
DIGITAL SIGNAL PROCESSOR TRENDS

ADVANCEMENTS IN DIGITAL SIGNAL PROCESSING TECHNOLOGY ARE ENABLING ITS USE FOR INCREASINGLY WIDESPREAD APPLICATIONS. DEVELOPERS WILL BE CHALLENGED TO USE THIS PROCESSING POWER TO ITS UTMOST, WHILE CREATING NEW APPLICATIONS AND IMPROVING EXISTING ONES.

Gene Frantz
Texas Instruments

..... During the past decade digital signal processors (DSPs) have hit critical mass for high-volume applications (Figure 1). Today, the entire digital wireless industry operates with DSP-enabled handsets and base stations. The mass-storage industry depends on DSPs to produce hard-disk drives and digital versatile disc players. Ever-increasing numbers of digital subscriber line and cable modems, line cards, and other wired telecommunications equipments are based on DSPs. Digital still cameras, hearing aids, motor control, consumer audio gear such as Internet audio are just some of the many mass market applications in which DSPs are routinely found today. More specialized DSP applications include image processing, medical instrumentation, navigation, and guidance.

With the growing importance of DSPs and their applications, it seems appropriate to look at the changes occurring in these devices and to hazard a few guesses about where DSP innovations will lead in the opening decades of the new century. The continued growth of DSP-enabled applications will depend on developments in several areas of technology: the underlying manufacturing processes, the DSP core and chip architectures, and the software for development and applications. An additional factor, and the most difficult one to anticipate, is innovation. In a few years, design-

ers will be dealing with DSPs that integrate hundreds of millions of on-chip transistors and deliver performance measured in trillions of instructions per second. (See 1999 IEDM short course on system on a chip by author, available for sale at <http://shop.ieee.org/store>.) Determining how to use that processing power effectively will require imagination that goes beyond conventional engineering methodologies.

Why have DSPs done so well in the last few years? The DSP phenomenon is part of the overall microprocessor success story, and it must be seen in that light. Like the high-end reduced instruction set computing (RISC) engines used in computers and the medium-range RISC microcontrollers in embedded systems, DSPs are becoming increasingly differentiated, designed to handle the processing tasks of specific types of applications. This trend will continue with all microprocessors in the years ahead, and it will be responsible for much of the future success of DSPs.

A specialized architecture

Although DSPs are similar to RISC engines in some respects, they're fundamentally different in other ways. These differences date from the earliest microprocessor architectures, and they'll continue to influence the development of DSPs and their applications in the years ahead. Essentially, DSPs are designed for

number crunching. Early computer theorists realized that many interesting mathematical functions could be performed by a series of high-speed multiplications and additions.^{1,2} Since many of these math functions are useful for transforming and manipulating analog signals in the digital realm, a machine that would perform them efficiently would be extremely valuable as a DSP. Accordingly, certain microprocessor architects designed their processors around hardware dedicated to performing multiply-accumulate functions, and DSPs were born.^{3,4}

Initial DSP designs borrowed another idea from early computer research as well. The first microprocessors,⁵ like the computer central processing units that preceded them, employed a von Neumann architecture,⁶ with a single bus and a unified address space for both data and instructions. However, at one time a research team at Harvard, in designing Eniac, had proposed a different architecture that used separate buses and address spaces for data and instructions. DSP designers seized on the Harvard architecture, with its separate buses, but they used the idea in a novel way. In addition to adding a bus for instructions, designers provided separate buses for each multiply-accumulate operand. Thus, data and instructions could be loaded and a complete multiply-accumulate performed during every cycle. Since designers accepted the value of unified address space, they didn't split instructions from data in the main memory, though caching schemes introduced later often keep small amounts of data and code separate in on-chip memory. This modified Harvard architecture has been an integral part of DSPs ever since, even though today's architectures may include a number of functions that the original computer researchers couldn't have imagined in their wildest fantasies.⁷

Deterministic operation

Since DSPs are used for processing continuous signals that come from, and often go back into, the real world, they're constrained to operate in real time. This constraint is another key difference between DSPs and other microprocessors, not only in application, but also in the underlying architecture. Every signal-processing task operating on a DSP must be deterministic. That is, the time it requires to finish

must be determined exactly, or it runs the risk of breaking up the signal processing. Any function that can disrupt the determinism must be eliminated from the architecture, or modified so that it's not disruptive.

Interrupts are the most notable example of a disruption. Signal-processing tasks simply cannot be set aside while the processor performs system functions. High-performance RISC engines cannot manage more than a light load with digital signal processing because they're interrupt driven. Fortunately, today's DSPs offer so much performance overhead that they can handle deterministic signal-processing tasks during regularly scheduled periods, then deal with interrupts and other non-real-time tasks during the intervals between these periods.

Architectural changes

Increasingly, DSPs and other types of microprocessors have borrowed structures from each other, so that the line sometimes seems blurred where one type of processor leaves off and another begins. DSPs have become more supportive of the types of functions traditionally performed by microcontrollers and high-end RISC microprocessors. Interrupt support, which is critical to multitasking in embedded control systems, is now a regular feature of many DSPs that are meant to combine control and signal-processing functionality in a single device. Direct memory access control and various types of input/output peripherals are also routinely integrated into DSPs to provide the system-level support needed in a single- or satellite-processor application.

Two-level cache memories have been adapted from high-end RISC engines for the special requirements of DSPs. The two-level cache architecture makes a relatively small on-chip memory look like a much larger one to the core—enabling extremely fast DSPs to operate without outstripping the data available at a given time. At the same time, the cache design, coupled with the sheer speed of the DSP, provides enough configuration flexibil-




Figure 1. DSP market size (source: Forward Concepts).




Figure 2. Combining software and hardware for the lowest cost system design. Cost can be defined in terms of financing, design cost, manufacturing cost, opportunity cost, power dissipation, time to market, weight, size.

ity and performance overhead that system designers can maintain the determinism they need for critical signal-processing tasks.

Greater parallelism

The most far-reaching recent innovation, though, is the introduction of very-long-instruction-word (VLIW) architectures to DSP cores. VLIW architectures are inherently parallel, providing multiple data paths for performing multiply-accumulates and other operations simultaneously. The introduction of Texas Instruments' TMS320C6000 core in 1997, the first DSP core based on a VLIW architecture, immediately raised the performance ceiling for DSPs by an order of magnitude. Top-flight DSP performance was no longer measured in hundreds of millions of instructions per seconds (MIPS), but in thousands of MIPS. A similar jump also occurred to million multiply-accumulates per second, the critical benchmark for number crunching.

When a VLIW architecture is supported by a carefully tuned C compiler, the powerful performance of the DSP engine becomes both highly efficient and easy to use. Programmers who have little familiarity with DSPs can then write code quickly without becoming familiar with the instruction set and underlying mechanics of the processor. A two-level cache memory also enhances ease of use by eliminating the need to micromanage the movement of data on and off chip. Since DSP assembly code is often seen as intimidating by noninitiates, the availability of straightforward compilers designed to use the underlying hardware most efficiently has made DSP development much more approachable for the vast pool of C programmers.

These changes have initiated a shift in DSP system development from hardware to software, a trend that will continue as DSP performance rises to much higher levels, and software tools become easier to use and familiar to larger numbers of programmers. Developers are finding that they can get more performance out of their systems earlier in the development cycle by using high-level languages than by doggedly handcrafting every routine in assembly to squeeze the last possible drop of performance from the DSP engine. Development time is already more valuable than MIPS, and the ratio is rising (Figure 2).

VLIW architectures have been criticized for enlarging programs by adding parallel instructions, but new DSP designs incorporate features that keep down code size. These features include single-instruction, multiple-data instructions and variable-length instructions that enable multiple instructions to be packed into the same stored word. Like performance, though, memory array sizes continue to increase geometrically, so the issue of code storage space will become less critical over time even though it will always be important.

Scalable increases in performance

VLIW architectures demonstrate that it's possible to continue to increase DSP performance by adding more multiply-accumulate data paths. Essentially, VLIW parallelism builds on the two structures—multiply-accumulates and multiple buses—that distinguished DSPs from other microprocessors from the very beginning. As long as the memory subsystem is designed to keep up with the core in throughput, and as long as the compiler is sophisticated enough to handle the complexities of a massively parallel pipeline efficiently, architects can keep adding extra multiply-accumulates and supporting buses to increase performance. Although core designs are far too complex to append data paths as merely modular additions, the overall effect is similar to just snapping on more pieces. Future DSP architectures will make use of this scalability as a straightforward approach to increasing performance.

Experts like to speculate about what new structure will introduce a performance boost comparable to the one provided by multiply-accumulates and multiple buses twenty years ago. Right now, though, there's no new, alto-

gether different architecture based on a new or rediscovered logic structure that suggests itself as the source of the next processing revolution. So added parallelism, with corresponding modifications in memory and code, will continue to be the main architectural technique to increase performance for some time to come.

Process advances

All of these architectural innovations indicate that DSPs are becoming more differentiated as the technology matures and new application areas are discovered. These considerations bring us back to my earlier question: Why have DSPs done so well in recent years? One part of the answer is that somewhere in the late 1980s, IC technology began to catch up with the potential offered by DSP architectures, just as it had begun to catch up with the potential of other types of processors a few years earlier.

Some numbers are revealing here. In 1982, a 50,000-transistor DSP offered 5 MIPS for \$150 and consumed 150 milliwatts (mW) per MIPS. A decade later, a 500,000-transistor DSP capable of 40 MIPS operated on just 12.5 mW/MIPS and cost \$15 (Table 1). These numbers show that, in the 1990s, DSPs were entering the realm of price, performance, and power consumption making them appropriate for high-volume applications. At the same time, markets appeared that demanded high signal-processing performance to open up more wireless channels, speed Internet delivery, and perform other needed services. It was a classic instance of the right technology arriving at the right time for the right applications.

Obviously, these trends are continuing. Current projections by Texas Instruments are that by 2002, a 5-million-transistor DSP that provides 5,000 MIPS will be priced at just \$1.50 and will consume 0.1 mW/MIPS. Ten years later, a DSP with 50 million transistors

Table 1. Two decades of DSP market integration (typical DSP figures).

	1982	1992	2002
Die size (mm)	50	50	50
Technology size (microns)	3	0.8	0.18
MIPS	5	40	5,000
MHz	20	80	500
RAM (words)	144	1,000	16,000
ROM (words)	1,500	4,000	64,000
Price (dollars)	150	15	1.50
Power dissipation (mW/MIPS)	150	12.5	0.1
Transistors	50,000	500,000	5 million
Wafer size (inches/mm)	3 / 75	6 / 150	12 / 300




Figure 3. Power dissipation trends. The Gene's Law (named by the author) trendline follows that of Moore's Law in that DSP power dissipation per MIPS halves every 18 months.

capable of achieving 50,000 MIPS will cost just 15 cents and run on 1 nanowatt (nW) per MIPS (Figure 3). During this time, operating frequencies are predicted to zoom to more than 10 gigahertz. These figures seem incredible, even in an industry accustomed to breathtakingly rapid changes.

**Table 2. DSP integration through the years
(typical device capabilities).**

	1980	1990	2000	2010
Die size (mm)	50	50	50	5
Technology (micrometers)	3	0.8	0.1	0.02
MIPS	5	40	5,000	50,000
MHz	20	80	1,000	10,000
RAM (bytes)	256	2,000	32,000	1,000,000
Price (dollars)	150	15	5	0.15
Power (mW/MIPS)	250	12.5	0.1	0.001
Transistors	50,000	500,000	5 million	50 million
Wafer size (inches)	3	6	12	12

Process challenges

In some ways, of course, these figures are indeed incredible. Road maps don't indicate the sweat, and sometimes panic, involved in going from one technology node to the next. Today, advanced DSP cores are manufactured with 0.15-micron transistor gate widths, and core operating voltages are at 1.5 V. Soon, gate widths will reach 100 nanometers (or 0.1 microns) and core voltages of around 1 V. According to data from Texas Instruments, road map projections call for gate sizes to diminish to 20 nm in a decade and core operating voltages to 0.2 V (Table 2 and Figure 4).

It's not yet clear how gates smaller than 50 nm will be made, since unwanted electron migrations through barriers at that scale are still a problem. Similarly, there are extremely complex problems to be addressed in the multiple layers of interconnect overlying the silicon. The capacitance and inductance caused by six or seven layers of metal conducting signals at hundreds of megahertz—soon gigahertz—is a big problem. The changeover to copper from aluminum interconnects has bought a few generations of security for on-chip continuity, though even the greater density of copper will not conduct reliably indefinitely as interconnect traces become thinner and thinner. These are only a few of the manufacturing challenges facing DSP suppliers as they look at the generations of technology ahead of them. Yet the physical limit of IC technology has always appeared to be about five years, or two process generations, in front of us. Chip technologists take it on faith that physicists will solve the materials problems by then. So far, their faith has been rewarded.

DSP optimization vectors

Manufacturing processes indirectly affect us all, but they aren't at the top of the list of concerns that a system developer has in evaluating a DSP for a specific design. While there are many considerations that enter into the evaluation process, the ones that matter the most are the three Ps: price, performance, and power consumption. System developers' requirements force DSP vendors to treat the three Ps as the key vectors of device optimization. Stated a different way, at any given process node, DSP vendors tend to optimize their products for low-cost, high-processing speed, or low-power operation depending on application needs. Taking any one of these vectors to an extreme means some degree of sacrifice from each of the other two.

For example, keeping costs down usually means keeping die sizes small by minimizing functional integration, which in turn tends to slow down throughput and hobble performance. Although a smaller chip may consume less power at a given time, if it takes longer to perform operations, it may consume more power overall than a larger chip. Another way of keeping costs down is by rescaling older, slower DSPs to gain the speed advantage of smaller transistors in a leading-edge process. But since simple rescaling doesn't optimize the design to take advantage of the new process node, performance, though improved, is not maximized, as it would be with a redesigned chip.

Optimizing for performance, power, consumption

The other two vectors, performance and power consumption, are inseparably linked at the transistor level. As CMOS process nodes advance, smaller transistors require less voltage to drive them, which means less power consumption. Lower voltages also tighten the gap between high and low state thresholds, enabling faster transitions that speed up switching and raise overall logic performance. In addition, since more small transistors can be packed in the same space than large ones, there's room on the chip for extra logic functions, larger memory arrays, additional buses, and so on that serve to increase performance.

The fastest transistors must achieve the absolute minimum in transition times between the on and off states. To accomplish this, the

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.