

## Reconfigurable Sensor Processing using Heterogeneous Architecture

D J Shand, C S Petrie, D J Houlston, C N Robinson  
Nallatech Ltd.  
Boolean House  
One Napier Park  
Cumbernauld  
Glasgow G68 0BH

### Abstract

*The use of complex sensors, and the processing of the data they produce, is a key element of successful defence systems. Whilst it has always been a challenging area of work, it can be argued that recently a number of factors have combined to raise the stakes in this area. The increasing commercial use of complex sensor systems has put pressure on the defence industry to align with these efforts and hence benefit from the apparent cost and time savings. At the same time, the requirements for defence systems have become much more fluid and less well defined. These, combined with the relentless pressure to reduce whole life costs of systems, mean that new ways are being sought to use all available technologies in a flexible way for the creation, use and maintenance of high performance, low cost, complex sensor systems. This paper discusses the contribution of hardware architectures to this problem space.*

*Keywords: Heterogeneous, FPGA, Sensor Systems, Processing, COTS, Obsolescence*

### Introduction

Complex sensor systems have always been insatiable consumers of data processing power. The clear aim in systems is to sense faster, more accurately, over a wider range of conditions and more reliably. All of these drive the user to incorporating an ever-increasing variety of increasingly complex sensors in greater and greater numbers. This causes an exponential rise in the amount of data being sensed. If any benefit is to be derived from this extra data, it must be processed. This leads to an exponential rise in processing requirements.

Traditionally, therefore, these systems must use a variety of techniques & technologies to meet the requirement. The practical outcome of this is that sensor based systems are often one-off custom products tailored specifically to the task for which they have

been designed. This is a technically successful approach in that a solution is produced. However, it suffers from a significant weakness. A custom system inevitably requires a custom knowledge base to understand and support it. This knowledge must be at all levels from the lowest understanding of what the system comprises, through how it works up to why it was designed the way it was. This knowledge exists in the original design and implementation teams, but as time progresses this is lost. As systems evolve and change this knowledge is difficult to maintain. This leads to them becoming effectively unmodifiable and unmaintainable. Whilst this weakness has always existed, the modern rapidity of change of both systems and the components underlying those systems turns it from a costly challenge to a fatal flaw.

Fortunately, the same rapid advances in technology that have precipitated this problem have also produced possible solutions. Complex sensor systems are no longer the sole province of the defence industry. Many are now developed specifically for the commercial sector. These are often developed for significantly lower cost and against more imprecise requirements than equivalent defence systems. It is reasonable to look to these areas for methods and technologies to help with the specific problems of the defence sector. In addition, technology itself has become more flexible with the advent of huge complex programmable logic devices in the form of Field Programmable Gate Arrays. These are heavily used in the commercial sector to assist in management of change and reducing design costs. They are increasingly being seen as a key part of any solution to reducing whole life cost in defence systems.

The task has been to look at these advances, and any others that can be predicted, to find how they can be used to define a base hardware architecture allowing engineers to design and maintain defence sensor processing systems more effectively.

### **Technology Usage – Present & Future**

The first task undertaken was a comprehensive audit of a number of factors directly influencing the direction and composition of sensor systems. This was with the intention of providing a number of benchmarks against which technology and techniques could be tested.

This was done in three stages. The first was a general user conference from which an overall view was garnered. This was followed up with a more detailed questionnaire sent to a wide variety of

sensor systems designers and users. The conclusions drawn from this were then tested in detail with a smaller number of systems designers.

The results were a series of opinions relating to both the future direction of technology used in processing systems and the challenges they are likely to have to overcome.

### **The Processing Task**

The first task was to look at the challenges now and in the future for sensor processing systems.

Sensor systems are very often mission critical (with a subset being safety critical). This has clear implications in the areas of reliability and fault tolerance requirements. There is no indication this will be less true in the future. In fact there are indications that in the future there will be more dependence on sensor based systems e.g. increased use of battlefield unmanned vehicles.

As may have been predicted, the pressure to use the latest advances in sensor technology will continue, leading to the inevitable increase in the amount of data processed. This will be coupled with equally continual changes in the types of data processing required, in order to get most effective use of new sensors or combinations of sensors.

One major “new” trend is the increasing use of multiple sensors. This may be multiple identical sensors (e.g. phased array systems) or multiple different sensors (e.g. multi-spectral seekers). Current design methodologies and more traditional sequential processing techniques are perceived to be less well suited to processing data from such systems.

Overall, processing is seen as the main bottleneck in developing more effective sensor based systems. Even at current

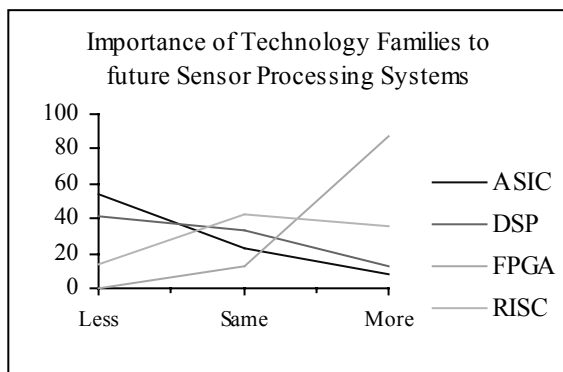
sensor and processing technology levels, upwards of 30% of systems down sample the raw sensor data because they cannot process it all. If the data cannot be processed, there seems far less incentive to develop and use more complex sensor based systems in the future.

### Processing Technology – The Options

Having established the problem space, we moved onto look at the variety of potential solutions (from a hardware perspective).

Currently, sensor processing is performed by a multitude of different technologies ranging from the generally available to the specifically designed. The range of choices for data processing is largely covered by digital signal processors, high performance processors, programmable logic (FPGAs) and application specific integrated circuits (ASICs).

We asked a variety of designers how important they believed these technologies would be for future sensor processing systems. The results were unexpectedly clear.



Clearly the industry view is that systems will increasingly be composed of FPGAs and RISC (or high end) processors. ASICs and DSPs are generally felt to be reducing in importance. Further investigation showed that the predicted decline in use of DSPs was the increasing ability of FPGAs to

perform the same function at reduced cost and greater efficiency. The reduction in ASIC usage was purely a cost decision based on the increasing cost of ASIC development as feature sizes reduce.

Regardless of the processing technology, it was clear from all responses that interconnectivity was the key to using the processor power. This is particularly true as multiple sensors come into systems, where it is essential to providing readily scaleable systems.

Currently VME is far and away the most popular “standard” form factor in use (within the defence industry), but its interconnectivity is considered woefully inadequate for most sensor processing systems. As a result, a variety of non-standard interconnection methods are used to supplement the basic level. These range from mezzanine modules and specialist cards to non-standard use of the VME backplane.

### Standards – Their Use and Abuse

The use of VME as a “standard” architecture that is used in a non-standard way by pretty much everyone is typical and important when talking about processing platforms of the future. When asked, most engineers profess a desire to use standards now and particularly in the future. However, when questioned more deeply, they have grave misgivings about the constricting nature of these standards given the extreme performance they are often required to deliver. They are also concerned about the long term maintenance of the standards and whether they (e.g. defence industry) will have sufficient influence. This leads to a tendency to use them as a base upon which custom solutions can be built, rather than a long-term structure.

In general, standards only work when people need them to work. This means at

points where two separate systems meet. In the commercial sector, where the total customer solution may involve many different companies and regulatory bodies there are many such interfaces, hence there are many standards. It is in the interests of all the small players to meet and maintain these standards so their product can become part of the whole solution, even if this requires extra effort on their part. In the defence industry, the whole solution is usually delivered by a very small number of companies (perhaps only one) so there are correspondingly fewer external interfaces. Therefore, there are fewer standards. The driver is delivery of the whole solution to requirement, and hence only the standards necessary to that objective are observed.

The use, or not, of standards is the output of the constant assessment of cost in terms of extra work, reduced functionality, increased complexity etc. vs. benefit in terms of interoperability, common understanding & resources etc. It may be that the future will see this balance tipping more towards the benefit side, with pressure to reduce overall lifecycle costs and be more flexible to changing requirements being the deciding factors.

### **Heterogeneous Architecture Concept**

The task, therefore, is to take the wide range of views and predictions of the stakeholders and turn these into proposals for what an effective Heterogeneous Architecture for sensor processing should be like.

It was immediately clear that proposing any “standard” in rigid terms was doomed to failure for a number of reasons

- Industry dislike of actually using standards
- Lack of appropriate means for maintenance
- Probably outdated before it got

published

- Too inflexible for the intense requirements of sensor systems.

Instead, it seemed more positive to propose a series of properties of a good hardware platform against which any existing or new platform could be assessed. This will have the effect of promoting any features that are particularly useful, perhaps even to the status of de-facto standards, and preventing some of the more avoidable errors.

Clearly a vital requirement for any hardware platform is flexibility. This must be thought of in a number of ways.

- Support for any type of processing system (ASIC, FPGA, RISC, DSP and others)
- Support for many types of form factor (e.g. PCI, cPCI, PC-104+, VME, Custom) and ruggedisation level from laboratory to battlefield.
- Able to adapt to external (e.g. I/O) standards changes.
- Configurable interconnection to allow scaleable solutions.

No less important is reliability and, hence, testability. Any hardware platform must be testable for full functionality both on the bench and in system.

The current, and likely future, balance of power in the electronics industry is heavily in favour of commercial developments. Therefore, it is vital that any proposed architecture is sympathetic to the direction of commercial development to maximise use of the massive economies of scale within the commercial marketplace.

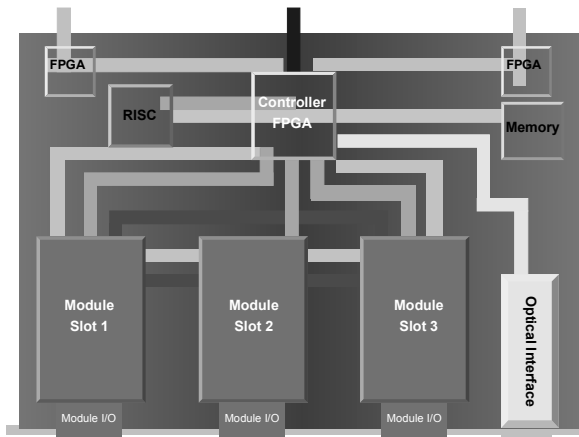
One way of meeting these requirements is a modular system based around high speed serial interconnections. If this is implemented using FPGA serial communications systems, this will produce a flexible system, capable of reconfiguration to suit the many requirements of sensor processing systems.

Modular architectures have a number of


other advantages. They can be incrementally redesigned and verified, allowing progressive enhancement or modification of a system. This means systems are in service longer, reducing lifecycle costs. The modularity also allows custom components to be included as custom modules without compromising the generality of the whole system. Modular systems also allow the luxury of late design partitioning based on actual performance, rather than early partitioning based on predictions.

### Specific Details

We have proposed the following generic motherboard/module combination. This shows a flexible ring/mesh topology to avoid unnecessary limitations in interconnectivity.



This will be populated with one of a range of application specific modules containing FPGAs, DSPs, ASICs, RISC processors or custom technology



The key to this architecture is the use of high speed serial connections between elements. This gives high data rates without the need to adhere to complex standards. Serial connections make inter-system connections much simpler as well as allowing easier conversion between physical form factors (compared to parallel or bus communications structures). Serial connections are also inherently more testable and reliable than equivalent parallel systems.

Generally, the trends within the wider commercial market are in line with this proposal. The two major FPGA manufacturers (Xilinx and Altera) are both delivering high speed serial interconnectivity between their devices. Using the latest product families in this architecture would give 100 Gigabits/sec communications paths between modules.

In addition, many RISC processors are also beginning to support various high speed serial standards (Hypertransport on the AMD Opteron being one example).

Even devices that do not directly support serial communications (Sharcs, Altivec etc.), can be interfaced into the fabric via a relatively inexpensive FPGA, which has the major additional benefit of allowing reconfiguration if changes are required to the topology.

## **Tools and Support**

Successful design and maintenance requires tool support to allow these hardware architectures to be used in a maintainable way. Generally, the design of systems is beyond the scope of this work. It is appropriate, however, to look at how tools could map designs onto such architectures. This requires the implementation of a standard inter-process communications structure at a system level (e.g. multi-device). This is, generally, poorly provided in the general market as it is specific to each hardware implementation and relatively low level. As part of this work we have looked at how this could be done in a transferable and verifiable way across a number of platforms.

## **Future Work**

The next phase of the work is to implement a reasonably complex system using the principles defined in this paper. This will be used in the implementation of a complex, representative sensor processing problem in a mixed technology architecture.

Specifically, this will require further work on the inter-process communications model, including integration of this into a higher level systems design tool.

The challenges associated with doing this will be used to assess and improve the architecture. The aim is to generate a proven and benchmarked reconfigurable heterogeneous sensor processing platform.

## **Acknowledgements**

Nallatech wishes to acknowledge the assistance of the other members of the EMRS DTC.