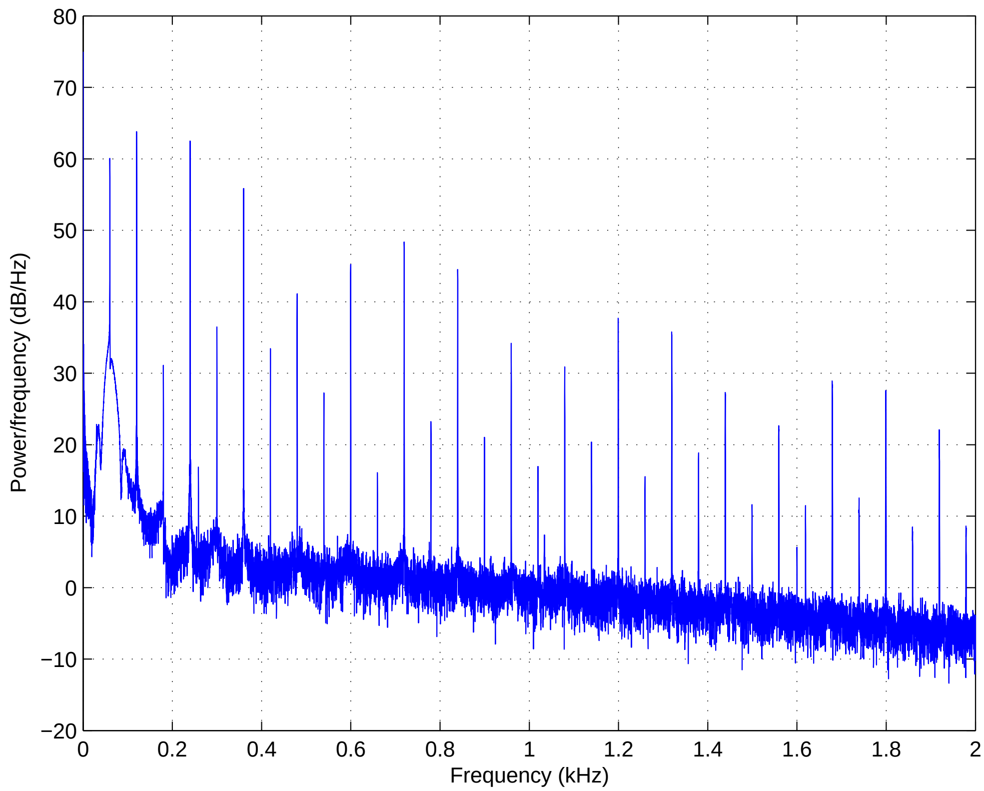

Fig. 22. Power Spectrum of the Startup Event

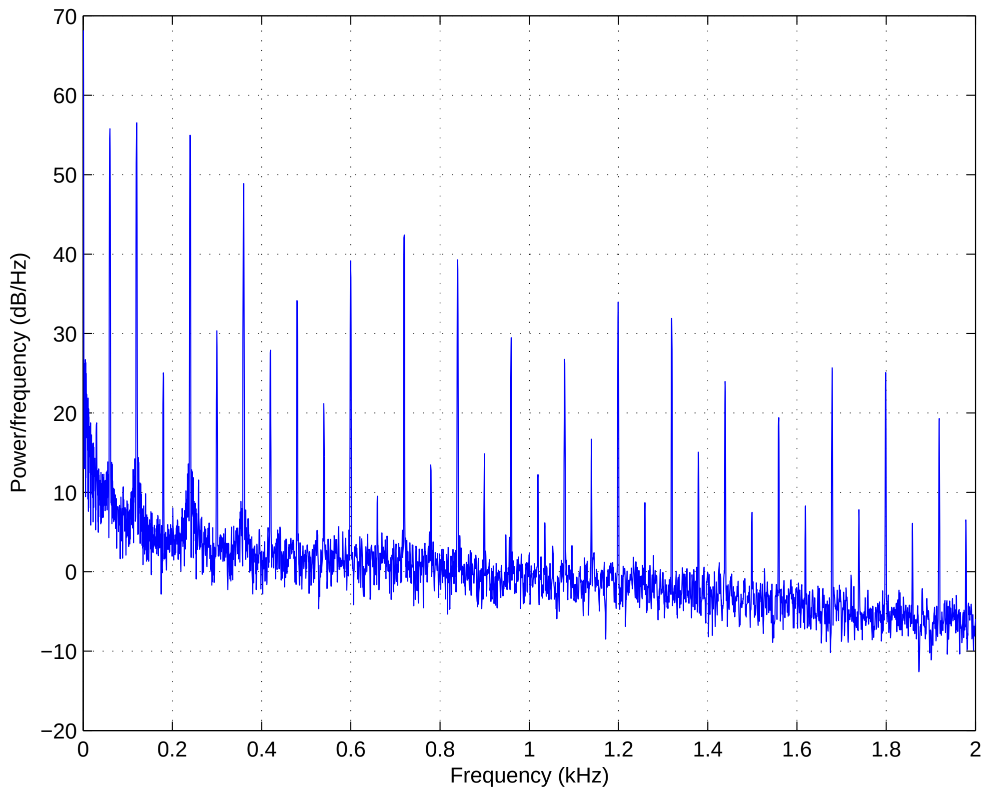

Fig. 23. Steady State Signal Through Transient Bridge Power Spectrum

plant and their results stayed much more consistent. Figure 15 diagrams the location of the test points and table II contains the resulting data.

Table II. Physical Plant Measurements

Point	Data Rate MBits/s	SNR dBm	Signal Strength dBm	Noise Level dBm	Transfer Time sec
A	11	33	-57	-91	15
A	11	35	-51	-86	11
B	11	30	-60	-91	18
B	11	19	-69	-88	10
C	11	33	-57	-91	-
D	11	32	-58	-90	17
D	11	20	-68	-88	11
E	11	16	-72	-89	11
F	11	15	-73	-88	10

It can be observed from the data that there is significant electrical noise in the physical plant, as expected. While this most directly effected the data rate, this rate still remained above the ethernet input to the module, thus maintaining the ethernet as the first bottle neck in the data transfer.

2. Sampled Signal Quality

a. Channels Sampled Verification

For the physical plant, the confirmation of the number of channels sampled is again conducted visually. In figure 24 one can observe the five sample signals and the sixth channel not connected.

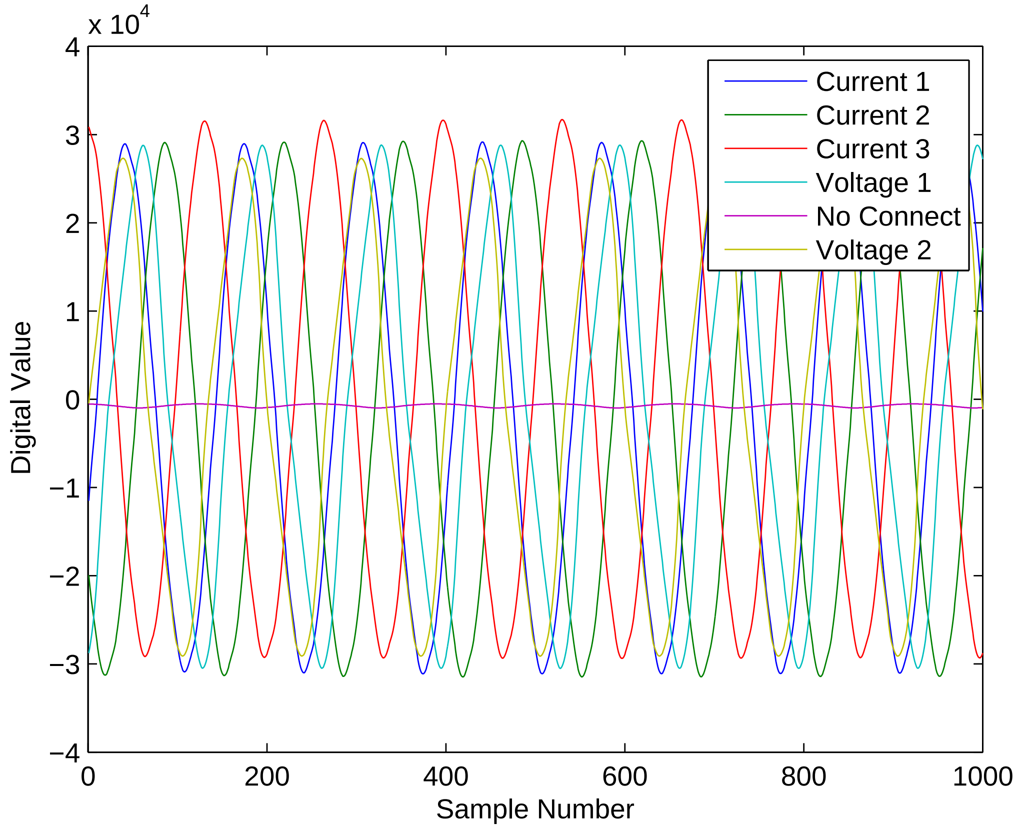

Fig. 24. Channels Sampled in the Physical Plant

b. Actual Sampling Rate

Confirmation of the sampling frequency is again conducted by inspecting the spectrum density resulting from an assumed sampling rate of 8 kHz. With the principle component located at 60 Hz, as expected, the sampling rate is confirmed. This can be seen in figure 25.

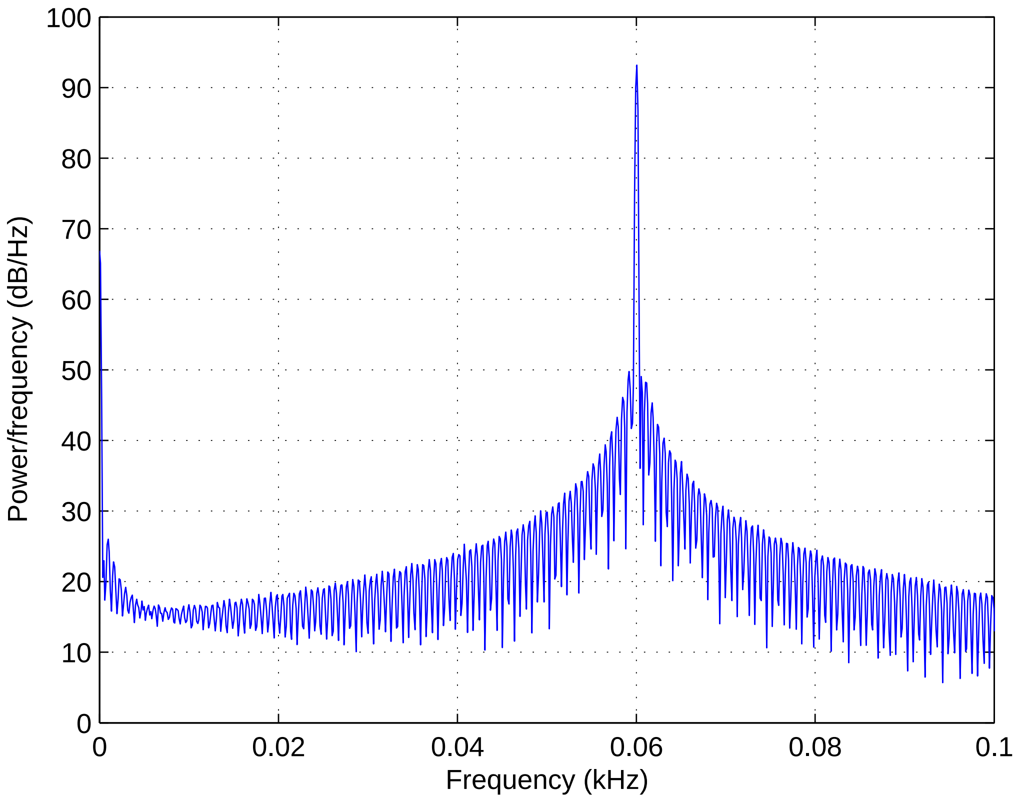

Fig. 25. Confirmation of Sampling Frequency in the Physical Plant

c. Steady State

Figure 26 shows the power spectrum of the sampled voltage signal as recorded by the endpoint located in the physical plant on the Texas A&M University campus.

Using a signal amplitude of 92 dB/Hz and a noise amplitude of -15 dB/Hz, one has a signal to noise ratio of 107 dB. Using equation 5.1, one can determine an effective bit resolution of 17 bits.

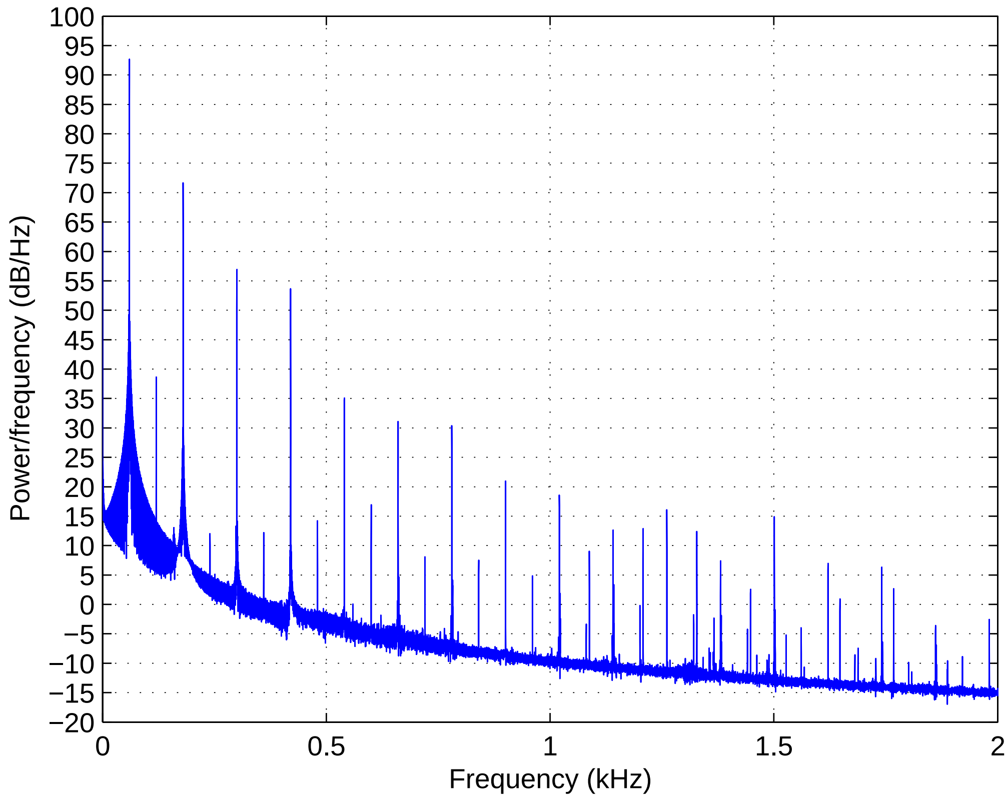

Fig. 26. Voltage Power Spectrum from the Physical Plant

Figure 27 shows the power spectrum of the sampled current signal as recorded by the endpoint located in the physical plant on the Texas A&M University campus. Using a signal amplitude of 93 dB/Hz and a noise amplitude of -15 dB/Hz, one has a signal to noise ratio of 108 dB. Using equation 5.1, one can determine an effective bit resolution of 17 bits.

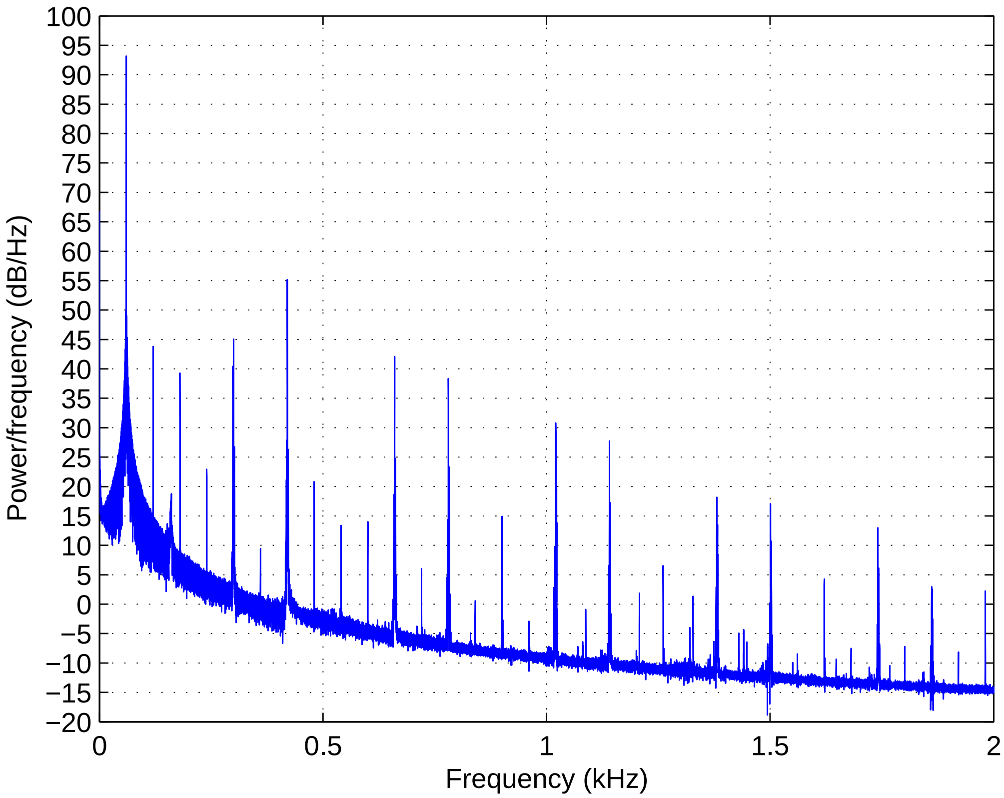

Fig. 27. Current Power Spectrum from the Physical Plant

Both the current and voltage signals demonstrate exceptional resolution. This is mostly due to the signals being fit to the maximum range of the A/D converters. The powered equipment in the lab was of a smaller scale, as a result, the current signals did not maximize the range of the A/D converters as well as the larger scale equipment in the physical plant.

d. Transient

At the time of writing, no start up data had been recorded at the physical plant.

CHAPTER VI

SUMMARY, CONCLUSIONS AND RECOMMENDATIONS

A. Summary

This thesis presented the development and performance evaluation of a data acquisition system for multi-phase electric machines. The system consists of an analog front end which includes a multi-mode switch and filtering. In addition to the analog front end, there is an analog to digital converter, a digital signal processor, and a wireless network interface. The data converter is capable of up to 64 kHz sampling on up to 6 channels. The wireless interface uses the 802.11b IEEE standard. The system connects to a server which receives the data for processing.

The evaluation of the system details the performance measurements in both the laboratory and the field. These measurements cover the effective resolution of the analog to digital conversion, the effectiveness of the multi-mode switch, and the wireless communications performance.

B. Conclusions

The work of this thesis demonstrated that it is possible to construct an endpoint using commercially available, off the shelf components that is capable of being used in continuous health monitoring of multi-phase electric machines.

C. Recommendations for Future Work

At the conclusion of the development process, several recommendations can be made for future work. To further improve the system, three components of the system can be changed:

- DSP
- Wireless
- Memory

The DSP selected was chosen to meet the requirement of the ethernet and the serial interfaces. Little attention was paid to the performance of the DSP. Since the Blackfin DSPs are fixed point devices, they are not capable of running the processing code currently running on the server. If a processor could be found that contained a floating point unit in addition to the required ethernet and serial interfaces, the processing could be off-loaded from the server and distributed to the individual endpoints. This would greatly reduce the bandwidth requirement of the wireless connection and increase the number of endpoints a server could handle.

There are several avenues available for improving the wireless subsystem. If a floating point DSP is available, one could change to one of the low bandwidth options discussed in this thesis, such as ZigBee. These units are much less expensive than the DPAC unit chosen. This would require the development of application specific software to handle the ZigBee network that would be created. Also, one would have to address the possible range limitations of this standard. A second option to improve the wireless subsystem is to upgrade to the 802.11g standard. This would provide additional bandwidth to an individual transmission, thus reducing the transmission time and increasing the number of endpoints that can be on a network. Another option for improving the wireless subsystem is to build an 802.11b/g radio from scratch. This has the potential to be less expensive in hardware than the DPAC module but would be a large development task.

The system that can be improved is a better evaluation of the amount of memory included with the endpoint. The design of this endpoint was based on an available reference design that included far more memory than what is used. It is probably safe to reduce this amount by at least half but the effect of this change would need to be propagated through the electrical and software design.

REFERENCES

[1] A. G. Parlos, K. Kim, and R. M. Bharadwaj, "Sensorless detection of mechanical faults in electromechanical systems," *Mechatronics*, vol. 14, no. 4, pp. 357–380, 2004.

[2] F. Eady, *Implementing 802.11 with Microcontrollers: Wireless Networking for Embedded Systems Designers (Embedded Technology)*, Burlingtion, MA; Newnes, 2005.

[3] P. Ferrari, A. Flammini, D. Marioli, and A. Taroni, "Ieee802.11 sensor networking," *IEEE Transactions on Instrumentation and Measurement*, vol. 55, no. 2, pp. 615–619, 2006, 0018-9456.

[4] C. P. Townsend, M. J. Hamel, P. Sonntag, B. Trutor, J. Galbreath, and S. W. Arms, "Scaleable, wireless web enabled sensor networks," in *2nd ISA/IEEE Sensors for Industry Conference*, 2002, pp. 172–178.

[5] R. E. Raygan and D. Callahan, "Wireless remote sensor data acquisition and management using linux os, murs and 802.11b," in *IEEE 55th Vehicular Technology Conference*, 2002, vol. 1, pp. 219–223 vol.1.

[6] M. Kohvakka, M. Hannikainen, and T. D. Hamalainen, "Performance of a pc/104 based wireless diagnostics system," in *The 29th Annual Conference of the IEEE Industrial Electronics Society*, 2003, vol. 1, pp. 860–865 vol.1.

[7] Crossbow Technology Inc., "Development kits, http://www.xbow.com/products /productdetails.aspx?sid=160," 2006.

[8] Honeywell International Inc., "Xbw-k2400 series, http://stevenengineering.com/ tech_support/pdfs/31wirexbw.pdf," 2006.

[9] Analog Devices Inc., "Ad73360," 2005.

[10] P. Su-Kang, L. Hong-Woo, C. Geum-bae, B. Hyung-Lae, L. Seong-Kil, C. Yong-Hae, and S. Jin-Youn, "A study on the reduction of in-rush current for energy saving of single phase induction motor," in *Proceedings of the Fifth International Conference on Electrical Machines and Systems*, 2001, vol. 1, pp. 67–71 vol.1.

[11] Microhard Systems Inc, "Mhx920a," 2006.

[12] C. Evans-Pughe, "Bzzzz zzz [zigbee wireless standard]," *IEE Review*, vol. 49, no. 3, pp. 28–31, 2003, 0953-5683.

[13] A. K. Arumugam, A. R. Nix, P. N. Fletcher, S. M. D. Armour, and B. S. Lee, "Scenario driven evaluation and interference mitigation proposals for bluetooth and high data rate bluetooth enabled consumer electronic devices," *IEEE Transactions on Consumer Electronics*, vol. 48, no. 3, pp. 754–764, 2002, 0098-3063.

[14] S. C. Wang, Y. M. Chen, L. Tsern-Huei, and A. Helmy, "Performance evaluations for hybrid ieee 802.11b and 802.11g wireless networks," in *24th IEEE International Performance, Computing, and Communications Conference*, 2005, pp. 111–118.

[15] Lantronix Inc., "Embedded device servers, http://www.lantronix.com/device-networking/embedded-device-servers/," 2006.

[16] DPAC Technologies Inc., "Wireless products, http://www.dpactech.com/ wireless_products/embedded-802.11-wireless-module.php," 2006.

[17] Robert Fugerer, "Simple techniques that lower cost and increase A/D converter resolution," 2003.

VITA

Douglas Andreas Goodsell

StarVision Technologies Inc.

1700 Research Pkwy Suite 170

College Station, TX 77845

Education

Master of Science, Mechanical Engineering, Texas A&M University, May, 2008

Bachelor of Science, Mechanical Engineering, Texas A&M University, December, 2002

College Teaching Experience

Spring 2005 Mechanical Measurements(MEEN 260), Texas A&M University, TA.

Fall 2004 Mechanical Measurements(MEEN 260), Texas A&M University, TA.

Spring 2004 Mechanical Measurements(MEEN 260), Texas A&M University, TA.

The typist for this thesis was Douglas Andreas Goodsell.