THE GENERIC ACOUSTIC PROCESSING SYSTEM (GAPS).

Nick Felgate

Naval Systems Department, Defence Science & Technology Laboratory, Farnborough, Hampshire, GU14 0LX , United Kingdom. NJFelgate@dstl.gov.uk

1. INTRODUCTION

Despite the end of the Cold War, submarines still remain a potent weapon both in terms of deployment by the UK and in UK defence against those submarines of other nations. In both cases, the main sensor for detection and localisation remains passive sonar i.e. exploiting the sound from the machinery and propulsion systems of the vessels themselves.

To make best use of the sonar systems deployed at-sea, it is useful to maintain shore-based facilities for the replay and analysis of recorded sonar data. This analysis has two aims. Firstly, to determine how well both the sonar system and the sonar crew performed and to make suggestions for improvement. Secondly, to undertake an in-depth analysis of the underwater acoustic signature of particular vessels to identify those features useful for detection and classification in other encounters.

Often shore-based facilities have just been at-sea systems re-engineered to take data from recorded media and a different one is needed for each sonar type employed at-sea. However, over the last three years the Naval Systems Department of the Defence Science and Technology Laboratory (part of the Ministry of Defence) in partnership with its sub-contractor Detica Ltd have been developing a generic system providing the replay and analysis facilities for a number of sonar systems. This system provides the analysis facilities at a fraction of the cost and risk than previous systems, provides additional functionality over that required at sea and provides better connectivity with the Department’s other facilities.

2. OVERVIEW

The major elements of GAPS are illustrated in Figure 1. They are:

- Stripping – extraction of the acoustic and supporting non-acoustic data from the replay media and conversion to a generic form capable of being processed.
- Beamforming – combination of signals from a number of hydrophone elements located in an array to provide a directional beam, both improving the signal-to-noise ratio of the acoustic signatures against the ambient background noise, and providing bearing-time information used in localisation.
- Spectral/Power Processing – estimation of the power in the underwater signature as a function of frequency or bearing against time using, principally, Fourier techniques.
- Display Processing – compression of the dynamic range of the spectral/power data, enhancing the significant acoustic features, in order to fit the dynamic range of the grey-scale of the graphical displays.
- Man-Machine Interface – display of the grey-scale and other acoustic and non-acoustic information, along with tools to help with its interpretation and to allow export in various forms to other applications.

Proceedings of the Institute of Acoustics

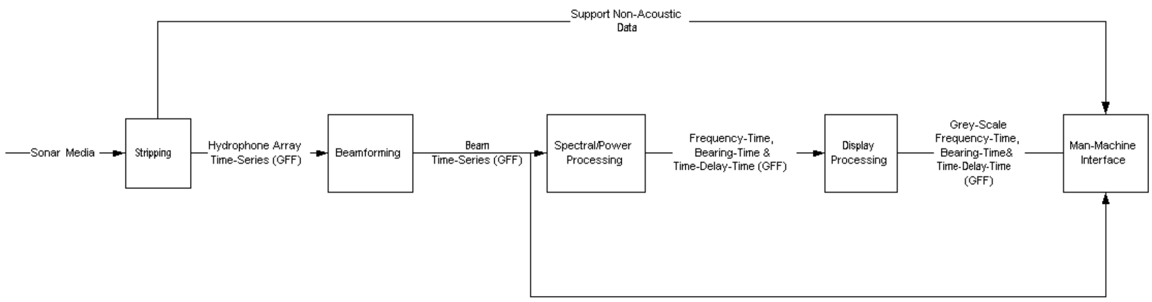

Figure 1 – GAPS Major Elements

In the development of GAPS, much effort has been expended in identifying the best algorithms from open-literature and classified sources, and a number of improvements have been developed locally. However, the most innovative aspect of GAPS is in the manner in which the system has been developed to be generic and expandable to the replay of a number of sonars.

The main mechanisms employed in GAPS providing this flexibility are:

- A modular processing chain.
- A distributed processing system employing commercial-off-the-shelf (COTS) hardware.
- A generic file format capable of storing data from various sensors at various levels of abstraction.

3. THE GENERIC FILE FORMAT (GFF)

3.1 History of the Generic File Format

The development of the Generic File Format (GFF) started prior to the development of GAPS. The format started life purely as a format for the storage of time-series with the aim of establishing one format for the data interchange between different sonar analysis and replay systems within the Department.

Many existing time-series formats were examined. Unfortunately most appeared to have restrictions on the sampling rates that could be represented (usually only integer, and sometimes only a pre-defined set), offered little support for multiple channels (many only stereo, and very few with support for channels with different sample rates) and offered no support for bandshifting. Moreover, few provided potential for expansion. It was therefore decided it would be necessary to define a new format employing the best elements from the formats examined, adapting them to be more flexible and adding a number of additional features to ensure data consistency.

With the advent of GAPS, it was realised that most sonar data is simply a time-sequence of more complex data items and that the initial time-series format could be adapted with little modification to store such data items. Support for other data types was provided by the definition of generic storage structures for these types within the time-series frames and by definition of fields to store details of the different processing involved at each stage. In this manner support was expanded to the various forms of time-series data (array element, beamformed, omni-directional), spectral data (complex spectrum, power spectrum, DEMON spectrum), broadband data (power and cross-correlation bearing-time history), non-acoustic support data (own-ship, tracker, sonar status) and grey-scale data (lofargram, directional-colour-coded lofargram, bearing-time history).

The Generic File Format is an open standard [1] with the UK's Ministry of Defence only reserving some control over modifications to the base specification.

Proceedings of the Institute of Acoustics

3.2 The Role of the Format Within GAPS

The Generic File Format provides the "glue" between all the processing elements within the system as shown in Figure 1. All data from the recording media is initially stripped into the format and all intermediate products between the main processing elements are stored in this format as well. The only information about the source data used by a processing element is that which it can extract from the header information of the GFF.

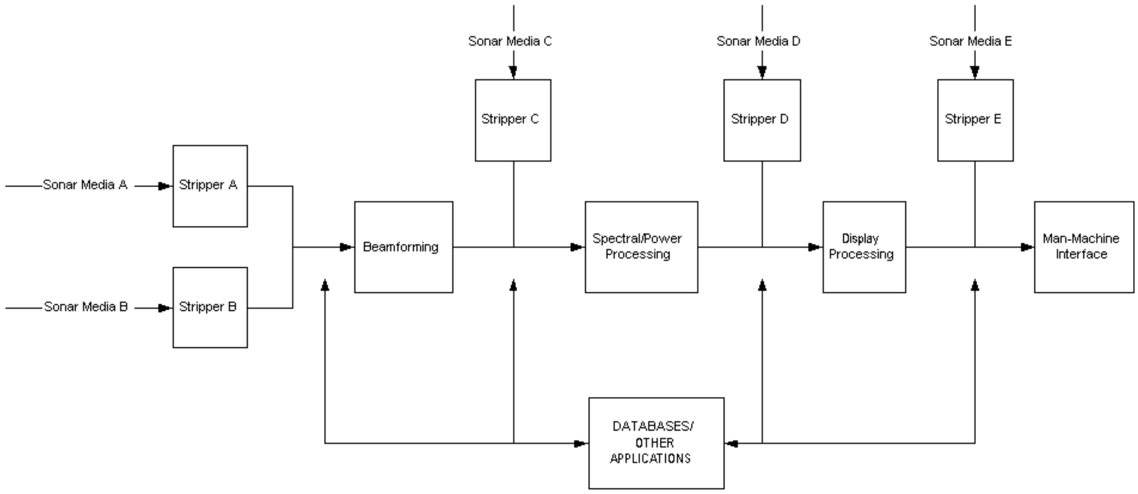

Figure 2 – GAPS Stripping & Interfacing

As illustrated by Stripper A and B in Figure 2, this allows support to be given to another sonar or recording media that provides nominally the same information simply by the development of a new stripper. Also, as illustrated by Stripper's B, C, D and E in Figure 2, it also allows support to be given to another sonar or recording media providing information at different levels of processing. Unless additional functionality is required then all other elements of processing need not be changed. In addition, as illustrated in Figure 2, it is also possible to exchange data between GAPS and other systems at different levels of processing. This also includes storage to or retrieval from a database system.

There are multiple benefits to this approach. Firstly, because of the reuse of previous elements of the system, there is lower risk, lower cost and smaller lead times in the expansion of the system to other sonars or media. There is lower overall maintenance costs as much of the software and hardware is shared between sonars. A better tool set can be developed through cost sharing and retrospective application of new algorithms to old sonar data. There is greater flexibility in the deployment of analysts through use of a common Man-Machine Interface for all sonars and therefore little overhead in retraining for new sonars.

The ability to store intermediate results means that it is possible to apply different processing without having to restrip from the original media. This means greater flexibility in allowing the user to re-examine the data using different processing parameters. It also means that most new algorithms can be tested, tuned and proved using recorded data prior to implementation in any at-sea system.

Proceedings of the Institute of Acoustics

4. GAPS COTS TECHNOLOGY

4.1 General Architecture

Figure 3 shows the hardware architecture on which the GAPS application currently runs within the Department. However the system being purely software is re-configurable to different hardware architectures and can actually be run on a single PC.

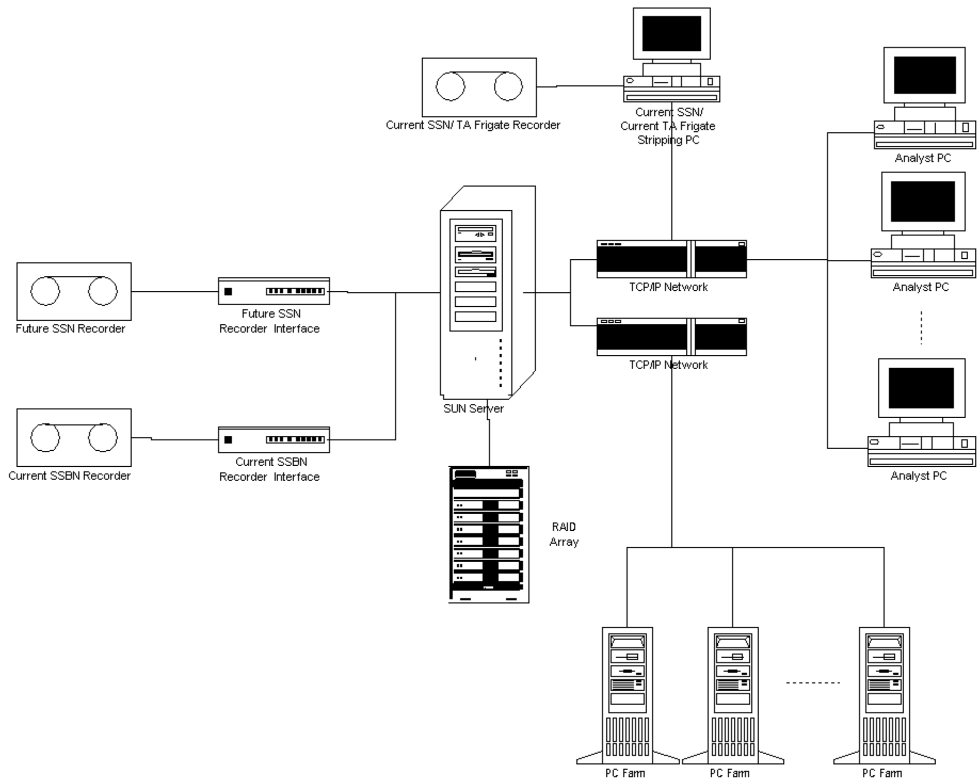

Figure 3 – Current GAPS Hardware Architecture

It is noted that for processing and display no bespoke hardware such as DSP cards is employed and the system relies entirely on commercial-off-the-shelf (COTS) equipment. Some previous work indicated that modern COTS processors could be just as fast as specialised DSP processors for spectral processing applications. Some later work has indicated that there may be some cost-performance benefit with regard to employing specialised Field Programmable Gate Arrays (FPGA) for high-throughput beamforming applications. However, development lead times, development complexity and risk have dictated that, at least in the short term, the system will still be implemented purely on general COTS hardware.

In addition to using the processors in the processor farm for distributed processing, the system can use the considerable spare capacity of the individual PCs on analyst desktops. This represents a significant saving in cost, both in the direct costs of hardware purchase and in accommodation costs. It does however have repercussions in terms of system design, because the processing distribution mechanism must be able to cope with processors of varying performance and with processors that may be loaded with other tasks. However, under normal conditions, with a user running low intensive tasks such as office applications, the extra work undertaken on the local machine for the GAPS system has largely gone unnoticed.

Proceedings of the Institute of Acoustics

The only bespoke hardware in the whole system is some interface boxes to replay recorders in the stripping system for a number of sonar. While there were COTS systems to fulfil this requirement they were expensive and with no alternative suppliers. In the manufacture of these boxes every effort has been made to make it supportable and adaptable.

4.2 Benefits of COTS Technology

The benefits of the use of COTS technology are numerous. The initial hardware costs tend to be lower due to the greater competition and there are alternative suppliers should the original manufacturer go out of business. There is usually a greater pool of developers and more advanced development tools, leading to lower costs in system development and maintenance costs. And invariably with the larger customer base the upgrade path is easier, sometimes with improvements in speed from new processors or networking capability available without even recompilation of source code. Developing to a standard PC operating system means that the GAPS application can be delivered to the analysts on the standard office desktop PC. Re-use of this display hardware provides cost-savings both in the direct cost of the display hardware and in accommodation costs. It also allows GAPS to be run side-by-side with standard office and other applications, leading to increased productivity.

5. DISTRIBUTED PROCESSING

5.1 Processing Tasks

The requirement for the GAPS system is that it should be possible to reanalyse data from any sonar at faster than real-time, desirably by at least a factor of eight. In order to achieve the processing throughput required to achieve these analysis rates using COTS equipment it is necessary to employ distributed processing in which a processing task is undertaken by a number of processors.

The main computing tasks to be undertaken by distributed processing are relatively few:

- Beamforming – use of delays (both integer sample delay and inter-sample interpolation) and weighted summation of time-series from different sensors to form time-series in which the signals from a particular direction have an enhanced signal-to-noise ratio.
- Spectral Processing – use of the Discrete Fourier Transform to derive the power of the underwater signature as a function of time and frequency, aimed principally at extracting the tonal features of the signature generated by a contact's internal machinery.
- DEMON Processing – extraction of the envelope of a frequency band followed by use of the Discrete Fourier Transform to derive the power of amplitude modulation of a carrier signal as a function of time and frequency, aimed principally at extracting the modulation of the noise of the contact's propeller cavitation.
- Broadband Processing – extraction of the power in a contact's underwater signature in a frequency band as a function of bearing and time for localisation purposes.
- Correlation Processing – use of the Discrete Fourier Transform and Inverse Discrete Transform to obtain the time-delay between different acoustic paths (between different sensors, or multiple paths to the same sensor) as a function of time for localisation purposes.
- Display Processing – spectral windowing, background noise estimation, normalisation and grey-level bin quantisation to derive lofargrams, demongrams, bearing-time-histories and correlogram displays. It is noted that in GAPS display processing can be performed using the distributed processing for fast screening purposes using standard settings, or in the analyst's local machine for deeper analysis with user-adjusted settings.

Proceedings of the Institute of Acoustics

These are relatively simple, single-input, single-output tasks with no inter-communication between tasks. The high processor load comes from the amount of data on which the analysis is to be performed in parallel (multiple hydrophones and beams, relatively high sample rates, long time periods).

5.2 Time-Slicing

The computing task could be divided in several ways for processing over the distributed network. For instance, certain processors could be allocated to certain processing tasks, or processors could be allocated to tasks for certain beams. However, taking a fixed approach to this allocation does not provide robustness with regard to processor availability; the distribution mechanism would need to be complex to cope with differences in processor speeds and the possibility that one processor may go down and its task need reallocating. Moreover, if there are greater computing resources available than directly employed then some resources might be redundant in a particular task leading to gross inefficiency in the system.

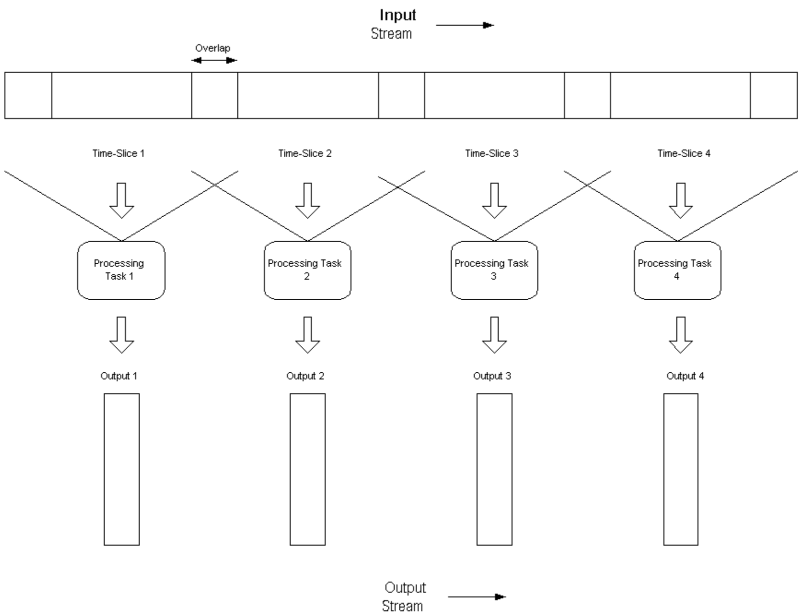

Figure 4 - Time-Slicing

To overcome these problems, GAPS employs time-slicing to further divide the data for distributed processing as illustrated in Figure 4. The data stream for a channel (element for beamforming, beam for spectral processing) is divided into time-chunks and the processing for these chunks allocated by distribution mechanism to an available processor. The distribution mechanism receives back the data from a processor and reconstitutes the output in a logical time-order (due to differing loads on each processor, or even different processor speeds, the tasks will not necessarily be received in the order in which they are sent out). This approach does lead to greater network traffic as often the same data is retransmitted to different tasks (e.g. overlap between consecutive Fourier Transforms in spectral analysis, filter initialisation in beamforming). It also leads to complexity in the processing needed to reconstruct the final output in proper time order. However, it does provide great robustness with regard processor availability as tasks may be retransmitted if a processing engine signals a failure or if the task results are not received within a specified time.

Proceedings of the Institute of Acoustics

For spectral, DEMON and correlation processing, the natural size of each time-slice is the size of the Discrete Fourier Transform employed as this simplifies the mechanism for the reconstruction of the output. However, in beamforming, there is no natural size for the time-slices and the size is very much based on experimental optimisation of the overall performance in terms of total processing speed and network traffic.

5.3 GAPS Distribution Mechanism

Given the relative simplicity of the processing tasks and the control needed, the distribution mechanism was implemented directly in GAPS itself using simple queuing mechanisms. While details of the implementation has changed slightly over the years as a result of experience and the need to expand the system, the mechanics of this system are basically the same as originally devised as illustrated in Figure 5.

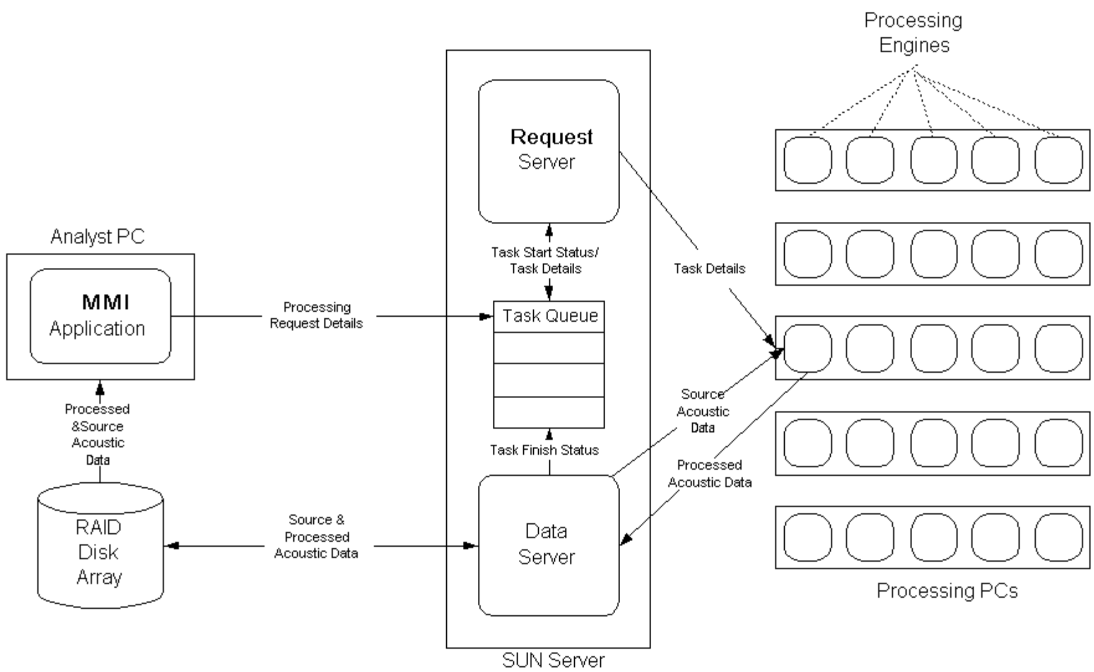

Figure 5 – GAPS Task Distribution Mechanism

The main GAPS application runs on the local PC. It handles the display of processed data and aural replay, and the generation of processing requests. The distribution mechanism has two server applications hosted on the SUN server. The request server is responsible for receiving requests from the Man-Machine Interface for processing, for allocating tasks to the individual processing engines and for monitoring a task's progress. The data server handles requests for data from the individual processing engines and handles the output from the processing task and reconstituting the results into the proper time order. The activities of the main application, request server and data server with regard to task allocation and processing are co-ordinated by means of the processing task queue.

Each processing engine is identical and is capable of carrying out all of the different processing tasks. It is noted that multiple processing engines are hosted on each individual processing platform. This effectively provides local load balancing within the memory and processing architecture of the individual platform and improves overall throughput. At any one time a particular engine may be waiting for data to be delivered, transferring information in, processing,

Proceedings of the Institute of Acoustics

waiting to transfer of output or transferring results out with different loads on the processor, memory and input/output elements of the host platform in each state. On average, the different engines should be in different states at any one time without significant clashes for resources. All communication between the processing engines and the servers takes place via SSL authenticated and encrypted TCP/IP sockets.

When a processing request is generated by the GAPS main application it is passed to the request server which places its details in a queue. When a processing platform is first switched on the processing engines are started up. These engines then constantly poll the request server for un-actioned tasks. If a task is available on the queue, the request server passes the task details to the processing engine and marks it in the queue as started. The processing engine then repeatedly requests the required data from the data server and when it is received from the RAID array it processes the request to completion. On completion, the engine passes the processed data back to the data server. The data server removes the task from the queue and stores the processed data back in the RAID array in the proper order. Depending on the processing request, a further process may be then applied or the data may be retrieved by the GAPS main application for display.

The current queuing mechanism is implemented by means of a number of binary files in various directories. A simple round-robin mechanism is used to process tasks in the same priority queue, while priority is obtained by maintaining separate directories for each queue and processing the queues in order. The detailed processing instructions and processing parameters for the tasks are held in simple text files.

6. SUMMARY

The Generic Acoustic Processing System is a system that provides replay and analysis facilities for a number of current sonars, and has the potential for providing such facilities for new sonars at low cost and risk. This has been possible through development of a distributed, modular processing system based on COTS hardware and development of a Generic File Format for the storage of acoustic data at different levels of processing. This, and the access via the GFF to real, at-sea recorded data at different levels of processing, makes it an ideal test-bed for the development of future sonar algorithms and techniques. There is scope for reconfiguring the system for other sonar uses including real-time and standalone systems. It is noted that the current system already achieves much faster than real-time processing. However, for real-time use, the distribution mechanism would need to be reconfigured to avoid latency in the delivery of a particular task. The system could also be adapted for other applications (e.g. engine vibration monitoring, aircraft noise parameterisation). There is also scope for adapting its distributed processing technology and Generic File Format for other applications.

REFERENCES

[1] Felgate N.J., “Standard for the Generic File Format Version 4.00”, Defence Science and Technology Laboratory, Technical Report TR01328 Issue 2.00, September 2003.