The Tempest Processing Engine

Power and Precision

Solid State Logic

SOUND || VISION

Dr Enrique Perez Gonzalez, Chief Technology Officer

*Some industry voices have suggested that CPU technology cannot on its own deliver the stability and power required for large scale audio processing. The **Tempest** engine which employs SSL's patented **Optimal Core Processing** proves that untrue. This document explains SSL's approach and why it works.*

In 2011 when contemplating the future of its digital product range, a group of SSL's R&D engineers spent a number of months evaluating the requirements of a platform on which they could base a new range of systems for a variety of professional applications at different scales. The platform was to be adaptable, have industrial strength and reliability, provide massive processing power and be capable of unmatched software and architectural flexibility.

SSL C-Series consoles (C10, C100, C200 and C300) are based on SSL's Centuri processors and the more recent Blackrock processors. Both employ arrays of Sharc DSP and FPGA chips and use an in-house operating system designed specifically for the purpose. The advantages of this approach, similar to those of other manufacturers of high-end digital audio consoles, are the deterministic behaviour of the DSP array and absolute control over the OS. On the downside there is the complexity and cost of scaling up the DSP resource, and for the software, the inability to leverage standard industry resources.

During the recent decade, as the technology of individual processors got close to some fundamental physical constraints and Moore's Law ran out of steam, the power of IT industry processing hardware continued to grow largely through the development of multi-core processors. The overall throughput of these new multi-core processors exceeded what was sensibly achievable with discrete DSP chips but as the IT industry is not by and large as concerned with real-time, deterministic performance as professional audio equipment has to be, the challenge becomes whether it is possible to manage multi-core processors to deliver the required performance and, if so, how much of the maximum theoretically available power can be utilised. We perceived that we could develop software much more efficiently using widely adopted programming languages and focused on which Real Time Operating Systems (RTOS) could best help us in our efforts to manage the processing resource.

Two notable approaches to using multicore processors compete: the first uses a "supervisor" program which enables the cores to run non-real-time operating systems in some cores while independently other cores run RTOS. The "supervisor" must give priority to the real-time cores over the non-real-time cores and regulate the priorities of low level interrupts and memory exchanges. The second option is to use two separate multi-core processors, one dedicated to processing audio the other to interface with the controls and displays. Using two separate multicore processors ensures that there is no interaction between non-real-time cores and real-time cores and secures robust real-time performance for very large processing systems. After lengthy experimentation and evaluation of both systems, we chose the latter due to its deterministic behaviour and reliability, considerations which are paramount in large scale professional audio applications.

A key constraint of multi-core processors is the way they access the available memory. Each core has a small amount of dedicated local (L-level) cache memory but the main external memory is shared. There is a fundamental trade-off between cache latency and hit rate given that larger caches have better hit rates but longer latency. For real-time audio signal processing, greater levels of performance can be achieved by managing the L-level memory tightly. Digital consoles need to allocate processing resource for a number of different tasks that occur simultaneously: filters, dynamics, delays, mixing, routing, clocking, metering, displays and control interactions. Generally it makes best use of the available processing resource to allocate these functions to specified cores.

After lengthy research we developed some technology, which we patented, to work with the RTOS to guarantee deterministic, real-time performance even with heavy memory exchange loads. This is called Optimal Core Processing and it optimises the ratio of the resource allocated to processing with that allocated to memory exchange. It enables us to do multiple 64-bit floating point operations that in turn enable us to execute complex, high resolution algorithms at very low latency. The code has also been engineered to deliver many thousands of multiply-add-and-accumulate instructions and doesn't therefore need to rely on additional FPGA bus-summing or any extensive external routing hardware.

The multi-core processors we use are industrial grade units designed to operate at a wider range of temperatures than 'ordinary' PC hardware. The Tempest processor uses a RTOS with our OCP software technology. The Audio Server, the name we give the controller, runs on a separate multi-core processor running an embedded operating system. This combination confers a number of benefits. The system is resilient; if anything happens to the Audio Server, the DSP continues and so does the audio. DSP boot time is fast. The Audio Server can connect to more than one control surface enabling multiple front panels to control the same engine. Standard IT industry topologies can be used to provide full mirrored redundancy. And in a programming language understood by many, with industrial strength libraries of useful code widely available, we have been able to accelerate our developments, accommodate changes and special requests more easily and increase the size of our programming team to deliver some exciting products.

The first of these were the **L500** and **L300** Live consoles which have proved hugely popular for their flexibility, their ease of use, their reliability and, of course, their sound quality. The Tempest engine is now being engineered for **System T**, our latest Broadcast Audio Production System. With System T the core technology is that used in the Live consoles but, in order to provide the high channel and bus counts required in today's major TV productions, has a version of the Tempest engine that delivers 800 paths and can route over 3,000 signals in and out. With dedicated broadcast specific architecture and algorithms, the software for System T also differs from that of the Live consoles and provides the features that large scale broadcast productions require with the robustness and security that Broadcasters demand.

Other Documents In This Series

- Professional AoIP for Broadcast - The Way Forward

More Information: The SSL Team bring decades of broadcast console design experience to the challenges of leveraging networked technology for broadcast. Email sales@soliddtsatelogic.com for help and advice.