ALLEN & HEATH XCVI

FROM DSP TO FPGA: A SIGNAL PROCESSING REVOLUTION

Getting to the core of the issue

A bit of 2010 web surfing would have undoubtedly unveiled the nearly ubiquitous use of DSP (Digital Signal Processor) chips as the processing backbone for the major digital mixing console players on the market. These DSP-based platforms typically use a number of DSP chips (aka a 'DSP farm') to process, mix and manage the required number of channels. In 2010's digital console world, DSP was just able to squeak by and meet the needs of most engineers, but as networked systems and processing and sample rate expectations grew in the following decade, DSP began to reach its limits.

For example, the 2010-era 48kHz Allen & Heath iLive used 11 dual-core DSPs. While 11 dual-core chips sounds impressive, this technology wasn't scalable enough to achieve the channel, mix and effects count and enhanced compressors and dynamic equalization that bleeding edge applications now demand. Adding in the sample rate increase from 48kHz to 96kHz, we would have required an eightfold power factor increase on those 11 dual-core DSP chips. The processor math didn't look good.

iLive 'Core II' DSP farm

The DSP farm approach also has drawbacks when it comes to latency and phase coherency, since audio is sent back and forth across several chips for processing and routing. DSPs are also inherently limited in the router size and capability, in the set of instructions audio algorithms can use, and in bit-depth, which is predetermined by the DSP family used. In short, DSP was seeming more and more like a Motorola Razr flip phone in an iPhone world. DSP technology didn't need a small evolution in power, it needed a complete revolution and a different paradigm entirely.

It's fun to stay at the FPGA

The necessary paradigmatic processing leap came in the form of the FPGA. Field Programmable Gate Arrays are a different sort of microprocessor altogether and bring about much more power per chip when used effectively. FPGAs are frequently used in computationally intensive applications such as AI accelerators, aerospace and defence avionics, high performance servers and super computers. To illustrate this jump in chip performance, the 11 iLive dual-core DSP chips were replaced with just one exponentially more powerful FPGA in the current generation Allen & Heath dLive.

> *"The I/O on the actual DSP chip was a non-starter for what we wanted to achieve, and that's not even taking into account going 96kHz: you've obviously got a fixed number of cycles in your DSP so you need twice as many of them, and lots of consequential problems. There wasn't really a processor on the market that could do everything we wanted to do. So we ended up moving into a custom FPGA processor."*

Jeff Simpson, R&D Engineer

ALLEN&HEATH

ALLEN & HEATH LIMITED, KERNICK INDUSTRIAL ESTATE,
PENRYN, CORNWALL, TR10 9LU, UNITED KINGDOM.
www.allen-heath.com

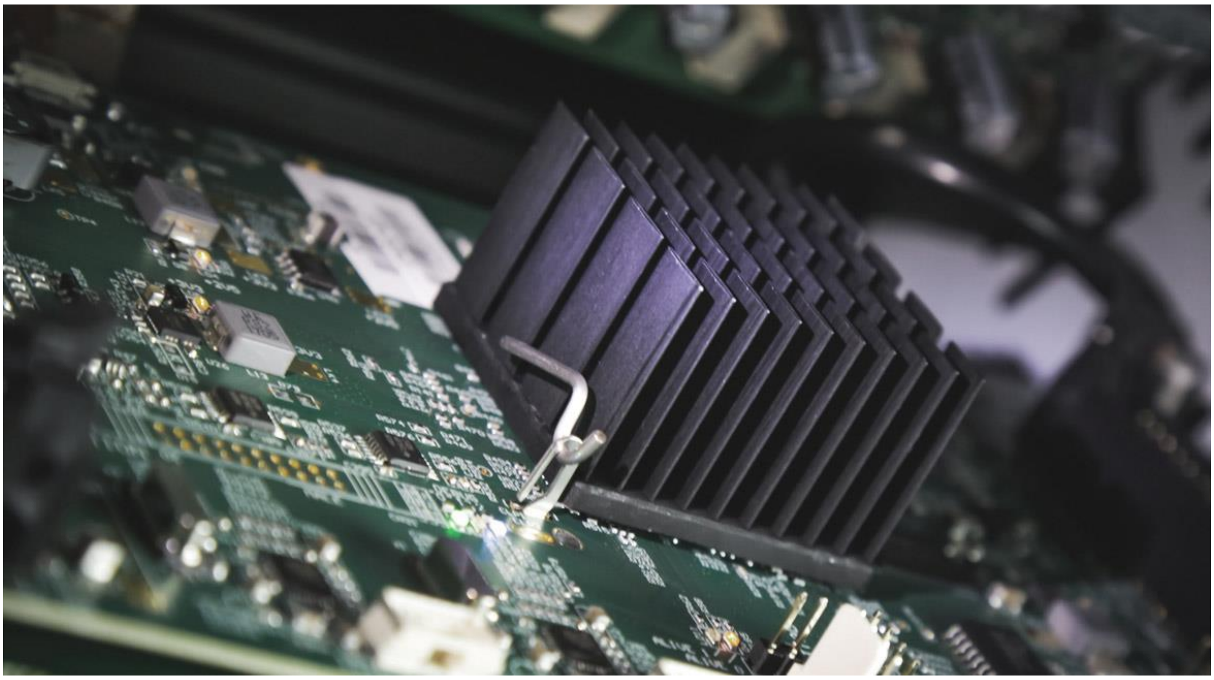

XCVI core in SQ mixers

This bespoke FPGA processor brings about immediate audio processing benefits: more mixing power, enabling a 96kHz sampling rate, low latency (0.7ms) and easier management of phase coherency. The chip runs parallel internal virtual processing cores for channel processing and routing. These kernels basically appear as miniature DSPs, but with extra power and flexibility. Dedicated virtual cores are also provided for specialized advanced dynamic (Dyn8) engines within the chip. Since FPGA technology doesn't set any bit-depth limitation, Allen & Heath engineers can pick the desired bit-depth for each processing block or algorithm ('variable bit depth'), resulting in superior precision and noise performance.

So why doesn't everyone use FPGA for processing? Because it's not a simple transition at all. FPGA programming is very 'low level' code. DSP programming is based on a set of instructions available for a certain DSP family which is already designed for audio signal processing, so you have a sort of running start at generating the chip instructions. But recreating these instructions in FPGA land (or creating new ones) is far from trivial. Dozens of man years of work went into crafting a custom set of FPGA instructions to aid efficient modelling, and a compiler tool to translate such instructions into low-level FPGA code for the Allen & Heath **XCVI** core. The good news was that the R&D team could port all of their existing algorithms over—and the even better news was that the algorithms could be further refined without the architectural limitations DSP had previously imposed. We now have the flexibility of a set of instructions we can add to or edit at will, rather than being constrained by what a DSP family can do.

To the best of our knowledge, Allen & Heath is the only company with this technological approach.

The technology is also scalable, as proven by the incredibly popular SQ series and now Avantis.

What's in it for me?

Technology for technology's sake is meaningless. Luckily, there are some solid benefits to our FPGA implementation.

Routing galore - On top of running channel processing and DEEP processing extensions, the XCVI core has parallel mixing engines (six in dLive) which calculate several thousand cross points per sample, while the FPGA router has capacity for thousands of audio paths. This enables fully flexible patching and routing, from compressor / gate sidechains and parallel compression on all channels, to advanced routing options such as Input to Matrix or Bus to Input, to Tielines for direct system patching of inputs to outputs, regardless of sample rate and format (analogue, digital, MADI, Dante, etc.) — this is over 800 system inputs and 800 system outputs on dLive, all freely patchable.

Latency - During training sessions we get asked a lot why we're so obsessed with latency. After all, a couple of milliseconds never hurt anyone, right? The problem with latency is that it's cumulative and in the modern era of digital processing, latency is everywhere. Wireless mics add latency, Dante transport adds latency, loudspeaker processors add latency, and so do IEM systems. So what starts as a negligible figure can quickly become unacceptable in a typical live audio device chain. Our latest generation mixers provide sub-millisecond latency, input to output, full stop.

Noise performance - To get the sonic performance right and keep noise at bay, it's very important to look at bit depth in all parts of the processing path. Processing introduces errors, which are manifested as THD+N. Fortunately this audio distortion can be controlled with appropriate design choices since resolution (precision) is essentially proportional to bit depth. We studied bit depth as part of the XCVI research and we ended up with much wider bit depth where it matters. A good example of this variable bit depth is the bus summing which employs 96-bit.

Remember all those debates on fixed point vs. floating point? Some DSPs offer floating point arithmetic to circumvent the issue of limited bid depths and overflow, however that comes at the expense of noise modulation. With variable bit depth, the debate becomes irrelevant. The ultimate solution is fixed point arithmetic, with more bits where you need them. Not quite solving world peace, but it's a start.

Phase coherency - In the analogue world you could take any route, any path inside your mixer, and your outputs would always be phase aligned. Try that on a digital desk, and the results can be less than optimal. Comb filtering comes into play, which is why we see inexperienced engineers at a lot of gigs populating their digital racks with long plugin chains, or getting creative with parallel compression, only to get a smeared mix which is totally out of focus and sounds like a phase shifter was added to the mains. There are only a couple

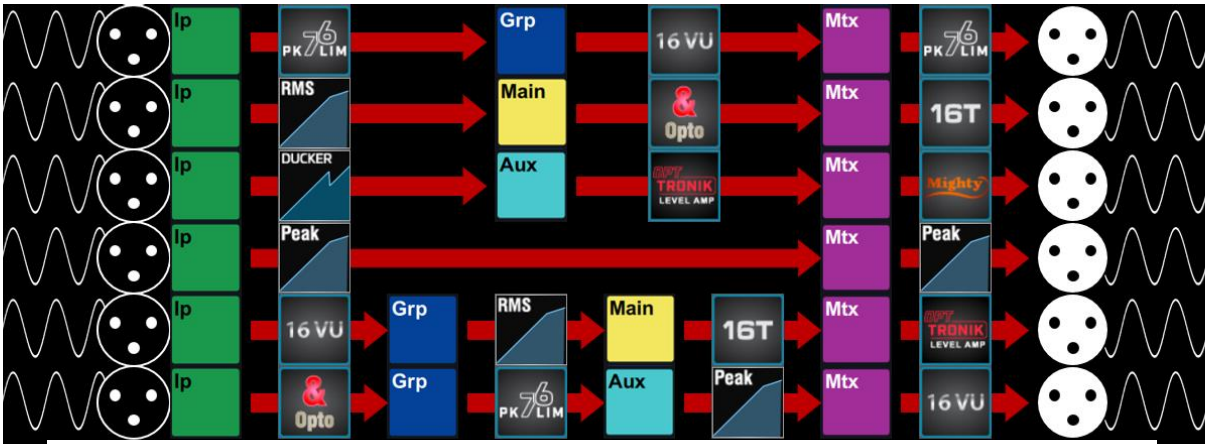

Mix coherency in dLive regardless of internal path or channel processing phase coherent mixing systems out there, and normally they come at the expense of a much higher system latency. With dLive, SQ and Avantis, no matter which path you take internally, input to mix, input to group to mix, input to matrix etc., all mix outputs of the same type are phase aligned. This approach contributes to their open and accurate sound.

A final word

In short, XCVI is a huge step forward from traditional DSP tech. We could have continued to 'duct tape' more and more DSP chips into our desks, but instead we chose to innovate and take the bolder leap, make the necessary financial investments, and place our bet on the wonderful world of FPGA.

Based on the feedback heard from fans across the world, the gamble seems to have paid off.

ALLEN & HEATH LIMITED, KERNICK INDUSTRIAL ESTATE,
PENRYN, CORNWALL, TR10 9LU, UNITED KINGDOM.
www.allen-heath.com