

# The WHOI Micromodem-2: A Scalable System for Acoustic Communications and Networking

Eric Gallimore, Jim Partan, Ian Vaughn, Sandipa Singh, Jon Shusta, Lee Freitag  
Woods Hole Oceanographic Institution  
Woods Hole, MA USA  
{egallimore, jpartan, jvaughn, ssingh, jshusta, lfreitag}@whoi.edu

**Abstract**—A successor to the WHOI Micromodem-1 underwater acoustic modem has recently been developed. The Micromodem-2 has the same compact form-factor as the Micromodem-1 and will support all of the existing applications for the Micromodem-1, as well as interoperate with the Micromodem-1. Existing acoustic communications protocols using phase-shift keying (PSK) as well as frequency-hopping frequency-shift keying (FH-FSK) are supported, as are navigation features including narrow-band and broadband long-baseline (LBL) navigation.

The Micromodem-2 is significantly more capable than the Micromodem-1 in computational ability and memory, bandwidth, non-volatile data storage, user expansion interfaces, and real-time clock precision. The expanded capabilities will allow new communications algorithms, modulation, error-correction methods, navigation features, and networking capabilities to be implemented. The improvements in processing capability and acoustic interfaces on the Micromodem-2 allow it to operate at acoustic frequencies from approximately 1kHz to 100kHz. The significant increases in available non-volatile storage enable the Micromodem-2 to capture data in-situ for diagnostic and research purposes.

The Micromodem-2's firmware architecture is similar to the Micromodem-1's firmware architecture, using a real-time operating system based on modular signal processing blocks. It has been improved to increase modularity and facilitate future portability, and it offers significant improvements in timing for use with navigation and networking applications.

## I. INTRODUCTION

The WHOI Micromodem [1] was originally designed as a compact, low-power modem with modest capabilities, intended for use on small underwater vehicles and to be integrated into oceanographic sensors for point to point or network applications [2], [3]. As it was developed, many capabilities were added, including navigation and high-rate phase-coherent transmission and reception.

Navigation features included in the original development phase included compatibility with REMUS broadband transponders to allow a vehicle to operate in a REMUS navigation network, and narrow-band transponder compatibility that supports many commercially-available deep and shallow-water transponders [4]. These features made it easier to add both acoustic communications and navigation to small vehicles by integrating these functions into one subsystem and transducer.



Figure 1. The 4.5" x 1.75" Micromodem-2 preserves the physical form factor of the Micromodem-1 while providing improved expansion connectors and I/O capabilities.

In the early 2000s, the communications capability of the modem evolved from support of the frequency-hopping frequency-shift keying (FH-FSK) open standard developed at WHOI for inter-operability among multiple modems to high-rate phase-coherent signaling. The FH-FSK standard is slow, only 80 bps, but reliable and simple. Originally the modem could only transmit PSK signals, requiring a more capable system, the WHOI Utility Acoustic Modem, for reception. However, in collaboration with NUWC-Newport, a very small floating-point co-processor was designed and software was ported to perform the adaptive equalization required for use of these signals in shallow water. The resulting system provides data rates from 80 to 5000 bps (burst rate), and supports not only shallow-water communications, but deep-water, direct-path links as well [5]. An additional improvement was addition of a multi-channel (4 or 8) analog input card that improves performance in multi-path environments. The adaptive decision feedback equalizer can optimize over multiple channels, greatly improving data rate in many situations.

Development of new applications on the Micromodem-1 has been constrained by several factors, in particular its limited memory, fixed analog front end, and simple expansion buses. With the Micromodem-2 design (Fig. 1), these constraints have been addressed, and we are now positioned to develop new capabilities for underwater networking, add expanded communication rates, and support acoustic communications research.

Table I: Specifications and Performance of the Utility Acoustic Modem, Micromodem-1, and Micromodem-2

	Utility Acoustic Modem	Micromodem-1	Micromodem-2
<i>Power</i>			
Hibernate	N/A	220 $\mu$ W @ 5V, 640 $\mu$ W @ 12V	165 $\mu$ W @ 5V, 455 $\mu$ W @ 12V
Low-Power Detect	N/A	<100mW ( [2], deprecated)	75-80 mW (preliminary)
Active	3W	158mW	300mW (using current software)
Receive PSK	3W	158mW + coproc	300mW + coproc (using current software)
Transmit <sup>1</sup> (185 dB re:1 $\mu$ Pa@1m)	30W nominal	8W-48W	8W-48W
Nominal Max. DSP Performance	30 MFLOPS (TI ‘C44)	160 MMACS (TI ‘C5416)	1066 MMACS (ADI BF548)
FPGA	N/A	N/A	13k LUTs, flash-based
<i>Volatile Memory</i>			
Data RAM	2 MB	128 kB	8MB (total, shared with Program RAM)
Program RAM	1 MB	2 MB	8MB (total, shared with Data RAM)
<i>Non-Volatile Memory</i>			
Flash	1 MB NOR	16 MB NAND	64 MB NOR
Configuration	N/A	96B, backed by RTC battery	8kB FRAM
Removable	N/A	N/A	32 GB microSD
Dimensions	3.5” x 8”	1.75” x 4.5”	1.75” x 4.5”
Input Voltage <sup>2</sup>		5V–16V	5V-24V ( $\pm$ 5%)
Boot Time		>1100 ms	<100ms
First designed	1996	1999	2009

The design goals, described in Section III, are derived from the extensive deployment experiences with the Micromodem-1 [1], while also taking advantage of technology advances. The Micromodem-2 preserves backwards compatibility with the Micromodem-1, in terms of its acoustic communication packet types, NMEA command set, and physical form factor, while providing a scalable path for future applications.

## II. PREVIOUS WORK

The first fully-embedded underwater acoustic modem from WHOI was the Utility Acoustic Modem (UAM) [6], designed in 1996. While quite capable, the UAM had a constant 3W power consumption due to its floating-point DSP, and it had a relatively large form factor.

The Micromodem-1 adopted a fixed-point DSP, reducing the power consumption by about 95%, and shrinking the form factor significantly, as detailed in Table I. The original Micromodem-1 design is over 11 years old, and, in that time, only minor hardware updates have been incorporated into the system. However, the Micromodem software has continuously evolved, and for several years, some development has been impeded by hardware limitations.

Based on this experience with the Micromodem-1, the Micromodem-2 has been designed with the expectation that it will need to be actively supported and used as a development platform for at least ten years. Table 1 compares the

specifications of the three systems.

With the Micromodem-1, insufficient processing power and RAM required the use of an add-on coprocessor card to receive PSK-modulated signals [1]. The coprocessor is based on a TI TMS320C6713B floating-point DSP, and consumes roughly 2W when equalizing and decoding PSK signals. It is turned off when not in use.

## III. MICROMODEM-2 DESIGN GOALS

The Micromodem-2 was designed to support networking and communications research by addressing the constraints of the Micromodem-1, while maintaining backward compatibility. An important requirement is scalability to allow the same system to be used in both power-constrained situations (e.g. multi-year ocean-bottom or glider deployments) and computation-constrained situations (e.g. short deployments for acoustic communications research or deployments on AUVs). A further need identified early in the re-design process was higher sampling and baseband processing rates, to support newer wideband transducers and higher frequencies.

Central to the plan to use the Micromodem-2 for the next decade is design flexibility and modularity. Advancements in components over the next several years will provide opportunities to increase capabilities as well as reduce power consumption, and the hardware and software have been selected and constructed such that these incremental hardware upgrades will be possible. Additionally, the hardware currently has significant headroom in terms of processing power, RAM, and storage, which will allow for extensive

<sup>1</sup> Transmit power depends strongly upon the transducer used, which in turn depends upon the frequency band, platform, and channel in question.

<sup>2</sup> Modems are commonly powered by power amplifiers which accept input voltages in the range of 5V to 48V.

future development.

A common difficulty encountered when testing Micromodem-1 systems was a lack of data that could be used forensically to determine the causes of good or poor communication. Statistics had to be recorded by the host system controlling the Micromodem, and therefore were often incomplete. To remedy this, the Micromodem-2 includes significant non-volatile storage that can be used to record channel statistics or make passband recordings for later analysis. These data can then be used to analyze the performance of the modem and characterize the propagation channel.

The following subsections provide background on design decisions including processor choice, power supplies, memory and storage, acoustic input and output stages, and user interfaces.

### A. Computational Capability

While the Micromodem-1 used a Texas Instruments TMS320VC5416 with a nominal maximum of 160 MMACS (million multiply-accumulates per second), the Micromodem-2 is equipped with an Analog Devices Blackfin ADSP-BF548 that provides over 1000 MMACS. Both processors are 16-bit fixed-point digital signal processors (DSPs), although the Blackfin incorporates more general-purpose microprocessor functionality and a larger peripheral set than the 'c5416. On both Micromodems, the DSP's clock frequency is dynamically scaled based on the processing load.

When selecting the processor for the Micromodem-2, only low-power (and therefore fixed-point) devices that were readily available at design time were considered. Our view is that, over the past decade, the market for fixed-point digital signal processors (DSPs) has shifted away from traditional low-power DSPs with modest memory spaces toward higher-power DSPs with features closer to general-purpose processors. This apparent shift can be seen by observing which DSP product lines have seen continued development, and which have stalled. We suspect that the shift has probably been driven by smart phones, which require complex operating systems yet can be recharged nightly. The TI 'c5x line of traditional DSPs (of which the Micromodem-1's 'c5416 is a member) appears to have had its development stalled for quite a few years (although TI's dual-core OMAP processors combining a 'c5x DSP and an ARM microprocessor are still being actively developed for the smartphone market). In contrast to the TI 'c5x product line, the Blackfin DSP product line appears to have a solid roadmap for the future. Thus, although processors were available that may have provided an easier software transition from the Micromodem-1, we chose to switch processor families and use a Blackfin due to its capabilities and a future upgrade path.

An important use of the increased computational capability will be the implementation of the adaptive decision feedback equalizer and error-correction software on the Micromodem-2, removing the need for the current floating-point coprocessor card. This will save a significant amount of power and reduce the cost of a complete system. However, in the interim, the

Micromodem-2 is compatible with the coprocessor card.

In the event that an application needs to scale to higher processing power, the Micromodem-2 has a high-speed expansion connector which will allow multiple Micromodem-2's (or other processors, such as a WHOI Optical Modem [7]) to be connected and used concurrently.

### B. Power Consumption

A trade-off associated with the improved computational ability of the Micromodem-2 DSP is an increase in the typical power consumption during packet decode relative to the Micromodem-1 (300mW vs 158mW, plus in both cases 2W for the coprocessor when equalizing and decoding PSK packets). To offset this increase in receive power, we improved the low-power states of the Micromodem-2. The Micromodem-2 operates in three modes, designated "active", "low-power detect", and "hibernate".

The *active mode* is used in standard applications where energy use is not a concern, and it runs multiple acoustic detectors and provides a full-featured user interface. In this mode, the Micromodem-2 draws about 300mW (measured at 25°C and 5V input voltage; this increases slightly with input voltage and temperature). In this mode as well as all others, the modem disables power to subsystems that are not in use, such as the microSD card and the analog output circuitry as illustrated in Figure 2.




Figure 2. Micromodem-2 Power Control System

*Hibernate mode* is primarily useful on long deployments or other applications where energy is very limited and the modem can operate with a low duty cycle. In this mode, the primary power supply on the modem is turned off, along with all the downstream components. Thus, while in hibernate, the modem cannot detect incoming packets or respond to user input. However, the modem can transition from hibernate to active mode based on a number of events such as a real-time clock alarm, a user signal on specific I/O pins, or a signal from another board using the modem's expansion bus. With the Micromodem-2, we decreased the hibernate power to 165 $\mu$ W versus 220 $\mu$ W for the Micromodem-1, with a 5V battery voltage. The hibernate state uses a linear voltage regulator, so power consumption will scale roughly linearly with battery voltage. For example, with a 12V battery voltage, the hibernate power consumption is 455 $\mu$ W and 640 $\mu$ W, for the Micromodem-2 and -1, respectively.

The *low-power detect mode* is a new feature of the Micromodem-2 that will run a matched-filter detector using a flash-based FPGA, while leaving the DSP powered off. The importance of a low-power detector has been emphasized in many papers [8], [9]. We are currently developing and testing several different detection algorithms based on matched filter FM sweep detection, as well as BPSK sequence detection [10]. Based on initial experimentation using the Micromodem-2, we estimate that a low-power detector state running a matched filter for an FM sweep will use 75mW-80mW. When a packet is detected, or user serial input is received, the Micromodem-2's improved boot time (0.1s vs. 1.1s) allows it to wake rapidly from the low-power detect state and respond. An alternative approach is analog tone detection, which uses significantly less power, though it is less robust than FM sweep detection. In addition, an FM sweep provides high-resolution timing synchronization.

Future revisions will incorporate hardware and software improvements to reduce power in the low-power detect and active modes.

### C. Memory

Easy implementation of new features is a key aspect of research and development systems such as the Micromodem, and restrictions on data memory made this difficult or in extreme cases (such as a PSK equalizer) impossible.

The wider memory address bus of the Blackfin DSP supports more data RAM, 8MB as opposed to 128 KB. Data RAM may be increased to 128MB in the future, when higher-density RAM becomes available. To maintain low-power capability yet provide large amounts of RAM, micropower asynchronous static RAM is used. This memory has a standby power of a few microWatts, at a cost of slower access speeds relative to high-power synchronous dynamic RAM.

The Micromodem-2 also adds significant non-volatile memory storage. The boot device is a 64MB parallel NOR flash, rather than a serial EEPROM, as was used on the Micromodem-1. This reduces the boot time and allows additional non-volatile user storage. The configuration parameters (such as node ID and communication parameters) are now stored in 8kB of FRAM rather than RAM backed by the real-time clock (RTC) battery. This means that the configuration parameters cannot be lost in the event of RTC battery failure, and no power is required to maintain the configuration parameters.

Perhaps most significantly, a slot for removable microSD cards with up to tens of GB of storage was added. This allows temporary storage of network packets, such as in a "data mule" application [11], logging communication performance such as bit error rates, transmitting research waveforms, and recording raw signals for purposes such as channel characterization or communications research.

### D. Acoustic Input and Output

There are a wide range of modem acoustic frequencies, depending upon the application. Frequencies down to 900Hz [12] and up to 100kHz or more are of interest. To support this,

the Micromodem-2 has a highly configurable fully-differential analog input, allowing frequency-agile operation by selecting among four bands, which are defined by high-pass noise rejection and low-pass anti-aliasing filters. It also has programmable gain and a 16-bit analog-to-digital converter (ADC) that can sample at up to 1MHz. The Micromodem-1 had fixed high- and low-pass filters, and a low-power 12-bit ADC which was not well suited to sampling above 80kHz. With our improvements in the analog front end, the Micromodem-2 is better suited to supporting multi-band network deployments [13], [3], [14].

To provide flexibility and capability beyond the single analog input on the modem board, both the Micromodem-1 and the Micromodem-2 can interface to existing multi-channel analog input boards. Future multi-channel analog input boards with improved capabilities will take advantage of the high-speed SPI or SpaceWire interfaces provided by the Micromodem-2.

Transmit power on the standard Micromodem-1 power amplifier is fixed, using a class D output stage for efficiency. For short-range networking experiments, a low-power linear power amplifier (at 25kHz, about 150dB re:1 $\mu$ Pa@1m, about 1W electrical power) was developed for experiments in transmit power control and spatial reuse in a network [15]. New interfaces available on the Micromodem-2 will allow use of a digital pulse-width-modulated (PWM) power amplifier driven directly by the modem's DSP.

### E. Various Improvements: Input Voltage Range, UARTs, Timing, and Temperature

Based on direct experience and user feedback a number of new features not present on the Micromodem-1 have been added. See Table II for a summary of some of these features.

We have added a wider input power range to support power buses up to 24V, with reverse-voltage and over-voltage protection. The primary power input from the cable-terminated edge connector is combined with an alternate power input on a stacking expansion connector using an pseudo-ideal diode OR circuit, which can be used to provide redundant power to the modem.

The Micromodem-2 incorporates improved I/O capabilities and is designed to reduce overall complexity, particularly in systems that required a separate microcomputer to interface simple sensors to the modem. Almost all I/O is routed through the modem's FPGA, which allows software configuration of almost all I/O on the board.

In contrast to the Micromodem-1's two RS232 serial ports, the Micromodem-2 has four asynchronous serial ports, including two standard serial ports with RS232 signaling levels, one serial port using 3.3V logic levels for low-power integration, and one which can be switched by software between RS232 and RS485/RS422 signaling. A logic-level (3.3V) serial port is also available via the high-density stacking interface connector on the modem.

The modem provides a total of 14 reconfigurable digital I/O pins to users, all of which operate at 3.3V and are 5V tolerant. In addition to the aforementioned asynchronous serial ports

(RS232/422/485), the modem offers I2C and SPI synchronous serial buses. For analog signals from user sensors, the modem includes an auxiliary low-speed analog to digital converter that is capable of up to 240 samples per second. The modem can also measure its own input voltage. The status of all of these inputs can be provided to users acoustically or locally via the control interface. Finally, a switchable 3.3V power output from the modem can provide up to 100mA to external circuitry.

Table II: Features of Micromodem-1 and Micromodem-2:

	Micromodem-1	Micromodem-2
Asynchronous Serial Ports	2 RS232	4 total: 2 RS232, 1 3.3V logic level, 1 switchable between RS232 and RS485/RS422.
Transmit timing	Transmit on pulse-per-second signal	Transmit on $n^{\text{th}}$ pulse-per-second, hardware line toggle, precision timer
System interfaces	1 general purpose input, 1 general purpose output	14 reconfigurable I/O pins, I2C, SPI, SpaceWire, auxiliary analog input, battery voltage monitor
High-speed expansion interface	8-bit parallel data bus at up to 8MHz, half-duplex; SPI	SpaceWire-based high speed bus capable of 100Mbps, SPI
Real-time Clock	Timekeeping accurate to about 20ppm, battery life <10 years at 25°C	Timekeeping accurate to 2ppm, battery life >15 years at 25°C
Operating temperature	0°C to +70°C	-40°C to +70°C or -40°C to +85°C (build options)

The Micromodem-2 includes a real-time clock with a precision of 2ppm (versus 20ppm for typical crystals) to provide better support for networks with synchronized sleeping and longer sleep times [16], [2], as well as improved navigation support [17], [18], [13], and tighter bounds on TDMA network timeslots [19], [20].

In some long-term deployments, the modem is configured to spend most time in the hibernate state and wake periodically to detect scheduled transmissions of acoustic data. The improved real-time clock accuracy results in further power savings in these applications, as shorter “active” periods are necessary to account for worst-case clock drift between the transmitter and receiver.

As with the Micromodem-1, the modem also supports an external pulse-per-second (PPS) input, though now it can also provide a PPS output if desired. The Micromodem-2 now has a dedicated external transmit trigger input as well as handshake signals for sharing a transducer with another acoustic system.

All Micromodem-2 components are specified from -40°C to

+85°C to support arctic deployments, though, to reduce cost, some versions include the commercial temperature range FPGA. Although the commercial-grade FPGA is only guaranteed from 0°C to +70°C, the manufacturer states that it should work properly to -40°C. In any case, we individually test and qualify units from -40°C to +70°C if required for specific deployments.

#### IV. EXPANSION INTERFACES

The original Micromodem-1 uses an expansion header with an 8-bit wide parallel interface to the floating-point coprocessor board. In addition, the receive hydrophone array is sampled by a multichannel analog input board that is connected to the main Micromodem-1 board with another expansion header, using a four-wire SPI synchronous serial interface. Additional connections are used to provide control signals to power amplifiers. These legacy connectors are included on the Micromodem-2 to provide backward-compatibility (Figure 3).

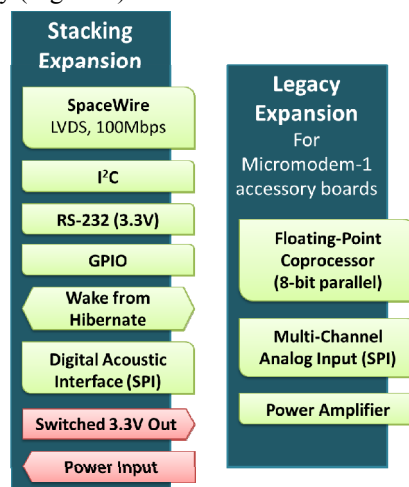


Figure 3. Micromodem-2 Expansion Interfaces

The Micromodem-2 provides a pair of high-density, high-speed stacking connectors to interface with other devices (Figure 3). While intended for future development, the connectors may be employed by users to eliminate wiring harnesses when connecting the modem to other systems. These connectors can provide power to the modem, regulated 3.3V power from the Micromodem-2 to other devices, asynchronous logic-level serial communication, wake-from-hibernate functionality, SPI and I2C synchronous serial buses, digital I/O, and a high-speed SpaceWire interface.

##### A. SpaceWire

The Micromodem-2 incorporates an interface based on SpaceWire (ECSS-E-ST-50-12C) for high-speed off-board communication. SpaceWire is a flexible yet simple-to-implement protocol originally developed by the European Space Agency for use in connecting subsystems aboard spacecraft [21]. It uses low-voltage differential signaling (LVDS) to connect devices with bi-directional, full-duplex, point-to-point links, and a simple protocol allows data packets

to be routed among devices without large processing overhead. Open-source implementations of this protocol are available, for example on [opencores.org](http://opencores.org). The Micromodem-2 complies with the electrical specifications of the standard, but it uses the high-speed stacking connector to pass the signals rather than the prescribed DB-9 connector. As implemented on the modem, it is capable of full-duplex wire speeds of 100Mbps. Intended uses include communication with coprocessors (including other stacked Micromodem-2 boards used as coprocessors), multi-channel analog input and output devices, vehicle computers, Ethernet or CAN bridges, sensors, and Gumstix-based single-board Linux control computers. SpaceWire (and the associated LVDS signaling) is robust enough to operate over distances of several meters, so it may also be used to transfer high-speed audio data to or from remote hydrophones or amplifiers.

## V. SOFTWARE ARCHITECTURE

The software architecture for the Micromodem-2 builds upon the real-time operating system (RTOS) developed for real-time signal processing on the Micromodem-1, previously described in [1]. The RTOS is a small and portable executive which enforces real-time read-locks and write-locks on streaming data buffers passing through processing modules, connected together as drawn in a typical signal-processing block diagram to build the modem and network applications.

At the application level, backward compatibility with the Micromodem-1 is maintained. However, new features are being added, including transmit queues, a streamlined user interface and built-in support for some wireless network functions.

### A. Transmit and Data Queues

Queues for transmit signals and payload data records are now possible due to the additional memory that is available. These will allow users (or the network stack) to enqueue signals, such as communication packets or navigation signals, with four different transmit timing methods. The queue for data records allows sensor samples to be collected before transmission and can be used as a temporary storage location for network data packets, for example in a “data-muling” application.

The transmit queue with the highest priority is a timer-driven queue used for enforcing turn-around times on transmissions such as “ping” packets and long-baseline (LBL) navigation signals, where a two-way travel time needs to be measured precisely using a specified turn-around time. An additional use of the timer-driven transmit queue could be for transmitting packets within appropriate network timeslots [19], [20]. The second-highest priority transmit queue is for signals which are to be transmitted on an external trigger provided by the user, often used for synchronous one-way navigation [17], [18], [13]. Third is a queue for transmitting signals at every  $n^{\text{th}}$  pulse-per-second (PPS), from either an internally-generated or externally-generated PPS, as required by, for example, periodic navigation signals. Finally, the default transmit signal queue is for signals with no detailed

timing requirements, such as standard communication packets. Each of the transmit queues allows enqueueing signals at the tail, at the head (next to be transmitted), or by sorted priority.

### B. User Interface

The standard user interface to the Micromodem-2 is an NMEA-style asynchronous serial interface (RS232, RS485, or logic-level UART), as is typical of many oceanographic instruments. Complete backwards compatibility for the Micromodem-1 NMEA commands is preserved. NMEA-style ASCII interfaces can be somewhat awkward to interact with, however, and they are inefficient at transferring binary data. While the modem is flexible enough to act as an instrument's primary controller, we are investigating various designs for a binary protocol intended to interact with typical host processors. Although more difficult for a human to read during debugging, a binary mode would simplify interfacing the modem to the current class of very low-power microcontrollers (e.g., TI's MSP-430s) that are becoming increasingly popular for low-powered instrument control. In the process of changing the interface, it will be generalized to all of the modem's interfaces, including the acoustic link itself.

### C. Applications and Integrated Networking Support

For user payload data, many Micromodem-1 users have been using the Compact Control Language [22]. A recent extensible encoding language payload data has also been developed, the Dynamic Compact Control Language [23]. The Micromodem-2 continues to support these methods of transferring user data.

Networking support [24] will be integrated directly onto the Micromodem-2. The goal is to develop and support a networking toolbox, which will allow users to build simple underwater networks quickly for specific applications.

## VI. CONCLUSIONS

The Micromodem-2 provides significant improvements in processing capability, storage, and bandwidth. These improvements overcome previous limitations present in the Micromodem-1, and will allow development of new capabilities with support for expanded communication rates, underwater network deployments, and research in acoustic communications and underwater acoustic networks.

## ACKNOWLEDGMENT

We would like to thank several individuals for their assistance with the design of the Micromodem-2. Keenan Ball and Peter Koski provided extremely helpful suggestions and design feedback, including feedback from experience with dozens of Micromodem-1 users, which helped greatly in prioritizing design decisions. Neil Forrester contributed to the FPGA design and prototype hardware. Our collaborators at Ultra Electronics-USSI, particularly Joshua Jarrard, provided feedback and assisted with testing of the initial version of the Micromodem-2.

## REFERENCES

- [1] L. Freitag, M. Grund, S. Singh, J. Partan, P. Koski, and K. Ball, "The WHOI Micro-Modem: An Acoustic Communications and Navigation System for Multiple Platforms," in *Proc. MTS/IEEE OCEANS 2005*, vol. 2, Washington, DC, USA, 17-23 Sept 17-23 Sept. 2005, pp. 1086-1092.
- [2] D. Frye et al., "An Acoustically-Linked Deep-Ocean Observatory," in *Proc. Oceans 2005 - Europe*, vol. 2, 20-23 June 20-23 June 2005, pp. 969-974.
- [3] M. Grund, L. Freitag, J. Preisig, and K. Ball, "The PLUSNet Underwater Communications System: Acoustic Telemetry for Undersea Surveillance," in *Proc. IEEE Oceans 2006*, Boston, MA, USA, Sept. 18-21 Sept. 2006.
- [4] L. Freitag, M. Grund, C. von Alt, R. Stokey, and T. Austin, "A Shallow Water Acoustic Network for Mine Countermeasures with Autonomous Underwater Vehicles," in *Underwater Defence Technology Conference*, Amsterdam, NL, June 2005.
- [5] S. Singh, S.E. Webster, L. Freitag, L.L. Whitcomb, K. Ball, and J. Bailey, "Acoustic communication performance of the WHOI Micro-Modem in sea trials of the Nereus vehicle to 11,000 m depth," in *OCEANS 2009, MTS/IEEE Biloxi*, Biloxi, MS, USA, 26-29 Oct. 2009.
- [6] M. Johnson, J. Preisig, L. Freitag, and M. Stojanovic, "FSK and PSK Performance of the Utility Acoustic Modem," in *Proc. MTS/IEEE Oceans 1999*, vol. 3, Seattle, WA, USA, 13-16 Sept. 1999, p. 1512.
- [7] N. Farr, A. Bowen, J. Ware, C. Pontbriand, and M. Tivey, "An Integrated Underwater Optical/Acoustic Communications System," in *MTS/IEEE Oceans Conf*, Sydney, Australia, 2010.
- [8] A. F. Harris III, M. Stojanovic, and M. Zorzi, "Idle-Time Energy Savings through Wake-up Modes in Underwater Acoustic Networks," *Ad Hoc Networks*, vol. 7, no. 4, pp. 770-777, June 2009.
- [9] J. Wills, W. Ye, and J. Heidemann, "Low-Power Acoustic Modem for Dense Underwater Sensor Networks," in *Proceedings of the First ACM International Workshop on UnderWater Networks (WUWNet)*, Los Angeles, California, USA, September, 2006, pp. 79-85.
- [10] J.P. Fumo and M.W. Ornée, "Implementation of an Advanced Ocean Transponder and Deckset Utilizing Complex Waveforms," in *OCEANS 2001 MTS/IEEE Conference and Exhibition*, vol. 3, 2001, pp. 1788-1793, In particular, Section V.A and Figure 4.
- [11] M. Dunbabin, P. Corke, I. Vasilescu, and D. Rus, "Data Muling over Underwater Wireless Sensor Networks using an Autonomous Underwater Vehicle," in *Proc. 2006 IEEE Intl. Conf. on Robotics and Automation (ICRA)*, May, 2006.
- [12] A. K. Morozov and D. C. Webb, "A Sound Projector for Acoustic Tomography and Global Ocean Monitoring," *IEEE J. Oceanic Engineering*, vol. 28, no. 2, pp. 174-185, 2003.
- [13] L. Freitag, M. Grund, J. Partan, K. Ball, S. Singh, and P. Koski, "Multi-Band Acoustic Modem for the Communications and Navigation Aid AUV," *Proc. MTS/IEEE OCEANS 2005*, vol. 2, pp. 1080-1085, 17-23 September 2005.
- [14] S. Shahabudeen, M. Chitre, and M. Motani, "A Multi-Channel Mac Protocol For Auv Networks," in *Proceedings of IEEE Oceans'07 Europe*, Aberdeen, 18-22 June 2007.
- [15] J. Partan, J. Kurose, B. N. Levine, and J. Preisig, "Spatial Reuse in Underwater Acoustic Networks using RTS/CTS MAC Protocols," University of Massachusetts Dept of Computer Science, UM-CS-2010-045, 2010.
- [16] A. Syed and J. Heidemann, "Time Synchronization for High Latency Acoustic Networks," in *Proc. IEEE Infocom*, Barcelona, Spain, April, 2006.
- [17] R. Eustice, H. Singh, and L. Whitcomb, "Synchronous-Clock One-Way-Travel-Time Acoustic Navigation for Underwater Vehicles," *Journal of Field Robotics*, in press, 2010.
- [18] S. Singh, M. Grund, B. Bingham, R. Eustice, H. Singh, and L. Freitag, "Underwater Acoustic Navigation with the WHOI Micro-Modem," in *Proc Oceans 2006*, Boston, MA, 18-21 Sept. 2006.
- [19] M. Molins and M. Stojanovic, "Slotted FAMA: A MAC Protocol for Underwater Acoustic Networks," in *Proc. IEEE Oceans'06 Conf - Asia Pacific*, Singapore, 16-19 May 2006.
- [20] X. Guo, M. R. Frater, and M. J. Ryan, "Design of a Propagation-delay-tolerant MAC Protocol for Underwater Acoustic Sensor Networks," *IEEE Journal of Oceanic Engineering*, vol. 34, no. 2, pp. 170-180, April 2009.
- [21] S. Parkes and J. Rosello, "SpaceWire ECSS-E50-12A," in *International SpaceWire Seminar*, Noordwijk, The Netherlands, 2003.
- [22] R. P. Stokey, L. E. Freitag, and M. D. Grund, "A Compact Control Language for AUV Acoustic Communication," in *Proc. IEEE Oceans 2005 - Europe*, vol. 2, 20-23 June 2005, pp. 1133-1137.
- [23] T. Schneider and H. Schmidt, "The Dynamic Compact Control Language: A Compact Marshalling Scheme for Acoustic Communications," in *Proc IEEE Oceans 2010*, Sydney, Australia, 2010.
- [24] J. Shusta, L. Freitag, and J. Partan, "A Modular Data Link Layer for Underwater Networks," in *Proc. IEEE Oceans 2008*, Quebec City, QC, 15-18 Sept. 2008.