

Lina J. Karam, Ismail AlKamal, Alan Gatherer,
Gene A. Frantz, David V. Anderson, and Brian L. Evans

Trends in Multicore DSP Platforms

Examining architectures, programming models,
software tools, emerging applications, and challenges

Multicore digital signal processors (DSPs) have gained significant importance in recent years due to the emergence of data-intensive applications, such as video and high-speed Internet browsing on mobile devices that demand increased computational performance but lower cost and power consumption. Multicore platforms allow manufacturers to produce smaller boards while simplifying board layout and routing, lowering power consumption and cost, and maintaining programmability.

Embedded processing has been dealing with multicore on a board, or in a system, for over a decade. Until recently, size limitations have kept the number of cores per chip to one, two, or four but, more recently, the shrink in feature size from new semiconductor processes has allowed single-chip DSPs to become multicore with reasonable on-chip memory and input/output (I/O), while still keeping the die within the size range required for good yield. Power and yield constraints as well as the need for large on-chip memory have further driven these multicore DSPs to become a system-on-chip (SoC). Beyond the power reduction, SoCs also lead to overall cost reduction because they simplify board design by minimizing the number of components required.

The move to multicore systems in the embedded space is as much about integration of components to reduce cost and power as it is about the development of very high-performance systems. While power limitations and the need for low-power



devices may be obvious in mobile and hand-held devices, there are stringent constraints for nonbattery powered systems as well. Cooling in such systems is generally restricted to forced air only, and there is a strong desire to avoid the mechanical liability of a fan if possible. This puts multicore devices under a serious hot spot constraint. Although a fan-cooled rack of boards may be able to dissipate hundreds of watts (an ATCA carrier card can dissipate up to 200 W), the density of parts on the board will start to suffer when any individual chip power rises above roughly 10 W. Hence, the cheapest solution at the board level is to restrict the power dissipation to around 10 W per chip and then pack these chips densely on the board.

Digital Object Identifier 10.1109/MSP.2009.934113

EMBEDDED PROCESSING HAS BEEN DEALING WITH MULTICORE ON A BOARD, OR IN A SYSTEM, FOR OVER A DECADE.

The introduction of multi-core DSP architectures presents several challenges in hardware architectures, memory organization and management, operating systems, platform software, compiler designs, and tooling for code development and debug. This article presents an overview of existing multicore DSP architectures as well as programming models, software tools, emerging applications, challenges, and future trends of multicore DSPs.

HISTORICAL PERSPECTIVES: FROM SINGLE CORE TO MULTICORE

The concept of a DSP came about in the mid-1970s. Its roots were nurtured in the soil of a growing number of university research centers creating a body of theory on how to solve real-world problems using a digital computer. This research was academic in nature and was not considered practical since it required the use of state-of-the-art computers and was not possible to do in real time.

It was a few years later that a toy by the name of Speak & Spell was created using a single integrated circuit to synthesize speech. This device made the following two bold statements:

- digital signal processing can be done in real time
- DSPs can be cost effective.

This began the era of the DSP. So, what made a DSP device different from other microprocessors? Simply put, it was the DSP's attention to doing complex math while guaranteeing real-time processing. Architectural details such as dual/multiple data buses, logic to prevent over/underflow, single cycle complex instructions, hardware multiplier, little or no capability to interrupt, and special instructions to handle signal processing constructs gave the DSP its ability to do the required complex math in real time.

"If I can't do it with one DSP, why not use two of them?" That is the answer obtained from many customers after the introduction of DSPs with enough performance to change the designer's mind set from "how do I squeeze my algorithm into this device" to "guess what, when I divide the performance that I need to do this task by the performance of a DSP, the number is small." The first encounter with this was a year or so after Texas Instruments (TI) introduced the first floating-point DSP, called the TMS320C30. It had significantly more performance than its fixed-point predecessors. TI took on the task of seeing what customers were doing with this new DSP that they weren't doing with previous ones. The significant finding was that none of the customers were using only one device in their system. They were using multiple DSPs working together to create their solutions.

As the performance of the DSPs increased, more sophisticated applications began to be handled in real time. So, it went from voice to audio to image to video processing. Figure 1 depicts this evolution. The four lines in Figure 1 represent the performance increases of DSPs in terms of instruction cycles per sample period.

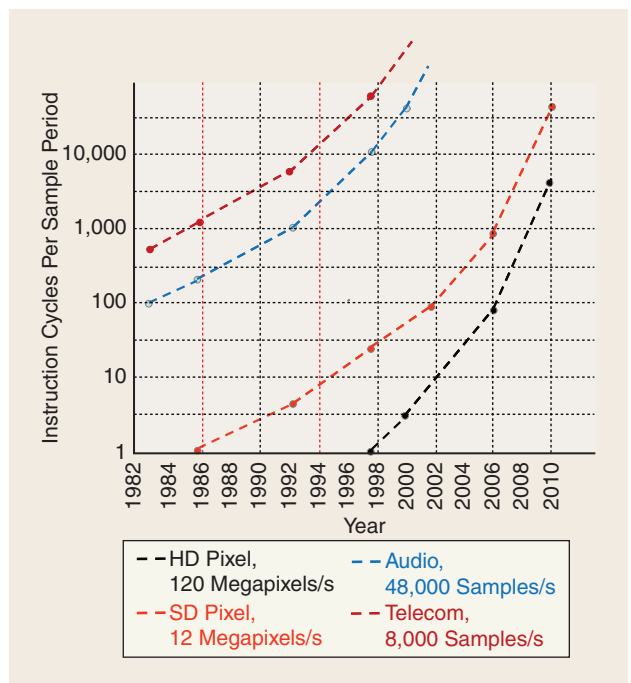
For example, the sample rate for voice is 8 kHz. Initial DSPs allowed for about 625 instructions per sample period,

barely enough for transcoding. As higher performance devices began to be available, more instruction cycles became available each sample period to do more sophisticated tasks.

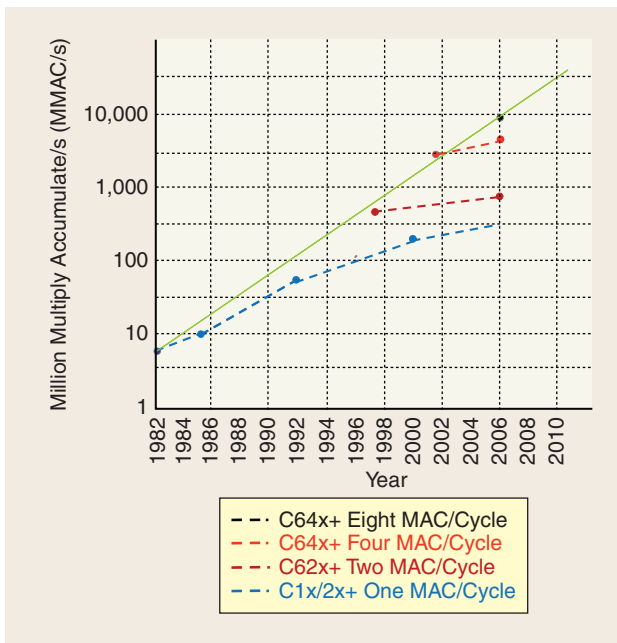
In the case of voice, algorithms such as noise cancellation, echo cancellation, and voice band modems were able to be added as a result of the increased performance made available. Figure 2 depicts how this increase in performance was more the result of multiprocessing rather than higher performance single processing elements. Because digital signal processing algorithms are multiply-accumulate (MAC) intensive, Figure 2 shows how, by adding multipliers to the architecture, the performance followed an aggressive growth rate. Adding multiplier units is the simplest form of doing multiprocessing in a DSP device.

For TI, the obvious next step was to architect the next generation DSPs with the communications ports necessary to matrix multiple DSPs together in the same system. That device was created and introduced as the TMS320C40. And, as one might suspect, a follow-up (fixed-point) device was created with multiple DSPs on one device under the management of a reduced instruction set computer (RISC) processor, the TMS320C80.

The proliferation of computationally demanding applications drove the need to integrate multiple processing elements on the same piece of silicon. This led to a whole new world of architectural options: homogeneous multiprocessing, heterogeneous



[FIG1] Four examples of the increase of instruction cycles per sample period. It appears that the DSP becomes useful when it can perform a minimum of 100 instructions per sample period. Note that for a video system the pixel is used in place of a sample.



[FIG2] Four generations of DSPs show how multiprocessing has more effect on performance than clock rate. The dotted lines correspond to the increase in performance due to clock increases within an architecture. The solid line shows the increase due to both the clock increase and the parallel processing.

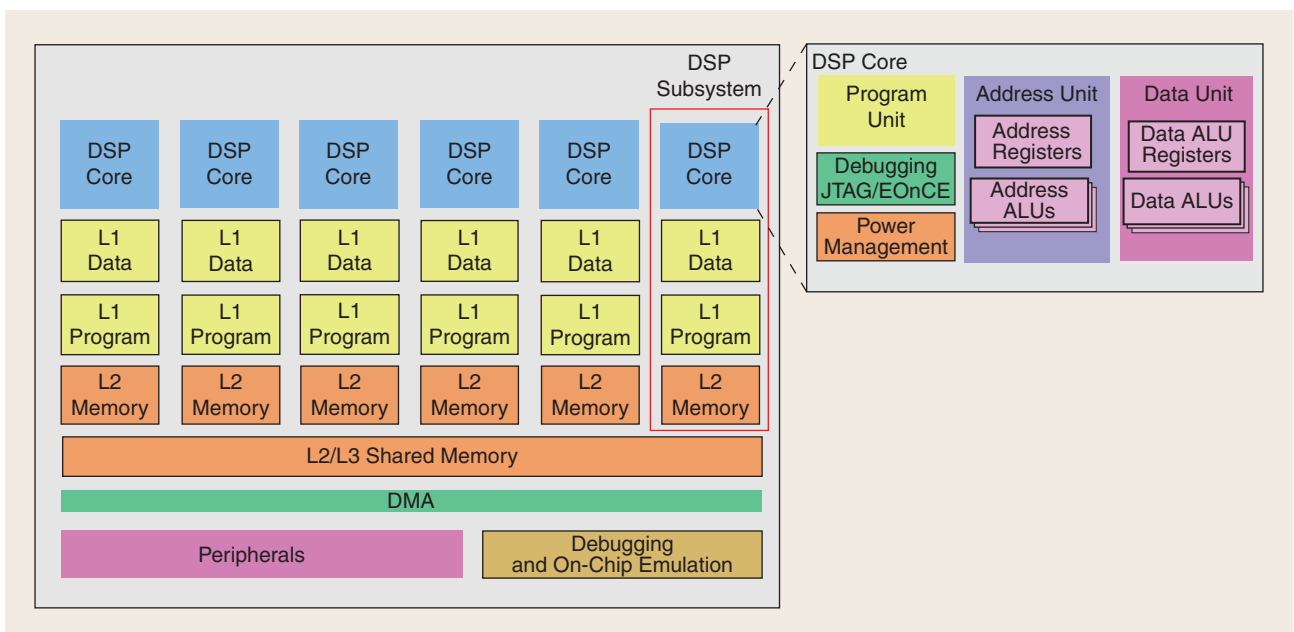
multiprocessing, processors versus accelerators, programmable versus fixed function, a mix of general-purpose processors and DSPs, or system in a package versus SoC integration. And then there is Amdahl's Law that must be introduced to the mix [1], [2]. In addition, one needs to consider how the architecture differs for high-performance applications versus long battery life portable applications.

ARCHITECTURES OF MULTICORE DSPS

In 2008, 68% of all shipped DSP processors were used in the wireless sector, especially in mobile handsets and base stations; so, naturally, development in wireless infrastructure and applications is the current driving force behind the evolution of DSP processors and their architectures [3]. The emergence of new applications such as mobile TV and high-speed Internet browsing on mobile devices greatly increased the demand for more processing power while lowering cost and power consumption. Therefore, multicore DSP architectures were established as a viable solution for high-performance applications in packet telephony, third generation (3G) wireless infrastructure and worldwide interoperability for microwave access (WiMAX) [4]. This shift to multicore shows significant improvements in performance, power consumption, and space requirements while lowering costs and clocking frequencies. Figure 3 illustrates a typical multicore DSP platform.

Current state-of-the-art multicore DSP platforms can be defined by the type of cores available in the chip and include homogeneous and heterogeneous architectures. A homogeneous multicore DSP architecture consists of cores that are from the same type, meaning that all cores in the die are DSP processors. In contrast, heterogeneous architectures contain different types of cores. This can be a collection of DSPs with general-purpose processors (GPPs), graphics processing units (GPUs), or microcontroller units (MCUs). Another classification of multicore DSP processors is by the type of interconnects between the cores.

More details on the types of interconnect being used in multicore DSPs as well as the memory hierarchy of these multiple cores are presented below, followed by an overview of the latest multicore chips. A brief discussion on performance analysis is also included.



[FIG3] Typical multicore DSP platform.

INTERCONNECT AND MEMORY ORGANIZATION

As shown in Figure 4, multiple DSP cores can be connected together through a hierarchical or mesh topology. In hierarchical interconnected multicore DSP platforms, data transfers between cores are performed through one or more switching units. To scale these architectures, a hierarchy of switches needs to be planned. Central processing units (CPUs) that need to communicate with low latency and high bandwidth will be placed close together on a shared switch and will have low latency access to each

others' memory. Switches will be connected together to allow more distant CPUs to communicate with longer latency. Communication is done by memory transfer between the memories associated with the CPUs. Memory can be shared between CPUs or be local to a CPU. The most prominent type of memory architecture makes use of Level 1 (L1) local memory dedicated to each core and Level 2 (L2), which can be dedicated or shared between the cores as well as Level 3 (L3) internal or external shared memory. If local, data is moved off that memory to another local memory using a non-CPU block in charge of block memory transfers, usually called direct memory access (DMA). The memory map of such a system can become quite complex and caches are often used to make the memory look "flat" to the programmer. L1, L2, and even L3 caches can be used to automatically move data around the memory hierarchy without explicit knowledge of this movement in the program. This simplifies and makes more portable the software written for such systems but comes at the price of uncertainty in the time a task needs to complete because of uncertainty in the number of cache misses [5].

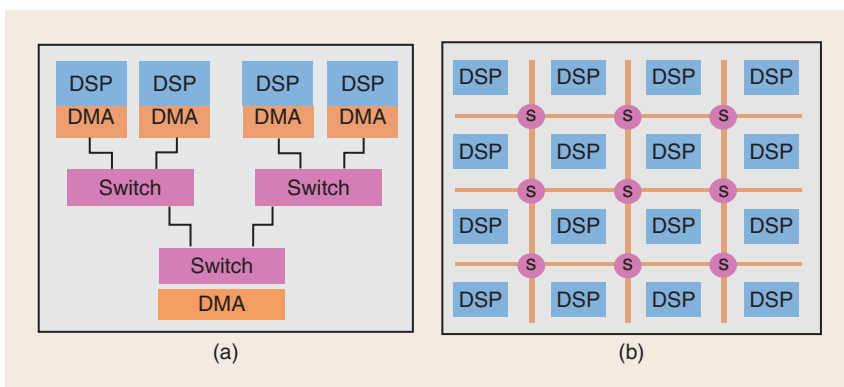
In a mesh network [6], [7], the DSP processors are organized in a two-dimensional (2-D) array of nodes. The nodes are connected through a network of buses and multiple simple switching units. The cores are locally connected with their "north," "south," "east," and "west" neighbors. Memory is generally local, though a single node might have a cache hierarchy. This architecture allows multicore DSP processors to scale to large numbers without increasing the complexity of the buses or switching units. However, the programmer generally has to

write code that is aware of the local nature of the CPU. Explicit message passing is often used to describe data movement.

Multicore DSP platforms can also be categorized as symmetric multiprocessing (SMP) platforms and asymmetric multiprocessing (AMP) platforms. In an SMP platform, a given task can be assigned to any of the cores without affecting the performance in terms of latency. In an AMP platform, the placement of a task can affect the latency, giving an opportunity to optimize the performance by optimizing the placement of tasks. This optimization comes at the expense of an increased programming complexity since the programmer has to deal with both space (task assignment to multiple cores) and time (task scheduling). For example, the mesh network architecture of Figure 4 is AMP since placing dependent tasks that need to heavily communicate in neighboring processors will significantly reduce the latency. In contrast, in a hierarchical interconnected architecture, in which the cores mostly communicate by means of a shared L2/L3 memory and have to cache data from the shared memory, the tasks can be assigned to any of the cores without significantly affecting the latency. SMP platforms are easy to program but can result in a much increased latency as compared to AMP platforms.

EXISTING VENDOR-SPECIFIC MULTICORE DSP PLATFORMS

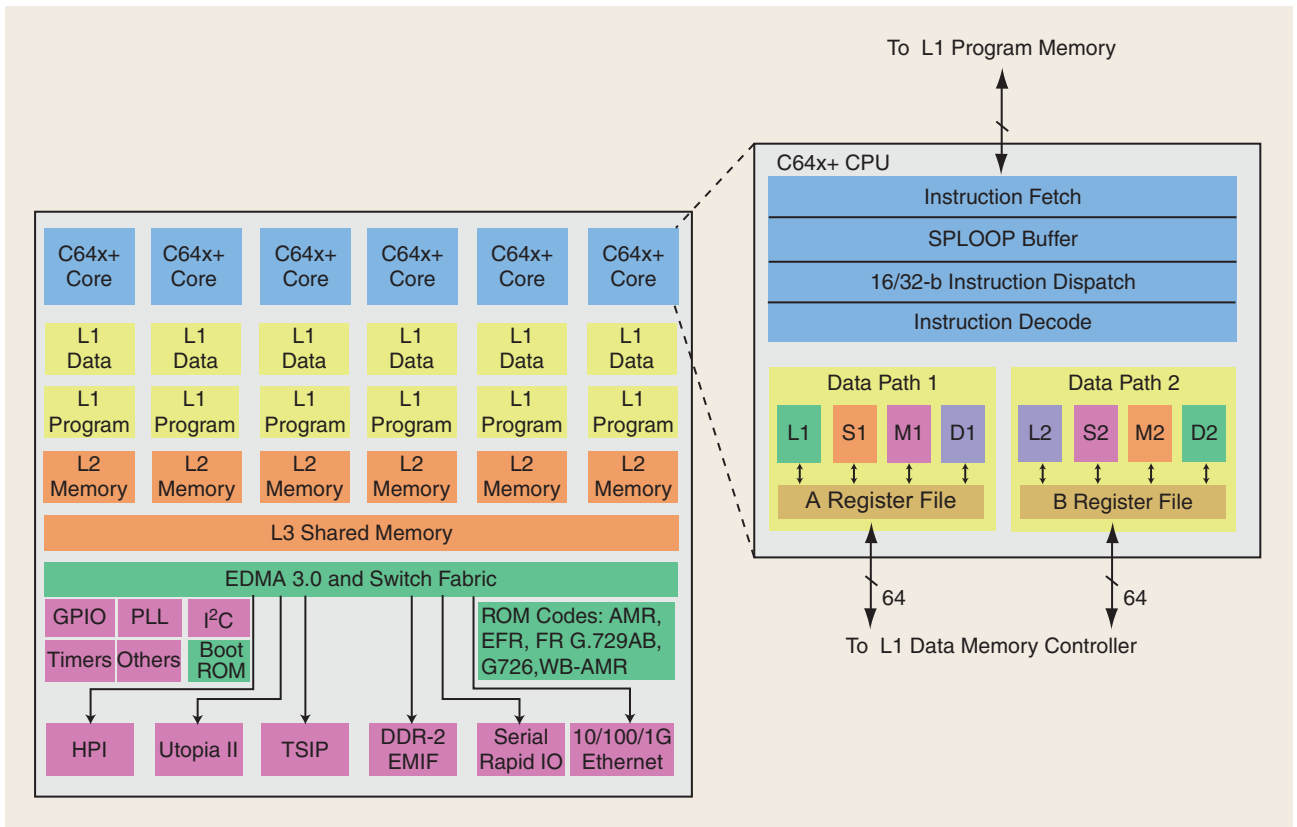
Several vendors manufacture multicore DSP platforms such as TI [8], Freescale [9], picoChip [10], Tileria [11], and Sandbridge [12], [13]. Table 1 provides an overview of a number of these multicore DSP chips.



[FIG4] Interconnect types of (a) hierarchical network and (b) mesh network multicore DSP architectures.

[TABLE 1] MULTICORE DSP PLATFORMS.

	TI [8]	FREESCALE [9]	PICOCHIP [10]	TILERA [11]	SANDBRIDGE [12], [13]
PROCESSOR ARCHITECTURE	TNETV3020	MSC8156	PC205	TILE64	SB3500
NUMBER OF CORES	HOMOGENEOUS	HOMOGENEOUS	HETEROGENEOUS	HOMOGENEOUS	HETEROGENEOUS
INTERCONNECT TOPOLOGY	SIX DSPs	SIX DSPs	248 DSPs AND 1 GPP	64 DSPs	THREE DSPs AND 1 GPP
APPLICATIONS	HIERARCHICAL	HIERARCHICAL	MESH	MESH	HIERARCHICAL
	WIRELESS	WIRELESS	WIRELESS	WIRELESS	WIRELESS
	VIDEO			NETWORKING	
	VOIP			VIDEO	



[FIG5] Texas instruments TNETV3020 multicore DSP processor.

TI has a number of homogeneous and heterogeneous multicore DSP platforms, all of which are based on the hierarchical-interconnect architecture. One of the latest platforms is the TNETV3020 (Figure 5), which is optimized for high-performance voice and video applications in wireless communications infrastructure [8]. The platform contains six TMS320C64x+ DSP cores each capable of running at 500 MHz and consumes 3.8 W of power. TI also has a number of other homogeneous multicore DSPs, such as the TMS320TCI6488, which has three 1 GHz C64x+ cores and the older TNETV3010, which contains six TMS320C55x cores, as well as the TMS320VC5420/21/41 DSP platforms with dual and quad TMS320VC54x DSP cores.

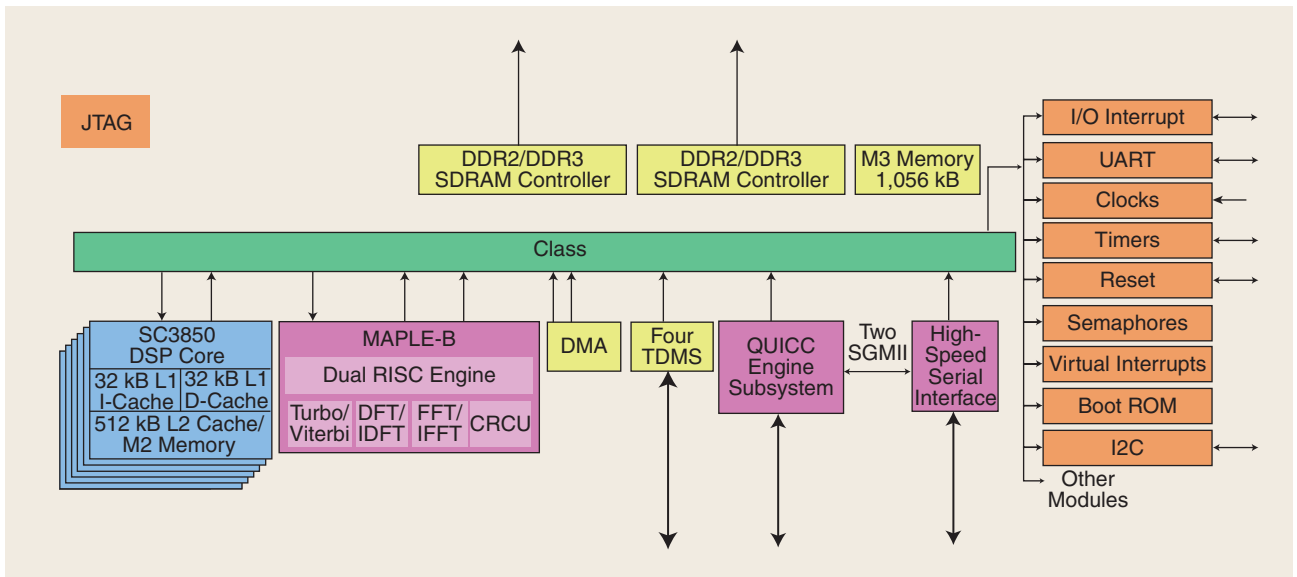
Freescale's multicore DSP devices are based on the StarCore 140, 3400, and 3850 DSP subsystems that are included in the MSC8112 (two SC140 DSP cores), MSC8144E (four SC3400 DSP cores), and its latest MSC8156 DSP chip (Figure 6), which contains six SC3850 DSP cores targeted for 3G-long-term evolution (LTE), WiMAX, 3GPP/3GPP2 and time division synchronous code division multiple access (TD-SCDMA) applications [9]. The device is based on a homogeneous hierarchical interconnect architecture with chip level arbitration and switching system (CLASS).

PicoChip manufactures high-performance multicore DSP devices that are based on both heterogeneous (PC205) and homogeneous (PC203) mesh interconnect architectures. The PC205 (Figure 7) was taken as an example of these multicore

DSPs [10]. The two building blocks of the PC205 device are an ARM926EJ-S microprocessor and the picoArray. The picoArray consists of 248 VLIW DSP processors connected together in a 2-D array as shown in Figure 8. Each processor has dedicated instruction and data memory as well as access to on-chip and external memory. The ARM926EJ-S used for control functions is a 32-b RISC processor. Some of the PC205 applications are in high-speed wireless data communication standards for metropolitan area networks (WiMAX) and cellular networks [high-speed downlink packet access (HSDPA) and wideband code division multiple access (WCDMA)], as well as in the implementation of advanced wireless protocols.

Tilera manufactures the TILE64, TILEPro36, and TILEPro64 multicore DSP processors [11]. These are based on a highly scalable homogeneous mesh interconnect architecture.

The TILE64 family features 64 identical processor cores (tiles) interconnected using a mesh network of buses (Figure 9). Each tile contains a processor, L1 and L2 cache memory, and a nonblocking switch that connects each tile to the mesh. The tiles are organized in an 8×8 grid of identical general processor cores and the device contains 5 MB of on-chip cache. The operating frequencies of the chip range from 500–866 MHz and its power consumption ranges from 15 to 22 W. Its main target applications are advanced networking, digital video, and telecom.



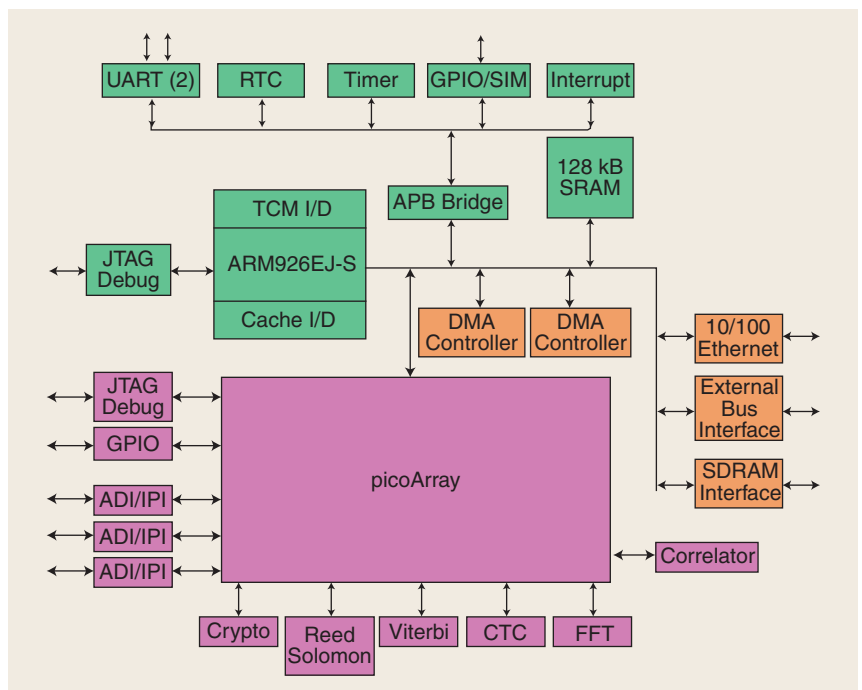
[FIG6] Freescale 8156 multicore DSP processor.

SandBridge manufactures multicore heterogeneous DSP chips intended for software-defined radio applications. The SB3011 includes four DSPs each running at a minimum of 600 MHz at 0.9 V. It can execute up to 32 independent instruction streams while issuing vector operations for each stream using an SIMD datapath. An ARM926EJ-S processor with speeds up to 300 MHz implements all necessary I/O devices in a smart phone and runs Linux OS. The kernel has been designed to use the POSIX pthreads open standard [14] thus providing a cross-platform library compatible with a number of operating systems (Unix, Linux, and Windows). The platform can be programmed in a number of high-level languages including C, C++, or Java [12], [13].

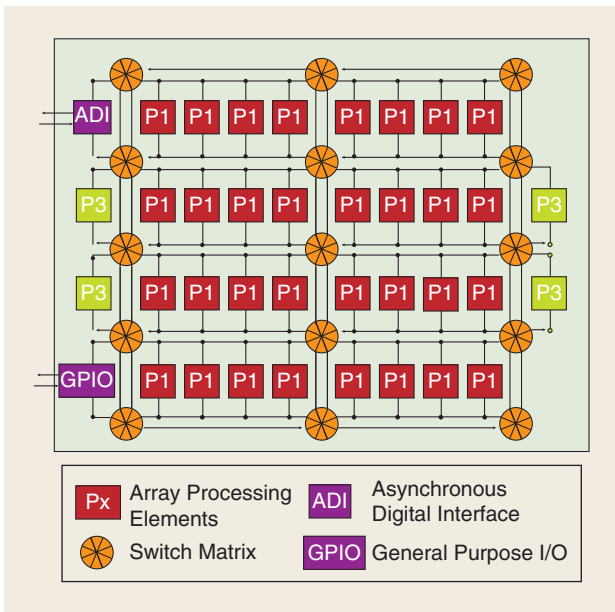
MULTICORE DSP PLATFORM PERFORMANCE ANALYSIS

Benchmark suites have been typically used to analyze the performance among architectures [15]. In practice, benchmarking of multicore architectures has proven to be significantly more complicated than benchmarking of single core devices because multicore performance is affected not only by the choice of CPU but also very heavily by the CPU interconnect and the connection to memory. There is no single agreed-upon programming language for multicore programming and, hence, there is no equivalent of the “out of the box” benchmark, commonly used in single core benchmarks. Benchmark performance

is heavily dependent on the amount of tweaking and optimization applied as well as the suitability of the benchmark for the particular architecture being evaluated. As a result, it can be seen that single-core benchmarking was already a complicated task when done well, and multicore benchmarking is proving to be exponentially more challenging. The topic of benchmark suites for multicore remains an active field of study [16]. Currently available benchmarks are mainly simplified benchmarks that were primarily developed for single-core systems.



[FIG7] picoChip PC205 multicore DSP processor.



[FIG8] The picoChip picoArray.

[TABLE 2] BTDI OFDM BENCHMARK RESULTS ON VARIOUS PROCESSORS FOR THE MAXIMUM NUMBER OF SIMULTANEOUS OFDM CHANNELS PROCESSED IN REAL TIME. THE SPECIFIC NUMBER OF SIMULTANEOUS OFDM CHANNELS IS GIVEN IN [17].

	CLOCK (MHZ)	DSP CORES	OFDM CHANNELS
TI TMS320C6455	1,200	1	LOWEST
FREESCALE MSC8144	1,000	4	LOW
SANDBRIDGE SB3500	500	3	MEDIUM
PICOCHIP PC102	160	344	HIGH
TILERA TILE64	866	64	HIGHEST

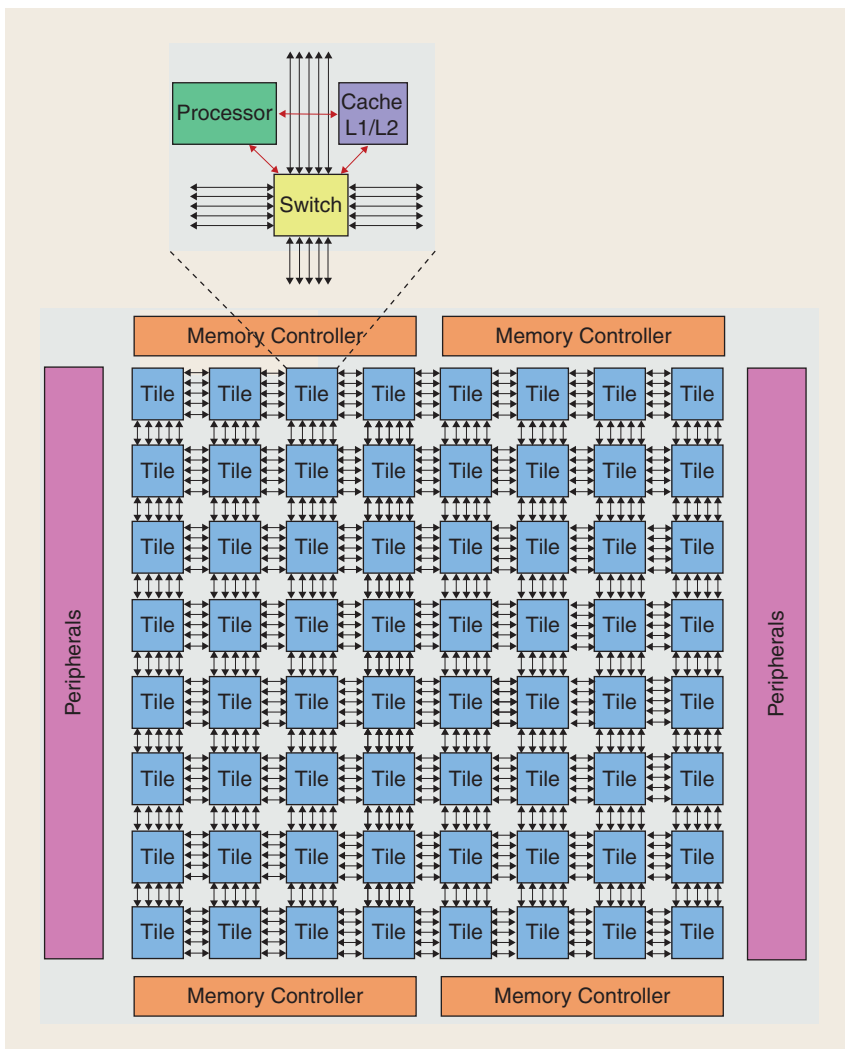
One such a benchmark is the Berkeley Design Technology, Inc. (BTDI) orthogonal frequency division multiplexing (OFDM) benchmark [17] that was used to evaluate and compare the performance of some single and multicore DSPs in addition to other processing engines. The BTDI OFDM benchmark is a simplified digital signal processing path for a fast Fourier transform (FFT)-based OFDM receiver [17]. The path consists of a cascade of a demodulator, finite impulse response (FIR) filter, FFT, slicer, and Viterbi decoder. The benchmark does not include interleaving, carrier recovery, symbol synchronization, and frequency-domain equalization.

Table 2 shows relative results for maximizing the number of simultaneous non-overlapping OFDM channels that can be processed in real time, as would be needed for an access point or a base station. These results show that the four considered multicore DSPs can process in real time a higher number of OFDM channels as compared to the considered single-core processor using this specific simplified benchmark.

However, it should be noted that this application benchmark does not necessarily fit the use cases for which the candidate processors were designed. In other words, different results can be produced using different benchmarks since single and multicore embedded processors are generally developed to solve a particular class of functions that may or may not match the benchmark in use. At the end, what matters most is the actual performance achieved when the chips are tested for the customer's desired end solution.

SOFTWARE TOOLS FOR MULTICORE DSPs

Due to the hard, real-time nature of DSP programming, one of the main requirements that DSP programmers insist on having is the ability to view low-level code, to step through their programs



[FIG9] Tiler TILE64 multicore DSP processor.

instruction by instruction, and evaluate their algorithms and “see” what is happening at every processor clock cycle. Visibility is one of the main impediments to multicore DSP

programming and to real-time debugging as the ability to “see” in real time decreases significantly with the integration of multiple cores on a single chip. Improved chip-level debug techniques and hardware-supported visualization tools are needed for multicore DSPs. The use of caches and multiple cores has complicated matters and forced programmers to speculate about their algorithms based on worst-case scenarios. Thus, their reluctance to move to multicore programming approaches. For programmers to feel confident about their code, timing behavior should be predictable and repeatable [5]. Hardware tracing with embedded trace buffers (ETB) [18] can be used to partially alleviate the decreased visibility issue by storing traces that provide a detailed account of code execution, timing, and data accesses. These traces are collected internally in real time and are usually retrieved at a later time when a program failure occurs or for collecting useful statistics. Virtual multicore platforms and simulators, such as Simics by Virtutech [19], can help programmers in developing, debugging, and testing their code before porting it to the real multicore DSP device.

Operating systems (OSs) provide abstraction layers that allow tasks on different cores to communicate. Examples of OSs include SMP Linux [20], [21], TI’s DSP BIOS [22], and Enea’s OSEck [23]. One main difference between these OSs is in how the communication is performed between tasks running on different cores. In SMP Linux, a common set of tables that reflect the current global state of the system are shared by the tasks running on different cores. This allows the processes to share the same global view of the system state. On the other hand, TI’s DSP/BIOS and Enea’s OSEck supports a message passing programming model. In this model, the cores can be viewed as “islands with bridges” as contrasted with the “global view” that is provided by SMP Linux. Control and management middleware platforms, such as Enea’s dSpeed [23], extend the capabilities of the OS to allow enhanced monitoring, error handling, trace, diagnostics, and interprocess communications.

As in memory organization, programming models in multicore processors include SMP models and AMP models [24]. In an SMP model, the cores form a shared

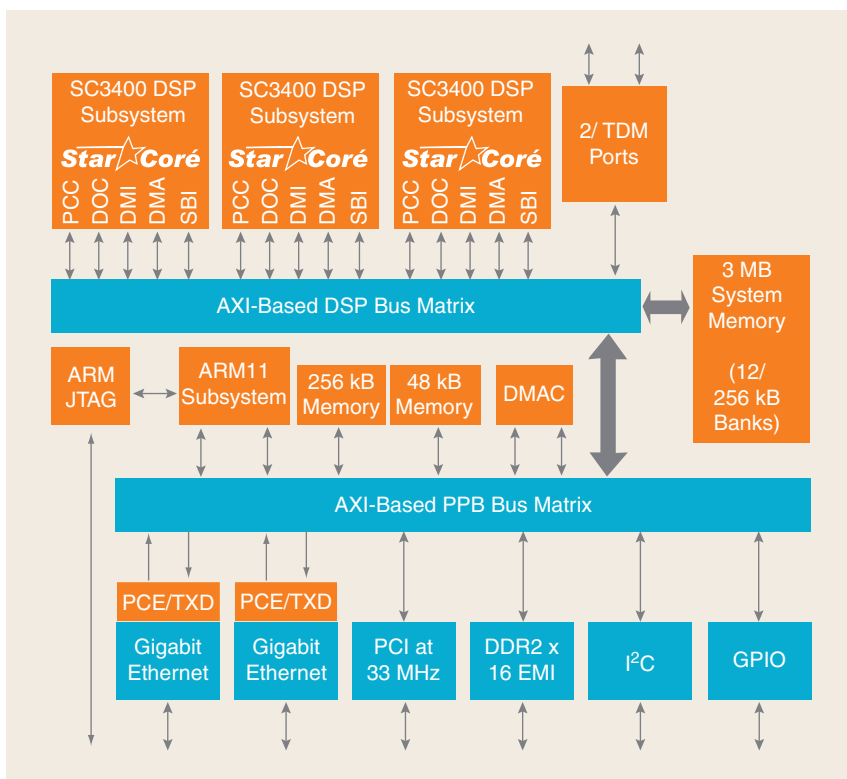
ADDING MULTIPLIER UNITS IS THE SIMPLEST FORM OF DOING MULTIPROCESSING IN A DSP DEVICE.

set of resources that can be accessed by the OS.

The OS is responsible for assigning processes to different cores while balancing the load between all the cores. An

example of such an OS is SMP Linux [18], [19], which boasts a huge community of developers and lots of inexpensive software and mature tools. Although SMP Linux has been used on AMP architectures such as the mesh interconnected Tiler architecture, SMP Linux is more suitable for SMP architectures (see the section “Interconnect and Memory Organization”) because it provides a shared symmetric view. In comparison, TI’s DSP/BIOS and Enea’s OSE can better support AMP architectures since they allow the programmer to have more control over task assignments and execution. The AMP approach does not balance processes evenly between the cores and so can restrict which processes get executed on what cores. This model of multicore processing includes classic AMP, processor affinity, and virtualization [23].

Classic AMP is the oldest multicore programming approach. A separate OS is installed on each core and is responsible for handling resources on that core only. This significantly simplifies the programming approach but makes it extremely difficult to manage shared resources and I/O. The developer is responsible for ensuring that different cores do not access the same shared resource as well as be able to communicate with each other.



[FIG10] The Agere SP2603.

In processor affinity, the SMP OS scheduler is modified to allow programmers to assign a certain process to a specific core. All other processes are then assigned by the OS. SMP Linux has features to allow such modifications. A number of programming languages following this approach have appeared to extend or replace C to better allow programmers to express parallelism. These include OpenMP [25], MPI [26], X10 [27], MCAP1 [28], GlobalArrays [29], and Uniform Parallel C [30]. In addition, functional languages such as Erlang [31] and Haskell [32] as well as stream languages such as ACOTES [33] and StreamIT [34] have been introduced. Several of these languages have been ported to multicore DSPs. OpenMP is an example of that. It is a widely adopted shared-memory, parallel-programming interface providing high-level programming constructs that enable the user to easily expose an application's task and loop-level parallelism in an incremental fashion. Its range of applicability was significantly extended by the addition of explicit tasking features. The user specifies the parallelization strategy for a program at a high level by annotating the program code; the implementation works out the detailed mapping of the computation to the machine. It is the user's responsibility to perform any code modifications needed prior to the insertion of OpenMP constructs. In particular, OpenMP requires that dependencies that might inhibit parallelization are detected and where possible,

AT THE END, WHAT MATTERS MOST IS THE ACTUAL PERFORMANCE ACHIEVED WHEN THE CHIPS ARE TESTED FOR THE CUSTOMER'S DESIRED END SOLUTION.

removed from the code. The major features are directives that specify that a well-structured region of code should be executed by a team of threads, who share in the work. Such regions may be nested. Work

sharing directives are provided to effect a distribution of work among the participating threads [35].

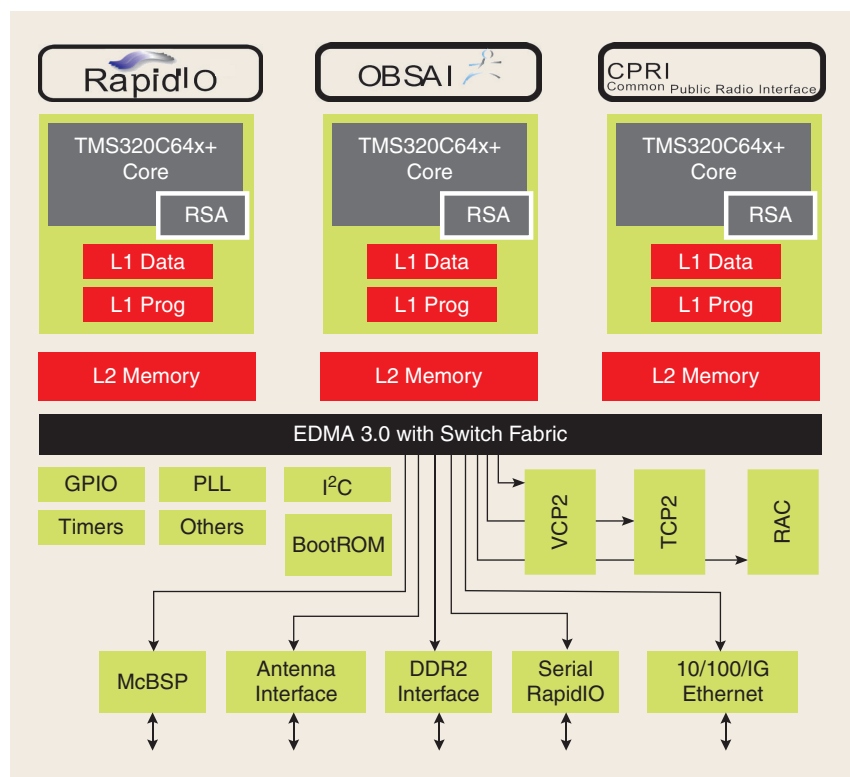
Virtualization partitions the software and hardware into a set of virtual machines (VMs) that are assigned to the cores using a VM manager (VMM). This allows multiple operating systems to run on single or multiple cores. Virtualization works as a level of abstraction between the OS and the hardware. VirtualLogix employs virtualization technology using its VLX for embedded systems [36]. VLX announced support for TI single and multicore DSPs. It allows TI's real-time OS (DSP/BIOS) to run concurrently with Linux. Therefore, DSP/BIOS is left to run critical tasks while other applications run on Linux.

APPLICATIONS OF MULTICORE DSPs

MULTICORE FOR MOBILE APPLICATION PROCESSORS

The earliest SoC multicore in the embedded space was the two-core heterogeneous DSP+ARM combination introduced by TI in 1997. These have evolved into the complex OMAP line of SoC for handset applications. Note that the latest in the OMAP line has both multicore ARM (symmetric multiprocessing) and DSP (for heterogeneous multiprocessing). The choice and number of cores is based on the best solution for the problem at hand and many combinations are possible. The OMAP line of processors is optimized for portable multimedia applications. The ARM cores tend to be used for control, user interaction, and protocol processing, whereas the DSPs tend to be signal processing slaves to the ARMs, performing compute intensive tasks such as video codecs. Both CPUs have associated hardware accelerators to help them with these tasks and a wide array of specialized peripherals allows glueless connectivity to other devices.

This multicore is an integration play to reduce cost and power in the wireless handset. Each core had its own unique function and the amount of interaction between the cores was limited. However, the development of a communications bridge between the cores and a master/slave programming paradigm were important developments that allowed this model of processing to become the most highly used multicore in the embedded space today [37].



[FIG11] TI TCI6487.

MULTICORE FOR CORE NETWORK TRANSCODING

The next integration play was in the transcoding space. In this space, the master/slave approach is again taken, with a host processor, usually servicing multiple DSPs, that is in charge of load balancing many tasks onto the multicore DSP. Each task is independent of the others (except for sharing program and some static tables) and can run on a single DSP CPU. Figure 10 shows the Agere SP2603, a multicore device used in transcoding applications.

Therefore, the challenge in this type of multicore SoC is not in the partitioning of a program into multiple threads or the coordination of processing between CPUs, but in the coordination of CPUs in the access of shared, non CPU, resources, such as DDR memory, Ethernet ports, shared L2 on chip memory, bus resources, and so on. Heterogeneous variants also exist with an ARM on-chip to control the array of DSP cores.

Such multicore chips have reduced the power per channel and cost per channel by an order of magnitude over the last decade.

MULTICORE FOR BASE STATION MODEMS

Finally, the last five years have seen many multicore entrants into the base station modem business for cellular infrastructure. The most successful have been DSP-based with a modest number of CPUs and significant shared resources in memory, acceleration, and I/O. An example of such a device is the TI TCI6487 shown in Figure 11.

Applications that use these multicore devices require very tight latency constraints, and each core often has a unique functionality on the chip. For instance, one core might do only transmit while another does receive and another does symbol rate processing. Again, this is not a generic programming problem. Each core has a specific and very well-timed set of tasks to perform. The trick is to make sure that timing and performance issues do not occur due to the sharing of non-CPU resources [38].

However, the base-station market also attracted new multicore architectures in a way that neither handset (where the cost constraints and volume tended to favor hardwired solutions beyond the ARM/DSP platform) nor transcoding (where the complexity of the software has kept "standard" DSP multicore in the forefront) have experienced. Examples of these new paradigm companies include Chameleon, PACT, BOPS, Picochip, Morpho, Morphics, and Quicksilver. These companies arose in the late 1990s and mostly died in the

FOR PROGRAMMERS TO FEEL CONFIDENT ABOUT THEIR CODE, TIMING BEHAVIOR SHOULD BE PREDICTABLE AND REPEATABLE.

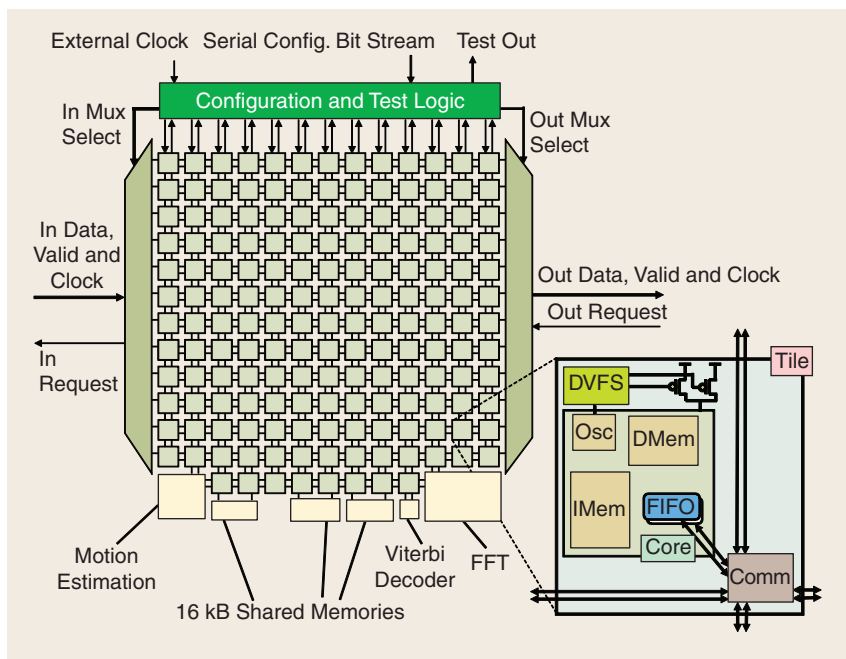
fallout of the tech bubble burst. They suffered from a lack of production quality tooling and no clear programming model. In general, they came in two types; arrays of arithmetic logic

units (ALUs), with a central controller, and arrays of small CPUs, tightly connected and generally intended to communicate in a very synchronized manner. Figure 8 shows the picoArray used by picoChip, a proponent of regular, meshed arrays of processors. Serious programming challenges remain with this kind of architecture because it requires two distinct modes of programming, one for the CPUs themselves and one for the interconnect between the CPUs. A single programming language would have to be able to not only partition the workload but also comprehend the memory locality, which is severe in a mesh-based architecture.

NEXT GENERATION MULTICORE DSP PROCESSORS

Current and emerging mobile communications and networking standards are providing even more challenges to DSP. The high data rates for the physical layer processing, as well as the requirements for very low power have driven designers to use application-specific integrated circuit (ASIC) designs. However, these are becoming increasingly complex with the proliferation of protocols, driving the need for software solutions.

Software-defined radio (SDR) holds the promise of allowing a single piece of silicon to alternate between different modem standards. Originally motivated by the military as a way to allow multinational forces to communicate [39], it has made its way



[FIG12] The AsAP processor architecture.

into the commercial arena due to a proliferation of different standards on a single cell phone (for instance GSM, EDGE, WCDMA, Bluetooth, 802.11, FM radio, and DVB).

CURRENT AND EMERGING MOBILE COMMUNICATIONS AND NETWORKING STANDARDS ARE PROVIDING EVEN MORE CHALLENGES TO DSP.

Sandbridge (see the section “Existing Vendor-Specific Multicore DSP Platforms”) has also been producing DSPs designed for the SDR space for several years.

Signal-Processing On-Demand Architecture (SODA) [40] is one multicore DSP architecture designed specifically for SDR applications. Some key features of SODA are the lack of cache with multiple DMA and scratchpad memories used instead for explicit memory control. Each of the processors has a 32×16 b SIMD datapath and a coupled scalar datapath designed to handle the basic DSP operations performed on large frames of data in communication systems.

Another example is the Asynchronous Array of Simple Processors (AsAP) architecture [41] that relies on the dataflow nature of DSP algorithms to obtain power and performance efficiency. Shown in Figure 12, it is similar to the Tiler architecture at a superficial glance, but also takes the mesh network principal to its logical conclusion, with very small cores (0.17 mm^2) and only a minimal amount of memory per core (128 word program and 128 word data). The cores communicate asynchronously by doubly clocked FIFO buffers, and each core has its own clock generator so that the device is essentially clockless. When a FIFO is either empty or full, the associated cores will go into a low power state until they have more data to process. These and other power-saving techniques are used in a design that is heavily focused on low power computation. There is also an emphasis on local communication, with each chip connected to its neighbors, in a similar manner to the Tiler multicore. Even within the core, the connectivity is focused on allowing the core to absorb data rather than reroute it to other cores. The overall goal is to optimize for data flow programming with mostly local interconnect. Data can travel a distance of more than one core but will require more latency to do so. The AsAP chip is interesting as a “pure” example of a tiled array of processors with each processor performing a simple computation. The programming model for this kind of chip is, however, still a topic of research. Ambric produced an architecturally similar chip [42] and showed that, for simple data flow problems, software tooling could be developed.

An example of this data flow approach to multicore DSP design can be found in [43], where the concept of bulk-synchronous processing, a model of computation where data is shared between threads mostly at synchronization barriers, is introduced. This deterministic approach to the mapping of algorithms to multicore is in line with the recommendations made in [44] where it is argued that adding parallelism in a nondeterministic manner (such as is commonly done with POSIX threads [14]) leads to systems that are unreasonably hard to test and debug. Fortunately, the parallelization of DSP algorithms can often be done in a deterministic manner using data flow diagrams. Hence, DSP may be a more fruitful space for the development of multicore than the general-purpose programming space.

CONCLUSIONS AND FUTURE TRENDS

In the last two years, the embedded DSP market has been swept up by the general increase in interest in multicore that has been driven by companies such as Intel and Sun.

One reason for this is that there is now a lot of focus on tooling in academia and also a willingness on the part of users to accept new programming paradigms. This industry-wide effort will have an effect on the way multicore DSPs are programmed and perhaps architected. But it is too early to say in what way this will occur. Programming multicore DSPs remains very challenging. The problem of how to take a piece of sequential code and optimally partition it across multiple cores remains unsolved. Hence, there will naturally be a lot of variations in the approaches taken. Equally important is the issue of debugging and visibility. Developing effective and easy-to-use code development and real-time debug tools is tremendously important as the opportunity for bugs goes up significantly when one starts to deal with both time and space.

The markets that DSP plays in have unique features in their desire for low power, low cost, and hard real-time processing, with an emphasis on mathematical computation. How well the multicore research being performed presently in academia will address these concerns remains to be seen.

AUTHORS

Lina J. Karam (karam@asu.edu) received the B.E. degree in computer and communications engineering from the American University of Beirut in 1989 and the M.S. and Ph.D. degrees in electrical engineering from Georgia Institute of Technology in 1992 and 1995, respectively. Since 1995, she has been on the faculty in the Electrical Engineering Department at Arizona State University, where she directs the Image, Video, and Usability and the Real-Time Embedded Signal Processing Laboratories. She was awarded the 1998 U.S. National Science Foundation CAREER Award. She is a Senior Member of the IEEE.

Ismail AlKamal (ismail.alkamal@asu.edu) received a B.E. degree in electrical engineering from Aleppo University in 2005 and an M.E. degree in electrical and computer engineering from the American University of Beirut in 2008. In 2008, he was a visiting researcher with the Image, Video, and Usability Group at Arizona State University. He also is the founder and lead system designer at Nawatt Labs, where he worked on several projects in embedded systems, data acquisition, industrial control and automation, vision systems, and ultrasound. He is a Member of the IEEE.

Alan Gatherer (gatherer@ti.com) is a Texas Instruments (TI) Fellow and the CTO for the High Performance Multicore Processor Businesses at Texas Instruments. He led the

development of high performance, multicore DSP at TI and is responsible for the strategy behind digital baseband modem development for 3G and 4G wireless infrastructure as well as high-performance medical equipment. He holds 60 awarded patents and is author of *The Application of Programmable DSPs in Mobile Communications*.

Gene A. Frantz (genf@ti.com) received his B.S.E.E. degree from the University of Central Florida (1971), his M.S.E.E. degree from Southern Methodist University (1977), and his M.B.A. from Texas Tech University (1982). He joined Texas Instruments (TI) in 1974, spending most of his career focusing on DSP, where he is a recognized leader both within TI and throughout the industry. He holds 45 patents and has written more than 50 papers and articles. He is TI's Principal Fellow and a Fellow of the IEEE.

David V. Anderson (dva@ece.gatech.edu) received his B.S and M.S. degrees from Brigham Young University and a Ph.D. degree from Georgia Institute of Technology (Georgia Tech) in 1993, 1994, and 1999, respectively. He is currently an associate professor in the School of Electrical and Computer Engineering at Georgia Tech and codirector of the Advanced Center for Embedded Systems. His research interests are in signal processing and embedded systems. He was awarded the 2004 National Science Foundation CAREER Award and the 2004 Presidential Early Career Award for Scientists and Engineers. He is a Senior Member of the IEEE.

Brian L. Evans (bevans@ece.utexas.edu) received a B.S. degree in electrical engineering and computer science from the Rose-Hulman Institute of Technology in 1987 and M.S. and Ph.D. degrees in electrical engineering from Georgia Institute of Technology in 1988 and 1993, respectively. From 1993 to 1996, he was a post-doctoral researcher in design automation for embedded systems at the University of California, Berkeley. Since 1996, he has been on the faculty at The University of Texas at Austin, where he is currently an electrical and computer engineering professor. In 1997, he won the U.S. NSF CAREER Award. He is a Fellow of the IEEE.

REFERENCES

[1] G. M. Amdahl, "Validity of the single-processor approach to achieving large scale computing capabilities," in *AFIPS Conf. Proc.*, Apr. 1967, vol. 30, pp. 483–485.

[2] M. D. Hill and M. R. Marty, "Amdahl's Law in the multicore era," *IEEE Comput. Mag.*, vol. 41, no. 7, pp. 33–38, July 2008.

[3] W. Strauss. (2009, Feb.). Wireless/DSP market bulletin. Forward Concepts [Online]. Available: <http://www.fwdconcepts.com/dsp2209.htm>

[4] I. Scheiwe. (2005, Nov.). The shift to multicore DSP solutions. DSP-FPGA [Online]. Available: <http://www.dsp-fpga.com/articles/id/721>

[5] S. Bhattacharyya, J. Bier, W. Gass, R. Krishnamurthy, E. Lee, and K. Konstantinides, "Advances in hardware design and implementation of signal processing systems [DSP Forum]," *IEEE Signal Processing Mag.*, vol. 25, no. 6, pp. 175–180, Nov. 2008.

[6] (2007, Apr.). Practical programmable multicore DSP, picoChip [Online]. Available: <http://www.picochip.com/>

[7] (2008, Aug.). Tile processor architecture technology brief, Tileria [Online]. Available: <http://www.tileria.com>

[8] (2007, Jan.). TNETV3020 carrier infrastructure platform, Texas Instruments [Online]. Available: <http://focus.ti.com/lit/ml/spat174a/spat174a.pdf>

[9] (2008, Dec.). MSC8156 product brief, Freescale [Online]. Available: http://www.freescale.com/webapp/sps/site/prod_summary.jsp?code=MSC8156&nodeId=0127950E5F5699

[10] (2008, Apr.). PC205 product brief, picoChip [Online]. Available: <http://www.picochip.com/>

[11] (2008, Aug.). Tile64 processor product brief, Tileria [Online]. Available: <http://www.tileria.com>

[12] J. Glossner, D. Iancu, M. Moudgill, G. Nacer, S. Jinturkar, and M. Schulte, "The Sandbridge SB3011 SDR platform," in *Proc. Joint IST Workshop Mobile Future and Symp. Trends in Communications (SympoTIC)*, June 2006, pp. ii–v.

[13] J. Glossner, M. Moudgill, D. Iancu, G. Nacer, S. Jinturkar, S. Stanley, M. Samori, T. Raja, and M. Schulte. (2005). The Sandbridge Sandblaster Convergence platform. Sandbridge Technologies Inc. [Online]. Available: <http://www.sandbridgetech.com/>

[14] (2004). POSIX: IEEE Standard 1003.1 [Online]. Available: http://www.unix.org/version3/ieee_std.html

[15] G. Frantz and L. Adams, "The three P's of value in selecting DSPs," *Embedded Syst. Programming*, pp. 37–46, Nov. 2004.

[16] K. Asanovic, R. Bodik, B. C. Catanzaro, J. J. Gebis, P. Husbands, K. Keutzer, D. A. Patterson, W. L. Plishker, J. Shalf, S. W. Williams, and K. A. Yelick. (2006, Dec.). The landscape of parallel computing research: A view from Berkeley. Tech. Rep. UCB/Eecs-2006-183 [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/Eecs-2006-183.pdf>

[17] BDTI [Online]. Available: <http://www.bdti.com/bdtimark/ofdm.htm>

[18] Embedded trace buffer, Texas Instruments eXpressDSP Software Wiki [Online]. Available: http://tiexpressdsp.com/index.php?title=Embedded_Trace_Buffer

[19] VirtuTech [Online]. Available: http://www.virtutech.com/datasheets/simics_mpc8641d.html

[20] H. Dietz. (1996, July). Linux parallel processing using SMP [Online]. Available: <http://cobweb.ecn.purdue.edu/~pplinux/ppsm.html>

[21] M. T. Jones. "Linux and symmetric multiprocessing: Unblocking the power of Linux SMP systems" IBM developerWorks, Mar. 2007 [Online]. Available: <http://www.ibm.com/developerworks/library/l-linux-smp/>

[22] TI DSP/BIOS [Online]. Available: <http://focus.ti.com/docs/tools/folders/print/dspbios.html>

[23] Enea [Online]. Available: <http://www.enea.com/>

[24] K. Williston, "Multicore software: Strategies for success," *Embedded Innovator*, pp. 10–12, Fall 2008.

[25] OpenMP [Online]. Available: <http://openmp.org/wp/>

[26] MPI [Online]. Available: <http://www.mcs.anl.gov/research/projects/mpl/>

[27] P. Charles, C. Grothoff, V. Saraswat, C. Donawa, A. Kielstra, K. Ebcioğlu, C. von Praun, and V. Sarkar, "X10: An object-oriented approach to non-uniform cluster computing," in *Proc. ACM OOPSLA*, Oct. 2005, pp. 519–538.

[28] MCAPI [Online]. Available: <http://www.multicore-association.org/workgroup/comapi.php>

[29] Global arrays [Online]. Available: <http://www.emsl.pnl.gov/docs/global/>

[30] Unified Parallel C [Online]. Available: <http://upc.lbl.gov/>

[31] Erlang [Online]. Available: <http://erlang.org/>

[32] Haskell [Online]. Available: <http://www.haskell.org/>

[33] ACOTES [Online]. Available: <http://www.hitech-projects.com/euprojects/ACOTES/>

[34] StreamIT [Online]. Available: <http://www.cag.lcs.mit.edu/streamit/>

[35] B. Chapman, L. Huang, E. Biscondi, E. Stotzer, A. Shrivastava, and A. Gatherer, "Implementing OpenMP on a high performance embedded multicore MPSoC," presented at the *Proc. IEEE Int. Parallel and Distributed Processing Symp.*, 2009.

[36] VirtualLogix [Online]. Available: <http://www.virtuallogix.com/products/vlx-for-embedded-systems/vlx-for-es-supporting-ti-dsp-processors.html>

[37] E. Heikkilä and E. Gulliksen, "Embedded processors 2009 global market demand analysis," VDC Research [Online]. Available: <http://www.electronics.ca/publications/products/Embedded-Processors-Global-Market-Demand-Analysis.html>

[38] A. Gatherer. (2008, Aug.). Base station modems: Why multicore? Why now? ECN Mag. [Online]. Available: http://www.ecnmag.com/supplements-Base-Station-Modems-Why_Multicore.aspx?menuid=580

[39] Software communications architecture [Online]. Available: <http://sca.jpeo.jtr.mil/>

[40] Y. Lin, H. Lee, M. Who, Y. Harel, S. Mahlke, T. Mudge, C. Chakrabarti, and K. Flautner, "SODA: A high-performance DSP architecture for software-defined radio," *IEEE Micro*, vol. 27, no. 1, pp. 114–123, Jan./Feb. 2007.

[41] D. N. Truong, W. H. Cheng, T. Mohsenin, Z. Yu, A. T. Jacobson, G. Landge, M. J. Meeuwssen, A. T. Tran, Z. Xiao, E. W. Work, J. W. Webb, P. V. Mejia, and B. M. Baas, "A 167-processor computational platform in 65nm," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1130–1144, Apr. 2009.

[42] M. Butts, "Addressing software development challenges for multicore and massively parallel embedded systems," presented at *Multicore Expo*, 2008.

[43] J. H. Kelm, D. R. Johnson, A. Mahesri, S. S. Lumetta, M. Frank, and S. Patel. (2008, Aug.). SCHISM: Scalable cache incoherent shared memory. Univ. of Illinois, Urbana-Champaign. Tech. Rep. UILU-ENG-08-2212 [Online]. Available: <http://www.crhc.illinois.edu/TechReports/2008reports/08-2212-kelm-tr-with-acks.pdf>

[44] E. A. Lee. (2006, Jan.). The problem with threads. UCB Tech. Rep. [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/Eecs-2006-1.pdf>