Technologies for Sonar Processing

Hugh M. South, David C. Cronin, Samuel L. Gordon, and Timothy P. Magnani

The Applied Physics Laboratory has been a leader in developing programmable sonar signal processors for 27 years. The Strategic Systems Department is significantly upgrading its sonar processing facilities both to capitalize on advances in processing hardware and to implement improved sonar processing algorithms. This new system also includes enhanced displays and modeling and simulation capabilities. (Keywords: Hardware, Signal processing, Software, Sonar.)

INTRODUCTION

Sonar processing has been a major task of APL's Strategic Systems Department (SSD) for the past 27 years. Most of this work has been performed within the Trident Sonar Evaluation Program (TSEP) sponsored by the Director, Strategic Systems Programs. Interdepartmental cooperative efforts with APL's Submarine Technology Department have led to the development of display systems for the Program Manager for Mobile Surveillance Systems of the Space and Naval Warfare Systems Command.

SSD staff members, along with our subcontractors, have made significant contributions to the development of recording and processing systems for sonar signals. For example, several of our displays are now commonly used in Navy sonars. This article focuses on the evolution of sonar processing technologies in SSD, with emphasis on the hardware and software, displays, sonar simulators, and recording systems.

Sonar is any system that uses acoustic means to detect, localize, track, or classify objects. The preeminent role of sonar is the detection and tracking of submarines and, to a lesser extent, surface ships operating in the world's oceans.

Submarines are highly capable weapons platforms that are difficult to detect when submerged. Because sound propagates relatively well in the ocean, the Navy has relied heavily on the use of acoustic detection systems for finding submarines. *Passive* sonar systems[1,2] detect sound radiated by a target of interest. *Active* systems launch pulses of acoustic energy and detect echoes from targets. By their nature, active systems signal that they are in operation, whereas passive systems can operate covertly. SSD works with both active and passive sonar, but this article discusses only the latter. In addition, a general sonar system consists of both a sensor array and a processor, but only our work on sonar processors will be discussed.

HARDWARE AND SOFTWARE DEVELOPMENT

Fundamentally, a sonar processor accepts acoustic signals (sound waves) detected by a sensor array in the

ocean, extracts the characteristics of those signals, and presents the characteristics on a visual display. Typical, sensor arrays processed in SSD are from 10 to 1000 ft in length and contain 50 to 1000 sensors or hydrophones. The signals received on such arrays must be characterized both by direction of arrival, called spatial processing or beamforming, and by time evolution, called spectral processing. Generally, the spatial and spectral processing can be performed separately with no loss in the detectability of the signals. From this description, one can view a passive sonar processor as a combination of a beamformer, a spectrum analyzer, and a display.

As Fig. 1 indicates, sensors are sometimes recorded on magnetic or other media rather than being processed directly. Most of the processors developed by SSD have been designed as shore-based systems for sensor data recorded on magnetic media during operational submarine missions. The output of the spectrum analyzer is stored on magnetic disk prior to display and analysis. Use of such intermediate storage allows the analysis of the data to proceed on a different schedule from the processing. This approach is characteristic of most of our passive sonar processor designs.

In addition to submarine sensor data, the passive sonar processor can accept input from a device called a front-end stimulator (FES). The FES, which will be described subsequently, can generate controlled, simulated signals to support processor test and calibration. The generic system shown in Fig. 1 could have many different implementations in hardware and software.

A general approach to building a sonar processor is suggested by the nature of the passive sonar problem. Spatial and spectral processing can be separated. In addition, processing for individual sensors and signal arrival directions can be separated. This separability suggests a distributed architecture: individual processors linked by a high-speed bus with some overall synchronization scheme. In such an architecture, the requirement to achieve real-time or faster than real-time processing rates can be met by using multiple
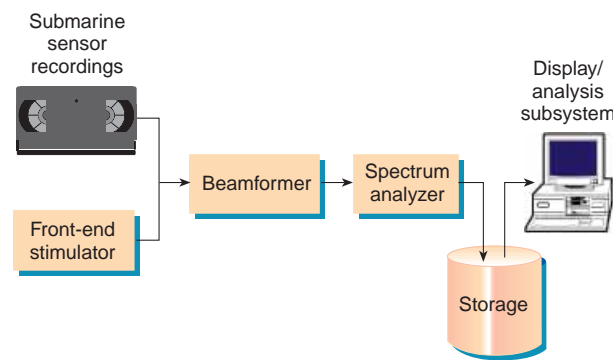
processors working on different data segments or different problems.

Since the various sonar processors may execute at different rates, the individual processors must also have access to memory buffers to smooth the flow of results among units. A distributed architecture naturally provides modularity, which eases system integration and allows new processors to be added as needed. Finally, any architecture must support the programmability of algorithms and algorithm parameters.

The first major SSD sonar processor, the Sonar Evaluation Program Analyzer (SPAN-A), became operational in 1979[1] and was followed in 1983 by a processor called SPAN-I, which employed similar technology.[2] SSD used SPAN-I until 1994.

By current standards, the commercial computers and signal processing hardware available in the early 1980s were slow. The fastest minicomputers then available, e.g., the System 32 by Systems Engineering Laboratories, ran at an internal clock rate of 1.67 MHz, roughly 1/125th of the speed of a personal computer today. Similarly, a typical high-speed floating-point processor of the time, the Floating Point Systems AP-120B, contained 8 circuit cards and could execute 12-million floating-point operations per second (Mflops), whereas a single Intel i860 chip today is nearly 7 times faster.

The lack of truly high-speed commercial signal processing hardware forced the developers of SPAN-A and SPAN-I to build the systems, both processing units and interconnecting bus, largely using custom hardware specifically designed for sonar signal processing. For example, SPAN-I contained 36 different board types, which were either wire-wrapped or double-sided printed circuit cards using standard transistor-transistor logic components. The use of custom fixed-point hardware made SPAN-I economical to build but significantly limited the types of algorithms it could run.

Both SPAN-A and SPAN-I were successful designs, but neither system could process the complete passive sonar suite developed for the new SSBN-726 (*Ohio*) class Trident submarine, which joined the Fleet in the late 1980s. Analysis of Trident submarine passive sonar data required a system that could process sonar arrays containing up to 1000 elements with input data rates to 30 MB/s, which was 6 times the maximum input rate of SPAN-I. The steady advances in microchip technology and the increasing speed of computers made it possible to implement SPAN-I functions in 32-bit floating-point arithmetic and still maintain the required processing throughput. In addition, the new processor, called the Trident Sonar Processor Analyzer (TSPAN),[2] was able to use much more commercial off-the-shelf (COTS) hardware than SPAN-I. TSPAN became operational in 1991 and is still in use today.

From the outset, the Navy had a preplanned product improvement plan for TSPAN to accommodate

Figure 1. Functional components of passive sonars developed in the Strategic Systems Department.

the steady evolution in Navy sonars and to capitalize on enhancements in commercial processing technology. In 1995, a team from SSD was formed to design a new system, provisionally called TSPANU (for upgrade), and development began in 1996. The overall design of TSPANU is given in Fig. 2.

The computing power of TSPANU is contained in two chassis. Each chassis contains 10 commercial array processing boards manufactured by Sky Computers; and each board contains 16 Pentium i860 processors providing 80 Mflops per processor. Therefore, the peak computational rate of TSPANU is 12.8-billion floating-point operations per second, about 3 times the rate of the original TSPAN without using any custom floating-point hardware. The need for custom-designed floating-point hardware limited the flexibility of TSPAN and made software development more difficult. Advances in commercial processing technology have removed those limitations on TSPANU.

TSPANU contains two custom-designed boards instead of the eight found in TSPAN. One of the boards, developed by APL's Technical Services Department, performs a specialized beamforming function called digital multibeam steering,[3] which allows very efficient beamforming of hydrophone data sampled at 1 bit. The other board, developed by SSD, buffers incoming data from the tape player and performs bit-level deformatting. Both boards are located inside processor chassis 1 (Fig. 2).

The principal data bus of TSPANU is a commercial product called Sky Channel, also a product of Sky Computers. Sky Channel was selected because of its high bandwidth (320 MB/s) and its ability to interconnect multiple chassis, providing for future expansion. In this interim TSPANU design, two other commercial buses support lower-speed data transfers: FDDI (fiber digital data interchange) and ATM (asynchronous transfer mode). The FDDI bus in particular allows much of the TSPAN software to be reused temporarily. Table 1 compares several of the SSD-designed sonar processors. By capitalizing on steady advances in processing technology, TSPANU has over 9 times the processing capability of the early

SPAN-I, provides much more flexibility, and costs about the same.

TSPANU will have the capability to perform virtually any current sonar signal processing or display function. Figure 3 gives a typical processing flow planned for the upgrade. Owing to its high throughput, TSPANU will be able to process two sonars simultaneously using the latest adaptive beamforming and display techniques.

For each of the four signal processing systems (SPAN-A, SPAN-I, TSPAN, and TSPANU), two factors affected the scope and complexity of the respective software development efforts. First, inclusion of COTS hardware was a major goal starting with the TSPAN system. We used COTS hardware when practical to reduce initial costs, facilitate integration, and reduce the expense of long-term maintenance and repair of the system. This approach has been carried forward to the TSPANU system.
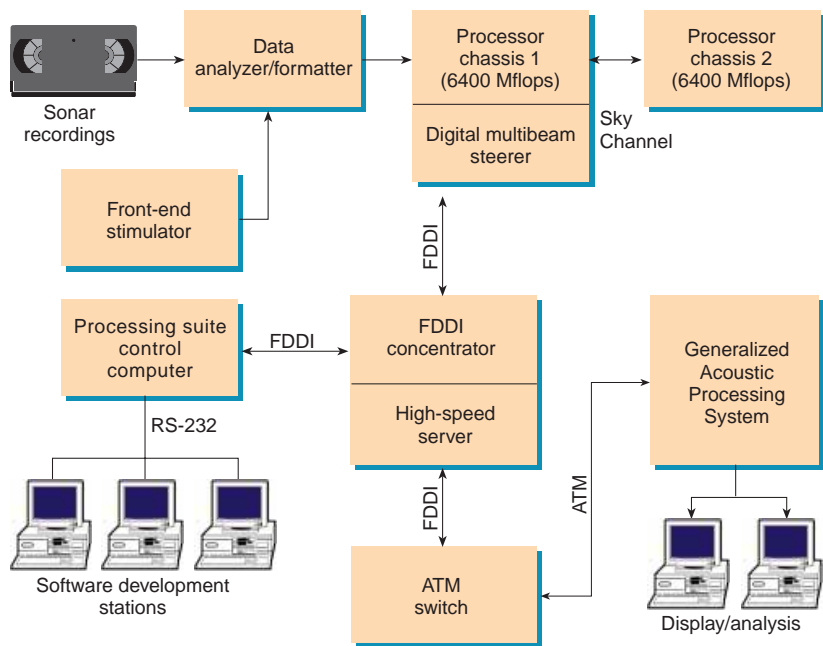
Figure 2. Initial architecture of the TSPAN upgrade signal processor (Mflops = million floating points per second, FDDI = fiber digital data interchange, ATM = asynchronous transfer mode).

Table 1: Evolution of SSD's sonar processor capability.

Feature	SPAN-I	TSPAN	TSPANU
Data bus speed (MB/s)	64	400	320
Computation rate (Mflops)	1,400	4,800	12,800
Custom-designed board types	36	8	2
Cost relative to SPAN-I		2.6	0.9
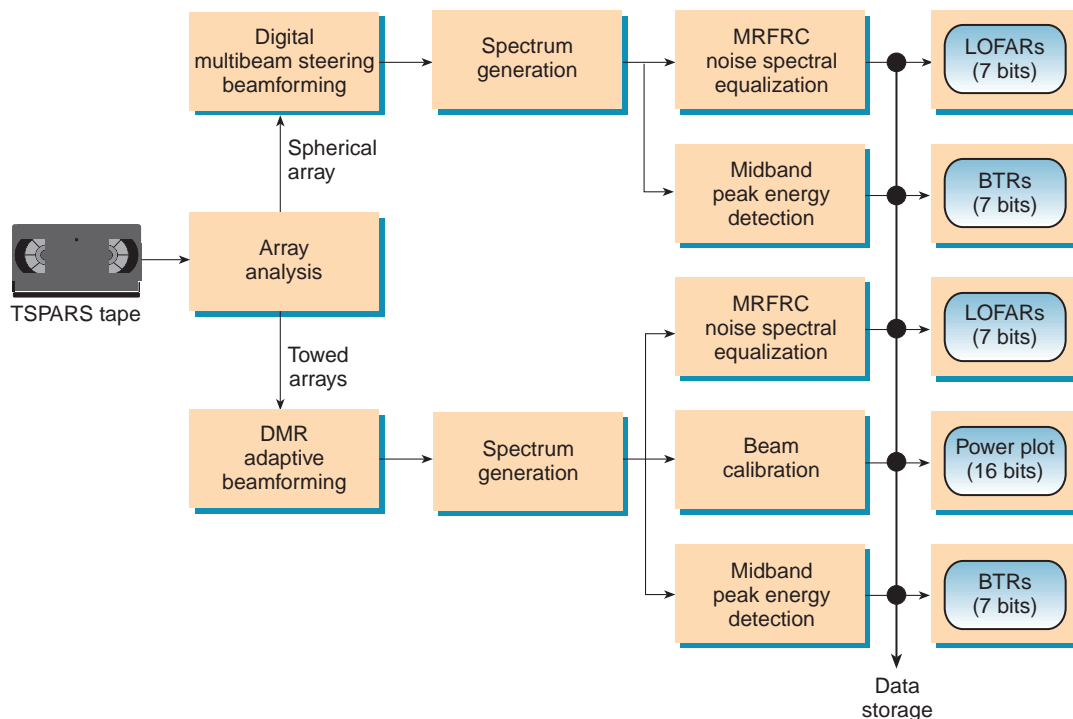
Figure 3. Planned signal processing flow for TSPANU (DMR = dominant mode rejection, MRFRC = mark-off random field rolling cylinder, LOFAR = low-frequency analysis and recording, BTR = bearing time recording).

The second factor resulted from increased expectations and ambitions regarding the ease with which each developed system should be reconfigurable. Even with the earliest system, SPAN-A, SSD staff realized that the development of a rigid processing system—one that could process only the sonar sensing systems deployed during the early 1980s—would be unacceptable. As a result, SPAN-A was built to be "programmable" to allow it to be reconfigured to process future sensors. But although the system was reconfigurable by the standards of the era, it was nevertheless difficult to change as a practical matter. Beginning with the TSPAN system, and advanced further for TSPANU, reconfigurability became a formally specified design requirement.

These two factors did affect the complexity of the software development efforts, but not entirely as one might expect. At first glance, it would seem that incorporation of COTS hardware would make the development task easier, whereas incorporation of increasingly "smart" software would make the development task harder. The latter proved to be true, but the former did not, because inclusion of COTS hardware in the TSPAN system held a surprise.

In the late 1980s, when hardware commitments had to be made for the TSPAN system, COTS choices were considerably fewer than we enjoy today. This shortage applied to raw floating-point processors as well as high-speed data transmission bus choices. (Today we have available standardized hardware such

as ATM and FDDI, and can also choose from software transfer protocols. Also, there are vendor-specific high-speed data transfer solutions, typically employing backplane-mounted crossbar switches.)

Since the TSPAN designers had fewer choices, they selected an input/output transfer computer product built by Aptec Computer Systems. The Aptec product allowed for the use of different hardware interface ports, which were chosen on the basis of the data transfer rate needed for a specific connection to a compute processor. Each port had its own unique programming requirements. Connected to these ports were different types of COTS processors, again chosen for specific properties such as throughput, size, and cost. The resulting system, although perhaps accurately described as being composed of COTS equipment, was a mixture of input/output processors, interface hardware, and dissimilar floating-point processors. The software development effort was unexpectedly complex, as it was necessary to program each type of processor uniquely and to separately program each type of hardware interface. The resulting software was composed of processor microcode, assembly, FORTRAN, and C code. The complexity of the system made it difficult to maintain.

The SPAN-I and TSPAN systems were clearly distinct. SPAN-I, with its custom hardware, required a single type of interface to mate virtually any processor to the custom-built ring bus. The associated system software development effort was correspondingly

small. TSPAN, on the other hand, minimized custom hardware development, but as a consequence, the software development effort was relatively large. One should not conclude, however, that the use of COTS hardware will increase current software development costs today. Instead, the lesson is that COTS inclusion at that unique period of time, although feasible, was not necessarily advantageous.

The TSPANU system will require development of only two custom hardware components, as previously mentioned. All other components, including linear beamformers, filtering units, and fast Fourier transform (FFT) processors, are implemented in COTS hardware built by Sky Computers. All hardware components are interfaced with the high-bandwidth Sky Channel bus, including the two custom-built boards. This bus defines an addressing space that supports TSPANU scalability requirements. The resulting homogeneous hardware allows the software development team to design and build common components and methods for data transfer, memory management, file management, and error handling that apply throughout the system.

Still left for discussion is the second factor described previously: future signal processing systems must be highly configurable. An instructive way to clarify this concept of a reconfigurable signal processing system is to describe its opposite, a rigidly constructed system. Early sonar processing systems had predefined capabilities. For example, a system might be able to process element data from only one or two sensing arrays, might always perform delay and sum beamforming, and then always follow a predetermined processing path. If analysis needs called for deviation, then a programmer would be required to laboriously divert data from the "usual" course, directing the data stream to newly implemented algorithms. In practice, the range of allowed deviations has been restricted because of software and hardware architectural limitations.

A usefully reconfigurable sonar system allows

- Easy redirection of the data stream
- Accommodation of new input sensor data
- Capture of alternate output results
- Incorporation of alternative algorithms
- Simple re-allocation of computer resources to eliminate detected bottlenecks
- Hardware expansion (or upgrade) without software changes

These qualities are so clearly beneficial that they were required in some form in every processing system developed for our objectives. As a result, all three of our sonar systems (SPAN-A, SPAN-I, and TSPAN) were designed to be and were described by their developers as "programmable" or "reconfigurable." Both descriptions had a basis in truth, but the goals

regarding flexibility were met with only limited success in early systems. The fact is that programmability is a matter of degree, and the more flexible a system is to be, the greater will be its up-front software development costs.

The flexibility of a sonar signal processing system will be the outcome of measured decisions and trade-offs, and will require careful design to be achievable to a degree both useful and affordable. The challenge during the ongoing development of the TSPANU sonar system is to make flexibility a reality by designing a software architecture that supports those goals. The following is a subset of the software design principles being implemented in TSPANU.

- A particular processing task will be easily scalable, without software changes, to use more or fewer individual processors as throughput requirements dictate. All data stream management and routing functions will be performed automatically by software.
- The concept of hardware scalability will extend beyond the chassis. Additional card-cages populated by additional array processors will be added to the system without software modification.
- A text table that is generated by a user will define the allocation of processors and the data interconnections between those processors. The parameters of all processes and interconnections will be defined symbolically in that table, preventing duplicate definitions of identical quantities.
- The system will be data-driven at speeds exceeding real time. When the data flow stops, the system will idle while waiting for more data.
- The recipient of all sonar outputs generated by the system will be a single console workstation. The interface to the workstation will allow upgrade or total replacement of the console without affecting the signal processing system.

The TSPANU system is designed to accommodate (1) hardware expansion or upgrade with few or no software changes, (2) future enhancement of its algorithms without major software modification, and (3) processing of sonar sensor systems that do not currently exist. To achieve these goals, the software architecture supports specific generalized concepts that are designed from the beginning. The use of a single COTS processor for almost all computationally intensive tasks, as well as an integrated high-bandwidth bus, is fundamental to meeting the software design goals.

SONAR DISPLAYS

The current SSD sonar processors employ a display and analysis subsystem separate from the signal processing subsystem (Fig. 2). Our sonar processors have produced many different sonar displays such as

LOFAR (low-frequency analysis and recording), a display of the energy at a given bearing and frequency range over time, and BTR (bearing time recording), a display of the broadband energy covering 360° over time (Fig. 4).

Before 1985, the sonar images produced by the signal processors were printed to thermal hard-copy paper and then analyzed. Each data tape generated about 1,000 sheets of hard-copy paper, or more than 100,000 sheets for a single SSBN submarine mission. Each image was printed on dry silver paper, which was quite expensive in the quantities required. Therefore, the goals of the first display and analysis system for the TSEP, known as DIANA, were to lower costs by reducing the need for hard-copy output and to provide online analysis capabilities. The resulting system could retrieve, unpack, format, and display a 500-KB 4-h full-screen LOFAR image in about 5 s and a 15-h BTR image in about 7 s.

Analysis using the DIANA approach demonstrated the power of display systems for organizing data and gave the users a wide range of optional tools for data selection and analysis. For example, DIANA was the first system in the sonar community to support linked LOFAR capability from a BTR display. Time slices from a set of LOFARs were abutted, forming a linked LOFAR image, which could be used to track an energy source moving through bearing space. On the first-generation DIANA, the user had to manually enter linked LOFAR definitions. Subsequent upgrades improved the manual capability and provided additional methods for defining and analyzing linked LOFARs. The result was that the increased analysis capability provided by the DIANA system was more valuable to the TSEP than the reduction of hard-copy paper costs.

As SSD signal processing systems were enhanced with the development of TSPAN, the existing DIANA system also needed improvements. In 1987, DIANA2 was initiated to store, retrieve, format, and display the large volume of data produced by the TSPAN processor; reduce the time to display a full-screen image; and improve the analysis product by providing more powerful tools and capabilities.

The new system was designed to accommodate data from seven TSPAN processing runs, each containing 6 h of processed image data from two sonar arrays. The DIANA2 system provided new predefined displays, such as polar plots and histograms of the signal-to-noise ratio found in frequency/time subsets. It also
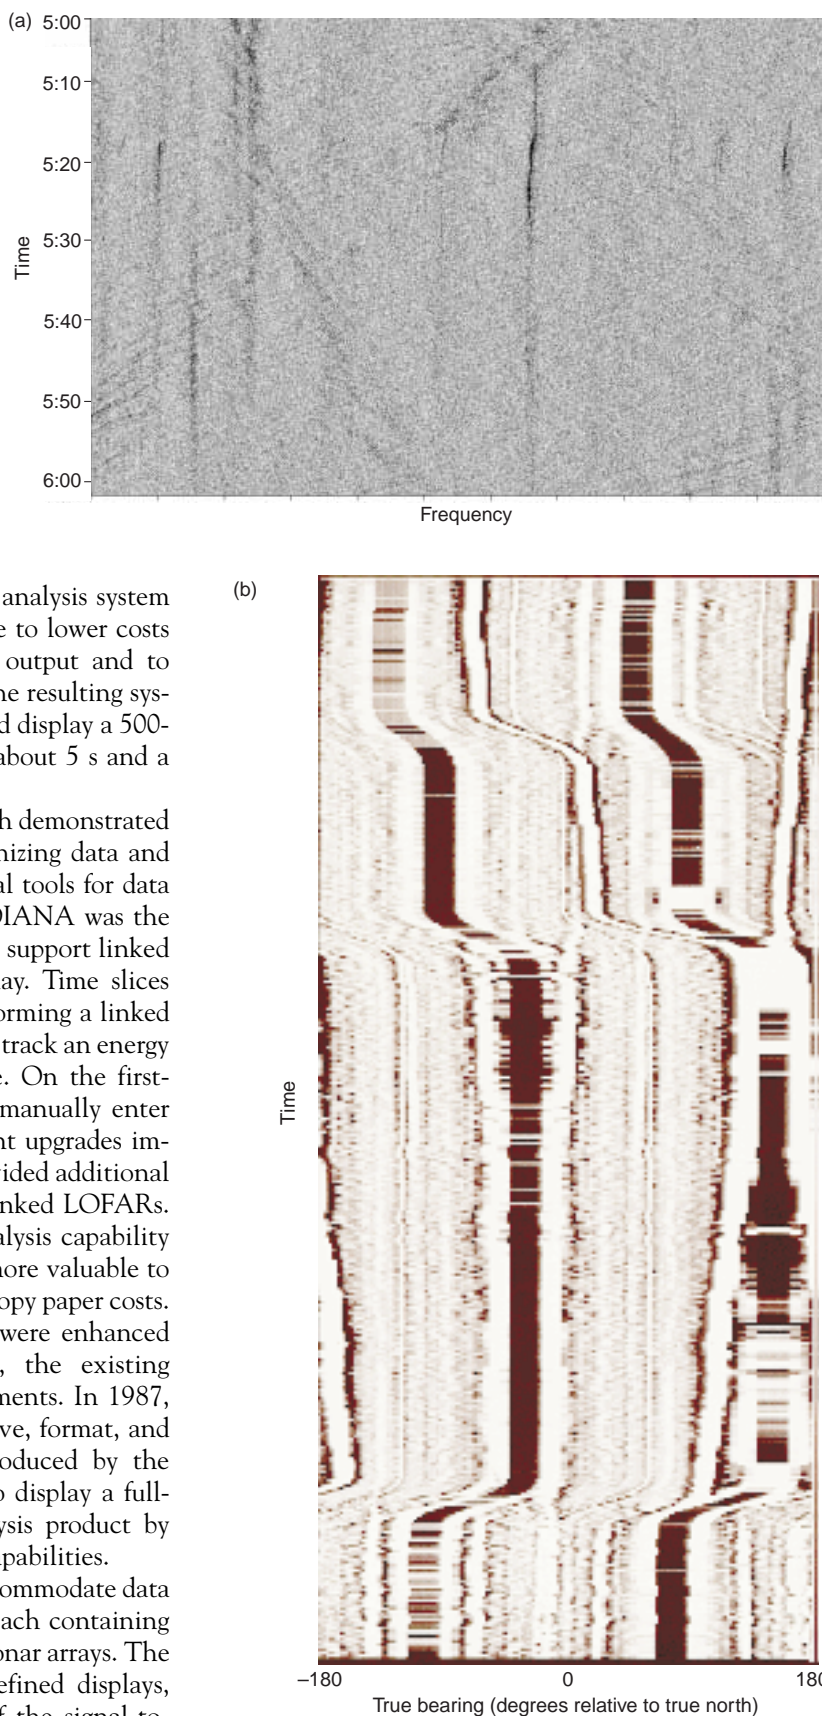
Figure 4. Sample (a) LOFAR and (b) BTR sonar display images.

allowed the user to define customized displays by selecting and organizing the data types to be shown. Data from other sources could be overlaid as well; for example, onboard logged contact information could be overlaid on BTRs. The DIANA2 system is still the primary display system for the TSEP.

In 1993, the Surface Towed Array Surveillance System (SURTASS) Program required a real-time signal processing and display system onboard a Navy surface vessel during an at-sea test of a prototype sonar array. SSD developed the Engineering Display Station (EDS) to meet this requirement. The system receives beamformed spectral data in real time from an external beamformer, processes the data producing 7-bit sonar display images, and stores the processed data to a run-time database containing up to 3000 full-screen images. The EDS gives the operator a wide range of displays and analysis functions for investigating the data during at-sea tests, e.g., full-screen LOFARs and BTRs, linked LOFARs, range-focused BTRs (Fig. 5), harmonic cursors, and other display types. The EDS system was deployed on 12 similar sea tests through 1997 and continues to be the primary display system for SURTASS research and development tests.

In 1994, SSD initiated another system development, the Generalized Acoustic Processing System (GAPS), as a TSEP special-purpose signal processing and display system that would receive data from TSPAN. A project that required use of signal processing, math, and data manipulation techniques and displays could use GAPS to support the needed features. Also, the system was envisioned as a workstation whereby new algorithms could be evaluated and tested before implementation on the TSPAN production system.

GAPS contains an extensible set of processing elements that are combined to perform complex functions. For example, power spectral estimates for a frequency range can be generated from time series data by a processing path comprising the complex mixer, lowpass filter, decimator, FFT, square law detector, sum and dump elements, logarithmic functions, and scale elements. Each element is then governed by a set of options. For example, FFT options include FFT size, overlap percentage, and window type. Another feature of GAPS is an extensible set of analysis functions, all selectable from a menu, including real-time audio and visual playback of time series data and tools for analyzing sonar array health such as wavenumber and hydrophone versus frequency plots (Fig. 6).

GAPS is implemented on single-monitor COTS UNIX-based workstations using X Windows and XView. For rapid prototyping, processing elements and analysis functions use the MATLAB or PVWAVE toolkits. The system emphasizes flexibility and ease of extension over processing speed. The GAPS upgrade (GAPSU) is currently under development to support new data types and formats, combine the DIANA2 functions with the special-purpose processing system, and add analysis functions and improve existing ones. The earlier production analysis systems, DIANA and DIANA2, were written for specific SPAN-A, SPAN-I, and TSPAN output data formats. GAPSU will accept these formats and will be extendible so that new formats can be added. This ability will allow data from sources other than TSPAN to be analyzed using tools developed for production analysis, and will allow analysis of the new data products expected from TSPANU and other potential sources.
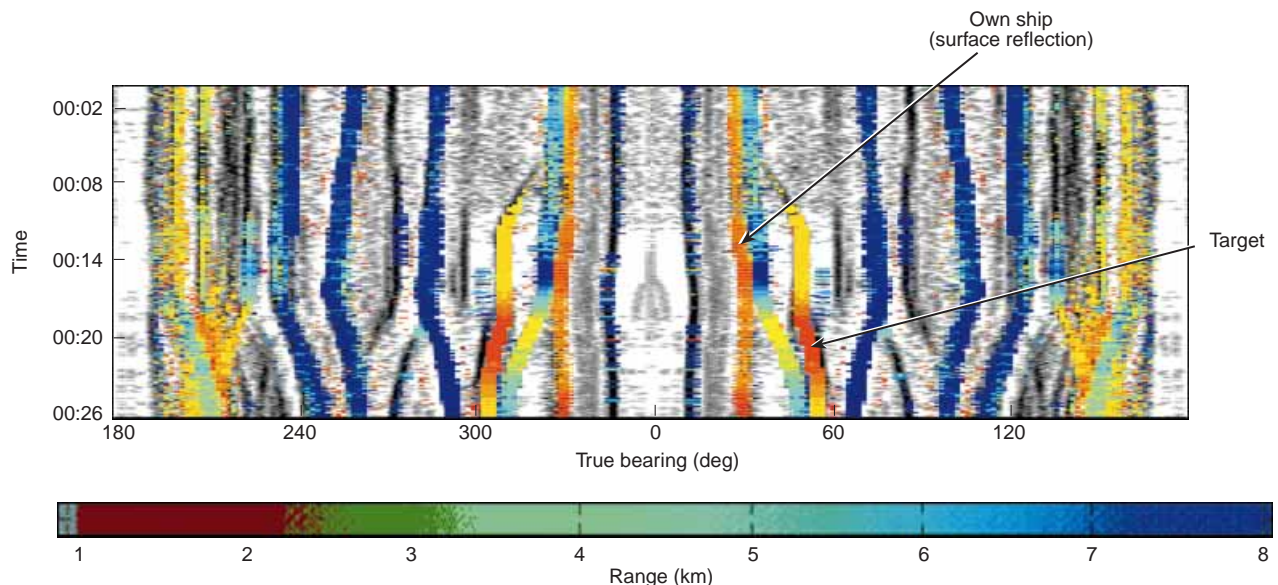
Figure 5. Color-coded multirange bearing time recording on the Engineering Display Station. W. Allensworth of Applied Hydro-Acoustics Research developed the signal processing algorithm that produced this display.
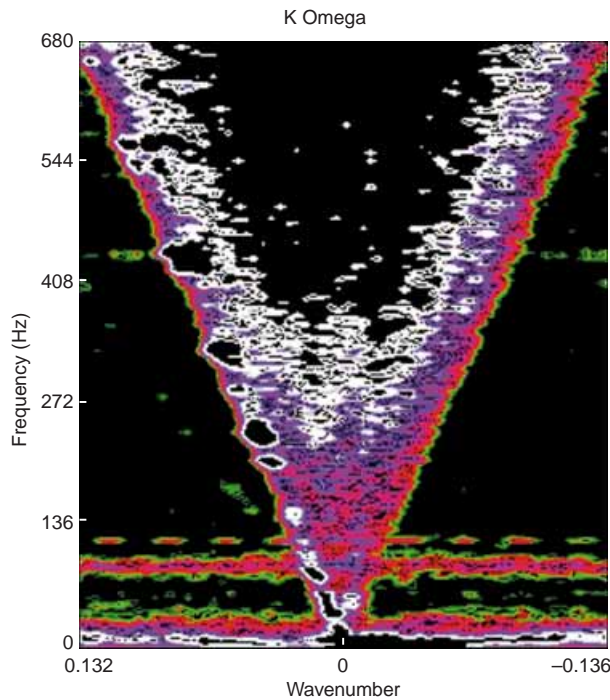
Figure 6. GAPS displays of wavenumber and frequency spectra allow TSEP analysts to verify proper operation of sonar arrays.

The design, development, and maintenance of the DIANA, DIANA2, EDS, GAPS, and GAPSU display systems have yielded valuable information for designing future display systems. Critical to the success of these systems was the interaction and feedback received from the end users, the sonar analysts, during the systems' full life cycles. Widely used COTS hardware and software products will greatly increase the ability to maintain, upgrade, and extend systems. COTS hardware and software on the GAPS and EDS systems have made it easy to upgrade to more powerful hardware and to extend the software for additional functionality.

FRONT-END STIMULATORS

The signal processors developed by SSD produce acoustic display data from prerecorded sonar array element data. Before being able to validate the processors' output using these prerecorded data of unknown quality, the signal processors must be tested under known input conditions. These conditions are provided by FESs, which have evolved in step with the signal processors from simple test tapes to autonomous, high-fidelity stimulators that replace the tape player input.

Today's needs for high-fidelity sonar stimulators are greater than ever. Sonar algorithms are becoming more computationally intensive and are often difficult to develop and initially validate with real-world data. Sonar stimulators provide a completely controllable testbed for the signal processor programmer, a means to inject known data with known characteristics into the signal processors. The information gleaned from the signal processor outputs is inherently matched to the sonar stimulator's inputs. Developers operate with stimulated input until they are satisfied with the performance of the signal processors against the stimulated data. They can then execute their programmed signal processor against operational data. As with signal processors, sonar stimulators are evolving in design and capabilities and progressing from rigid architectures to flexible, scalable architectures. Rigid architectures have limited potential for growth. Flexible, scalable architectures have the ability to grow to support changing requirements.

The original SSD-developed FES became part of the TSEP during the procurement of TSPAN. It is a VME solution and contains a control computer, two Mercury 10-MHz array processors, a data acquisition and monitor board, and a spherical array inverse beamformer board and formatter board. All components are COTS, except for the last two boards. The simulation capability of the FES is basic: a simple acoustic contact that contains narrowband, broadband, and modulation components; a flat, lossless ocean model that propagates the acoustic wavefront in a direct path to the sonar; and an isotropic noise injector that adds background noise into each hydrophone. Array elements are assumed to be static and operating correctly.

The FES has serious limitations. First, the TSEP must process towed arrays whose characteristics cannot be modeled by the current FES. These towed arrays do not fall within the constraints for aperture length, hydrophone count, and hydrophone geometry. Second, processing modes must be developed that operate on signal content the FES does not provide. Finally, in the long term, there is a need to ensure that the FES design is more adaptable to changing requirements. The FES upgrade (FESU), started in conjunction with TSPANU development, is being built with the objective of overcoming these deficiencies and improving long-term versatility. Its design will incorporate an open architecture that supports growth and change as well as the new features required by the next evolution of signal processors. The new features being directly incorporated into the FESU are all applicable to towed array simulations, including

- Synthesis of four simple contacts
- Specification of an order of magnitude higher number of towed array hydrophones in any three-dimensional configuration
- Application of shape distortion models to the towed array hydrophone positions
- Application of failure models to the towed array hydrophone signals
- Simulation of contact signals from four arrival angles
- Planar (infinite range) and curved (finite range) wavefront models

The FESU architecture closely parallels that of TSPANU and GAPSU. The commonality of hardware components among systems reduces maintenance costs, allows for software reuse and sharing among systems, and extends the homogeneity of the TSEP processing suite. FESU uses a Themis 20MP computer for process control and for the operator–machine interface. There is also one SKYBOLT II 9U processor with two daughtercards, each containing four 40-MHz i860 processor nodes. This configuration provides a total of 640 Mflops, over an order of magnitude larger than the numerical capability of the original system. The increase in processing power is needed to support the increase in functional requirements.

The baseline FESU is scheduled for completion in 1999. It will provide the signal content required to develop newer sonar signal processing algorithms and the formats of the sonars new to the TSEP. Just as the requirements of the TSEP signal processing established the requirements for the FESU, future TSEP requirements will likely drive FESU enhancements. They will be modifications to an established system designed to grow to support changing requirements.

RECORDING SYSTEMS

As mentioned earlier, most SSD passive sonar processors have been designed primarily to process array element data recorded on magnetic tape. Most of these tapes have been recorded on systems also developed by SSD. The original recording system, SPARS (Sonar Evaluation Program Acoustic Recording System),[1,4] was the state-of-the-art in recording systems at the time, but compared to today's technology, its capabilities were limited and its size bulky. Its primary features included serial, synchronous shipboard data recording rates of 9 to 10 MB/s on 22-lb reels of tape that held 4.6×10^{11} bits with a nominal bit error rate of 5 per 10^6. Playback capability was limited to less than 40-million bits per second.

The next-generation recording system was TSPARS (Trident SPARS). SSD's experience with the implementation of SPARS provided the basis of specific design scenarios for the TSPARS interface and processing hardware fabricated by Interstate Electronics Corp. (IEC). SSD staff worked closely with IEC to incorporate features that would enhance data handling as well as physical handling of the systems onboard submarines. The TSPARS tape recorder is COTS equipment, as was SPARS. The outcome is a system that maintains individual interface characteristics through asynchronous recording and has shipboard data recording rates of 25 to 100 MB/s on 5-lb tape cassettes that hold 7.9×10^{11} bits with a bit error rate better than 1 per 10^{10}. Playback capability ranges from 10- to 250-million bits per second. Figure 7 shows a diagram of TSPARS hardware for AN/BQQ-5E submarine sonar.

The system has five acoustic data sources. The primary difference between TSPARS and its predecessor is the ability to interface spherical array hydrophone data. This interface captures the 1-bit digital multibeam steering data, accounting for nearly 24 MB/s of input. The various towed array hydrophone data, noise and vibration monitoring hydrophone data, and active intercept data comprise the remainder of the 37.5-MB/s recording data rate.

TSPARS data are obtained as near to the waterborne signal as possible without degrading the submarine signal transmission and before significant onboard processing is initiated. Most TSPARS interfaces receive digital inputs obtained from transmissions between the submarine's arrays and inboard electronics, although TSPARS can perform analog-to-digital conversion if necessary. The interface boxes provide basic transistor-transistor logic electronics to capture signals, provide system isolation near the submarine source, and retransmit data to TSPARS units for further processing. Digital signals, once provided to TSPARS interface circuit boards, are further processed in preparation for recording.

The TSPARS design encompasses a VMEbus architecture for control and low-rate data transfer. Each drawer has multiple power supplies to provide redundancy, a central processing unit, and a programmable read-only memory circuit board. In addition it is modularized with a backplane interconnection. Individual interface circuit boards contain software and firmware configured appropriately for the interfaced signals. In addition to processing acoustic data, the interface circuit boards extract submarine system status
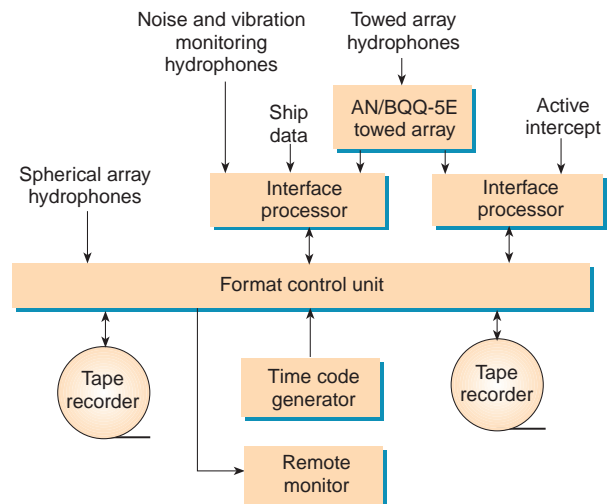
Figure 7. Block diagram for the second-generation SSD-designed sonar data recording system (TSPARS) for the AN/BQQ-5E application.

data, monitor themselves, and provide combined status and configuration information for recording on the VME-bus.

Redundancy of critical, single-point circuit boards is used in TSPARS to achieve the desired mean-time-between-failure rate. Since TSPARS must have a minimal impact on submarine crew members during operation, the redundant cards are housed in the format control unit and are activated by a single switch. The TSPARS components associated with individual interfaces are isolated from one another. Should a single interface failure cause corruption of the data stream, each individual interface can be secured and eliminated from the recorder data bus. The null block data rate will adjust automatically to compensate for this loss and maintain the constant recorder data rate.

As mentioned earlier, the data rate capability of TSPARS is significantly greater than that of SPARS, but with advances in technology that permitted the increase in capacity to result in less space required aboard the submarine. The increased capability of TSPARS has provided for the first time the ability to collect all major acoustic sensors for all SSBN configurations.

The current TSPARS tape recorder is a ruggedized, 19-in. rack-mounted chassis with shock-mounted tape deck. Its recording format is helical scan, using the industry standard ANSI ID-1 on large COTS D1-L tape cassettes widely found in the digital video broadcast business. These recorders use a powerful, interleaved Reed-Solomon encoding scheme to guarantee a 10^{-10} corrected bit error rate. The inputs consist of one wideband channel of 8-bit parallel, emitter coupled logic–transmitted data from the format control unit and one longitudinal auxiliary channel of IRIG-B time code from the time code generator/translator. A second longitudinal channel for voice annotation is not used. The recorder is controlled through the format control unit using an RS-422 serial link.

Beginning this year, the definition and specification of TSPARS 2000 will begin. This system is envisioned to be an acoustic recording system plus workstation for submarine sonar operators. Current TSPARS interfaces capture sensor data before extensive submarine system manipulation has occurred and, as such, provide a source of data suitable for off-line, onboard examination in parallel with recording for later analysis. A combination of physical size and processing requirements will determine the display information made available. For example, displays providing the relative amplitudes of sensor outputs for all arrays could be provided to give the submarine crew real-time assessments. The workstation portion of TSPARS 2000 would provide the submarine with additional tools rather than duplicating those functions already performed by existing submarine systems.

FUTURE DIRECTIONS

Within the next 5 years, SSBN sonars will gain significantly increased processing capability as part of the Navy's Acoustic Rapid COTS Insertion Program. These changes will require improvements to existing analysis tools available on TSPANU and associated TSEP equipment. However, the basic architecture of the new TSEP systems will accommodate these changes in a straightforward, cost-effective way. Plans are already under way to upgrade the processor technology of the current TSPANU and to accommodate higher recording bandwidth in TSPARS 2000. These efforts will ensure that state-of-the-art SSBN acoustic performance is maintained.

REFERENCES

[1]South, H. M., "Digital Recording and Signal Processing Systems for Hydrophone Arrays," *Johns Hopkins APL Tech. Dig.* **4**(3), 212–218 (1983).
[2]South, H. M., "High-Speed Processors for Sonar," *Johns Hopkins APL Tech. Dig.* **11**(1–2), 105–113 (1990).
[3]Anderson, V. C., "Digital Array Phasing," *J. Acoust. Soc. Am.* **32**, 867 (1960).
[4]South, H. M., "High Data Rate Recording and Processing Systems for Hydrophone Arrays," in *Oceans '80: An International Forum on Ocean Engineering in the '80's*, IEEE 80CH15727, Piscataway, NJ, pp. 263–266 (1980).

THE AUTHORS

HUGH M. SOUTH is a program manager in the Undersea Systems Program Area. He received a B.S. degree in electrical engineering from Rice University in 1971 and a Ph.D. in the same field from The Johns Hopkins University in 1981. Dr. South came to APL in 1975 and became supervisor of the SPAN Signal Processing Laboratory in 1979. He later became supervisor of the Signal Processing Group in 1985 and program manager for Surveillance and Acoustics in 1997. Dr. South served as technical sponsor for the developments of the SPAN-I and TSPAN sonar processors, and he is currently managing APL's work on the Advanced Deployable System. His e-mail address is hugh.south@jhuapl.edu.

DAVID C. CRONIN is a member of APL's Principal Professional Staff. He received a B.S. degree in aeronautical engineering from the University of Maryland and an M.S. degree in mechanical engineering from the University of Southern California. In 1967, he joined APL as an analyst in the Strategic Systems Department's Launcher/Ship Group. In 1975, he transferred to the Sonar Evaluation Program and was instrumental in the development of that program. Today, Mr. Cronin is the TSPARS project manager. His e-mail address is david.cronin@jhuapl.edu.
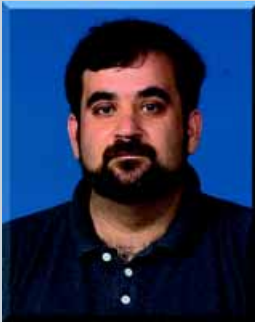
SAMUEL L. GORDON received his B.A. in mathematics from the University of Maryland, Baltimore County, in 1980. He joined APL in 1991 and is a senior staff member of the Signal Processing Group. Mr. Gordon has worked on a variety of graphics systems for sonar analysis. He is currently the lead engineer on the Generalized Acoustic Processing System. His e-mail address is samuel.gordon@jhuapl.edu.

TIMOTHY P. MAGNANI is a software engineer in the Strategic Systems Department's Signal Processing Group. He received his B.S.E.E. degree from Virginia Tech in 1980 and his M.S.E.E. degree from The Catholic University of America in 1987. Mr. Magnani came to APL in 1996 with a background in sonar simulation to head the Front End Stimulator Upgrade effort. His tasks include systems design, algorithm and software development, and hardware design. He has also assisted in the design and development of the Multiple Image Coordinate Extraction (MICE) platform to provide remote targeting capabilities from unmanned aerial vehicle telemetry and imagery. His e-mail address is tim.magnani@jhuapl.edu.