

**Woods Hole
Oceanographic
Institution**



**High Continuous Bandwidth Multichannel
Acquisition System**

by

Keith von der Heydt

June 1991

Technical Report

Funding was provided by the Office of Naval Research
under Contract No. N00014-91-J-1296.

Approved for public release; distribution unlimited.

DOCUMENT
LIBRARY
Woods Hole Oceanographic
Institution

WHOI-91-14

**High Continuous Bandwidth Multichannel
Acquisition System**

by

Keith von der Heydt

**Woods Hole Oceanographic Institution
Woods Hole, Massachusetts 02543**

June 1991

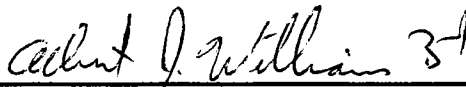
Technical Report

Funding was provided by the Office of Naval Research
under Contract No. N00014-91-J-1296.

Reproduction in whole or in part is permitted for any purpose of the United States
Government. This report should be cited as Woods Hole Oceanog. Inst. Tech. Rept.,
WHOI-91-14.

Approved for public release; distribution unlimited.

Approved for Distribution:



Albert J. Williams 3rd, Chairman
Department of Applied Ocean Physics
and Engineering



High Continuous Bandwidth Multichannel Acquisition System

Keith von der Heydt
Woods Hole Oceanographic Institution

June 12, 1991

1 Abstract

Multichannel data acquisition has been a keystone of 7 ONR sponsored Arctic acoustic research programs conducted jointly by WHOI and MIT investigators from 1978 through 1989. This report describes the status and capability of the most recent system developed at WHOI for the purpose of acquiring digital data from up to 64 channels at sampling rates up to 20 kHz per channel with data bandwidth to 5120 Hz. ONR funded the development of and use of this system and its prototype for 2 Arctic field experiments, PRUDEX 87 and CEAREX 89. It was most recently used during the Heard Island Feasibility Experiment in February 1991. Of note are *the auto-gain ranging capability offering a dynamic measurement range of greater than 120 dB*, the continuous storage capability of up to 200,000 samples per second to a Small Computer System Interface (SCSI) device, typically optical disk, and easy expandability with additional identical channels connected in parallel.

Contents

1	Abstract	1
2	Introduction	3
3	The Gain Ranging Amplifier	3
3.1	Hardware	4
3.2	The Gain Ranging Amplifier Design Topology	5
3.2.1	Front End	6
3.2.2	Highpass Filter	6
3.2.3	Lowpass Filter	6
3.2.4	Offset Cancelation	7
3.2.5	Auto-Ranging Amplifier	7
3.2.6	Analog to Digital Conversion	7
3.2.7	Controller	7
3.2.8	Data Transmission	8
3.2.9	Power Requirements, Packaging, Interconnections	8
3.3	Software	9
3.4	Test Data	13
3.4.1	Noise Floor	13
3.4.2	Low Pass Filter Response	14
3.4.3	Distortion	14
3.4.4	Crosstalk	15
4	Acquisition System	15
4.1	Storage Medium	15
4.2	Computer	16
4.3	GRA Interface	16
4.4	Acquisition Software	18
4.5	Optical Disk Storage Details	19
4.5.1	Optical Disk Header	19
4.5.2	Optical Disk Directory	20
4.5.3	Data Files	21
5	Final Comments	23
6	Acknowledgments	23
7	Figures	24-52
8	Appendices	53-61

2 Introduction

The techniques for doing multichannel seismic work in science and industry have usually focused on high dynamic range and high speed acquisition for short periods. The intermittent nature of reflection and refraction data acquisition often permits a system to buffer incoming data in memory for the duration of an event and store it to an archiving medium between periods of data generation.

Our acoustic work has required multichannel data acquisition at bandwidths similar to seismic applications but for extended periods (hours). We have developed a system of modest cost that offers both the high dynamic range capability often required of acoustic and seismic systems and high continuous storage bandwidth to a medium that affords easy access during processing.

This report describes the features of a system that has satisfied these requirements. The work has come about as a function of scientific needs of our joint MIT/WHOI Arctic acoustic work. The trend, as might be expected, has been to more channels, higher signal bandwidths and the ability to deal with greater volumes of data. Thus an objective has been to design acquisition systems that are expandable and have data generating capabilities consistent with that of practical storage devices. As it becomes easier to acquire large data sets, increasing emphasis must be placed on the need to monitor and verify data quality in near real-time to minimize the chance of returning from an experiment with bad data.

Three sections are presented in this report:

1. A discussion of and documentation regarding the Gain Ranging Amplifier (GRA) developed for use during ONR sponsored Arctic acoustic research programs at WHOI
2. A discussion of and documentation describing the multichannel acquisition hardware and software developed to store data generated by a bank of parallel GRA's. An update to WHOI Technical Report #87-49 regarding storage formats used for data on the optical disk is included.
3. Some thoughts on future data acquisition efforts

3 The Gain Ranging Amplifier

We have been involved with projects having data acquisition requirements focused on high dynamic range digital recording of seismo-acoustic signals. The character of such signals often includes explosive events, widely varying background noise that can mask signatures of interest, and generally unpredictable conditions.

In this discussion, the term *dynamic range* refers to the range of signal amplitudes over which the system can automatically adjust to prevent overloading and represent with the combination of a 13 bit analog-to-digital converter word and a 3 bit gain word. We are using the term *dynamic range* synonymously with system *measurement range* to describe the difference in spectral level between the maximum signal that can be represented and the noise floor of the system. This we distinguish from the 13 bit binary *precision* of the system

The GRA functional diagram in Figure 1 consists of 6 sections: input buffer amplifier, optional highpass filter, lowpass filter, cascaded amplifier chain, bank of 5 ADC's and controller. Each section will be described in detail.

3.2.1 Front End

In our experience it has often been necessary to distribute sensor signals to more than one recording system. We have chosen to provide separate, low noise single or differential input amplifiers that have multiple, individually buffered outputs. This led to a single ended input on the GRA card. A PA has its fixed gain set under software control to be as high as possible for the expected maximum input signal to avoid the noise floor of the LPF. Figures 2 and 3 with fixed gains of 1 and 16 respectively show the effect of PA gain on the effective system noise floor. There is little decrease in the level of the noise floor using PA gains greater than 16.

3.2.2 Highpass Filter

The highpass filter is used to limit low frequency motion, often of very high amplitude, from hydrophone sensors when towed or suspended in water. Its noise contribution is unimportant compared to that of the lowpass filter.

3.2.3 Lowpass Filter

A software controlled switch is connected such that the input to the LPF can be selected from the either from the HPF, or directly from the PA or to ground during offset calibration periods.

The LPF is a large Frequency Devices module, M/N 848P8B, with noise specified at 50 μ volts over the band 5 Hz to 50 kHz. This broadband noise floor sounds a bit more ominous than it is, as can be seen in Figures 7 through 13 showing noise at various signal levels. (The spectral levels indicated are very close to "per Hertz" normalization because the FFT bin bandwidth is almost 1.) Clearly, the LPF was a design compromise between achieving filter cutoff programmability and a reasonable noise level. The system noise floor is dominated by this filter and it is due to the bi-quad, op-amp intensive configuration of the filter design. These op-amps, TL084's, are moderate in bandwidth but not particularly low in noise level and at this time somewhat aged. These filters were the best available when this card was designed in late 1988.

The rolloff rate of the LPF is 48 dB/oct, samples of which are shown in Figures 4, 5, & 6. The phase and gain match of these filters is $\pm 2.0^\circ$ and $\pm 0.2dB$, max. Distortion from the LPF, 0.004% ($-88dB$) max, is low compared to the distortion as a function of the gain-ranging technique which will be discussed later. These modules, about \$250 each in small quantities, are a significant part of the total cost of a card.

3.2.4 Offset Cancellation

The AC coupling following the LPF is used to minimize the DC offset applied to the cascaded amplifier chain. We are often interested in very low frequency signals hence this coupling resides at .01 Hz via a 470 μ f aluminum electrolytic capacitor and a 33.2k metal film resistor. The leakage current of these capacitors is specified at about 15 μ V max but in practice it is orders of magnitude lower, making them effective in this application. The objective is to prevent the offset in the amplifier chain from using a significant portion of the $\pm 10V$ swing available at the higher gain outputs. Most of the remaining offset is subtracted out numerically by the controller. This will be discussed later. A side effect of the very low cutoff of the AC coupling is that it takes a few minutes to allow the offset calibration process to arrive at a "quiescent" level after amplifier turn-on.

3.2.5 Auto-Ranging Amplifier

The cascaded amplifier method was chosen because it offers a gain selection method rather than a gain switching method, eliminating the possibility of overload due to inadequate auto-ranging rate. At high signal levels, the higher gain amplifiers are in saturation but quickly recover when the signal decreases. The choice of opamp here is important. Some amplifiers exhibit particularly disagreeable behavior when subjected to abuse such as this. The Linear Tech 1007's have sufficient gain-bandwidth and slew rate to give minimal skew from one gain selection to the next. Their voltage noise level is very low and current noise is not an issue as the source impedances are also low. An important feature is that the controller selects the gain stage to use based not only on whether its level is within scale but also on having been on scale for a minimum number of previous samples. The amplifiers in the chain are DC coupled. Individual offset nulling is not necessary because of both the low initial offset of the devices and the dynamic nulling described below.

3.2.6 Analog to Digital Conversion

The ADC's are MicroLinear 2230's, atypical in that they are 12 bits plus sign, have an on-board S/H function and are self calibrating. Unfortunately, they draw about 200 mw from $\pm 5V$ power rails. The input swing is specified at $\pm 5V$ but they can withstand $\pm 12V$ swings without damage. The 13 bit word from this part fits conveniently into a 2-byte sample width when 3 bits are used to specify the gain at which the sample was taken.

3.2.7 Controller

An 87C51 micro-controller provides the following functions:

- configuration capability via an onboard UART
- implements various test and diagnostic modes

- runs offset calibrations
- makes the gain selection after all ADC's sample
- numerically nulls the offset in the selected sample based on a stored measured offset value for each gain range
- makes the sample available on the interface data bus

With a 12 MHz clock, the controller can process a sample in 43 μ s. The ADC's sample in 32 μ s (about 30 kHz). If a 16 MHz clock was used, the controller would be perfectly matched with the ADC. At 12 MHz, the maximum sampling rate of the GRA is about 22 kHz which is more than adequate with the maximum bandwidth of 5120 Hz from the LPF. Typically we use a sample rate that is 3 to 4 times the selected LPF cutoff, unless a greater rate is useful for subsequent beamforming operations.

Each card has an address programmed with an 8 bit DIP switch allowing as many as 256 cards to be bused. (A more practical limit due to line driving considerations is probably no more than half that). The host computer uses a standard communications (COM) port via an optically isolated interface to communicate with individual cards. This serial link is bused also so that each card responds only to commands preceded with its unique address. The protocol used is similar to the familiar Serial ASCII Instrumentation Loop (SAIL, IEEE 997-1985) format of preceding commands with "#aa" where "aa" represents a hex address.

3.2.8 Data Transmission

Approximately 75 μ s after reception of a sample pulse, the new 2 byte sample is loaded into latches whose outputs reside on the system data bus (Figure 26). At high sampling rates, the controller is processing the previous sample as the new one is being acquired. There is a "RDY" line that is bused such that when all cards have a new sample, it goes to a high level, indicating to the GRA interface in the host that a sample suite is ready to be off-loaded. To read a sample from a card, the host must issue the card's 8 bit address on the system address bus, which enables the data latch outputs for that particular channel. This transaction will be described in detail later.

3.2.9 Power Requirements, Packaging, Interconnections

The GRA requires 4 power supply voltages. Thirty-six channels in 3 chassis consume 5 amps @ +5V, 3 amps @ -5V, 2.5 amps @ +12V, and 2.5 amps @ -12V. This adds up to about 2.8W per channel. The +5 volt supplies all digital devices and the +5V analog requirements of the ADC's. All other analog power is \pm 12 V because of the input signal swing limitations of the ADC's. Separate digital and analog grounds are run from each chassis of 12 GRA's to the supplies. Power packs were made up of 3 amp +/- 12 volt and 6 amp +/- 5 volt units. Each supply is metered for voltage and current and is liberally cooled. A single power pack can run 3 full chassis of 12 GRA's.

A single 12 channel chassis is a standard 19 inch rack mount, 7 inches high. Since all connections but the analog input to each signal are bused, a printed circuit backplane was designed to accommodate 12 - 100 pin edge connectors and a 50 pin header for a flat cable connection to the rear panel. This is just visible on the left of Figure 42 showing the front view of a chassis. All signals on the backplane have ground runs between them to minimize interference and noise. Particular attention was given to isolation of the analog inputs to minimize 60 Hz contamination. Power supply traces are doubled.

The rear panel connections shown in the photograph of Figure 43 are:

- The 12 isolated BNC inputs at the top are for analog signals. Connections to the card edge connectors are coax.
- The 50 pin blue ribbon connector is the system data/address/control bus. All chassis are connected on the same 50 pin flat cable.
- The power supply connections are made using forked lug crimp connectors on twisted pair, #16 wire. A separate set of power cables from each chassis connects up to 3 to a power pack.
- One chassis in a system must power the optically isolated serial interface from the "+5D" BNC shown. All chassis have the "C" BNC (for communications) bused to the serial interface. The "F" BNC's on all chassis are bused to the sample pulse source. The "A" BNC corresponds to the "DCCAL" input on the GRA cards and is used for DC calibration.

3.3 Software

The 87C51 processor was chosen for the GRA controller not so much for unique characteristics as for previous experience with it in other systems. It is probably one of the 2 best known 8 bit controller types, the second being the 'HC11 family from Motorola and elsewhere.

The code for the GRA function is in native assembly for speed. It was expected that the maximum sampling rate was going to be determined by the rate at which the controller could determine gain and make the offset correction. A core algorithm for determining gain executes in a maximum of 25 μ s. Up to 18 μ s more is required to make the offset correction, make the new sample available and display status in the form of the "RDY" and gain indicator LED's.

About 3.5k bytes of the available 4k bytes of code space are used. About 25% of that is allotted to the sampling-with-continuous-offset-update and fixed gain modes of operation that have not in fact been used much in practice.

GRA functions are invoked in 2 ways: i) when more than 1 card is on the bus, the *System Level Commands* are usually used, whereby the card address is part of the command, ii) when a single card is used, or when it is desirable to quickly output the same command to all cards on the bus, *Global Commands* can be used, whereby no card is specified in the command, and all cards respond in sequence based on their address. In response to all

system commands, a card will respond with 3 ASCII characters; a dash followed by 2 hex digits representing the card's address.

Serial communications are conducted on an interrupt basis so that no time is wasted polling the serial port for commands while acquisition is in process. Since there are few circumstances when interrupts are disabled, there is little chance that an amplifier will fail to recognize a serial command. The default and maximum baud rate is 4800. When a character arrives via the serial port, the interrupt routine checks for a "#". If received, a command is expected and the card address is read. If it's the correct address, the command is read and executed, otherwise, the command is ignored and the search for "#" is resumed.

The features and operational modes supported by the GRA software are described below. In many cases, there is a one-to-one correspondence between System and Global commands. In all descriptions, "aa" is a card's address and the hex characters "Fn" are for *System commands* when preceded by a card address and for *Global commands* without an address).

GRA Functions:

- Soft Reset; System code #aaF0, Global code #F0. All functions of the GRA are reset, however the channel address DIP switch is not reread. (The only way to get that to occur is a hard reset with the manual button on each card, or a power-on reset.) The GRA performs the following functions:
 - reset, clear stack
 - blink the green RDY LED, blink the 4 red GAIN LED's (delayed briefly as a function of the card address)
 - recalibrate the 5 ADC's
 - load a default value of 80 Hz into the LPF
 - de-select the HPF
 - set the PA gain to 1
 - set the baud rate to 4800
 - clear diagnostic codes
 - invoke offset calibration
- Enable for Sampling; System code #aaF1, Global code #F1. After reception of this command, a channel will respond to a falling edge on the sample input by taking a sample. Without this command, a card will not take samples. This is normally the last command issued to cards prior to initiation of sampling. All LED's will be out until sampling starts. During sampling, the green RDY LED glows as a function of sample rate. Cards with higher addresses will show a brighter RDY LED because it is on from the time the sample is ready until it is read by the host computer. If the red ERROR LED turns on, it indicates that at least 1 sample was missed. It remains on until another control command is executed.
- Enable Sampling, Update Offset; System code #aaF2, Global code #F2. This function is similar to the above, however the samples are summed continuously to provide

an updated offset value every 64k samples. All LED's on selected channels will be out until sampling starts. Status LED's will then display as described for command "F1". This mode can be invoked for sample rates as high as 2500 Hz from the standpoint of computational overhead. The benefit over the normal sampling mode is that at lower sampling rates, acquisition runs may be prolonged (many hours), and offset may drift. This technique can accommodate that by regularly summing 64k samples, computing the average and using the residual as the offset value. This mode has not been used because of concern that the variability with as few as 64k samples is nearly .5% of the 64k sum however we think that the concept is reasonable if a factor of 10 or more samples is used.

- Set LPF cutoff; System code #aaF3nn, Global code #F3nn. The LPF cutoff (-3dB point) is programmed using this function. The 2 hex digits nn set the cutoff to (nn+1)20 Hz. The minimum value is 20 Hz and the maximum is 5120 Hz. Once set, the filter remains at that cutoff until this command is used to change it or a hard or soft reset occurs. The 4 gain LED's will blink in response to the Global code but there will be no visible response to the System code.
- Halt Sampling; System code #aaF4 Global code #F4. After this command is received, the card ignores further transitions on the sample input and commences the offset calibration process. No other parameters are changed....the LPF, PA gain, etc are unchanged. The 4 gain LED's will blink in response to the Global code but there will be no visible response to the System code. Neither the red ERROR LED or the green RDY LED should be on. If the red LED is on, it means that 1 or more samples were missed. If the green LED is on it means that when sampling was halted, the last sample was not taken from the channel (not a problem but this should not normally occur). The ERROR LED is the most important indicator of successful operation.
- Set Preamp Gain; System code #aaF5n, Global code #F5n. The on-board fixed PA has 8 gain settings in binary steps. The gain code is the exponent of 2 that equals the resultant gain. Valid codes are n = 0 through 7 corresponding to linear gains of 1 through 128. The 4 gain LED's will blink in response to the Global code but there will be no visible response to the System code.
- Set Baud Rate; System code #aaF6n, Global code #F6n. As with the PA gain set command above, the baud rate is set according to the hex byte code "n". Valid codes are n = 0 through 4 corresponding to 300, 600, 1200, 2400, 4800 baud. The default rate is 4800. After the rate is changed, the host serial port must also change to continue communications. The 4 gain LED's will blink in response to the Global code but there will be no visible response to the System code.
- Set GRA Gain; System code #aaF7n, Global code #F7n. Normally cards are used in the default autoranging mode. They can also be configured to run in fixed gain mode with one of the 5 available gains selected. The codes 1 through 5 correspond to gains of 1, 8, 64, 512, & 4096 respectively. A code of 0 requests ranged gain. The 4 gain LED's will blink in response to the Global code but there will be no visible response to the System code.

- Set Slew Delay; System code #aaF8n, Global code #F8n. The hex value n is the slew delay in number of samples. During auto-ranging, when the level at any of the upper 4 gains is above 82% (upper 3 bits of the word set) that gain cannot be selected. If the level at a gain is not above 82% of full scale, a check is made as to how many samples have occurred since that gain was found to be onscale. This is referred to as the slew delay and defaults to 3 samples. That means that in addition to being below 82% of full scale, a gain must have been on scale for at least the 2 previous samples before it can be selected for the current sample. In other words, gain will be decreased immediately but increased only after a gain range has been onscale for a minimum time related to the sample rate. At a 20kHz sample rate, a delay of 3 means that the output at a given gain range must have been onscale for a minimum of 100 μ s. The 4 gain LED's will blink in response to the Global code but there will be no visible response to the System code.
- Disable Channel; System code #aaF9, Global code #F9. This command is used to deselect channels so that they are effectively removed from the bus during data acquisition operations. In this mode, the green RDY LED is constantly on. During acquisition runs, cards that are not used should be left in this condition. Normally this command is issued to all cards followed by the "F1" or "F2" command to enable those that are to be active during subsequent data acquisition.
- Diagnostic Sampling 1; System code #aaFAn, Global code #F9n. These diagnostic operations are used for test purposes to verify basic functions of a GRA. This command with $n=0$ allows continuous testing of the controller's interaction with the ADC's by configuring the card for normal sampling but latching a count down hex value starting at 0xffff and proceeding to 0 as sample edges arrive. The same command with $n=1$ will cause the card's address to be placed in the output latches after each sample period. The ADC's are setup normally and in fact do sample. The 512X ADC is monitored as usual to know when sampling is complete and its data is read out to enable the next sample.
- Diagnostic Sampling 2; System code #aaFBn, Global code #FBn. For diagnostic purposes to check that offset calibration and sampling of the normal and DCCAL inputs are functional. The following n codes are available:
 1. $n=0$, normal offset calibration process is invoked with the outputs of all 5 gain ranges displayed on the screen with the resultant complemented correction factor shown as well.
 2. $n=1$, continuously display the value of all 5 ADC's with the LPF input shorted. This is a means for making the only offset adjustment on the board which is a pot on the AC coupling amp before the amplifier chain. Once this is done during the initial card setup, it will probably no longer be required. No offset correction value is computed or modified during this test.
 3. $n=2$, sample the standard input at an internal rate, which is slow enough to display...only a couple of Hz. The offset corrected values from all 5 ranges are displayed.

4. $n=3$, sample the standard input at an external rate applied to the sample period input. The offset corrected values from all 5 ranges are displayed.
 5. $n=4$, sample the DCCAL input and display the resultant value after DC offset correction. The DCCAL input is applied to the amplifier chain DC coupled so it can be used to calibrate the ADC's with DC levels at minimum and maximum values.
- Request Pulse Calibration operations; System code #aaFC, Global code #F9. One of the concerns of a multichannel system is the inability to test with confidence the operation and accuracy of all the cards quickly enough that it can be done routinely in the field. It can be difficult to identify a card that is misbehaving. The intention here is to be able to apply a single "boxcar" pulse of some fixed duration and amplitude, to the amplifier input. Samples are acquired by the host at a selected sample rate. At the host, the spectrum of the impulse would be computed and energy levels compared at specific frequencies with a stored replica of what the response should be at one or more LPF settings. In this way a complete system test can be accomplished in a short time that can be used as a "go-nogo" criterion at any time and be a rigorous test of all sections of the GRA.

This feature has never been fully implemented, partly because there is no easy way to get the same pulse input to all channels without operator intervention at the rear panel DCCAL input. Thought has been given to committing the DCCAL input to a replica of an accurate portion of the ADC voltage reference. This would only have to be done to one card in each chassis as the DCCAL input is bused.

3.4 Test Data

Three sets of spectral plots are shown to demonstrate the amplitude characteristics of the GRA. They are: 1) noise floor, 2) dynamic range & harmonic distortion, 3) Lowpass filter. The Y-axis on all plots is dB relative to 1 Volt (dBV), un-normalized "per Hertz". The per Hz correction can be had by subtracting $10\log(\text{bin-bandwidth})$. In all cases the FFT size was 512 complex, therefore with a 500 Hz full scale bandwidth, the dBV numbers are essentially per root Hertz. Appropriate figures are cited in each section.

3.4.1 Noise Floor

The same channel was used for all plots and can be considered representative. The input was DC coupled and open (it is terminated on-board with a 33.2K resistor). The sampling parameters are given on each plot. Normally an average of 8 spectral frames is shown. At a PA gain of 1, the input can swing a maximum of ± 10 volts or approximately 7Vrms (+17 dBV). Spectra of the broadband noise floor at PA gain settings of 1 & 16 are shown in Figures 2 & 3. The LPF was set at 80Hz. Note that the effect of the on-board PA is to position the measurement range of the system relative to the noise floor of the LPF, i.e., if input signal levels will never exceed -20 dBV it may improve data quality to apply up to 24 dB gain at the PA.

Gains of 32, 64 and 128 are available but gain beyond 16 are rarely useful as the inband noise floor doesn't decrease significantly beyond about -140dBV. It can be seen that by themselves, the GRA's do not suffer from 60 Hz contamination. The HPF was not included but has no significant contribution to noise compared to the LPF.

3.4.2 Low Pass Filter Response

Sample spectral plots are shown to demonstrate the LPF amplitude response at 3 cutoffs, 4000 Hz (Figure 4), 260 Hz (Figure 5), and 80 Hz (Figure 6). The input for these tests was digitally generated white noise flat to about 20kHz. Again, these plots are not normalized to 1 Hz bandwidth. The phase characteristics of these filters are classical 8 pole Butterworth with nearly linear group delay to about half the cutoff frequency.

3.4.3 Distortion

A series of tests were run to demonstrate the effect of gain ranging on harmonic distortion as a function of sampling rate, signal level, and signal frequency. The HPF was not connected for these tests. A representative selection of the results are shown. The test signals were tones from a Krohn-Hite model 4400A, specified to have harmonic distortion levels of -100 dB so it is assumed that the plots show the cumulative effects of the GRA signal conditioning. The signals were acquired with a test program using the same routines as the acquisition system discussed later.

As signal amplitude and the ratio of sample rate to signal frequency increases, samples are acquired from more gain ranges. This occurs because there is a greater opportunity for gain increases as a signal passes through zero amplitude. The slew delay for all tests was 3. It can be argued that a greater slew delay could make distortion due to ranging a non-issue. What is shown is the case where very high amplitude tones, low in frequency relative to sample rate, cause maximal gain changes, and hence worst case distortion levels. Similar conditions are not often encountered in real data.

In Figures 7 through 13, a 100 Hz tone is the input signal at levels of approximately 7 Vrms (maximum input signal) down to -60 dBV in 20 dB steps. The LPF is at 260 Hz, the sample rate is 10kHz and the PA gain is set to 1,16 and 128 as labeled. The LPF is specified to have $\leq .03\%$ distortion at 3.5 Vrms. At the maximum signal level, the harmonic components observed beyond the filter cutoff of 260 Hz are not a function of the ranging amplification as they are observed in the fixed gain example of Figure 21. Figure 8 at 100 Hz and 0 dBV, shows the 2nd harmonic about -68dB with ranging and disappears at fixed gain shown in Figure 22. As the signal level decreases, higher gain ranges are used. At the higher signal levels, it is possible that the 36dB gain range is used but usually only the 0 and 18 dB stages are sampled. It is apparent in Figures 12 & 13 that there is little benefit to PA gains beyond 16 since the passband noise floor at 128 is only a few dB lower than at 16. Increasing PA gain beyond 32 simply decreases the measurement range of the amplifier.

Figures 14 through 20 are similar but the signal is at 1 kHz and the LPF is set at 3200 Hz. At 7 Vrms, the highest distortion component is about -64 dB from input. At these high signal levels, the LPF is responsible for the other components that don't appear to

be harmonically related. The plot looks virtually the same at a fixed gain of 1, not shown because the gain at such high signal levels usually remains at unity. Again, there is little difference in the passband between PA gains of 128 and 16.

3.4.4 Crosstalk

Figures 23 and 24 show the effect of crosstalk on a card with an open input adjacent to, and some distance away from a channel with an input at maximum level. There is about a 20 dB peak at the signal frequency in the adjacent channel but it is about 110 dB down from the input on the channel with signal and somewhat smaller on the more distant channel.

4 Acquisition System

In this section, the acquisition system hardware, software, data storage and directory formats are described. The objective of this system design was to be able to acquire and store data at relatively high continuous rates using an inexpensive PC. We wanted to be able to acquire data limited only by the continuous storage bandwidth of the storage device, allowing upgrades as faster devices became available. To permit continuous recording to an archivable medium, it had to be removable and high density.

The system currently has 40 of the GRA channels just discussed, a PC/AT, Industry Standard Architecture (ISA) bus computer, GRA interface and one or more Write-Once Optical Disk (WORM) drives. The system configured as shown in Figure 25 has been in the field and successfully recorded data on 2 occasions; the CEAREX experiment in 1989 and the Heard Island Feasibility Experiment in 1991. The acquisition software was upgraded in 1990 to incorporate realtime display-after-write features.

4.1 Storage Medium

The optical disks are Optimem 1000M drives, using 12 inch, 2 - sided platters storing 1.0 or 1.2 gigabytes (GB, $1 * 10^9$ bytes) per side. The interface is a single ended version of the Small Computer System Interface (SCSI), supported by a \$250 host adapter in the PC. The drives are now about 5 years old but are still among the fastest available, capable of storing data at 420 Kbytes/sec *continuously*. Though the newer Optimem 4000 drive with 2.0 GB per side doesn't offer the highest density (other WORM's can now store up to 3.5 GB per side), it offers about a 50% greater *continuous* storage rate. The drives we are using cost about \$12k and the newer models are about \$16k. The media cost can often be the decisive factor in storage method selection. The Optimem disks cost \$500 to \$700 which is at least 50% greater than the higher density media used by the slower 12" Sony drives.

The Exabyte 8mm tape drive could be used in place of the optical disk drives. The media cost is very low and the new 8mm units have similar storage bandwidth however the random accessibility of the data is very much slower.

4.2 Computer

The ISA bus computer is an inexpensive 286 or 386 machine with a standard 8 MHz bus clock and 8 Mbytes of memory. The GRA controller was designed to use the host direct memory access (DMA) capability which can access extended memory. Many of the cloned AT machines use chip sets that allow the DMA "SCLK" to be increased from 4 to 8 MHz. At this rate a 16 bit DMA transfer is accomplished in $.5 \mu\text{s}$, i.e. up to 4 MB/s transfer rate across the bus to/from 16 megabyte memory space. This, with a SCSI host adapter capable of 2 MB/s transfers to the optical disk makes it possible to achieve 1 MB/s net acquisition bandwidth.

The host communicates with the GRA's directly for configuration purposes via a standard COM port and an optically isolated link. The rationale behind the isolation is the minimization of noise contamination of the analog signals by the host. It is possible that this precaution is unnecessary, which would simplify the serial connection.

4.3 GRA Interface

The GRA data/address/control bus, Figure 26, is connected to the GRA interface, a 16 bit ISA bus card designed specifically for this application. A functional diagram is shown in Figure 28 and a photograph of the card is given in Figure 41. The ISA bus is reproduced in Figure 29 and the connections to the acquisition parameter panel are shown in Figure 27. Figure 30 is a synopsis of pinouts from the GRA to the interface, including connections suitable for use with an earlier model GRA. The interface consists of the following sections shown on the interface schematic in Figure 31.

- A sample buffer implemented in a 1k by 18 bit FIFO
- A sequencer, implemented in a 1k by 18 bit FIFO
- Timing and sequencer control implemented in an ALTERA EP1800 PLD
- IRIG-B time code reader and realtime clock synchronizer
- Acquisition parameter generators and realtime clock implemented in 82C54 counter-timers

The GRA controller is a combination of a first-in first-out (FIFO) based state machine and software that allows the runtime configuration of up to 12 I/O lines to control a data source, normally but not limited to a bank of GRA's. Eight of these lines comprise the GRA address bus that selects a channel for data output. Four lines are used for internal control and 4 are unassigned. The assertion of a high level on the RDY line initiates a sequencer cycle that steps through each active GRA address and clocks 2 byte data values into a 2nd 1k X 18 FIFO that is used as a data buffer between the GRA's and the DMA transfer to host memory. The timing for this is shown in Figures 33 & 34 for 1 channel and "n" channels respectively. The sequencer is clocked at 4 Mhz, allowing 250ns for each state. Each channel is interrogated for its data value in 4 cycles for a maximum transfer rate from

GRA's to interface of nearly 2 MB/s. The coding created by the acquisition program and loaded into the sequencer is shown in Appendix A.

A bank of 3 - 82C54's are programmed at the start of an acquisition run to provide a variety of synchronized parameters. The available functions are:

- A sample period generator with 1 μ s resolution and 64k μ s maximum period with software enable
- A 2nd sample period generator with external input and 64k period maximum with software enable
- "Master" rate with millisecond or second resolution and 64k ms or sec maximum period with software enable
- A "Delay Window from Master" with millisecond or second resolution, and 64k ms or sec maximum period with software enable
- External "Master" input, rising or falling edge
- 1Hz and 1 kHz outputs
- A realtime clock with 1 μ s resolution
- An IRIG-B synchronizing capability allows quick reading of the timecode to set the realtime clock and sync the system to within about 50 μ s of the IRIG source.
- All outputs are buffered and can be programmed for falling or rising leading edge.
- External edge sync capability

IRIG-B timecode is used establish realtime synchronization with millisecond or better accuracy. Compared to the hardware needed to read many BCD digits from an external clock it is much simpler to read IRIG-B and much more likely to be available since it can be obtained from many sources including satellite based systems such as GOES. A software routine was written to read the timecode using the circuit in Figure 32. Synchronization is easily achieved over a wide range of timecode amplitude. Accuracy of the synchronizing routine can be checked and "tweaked" by comparing the timing of 1 Hz clocks between an IRIG source and the acquisition system time base. When synchronized, the leading edge of the 1 ms width pulse at the 1 Hz output should be lined up with the leading edge of the 1 Hz output from the synchronizing source, as well as the rising edge of the GRA interface 1 kHz output.

Typically, a rubidium timebase is used to assure stability. A UPS (battery backed power supply) is used to run the timebase and the source of the IRIG standard if retention of absolute time sync is critical.

The Altera EP1800 PLD was used to eliminate much of the "glue" logic needed to implement many of the control functions. This 48 cell part replaced about 30 devices and allowed easy design modification, important to a unique system. Without doubt there is a cost in time and effort to learn use of the PLD but in the long run it was clearly worthwhile. Figures 35, 36, & 37 represent the functional internals of the part as programmed.

Synchronous operation is the preferred mode for these devices and it is difficult to achieve timing relationships relying on device propagation delays as one can do with discrete gates. In this design there are 2 unused cells and 3 partially committed. The real advantage is that this controller board could be printed (the present unit is only wire-wrapped) yet remain flexible because even the hardware is somewhat programmable. Figures 35 and 36 are primarily devoted to sequencer, DMA reload and realtime control. The circuitry in Figure 37 is related to acquisition parameter control.

4.4 Acquisition Software

The acquisition software is entirely in C, using Borland's Turbo C, and consists of a main program, and 3 sets of routines related to acquisition functions, SCSI transaction functions and graphics functions for realtime displays. There are about 2000 lines of code. In addition a number of utilities have been written for test and diagnostic purposes and for data quality verification. The challenges of the acquisition task were:

1. conduct the data acquisition and recording in a pseudo background mode since it's beyond the direct response capability of the processor and DOS
2. maintain precise time stamping of and checking for incomplete data packets
3. under most conditions, permit some modest display of data for quality assurance.

The software configurable aspects of the GRA controller are key to satisfying these requirements.

An ASCII file that is the template for an acquisition run is edited beforehand to contain all necessary information about all GRA channels to be used. Included is information about sample rate, filtering, sensors, the number of channels connected and which are to be used during the pending recording session. The program reads this file, displays portions on the screen for verification, configures the GRA's, programs the acquisition parameters and the sequencer on the GRA controller, synchronizes to the IRIG timecode if requested and awaits a keyboard hit to start sampling. It can be set up so that everything is done by invoking a single short batch file. Figure 38 is an abbreviated version of this template. The acquisition program keys on the colons to read parameters.

The host DMA and interrupts are used to handle the time critical functions of the process. DMA addressing includes use of a page register that permits access to the extended memory region beyond 1 MB. One DMA channel is used by the SCSI host adapter for transferring data from memory to disk and one is used by the GRA controller for writing data to memory. To have sufficient compliance with SCSI transfers, multiple buffers are setup in memory with arbitrarily large size. An important motivation for large buffer space is that the Optimum drives may execute an auto-focus procedure at any time after an hour's operation, that takes up to 1 second during which no SCSI transfers take place. Buffers can be put anywhere in memory but usually larger is better and 4 to 6 - 1 megabyte (MB) chunks are used. The DMA facility of the PC/AT is constrained to a maximum of 64k words per transfer, 128 kilobytes, (KB) so the DMA registers must be reloaded 8 times to fill a 1 MB buffer. To accomplish this while minimizing host interaction, an interrupt

routine is executed when the terminal count (T/C) line on the AT bus goes active. This routine reloads the DMA registers with predetermined pointers and updates pointers to the next sub-buffer. When a "last" sub-buffer is enabled, this interrupt routine sets a flipflop so that at the next T/C , a second interrupt routine reads the realtime clock to get the time to the μs of the first sample of the next 1 MB buffer.

The 1 MB buffers are filled circularly. The foreground activity of the program is to monitor buffer changes and initiate buffer writes, again via DMA, to the optical disk. At each buffer write, status information is written to the screen as to record number, start time of buffer, which buffer was written and which is being filled. The elasticity of the FIFO which buffers data from the GRA's relaxes the timing requirements of the interrupt driven routines. Transfers to memory from the FIFO occur whenever it is \geq half full. This arrangement was accomplished without resorting to timing-critical software, no assembly language, and has worked flawlessly to date.

For data recording well below maximum rates, the acquisition program can be used in a graphics mode that allows 1 channel at a time to be extracted from the last record written to the optical disk and a time series segment or a spectral average displayed on the screen. During the Heard Island experiment a simple delay & sum beam was formed and spectra displayed using data from a vertical array. Because of the large amount of buffering available in extended memory space, these displays were achieved with data rates as high as 320 KB/sec. At this rate however display can only occur about once every minute since time must be allowed for buffer write operations to "catch up".

4.5 Optical Disk Storage Details

An earlier WHOI technical report #WHOI-87-49, described the initial application of the Optimem disk drives to our data acquisition efforts. The formats used to store data on the disk and the directory structures developed were discussed. They are *not compatible* with any file system and therefore must be read *raw* for analysis. Essentially, those methods continue to be used with, however, a few important but compatible changes which will be noted.

4.5.1 Optical Disk Header

Each optical disk starts and ends with a disk header that is one sector long. There is normally a disk header sector at the first logical block (LBA 0) and the last logical block, (LBA 999999 for 1.0 GB per side disks and LBA 1172499 for 1.2 GB per side disks). The "C" structure of the optical *Disk Header* is shown below. Note that "longs" are 4 byte integers. A sample disk header is printed out in Appendix C.

```
struct disk_header
{
    unsigned char  dhkey[4];      /* disk header key, "ODHD" */
    unsigned char  sn[16];       /* platter S/N, 15 chars, leading */
                                /* zeroes, terminated with null */
}
```

