



Applying fiberoptic data links to instrumentation
by Michael Steven Beer

A thesis submitted in partial fulfillment of the requirements for the degree of MASTER OF SCIENCE in
Electrical Engineering
Montana State University
© Copyright by Michael Steven Beer (1980)

Abstract:

The subject of this thesis is the application of a fiber optic data link for use in instrumentation applications where conventional information transmission techniques cannot be used due to excessive electrical noise. A technique was developed for converting a direct current voltage to a pulse string whose frequency was proportional to the voltage and transmitting this pulse string over the data link.

This technique also allowed recovery of the information at the receiving end in digital form using a microprocessor-based data acquisition system. As a demonstration of the validity of this technique, sensors were constructed to enable use of the data link to measure alternating currents, direct currents, and temperature. Mathematical transfer functions were developed for each of these sensors to aid the microprocessor-based system in providing accurate measurements.

STATEMENT OF PERMISSION TO COPY

In presenting this thesis in partial fulfillment of the requirements for an advanced degree at Montana State University, I agree that the Library shall make it freely available for inspection. I further agree that permission for extensive copying of this thesis for scholarly purposes may be granted by my major professor, or, in his absence, by the Director of Libraries. It is understood that any copying or publication of this thesis for financial gain shall not be allowed without my written permission.

Signature

Michael Beer

Date

July 11, 1980

APPLYING FIBEROPTIC DATA LINKS TO INSTRUMENTATION

by

MICHAEL STEVEN BEER

A thesis submitted in partial fulfillment
of the requirements for the degree

of

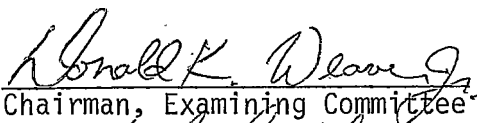
MASTER OF SCIENCE

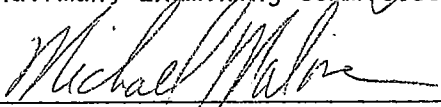
in

Electrical Engineering

Approved:


Head, Major Department


Chairman, Examining Committee


Graduate Dean

MONTANA STATE UNIVERSITY
Bozeman, Montana

July, 1980

ACKNOWLEDGMENT

The development and application of the fiberoptic data link draws on information learned through several other data acquisition projects developed at Montana State University. The author extends appreciation and thanks to Dr. Daniel N. March, who contributed the fiberoptics used to build the data link, and Western Telecomputing Corporation, that has provided assistance with hardware and its support. Of course, the author is especially grateful to Dr. Donald K. Weaver for his assistance and support of the project.

Michael S. Beer

TABLE OF CONTENTS

Chapter	Page
I. INTRODUCTION	1
II. THE FIBEROPTIC DATA LINK	7
The Fiberoptic Transmitter	9
The Fiberoptic Receiver	9
III. THE INTERFACE CIRCUITS AND TRANSDUCERS	14
The Voltage-to-Frequency Converter	14
The Current Transducers	16
The Rectifier and Filter	19
The Temperature Transducer	22
IV. SOFTWARE DEVELOPMENT	24
Software Routines	25
The Coordinating Program	28
The Hardware Drivers	31
The Teletype Driver	31
The Frequency Board Driver	33
The A-to-D Converter Driver	35
V. EXPERIMENTAL RESULTS	37
VI. SUMMARY	49
BIBLIOGRAPHY	55
APPENDIX	58
Appendix A:	59
Analysis of Filter/Rectifier Stage..	60
Appendix B:	63
Interfacing PLM/80 and Assembly	
Language Programs	64
Appendix C:	65
Program for Determining Minimum	
Software Delay	66

Appendix D:	68
Coordinating Program PLM/89 Listing.	69
Coordinating Program PLM/80 Listing with Assembly Mnemonics	74
Appendix E:	88
Parameter Passing Specifications for Teletype Drivers	89
Teletype Drivers PLM/80 Listing	91
Teletype Drivers Assembly Listing ..	96
Teletype Drivers PLM/80 Listing with Assembly Mnemonics	100
Appendix F:	113
Parameter Passing Specifications for Frequency Board Drivers	114
Frequency Board Drivers Assembly Listing	116
Appendix G:	123
Parameter Passing Specifications for A-to-D Drivers	124
A-to-D Drivers Assembly Listing	125

LIST OF TABLES

Table	Page
1. Hall Effect Device Coefficients	41
2. Current Transformer Coefficients	42
3. Temperature Sensor Coefficients	44
4. Coefficients for Voltage-to-Frequency Converter..	48

LIST OF FIGURES

Figure	Page
1. System Signal Flow	4
2. Transmitter	10
3. Fiberoptic Receiver	11
4. Voltage-to-Frequency Converter	15
5. Rectifier/Filter	21
6. Temperature Sensor	23
7. Software/Hardware Interface	26
8. WTC 520 Board	34
9. Transfer Function for Hall Effect Device & Data Link	40
10. Transfer Function for Current Transformer & Data Link ...	40
11. Transfer Function for Temperature Sensor	43
12. Transfer Function for Voltage-to-Frequency Converter	47
1A. Rectifier/Filter	61

ABSTRACT

The subject of this thesis is the application of a fiber optic data link for use in instrumentation applications where conventional information transmission techniques cannot be used due to excessive electrical noise. A technique was developed for converting a direct current voltage to a pulse string whose frequency was proportional to the voltage and transmitting this pulse string over the data link. This technique also allowed recovery of the information at the receiving end in digital form using a microprocessor-based data acquisition system. As a demonstration of the validity of this technique, sensors were constructed to enable use of the data link to measure alternating currents, direct currents, and temperature. Mathematical transfer functions were developed for each of these sensors to aid the microprocessor-based system in providing accurate measurements.

CHAPTER I

INTRODUCTION

This thesis describes a research activity directed towards applying fiberoptic data links to instrumentation. The intent of this thesis is to describe the equipment necessary to achieve this, and verify that the instrumentation/data link combination operates properly. There will be no attempt to compare this method with other instrumentation techniques, and no efforts will be made to optimize the results presented here with respect to a certain set of conditions. The techniques advanced in this thesis have application in areas where high voltages and electromagnetic interference make electrical methods of information transmission impossible or impractical.

Fiberoptic data links have been much publicized by Bell Laboratories because of the large quantity of information that they can carry. Fiberoptics have useful properties other than large bandwidth. The properties of fiberoptics that this research exploits involve electrical noise immunity and electrical isolation. Another benefit that the methods developed have is that they digitally encode the information that is to be transmitted. As a consequence, these methods work particularly well in, but are not limited to, systems that employ microprocessors and other digital equipment.

Fiberoptics involves the transmission of signals using visible and near-visible light in transparent waveguides. Due to the short wave-length of light relative to the other forms of electromagnetic

radiation, these waveguides are much smaller than conventional waveguides. Typical fiberoptic fibers have cross-sections approaching that of a human hair. Because of this, the cables encasing the fibers are small diameter (1/4") and very flexible. The fiberoptic fibers themselves are constructed of either glasses or plastics, with glasses being preferred over plastics due to their lower attenuation. The fibers themselves are usually encased in several layers of plastic to form the fiberoptic cable. Since both glass and plastic are excellent insulators, the resulting fiberoptic cable provides excellent electrical isolation between transmitter and receiver. Since light waves are used rather than electric waves, the signals in the fiberoptics are immune to those type of interference that degrade electrical signals.

Information can be sent through fiberoptics in either analog or digital form. It was decided to use digital encoding in this project. The form of digital encoding chosen was pulse modulation, in which the frequency of the pulse string coming out of the analog-to-digital (A-to-D) converter is nearly a linear function of the input voltage. This type of A-to-D conversion will be referred to as voltage-to-frequency conversion and the A-to-D converter will be referred to as a voltage-to-frequency converter throughout the rest of this thesis.

The instrumentation technique advanced in this thesis involves converting the quantity to be measured to a voltage. This voltage is then applied to a voltage-to-frequency converter and the pulse string

from this converter is used to turn a light emitting diode (LED) on and off. The light from this LED is then launched down the fiberoptic cable. At the receiving end of the cable, the light signal is transformed into an electrical signal that conforms to the voltage levels established for Transistor Transistor Logic (TTL) circuits. This signal's frequency is then proportional to the quantity initially measured, the exact proportionality coefficient(s) being dependent on the transfer function of the system as a whole. This type of system is diagrammed in Figure 1.

Two types of sensors were developed to interface with the fiberoptics data link. One type of sensor measures temperature and the other type measures electrical current. The temperature sensor is an integrated circuit that produces a current proportional to the temperature of the device. Additional circuitry changes this current into a voltage and allows amplification and offset to be set so that the temperature range of interest can correspond to the voltage range of the voltage-to-frequency converter.

The current sensors developed for the data link evolved along two lines of thought. One line was aimed at measuring current levels in conventional A.C. lines and the other line was aimed at measuring currents that could be either D.C. or A.C. The first sensor was developed to measure these A.C. and D.C. currents. It is based on a Hall effect device and generates an output voltage proportional to the in-




Figure 1 System Signal Flow

stantaneous intensity of the magnetic field surrounding the conductor and the excitation current supplied to the Hall effect device. The second current sensor is based around a conventional current transformer similar to those used in most power applications. This sensor is, of course, limited to A.C. currents with frequencies near 60 Hz. The current transformer outputs a current proportional to the A.C. current being measured. This A.C. current is passed through a resistor whose size was chosen so that the voltage drop generated corresponds to the voltage produced by the Hall effect device under similar conditions. The signal from either current sensor is processed by circuitry that rectifies and amplifies the signal producing an output that is compatible with the voltage-to-frequency converter.

No matter which sensor is used, the frequency of the voltage-to-frequency converter has to be determined. A microprocessor based data acquisition system (DAS) was used to determine the frequency of the pulse string generated by the fiberoptics data link receiver. The method used involved counting the number of pulses from the receiver over a known period of time. This method is generally referred to as a gated counter. The gated counter is controlled by a program that the DAS microprocessor executes. The program can direct the results to be printed on an output device or saved in memory.

This thesis describes the construction of the fiberoptic data link, the construction of the voltage-to-frequency converter, the con-

struction of temperature and current transducers compatible with this equipment, and it describes the software used to control and calibrate the system.

CHAPTER II

THE FIBEROPTIC DATA LINK

The fiberoptic data link consists of a transmitter and receiver separated by 100 meters of Valtec PC10 FIBERdata single-fiber fiberoptic cable. This data link is not a high speed data link, although the speed limitation is in the electronics and not the fiberoptics. It was decided to design the link to handle a square wave at a maximum frequency of 100 kHz. The link had to also be able to perform down to D.C. conditions. These specifications were arrived at after considering the frequency ranges available in hybrid and monolithic voltage-to-frequency converters. Since these specifications were set, several commercially available data links have come on the market with virtually identical specifications. This frequency range appears to be the industry standard for low data rate fiberoptic links.

The specification that the link be capable of operating down to D.C. completely precludes the use of capacitively coupled amplifiers at any point in the transmitter or receiver. It was also decided from a point of size, power consumption and operating environment that semiconductor sources and detectors would be used in the data link.

The light sources chosen were conventional LEDs. These sources are by far the easiest sources to bias, operate and align. Another factor in this selection was the large amount of work that has gone into improving LED reliability and operating life. LEDs are also much easier to obtain than other types of solid emitters. The main

disadvantage is that the devices require large amounts of power relative to the amount of power coupled into the optical fiber. This is a function of the broad area over which the light is emitted from the LED package and the small area that the terminated fiber optic has. This inefficiency is a trade-off that must be accepted to keep the ease of alignment attribute. The sources used in this data link nominally consume 100 milliwatts of power and the author has calculated that the amount coupled into the fiber is on the order of a few microwatts. Both red and infrared sources were used in the data link.

The detectors chosen were Positive Intrinsic Negative (PIN) diodes. These devices are readily available and operate at much lower bias voltages than avalanche type detectors. This is due to their internal construction which more effectively utilizes the voltages generated in the space-charge layer (11). The amount of current developed by one of these diodes in its reverse operating mode is linearly dependent on the incident light. This is true over the voltage range that extends from the reverse breakdown voltage of the diode to the point at which the generated current and the input resistance of the first amplifier produce enough voltage to remove the diode's reverse bias.

There was also the choice between single fiber light guides and multiple fiber light guides. The multiple fiber light guides are undoubtedly easier to use in a data link since they carry more light than single fiber guides, but they are more costly and also more difficult

to properly terminate. Because of this, single fiber cable was chosen for this project. Again, it appears from advertising in electronics magazines that the choice of single fiber cables over multifiber cables has been made by nearly every data link manufacturer.

The Fiberoptic Transmitter

The schematic for the transmitter for the visible LED was provided by Valtec, the fiberoptic cable manufacturer. This was modified slightly for the infrared LED which required a different drive current from the red LED. These circuits are shown in Figure 2. The NAND gate provides hysteresis and drive current for the transistor, which can supply the higher current required by the light emitting diode. The different LED's used are shown along with the required current limiting resistors. The transmitter has worked well from the beginning of the evaluation, and this circuit has not been changed.

The Fiberoptic Receiver

This circuit was the most difficult of all of the circuits to design. The chief reason for this was the very small signal produced by the PIN diodes. The fiberoptics cable manufacturer had provided a schematic for a receiver with the cable that was used in the project. This receiver was constructed and it did work, but it was very sensitive to just about any change in operating conditions. Most of these pro-




Figure 2 Transmitter

blems were found to be coupling problems between the power supply and various parts of the circuit. To avoid these problems, a more conservative design was adopted. This design is presented in Figure 3. Great care was taken to not only decouple the power supply from the operational amplifier supplies, but to also decouple the various parts of the circuit from one another.

The input signal is very small, having been measured at .5 microamperes for the red LED/detector pair. The input resistance of the op amp A1 is on the order of 70 megohms, so virtually all of this current passes through the 510 kohm resistor. An input resistance this large may be criticized as introducing thermal noise and converting any

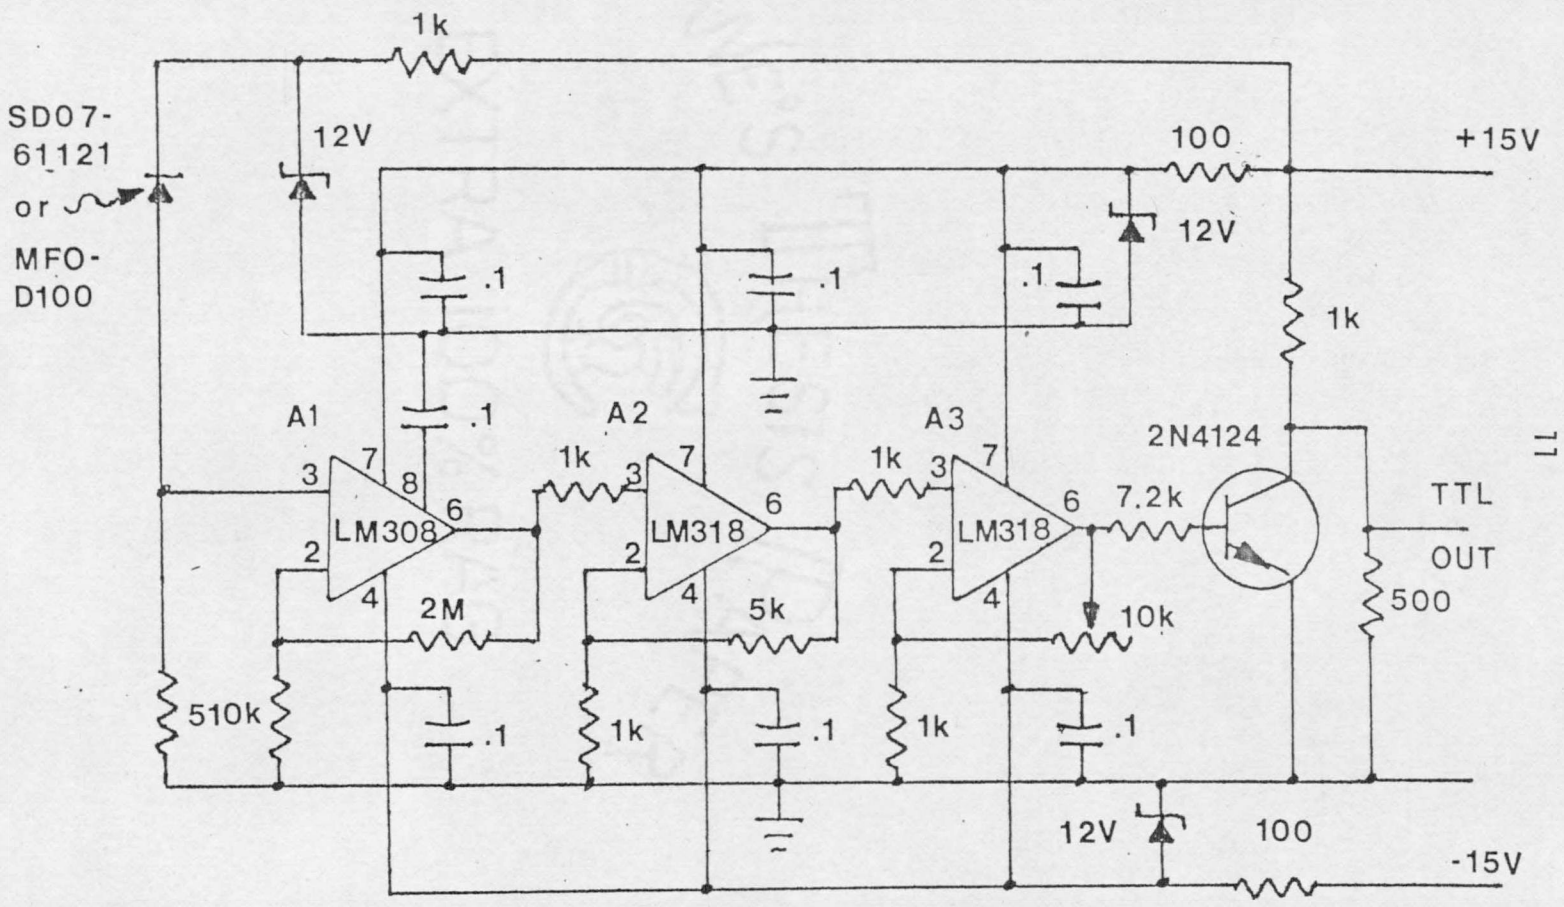


Figure 3 Fibernoptic Receiver

stray currents into voltage drops, but as long as this introduced noise does not create spurious logic transitions at the output, it is a small price to pay for decreasing the amount of voltage gain the op amps must supply. The op amp A1 must have very small offset current and bias current due to the already mentioned small signal. The LM308 op amp selected for this task meets these requirements easily with .4nA and 3nA respectively (8). The voltage gain was kept low in this stage since the LM308 does not have a gain-bandwidth product much better than 1 MHz and the circuit must have an overall bandwidth of 100 kHz. Once the signal has been amplified by the LM308, there is sufficient drive current, and the bias and offset current specifications can be relaxed somewhat. For the next two amplifiers, LM308 op amps were used. These amplifiers have a gain-bandwidth product of 15 MHz (8), and can easily be pushed to higher voltage gains than the LM308. The last of these two amplifiers has a trimpot in the feedback loop to provide an adjustable gain. This is to compensate for variable input levels caused by different detector/source pairs or losses due to variable fiberoptic lengths. The final stage of the receiver is the output transistor. This transistor accomplishes two things. It not only provides some wave shaping that removes noise introduced in the amplification and cleans up the rising and falling edges of the waveform, but provides a signal inversion to make up for the inversion introduced by the fiberoptics transmitter.

There were two source detector pairs that were used in the data link. The pair that operates in the red end of the spectrum is a Spectronics SD0761121 PIN diode and a Fairchild FLV104 LED. This was the pair supplied by the fiberoptic cable manufacturer. The other detector/source pair operates at infrared wavelengths. These are both Motorola devices and the part numbers are MFOD100 and MFOE200 respectively. The infrared pair gives a slightly stronger signal, which is probably due to the higher semiconductor efficiency at this wavelength (11). It is interesting to note that either detector will work with either LED, but it is often necessary to readjust the gain in the receiver to get maximum performance when this is done. If the detector and LED are matched, there is no need to readjust the gain when going from one detector/source pair to the other.

CHAPTER III

THE INTERFACE CIRCUITS AND TRANSDUCERS

Once the data link was operating, interface circuitry was constructed to allow generation of a pulse train from an input voltage. This pulse train could then be fed into the fiberoptic data link. The input voltage was to be derived from some physical quantity that a transducer could change into an electrical signal. This electrical signal had to be transformed into a voltage, since the only A-to-D conversion available is a voltage-to-frequency converter. Provisions had to be made to scale and offset this voltage so that the maximum output of the transducer would correspond to the maximum allowed input voltage to the voltage-to-frequency converter.

The voltage-to-frequency converter will be the first interface circuit discussed. The description of the current transducers and their interface circuitry will follow, and the chapter will close on a similar description for the temperature transducer.

The Voltage-to-Frequency Converter

The voltage-to-frequency converter used for the analog-to-digital conversion is based around an Analog Devices AD527 monolithic voltage-to-frequency converter (1). This circuit was selected over other integrated circuits because of the small number of additional components necessary for operation (6 additional components), its good stability and linearity, and the good test and application documentation provided

provided by the manufacturer. A monolithic voltage-to-frequency converter was chosen over a hybrid or discrete component device. Integrated circuits recently introduced on the market are very much less expensive in terms of cost and power consumption than either of the other types, at very little sacrifice in performance. The schematic for the circuit used is given in Figure 4. This circuit is based on circuit suggestions provided by Analog Devices.




Figure 4 Voltage-to-Frequency Converter

The internal operation of the device is not dramatically different from other voltage-to-frequency converters. A current source controlled by the input voltage feeds charge into an integrator composed of an external capacitor and an internal operational amplifier. The voltage

developed across the capacitor is compared to a reference voltage derived on the chip. When this voltage exceeds the reference voltage, an astable multivibrator is triggered. The output from the astable multivibrator in turn drives an output amplifier which provides higher driver current. The frequency of operation is determined by the size of the external capacitor and also the size of the external resistors used to adjust the magnitude of the current from the voltage-controlled current source.

The linearity and temperature stability of the circuit is greatly dependent on the quality of the external resistors and capacitor. These components must have very low temperature coefficients and the timing capacitor must have a very low effective series resistance if the circuit is to meet the manufacturer's specifications.

The Current Transducers

Two types of transducers were used to make the current level measurements on a conductor. One of these transducers was an Ohio Semitronics CT50L DC to 5 kilohertz transducer (9). This transducer is based on the well known Hall Effect (4). Unlike conventional current transformers which depend on the change of the current and a change in its accompanying magnetic field, this transducer directly measures the intensity of the magnetic field. Thus, its time response is not limited by the inductance present in a sensing coil, but only by the semi-

conductor transport processes and the properties of the material used to concentrate the magnetic field around the Hall Effect sensor. This enables the transducer to have a response linear to within .5% over a range from D.C. to 5 KHz. The upper figure is a function of this particular Hall Effect device and not the Hall Effect itself.

There are two disadvantages to this type of transducer. The first disadvantage is that it must be supplied with a constant excitation current (nominally .1 A). Variations in this excitation current show up as unwanted variations in the transducer output. The second disadvantage is that the Hall Effect device used is a semiconductor device and the mobility of the charge carriers in the semiconductor is dependent on the material temperature. Since the voltage produced by the Hall Effect is dependent on the mobility of the charge carriers, variations in temperature may also cause unwanted output variations.

Neither of these disadvantages is insurmountable. The excitation current can be derived from a variety of sources. Possible sources are batteries, power supplies, or in the case of measuring high A.C. current, the excitation current could be transformed right off the power line itself. The problem of temperature dependence is also easily solved. Since the data is taken in digital form, the temperature dependence can be removed mathematically. Methods to do this are well known (3) and usually consist of fitting experimentally derived error data to a mathematical function. This function is either

directly evaluated by the microprocessor or is built into a lookup table that the microprocessor can reference. Both of these error removal methods require that the temperature of the transducer be known. Temperature measuring methods are discussed later in this section.

The second transducer used was a conventional current transformer, specifically a Westinghouse Current Transformer type CMS. This is a transformer similar to those presently used by the power companies in determining current levels in transmission equipment. The only modification required is the insertion of a small shunt resistance across the transformer output terminals. This shunt resistor (.024 ohms) was chosen so that the voltage drop across it would be compatible with the output voltage developed by the Hall Effect device.

Both transducers are operated so that a current of 50 amps causes a 30 millivolt peak output. The 50 amp figure was chosen for experimental convenience. Hall Effect transducers are available to measure currents up to 8000 amperes (9). Current transformers are available for a similar range (15).

A 50 amp A.C. transmission line was simulated by passing a 5 amp A.C. current through a ten turn coil wound through the active area of each transducer. An A.C. current source consisting of a variable transformer and a 1 ohm load resistance provided the current to the transducers. The current level in the "transmission line" was

then measured by recording the voltage dropped across the one ohm load and multiplying by ten.

The Rectifier and Filter

The signal coming from the current transducers has to be rectified and amplified before the signal could be used by the voltage-to-frequency converter, which requires a D.C. input voltage. Since the current transducers have a 30 millivolt peak output, direct rectification by a diode bridge of any sort will introduce an error. Suppose that the signal is first amplified to 10 volts peak prior to rectifying. If the signal passes through a full wave bridge rectifier, two forward diode voltage drops are subtracted from the signal. For silicon diodes this translates to somewhere in the vicinity of 1.2 to 1.5 volts subtracted from the peak voltage. For voltages less than peak, the error will be less than this, but it will be a nonlinear proportion of the maximum error due to the nonlinear transfer characteristic of the diodes. The 1.5 volt maximum error corresponds to a 15% error, which is much too high for the accuracy desired. Amplifying the peak output voltage to 100 volts would reduce the percentage error by a factor of 10, but these higher voltages have gotten out of the voltage range of most common integrated circuit amplifiers. This argument precludes using the "brute force" methods of rectification common to power supplies. The solution to this problem is to use a rectifier that is configured to

utilize the diodes only to block the portion of the signal that is not wanted rather than block and pass the waveform. The circuit in Figure 5 is based around a circuit idea from a National Semiconductor application note (7). This circuit not only provides full wave rectification, but also contains a portion to smooth the rectified waveform and amplify the resulting D.C. voltage to match the voltage range of the output to the voltage range of the voltage-to-frequency converter. Although this circuit substantially alters A.C. waveforms, D.C. waveforms are just amplified and smoothed to an average D.C. level. The circuit operation and a derivation of its transfer function are provided in Appendix A.

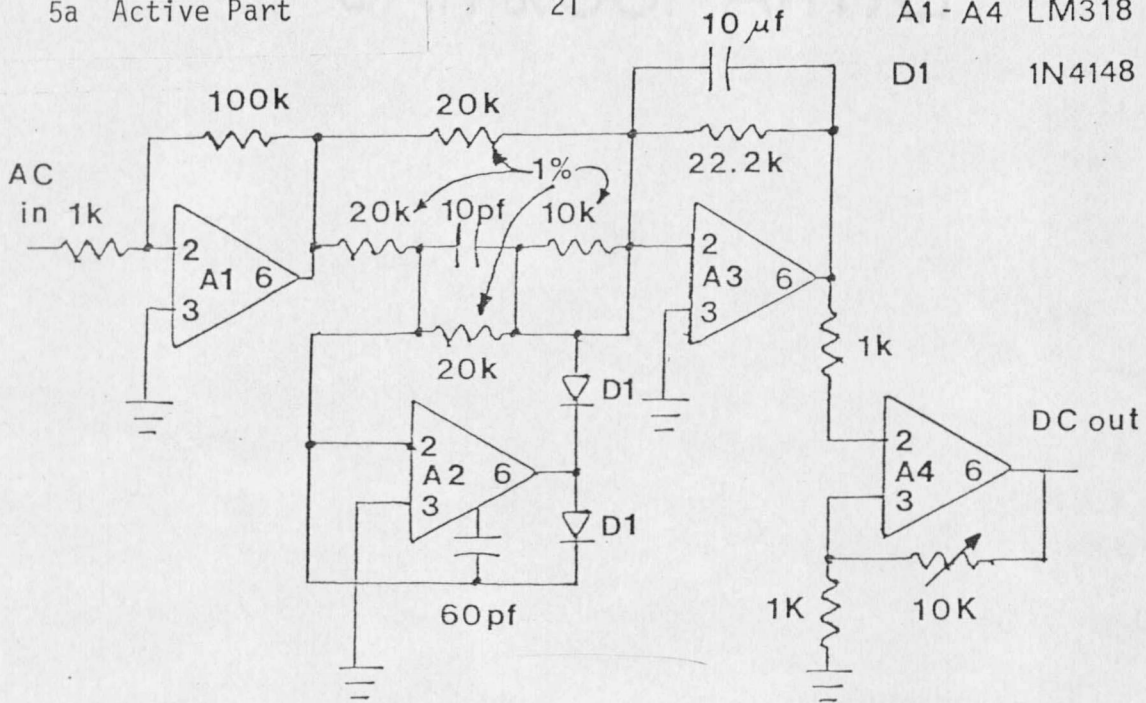
The schematic for the amplifier and precision rectifier is provided in the two parts of Figure 5. One of these parts depicts the active circuitry, while the other part depicts the power supply decoupling scheme (2). This decoupling scheme may seem rather elaborate at first, but the accuracy of the entire instrument depends on the proper functioning of this one circuit. The circuit will be operating in close proximity to the voltage-to-frequency converter which can introduce switching spikes on the ground and power lines of the rectifier. The decoupling will help to absorb these spikes before they can get into the op amp's circuitry. In the case of an application such as an A.C. current monitor, this circuitry would be operating on a power line and be totally immersed in the 60 cycle

5a Active Part

21

A1 - A4 LM318

D1 1N4148



5b Decoupling Part

R = 50 ohms

C1 = .1µf

C2 = 100 pf

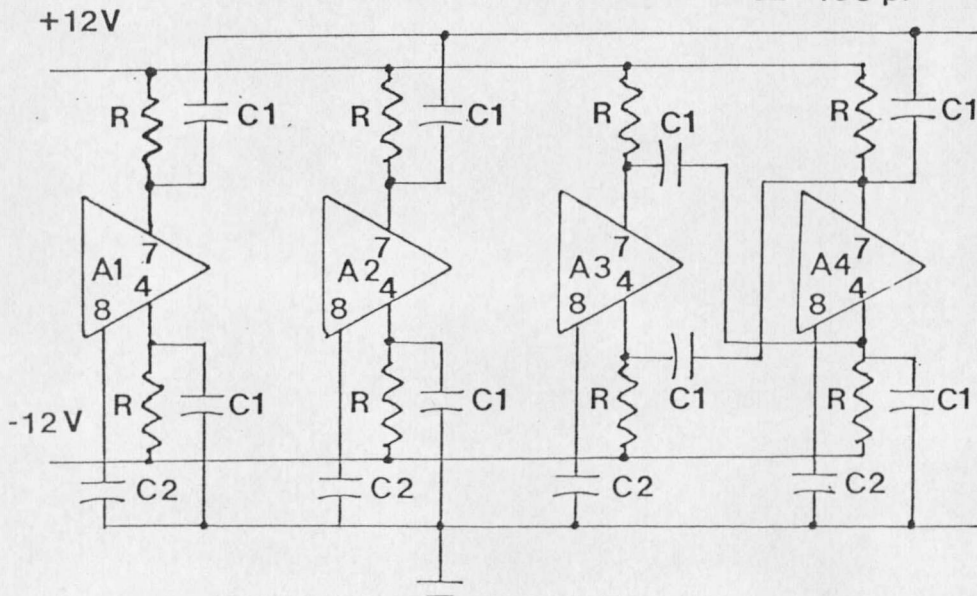


Figure 5 Rectifier/Filter

electromagnetic field of the power line. These fields may induce noise in the circuit at the 60 Hz line frequency. This noise is totally indistinguishable from a valid signal if it gets onto either the input or the output of any of the amplifiers, so it is essential that it be kept away from the active elements at all costs.

The Temperature Transducer

Since operation of the Hall Effect Device may require an accurate knowledge of the temperature, a temperature transducer would be required. Several different integrated circuit transducers were acquired, connected to whatever supporting circuitry was required, and this combination tested in an environmental chamber. On the basis of these tests, the Analog Devices AD590 two terminal IC temperature transducer was selected as the transducer to be used (1). The parameters used in selecting this transducer were ease of use with the voltage-to-frequency converter and accuracy.

The temperature transducer has only two types of errors in the current output that is the response to temperature. These errors are offset errors and nonlinearity errors. There are several grades of this device available. In general, a more expensive grade has a smaller offset voltage error than a less expensive grade. The nonlinearity errors are the same from one grade to the next.

The circuit that the AD590 is used in is shown in Figure 6.

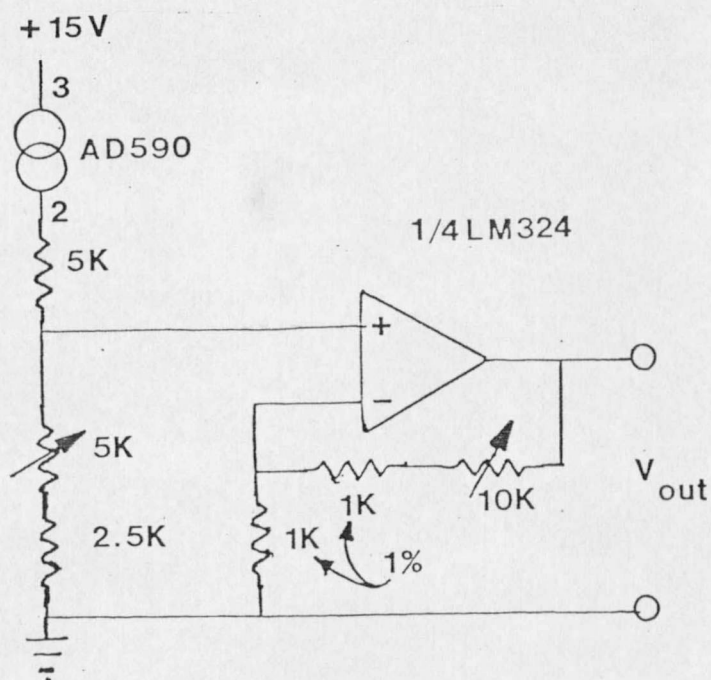


Figure 6 Temperature Sensor

There are two variable resistors in the schematic. The 5k resistor is used to remove the offset inherent in the device. The 10k variable resistor in the op amp feedback loop is provided to match the output voltage produced by the maximum anticipated temperature to the maximum input voltage that the voltage-to-frequency converter is set for. This allows the desired temperature range to be spread over a larger fraction of the voltage-to-frequency converter's operating range.

CHAPTER IV

SOFTWARE DEVELOPMENT

Before discussing the software to control the hardware that has been developed, it is necessary to describe the hardware environment in which the software operates. This environment consists of four pieces of equipment, all functioning as one microprocessor-based data acquisition system. The data acquisition system need not be this complex, this organization primarily reflects the equipment available for this project. The equipment consists of an INTEL PROMPT80 microcomputer, a Western Telecomputing Corporation WTC-800 computer system, a WTC-700 data acquisition system contains a four channel A-to-D converter (used to calibrate the data link instrumentation), a four channel event counter and the controller card that is connected to the WTC-800 system. This controller card provides all of the bus timing and addressing for the WTC-700 system. This system has a bus structure that is markedly different from the bus used in the WTC-800 system. The WTC-700 card cage contains slots for the controller card (already mentioned), two power supply cards, and 16 slots for data acquisition cards. Each of these slots is separately addressed, so the position of a specific card in the bus determines the software address of the card. Each of the cards in this system has as its output an ASCII string representing the quantity measured. The controller card enables this string to be read off one byte at a time.

Software Routines

The software written for this research project consists of two basic types of programs. These are the programs that interface directly with the hardware (called hardware drivers) and the program that coordinates these interface programs. This modular program structure (6) and how it interfaces with the hardware is illustrated in Figure 7.

This type of structure was chosen for its flexibility. If the hardware/software system is to do a different task, the hardware drivers do not need to be changed. The coordinating program is the only module that has to be altered, and then most of the changes will only be in the order that the supporting modules are called by the coordinating program. Another advantage of this type of structure from the programmer's point-of-view is that the programmer is not constrained to any one programming language. Routines that lend themselves to register manipulation, or that must execute very rapidly may be written in assembly language. Routines that involve data processing may be written in a higher level language that allows arrays and/or data structures (i.e., BASIC, PLM/80 or PASCAL). Routines that involve mathematical processing may be written in a formula oriented language such as FORTRAN or ALGOL.

The routines presented in this thesis are written in a combination of PLM/80 and INTEL 8080 assembly language. The choice of a certain language for one of the routines reflects not only the above mentioned considerations, but also the fact that at the beginning of the project

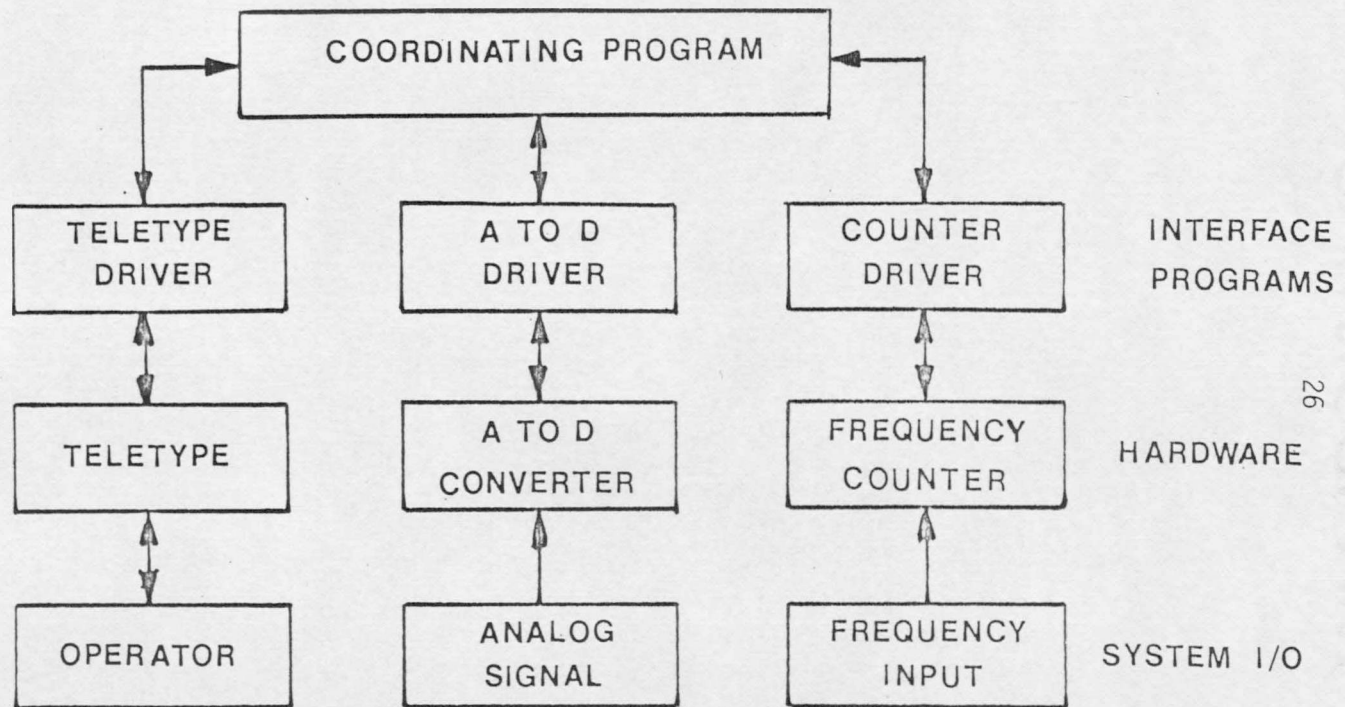


Figure 7 Software/Hardware Interface

assembly language was the only language resident in the MDS 230 software development system used to develop all the software in this project. The author also recognizes that the reader of this thesis may not have access to a PLM/80 compiler or be familiar with the PLM/80 syntax. Hence, all PLM/80 routines are presented in the appendix in two forms. The first of these forms is a listing of the PLM/80 source code. The second form is a listing of the same source code with the compiler generated assembly language mnemonics included. Appendix B also contains a description of how PLM/80 and assembly language programs are interfaced. This part of the appendix should make the register and stack manipulation used prior to and after calls easier to understand. The program listings are contained in appendices D through G.

The software routines for this project must be able to:

- (a) receive input from the operator as to what sensor(s) to read and when to read them
- (b) read the sensors
- (c) process the readings
- (d) report the results to the operator

These requirements demand that the coordinating program be able to interface with drivers for the following hardware:

- (a) A teletype
- (b) The frequency counter connected with the fiber optics data link
- (c) An A-to-D converter (for calibration only)

Before discussing individual hardware drivers, the operation and

construction of the coordination program will be considered.

The Coordinating Program

This program first interrogates the user to set up parameters that determine the initialization of the frequency counter, the number of times the readings will be taken, and the time delay between the readings. The program checks the initialization values provided by the operator in two ways. The values are first checked to see if they correspond to the format given in the interrogation message. If the values pass this first test, the values are converted from the ASCII strings that came from the teletype to either binary numbers (for the delay and repetition parameters), or binary coded decimal (BCD) numbers (for the counter initialization parameter). The magnitudes of the binary numbers are checked to insure that they fall within the range prescribed by the interrogation message. It is not necessary to do this for the BCD numbers since if the input number string passes the first test, it will fall within the proper range. Should any of these values fail either of these tests, the appropriate interrogation message is repeated until the operator can provide a satisfactory response. There is only one assumption made in this program. It is assumed that the operator will enter numbers on the teletype and not some other character. The program will function should characters other than numbers be supplied, but no error

message will be generated and the results will in general be different from those anticipated.

Once the initialization parameters have been set, the coordinating program will call the hardware drivers in the proper order to read the sensor and print the results on the teletype. The program then enters a software timing loop to generate a delay between readings. The duration of this timing loop was one of the parameters specified in the initialization sequence. When this loop times out, the program will make another reading if it has not completed the specified number of readings. When the program has finished these readings, it re-enters the initialization sequence and waits for further operator input.

At this point, it is necessary to make some additional comments on the actual time delay between readings of the frequency counter.

This delay is determined by three factors:

- (1) The length of time the A-to-D converter requires to make a conversion. This also includes delays in the accompanying software routines.
- (2) The length of time required to print out or store the results of a conversion.
- (3) The delay introduced by the software controlled timing loops.

The user has very little control over the first of these factors. The delay caused by the second factor can be minimized, and the delay caused by the third factor is completely controllable.

The delay caused by the A-to-D conversion is directly proportional to the number placed in counter zero to specify the length of time counter one downcounts. The delay is also dependent on the clock rate fed into counter zero. The hardware configuration chosen for the WTC-520A board gives a clock period of 25.6 microseconds. At this clock rate, if counter zero is initialized with a count of 8200, counter one will contain a count corresponding to 1/4 the actual frequency of the incoming pulse string. These two factors imply that the length of time required for one conversion is .21 second. If there were no other appreciable delays in the system, the maximum sampling rate would then be approximately 4.76 conversions per second.

When the delays caused by the teletype are accounted for, this rate drops dramatically. If the teletype does no more than print out the four digits of the conversion, a carriage return and a line feed, then the total delay is .81 second and the rate is reduced to 1.235 conversions/second. It is quite likely that more information than this would be printed on the teletype, with a corresponding reduction in sampling rate.

The maximum sampling rate of 4.76 conversions per second could be very nearly approached if the data were stored directly in memory, instead of being immediately printed on the teletype. It was assumed in calculating the above sampling rate that the time required by the software is negligible compared to the hardware time. A routine to

place the data directly in memory would presumably wait in a loop until the conversion was completed, read the appropriate counter, store the BCD information in memory, store a data delimiter in memory, check to see if the counter has been read the required number of times and, if not, begin another conversion and jump to the wait loop. A portion of code to do just this is presented in Appendix C. This routine has not actually been run on the DAS, as it's purpose is merely to illustrate a point. This point is to provide an idea of the length of time a similar program would require to store the information in memory, and see if this time is significant compared to the hardware conversion time.

Each pass through this routine takes a maximum of 220 clock cycles. With a clock rate of 2 MHz, this is an execution time of:

$$220 \times 1/(2 \times 10^6) = .11 \times 10^{-3} \text{ sec}$$

Since the A-to-D conversion time alone is 210×10^{-3} sec., this additional processing time is relatively inconsequential. This example justifies the earlier assumption that the time consumed in the hardware interfacing routines is not significant compared to the conversion time.

The Hardware Drivers

Now that the overall operation of the coordinating program is understood, the program modules that the coordinating program inter-

faces to will be considered. Details on interfacing software to these routines and the listings of the routines are contained in the Appendices D through G.

The Teletype Driver. The teletype is the main I/O device in the system. Contained in the PROMPT 80 computer used in this project is an INTEL 8251 Universal Synchronous/Asynchronous Receiver/Transmitter (USART) (5) and associated circuitry to drive a teletype on 20 ma. current loops. All of the communications to and from the teletype are done via this USART. The USART appears to the microprocessor as two input/output ports. The port at address 236 serves as the port through which the characters are transmitted to and from the teletype. These routines are derived from a set of assembly language routines prepared earlier by L. K. Smith and D. K. Weaver (10). This teletype driver does not use interrupts and thus must stay in a waiting loop during character transmission or reception. These routines are now written in a mixture of PLM/80 and assembly language and have been modified both in logic structure and buffer structure. Assembly language was used to decrease size and speed execution of some of the routines.

The routines require that a buffer area in memory have previously been defined. This buffer area is up to 256 bytes long. The first byte in the buffer (lowest address) is reserved for the total number of filled bytes in the buffer. It is zero if the buffer is empty. There are no restrictions on the contents of the rest of the buffer.

The routines themselves are presented in Appendix D, along with descriptions of the parameters required by each routine. There are provisions for calling the important routines from either assembly language or PLM/80. The routines not only enable the teletype to be used as an I/O device, but also allow rudimentary line editing to be done on the line presently in the buffer. Characters may be deleted beginning with the last character entered by pressing the rubout key. The contents of the entire buffer may be erased leaving an empty buffer by using CONTROL-X. If the line being entered exceeds the buffer length, the routines will quit echoing the characters after the buffer is full. The carriage return key terminates the line input.

The Frequency Board Driver. The next interface routine to be discussed will be the driver for the WTC-520A frequency board shown in Figure 8 (14). The WTC-520A frequency counter board is designed around the INTEL 8253 counter (5). There are three sixteen bit down-counters on this chip. Counters 0 and 1 are used on the WTC-520A board to form a gated frequency counter. Each counter has a gate enable input that controls when the input pulse string is to be counted, and an input line that indicates when the counter reaches zero. Counter zero on this chip and a D flipflop are used to control the length of time that counter one may count the incoming pulse string. To successfully operate the frequency board it is necessary to:

- (1) Reset the output of counter zero.

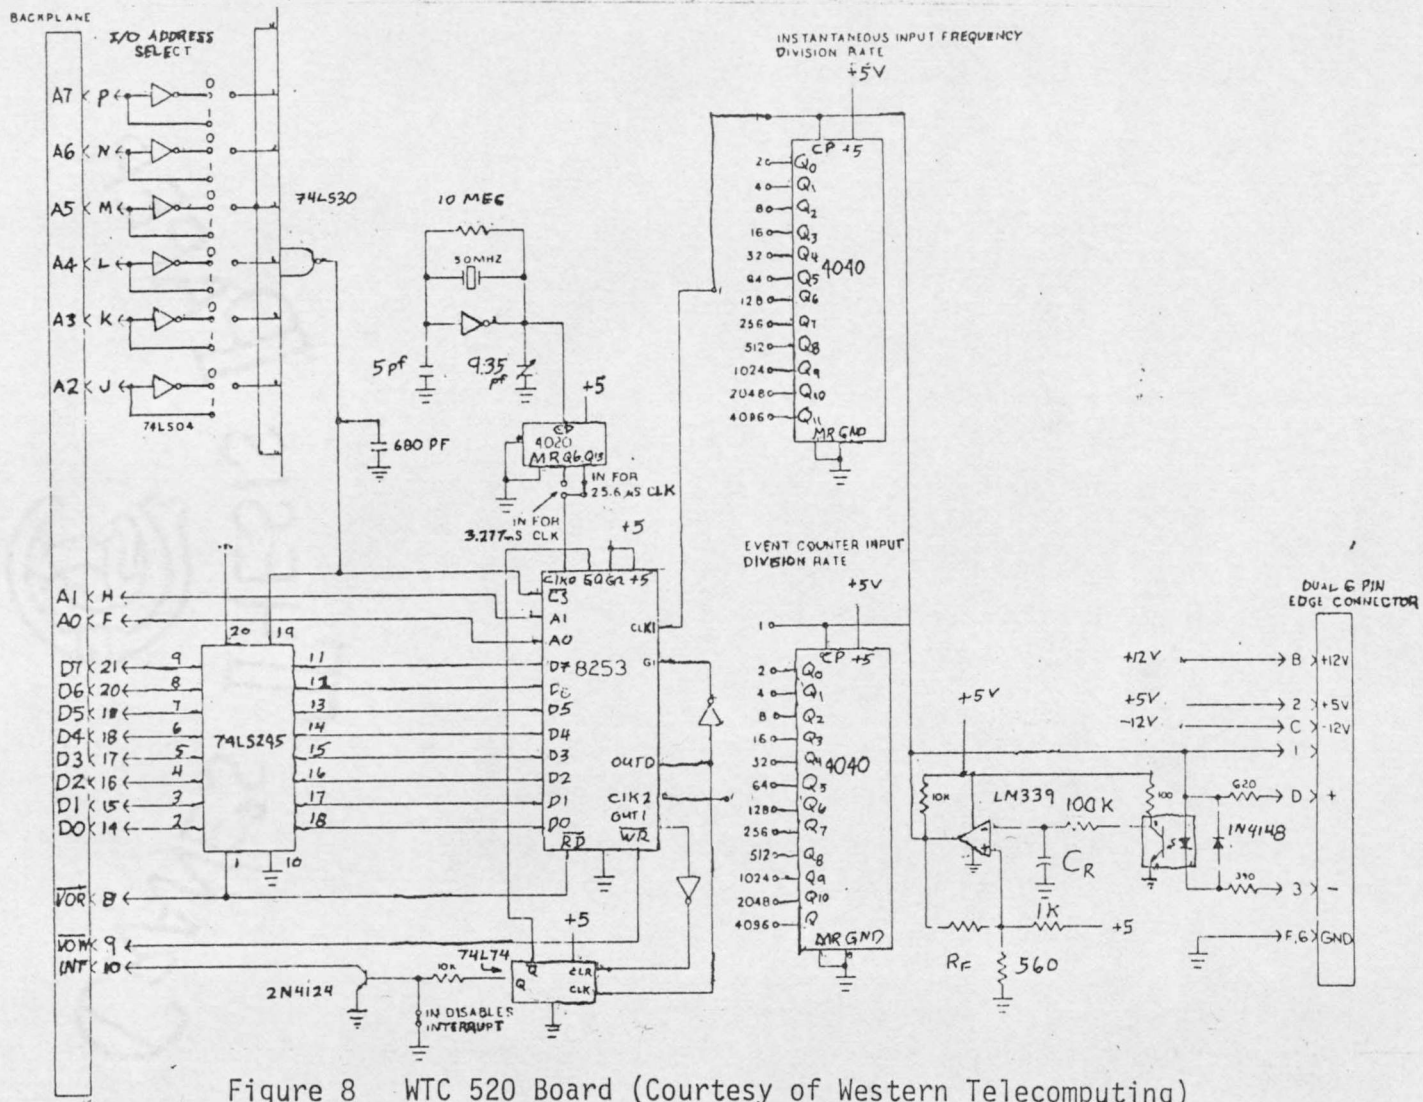


Figure 8 WTC 520 Board (Courtesy of Western Telecomputing)

- (2) Preset the D flipflop by forcing counter one's output high, then low. At this point, both counters are ready to begin counting.
- (3) Load counter one with the number corresponding to the maximum number the counter can hold in the mode it is being operated in.
- (4) Load counter zero with the number corresponding to the time interval that counter one is to count for. This starts both of the counters.

When counter zero reaches zero, the D flipflop will be triggered, and both counters will stop counting. The counters may then be read at any time. Counter zero will, of course, contain zero, but counter one will contain a value corresponding to the number of counts subtracted from the number initially loaded into this counter. It is the responsibility of the programmer to determine that the counting interval is always short enough to insure that counter one will not count more pulses than the number initially loaded into the counter.

Since in this application the microprocessor-based system has nothing to do until the frequency counter is through making it's conversion, the driving routine merely waits in a loop checking to see when counter zero's contents are zero. When this is true, the contents of counter one are read and subtracted from the value initially loaded into counter one. The result of this subtraction is equal to the number of pulses input to the board during the time it took counter zero to reach zero. This number is changed into an ASCII string and placed in the appropriate buffer.

The A-to-D Converter Driver. The third interface driver controls a WTC-700-120A analog to digital conversion board (13). This board resides in a WTC-700 Data Interface System connected to the WTC-800 156 Microprocessor Controller (13). This controller uses the INTEL 8255 parallel interface. The 8255 must be set into the proper mode to interface with the support electronics on the WTC-700 156 board. When the mode of the 8255 is initially set, the entire WTC-700 system is reset and then the individual cards in the DAS may be communicated with through the 8255.

The WTC-700-120A analog-to-digital converter may be read by signaling the card to make a conversion, waiting for the conversion to be completed, and then reading the results of the conversion. The results are in the form of a five byte ASCII string and are read from the WTC-700-120A using a byte by byte handshake. The details of this protocol are provided in the program listing contained in Appendix G.

CHAPTER V

EXPERIMENTAL RESULTS

After the hardware had been constructed, a series of tests were run to calibrate the operation of the sensors, fiberoptic data link, and DAS operating as one system. The process used in calibration involved two stages. The first stage involved individually adjusting the response of the sensors, voltage-to-frequency converter, data link, and software parameters until they responded reasonably close to design specifications. The second state involved deriving mathematical formulas to correct for remaining errors in the system.

To accomplish the first stage of calibration, adjustments that had been built into each sensor's circuitry were used to bring the response of each sensor reasonably close to the design specifications. An exact setting was not required or attempted, as the mathematical portion of the calibration procedure compensates for any residual errors. Next, the gain of any required interface stages was set so that the maximum voltage that would be delivered to the voltage-to-frequency converter was +10 volts. Then the voltage-to-frequency converter was set so that a 10 volt input would produce a 10 kilohertz output. Finally, a 10 kHz square reference was used to adjust the gate time on the gated counter so there was a simple relation between the value in the frequency counter at the end of the counting period and the frequency being counted. At this point, the first stage of the calibration was finished.

The second stage of the adjustment procedure uses formulas to correct for nonlinearity and remaining adjustment errors since the DAS contains a microprocessor that is capable of performing numerical calculations. A series of measurements are made over the range of operation of the sensor. As these measurements are made, they are transferred over the data link to the DAS where they are recorded. At the same time, the same measurements are made with another set of instruments that are known to be accurate. A mathematical formula may then be constructed relating the values recorded with the DAS to the actual quantity that was measured. The derived formula is a representation of the transfer function for the system as a whole.

There are several methods to construct this formula (3). The method chosen in this case involves adjusting the coefficients of a polynomial so that the polynomial corrects for the errors in the system. The values calculated with this polynomial are expected to produce more accurate measurements than the results without correction. A BASIC program was provided by Dr. Donald Weaver to calculate the polynomial coefficients required for this method. To demonstrate the increased accuracy of the polynomial method over the linear approximation method, the author wrote a BASIC program to provide a least squares fit of the same data to a linear relation. Printouts of the data used to calculate the coefficients for this polynomial and the coefficients calculated for a seventh degree polynomial are provided in a table for each of the sensors. A graph of

the calculated transfer function is also given, along with a graph of error between the linear approximation and the seventh degree polynomial. In interpreting these graphs, it is important to remember that the polynomial representation of the transfer function is not exact, but has a waviness to it due to the fact that a finite number of terms were used in the polynomial. This waviness is apparent in the plots of the error function which is the difference between the polynomial and linear representations. While the minor waviness is a mathematical construction, the plots do indicate the increase in accuracy that is available with the polynomial method.

The two transfer functions for the current sensors (Figures 9 and 10 and Tables 1 and 2) plot the current measured in the conductor versus the frequency of the signal measured from the data link. The circuitry had been adjusted so that 5 amperes through the conductor should produce a 10 kilohertz reading with the DAS gated counter. Neither circuit produces exactly this amount. These discrepancies are due to the method used in the initial hardware adjustments. Readjusting the hardware to eliminate these discrepancies is unnecessary since the polynomial representation of the transfer function for the readjusted hardware would be just as accurate as the present polynomial.

The transfer function for the temperature sensor is given in Figure 11 and Table 3. This transfer function differs from the other transfer functions in that the output of the device was not adjusted so




Figure 9 Transfer Function for Hall Effect Device & Data Link




Figure 10 Transfer Function for Current Transformer & Data Link

DATA COEFFICIENTS ARE COMPUTED FROM

I	X(I)	Y(I)	I	X(I)	Y(I)
1	.29	.214	2	.544	.292
3	.984	.5355	4	1.494	.8168
5	2	1.094	6	2.538	1.373
7	2.988	1.644	8	3.48	1.915
9	4.048	2.231	10	4.504	2.485
11	5.044	2.787	12	5.51	3.045
13	6.04	3.355	14	6.514	3.613
15	6.98	3.884	16	7.52	4.181
17	7.978	4.439	18	8.528	4.735
19	9.038	5.03	20	9.524	5.29
21	9.944	5.548	22	10.508	5.858

COEFFICIENTS OF POLYNOMIAL SUMMATION $A(I)*X^{**I}$

I	A(I)
0	.106290860395553
1	.27675797880948
2	.230824938867842
3	-.0944872305830564
4	.020908734981913
5	-2.5305582678523D-03
6	1.57327492499476D-04
7	-3.9254887314146D-06

Table 1 Hall Effect Device Coefficients

DATA COEFFICIENTS ARE COMPUTED FROM

I	X(I)	Y(I)	I	X(I)	Y(I)
1	.274	.1187	2	.538	.2413
3	1.034	.4723	4	1.55	.7123
5	2.08	.9535	6	2.53	1.156
7	3.04	1.403	8	3.518	1.623
9	4.04	1.87	10	4.5	2.081
11	5.024	2.329	12	5.564	2.581
13	6.054	2.813	14	6.504	3.019
15	7.104	3.303	16	7.518	3.497
17	7.988	3.716	18	8.51	3.974
19	9.05	4.219	20	9.46	4.426
21	10.04	4.697	22	10.512	4.916

COEFFICIENTS OF POLYNOMIAL SUMMATION $A(I)*X^{**I}$

I	A(I)
0	-.0142191448561888
1	.489367371176304
2	-.0279904147401379
3	.0118919854596125
4	-2.38428291814793D-03
5	2.407419265383D-04
6	-1.14098596465538D-05
7	1.88465239576853D-07

Table 2 Current Transformer Coefficients

