THE ROLE OF PROGRAMMABLE DIGITAL SIGNAL PROCESSORS (DSP) FOR 3G MOBILE COMMUNICATION SYSTEMS

Sorin ZOICAN
Politehnica University of Bucharest
Bdul Iuliu Maniu nr. 1-3, Phone: +4021 402 48 03, Fax: +4021 316 96 32, Email: sorin@elcom.pub.ro

Abstract: This paper illustrated the role of Digital Signal Processors (DSP) for third generation mobile systems. The currently deployed wireless infrastructure supports the third generation of mobile communications comprised of many distinctly different standards. As a result, interoperability among these standards is nearly impossible with respect to the cellular base station. Many solutions for base station or mobile station have been implemented over the years, and each solution required a combination of two components, ASICs (Application Specific Integrated Circuits), and DSPs (Digital Signal Processors). This two-chip solution partitions the processing tasks between the ASIC and DSP, respectively. Although this solution is functionally acceptable, its system cost and flexibility are not completely optimized. Global objectives and attributes that include worldwide roaming, universal connectivity, high data transmission rates, location service capability, and support for high-quality multimedia services are now required. The paper is organized as follows: the 3G requirements, advanced technologies that involve DSP, architecture examples and conclusions.

Key words: *Digital Signal Processors, 3G Mobile Communications Systems*

I. THE THIRD GENERATION OF MOBILE SYSTEMS REQUIREMENTS

As more and more applications require audio, video and communications processing capabilities, the requirements placed on processors used in base station and mobile stations (portable devices and edge-client devices) have become more computationally and bandwidth intensive. Both RISC microcontrollers (MCU) and DSPs have served these applications.

While RISC processors are traditionally architected to enable efficient asynchronous control flow, DSPs are architected to perform well for synchronous, constant-rate data flow (for example, audio or voice-band applications).

Because so many embedded applications have intense requirements for both control and media processing, engineers have typically used DSPs and MCUs together, either at the board level or in system-on-chip (SoC) integration. Together, the respective functional aspects of RISC processors and DSPs unite as the perfect processing engine for a wide variety of multimedia applications and products, such as cellular telephones, digital cameras, portable networked audio/video devices, and so on.

Key base-station areas that require high-performance DSPs will include:

- Antenna Arrays with Adaptive Digital Beam-Forming (in BS- Base Station)
- Power Control (in both BS and MS – Base and Mobile Stations)
- Voice Processing (in BSC: Base-Station Control)
- Base Band Modem (in BTS: Base Transceiver Station)

Digital signal processors are required both in BS and MS as we can observe in figures 1 and 2. Nowadays, there are some emerging technologies such as:

- DSP – based Internet telephony which bridge between PSTN and packet network (VoIP gateway); the DSP advancements in processing power, smaller footprint, and reductions in power dissipation have expanded number of channels carried on VoIP gateways.
- ADSL market
– Software Radio
– Space-Time Processing

The requirements of 3G mobile systems are presented below. These requirements are supported by novel and enhanced DSP architecture that will be present in section II.

Smart Antennas

Digital beam-forming algorithms are designed to target source locations in a noisy environment. They rapidly compare responses of several spatially deployed antennas; the result of the computation is a signal that is believed to have originated from the target direction. Basically, they compute a correlation function that compares the signals and gives a measure of how close the desired and received signals are. Due to the many factors involved in the algorithm, and their wide dynamic range, floating point multiply accumulate operations are used almost exclusively to minimize roundoff errors. The target is mobile, and could be moving at a significant speed, this adds another dimension of complexity to the computation.

Manuscript received September 1, 2008; revised October 5, 2008

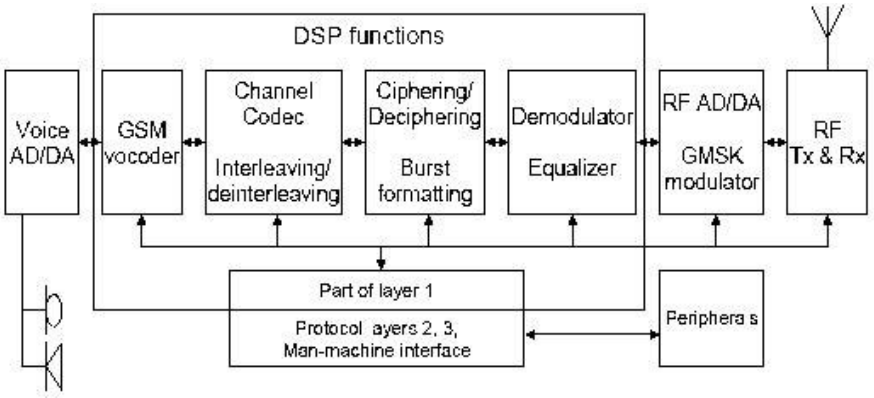

Figure 1. The GSM Mobile station

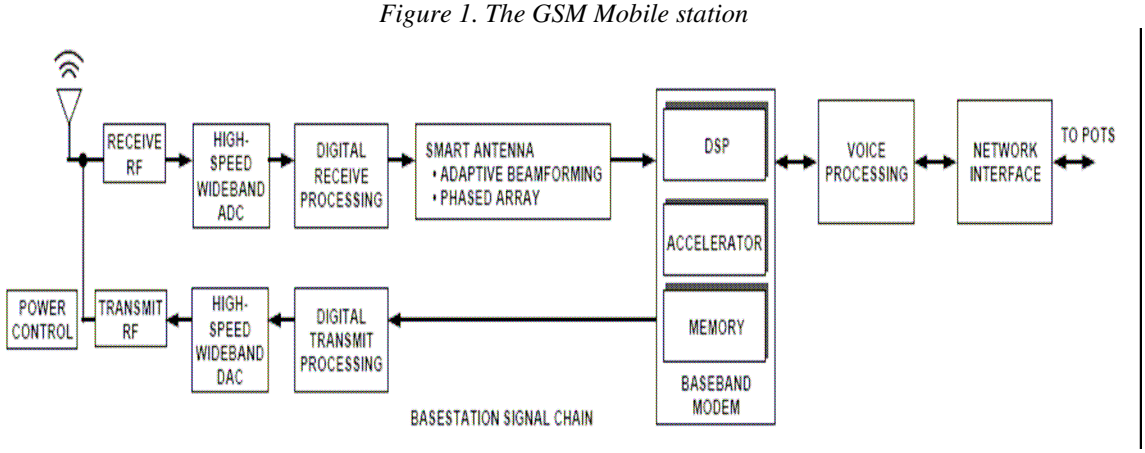

Figure 2. The base Station

Adaptive beam-forming makes use of additional information to continually track the mobile target. Beam-forming in 3G systems may be integrated with the Rake receiver, where the signal is operated on to combat fading and multipath effects. For these algorithms, the TigerSHARC's rapid performance of floating-point computations makes it an excellent fit.

Smart antennas are being regarded by many as the key solution to increasing the spectral efficiency and improving the system performance in mobile communication. In principle, an antenna is smart only when it can recognize and track the signal of a particular mobile telephone while suppressing interfering signals.

Power Control: In the code-division multiple-access (CDMA) systems proposed for 3G, base-station-initiated power control of remote-unit transmitters (uplink) is critical to compensate for fast fading, peaks in transmission power, and to avoid near-far problems3. This is necessary to reduce inter-cell interference. The computations required for power control are multiply-accumulate intensive, requiring high performance digital signal processing to meet delay time requirement in 3G systems. Base stations may also implement the feedback mode transmit diversity (FMTD) algorithm, which is a power control/beam forming application that uses multiple antenna transmissions with varying weights. The computation is multiply-accumulate intensive. For such applications, ADSP-21065L SHARC is a processor of choice.

Voice Processing: DSPs are the traditional choice for speech processing within the cellular system. The phone user's opinion of the quality of the system is directly dependent on the performance of the speech coder, and this has a strong influence on the channel density. Several speech coders are in use today in current 2G systems and must be supported in 3G systems. Although lower codec bit rates increase equipment capacity, they worsen the speech quality. The critical DSP characteristics for high-quality voice processing combine large on-chip RAM and high processing capacity to support fast context switching and high channel density. The ADSP-21mod980, with its 8 DSP cores, capable of 600 MMACS (Million MACs per second), is the ideal candidate for this portion of the signal chain.

Base Band Modem: The 3G standard is expected to be an essential factor that enables applications involving the transmission of wideband signals. Accordingly, the base band modem (BTS) must be designed and implemented with the ability to intermix high bandwidth applications and low bandwidth voice and paging. In the downlink, the base transceiver station packages parallel transport-block streams into physical channels; and in the uplink, it recreates the transport blocks from the base band signal.

Figure 3 shows a typical base band modem section of the 3G base station for both uplink and downlink configurations. During downlink, error-coding schemes are first applied to the transport block. Then the blocks are

reordered and recombined with other channels before being sent off to the radio. For the uplink, the rake receiver is first used to sort out multipath effects and possibly to combine the data from several antennas. The blocks are then restored to their original order and channels before forward error correction is applied. In the next section, the partitioning of the base band modem provides insight as to where a designer might choose to use a DSP.
An optimum must be sought between minimizing the performance cost and maximizing the flexibility of the system to handle future design iterations.

Rake, Channel-Encoding/Decoding Hardware-Software Tradeoffs: An overview of the functions depicted in the figure 3 is needed in order to see exactly where DSP is more appropriate than other alternatives. The interleaving, channel segmenting, and rate matching are I/O intensive operations, which combine data from several sources and reorganize data to minimize the effects of errors. Because of the variability of the parameters, data-rates, and memory-referencing, these functions are ideally suited to DSP for manipulation; they would be difficult to implement cost-effectively in an ASIC. The error-coding and -correction algorithms involve significant bit manipulations that—properly implemented - can be implemented in the DSP. The error-correction algorithms also represent an area of the modem that can provide equipment manufacturer differentiation. The encoding standards have been fixed, while decoding is left to implementers to design using their own intellectual property. Companies that have a strong ASIC capability might choose a hard-wired design, while others that are strong in programming and desire flexibility will choose the DSP approach. Analog Devices TigerSHARC DSP provides all the processing capacity to enable a single high speed 3G data channel. 3G systems, employing spread spectrum communications, will utilize CDMA spreading codes in order to provide greater use of available bandwidth. The spreading and dispreading algorithms are multiply-accumulate intensive, but at extremely high data rates. The rake receiver takes its name from the fact that its diagram resembles a garden rake. Each finger tries to correlate the incoming data with the expected spreading code. As a result, the rake receiver needs to be able to process K times the determined bandwidth, where K is the number of fingers in the rake. In addition to that, the receiver must operate at the frequency set by the spreading codes. DSP technology today doesn't cost-effectively support the bandwidth required for the spreading and despreading in 3G systems. However, production systems will not be manufactured and installed for several years, so there is every incentive for this situation to change. A clear examination of the upgrade path and projected performance of DSPs may show that, by the time these systems are deployed, these functions can be handled cost-effectively.

Glueless Homogeneous and Heterogeneous Multiprocessing: Regardless of the technology used to implement each section of the base band modem, a significant amount of data must be moved around the system. In considering the design, components (or groups of components) that support high-bandwidth communications must be used. The TigerSHARC DSP provides several options for high-speed communication, including on-chip DMA (direct memory-access) and SDRAM support, along with dedicated user programmable link ports. In multiprocessing designs, a high speed cluster bus can be used to connect as many as eight TigerSHARC DSPs without additional logic.

Advanced Technologies: The advanced technologies take into consideration the enhancements provided by multiple antenna (MIMO) processing in physical layer. Traditionally, the speed of a wireless link is limited by the radio resource (power, bandwidth). When the transmitter has antennas and the receiver has antennas, the link speed increases linearly with given the same power and bandwidth budget. Multiple antenna introduces Spatial Dimension into the radio resource set. Some technologies are briefly illustrated below. These technologies are suitable for using of DSP.
- *BLAST (Bell-Labs Layered Space-Time Architecture):* At the transmitter, n independent data streams are transmitted out of the n antennas on the same bandwidth. At the receiver, each receive antenna "sees" all of the transmitted sub-streams superimposed, not separately. If multipath scattering is sufficient, these n data streams have different spatial signatures to each of the n receive antennas are separable.
- *Orthogonal Frequency Division Modulation (OFDM):* Due to to multi-path effects that produce the inter-symbol interference, the transmission of a wideband signal is more difficult than a narrowband signal. In order to overcome this effect complex equalizer is needed. OFDM is an effective technique to transmit wideband signal. Serial input data is split into N parallel streams, each stream is fed into a "frequency channel". With respect to each data stream, it is a narrowband signal, and could get through the channel easily without equalizer. At the receiver, the N parallel streams are combined.
- *Utilization of Feedback:* Wireless channel is time varying from good to bad signal very quickly, so adaptive transmission is can be used to exploit the time varying nature. High throughput transmission mode is employed when channel is good; when channel is bad, low throughput transmission mode is employed for better protection.
Need the feedback of channel condition to the transmitter.
- *MAC Layer Enhancement (Jointly Adaptive MAC Scheduling):* For voice service, data source is quite steady. Optimal strategy is to power control users so as to act against fading and maintain a steady channel throughput. For packet data service, data source is in burst. Optimal strategy is to act in line with fading. Increase power when channel is good. Decrease power when channel is bad. For multi-user network, MAC layer is used to coordinate resource allocation among different users. Between a user and a base station, the link throughput (speed) is adaptive based on the link condition.

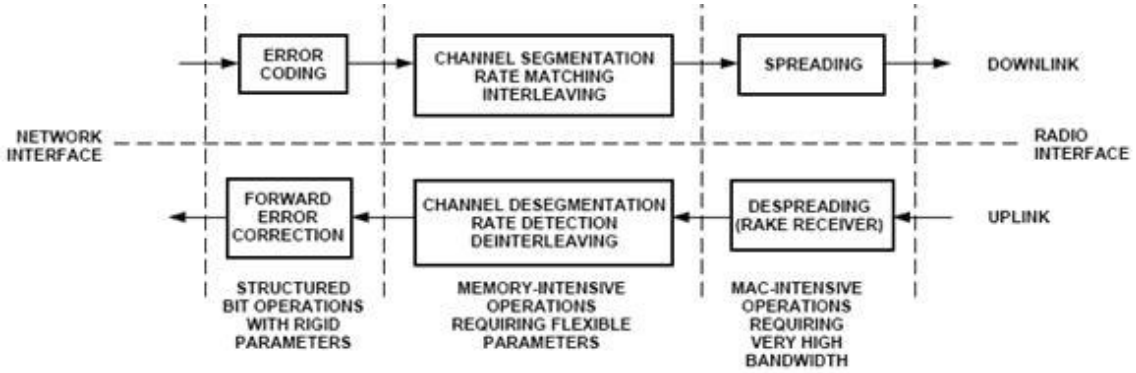

Figure 3. Block diagram showing the baseband processor's signal chain.

The key is to have a jointly adaptive MAC layer and physical layer. Priority is given to users with good channel conditions. (because they could use the radio resource more effectively).

II. DSP ARCHITECTURES FOR 3G MOBILE COMMUNICATIONS SYSTEMS

The algorithms employed in these functional blocks are MAC intensive (i.e., they employ many steps of multiply-and-accumulate). MAC-intensive functions for 3G include FIR, correlation, and equalizer functions. The more rapidly these algorithms are performed, the better the quality and performance of a base station The choice of a DSP to obtain the required computation speed is not a straightforward matter of specifying the highest clock speed.

Architecture and instruction sets greatly affect the speed of algorithm execution. "MIPS" (millions of instructions per second) is also not a valid measure, since each manufacturer counts instructions differently. A highly useful recommended measure, more closely related to algorithm execution, is the peak million multiply- accumulates-per-seconds (MMACS). This calculation is the product of the clock speed and the number of MACs the DSP is capable of executing per clock cycle. Another aspect to consider is the class of DSP architecture employed. Two recently introduced new classes to consider are: very long instruction word (VLIW) and static superscalar.

VLIW attempts to reduce cost and increase execution speed by reducing hardware complexity. The sequencing mechanism in VLIW relies on an instruction format wherein every single execution unit in the chip is under direct programmer or compiler control. Unfortunately, VLIW has little or no hardware support for maintaining the integrity of data dependencies or avoiding scheduling hazards associated with real-time processing. In VLIW, all operation latencies in a particular implementation are fully exposed to software. The TMS320C6x series from Texas Instruments is an example of VLIW architecture.

Static superscalar architectures enforce a consistent and functionally well-defined programming model, and the schedule is determined prior to run time. It incorporates static scheduling techniques like those found in VLIW, but it retains many superscalar and RISC attributes, enabling real time systems. Consequently, code can be written directly in assembly without requiring sophisticated timing prediction. The TigerSHARC™ DSP from Analog Devices is an example of a static superscalar architecture.

Many solutions for base station or mobile station have been implemented over the years, and each solution required a combination of two components, ASICs (Application Specific Integrated Circuits), and DSPs (Digital Signal Processors). This two-chip solution partitions the processing tasks between the ASIC and DSP, respectively. Although this solution is functionally acceptable, its system cost and flexibility are not completely optimized. ASICs, custom devices designed to execute a specific set of tasks within the baseband processing, typically handle chip rate processing functions. DSPs handle the more complex functionality of the symbol rate processing. The distinct difference between these two types of components is the level of programmability or flexibility they offer the system designer. Once released to production, ASIC functionality cannot be enhanced, increased, or changed without entailing a severe and costly re-design. ASIC re-design can cost an equipment manufacturer dearly on two fronts: cost – engineering expenses invariably run on the order of multiple millions of dollars, and time-to-market, which can take the better part of a year, causing a potential market opportunity to be missed. The DSP, however, is a completely programmable device lending itself quite nicely to a global market of evolving and competing communications standards. In order to integrate the chip-rate and the symbol-rate processing capabilities onto a signal piece of silicon, the device architecture must perform the high-value, complex functions associated with symbol-rate processing and the cumbersome, low-value, high speed operations found within the chip rate processing

When deciding how to partition the modem, the nature of the algorithms and data rates become key factors in deciding what should be processed with an ASIC and what should be performed in the DSP. The next section describes some of these algorithms in detail, and explains the tradeoffs. This guideline depends upon the cost to process N channels of a specified bandwidth. For maximum flexibility, the entire structure could be implemented using a cluster of DSPs. On the other hand, a fixed specification may be most-

economically implemented in an ASIC. When evaluating the most suitable approach, the flexibility criterion demands that an engineer consider how the design supports: quickly upgrading the parts of a system to newer technology, scaling the system to improve performance, product differentiation through the addition of new features. The DSP Advantages are the following: they can adaptively mate with different RF front-ends to fit different markets and standards and alllow product differentiation with spare DSP bandwidth.
The following sub-sections will present two DSP architectures that have all the requirements presented above for 3G mobile communications systems.

The TigerSHARC Architecture

The TigerSHARC DSP is a VLIW (very long instruction word) device that permits multiple instructions per line, and therefore reduces the overall cycle count required to perform 3G related functions such as channel decoding, despreading, and path searches. Since the acceleration capabilities reside in software rather than in static hardware blocks or coprocessors, TigerSHARC DSPs provide the flexibility, scalability, and interoperability needed in today's highly competitive market. These characteristics and the directly associated benefits elevate the TigerSHARC and this "Software Radio" approach as the optimal baseband processing solution available today.

When equipment manufacturers design a software-based solution, it enables the adoption of a single, highly flexible platform, providing the ability to adapt to the 3G standards, thus dramatically reducing time to market. This increased flexibility is not lost on the network operators either. The deployment of a completely scalable base station platform not only equates to easier potential updates in the future, but also allows for the immediate interoperability of two or more competing standards, providing greater support for roaming terminals from foreign networks.

Reduced system cost is also addressed by the Software Radio approach. If the complete baseband processing functionality is contained in one device, reduced system cost follows. The ADSP-TS101S TigerSHARC DSP, through a technique called "load balancing," dynamically shifts its processing power between chip-rate and symbol-rate requirements, according to the channel profile at any given time. "Load balancing" eliminates the costly underutilized processing bandwidth found in currently deployed two-chip baseband processing solutions. Additionally, the existence of an optimized instruction set for wireless communications, used by the TigerSHARC DSP, accomplishes the required tasks found in baseband processing with a higher degree of efficiency, enabling remaining processor instruction bandwidth to handle added functionality in the future.
The TigerSHARC architecture is presented in figure 4.

The Blackfin Architecture

Future mobile architecture requires a new combination of DSPs and microcontrollers such as unified micro architectures.Various MCU makers have incorporated some signal processing functionality, such as instruction-set extensions and MAC units, but even this approach lacks the essential architectural basis required for advanced media processing applications. It is important to note that engineers are attracted to these kinds of multiprocessor or multicore design techniques for one reason only: No single processor has the processing power or instruction characteristics to meet the requirements of their applications. This especially rings true for high performance, media-rich embedded applications. Fundamentally, multiprocessing is only employed in the absence of a viable single-processor alternative. Despite the seemingly simple solution offered by multiprocessing, the reality is that no one would suffer the complexities of a multiprocessing environment except as an absolute necessity to meet the needs of a particular application. The need for a "unified" microprocessor for embedded media applications has long been evident.

However, it was not until Analog Devices and Intel jointly developed the high performance Micro Signal Architecture (on which all Blackfin devices are based) that a single architecture was powerful enough, inexpensive enough, and truly optimized both for the complex, real-time world of media data flow and for the control-oriented tasks typically handled by RISC processors. This unified approach to the increasingly media-rich embedded application space is clearly an ideal replacement for previous heterogeneous DSP/MCU integration techniques. The figure 5 illustrate the Blackfin architecture.

Blackfin takes the unique step of architecturally combining media processing attributes like dual MACs (multiply-accumulate engines, commonly used for high performance DSP applications) and classic RISC characteristics like a memory management capability that facilitates simplified, enterprise-level programming modes and styles. The device has DSP features not found on any RISC microcontroller and important microcontroller characteristics not typically on DSPs.

DSPs and microcontrollers have classically remained separate because they are each architecturally optimized for very different kinds of tasks. DSP applications usually focus on performing as many arithmetic computations (such as MAC operations) as possible in the fewest number of core clock cycles. To accomplish this, DSPs often use esoteric VLIW (Very Long Instruction Word) instructions, leaving code density as a secondary concern. Also, DSP applications are data-bandwidth intensive; thus, they often have large memory data buses and DMA engines to reduce the data movement load on the core processor. For media-processing applications and formats, DSPs require ancillary processors to make up for a lack of flexibility. However, in addition to native support for 8-bit data (the word size common to read red-green-blue pixel-processing algorithms), Blackfin's enhanced video instructions and video ALUs process rich-media bit streams at up to 10 times the performance of DSP-only implementations. The unified architecture is designed to support software that can efficiently execute video compression, motion estimation, and Huffman coding algorithms used by video and image-processing standards

such as MPEG2, MPEG4, and JPEG.

Implementing video algorithms in software allows OEMs to adapt to evolving standards and new functional requirements without hardware changes. The core architecture allows for support of algorithms such as MPEG2, MPEG4, and JPEG compression. The integrated video instructions also eliminate complex and confusing communications between the main processor and a separate video CODEC. These features help lower overall system cost while improving the time to market for the end application. The control functions that are typically the focus of microcontrollers, on the other hand, involve many conditional operations, with frequent changes in program flow. These programs are most often written in C or C++, and code density is vital, making variable-length instructions a crucial architectural feature. The unified Blackfin architecture realizes the benefits of both approaches. Its variable-length instruction set extends all the way up to 64-bit opcodes used in DSP inner loops, but the Blackfin instruction set is optimized so that 16-bit opcodes represent the most frequently used instructions. Thus, compiled Blackfin code density figures are competitive with those of popular microcontrollers, the Blackfin architecture is optimized for use with a C/C++ compiler.

Its fully interlocked pipeline and algebraic syntax assembly language make it easy for developers to write in either high level language or assembly with equal ease. Thus, Blackfin's architectural nature is not optimized for either media or microcontroller functions at the expense of the other—rather, it is truly optimized for both, and this is the crux of the breakthrough nature of the Blackfin devices. Because they have system-level responsibilities, true embedded media processors cannot simply operate in the same egocentric fashion as a focused, number-crunching DSP. They must instead have a full set of enterprise- level security characteristics like memory-management capabilities that define separate, freely accessible application-development spaces while keeping distinct code sections safe from being overwritten and the overall system intact. Blackfin, like MCUs, that support full-featured embedded operating systems, has both protected and unprotected operating modes. Respectively, Blackfin calls these "User mode" and "Supervisor mode." User mode prevents users or code from doing anything that affects system-level resources. These kinds of security and system integrity issues are simply not needed by nor architected into traditional DSPs. Like a microcontroller, Blackfin allows asynchronous interrupts and synchronous exceptions. Both kinds of events cause pipelined instructions to suspend execution in favor of servicing the triggering event.

Blackfin's mappable interrupt priorities are a feature common to microcontrollers, but not usually found in DSPs. The chip's exception-handling capabilities are enterprise-class features that protect a Blackfin-based embedded system from invalid or illegal programming. In today's security-focused environment, Blackfin's exceptions are an important means of ensuring that no one succeeds in

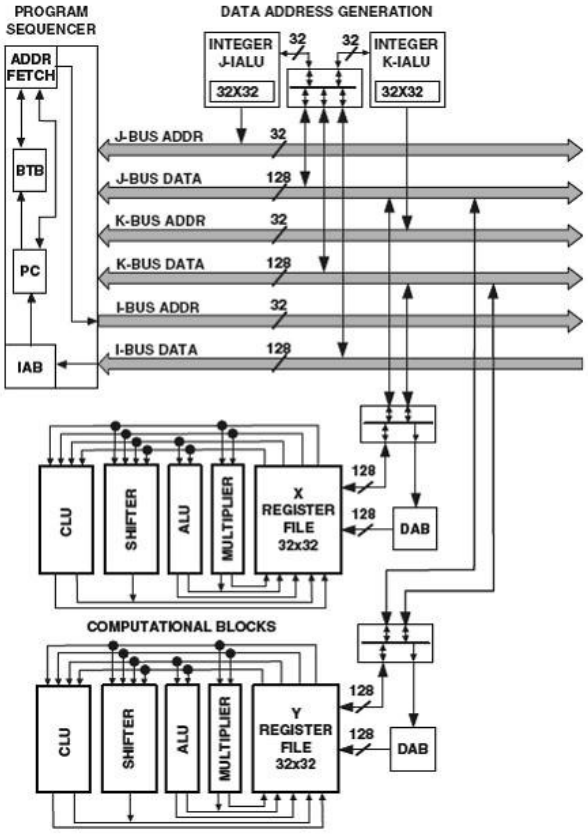

Figure 4. Tiger SHARC block diagram (Analog Devices)

executing an instruction that doesn't exist, or accesses memory that's off limits or not defined (DSPs, on the other hand, allow much easier raw access to the "metal" with few protections). Like other embedded processors, Blackfin can generate interrupts and exceptions in software, facilitating an interface from User mode into the unprotected domain of Supervisor mode. Blackfin's rich set of on-chip peripherals makes the devices ideal for the control side of embedded applications. Blackfin devices integrate a real-time clock, a watchdog timer, general-purpose timers, bidirectional flag/interrupt pins, SPI-compatible ports and UARTs. Some members of the family include PCI and USB connectivity. These peripherals are typical of control-oriented RISC processors. Blackfin devices also provide streaming- media-oriented peripherals, such as a parallel peripheral interface for connecting to high speed video and data converters, and synchronous serial ports for connecting to high resolution digital audio devices and high speed telecom interfaces. One of the problems many developers encounter when trying to apply MCUs to media-processing applications is supporting memory bandwidth fast enough for streaming data. Embedded media processors unequivocally must have "wide open" DMA capabilities for getting data blocks on and off the chip. With the large amounts of data movement required by most media-processing applications, data movement cannot be permitted to cause processor interrupts that would affect real-time performance, and it is inefficient for the core processor to be involved in data movement.

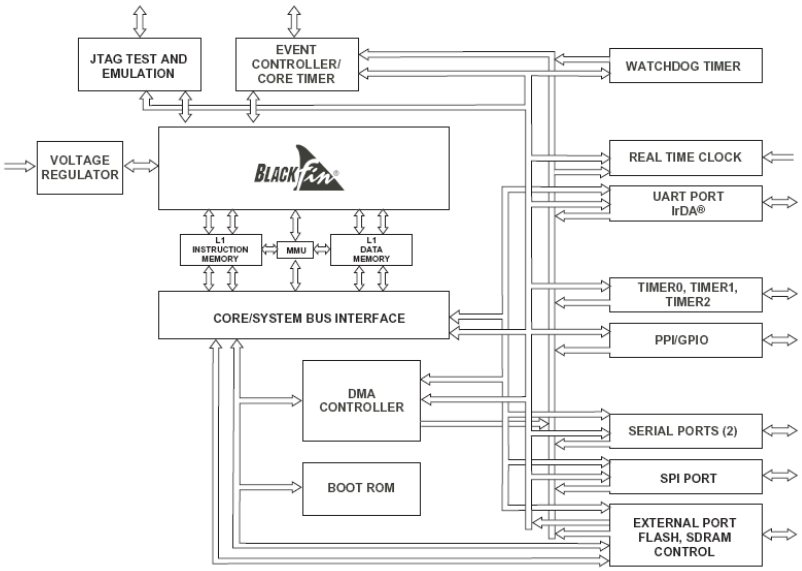

Figure 5. The Blackfin architecture

The DMA channels are independent of the processor core, allowing it to operate deterministically and allowing the DSP data flow in any embedded application to be completely independent of the control processes.
In typical applications, raw data gets DMA'd into the media processor from the chip's peripheral(s), then gets DMA'd to and from external memory during media processing, and then the processed data gets DMA'd back out to the peripheral(s) or to system memory. Blackfin has 16 high speed DMA channels to support bidirectional streaming-media channels between peripherals and memory, virtually eliminating data-movement responsibility from the processor. Blackfin is not a DSP with an enhanced instruction set, nor is it a microcontroller with DSP extensions. The device is an equally high performance media processor and compiler-friendly processor that will be familiar and satisfying to both classes of developers. It allows simple bit-level manipulation, the use of algebraic assembly syntax, and an SRAM memory model for low latency access to assembly modules. Dedicated L1 memory for stack and heap pace, dedicated stack and frame pointers, enhanced address generators, and a very large linear address map make Blackfin as good a compiler target as any RISC processor on the market.

The Blackfin architecture allows L1 fast system memory to be defined as either cache or SRAM. This is another example of how Blackfin allows programmers to flexibly tune and trade off performance versus power consumption through the ability to turn on and off unused chip resources. The Blackfin MMU allows developers to protect selected areas of memory and manage system resources (cache and other memory subsystems) in an enterprise-class embedded environment where not everyone has access to all sections of memory and system resources. Like most microcontrollers, Blackfin has on-chip hardware support for software exceptions, hardware breakpoints, performance counters, and execution trace, plus complete control of the target hardware through a single JTAG port.

The Blackfin software development model enables high performance DSP functionality within a framework matching that of typical RISC devices. System-level and product-level application code can be written in C/C++ and layered on top of a standard real-time operating system (RTOS). Lower-level code, such as media operations, can be optimized in mixed assembly code and C/C++ code, utilizing hand-tuned, assembly libraries of high performance media functions. Blackfin can be just as easily programmed in assembler, compiled C/C++ code, or a mixture of both. Control of power dissipation has long been a major concern for embedded applications that are frequently designed for portable and power-constrained environments. And when embedded systems have needed DSP functionality, choices have been few. When separate MCU and DSP cores are combined on one chip, power control becomes even more difficult and complex.

As an embedded media processor, Blackfin's integrated dynamic power management (DPM) controller can optimally address the needs of a given embedded application. Multiple power modes support a range of system performance levels, and the device is designed to allow selective disabling of clocks to unused peripherals and L2 memory. The PLL frequency can be adjusted over a wide range to save power based on various processing needs, and Blackfin's voltage can be adjusted to enable exponential savings in power consumption. The fast clock rates needed to meet the computational complexity and performance requirements of an embedded media processor make it difficult to apply tactical power-saving design schemes. Blackfin's dynamic power management capabilities optimize performance versus power for specific tasks, supporting a multitiered approach to power management that adjusts performance based on system needs. Going one step

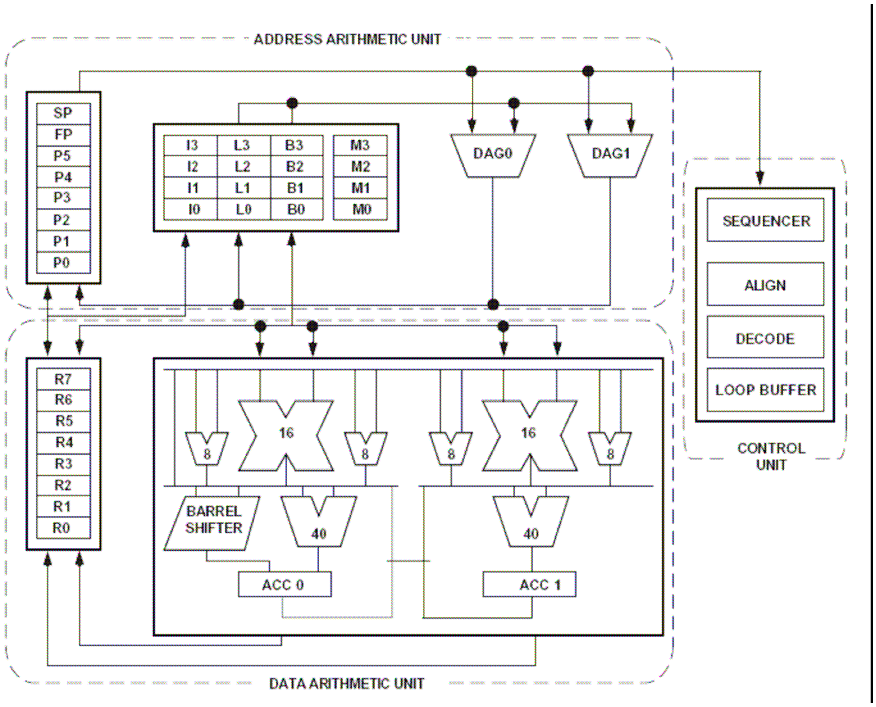

Figure 6. Blackfin processor core

further, Blackfin also allows the core voltage to be changed in concert with frequency changes, so less power is consumed when running a section of code at a lower frequency and a lower voltage, even if execution time is extended. The advantages of using a single, unified core to replace a separate microcontroller and DSP are many. But the key benefits include reduced cost of ownership and faster time to market. Both these benefits are driven by the ability to use a single tool chain to develop code on a single, unified platform. Instead of having to learn two instruction sets and tool chains, developers can learn one instruction set and maintain a single code base running on a single operating system (where applicable). In fact, because they are architected from the ground up for true embedded media processing, these devices actually define a new domain of "media instruction set computing." A unified core has one set of software application programming interfaces (APIs) and drivers, one debugger, one loader, one linker, one language, and so forth. However, the biggest advantage to using a unified processor that delivers the "best of both worlds" is the applications it can enable at performance and price points previously unattainable.

III. CONCLUSION

Designers of 3G base stations will make use of the DSPs in order to achieve the high performance and flexibility needed for tomorrow's voice and data applications. Flexibility at all levels will drive the need for scalable technologies, such as static super scalar architectures and glueless interconnection of system components. Effective embodiment of these design principles will fulfill the promise of 3G to provide the foundations of the kind of wireless infrastructure necessary for tomorrow's killer applications

REFERENCES

[1]. Robert B. De Robertis , Rasekh Rifaat1, "DSPs Enhance Flexible Third-Generation Base- Station Design", Analog Dialogue 33-10, 1999, pp. 6-8

[2]. Gatherer, E. Auslander, editors, "The Application of Programmable DSPs in Mobile Communication", Ed. Wiley, 2002, 430 pages

[3]. Michael Long, "Rethinking Base Station, Baseband Processing for Wireless Communication", Analog Devices, DSP Connections, 2002, pp. 43-45

[4]. Vincent Lau, "A Survey of Enabling Technologies for 3G+ Wireless Systems", Lucent Technology, Bell Labs Innovations, 2002, 22 pages

[5]. Analog Devices , "Unifying Microarchitecture for Embedded Media Processing", 2004, 8 pages

[6]. Analog Devices, "ADSP-TS201 TigerSHARC Processor Hardware Reference", 2005

[7]. Analog Devices, "ADSP-TS201 TigerSHARC Processor Programming Reference", 2005

[8]. Analog Devices "ADSP-BF53x Blackfin Processor Hardware Reference", 2007

[9]. Analog Devices "ADSP-BF53x Blackfin Processor Programming Reference", 2007