

DISTRIBUTED VS. CENTRALIZED AUDIO PROCESSING

Peak Audio founders say the distributed DSP model is obsolete.

By Don Kreski

There's a debate going on today that some audio professionals thought was settled years ago. Is it better to centralize digital signal processing in a large audio system or distribute it over many processors linked via the network?

Most sound companies today argue for distributed processing, but Rich Zwiebel and John Britton, the founders of engineering firm Peak Audio, which originated the concept of configurable audio processing back in 1992, say it's an obsolete approach.

The biggest argument typically made for distributed processing is that it's more reliable than centralized. If one component fails, the remainder of the system can continue to function.

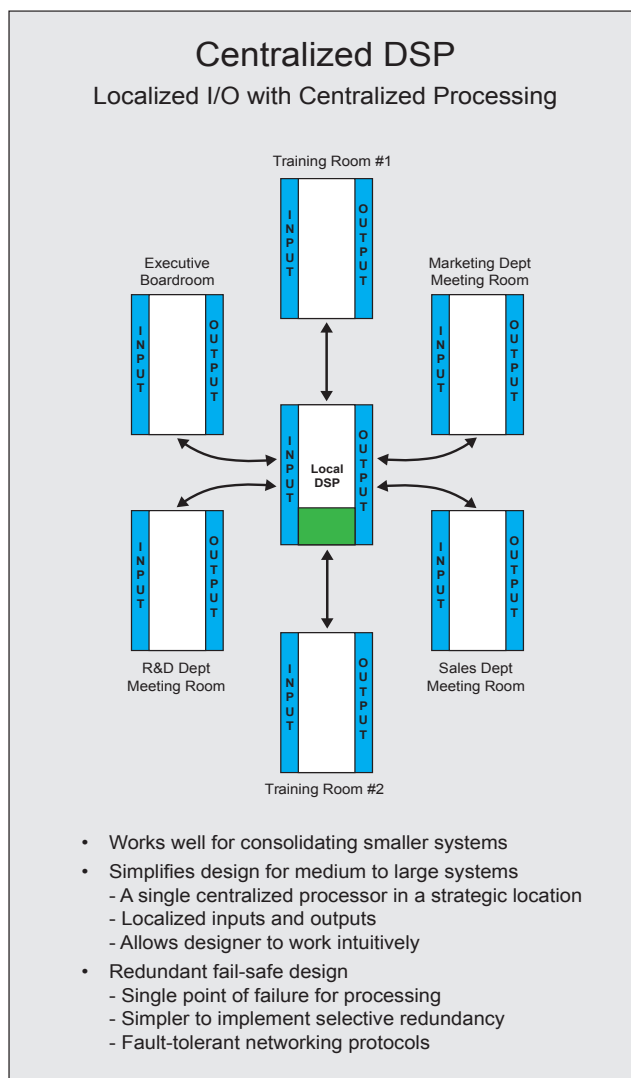
Not so, says Zwiebel. "Our engineering team has done it both ways, and we know that distributed is far more problematic, because it is in so many boxes. That is, you go from a single point of failure to many points of failure, any one of which can end the event." Yet with today's technology, you can safely rely on a single processor, because you can back it up with a redundant DSP that takes over automatically should your primary unit develop a glitch.

THE CONCEPTS BEHIND DISTRIBUTED DSP

Britton and Zwiebel's engineering team originally began working on DSP more than 20 years ago, creating the hardware and software designs for Peavey's MediaMatrix product line. Introduced in 1992, MediaMatrix was a centralized processing platform that could use a flexible number of DSP cards in a multi-slot cardframe.

"Our original idea was to be able to support smaller installations but then get the horsepower up by adding processing cards as needed for larger systems," Britton explains. The multi-card design proved to be very popular.

After Britton and Zwiebel's team developed CobraNet in the mid-90s, they began thinking about putting their DSP chips into



separate boxes and linking them via the network. Peak Audio finished engineering a new product, MediaMatrix Nion, in 1999, and it was even more scalable and flexible than the earlier designs.

Q-SYS CENTRALIZED PROCESSING EXPLAINED

Q-Sys, developed by John Britton, Rich Zwiebel, and other engineers at QSC Audio, is not the only centralized audio processing platform, but it is one of the most interesting.

The Core is the brain of the system, handling all audio processing, signal routing, and control functions. Cores come in five sizes, handling from eight to 1024 channels of audio, and Cores can be combined to create a larger system of systems. Q-Sys Cores use standard Intel multi-core processors running a hard realtime Linux-based operating system, while the signal path employs a floating-point format to preserve dynamic range. Q-Sys employs standard off-the-shelf network interfaces as well.

Q-Sys I/O frames handle audio inputs and outputs, as well as GPIO and serial control. Each frame can house up to four I/O cards, and each card can handle up to four channels of analog input or output, with higher channel counts for the digital audio cards. I/O frames are normally located near the audio sources, mixing desks, amplifiers and loudspeakers, and connected back to the Cores via a Gigabit Ethernet. With the appropriate I/O cards, Q-Sys is compatible with any analog audio device. Digital audio I/O cards provide a bridge to other systems.

Users can control a Q-Sys audio system from

a PC or Mac, an Apple iPad or iPhone, from QSC touchscreens, or from Ethernet-based control systems by Crestron and AMX.

Cores, I/O frames, and control devices are linked via a standard, Layer 3, Gigabit IP network which can be configured to simultaneously carry data, streaming video, and telephony traffic—a converged network.



There are times where distributing processing within a facility makes sense. An example might be a theme park, where each LAN is maintained by its own staff. Yet centralized sources, such as background music or emergency pages may need to be fed to the entire park. In this case, Q-sys supports a "system of systems" approach, with multiple Cores distributed on an enterprise network. This hybrid approach keeps all the benefits of centralized processing while distributing specific processing systems

throughout the facility.

Latency in a Q-Sys network is very low. The total end-to-end latency, including analog-to-digital conversion, transmission across the network to a Core, processing in the Core, a second pass across the network to the destination, and the final digital-to-analog conversion, is just 2.5 milliseconds. That's roughly equivalent to the time it takes for sound to travel to your ear from a person talking 2.5ft. away—an imperceptible delay. "In addition to this very low latency, all outputs are exactly in sync from any input," Zwiebel explains.

This mainstream, standardized approach to networking has been very well received by IT managers. Since Q-LAN is IP-compatible, IT departments can manage and monitor the Q-Sys network and system components using the same best practices as they would converged IP networks. "Design consultants and sound contractors love Q-sys as well," Zwiebel says. "For the first time, they can design and program a system of almost any size without any sub-segmentation due to processor constraints. This ability to freely design a system can be a big benefit in getting exactly what they envision. And with all the routing in one box, it's easy to switch any input to any output in the facility."

Processors could be placed throughout a facility close to input and output devices, yet still be linked for facility-wide functions such as paging. The team initially felt that system reliability might improve as well, since if any of the processors failed, the rest of the audio system could go on functioning.

In 2001, Zwiebel and Britton sold Peak Audio to Cirrus Logic and by 2006 they accepted management positions at QSC Audio. The team they organized there developed a new platform, Q-Sys, modernizing the DSP concept with all audio processing centralized in a single core. The core, based on an Intel processor, is linked to inputs and outputs via off-the-shelf gigabit network components.

CENTRALIZED PROCESSING DEBATED

Zwiebel says when he shows Q-Sys to potential buyers, they often question its divergent design. "The most common negative comment I get is that customers feel they need to put their processing in many locations," he explains. "To that I always ask, "Is it the processing or the I/O that they need to distribute?" It is almost always the I/O.

"Once they think about it, they usually admit that they do not really care where the processing lives, but to that I reply they actually do. Would they rather have all of the processing located in a secure, air-conditioned accessible space, or spread around all over the

place, say up on a catwalk with the amplifiers, or in a legislative chamber that cannot be accessed during use? They always agree that yes, they would prefer to put it into their own equipment room.

"The other negative I hear is that the distributed system is more reliable, because if they lose a box, they only lose that part of the system. But in reality, when you lose a box, you usually have major problems. What if the box you lose has the input from the announcer at a stadium? Pretty much a show stopper."

Zwiebel gives the example of a sound system that Peak Audio designed for the United States Senate back in 1999. They had five processors in the system and so, if one failed, as

it did during a committee hearing he remembers, they lost only 20% of the system. “Yet that too was a show stopper. It was not possible to say, ‘Oh, you Democrats cannot talk in this hearing, but all of you Republicans can.’” The hearing had to be cancelled.

Zwiebel explains that, in an audio system, it’s the cables and connectors that are most likely to cause problems. Using five separate boxes rather than one means there must be at least five network cables and 10 network connectors added to a crucial part of the system. In a centralized system, it’s possible to eliminate these cables by putting all of the processing on one chip.

“Another benefit of centralized processing is that it simplifies design and design file management,” he adds. When a contractor deploys multiple DSP boxes, each requires its own design file or layer and the designer has to keep track of which box has what functionality. “This introduces additional complexity and can result in painfully slow system deployments, updates and compiling times, while also making troubleshooting more difficult. The centralized approach we’ve taken with Q-Sys uses one master design file for the entire system which simplifies development, deployment, management, and updating.”

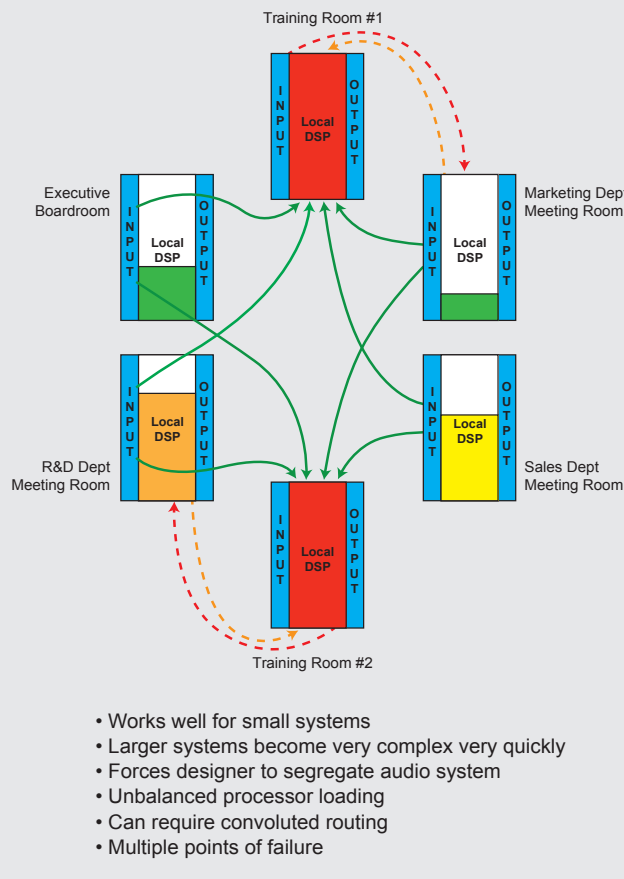
REDUNDANCY

Given the possibility of failure of any sound system, many designers have tried to build redundancy into mission-critical distributed systems, but it’s a difficult, expensive process. Britton says it is far simpler, less expensive, and more reliable to duplicate components in a centralized system.

For example, with Q-Sys, processing cores can be set up in pairs on the network. Should the system detect a fault in the primary core, the secondary takes over automatically within seconds. When an engineer adjusts settings, the changes are made simul-

Distributed DSP

Distributed I/O with Distributed Processing



aneously on both cores, so the secondary is always ready to take over. Every Q-sys product has a primary and a secondary network port. Two completely separate networks can be installed, so that if any part of either network should fail, whether it is a switch or a cable, the other network automatically and instantaneously takes over.

In addition, a centralized system can give the designer the option of adding redundant I/O frames for high-priority areas and even redundant amplifiers.

IT COMPATIBILITY AND AVB

Zwiebel says there has been a lot of discussion recently about AVB, or Audio-Visual Bridging, that many in the industry advocate as a network standard. One of the advantages

of AVB is its ability to transfer data quickly. The AVB standard requires that total latency, or delay, introduced by the network be limited to 1.024 milliseconds.

“But standard gigabit Ethernet, using off-the-shelf components, is also fast enough to limit network latency to 1 millisecond,” Zwiebel explains. “AVB is, from an IT perspective, non-standard, requiring Layer-2 networking versus the Layer-3 standard that is Internet Protocol. As a consequence, AVB requires the use of proprietary AVB-capable switches, limiting its ability to take advantage of an existing IT infrastructure. Data traveling on an AVB network cannot pass through Cisco or HP routers, which are installed in at least 80 percent of the Ethernet networks around the world.”

Zwiebel and Britton argue that audio networks today should be completely compatible with commonly used network equipment, as Q-Sys is. In addition to low-latency audio LAN distribution, they add, an audio processing platform should support long-haul streaming over wide-area networks or the Internet.

A centralized processing platform, running on an IP network, should also follow the best practices used in the computer industry. “Audio traffic should not require a separate network infrastructure,” Britton adds. “The IT department will need to use the same quality of service settings they would use for voice over IP or IPTV, to make sure audio traffic has proper priority, but those techniques are well known now.”

Is centralized processing the wave of the future? “We think this is the way to go,” says Britton. “The centralized approach we’ve taken has yielded a very powerful processing platform, yet at the same time reduced development effort, reduced system complexity, and it will give the software a very long life as ever more powerful processors and network interfaces are brought to market by the mainstream computer industry.”