Overview of TAPS-8

The TAPS-8 is a lineal descendent of a line of multi-frequency echosounders developed to measure acoustic backscattering from zooplankton. The essential characteristics of this device include:

- Measures absolute volume scattering strengths
- Computes echo statistics to qualify data sets
- Measurement volume is adapted for zooplankton
- Measures at a suitable set of frequencies
- Stores data internally and outputs in real-time

All of these factors should be addressed in the design of similar acoustic sensors. The purpose of this section is, first, to describe the characteristics of TAPS-8 and, second, to provide a framework for design of similar systems in the future.

Volume scattering strength (S_v) is defined as the backscattering from a volume of water containing a 'large number' of 'randomly distributed' scatterers, adjusted to unit volume. It is the ratio of the backscattered intensity from a unit volume of scatterers to the incident intensity, where the backscattered intensity is referred to unit distance from the volume. [The units of volume scattering work out to be m^{-1}, which entices some to explain this as backscattering cross-section (an area) per unit volume. This is specious reasoning and should be avoided.]

This definition provides an absolute value for measurements of volume scattering. Assuming, of course, that one can accurately specify the incident intensity and measure the backscattered intensity. This is possible only from a system that has been calibrated, including specification of the output intensity (Source Level, or SL), the receiving sensitivity (RS), and the properties of the system that influence the effective ensonified volume (beam pattern, range, and pulse length), as well as the transmission losses due to geometrical spreading and absorption.

Volume scattering is a random phenomena. That is, the measured value on each ping is a random variable and is a poor estimate of the 'true' value of volume scattering. A better estimate requires several measurements -- either of adjacent volumes or on subsequent ping cycles.

Note also that volume scattering is only defined when certain characteristics apply to the scattering volume. There must be a 'large number' of scatterers that are 'randomly distributed'. Rather than re-hash the ins and outs of these asssumptions, the reader is directed to the discussion of volume scattering (***Basic Acoustics***) that is provided with these notes. Suffice it to say that volume scattering is an important property because it is the only situation under which we can associate the backscattered intensity with the number and types of the scatterers.

One occurrence that can cause the data to depart from volume scattering is the presence of a single strong scatterer in some of the data or a strong scatterer that has directional properties and varies it's aspect angle from ping to ping. In such a case, even when there is a background of volume scattering, the mean value of the backscattered intensity will not usually correspond to any rational sum of the individual scatters. This is because the statistics of the scattering from a single scatterer -- especially one that is not present on every ping -- differs radically from that of volume scattering. The sum of random variates from dissimilar distributions is, well, messy.

A typical situation where one might run into this problem would be the downward-pointing echosounder pinging at a layer of zooplankton. The most numerous constituents would probably be copepods and similar sized and shaped zooplankters. Less common might be mysids or pteropods or larval fish. These latter, however, are strong scatterers compared to the zooplankton and would easily bias the backscattering intensity. If the actual ensonified volumes in this layer are such that the volume contains, on average, only 1 or 2 or so of the stronger scatterers, the echo statistics will be mixed and volume scattering cannot be presumed. In this case, the use of echo intensities to estimate biomass is not robust.

The TAPS sensors address this problem by constraining the sample volume to about 2 liters. With a sample volume of this size, one would expect to see the larger scatterers only rarely, if at all. The more numerous scatterers will normally provide adequate numbers to obtain volume scattering.

TAPS-8 has a mode where the pulse length is lengthened and data are sampled out to 16m range. In this case, we are trying to obtain scattering from the less numerous zooplankters. Since the sample volumes are still not all that large, however, we need some way to decide if the measurements can be treated as volume scattering or not. We do this by computing a ratio of two of the moments of the (measured) echo intensity distribution (see the volume scattering notes for details). By inspecting this estimate, we can select only those data that appear to be volume scattering for further processing.

We process the measurements of volume scattering to provide estimates of the size-abundance for the scatterers. A model or models is required to sort the data against. We use a version of least-squares fitting to obtain a functionally optimal estimate of the size-abundance that best fits the measured data given the model(s) supplied. This process has been described elsewhere. Of interest in these notes, however, is the choice of the frequencies and the number of frequencies to be employed in these estimates.

The figure below shows the mean frequency (of a pair of channels) that would best estimate copepod-like scatterers of a given size. For example, to optimally estimate the abundance (and reliably estimate the ESR) of 0.5 mm ESR copepods (roughly 2.5mm long), one should have a set of frequencies that are centered on about 350 kHz.

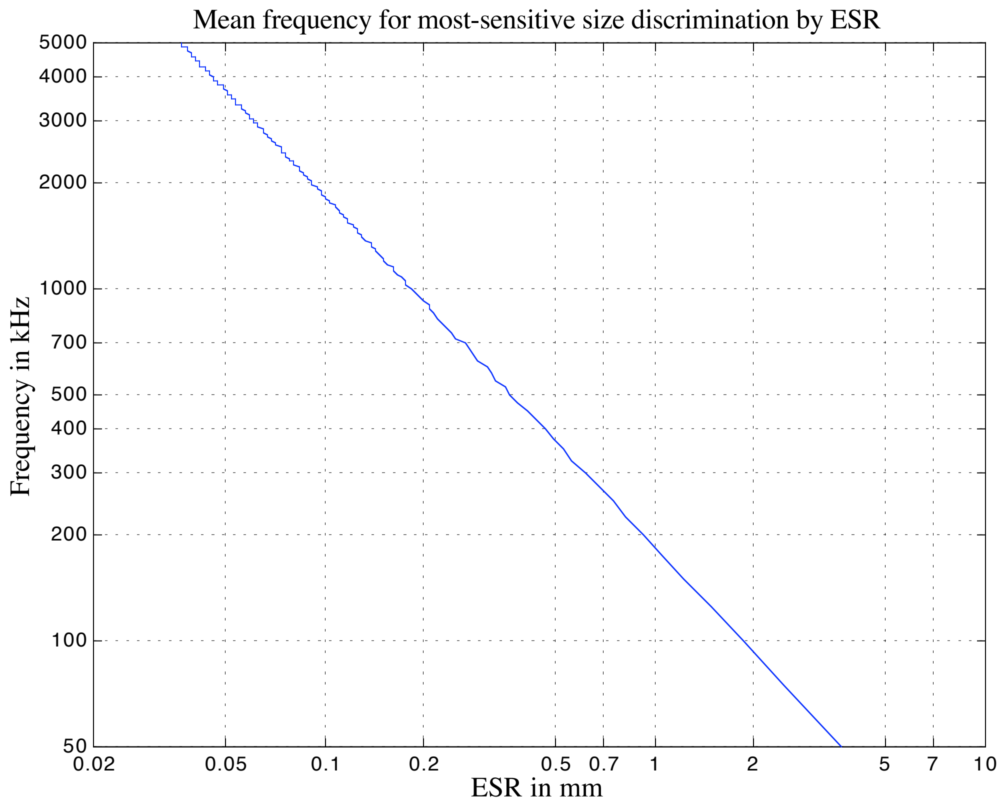

The number of frequencies is usually determined by hardware and power considerations. From an estimation point of view, one should keep in mind a couple of factors. Statistically, one would expect to be able to distinguish no more than N-1 size classes (plus the total biomass) given data at N frequencies. Experience shows that inversion using many (100 or more) size classes improves the fit of the estimates to the data, with many contiguous size bins containing abundance estimates. Invariably, however, these contiguous bins are really estimates of a single size (or, occasionally, two sizes) -- the spread around the peak is actually an estimate of the precision of the size estimate, caused perhaps by a range of sizes in the sampled organisms. But, withal, estimation of more sizes requires more frequencies.

However, the ability to distinguish separate sizes of scatterer is not wholly a function of the number of frequencies. It also depends on the shape of the target strength function with frequency. If this is a gentle, slowly-varying function (as the truncated fluid sphere model is), then adding frequencies between existing frequencies may not add more effective size bins to your estimates. See Greenlaw and Johnson 1983 for more discussion on this subject.

The sample volume is a parameter that must be chosen as well. If the scatterers of interest occur at abundances of, say, 1000/m^3 or higher, then to obtain a mean number of scatterers in the sample volume of, say, 10, the sample volume must be at least 10 liters.

Given an estimate of the desired sample volume, one can do an engineering assessment of the free variables to obtain this value in practice. The sample volume consists of a section of a spherical shell of thickness $c\tau/2$, where c is the sound speed and τ is the pulse length, at range R. The incident intensity is not uniform over this section but varies according to the directivity pattern of the transducer (this is explained in excruciating detail in ***Basic Acoustics***). For a circular disc transducer such as those used in TAPS, the effective sample volume can be computed from

$$V_e = R^2 \frac{c\tau}{2} \frac{4.853}{kD}$$

where D is the diameter of the disc and k is the wavenumber ($2\pi/\lambda$). Thus, one can choose any of these 'free' parameters to obtain the desired sample volume. For TAPS, we have selected transducers with approximately 8-10° beamwidths (DI ≈ 30 dB), the pulse length is 336 μSec, and range is 1.25 m. These choices provide an effective sample volume of about 1-2 liters.

TRANSDUCERS

Transducers are electro-mechanical devices that convert electrical oscillations into mechanical oscillations (sound) and vice versa. The usual material for transducers used in water is one of the piezo-electric ceramics; in the case of TAPS-8, a lead-zirconate-titanate mixture known as Navy Type I. Each channel of TAPS-8 uses a separate transducer for both sound generation and sound reception.

TAPS transducers are disc-shaped, with electrodes on the planar faces of the discs. Applying a sinusoidal electrical drive to these electrodes causes the disc thickness to vary in synchrony with the drive, causing sound energy to be emitted from each face. If one face (the back) is terminated into a low mechanical impedance (such as air), then the sound energy that would be radiated from that face is essentially reflected (and inverted in phase). If the drive frequency is such that the reflected sound arrives at the front face in phase with the sound that is being generated at that face, then the signals add together to approximately double the sound amplitude. This occurs when the thickness of the disc is approximately one-half wavelength. This frequency is known as the resonant frequency for the transducer.

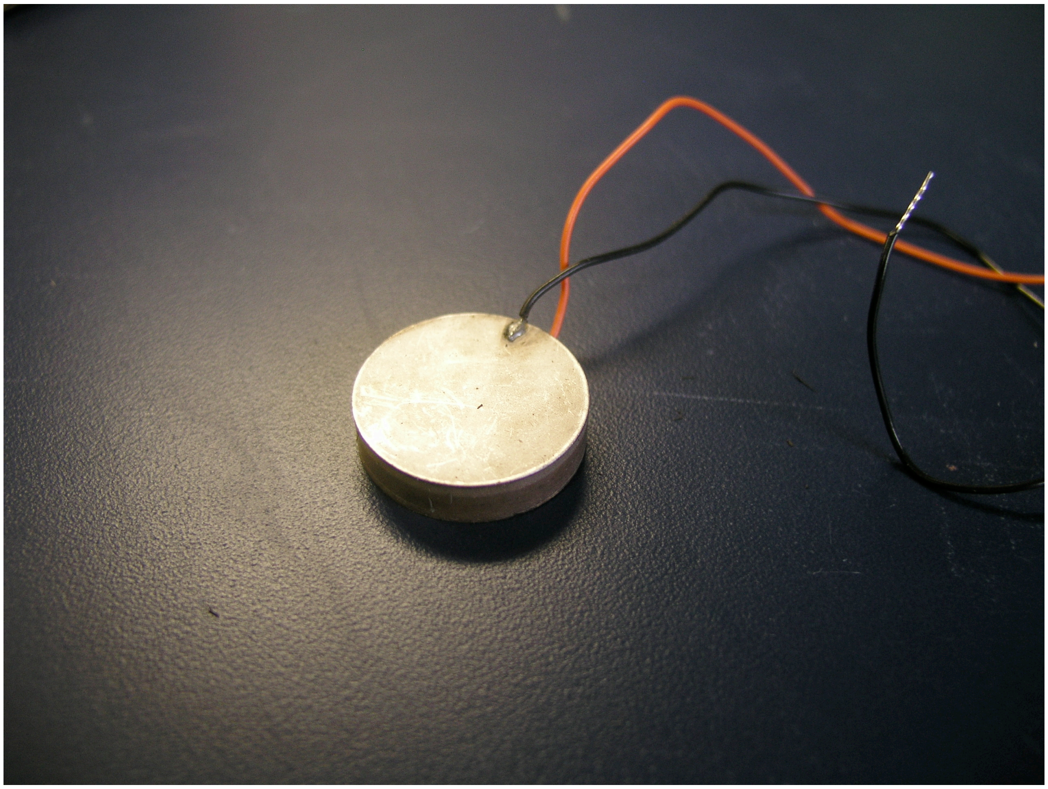

Piezo-electric ceramic transducers appear as complex electrical loads to a power amplifier. At low frequencies, the transducer looks rather like a resistor in series with a capacitor. In the vicinity of the resonant frequency, the value of the resistor decreases substantially and a minimum value of the magnitude of the impedance, |Z|, is found. As the frequency increases somewhat, the impedance magnitude increases to a peak value at a frequency called the anti-resonance. In the region between resonance and anti-resonance, the impedance can change from capacitive to inductive (the phase shifts from - to +). At frequencies above the anti-resonance peak, the transducer again approaches the impedance of a resistor in series with a capacitor.

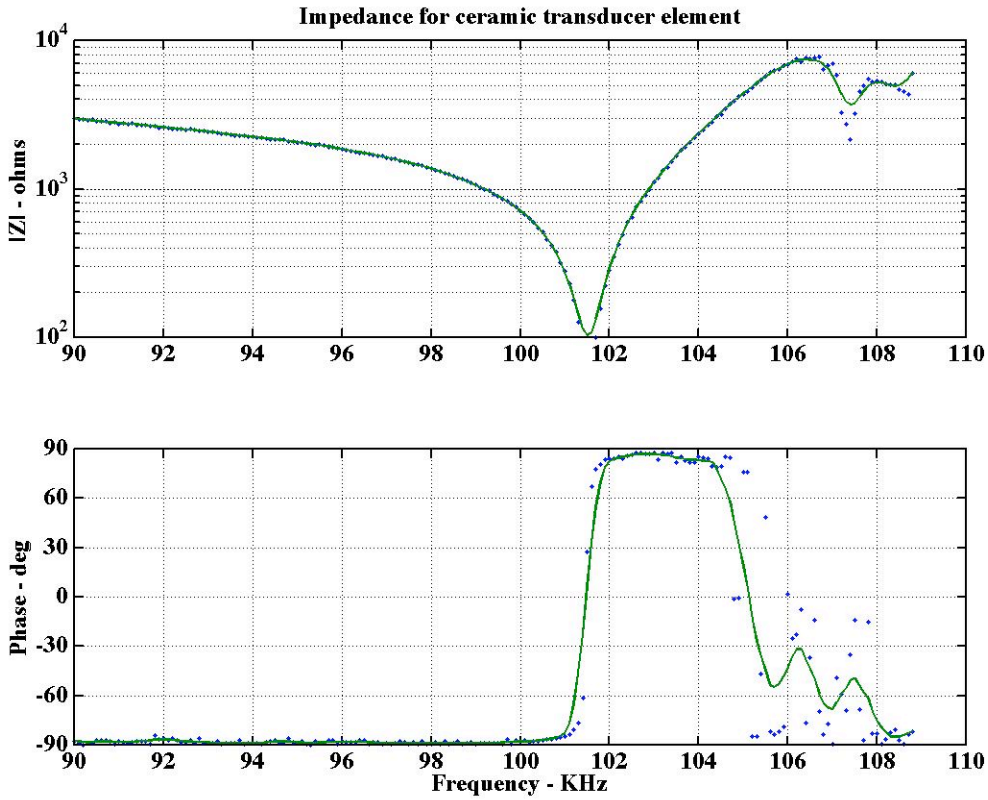

The receiving efficiency of the transducer -- the voltage developed between the electroded surfaces in response to an impinging sound wave -- tends to peak at the anti-resonance frequency. Below this point, the efficiency is essentially constant. Above this point, the efficiency is inversely proportional to frequency. The transmitting efficiency, in terms of unit current drive, increases with frequency up to the resonance point and decreases as the frequency-squared below this point. See the figure on the next page.

If a transducer is intended for operation at a single frequency, as are the TAPS transducers, optimum response is usually obtained by selecting that frequency to lie between the resonance frequency (where output power is greatest) and the anti-resonance frequency (where receiving response is the greatest).

Laboratory transducers and transducers intended for shallow use can obtain a close approximation to the air-backed condition by glueing a cork-rubber material (Corprene) to the back and sides. This material is not rigid enough to stand much ambient pressure, however. The transducers in TAPS-8 are of this construction and, thus, cannot be used at depths much over 20-30 m.

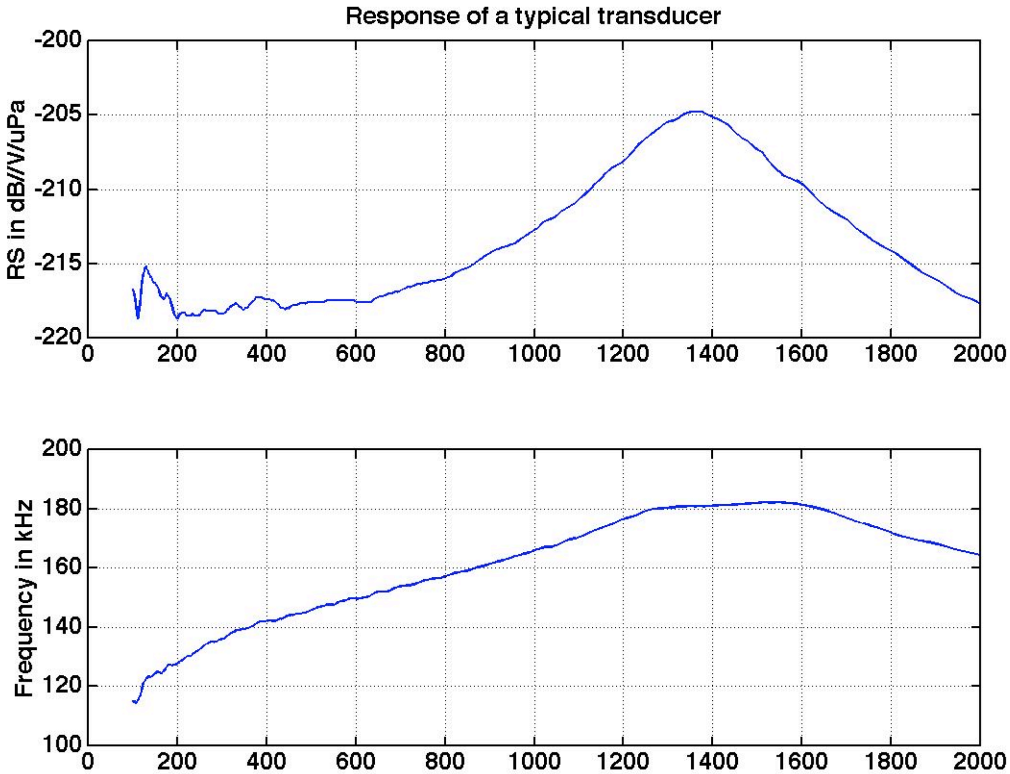

Deep-submergence transducers can be 'fooled' into seeing a low mechanical impedance at the back face by employing a mechanical transformer on the back side. We typically use a compressed paper-mylar material called Copaco as a matching layer between the ceramic and an aluminum housing. The thickness of the Copaco is selected to be approximately one-quarter of a wavelength ($\lambda = c / f$ where λ is the wavelength, c is the sound speed -- 929 m/sec for Copaco -- and f is the frequency in Hz). The combination of the Copaco layer and a large mass of aluminum creates an effective low impedance at the back face of the transducer.

Sound waves are reflected whenever they encounter a discontinuity in acoustic impedance (defined as the product of density and sound speed - ρc). The sound radiated from the front face of the transducer is partially reflected (back into the ceramic) if the impedance seen there is significantly different from the impedance of the ceramic, typically around 34 MRayls (10^6 kg m^{-2} sec^{-1}). The impedance of water is about 1.5 MRayls. A better match between these impedances can be had by interposing a quarter-wavelength layer of material with an intermediate impedance. This layer, the acoustic window, must also provide water-proofing for the transducer electrodes, which will corrode rapidly in seawater.

Optimum results will obtain when the intermediate layer impedance is the geometric mean of the ceramic and water or about 7 MRayls. We use an two-part epoxy material (Circalok 6007 Black encapsulant with RT-7S epoxy hardener from Lord Corporation) for our transducer windows. This material has a cured impedance of about 4.5 MRayls, close to the desired impedance. Equally important, this material has the same coefficient of thermal expansion as 6061-T6 aluminum, reducing the chance of separation of the epoxy from the aluminum housing with resultant leaks of seawater that will destroy the transducer.

We have found that elastomers such as urethanes make satisfactory acoustic windows (although the impedance matching properties are slight) at low frequencies but the loss through these materials is excessive at high (MHz) frequencies. Thus we use hard epoxy encapsulants almost exclusively.

MATCHING THE TRANSDUCER TO THE POWER AMPLIFIER

All power amplifiers work best into a specific load impedance, often 50 Ω. The complex nature of the transducer impedance means that it is very unlikely that the transducer will appear as a matched load to the power amplifier. It is straightforward to construct a network of electrical elements that will transform the complex transducer impedance into a purely-resistive 50 Ω, at least at one frequency. The procedure we use is based upon the narrow-band reactive coupling network derived from the series/parallel impedance equivalence. Our power amplifiers are designed for operation into 50 Ω, so that is the impedance to which we will match the transducer.

Suppose we measure the impedance of the transducer at the operating frequency as $R_s + j\,X_s$. The series resistance is R_s and the series reactance is X_s. We can transform this to a pure resistance by inserting a series element with reactance $X_1 = -X_s$. Generally, the transducer impedance is capacitive (negative) so the first matching element will be an inductor. At this point, we now have a resistive load with $R = R_s$. Note that this resistance is both the series and parallel resistance; let us refer to it as R_t, the tuned resistance.

Now suppose we insert a second reactive element (X_2), this time in parallel with the tuned transducer. By the transformative laws of impedance, we can calculate that the series resistance of this parallel combination will be

$$R_s^{'} = \frac{R_t X_2^2}{R_t^2 + X_2^2}$$

If we desire R' to be 50 Ω, we can solve this equation for the unknown X_2 to obtain

$$X_2 = \sqrt{\frac{50 R_t^2}{R_t - 50}}$$

where it is assumed that $R_t > 50\ \Omega$. This reactive element can be either capacitive or inductive. We generally select a capacitor to provide a low-pass characteristic to this matching network.

Finally, the impedance of the two-element network and transducer are measured again. If we have selected the parallel tuning element correctly, the impedance will be $Z = 50 + j X$, where X is some (capacitive) reactance. We complete the network by adding X_3, a series inductor, to tune out the measured reactance, resulting in a transducer that is tuned to 50 Ω at the operating frequency.

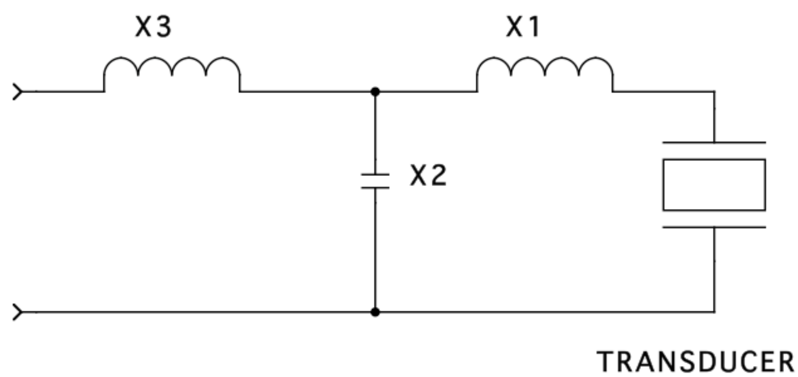

TRANSMITTER

Schematics for this card are contained in the folder *TAPS-8 Schematics file*. See *tbs*.

The transmitter used in TAPS-8 is a class-D switching power amplifier. Logic signals -- a transmit gate pulse and a frequency input -- are combined to produce two out-of-phase drive signals for power MOSFET transistors. These transistors drive the push-pull input of a matching transformer. The output of this transformer will drive a 50 Ω resistive load at approximately 150 watts. An inhibit logic input is provided to prevent transmissions when desired (i.e., to obtain noise files).

TUNING CARD

The transmitter drives a network of tuned elements that, in turn, drive the transducers. Each transducer has an attached network which provides both the matching function and a band-pass characteristic so that each transducer can operate only on it's assigned frequency.

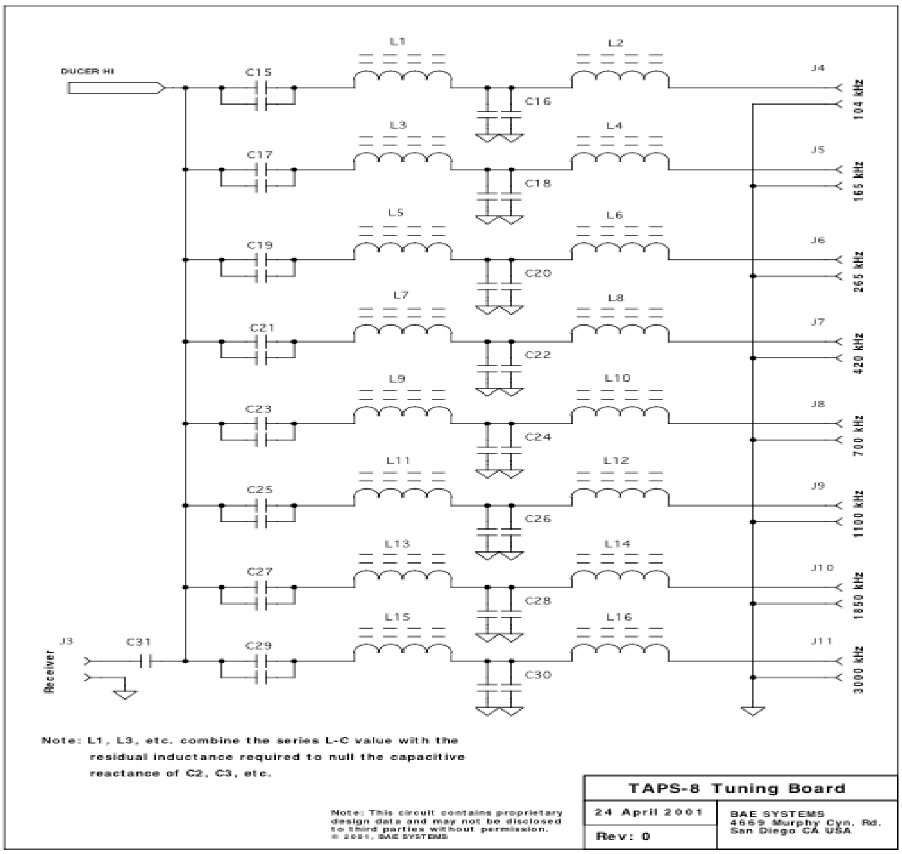

The transducer matching is done with a 3-element network as described above. A series inductor tunes the (capacitive) reactance of the transducer. A parallel capacitor creates a series resistance of 50 Ω. And a series inductor tunes out the residual reactance of the first two elements.

Tuning consists of a series L-C circuit designed for a suitable bandwidth driving 50 Ω. The center frequency is, of course, the center frequency of the channel. Note that the inductance of this bandpass network is incorporated into the last tuning inductor.

When complete, the network should show relatively high impedance except in the vicinity of the channel frequencies*, where the impedance should approximate 50 Ω. and the phase should be zero. Of course, a real network with complex loads will not behave precisely like this but there should be regions for each channel where the load impedance is approximately correct, such as the (actual) impedance sweep shown below.

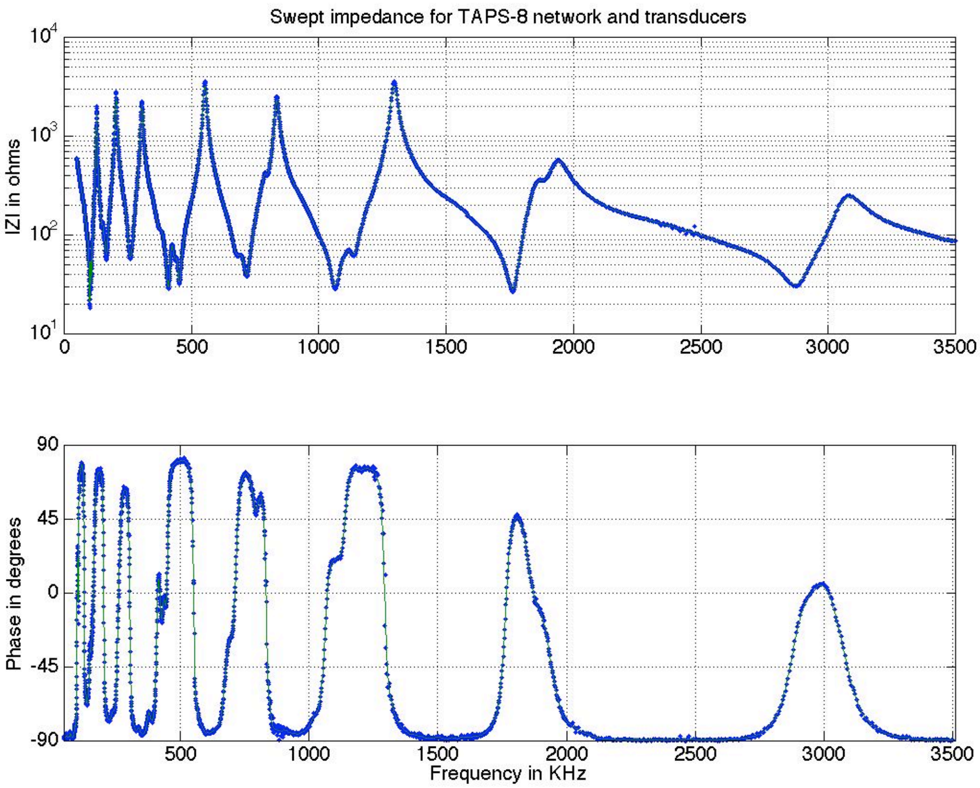

Note that, generally, the point of zero phase corresponds to an impedance magnitude of about 50 Ω. The impedance at 1850 and 3000 kHz was chosen to be 20-30 Ω to reduce the power applied to these transducers.

* Channel frequencies for TAPS-8 are nominally: 104, 165, 265, 420, 700, 1100, 1850, and 3000 kHz. Some units may have a slightly different highest frequency.

RECEIVER

Schematics for this card are contained in the folder *TAPS-8 Schematics file*. See *TAPS8 Rcvr page 1*, and *TAPS8 Rcvr page2*.

A single receiver is used in TAPS-8 to process echoes from all channels. This receiver is a single-conversion superheterodyne design with time-varied-gain and a linear detector output.

Echo signals from the transducer tuning network are input to a Transmit/Receive (T/R) switch composed of diodes D1-2, Q01, and associated components. The two diodes are connected back-to-back (cathode to cathode). During transmit, the high voltages applied to the transducers is strongly attenuated (by 40-60 dB), thus blocking them from the receiver input.

The transmit gate is applied to U09, a dual one-shot multivibrator. The first stage is triggered on the falling edge of the transmit gate and produces a short (ca. 100 μSec) pulse. The trailing edge of this pulse is used to trigger a longer gate pulse, ca. 45 mSec, which is applied to the gate of Q01. When Q01 conducts, current flows through the diodes D1 and D2. With current flowing in the diodes, low-level signals -- such as echoes -- at the input of the T/R switch are capacitively coupled to the output, thus 'opening' the receiver for signals.

The output of the T/R switch is loaded by a 100 kHz high-pass filter, designed for a characteristic impedance of 500 Ω. We discovered many years ago that a typical medium-frequency disc transducer tends to produce a long tail of oscillations after the transmit drive signal ends. This tail could be several milliseconds long, making it impossible to detect echoes at short range. The frequency content of this tail revealed that it was composed of damped oscillations produced by circumferential and face-shear modes of vibration in the disc. Since the resonant frequencies of these modes is much lower than that of the primary thickness mode, preceding the receiver with a high-pass filter effectively removed this short-range interference, allowing detection of echoes almost immediately after the end of the transmit pulse.

The first stage of amplification consists of a X2 transistor amplifier stage (Q02-3) which largely serves as an impedance transformer between the 500 Ω high-pass filter and the 100 Ω input impedance of U01, a broadband variable-gain amplifier. U01 comprises a fixed-gain stage and a voltage-variable attenuator which is linear in dB. That is, a given voltage step produces a fixed gain change in dB.

This stage is used to provide the spreading loss gain correction that is correct for volume scattering (also called time-varied-gain or TVG). The voltage ramp (which is non-linear since the amplifier is linear in dB) is developed on the Controller board and applied to a buffer amplifier (U02) and thence to the TVG amplifier, U01. The AD603 amplifier requires a specific voltage difference at the gain pins (1 & 2): -0.5V to +0.5V. Stage gain at -0.5V is -10dB while gain at +0.5V is 30dB. A 0-1V TVG gain voltage signal is provided by the controller card. A 0.5V reference voltage,produced by the precision voltage reference, U03, and the resistive divider R14-5, is added to this signal to obtain the ±0.5V range to U01.

The output of U01 is terminated in 200 Ω in RF transformer, T1. The secondary of T1 drives the (50 Ω) RF input of M1, a passive doubly-balanced mixer. The LO (local oscillator) input of M1 is driven by a sine wave developed on the Controller board and buffered by U04 and T2. This input to the mixer is 50 Ω, also.

A balanced-mixer produces an output that consists of a signal at the sum and at the difference of the two input frequencies. Thus the signal at TP4 would be a composite replica of the input signal at |Fo + Flo| and |Fo - Flo|. We used a LO frequency that is 35 kHz higher than the transmit frequency, so one of the signals is centered at 35 kHz and the other is at Fo + 35 kHz. This composite signal is amplified in broadband amplifier U06 and then applied to a passive bandpass filter, FL2. This filter passes only the difference-frequency signals to the following stages.

A second fixed-gain amplifier, U06, amplifies this filtered signal. The output is applied to U07, an instrument amplifier with digitally-controlled gains of 1, 2, 4, 8, and 16 X. Control signals are provided by the controller card to set the gain of this stage according to the channel sensitivities measured during calibrations.

The final receiver stage is detection of the AC signals in U08, a linear detector. The output of U08 is a slowly-varying DC signal that represents the RMS envelope of the input signal. This output is provided to the Controller card for digitization and further processing.

Power for the Receiver card is obtained from a power switch on the Controller card. 24V DC is applied to a 15V positive regulator. This voltage is further regulated to 12 and 5 V for use in the receiver. A DC-DC converter, U13, changes the +15V to -15V, which is also regulated to -12 and -5V. Test points for the various voltages are provided on the PCB. The power supply circuits are contained in a shielded enclosure to reduce interference from the TC962 switching inverter.

CONTROLLER

Schematics for this card are contained in the folder *TAPS-8 Schematics file*. See *TAPS8 CPU*, *TAPS8 memory*, and *TAPS8 power*.

The controller card is responsible for connecting power to the acoustic system, tracking time and date, generating transmit gates, setting transmit and LO frequencies (on the DDS card), selecting receiver gains, generating the TVG voltage, and digitizing the echo envelopes. The CPU chosen for this card is the Motorola 68HC11, an 8/16-bit cpu. This particular version of the 68HC11 comes with Forth language embedded in ROM (68HC11V35 from New Micros, Inc.). Thus, the operating program is largely programmed in Forth.

The CPU, U01, together with glue logic U02-5 and memory IC's U06-7, comprise a simple microprocessor system. U06 is used as system RAM (32 Kb) while U07 is a 32Kb EPROM holding the operating program. The system RAM chip is a Dallas Semi DS1644, which has an accurate real-time-clock embedded in the chip along with batteries and power selection circuits to make this RAM non-volatile.

Communications with the cpu are via an RS-232 serial port buffered by U08. This chip senses connection to a conforming RS-232 port on the output and turns on and off accordingly, thus reducing power drain when no serial port is connected (i.e., during normal operations).

The basic memory map of the 68HC11 is 64 Kb in size. The lower 32 Kb contains on-board registers and the RAM in U06. The upper 32 Kb contains the operating program in EPROM (U07) and the Forth code in ROM on the cpu chip. A complete memory map is shown below. All addresses are in hexadecimal.

FROM	TO	CONTAINS
0000	0100	FORTH variables & stack
0101	01FF	Text Input Buffer
0200	103B	misc FORTH variables
103D	7EFF	RAM
7FF8	7FFF	RTC registers (in RAM)
8000	AFFF	PROGRAM (EPROM) space
B000	B3FF	Registers & Ports
B400	B5FF	PROGRAM (EPROM) space
B600	B7FF	EEPROM
B800	BFFF	PROGRAM (EPROM) space
C000	C0FF	External memory/latches
C100	DFFF	PROGRAM (EPROM) space
E000	FFFF	FORTH ROM

In order to store large amounts of data, external data memory is provided. Up to four external NVRAM chips may be installed, each holding 2 MB of data. Access to this memory is via memory-mapped pages, as explained below.

The limited number of I/O lines on this processor requires additional ports be constructed externally to the cpu. This is done by memory-mapping a section (256 bytes) of high memory (at $C0xx) for use external to the cpu. U12 is an 8-bit comparator which compares the upper address lines to a fixed value ($C0). Whenever this address space is enabled (in a read or write operation), the output of this chip goes low. This signal is diode-ORed to provide a MEMDIS signal to the cpu which disables normal memory access. This signal also enables U13, a 3-line-to-8-line decoder. This chip decodes the lower 3 bits of the address lines to select one of 8 possible outputs.

If the address is $C007, then the \EXP output of U13 is enabled. This line is condtioned with R/W from the cpu and used to latch an 8-bit latch, U16. The inputs to U16 are the 8 data lines; latching this chip moves the values on the data lines to the output lines. These outputs are used as a static digital port controlling channel selection, receiver gain, and selecting the ADC or DAC for data transfer.

If the address is $C000-2, one of the \xADR lines is enabled. These lines control 8-bit latches, U17-9, which are used as a 20-bit address latch for the extended data memory composed of NV-RAM chips U21-4. Prior to reading or writing to this extended memory space, the complete 20-bit NVRAM address must be setup in these latches by three writes to $C000 (high), $C001 (middle), and C002 (low).

If the address is $C003-6, one of the \RAMx lines is enabled. These lines select one of the four NVRAM chips, each of which hold 2 Mb of data. Thus a 20-bit address is sufficient to address 2 Mb of data in each chip since separate lines are used to enable the particular chip being addressed.

(One could of course, use a 24 bit address and decode the upper bits to select the NVRAM chips. This would not be significantly simpler than the scheme used, although it might simplify the code slightly.)

Another 8-bit comparator, U25, is used to memory map a second RTC to $C1xx. This RTC, U26, is included to allow the use of an interrupt to pace data collection. When enabled, the \IRQ line can be used to wake up the cpu and cause it to collect a set of data. When data collection is complete, the cpu is put in a sleep mode to conserve power until the next time data is required.

The 68HC11 cpu has a dedicated sychronous serial port which is used to drive two peripheral chips -- an 8-channel, 12-bit ADC, U10, and an 8-bit DAC, U09. Chip select lines for these IC's are provided by the memory-mapped data port from U16. Three lines are required for the serial peripheral interface (SPI) bus: SCLK, MISO, and MOSI. These acronyms stand for system clock (SCLK), Master-in/Slave-out (MISO), and Master-out/Slave-in (MOSI). An SPI peripheral chip uses the SCLK to clock in data on the MOSI line and, simultaneously, clock data out on the MISO line. The chip select lines control which peripheral is active on these lines.

The ADC chip requires a byte of setup data to select the input channel and certain operating parameters. During transmission of this byte, data sampling begins. Two more bytes are sent to the ADC to clock out the digitized channel voltage.

The DAC chip requires two bytes of data to transfer the 8-bits of data and certain setup bits. The chip does not send any data in return. The output of this DAC is converted to a 0-1V voltage signal with a resistive divider.

SPI signals are also sent to a connector for use on the DDS card.

One digital I/O port on the cpu, PORT-A, is used for certain control functions. Bits 3-7 of PORT-A are available as output bits and are inverted and buffered in U28 for use on and off the Controller card. One of these bit lines (PA3) controls power to the peripheral cards through the power switch composed of Q1-2. Q2 is a p-channel power MOSFET transistor. Bringing the gate voltage low (via Q1) turns this transistor on and provides 24VDC power to the external cards.

PA7 is used as the transmit gate signal. PA4-5 are used on the DDS card as a chip select signal (PA4) and to select one of the two available frequencies (PA5). PA6 is used as the INHIBIT signal to the transmitter.

The controller card is provided with 24V power from one of two possible input sources (typically batteries). These sources are diode-ORed so that the higher-voltage source is drained first. This voltage is applied to a linear regulator chip, U27. This chip has a shutdown pin that causes the chip to open the current path whenever this pin is high. Installing a shorting plug on the case of TAPS-8 shorts the bottom of resistor R22 to ground, pulling this line low and turning the regulator chip on. This provides +5V to the cpu and associated circuitry, causing it to run the startup sequence and begin operations.

DDS

Schematics for this card are contained in the folder *TAPS-8 Schematics file*. See *tbs*.

Transmit and local oscillator signals are produced by a Direct-Digital-Synthesizer (DDS) chip on the DDS card. The DDS function is incorporated in an integrated circuit, U01, an AD9835 from Analog Devices. This chip is provided with a clock input (50 MHz from U03) and digital inputs via the SPI bus. Inside the DDS chip are two registers that hold digital values. On each clock cycle, one of these registers' value is added to an accumulator and that value used to lookup a value from a ROM table of sine values. This value is sent to a DAC. The output is a current waveform -- a sinewave at a fixed frequency. Changing the value in the register, or changing between the two registers, changes the output frequency immediately.

The output of the DDS chip is a current that is converted to a voltage in precision resistor, R2. This sinewave signal is buffered in U2 and provided to an SMA connector on the card. This signal is used as the LO frequency input to the receiver.

The sinewave signal is also provided to a wide-band comparator, U4, where it is converted to a TTL square-wave. This square-wave signal is buffered in the HCMOS-TTL chip, U5, and sent to J3 for use as the transmit frequency in the transmitter. Recall that the transmitter requires an input frequency four times the desired output frequency.

Switched 24V is regulated to +12V and then to +5V on the DDS card. The +12V is inverted in U7 to provide -12V for the sinewave buffer amplifier, U2.

The frequency source for the DDS chip is a programmed oscillator chip, U3. These can be obtained in various frequencies from the vendor (DigiKey, in this case). The values sent to the DDS are based upon this clock frequency and must be adjusted if a different clock frequency is provided.

CABLING AND INTERCONNECTS

tbs

Trouble-shooting

tbs

Calibrations

tbs under separate contract