SONAR SIGNAL PROCESSING BASED ON COTS COMPONENTS

D. Maiwald, U. Fass, M. Heumann, R. Kettenburg, and B. Rose

ATLAS ELEKTRONIK GmbH
Sebaldsbrücker Heerstr. 235, D-28309 Bremen
email: maiwald@atlas.de

ABSTRACT

This paper addresses the application of commercial off-the-shelf (COTS) components to sonar signal processing. The Active Towed Array Sonar Sytem (ACTAS) is introduced, the COTS components used for setting up the signal processing are described, and key experiences with the use of COTS are illustrated. Furthermore, the main software components necessary for system set-up are described. CORBA is used for communication and data transfer between software components.

1. INTRODUCTION

This paper deals with the application of commercial off-the-shelf (COTS) components to sonar signal processing. The technology is used in a low frequency Active Towed Array Sonar System (ACTAS) which is described in section 2.

Customers of sonar systems have several requirements relevant to the overall signal processing design: They want their system to be supplied e.g. by spare parts for more than 20 years. A certain amount of growth potential is expected mainly because new algorithms can be foreseen during the long time of operation. Technologies like multistatic operation are e.g. now tested in sea trials and will be available in a few years. In summary, the customer expects a flexible system with interfaces which can be changed. Finally, shrinking budgets of the navies lead to a situation where the sonar companies have to use COTS products in order to realize systems well-priced.

In the past, sonar signal processing has been almost always set-up by dedicated hardware and by digital signal processors. Recently, the improvements of general purpose processors have been considerable because of their use in communications systems and the consumer market. It is possible since a few years to set-up sonar signal processing by using these general purpose processors. Reasons for breaking this new ground are:

- General purpose processors like the Intel Pentium processor 4 can be bought of the shelf. Because of their use in the consumer market, there is a continuous improvement in processing power and reduction in power consumption. Furthermore, the presently available processors and their successors will be available for the next decades because of their distribution and the enormous amount of software already written for them.
- There are several vendors of the corresponding boards even for ruggedised versions. This means that the arguments given above for the processors are also valid at the board level.
- Software development tools for general purpose processors are very well developed which reduces debugging

and integration time drastically. Furthermore, software developers e.g. for C++ or JAVA can be more easily recruited then specialists for DSP programming.

In summary, the form factor defined as processing power per volume is sufficiently high for the realization of certain sonar signal processing systems.

The intentions of this paper are:

- To give an introduction to the sonar system in which the described technology is used.
- The key COTS components for setting up the system are depicted.
- Experiences and results of using COTS components are given. Although it is quite simple to realize some signal processing by PCs connected via ethernet, the effort to build a product is considerable.

2. ACTIVE TOWED ARRAY SONAR SYSTEM

2.1 System Components

The active towed array sonar system is illustrated in Figure 1. The complete system consists of three sub-systems: the towed system, the winch and handling system (WHS), and the sonar system. The main wetend components are the towed body (TB) and the towed array (TA). Both are towed by a ship which moves with a velocity V_P. The towing depth of the body and the receiving array are approximately identical and can range from several ten to several hundred meters. It is controlled by the tow cable length between body and towing platform and the towing platform speed. One great advantage of active towed array sonar systems compared to the more traditional hull mounted sonar systems (HMS) is the potential of adapting the sonar operating depth according to the sound velocity profile of the operating sea area. The other advantage is the lower operating frequency compared to the HMS which gives three to four times greater detection ranges. Of course, a HMS is always available for operation whereas a towed system has to be deployed before operation.

2.2 Towed System

The active transmission unit is placed in the TB. Horizontally, the transmission array supplies an omnidirectional pattern. Vertically, it avoids the insonification of the sea surface and the ocean bottom. A receiving line array consisting of N equi-spaced identical sensors is attached to the body with a tow cable of several hundred meter lengths. The array can be a single line array, a triplet array, or a twin array depending on the purpose of the sonar system. A triplet or a twin array are needed for left/right discrimination. At low frequencies a twin array outperforms a triplet array especially at high tow-
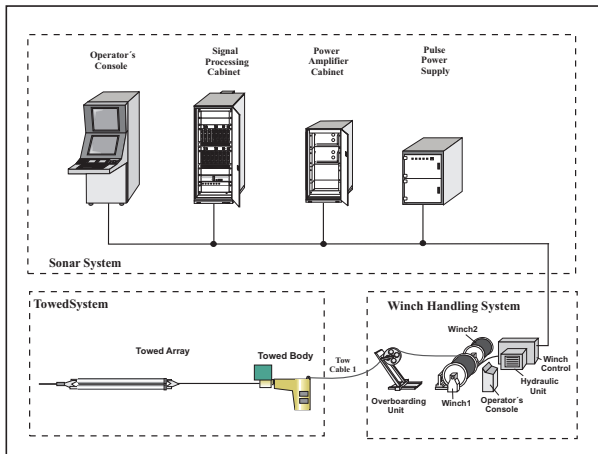
Figure 1: ACTAS system components: Wetend, Winch and Handling System, and Sonar System.

ing speeds. Further information on twin-line array design can be found in [10], [5].

Active sonar systems are operating in a sensitive marine environment. Designers and users of active sonar systems are aware of the potential threat to the marine environment. They take account of the influence of artificial underwater sound on e.g. mammals like whales, dolphins and porpoises. The sonar system addressed in this contribution follows rigorous mammal protection procedures in order to exclude any mammal injury due to active sonar operation. Further information on this topic can be found in the internet, c.f. [7].

2.3 Sonar System Hardware

The main components of the sonar system are the console, the signal processing cabinet, the pulse power supply, and the power amplifier. The last two units are necessary to generate the appropriate transmission pulse with the corresponding transmission power. The components described in this paper are mainly located in the signal processing cabinet. Two VME racks are placed in this cabinet together with central components for power distribution.

3. SIGNAL PROCESSING REQUIREMENTS

The main signal processing functions are: beamforming, and active and passive signal processing.

Auxiliary functions are given by: audio processing, raw data recording, compass processing, display processing, and passive and active target tracking.

The beamformer is realized using the FFT algorithm, cf. eg. [6]. Since a twin-array is used for ACTAS, the active frequency band is processed with maximum left/right discrimination. The passive output of the beamformer is optimized for low SNR signal detection.

In this system, three pulse types are typically transmitted: two FM signals and one CW signal. This allows for Doppler evaluation for both, frequency modulated (FM, cf. [3]) and continuous wave (CW) signals. In active signal processing, pre-processing (correlation, filtering, subsampling) together with post-processing (alarm search, connectivity analysis) is implemented. Pulse lengths of up to 24 s are used. Further-

more, the Doppler is calculated by joint analysis of the two FM signals. In passive signal processing, broadband and line detection (BDT, LDT) as well as target analysis (spectrum, low frequency analysis, detection of envelope modulation on noise) are realized, cf. e.g. [1].

4. COMPONENTS

4.1 Signal Processing Platform

Single board computers (SBC) are used for signal processing. These are 6U VME-bus (versa module eurocard) CPU boards with Intel Pentium processors. The processor speed varies in general between 500 MHz and 2.0 GHz depending on the type and the manufacturer, for a description of examples see [9], [8]. The boards are equipped with two gigabit ethernet ports and a 64-bit/66 MHz PCI (peripheral component interconnect) bus. A PMC (PCI on mezzanine cards) interface can be put on the board for further extensions. VME transfer is possible e.g. between a SBC and a digital signal processor (DSP). Most of the boards have front panel I/O with keyboard, mouse, and monitor. This is very convenient during system integration, debugging, and service, because the developers can monitor the behavior of their programs on-line in the system.

The ruggedised versions of the boards have an increased temperature range, and shock and vibration immunity.

A signal processing cluster in the signal processing cabinet (Figure 1) is equipped with a 20-slot VME backplane. Therefore, a cluster can contain up to 20 SBCs. No on-board ventilation is necessary, the cooling is just done by air ventilation through the cluster.

In the system described in this paper, Windows NT embedded and Windows XP embedded is running on the boards.

4.2 Network

A digital array concept is used and signals are sampled and digitized directly at the receiving hydrophone. They are transferred via optical cables to the electronics cabinet. There data are received by a digital signal processor and transferred to a first SBC. Automatic gain control is running on the digital signal processor and sensor blanking is implemented at this stage. The complete remaining signal processing is set-up by SBCs and gigabit ethernet is used to distribute data in the system.

The basic network structure is presented in Figure 2. In this figure, four rows of processors (B to E) are shown. Each row consists of a server and up to n processing boards (Proc) and can be imagined to constitute a signal processing function (e.g. beamforming). The letters (A to F) at the SBCs indicate the corresponding ethernet subnets. The BootServer got a connection to all subnets. There are SBC which have two gigabit ethernet ports in use. First, the servers need two ports to improve the data throughput. Secondly, certain processing boards possess connections to the console subnet (F) to be able to send processing results to the console. The main motivation for using the subnet structure is the improved data throughput of the servers by using two ethernet ports. A prerequisite for this is a 64-bit PCI bus on the SBCs to allow for a sufficient transfer bandwidth. Furthermore, if a broadcast is needed e.g. in one subnet, this will not lead to a performance reduction on the boards in other subnets.
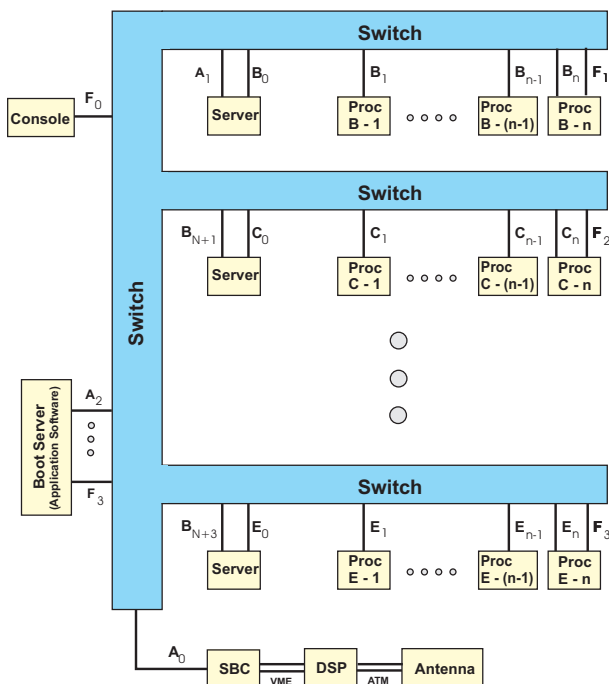
Figure 2: Simplified network structure: The switch can be set-up physically by several stacked switches. The Servers and Procs are all single board computer. The BootServer is an industrial PC system.

Function	R_0	L	N_p	
Beamforming	4.0	1.56	$\lceil 6.6 \rceil$	$= 7$
FM Correlation (1 channel)	4.0	0.25	$\lceil 1.4 \rceil$	$= 2$
CW Spectrum	0.5	0.125	$\lceil 0.1 \rceil$	$= 1$

Table 1: Processing throughput and number of boards of various algorithms ($R = 20$ Mbyte/s). Performance data recorded with a 1.2 GHz Pentium 4 processor using the Intel Integrated Performance Primitives Library, see [2]. Units for R_0 and L are Mbyte/s and s/Mbyte, respectively.

5. DATA TRANSFER

The hydrophone data rate is about $R_0 = 4$ Mbyte/s. The data rates at the passive and active signal processing channel input are determined by the number of beams, the corresponding processing bandwidths, and the data formats used. For the passive and active channel one obtains about 6 Mbyte/s and 4 Mbyte/s, respectively. Although these data rates are quite moderate, a fast ethernet connection is needed for data transfer. Nominal data rates of ethernet connections give only a first hint of system performance. More important are the transfer times for single data packages.

For illustration purposes consider the following example. A $K = 1$ M byte data package takes 50 ms to be transferred from one computer to the other at a data rate of $R = 20$ Mbyte/s. After receiving this data package the SBC can perform the corresponding signal processing operation and finally pushes the processing result to another SBC. Let us assume that this data package is of the same size as the input data package, then it takes again 50 ms to push the data to a server. After this, the SBC fetches the next data package and so on. At an input data rate of $R_0 = 4$ Mbyte/s a data package of 1 M byte has to be processed every $T_{max} = K/R_0 = 250$ ms to avoid an overflow and therefore data loss. This means that the processing SBC gets (250-2 * 50) ms = 150 ms for signal processing of the corresponding data package. The signal processing availability a_V of the processor, defined as the ratio of available processing time to input time for a certain data package, is therefore given by $a_V = 1 - 2R_0/R$. For the parameters $R_0 = 4$ Mbyte/s and $R = 20$ Mbyte/s one obtains $a_V = 0.6$, i.e. the processor is available for signal processing tasks only for 60% of the time although the transfer data rate is four times the input data rate.

If the algorithm takes longer than the available processing time on one SBC, then more SBCs have to be used to perform the job. This can be accomplished e.g. by a time slice processing scheme. One processing board pulls the first data piece, a second the next one and so on. Let us assume that N_p boards are used in this way for a processing task to be completed and that all boards are of the same type. After the board N_p has got the corresponding data package, the first board must be finished with its processing and be waiting to obtain its next data piece. Otherwise, there will be data loss. Let L be the time it takes to process one byte (s/byte), then it takes KL seconds to process a data package of size K and the number of boards necessary to perform the algorithm is given by $N_p = \lceil R_0(L + 2/R) \rceil$, where $\lceil x \rceil$ indicates the smallest integer number greater than x. Table 1 gives some experimental results. Worst case scenarios concerning the parameters like e.g. the FFT length are used. The beamformer calculates in total 576 beams for 256 input channels. Of course, processing functions which do not occupy a complete board can run conjoindly on a single board, so that the average board load is almost identical for all SBCs.

6. SOFTWARE

6.1 Client/Server Concept

The main software communication element used is the client-server-concept. Client and server are both software programs running on the same or different computers and are communicating with each other. The server and client are the passive and active component, respectively. Of course, in this context the words active and passive are used to describe the program behavior and are not connected to sonar signal processing. This means that a client program addresses a server program for some operation.

In this contribution we use this concept, where the server is a data storage component and the client is a processing node. In the sonar context, the server stores e.g. hydrophone data (StaveServer), a client like a beamformer is performing the beamforming operation, and finally sends beam data to a beam data server (BeamServer). The StaveServer and the BeamServer are therefore passive components waiting for some action to go on and the beamformer is the active component asking for and sending processed data.

6.2 CORBA

The Common Object Request Broker Architecture (CORBA) is a world wide standard for communication in distributed software systems. It is developed and published by the OMG (Object Management Group, cf. [4]) with a large amount of

industry support. CORBA is mainly designed for distributed applications running in heterogeneous systems. Nevertheless, even in homogeneous systems like the one discussed in this paper it has advantages to use CORBA because it allows to develop distributed application in a non-distributed way. This saves a lot of development effort in programming the interprocess communication.

In our application we use the gigabit ethernet for data transfer. CORBA is used to reduce the development costs in programming the data transfer between SBCs. This communication layer (middleware) can just be bought which saves a lot of debugging. CORBA is well-proven even in the defense industry.

Nevertheless, even if CORBA reduces the inter-process communication programming drastically, this does not mean that the system integration is not tedious and time consuming.

6.3 Server Implementation

Figure 3 shows the basic server concept. Various clients are pushing data to and pulling data from a server. For each client a separate thread (Windows NT/XP) is started, the server has to manage these threads in a thread-safe way.

The servers have to manage various non-acoustic data. One example is the start of a transmission period which is important in active signal processing in order to calculate the exact range of a target. The servers are designed to fulfill this requirement.
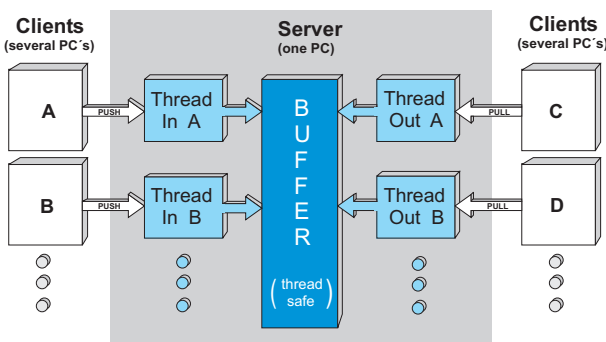

Figure 3: Data server concept: Several SBCs are pushing data to and pulling data from a server. The arrows in the figure indicate that the clients are always the active part in the communication between client and server although the main data flow may be in the opposite direction, cf. section 6.1.

6.4 BITE and Loader

A Built-In Test Equipment (BITE) is incorporated in the system. It consists of both, hardware and software items. Each major unit, like e.g. the console, the signal processing cabinets, or each function module (e.g. the beamformer module) is equipped with a controller card. These cards are interconnected via a controller area network bus (CAN). All system states and hardware and software errors are collected and transfered via CAN bus to a central point where a test control program is running.

BITE indicates the system status to the operator and gives error messages in case of broken components like e.g. mal-

function processing boards or missing ethernet connections. Error messages of applications are collected by BITE and provided to the operator for recording and system management. On-line and off-line tests are available for system testing. Furthermore, BITE controls the switching on and off of the system components like e.g. the electronic cabinets.

Furthermore, BITE distributes a function identification number among the SBC during the system booting phase and defines which application runs on a certain board. All SBCs got the same flashed operating system and the application software is downloaded from the BootServer.

Because no input device like keyboard or monitor is connected to the SBCs, it is necessary to have a program for SBC control. Therefore, the functionalities supplied by the so-called Loader which is running on the BootServer are: remote access to a selected SBC, re-start and re-boot of selected boards, and shutdown of application software during system shutdown.

The BITE features described above enable a neat way to implement redundancies. If e.g. a signal processing function has to be switched from one set of boards to another, this is done by changing a correspondence table (application to board number) before restarting the affected boards. Furthermore, spare boards are running in the system which can be activated in case of board malfunction.

7. SUMMARY

This contribution describes the use of COTS components in a sonar signal processing system. Over several years, we got good experiences with this approach. Usually, system set-up and software development can be done in a shorter time than using dedicated hardware or DSPs. Nevertheless, system integration is still a challenging and time consuming task.

REFERENCES

[1] W.S. Burdic. *Underwater Acoustic System Analysis*. Prentice-Hall, Englewood Cliffs, New Jersey, 1991.

[2] Intel Corporation. http://www.intel.com.

[3] F. Ehlers and D. Tielburger. Verfahren zum Bestimmen von Zieldaten mit einer Aktivsonaranlage, Patent DE 103 32 886, 2003; Method for Determining Target Data using an Active Sonar, European Patent 1 500 953, 2006.

[4] Object Management Group. http://www.omg.org.

[5] H. Höstermann and H. Schmidt-Schierhorn. The design of the ACTAS twin-array. In *Proc. Underwater Defense Technology*, 1999.

[6] B. Maranda. Efficient digital beamforming in the frequency domain. *J. Acoust. Soc. Am.*, 86(5):1813–1819, 1989.

[7] NATO Undersea Research Centre (NURC). http://solmar.nurc.nato.int.

[8] Concurrent Technologies Plc. http://www.cct.co.uk.

[9] Inc. SBS Technologies. http://www.sbs.com.

[10] H. Schmidt-Schierhorn and T. Warhonowicz. Towed twin line array: Design and results of sea trials. In *Proc. Underwater Defense Technology Pacific*, 1998.