

MIDAS-W: a workstation-based incoherent scatter radar data acquisition system

J. M. Holt¹, P. J. Erickson¹, A. M. Gorczyca¹, T. Grydeland²

¹ Massachusetts Institute of Technology, Haystack Observatory, Westford, MA 01886, USA

² University of Tromsø, Tromsø, Norway

Received: 28 December 1999 / Revised: 19 June 2000 / Accepted: 21 June 2000

Abstract. The Millstone Hill Incoherent Scatter Data Acquisition System (MIDAS) is based on an abstract model of an incoherent scatter radar. This model is implemented in a hierarchical software system, which serves to isolate hardware and low-level software implementation details from higher levels of the system. Inherent in this is the idea that implementation details can easily be changed in response to technological advances. MIDAS is an evolutionary system, and the MIDAS hardware has, in fact, evolved while the basic software model has remained unchanged. From the earliest days of MIDAS, it was realized that some functions implemented in specialized hardware might eventually be implemented by software in a general-purpose computer. MIDAS-W is the realization of this concept. The core component of MIDAS-W is a Sun Microsystems UltraSparc 10 workstation equipped with an Ultrarad 1280 PCI bus analog to digital (A/D) converter board. In the current implementation, a 2.25 MHz intermediate frequency (IF) is bandpass sampled at 1 μ s intervals and these samples are multicast over a high-speed Ethernet which serves as a raw data bus. A second workstation receives the samples, converts them to filtered, decimated, complex baseband samples and computes the lag-profile matrix of the decimated samples. Overall performance is approximately ten times better than the previous MIDAS system, which utilizes a custom digital filtering module and array processor based correlator. A major advantage of MIDAS-W is its flexibility. A portable, single-workstation data acquisition system can be implemented by moving the software receiver and correlator programs to the workstation with the A/D converter. When the data samples are multicast, additional data processing systems, for example for raw data recording, can be implemented simply by adding another workstation with suitable software to the high-speed network. Testing of new data processing software is also greatly

simplified, because a workstation with the new software can be added to the network without impacting the production system. MIDAS-W has been operated in parallel with the existing MIDAS-1 system to verify that incoherent scatter measurements by the two systems agree. MIDAS-W has also been used in a high-bandwidth mode to collect data on the November, 1999, Leonid meteor shower.

Key words: Electromagnetics (instruments and techniques; signal processing and adaptive antennas) – Ionosphere (instruments and techniques)

1 Introduction

The Millstone Incoherent Scatter Data Acquisition System (MIDAS) is based on an abstract model of a general incoherent scatter radar. This model is implemented in a hierarchical software system which serves both to isolate different levels of the software from each other, and to hide hardware implementation details from higher levels of the system. Inherent in this is the idea that portions of the software and hardware can easily be changed in response to technological advances, and the MIDAS software and hardware has, in fact, evolved while the basic software model has remained unchanged. MIDAS-1, the production system at Millstone, has undergone numerous changes since it became operational in 1992, and MIDAS-C, which was built for bistatic observations using the Algonquin Observatory antenna in Canada, included further significant upgrades.

However, the digital signal processing sections of MIDAS-1 and MIDAS-C, which compute lag profiles (Lehtinen, 1986; Holt *et al.*, 1992) from digitized samples, have remained largely unchanged. The received incoherent scatter signal is converted to a low intermediate frequency (IF), which is two or four times the final complex sampling rate. This low IF is digitized by a custom analog to digital (A/D) module, and then filtered

and converted to baseband complex samples by a custom DSP module. The only major difference between MIDAS-1 and MIDAS-C is that MIDAS-C uses a newer, faster array processor.

From the beginning, it was realized that many MIDAS signal processing functions could eventually be implemented in software. MIDAS-W is the realization of this concept. The digital processing "section" of MIDAS-W consists of a commercial A/D board in a Unix Workstation, along with several additional workstations which communicate with each other over a dedicated high-speed local area network (LAN). The A/D board samples an IF at a fixed rate, and these samples are retransmitted in digital form on the LAN. With appropriate software, any computer on the network becomes a receiver. All functions handled by DSP hardware in MIDAS-1, such as filtering, decimation and correlation are carried out in software on one or more of the workstations. Additional functions, such as storing or displaying the raw samples, can also be easily implemented in software.

Digital receivers which sample an intermediate frequency and process this signal in programmable digital hardware are now very common. For example, this technique is used in cellular phones where costly and bulky analog hardware must be minimized, and where multiple standards (for example, TDMA, GSM and CDMA) favor systems which can be adapted to different standards through software changes. MIDAS-W differs from most digital receivers in two respects. First, all processing is carried out in general-purpose computers rather than specialized digital signal processing (DSP) hardware. Second, the digitized signal is rebroadcast on a LAN for processing by as many computers as needed by a particular application. This provides a high degree of flexibility, which is ideally suited to a research environment such as an incoherent scatter radar. Unit engineering costs, which are very high for low-volume systems, are also greatly reduced. The system is also very easy to upgrade, simply by changing workstations or the network as newer, faster hardware becomes available.

2 MIDAS software architecture

The core of the MIDAS software system is an abstract model of a general incoherent scatter radar. The higher-level MIDAS software interacts only with the abstract radar, which permits changes to be made in hardware and associated low-level software with little or no impact on the high-level software. The abstract model is embodied in a data structure, known as a command/status or CMST structure, which is a complete description of a radar integration period.

The CMST structure allows a number of important possibilities to be realized. For example, the structure can be used not only as a data record, but also as a specification for an integration period, that is, a radar command structure. In addition, by updating parameters such as antenna pointing angles and transmitter

peak power during an integration period, the structure can be used as a radar status specification. If properly implemented, the data structure approach also provides several other advantages. A distributed data acquisition system can be implemented by sending the structure to the various processors and processes which control the radar. And, if the structures are implemented in a machine-independent fashion, the resulting structure can be interpreted and processed by any computer, eliminating the common problem of having to convert a data tape to the format of a computer other than that on which the data were collected. In MIDAS, these goals have been achieved by implementing the structures in External Data Representation (XDR) language. XDR was developed by Sun Microsystems and is the basis of their Network File System (NFS), which is a *de facto* standard in the computer industry. The software libraries needed to implement XDR are available on all modern computers, and are in the public domain, making it easy to port the system to new processors if necessary. This public domain software also makes it easy to define remote procedure calls (RPC), which allow the structures to be passed back and forth between separate processes, between separate closely coupled processors, between computers on a local network, and between computers which can communicate via the Internet or serial communications lines.

The structure contains three parts. First there is a header containing basic information such as the site ID. Like the other parts of the structure, the header includes a variable length character string which could contain notes, log excerpts, source-code or any other desired information. The second part of the structure is the integration period specification. This contains all information required by the radar to execute the integration period, such as antenna pointing commands, timing information, waveforms etc. The third part of the specification contains the integration period data, such as actual antenna pointing directions, transmitted power, and the measured lag-profiles.

The structures are actually encoded in RPC language. This is a superset of XDR language which defines the structure itself. The full RPC language includes additional features needed to define the remote procedure calls used to transfer the structures between processes. XDR structures are also continuous data streams, unlike C language recursive structures which are spread throughout computer memory by system memory allocation routines. This permits the XDR structures to be stored on tape or disk as data records, and to be placed in shared memory regions to be shared between processes. At the end of each integration period, the entire CMST structure is written to disk. It is then immediately available for further processing.

Figure 1 shows the high-level software model for Millstone Hill Observatory data acquisition. The hierarchical organization of the software system serves to isolate the lower levels of the system, which are responsible for time-critical tasks, from unpredictable demands from higher levels. This organization also performs an important security function by isolating

MIDAS SOFTWARE MODEL

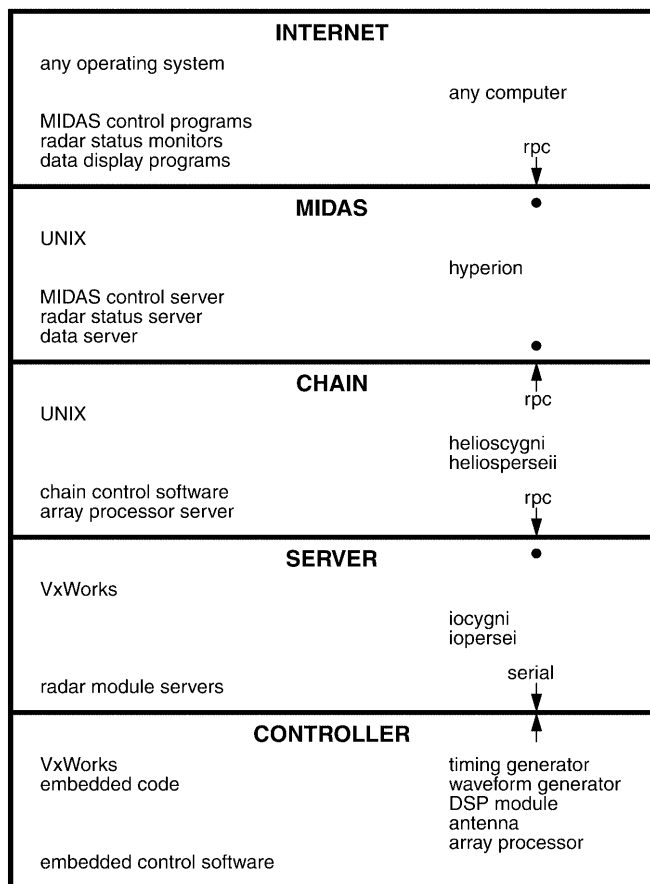


Fig. 1. Detailed model of the high-level software used in MIDAS, showing its hierarchical structure (see text for further details)

radar hardware such as transmitters and antennas from unauthorized control.

The lowest level of this model is known as the controller level and consists of the software which runs on individual radar modules such as the timing generator or correlator. The code at this level is module specific. The timing and waveform generators are controlled by small computers running a compact real-time operating system such as VxWorks. The MIDAS-1 correlator is an array processor programmed in a specialized array processor language supplied by the manufacturer, whereas the MIDAS-W correlator is implemented in C-language software on a workstation.

The server level contains software modules which provide an interface between the controller level and higher levels of the hierarchy. These radar module servers hide the details of the radar hardware modules from higher levels of the software system. Thus this layer is a kind of glue, connecting the general description of radar experiments which exists at higher levels from the specific details of the radar itself. This is how the differences in the MIDAS-1 and MIDAS-W correlators are completely hidden from higher levels.

The chain level contains the controlling software for an individual receiver chain. Processes at this level

communicate with the server and observatory levels via remote procedure calls to the higher and lower levels. Thus the chain level software has complete control of all activity at that level and is protected from several outside processes simultaneously seeking status information which might interfere with the server’s primary duty which is to control the chain. Critical radar functions such as transmitter timing are controlled from this level and can only be changed by someone with direct access to this level. For example, when an experiment is started, the chain level software is given a list of permissible radar timing options. Timing options may then be selected from this list by a process running at a higher level of the hierarchy. Several receiver chains may run simultaneously under control of a single instance of the observatory level software. These chains might process different frequency bands during experiments which utilize more than one transmitter frequency.

The observatory level contains RPC servers which provide a buffer between the chain level and the Internet level. Experiment descriptions, experiment status, and data from the experiment are stored at this level. Experiment control and monitoring programs interact with the system by making RPC calls to servers which reside at the observatory level. Several instances of the observatory level software may run simultaneously, as during bistatic observations involving both MIDAS-1 and MIDAS-C. This mode will also be used during the transition period from MIDAS-1 to MIDAS-W, when both systems will collect data.

The highest level of the software hierarchy is the Internet level. This software can run on any computer on the observatory network or any computer connected to the observatory via the Internet. Functions include high-level experiment control, experiment status monitoring and data browsing. Monitoring and browsing are available without restriction to anyone on the Internet. Experiment control must be protected by passwords or other means, and is limited to “benign” control functions such as specifying the start time of a predefined radar operating mode. Most of our Internet level software was originally custom-built, and ran only on Sun workstations. Today, the Internet software is usually a Web browser running JAVA programs or communicating with a Web server running at the Observatory level.

3 Digital receivers

MIDAS-1 implements a form of bandpass sampling. In a traditional analog quadrature detection scheme (Fig. 2), a real bandpass signal is converted to baseband by multiplying it by two sinusoids which are 90° out of phase. The two resulting signals are the real (in-phase) and imaginary (quadrature) components of the complex baseband signal. These signals are then low-pass filtered and digitized by two A/D converters. In MIDAS-1, on the other hand, the real bandpass signal is converted to a IF, which is directly sampled by a single A/D converter (Fig. 3). In the figure, the low IF is 50 kHz, which is




Fig. 2. Bandpass sampling by a traditional analog quadrature detection scheme (see text for explanation)

commonly used for single-pulse experiments at Millstone Hill. Several other low IFs and sampling rates are also supported. After down conversion, the signal is passed through a lowpass filter. This filter is actually redundant in the current version of MIDAS-1, because the output of the 2 MHz IF is passed through a 90 kHz bandpass filter prior to being converted to 50 kHz. Had a suitable tunable bandpass filter been available when MIDAS-1 was constructed, that could have been used in place of the tunable lowpass filter, and the bandpass filter at 2 MHz eliminated. As we shall see later, the key requirement is that the real signal be bandlimited to a bandwidth equal to the Nyquist frequency corresponding to the A/D sampling rate. In Fig. 3, the data are sampled at 5 μ s and the signal must be bandlimited to 100 kHz. The A/D samples are then passed through a pair of finite impulse response (FIR) filters which produce the in-phase and quadrature components and band limit the complex signal to 25 kHz. Finally, the samples are decimated by four to produce the final complex samples at 20 μ s intervals.

The MIDAS-1 scheme involves complex, specialized hardware, and also suffers from a lack of flexibility. For example, the FIR filter length is limited to 32 taps. Given the impressive speed of modern workstations, we set out to determine whether a more flexible, higher-performance workstation-based approach was possible. Bose *et al.* (1999) have followed a similar approach in demonstrating a workstation-based communications system, which they refer to as a virtual radio. Sun Microsystems UltraSPARC workstations were selected because of their impressive floating point performance and our familiarity with Sun workstations and the Solaris operating system. Other choices, such as Linux, are possible given suitable A/D board drivers.

The Nyquist Sampling Theorem states that a band-limited signal with bandwidth less than F Hz is uniquely determined by its values at uniform intervals less than 1/2F seconds apart. For example, a signal limited to frequencies below 500 kHz, can be reconstructed from samples at 1 μ s intervals. If frequencies outside the Nyquist limit are present, aliasing occurs and the sampled signal is ambiguous. As illustrated in Fig. 4, if a signal limited to frequencies below 500 kHz is shifted upward by an integral multiple of 500 kHz, it will yield the same 1 μ s sampled values as the original signal. Put another way, the spectrum of the sampled signal is periodic. When a lowpass signal is sampled, it is customarily first passed through a lowpass filter with a bandwidth equal to or somewhat less than the Nyquist frequency to ensure that aliasing does not occur. Any frequencies above the Nyquist frequency are then lost, but the lower frequency components are recovered unambiguously. In the case of incoherent scatter radar, the returned signal is usually narrower than the Nyquist bandwidth, and is unaffected by the filter. However the filter is still necessary to remove interfering signals and higher frequency noise, which would decrease the signal-to-noise ratio if the filter were omitted. Alternatively, if




Fig. 3. The MIDAS-1 bandpass sampling scheme (see text for explanation)

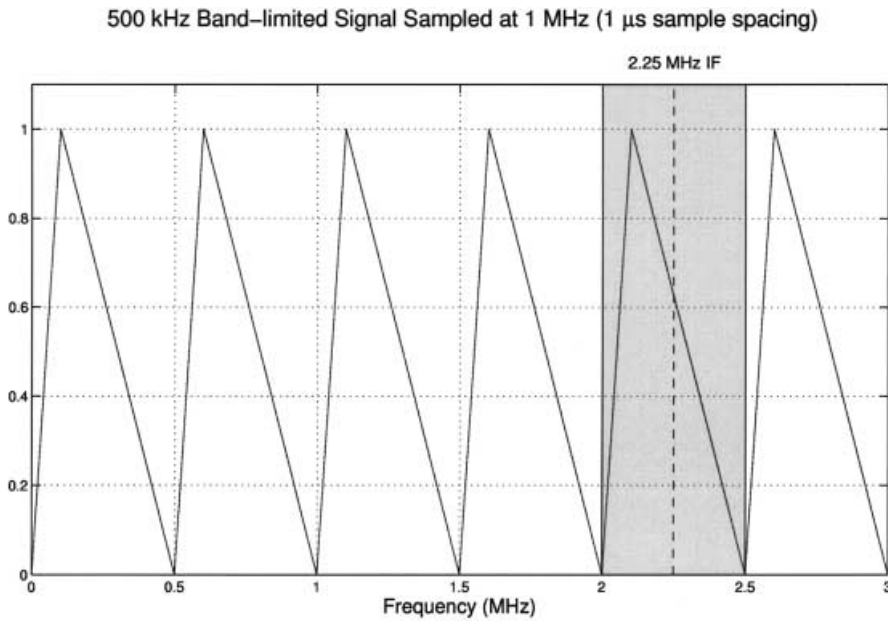


Fig. 4. The 500 kHz band-limited signal sampled at 1 MHz (1 μs sample spacing) as explained in the text

the signal is limited to one of the higher-order Nyquist bands, one can sample it directly with the same sampling interval as if the signal were lowpass. When such a bandpass signal is sampled, it is customarily first passed through a bandpass filter with a bandwidth equal to or somewhat less than the Nyquist frequency to ensure that aliasing does not occur. In the example shown in Fig. 4, if the signal is confined to frequencies between 2 and 2.5 MHz, it can be recovered by sampling at 1 μs intervals. Thus in the case of an incoherent scatter radar, we could digitize the signal by down converting the radar return to 2.25 MHz, passing this signal through a 500 KHz bandpass filter centered at 2.25 MHz, and then sampling at a rate of 1 MHz.

MIDAS-1 uses a 2 MHz second IF. It would be convenient to directly sample this while testing the bandpass sampling technique described here in parallel with the current data acquisition system. Unfortunately, this is not possible to do with a convenient sampling rate due to the relationship between sampling rate and bandwidth illustrated in Table 1. The IF must be an odd multiple of 1/2 the bandwidth, or equivalently, the sampling rate must be of the form $T_s = n_{odd}/4f_c$, where f_c is the center frequency. As shown in Table 1, none of the possible sampling rates are equal to or can be

Table 1. Relationship between sampling rate and bandwidth

n_{odd}	T_s (μs)
1	0.125
3	0.375
5	0.625
7	0.875
9	1.125
11	1.375
13	1.625
15	1.875
17	2.125
19	2.375

decimated to a sampling interval which is an integral number of microseconds.

The 1 MHz sampling rate used in these examples is very convenient because it is an integer multiple of any sampling rate whose sample interval is an integral number of microseconds. An incoherent scatter data acquisition system which sampled at 1 MHz and then decimated these 1 μs samples to any specified final sampling rate would be highly flexible. This sampling rate is also very well-suited to the Millstone Hill UHF transmitter, which has a 1 μs rise time. Finally, a 1 MHz sampling rate is well within the capabilities of current workstations and networks.

Therefore, we selected this sampling rate for the initial implementation of MIDAS-W. A block diagram of MIDAS-W is shown in Fig. 5. As discussed later, the A/D board used in MIDAS-W is capable of sampling at rates up to 80 MHz, and in the near future we hope to increase the sampling rate to 2–4 MHz, sampled at the 1st IF near 30 MHz.

It is worth noting that the actual Doppler bandwidth of the incoherent scatter ion line is usually less than 10 kHz at 440 kHz. The 500 kHz MIDAS-W bandwidth is needed to accommodate the maximum bandwidth of the transmitted waveform, and hence to achieve the best possible range resolution. In the future we intend to use MIDAS-W for incoherent scatter plasma line observations. The bandwidth of a plasma-line profile can exceed 10 MHz, which would require a sampling rate of 20 MHz. The MIDAS-W A/D board is already capable of sampling at this rate and general-purpose computers should be capable of the necessary processing within a few years. In the meantime, it will be possible to measure the plasma line over a limited altitude range, for example, near the F-region peak.

In order to verify the MIDAS-W design, we first wrote a MATLAB script demonstrating a practical algorithm for a software receiver and correlator. First,




Fig. 5. MIDAS-W layout as explained in the text

lowpass FIR filter coefficients and simulated data are read from disk files. The data are simulated incoherent scatter data generated by passing Gaussian white noise through a filter and adding white noise. This signal is sampled at $1 \mu\text{s}$ intervals and written to disk. The script simulates an experiment with 200 sweeps, each containing 5000 samples (corresponding to 750 km). The data from each sweep are read and converted to $20 \mu\text{s}$ complex baseband samples. The algorithm used to accomplish this is equivalent to but significantly different from what would be done in dedicated DSP hardware. First, since only $1/20$ of the samples will be retained, these samples are the only ones processed. The MIDAS-1 DSP module, on the other hand, computes all the filtered samples, even though three out of four will not be retained. This comes at no cost, because of the pipelining used by the FIR filter chip in the DSP module. In a general purpose computer, the total number of floating point operations is what counts, and logic can be utilized to ensure that only needed computations are carried out. Conversion of the real signal to its complex baseband form is accomplished as part of the filtering operation. When the complex exponential which converts the signal to baseband is sampled at $1 \mu\text{s}$ intervals, its values are 90° apart, and take on the values:

$0.0 + 1.0j$
 $-1.0 + 0.0j$
 $0.0 - 1.0j$
 $1.0 + 0.0j$
 $0.0 + 1.0j$
 $-1.0 + 0.0j$
 $0.0 - 1.0j$
 $1.0 + 0.0j$
 $\dots\dots\dots$

As a result, the multiply and accumulates required by the FIR filter need only be carried out once for each filter coefficient. In the MIDAS-1 DSP module, on the other hand, all operations are carried out for both the in-phase and quadrature components. In the workstation implementation, four separate sums are computed and then two are combined to form the in-phase component and two are combined to form the quadrature component. This involves a fair amount of bookkeeping, but that is easy and inexpensive on a general-purpose computer. In the future it will be easy to implement more novel signal processing schemes

which would be difficult to accomplish on special-purpose DSP hardware.

4 Simulation benchmarks

A C-language receiver program was written which implements the same filtering algorithm as the Matlab script. This was used to test the performance of the software receiver algorithms, and was also the basis of the software receiver in the MIDAS-W system itself.

Table 2 shows the performance of the simulated software receiver for several combinations of pulse length, lag spacing and maximum lag. In all cases, an FIR filter with 129 coefficients was used to filter the $1 \mu\text{s}$ samples prior to decimation. Performance was measured in terms of the average Mflops (million floating point operations per second, where a flop is taken to be a multiply and accumulate) achieved and the maximum possible duty cycle. A standard $320 \mu\text{s}$ pulse with $20 \mu\text{s}$ lag spacing achieves a duty cycle of 27.1%, more than four times the maximum possible for the Millstone transmitter. Longer pulses yield a higher duty cycle for a given number of lags calculated, because the pulse repetition frequency (PRF) is lower. Even when lags are calculated out to the pulse length, the program has an ample performance margin. As the lag spacing decreases, the performance margin decreases. Nevertheless, the maximum duty cycle exceeds 6% in every case except a $320 \mu\text{s}$ pulse with $4 \mu\text{s}$ lag spacing. Short lag-spacing (wide-frequency) modes like this would normally be used

Table 2. Software receiver performance

Pulse length (μs)	Lag spacing (μs)	Maximum lag (μs)	LP filter coefficients	Mflops	Maximum duty cycle (%)
320.0	4.0	320.0	129	71.0	4.2
640.0	4.0	320.0	129	71.2	8.5
640.0	4.0	640.0	129	89.3	6.6
2000.0	4.0	320.0	129	71.2	26.5
2000.0	4.0	640.0	129	88.8	20.6
320.0	8.0	320.0	129	56.4	10.3
640.0	8.0	320.0	129	55.9	20.4
640.0	8.0	640.0	129	67.3	16.8
2000.0	8.0	320.0	129	55.7	63.5
2000.0	8.0	1000.0	129	77.0	46.4
320.0	20.0	320.0	129	40.4	27.1
640.0	20.0	320.0	129	40.4	54.2
640.0	20.0	640.0	129	46.3	48.9
2000.0	20.0	320.0	129	40.1	168.1
2000.0	20.0	640.0	129	47.0	155.0

with longer pulses, which have higher maximum duty cycles.

For the 320 μs pulse with 20 μs lag spacing mode, the number of operations required for filtering and correlating are approximately equal. As the lag spacing decreases, the correlation takes up an increasing proportion of the available time. In fact, the filtering can be essentially independent of lag spacing, since while more final samples must be computed, the filter length can become proportionally smaller.

5 MIDAS-W implementation

The core component of the initial implementation of MIDAS-W is a Sun Microsystems Ultra 10 with a 333 MHz UltraSPARC-IIi CPU, equipped with an Ultraview Corporation Ultrarad-1280 PCI bus A/D converter board. A 2.25 MHz IF is bandpass sampled at 1 μs intervals, and these samples are multicast over a 100 Mb/s fast Ethernet for further processing.

The Ultrarad-1280 includes two simultaneously sampling 12-bit A/D converters, two simultaneously updated 12-bit D/A converters and 4 million 32-bit words (16 Megabytes) of dual-ported shared memory. All A/D converters are capable of operating simultaneously at up to 40 Megasamples per second each. The two A/D channels can be combined into a single 80 Megasamples per second channel. Sampling is controlled by a high-speed automatic timebase which is software settable in increments of 25 ns. The board includes 8 TTL input lines which are simultaneously sampled at the A/D rate. These are used to provide radar timing signals to the board. The D/A converters and eight associated TTL output lines potentially allow the board to be used as a waveform generator and radar controller.

Since the MIDAS-1 system still provides timing and control signals for radar operation, we utilize these in the current development system to provide MIDAS-W with contextual information. (Eventually, this data will be generated by MIDAS-W itself when it assumes control functions.) Six TTL level MIDAS-1 signals, sampled simultaneously with the analog IF signal by the Ultraview ADDA board, provide information about the start of the transmitted pulse, mode sequence number (for coded pulses), transmitter RF blanking, receiver antenna sense, and timing mode. Time-stamping of each received pulse is done automatically by the Ultraview board, through using the transmitted pulse start signal to serially write a hardware clock value in the ADDA circular sample buffer. This allows MIDAS-W to sense any hardware timing problems which might cause a missed pulse.

Within the MIDAS-W framework, client programs must also be provided with the contextual information contained within the TTL timing and control signals in order to properly sort and process the IF signal. We therefore prepend a short IF data header to the 1 μs analog IF sample block before multicasting occurs over the 100 Mb/s network. For each pulse, the IF header describes the absolute and relative pulse number,

transmitter mode, RF blank status, receiver configuration, receiving antenna pointing, and sampled data layout. Accommodations are made for multiple, non-contiguous sampling rasters as well as for compact packing of lower bit resolution samples.

A block diagram of MIDAS-W is shown in Fig. 6. It includes the following software components:

- a. *Distributor* that reads samples from the ADC, identifies pulse boundaries, forms and distributes data packets over the local, high-speed network.
- b. *Rti, scope* that reads data from the network and displays either individual IPPs as a scatterplot (scope) or range/time/intensity as a color plot (rti). Scope uses a prepackaged Fermilabs widget, while rti uses the XGL accelerated graphics toolkit.
- c. *Correlator* that reads data from the network, performs digital filtering mixing and decimation, forms lag profile matrices and integrates for a specified period. Writes results to a file.
- d. *Integrator* that is same as correlator, except that it reads data directly from a disk file in which dumper has previously recorded data.
- e. *Dumper* that records all the samples that appear on the network directly to a file.
- f. *Replay* that plays back samples as recorded by dumper. This enables us to "repeat" an experiment. This is useful for development and testing.

6 Results

The high measurement bandwidth of MIDAS-W compared with MIDAS-1 will greatly extend the capabilities of the Millstone Radar. Incoherent scatter measurements will still constitute the backbone of the system, and it is important to verify that this capability has been correctly implemented. In section A, initial MIDAS-W incoherent scatter measurements are presented and compared with simultaneous MIDAS-1 measurements. In section B, measurements made during the 1999 Leonid meteor shower are presented. These measurements would not have been possible with MIDAS-1, and illustrate the expanded capabilities of MIDAS-W.

A MIDAS-W incoherent scatter measurements

Simultaneous observations by MIDAS-1 and MIDAS-W were carried out around noon on 19 November, 1999. The measurements utilized the Millstone Hill 68 m zenith antenna and 300 μs pulses. The receiver 30 MHz intermediate frequency was split and down converted to 2 MHz for processing by MIDAS-1 and 2.25 MHz for processing by MIDAS-W. The MIDAS-1 measurements utilized a fixed 30 s integration period. The MIDAS-W 1 μs samples were saved to disk for later processing.

Figure 7 shows the MIDAS-W normalized lag-profile matrix resulting from filtering and decimating the raw measurements to 20 μs samples and computing lag products from these samples. The matrix has the

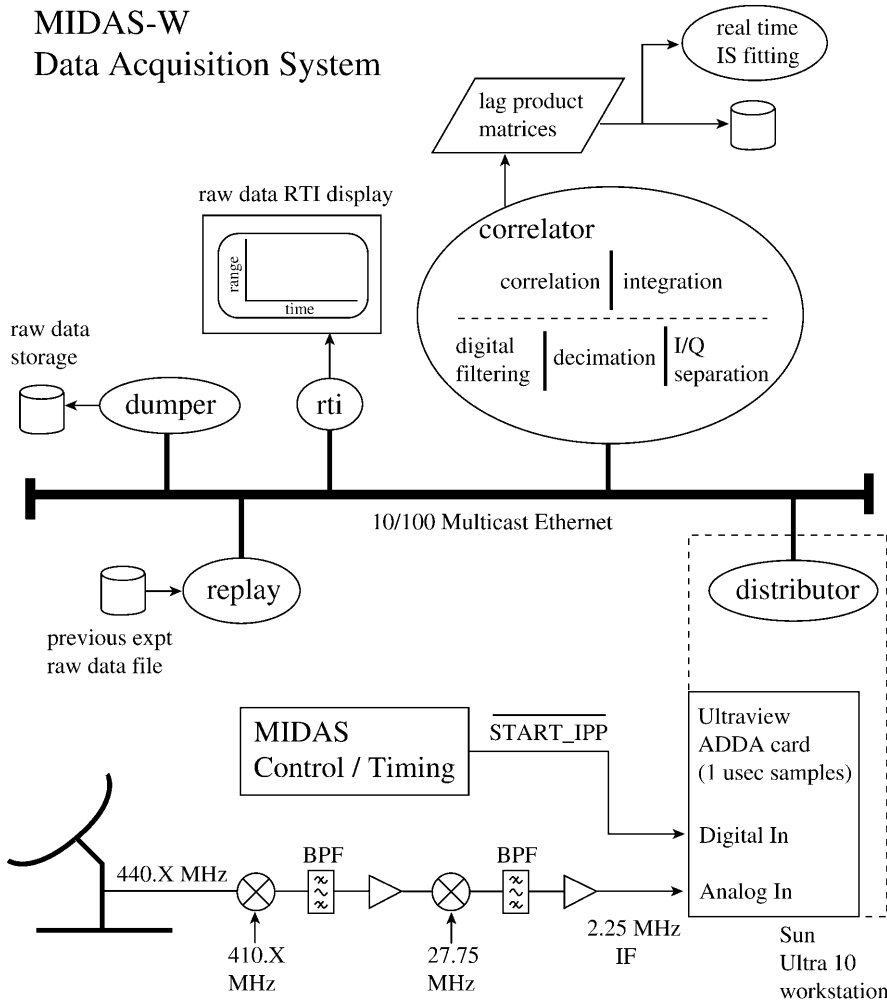


Fig. 6. Complete layout of the MIDAS-W system (see text for explanation of terms used)

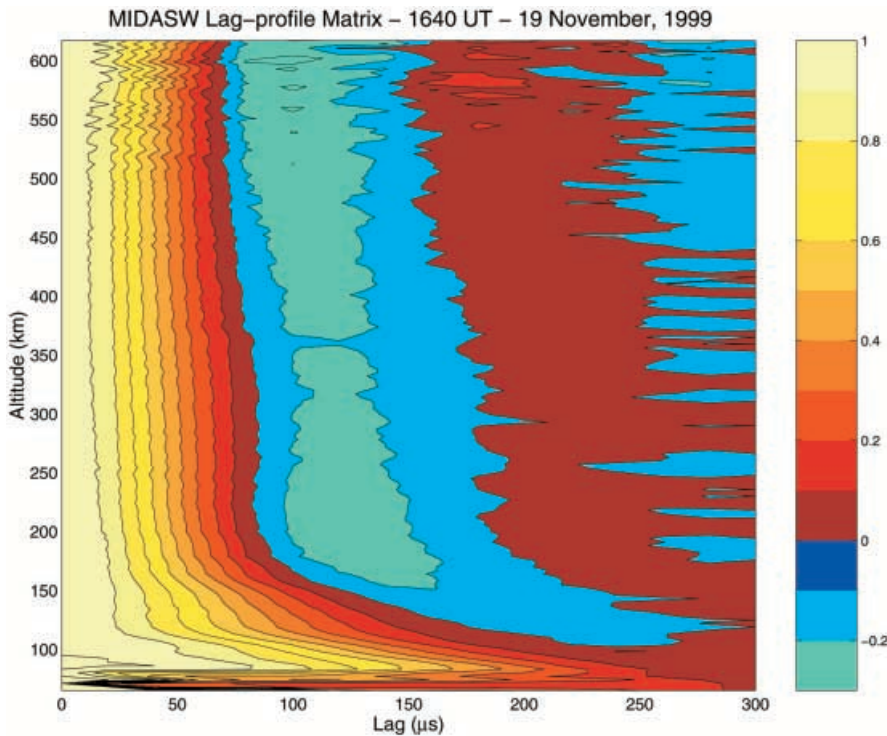


Fig. 7. Actual example of MIDAS-W normalized lag-profile matrix for 1640 UT on 19 November, 1999. (For explanation, see text)

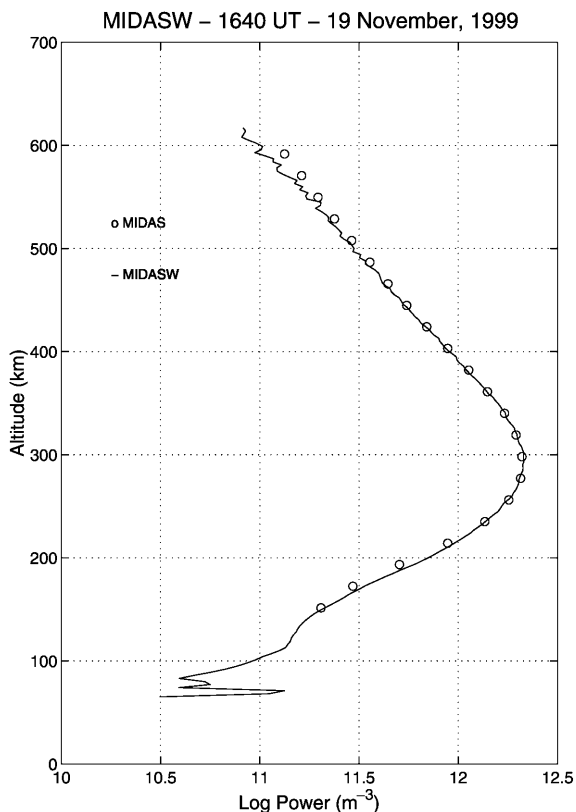


Fig. 8. Comparison of MIDAS-1 and MIDAS-W power profiles (see text)

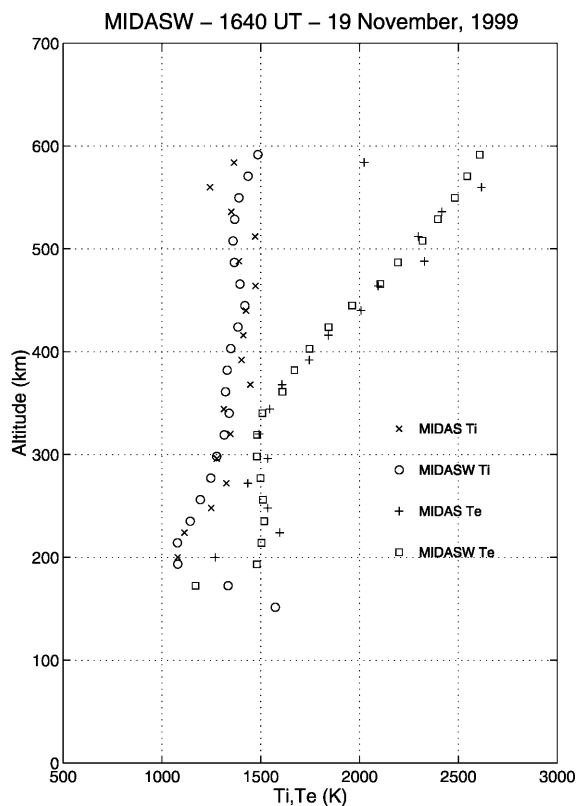


Fig. 9. Comparison of MIDAS-1 and MIDAS-W temperature profiles (see text)

expected form. The zero crossings move to lower lags as altitude increases, corresponding to increasing ion temperature. The lowest altitudes are contaminated by ground clutter, and the data become more ragged at higher altitudes where the signal to noise ratio is small.

Figure 8 compares MIDAS-1 and MIDAS-W power profiles. The MIDAS-W profile is actually the zero lag of the lag-profile matrix, that is the normalization factor for Fig. 8. The MIDAS-1 profile is after summation with a trapezoidal summation rule. The two profiles are in as good agreement as can be expected given the differences in processing. We plan to carry out a comprehensive cross-validation of the two systems after further development of MIDAS-W. Figure 9 compares MIDAS-1 and MIDAS-W temperature profiles. The MIDAS-1 temperatures were derived using the standard INSCAL incoherent scatter analysis program used for routine data processing at Millstone Hill. The MIDAS-W temperatures were derived using a simplified program which neglects some second order corrections, such as the effect of the receiver impulse response on the measured spectrum. Also, the MIDAS-W data were windowed rather than summed with a trapezoidal summation rule, which results in larger error bars. Nevertheless, the agreement is very good. When MIDAS-W has been fully integrated into the MIDAS software system, it will provide integrated lag-profiles in the same CMST format as MIDAS-1, which will permit analysis by INSCAL. This will presumably yield even better agreement.

B Leonid meteor shower observations

The arrival of the yearly Leonid meteor shower during November 16–17, 1999, provided an opportunity to use MIDAS-W in a more computationally intensive experiment. UHF radars such as Millstone Hill can study meteor influx into the upper atmosphere through recording meteor head echoes (McIntosh, 1962). These result from scattering of the incident radio wave from the plasma cloud surrounding an inbound meteor as it ablates in the 80–120 km altitude range. Meteor head echo experiments place a large amount of strain on data acquisition systems for two reasons. First, due to the narrow 1.2° field of view of the Millstone Hill antennas and the high speed of the meteor target (ranging from 10–15 km/s for sporadic meteors to 71 km/s for Leonid shower meteors), a typical meteor generates head echoes only for 30–50 ms, mandating the use of a recording mode capable of saving power profiles from individual radar pulses. Additionally, the large line-of-sight target velocity causes a frequency shift of up to 200 kHz at 440 MHz, necessitating the use of a wideband mode with a correspondingly large signal processing burden.

For the 1999 Leonid period, we used the 68 m Millstone zenith antenna with an experimental mode first employed at the EISCAT facility (Wannberg *et al.*, 1996). By transmitting a 13 baud Barker code and examining each decoded power profile for Doppler decorrelation effects, individual meteor echoes yield information not only on amplitude and range but on

line-of-sight velocity. We used a 4 μ s baud on transmission and 2 μ s sampling on reception to allow a maximum velocity in excess of 50 km/s while still retaining all essential frequency information for Barker decoding. Factoring in the changing position of the meteor radiant (the place where Leonid shower meteors appear to originate) relative to the zenith antenna on the nights in question, this allows us to see Leonid meteor head echoes for the entire period of time when the radiant is above the horizon (with some attenuation and decoding errors caused by the analog anti-aliasing filter for those echoes with the fastest velocity).

We note that this experiment would not have been possible using the MIDAS-1 system. The MIDAS-1 front end contains a lowpass filter before the A/D converter with a maximum bandwidth of 192 kHz, which would limit the final sampling rate and thus cut out significant energy for highly Doppler shifted echoes such as those from Leonid meteors. Even if the required 2 μ s (or 300 m range) sampling rate were possible with MIDAS-1, the desired observation altitude range of 70 to 160 km requires a total decoding, processing, and storage speed much too large for the slow Sun-3E VME bus system and Mercury array processor. In fact, this rate is also too large for any single Ultra 10 workstation to process. However, MIDAS-W's flexible signal processing system allowed us to simply split the altitude raster in half, using one workstation for 70–115 km altitude computations while the other processed 115–160 km.

We collected 74 GB of power profile data over 44 h for the 1999 Leonid meteor period, from 19 UT on

November 17 to 15 UT on November 19. Figure 10 presents two typical meteor echo sequences as seen by MIDAS-W at approximately 1057 UT on November 18, 1999, during the dawn period when meteor counts peak (McKinley, 1961; Zhou *et al.*, 1995). Plotted in each column of Fig. 10 are a number of Barker-decoded power profiles with 300 m range resolution as a function of altitude, arranged in order of increasing time from top to bottom. (Only the bottom half of the power profiles, processed by the first of two MIDAS-W “correlator” programs, are presented.) Each power profile is self-normalized, with the bar graph at the right hand edge of the column indicating relative peak power. Due to network bandwidth issues which will be resolved in future versions, MIDAS-W was capable of recording only 80% of the incoming pulses, resulting in some one-pulse gaps in time as can be seen by examining the left hand axis.

The left hand column shows a relatively low speed sporadic meteor, traveling mostly perpendicular to the radar beam and lasting for approximately 42 ms. The Doppler decorrelation signature in the power profiles centered at 105 km is consistent with the small amount of range motion seen. In contrast, the right hand column captures a high velocity Leonid meteor with a large component parallel to the beam, remaining in view for over 80 ms. The echo moves approximately 5 km downwards in range over 0.078 s, resulting in a time-of-flight velocity of 65 km/s consistent with the Doppler decorrelation signature seen in the individual echoes. This speed also agrees well with the line-of-sight velocity predicted for a 71 km/s Leonid meteor at the radiant




Fig. 10. Two typical meteor echo sequences as seen by MIDAS-W at about 1057 UT on 18 November, 1999 (see text for details)

elevation of 63° associated with 1057 UT at Millstone Hill's location.

7 Conclusions

MIDAS-W represents a new generation of digital receivers and DSP systems, which will be implemented entirely in software in a general-purpose computer. This approach has several important advantages over both analog receivers and digital systems which rely on special-purpose hardware:

1. Simple: MIDAS-W requires an analog input and a few control signals, and uses exclusively off-the-shelf components. The prototype, capable of handling a limited range of simple long-pulse experiments, was constructed by a single programmer in less than two weeks.
2. Maintainable: it is very simple to modify, inspect, debug and test, since almost everything is implemented in software.
3. Flexible: it is very easy to experiment with alternative signal processing techniques by simply adding a node to the network. This node can process the samples without disturbing the production system. This also accommodates direct comparison between different techniques, since they can operate in parallel.
4. Customizable: if the existing components are not suitable for an experiment, new components can easily be implemented in software.
5. Scalable: additional processing power can be achieved simply by adding more computers to the network.
6. Upgradeable: the performance of the system can be improved simply by upgrading the workstations and high-speed network.

MIDAS-W has been tested both with standard Millstone Hill incoherent scatter operating modes and high-bandwidth modes which were not possible with

MIDAS-I or MIDAS-C. In the near future it will be integrated with the higher-level MIDAS software system, after which it will become the production Millstone Hill data acquisition system. The low-level MIDAS-W hardware and software, shown in Fig. 6, is not dependent on the higher-level MIDAS software and, as illustrated by the examples discussed, can be operated independently. In fact, a second system, using identical hardware and software is being built for use with the EISCAT incoherent scatter radars.

Acknowledgements. Observations and analysis were supported by the National Science Foundation through a Cooperative Agreement with the Massachusetts Institute of Technology.

Topical Editor M. Lester thanks J.D. Sahr and N. Wade for their help in evaluating this paper.

References

- Bose, V., M. Ismert, M. Welborn, and J. Guttag,** *Virtual radios, IEEE JSAC issue on Software Radios*, April, 1999.
- Holt, J. M., D. A. Rhoda, D. Tetenbaum, and A. P. van Eyken,** Optimal analysis of incoherent scatter radar data, *Radio Sci.*, **27**, 435–447, 1992.
- Lehtinen, M. S.,** Statistical theory of incoherent scatter measurements, PhD Thesis, University of Helsinki, Helsinki, Finland, 1986.
- McIntosh, B. A.,** The meteoric head echo, *J. Atmos. Terr. Phys.*, **24**, 311–315, 1962.
- McKinley, D. W. R.,** *Meteor science and engineering*, McGraw-Hill, 1962.
- Proakis, J. G., and D. G. Manolakis,** *Digital signal processing: principles, algorithms and applications*, Prentice Hall, 1995.
- Wannberg, G., A. Pellinen-Wannberg, and A. Westman,** An ambiguity-function-based method for analysis of Doppler decompressed radar signals applied to EISCAT measurements of oblique UHF-VHF meteor echoes, *Radio Sci.*, **31**, 497–518, 1996.
- Zhou, Q., C. A. Tepley, and M. P. Sulzer,** Meteor observations by the Arecibo 430 MHz incoherent scatter radar-I. Results from time-integrated observations, *J. Atmos. Terr. Phys.*, **57**, 421–431, 1995.