COMPUTERWORLD March 12, 2001

TECHNOLOGYQUICKSTUDY

HOT TRENDS & TECHNOLOGIES IN BRIEF

Digital Signal Processor

BY TOM THOMPSON

66

VERY MODERN computer has a microprocessor in it, but not many have a digital signal processor (DSP). Since the CPU is a digital device, it clearly processes digital data, so you might wonder what the difference is between digital data and a digital signal. Basically, signal refers to communications - that is, a continuous stream of digital data that might not be stored (and thus might not be available in the future) and that must be processed in real time.

The digital signals can come from almost anywhere. For example, downloadable MP3 files store digital signals that represent music. Some camcorders digitize the video signals they generate and record them in a digital format. And the more sophisticated cordless and cellular phones typically convert your conversation into a digital signal before broadcasting it.

Variations on a Theme

A DSP differs markedly from the microprocessor that serves as the CPU in a desktop computer. A CPU's job requires it to be a generalist. It has to orchestrate the operation of diverse pieces of computer hardware, such as the hard disk drive, the graphics display and the network interface, so they work together to perform useful tasks.

This agility means that a desktop microprocessor is complex — it must support key features such as memory protection, integer arithmetic, floating-point arithmetic and vector/graphics processing.

As a result, a typical modern CPU has several hundred instructions in its repertoire to support all of these functions. This requires that it have a complex instruction-decode unit to implement the large instruction vocabulary, plus many internal logic modules (termed *execution units*) that carry out the intent of these instructions. As a result, a typical

DEFINITION

A **digital signal processor (DSP)** is a special-purpose programmable microprocessor designed to manipulate in real time a communications stream of large amounts of digital data in order to improve its quality or modify it in specific ways. DSPs are widely used for processing audio, video and graphics streams.

desktop microprocessor contains tens of millions of transistors. In contrast, a DSP is built to

be a specialist. Its sole purpose is to modify the numbers in a digital signal stream — and do it quickly. A DSP's circuits consist mainly of high-speed arithmetic and bit-manipulation hardware that can rapidly modify large amounts of data.

As a consequence, its instruction set is much smaller than that of a desktop microprocessor — perhaps no more than 80 instructions. This means that the DSP needs only a slimmed-down instructiondecode unit and fewer internal execution units. Moreover, any execution units that are present are geared toward highperformance arithmetic operations. Thus, a typical DSP consists of only several hundred thousand transistors.

As a specialist, a DSP is very good at what it does. Its myopic focus on math means that a DSP can continuously accept and modify a digital signal, such as an MP3 music recording or a cell phone conversation, without stalling or losing data. To help improve throughput, DSPs have extra internal data buses that help shuttle data among the arithmetic units and chip interfaces faster.

In addition, a DSP might use a Harvard architecture (maintaining completely physically separate memory spaces for data and instructions) so the chip's fetching and execution of program code doesn't interfere with its data processing operations.

Why Use DSPs?

A DSP's data-wrangling capabilities make it ideal for many applications. By using algorithms steeped in the mathematics of communications and linear-system theory, a DSP can take a digital signal and perform convolution operations to enhance or reduce specific characteristics of that signal.

Certain convolution algorithms enable a DSP to process an input signal so that only desired frequencies appear in the



data and address buses help boost throughput by enabling simultaneous code and data fetch operations. A secondary data bus enables the ALU to read two data values, multiply them and perform an accumulate operation in one clock cycle. processed output, implementing what's called a filter.

Here's a real-world example: Transient noise often appears as high-frequency spikes in a signal. A DSP can be programmed to apply a filter that blocks such high frequencies from the processed output. This can eliminate or minimize the effects of such noise on, say, a cell phone conversation. DSPs can apply filters not only to audio signals but to digital images as well. For example, a DSP can be used to increase the contrast of an MRI scan.

DSPs can be used to search for specific patterns of frequencies or intensities in a signal. For this reason, DSPs are often used to implement the speech-recognition engines that detect specific sequences of sounds, or phonemes. This capability can be used to implement a hands-free phone system in a car or allow your child's robotic pet dog to respond to voice commands.

Because they have far fewer transistors than a CPU, DSPs consume less power, which makes them ideal for batterypowered products. Their simplicity also makes them inexpensive to manufacture, thus they're well suited for cost-sensitive applications. The combination of low power use and low cost means that you can often find DSPs in both cell phones and that robotic pet.

At the other end of the spectrum, some DSPs contain multiple arithmetic execution units, on-chip memory and extra data buses, allowing them to perform multiprocessing. Such DSPs compress realtime video signals for transmission over the Internet and can decompress and reconstitute the video at the receiving end. These expensive, highperformance DSPs are often found in videoconferencing equipment.

Thompson is a training specialist at Metrowerks. Contact him at Thompson@ Metrowerks.com.