SONAR TECHNOLOGY - PAST and CURRENT.

Tom Curtis

Curtis Technology (UK) Ltd, Weymouth, Dorset, DT4 7BS

tomc@curtistech.co.uk

1. INTRODUCTION.

The changes in the political climate in recent years have significantly changed the way in which military sonar systems are designed and procured.

Prior to 1990, military requirements were a major technology pull for technology developments, with large military research budgets funding semiconductor, system and software research. However, the end of the cold war brought major changes in the perceived threat that Navies needed to counter and, at the same time, the so-called "peace dividend" resulted in major changes in military research funding in many countries. As a result, military systems now rely much more heavily on available commercial technology. In recent years, sonar system design has migrated from custom silicon technology, using "hand crafted" systems and software, to the use of open architecture COTS systems (commercial off the shelf), using third party operating systems and software.

This paper briefly reviews the technology deployed in some of the military sonar applications over that period and outlines how current commercial technology realisations differ from those of the previous generation and provide improved system performance, with a reduced cost of ownership.

2. TECHNOLOGY BACKGROUND

Although the basic principles of passive sonar were understood in the sixteenth century, the technology that allowed that understanding to be translated into the active and passive sonar systems we know today did not appear until much later. The fundamental understanding of underwater acoustics and the initial development of transducers and hydro-phones stem mainly from work carried out in the period 1820 -1920. World War I gave a major impetus to sonar system development and large numbers of passive systems were deployed during the course of the war: in all some three thousand hull-mounted and towed array systems being fitted in this period. Active sonar was first demonstrated against operational targets in 1918 when Langevin detected submarine echoes from distances up to fifteen hundred metres, using piezo-electric transducers and vacuum tube amplifiers.

The next twenty years saw steady progress in the understanding of underwater acoustics and in the application of electronics to military sonar systems. The main interest in military systems at that time was towards high frequency sonar with compact transducer arrays but many other system components, such as sonar domes and range recorders, were developed during that period. On the commercial front, depth sounders and fathometers became available from a number of sources in the UK and the USA.

The start of World War II brought about a further increase in activity and acoustic mines, acoustic homing for torpedoes and high frequency scanning sonar were all first developed around that

time. A variety of sonar countermeasures and counter-detection techniques were also demonstrated. For example, the Germans developed acoustic cladding to reduce submarine detection. Major milestones in recent years have included the 'rediscovery' of towed array systems for long-range surveillance in the late 1960s and the development of long range, low frequency active systems.

Most of the recent advances in sonar system capability have been underpinned by improvements in electronic technology, with sonar engineers quickly exploiting new devices as they become available, firstly with the advent of thermionics, then the transistor and, more recently, integrated circuits and digital signal processing. These technologies, and the improvements in system performance that they have provided, have steadily transformed sonar system design from the analogue domain of linear circuit analysis and synthesis into the digital domain of computing and signal processing algorithms.

2.1 Signal Processing Algorithms.

In a very general sense, the basic processing functions used in sonar systems are similar to those employed in radar or telecommunications and are well reported in the literature [1,2]. Some of the algorithms routinely implemented on sonar systems are listed in Table I. Current systems use mainly the linear vector operations listed in the Table. They exploit, for example, transform and correlation/convolution methods for conventional detection and the more processor intensive matrix processing algorithms for adaptive and high-resolution systems.

DSP ALGORITHMS

MAINLY TRANSFORM TECHNIQUES, CONVOLUTION/CORELATION,
FILTERING AND LINEAR ALGEABRAIC METHODS

TRANSFORMS	-	FFT, various number theoretic transforms, Hough, Hadamard, etc
FILTERING	-	correlation/convolution, FIR/IIR filters, Interpolation/decimation, median filtering, Kalman, LMS adaptive, etc.
MATRIX ALGEBRA	-	Matrix multiplication/inversion/rotation, SVD, eigenvalue/eigenvector techniques, etc.
MISCELLANEOUS	-	mapping (grey scale, etc), sorting, etc.

TABLE 1 – Some Typical Sonar DSP Functions

The detailed implementations needed for sonar DSP are however often very different from those used in other fields: simultaneous surveillance cover over a wide field, combined with good detection and localisation performance, requires the use of large sensor arrays and multi-channel processing. A number of different detection processes are often provided simultaneously on the same data sets, to ensure a higher probability of detecting a variety of targets with different characteristics over a range of operational conditions. Typically for a passive system, broadband

correlation, energy detection, narrow band spectral analysis, vernier detection, DEMON analysis, transient analysis, etc, are all required simultaneously for target detection, classification and tracking. As a result, sonar systems require overall throughputs comparable to, and often well in excess of, those used in other digital signal processing applications.

PASSIVE SYSTEM

Say –
256 element array, 0 – 1 kHz
256 beams
Narrow band surveillance, vernier,
broadband, transient, DEMON, etc

BEAMFORMING	- 200 million multiply-accumulates/second
SIGNAL PROCESSING	- 130 million multiply-accumulates/second

TABLE 2 – Processing Throughput for Baseline Passive System

Even when using relatively simple transform and correlation based processing algorithms, current generation sonar processor throughputs are daunting: Table 2 indicates the processing power required for the various operations in a baseline passive system. In more realistic applications, with large area arrays and more sophisticated receive processing, these figures increase by two or three orders of magnitude. Consequently, although the per-channel bandwidth for sonar processing systems may be low, the large number of processing channels, and the need to provide a range of different detection and localisation processes simultaneously, requires systems with very high throughputs, typically as high as 10^{12} arithmetic operations per second. Achieving this level of performance with commercial DSP parts and standard architectures is an interesting exercise.

2.2 Digital Technology Evolution

Digital signal processing methods have been employed in sonar systems since the late 1950s [3]. Early schemes used fairly crude signal processing algorithms and the operational performance they achieved was limited by the technology around at the time [4]. Initial designs were limited by the lack of compact storage technologies and by the limited functionality provided by the arithmetic processors available at that time. Many techniques were investigated to provide acceptable system performance within these technology constraints [5].

Random access and read-only memories, arithmetic/logic units (ALU's) and simple microprocessors became available in the early 1970s. Whilst these parts were available commercially, much of the underlying technology was funded either directly or indirectly from US space and defence budgets. These devices allowed programmable signal processing architectures to be developed for a range of conventional sonar processing applications [6]. Within a few years, a wide spectrum of LSI devices were available and by the early 1980s, high speed arithmetic processors, single chip micro-computers and high density RAM were available that provided sufficient processing power to build powerful conventional and small real-time adaptive systems.

With the introduction by IBM of the PC in August 1981, massive commercial investment in the areas of memory and processor development soon followed. Rapid progress was made in the area of high-density silicon technologies, driven primarily by this commercial PC market. Initially much of this development was carried out in the US but investment in the Far East, particularly in Japan, soon started to make inroads into the US technology lead and to gain market share in the commercial memory area. As a counter to this perceived shift in technology expertise, the US military funded a technology program, the Very High Speed Integrated Circuit (VHSIC) program, aimed at making sub-micron silicon technology available to the US military systems houses. A similar programme, the Very High Performance Integrated Circuit (VHPIC) programme [7], linked to the more commercial Alvey IT Programme, soon followed in the UK. These two programmes were the last major technology initiatives funded by the military: they were expensive programmes, started during the Cold War, aimed at providing silicon devices that could counter the threats considered valid at that time. By the time they were completed, the world order had changed: the threats the technology had been developed to counter had changed substantially. In retrospect, although both programmes achieved some very capable silicon, the technology pull from continuing commercial investment and the change in threat due to the end of the Cold War resulted in very limited application of the resultant silicon devices. However, much of the process technology developed by the US foundries was carried over into their commercial process lines and certainly helped the rapid emergence of a viable sub-micron technology, which formed the starting point for the commercial deep-sub-micron processes currently available.

So, by the beginning of this century, deep sub-micron processes were available from many foundries word-wide. With this level of technology, various standard DSP parts are available from a number of manufacturers [8,9,10], with geometries in the range 250 to 130 nm. Typically, these parts contain programmable micro-code ROM store, RAM for data storage and 24- or 32-bit ALU's (either floating point and fixed) and operate at throughputs of several hundred million arithmetic operations per second. Devices have generally been designed for single processor systems, with internal chip architectures optimised for FIR or FFT type operations, using real-value arithmetic. As a result, they are sometimes difficult to use in applications needing complex vector operations and in multi-processor systems, although many devices now being developed have the hooks for multi-processor operation built in.

Leading edge application specific integrated circuit (ASIC) technology currently provides device geometries down to around 90 nm, with every prospect of the feature size being scaled to below 65 nm within the next year or so. Current silicon ASIC geometries support complexities in excess of ten million gates on a single die. This level of integration allows complete system-on-chip (SoC) designs to be considered. The cost of custom design for small geometry SoC has always been a problem but is cost effective in high value-added or large user volume applications. However, the improvements in performance of SoCs, achieved by transistor geometry scaling in recent years, are starting to hit some fundamental limits. Recent geometry scaling from 130 nm to 90 nm have not produced the increase in operating speed that previous scaling steps have realised. Excessive transistor leakage at the smaller feature sizes on some processes has resulted in die where the resultant dc dissipation has caused major thermal problems - this sets the limit to device operating frequencies rather than the more normal dynamic power dissipation limits that dominated earlier larger geometry processes [11].

In order to reduce these process problems, exotic technology tweaks like using strained silicon and high-k dielectric gate materials are being incorporated into the process flow by the major foundries. Consequently, it is becoming increasingly difficult for the small systems houses to get access to these technologies and develop custom DSP parts on "leading edge" processes: the up-front costs for computer aided design (CAD) tools and the foundry charges for that level of design are just too prohibitive.

The way around this problem in the past was to develop custom hardware at printed circuit board (PCB) level using "bit-slice" parts [12], for example using high-speed static memory and various multiplier or multiplier/accumulator devices. However, with increasing chip complexity, these low complexity DSP "building block" parts have fallen off the technology "trailing edge". However, recent advances in FPGA technology have started to fill that gap.

FPGAs with random access memory (RAM) based configurable logic blocks (CLBs) have been around for some time, from the major players such as Xilinx [13] and Altera [14]. The combination of programmable logic and the ability to use the CLB RAM blocks as distributed memory made these devices immediately attractive in many DSP applications.

More recently the manufacturers have started to add features to their basic FPGA architectures in order to differentiate their products in the market place. Xilinx, for example, initially added blocks of larger fast dual port RAM to their devices, making the chips useful in areas such as corner turning for DSP. Arrays of embedded dedicated multipliers were soon added to this block memory feature and current generation FPGA devices are available with hundreds of 4kbyte memory blocks and 18x18 bit multipliers, together with upwards of 100,000 programmable logic slices.

Using these programmable parts, the system designer has much more freedom in architecture design and the vast amount of DSP-related resources available on chip allows processing engines with massive throughput to be considered.

The emergence of these technologies and, more particularly, the widespread availability of the CAD design tools to support them, requires a careful review of the techniques that have been employed in the past to develop digital processing systems, if full advantage is to be made of technology improvements. The main question in the exploitation of these techniques becomes not *'what can the technology support?',* but *'how* do *we utilise this level of processing power in a cost-effective, timely way?'.* In particular, new algorithms must be developed, or existing algorithms massaged, to provide a better match between the system requirements and the enabling technology. Distributed processing and communications architectures must be developed to utilise the parallelism inherently available on silicon to support the high throughput processors needed for current and future applications.

2.3 Distributed Heterogeneous Systems.

In the past, high throughput systems have been implemented using either dedicated custom hardware processors or programmable software-based multiple array processors. Both of these approaches have their inherent advantages and disadvantages. Dedicated hardware can achieve very high throughputs and wide bus transfer bandwidths but lack flexibility and are difficult to modify or reconfigure to meet changing operational needs. Early hardware systems used many different card types and were costly to design and support, especially in military applications. Software based systems are programmable and require only a limited numbers of card types but system software has proved problematical, particularly in large, real time DSP applications. System level design tools were limited and excessive software complexity and high integration costs common in many systems.

Consequently, it has been difficult to develop high throughput signal processors using either hardware-or software-based technologies. Both approaches have required long and costly development programmes and often produced inflexible systems that were difficult to modify and upgrade during their operational lifetimes. Processing schemes, with architectures between dedicated hardware and totally flexible software-based systems that exploit silicon ASIC technology have been developed for military applications. These aim to reduce system development time-scales and cost using a combination of COTS and FPGA technology and use improved CAD, to provide sufficient flexibility and processing power for next generation DSP applications.

Developing system level architectures capable of integrating large processor arrays without losses in individual processor efficiency is a major problem area. The decomposition from system specification onto silicon resources is straightforward in some applications, for example in synthetic aperture radar systems, where the required functions have a high degree of granularity. Here the problem can be decomposed into well-defined individual tasks and implemented using a highly pipelined collection of autonomous processors. Also, when the system functions required are sufficiently well-ordered, they can be mapped onto a regular array of identical closely-coupled processors, as for example in systolic processing schemes [15]. However, most sonar signal processing problems are not well structured at the system level. They often require 'long range' interactions between various sub-systems: flow graph network architectures [16], which support these interactions, are more attractive than the more regularly structured architectures that support only 'short-range' interconnection paths. In many applications, the overall system function can be built up from a limited set of sub-functions or primitives. For example, many practical sonars can be configured using 'building blocks' performing say beamforming, filtering, transform processing, matrix manipulation, normalisation, display, etc. This allows a relatively straightforward mapping from a conventional system block schematic onto modular hardware (or software) via a signal processing flow graph that describes the overall sonar function.

In earlier systems, this flow graph architecture was supported by custom designed bus and control structures, but in recent years Component Peripheral Interconnect (PCI) methodology [17], developed for commercial PCs, has been used. Using PCI (and its derivatives) has a number of advantages: firstly, standard chip sets are available that look after the detailed bus transfer protocols and secondly, migrating signal flow onto open architecture PCI allows system software to be developed using “industry standard” operating systems and software.

3. SYSTEMS REALISATIONS.

The following paragraphs outline briefly two implementations of a similar generic sonar surveillance system. The first implementation was developed in the early 1980s: it is a cabinet-based system, housing upwards of eighty custom-designed 6U PCB cards and MSI TTL components. The second, developed in 2002, is a single shelf system using four COTS cards housing programmable DSP and FPGA components.

3.1 1980s System

A photograph of the earlier system is shown in Figure 1. The cabinet contains three shelves, each housing twenty eight 6U PCBs, custom back-planes and power supplies. The lower shelf houses the beamforming subsystem, the middle shelf, the surveillance and vernier processing and the upper shelf the display and history archive system. The control and BITE systems were driven by TI TMS9900 microprocessors, one per shelf, and the system configured via dedicated tracker-balls, switches and a touch panel, housed in the display console.

One of the main constraints in the system design was to keep the overall control mechanisms as straightforward as possible and to minimise the cost of system software. To this end, fairly simplistic processing algorithms were developed that used the hardware efficiently but needed minimal control overhead, both in terms of hardware and code.

Analogue data from the array was digitised using successive approximation ADCs, with thick-film hybrid anti-aliasing filters. Because of the limits imposed by this analogue filter/ADC technology combination, a slow AGC system was used prior to digitisation to maximise signal dynamic range.

Figure 1 – System Photograph (circa 1984)

The beamforming system used a direct time-domain space/time memory design, with selectable array weighting and inter-sample interpolation to minimise beam side-lobe levels. It was built using TRW MSI multiplier/accumulators [18], static RAM and EPROM.

The signal processing level again used mainly TRW parts and a straightforward radix-32 based Fourier transform algorithm was developed to implement 1024-point complex FFTs for the surveillance processing. Vernier processing was implemented using a frequency-domain "interpolation and over-lapped block sum" algorithm, again with minimal control overhead as one of the main design criteria.

The display and archive level provided LOFAR history storage, using large numbers of 64k DRAM and a custom designed four display-head graphics system.

The system power requirements were around 1 kilowatt, the cabinet stood just over a metre high. Production systems were built by Ultra Ocean Systems, Weymouth and sold for around £300k in 1982.

3.2 2000s System.

A photograph of this system is shown in Figure 2. The single shelf system provides full beamforming, surveillance and vernier processing on a variety of passive arrays, as well as broadband and transient analysis. Update rates and analysis frequency resolution options are significantly better than those implemented on the previous system, with faster than real-time operation available on archived data. The system contains comprehensive BITE and confidence check facilities. The DVD-RAM recorder provides in excess of 12 hours of raw data recording on a typical passive system and the MO drive allows over 1500 high resolution screen dumps to be archived.

Individual units are networked together via T-base 100 ethernet to provide multiple operator stations, with raw and processed data from each unit being accessible across the network.

The unit is based on a standard 4-slot, 3U Compact PCI chassis. It is configured around a commercially available single card computer, the Gespac PIII-SYS 650E [19]. This provides the standard PC system interfaces and resources (IDE and SCSI, mouse, RS232 and Ethernet Tbase10/100 interfaces, an XVGA video adapter - 1280x1024 pixels by 16M colours, memory, RTC, etc). Apart from the CPU card, the system uses three other OEM COTS boards, viz. an array interface card, a programmable DSP card and an AC97 compatible audio sound card. Also included in the rack are standard bulk storage devices for raw data archiving, data logging and for storing screen dumps.

The array interface card (the TAXI I/F board) connects the array terminal unit to the DSP system via dual FDDI (175 MBPS) or HOTLINK (400 MBPS) serial links. It also provides 64 Mbytes of high speed buffer memory and some supporting control and processing logic, using FPGA devices.

The DSP card (the BCVP) uses the Sharp LH9124/9320 vector processor DSP chip set [20,21], again with high speed buffer memory and FPGA logic. This card is described in more detail in Reference[22].

The AC97 audio system provides sonar aural output and also contains system interfaces to a DVD-RAM drive for raw data archiving and to an MO drive for screen dumps. This combination also provides DVD movie and CD playback for training (or other) purposes.

System software was written using a commercial RAD environment (Borland DELPHI) and the system OS used is Microsoft NT4. NT4 is not strictly a real time OS, so various tricks, for example setting suitable processing priorities in the software modules/drivers and the judicious use of hardware FIFOs, were employed to ensure signal integrity. Using NT4 and Delphi essentially provides complex control structures for free, so the main design requirement on this

newer system was to maximise hardware efficiency, rather than minimise control overheads. Because the prime “number-cruncher” used in the system was a programmable DSP engine, most of the processing algorithms were massaged to fit onto the architecture of this engine and consequently most are implemented in the frequency domain.

Analogue data from the array is digitised using high dynamic range sigma-delta ADCs. These ADCs provide better dynamic range and channel-to-channel matching than the analogue filter/ADC technology combination used in the older system, so AGC prior to digitisation to maximise signal dynamic range is no longer needed.

The beamforming system uses a frequency-domain fractional DFT algorithm [23], realised on the BCVP card, to maximise hardware efficiency. The signal processing uses a radix-16 based Fourier transform algorithm to implement 4096-point complex FFTs for the surveillance processing. Vernier processing uses a frequency-domain zoom process. Both of these algorithms were realised on the BCVP card.

The display processing and archive history storage are implemented using the standard PC system memory and the XVGA graphics system on the CPU card. Other processing algorithms and operator aids are provided on the new system that were not available on the earlier one. The system power requirements are around 60 watts and the system sells for around £18k currently.

Figure 2 – System Photograph (circa 2002)

SUMMARY

This paper has outlined some of the changes in enabling technology, processing algorithms and systems architectures that have affected sonar development in recent years: it has demonstrated the size and cost reductions possible using current generation COTS hardware compared to the custom design approach of the 1980s.

REFERENCES

[1] Kay, S.M. and Marple, S.L.: "Spectrum Analysis - A Modern Perspective", Proc IEEE, Vol 69, No 11,1981.

[2] Knight, W.C. et al: "Digital Signal Processing for Sonar", ibid.

[3] See for example:

Faran, 3.F. and Wells, R.: "Correlators for Signal Reception", Harvard Acoustic Lab Tech Memo, No 27,1952.

Anderson, V.C.:"DELTIC Correlators", Harvard Acoustic Lab Tech Memo, No 37, 1956.

[4] Anderson, V.C.: "Digital Array Phasing", JASA, vol 32, 1960.

[5] Remley, W.R.: "Some Effects of Clipping in Array Processing", JAS A, Vol 39, 1966.

[6] Proceedings of Conference on Real-time, General Purpose, High-Speed Signal Processing Systems for Underwater Research CP-25, SACLANT ASW Research Centre, La Spezia, 1979.

[7] “VHPIC and the UK Alvey Programme”, W Fawcett, In AGARD The Impact of Very High Performance Integrated Circuits on Radar, Guidance and Avionics Systems, 1985.

[8] “ADSP-BF531/ADSP-BF532/ADSP-BF533: Blackfin® Embedded Processor Data Sheet”, Analog Devices, Norwood, MA.

[9] “TMS320C6414, TMS320C6415, TMS320C6416 Fixed-Point Digital Signal Processors”,

Texas Instuments, Houston, Texas.

[10] “MSC8102 Quad Core 16-Bit Digital Signal Processor Technical Data”, Motorola.

[11] “Low Power SoC Design”, EDA Weekly Review for May 24, 2004.

[12] "Control Ordered Sonar Hardware – COSH: A Distributed Processor Network for Acoustic Signal Processing", T E Curtis, A G Constantinides and J T Wickenden, Part F, Proc IEE, 1984.

[13] “Virtex-II Platform FPGA Handbook”, Xilinx, San Jose, CA.

[14] “Stratix II Device Family Data Sheet”, Altera Corperation, San Jose, CA.

[15] “Systolic arrays for high throughput front-end signal processing”, P J Hargrave, International Specialist Seminar on the Design and Application of Parallel Digital Processors, Lisbon , Portugal, 1988.

[16] “A Common Operational Software (ACOS) Approach to a Signal Processing Development System”, Proc ICASSP, Boston, 1983.

[17] “PCI Hardware and Software”, E Solari and G Willse, pub by Annabooks, San Diego, USA, 1996.

[18] “TRW LSI Multiplier-Accumulators”, TRW Inc, 1979.

[19] “Gespac PCISYS-PII/III Processor”, Data Sheets pub by Gespac, Geneva, Switzerland.

[20] “LH9124 Digital Signal Processor Users Guide”, Sharp Electronics Corporation, Camas, WA, USA, 1992.

[21] LH9320 Memory Management Unit, Users Guide, Sharp Electronics Corporation, Camas, WA, USA, 1995.

[22] “Compact Sonar Surveillance Processing Systems”, T E Curtis and M J Curtis, IOA Conference on Sonar Signal Processing, Loughborough, 2004.

[23] See for example

“Principles of Sonar Beamforming” at http://www.curtistech.co.uk/papers/beamform.pdf