

DSP brings base station SDR reality

Developments in digital signal processing (DSP) and data conversion are enabling the commercialization of software-based wideband receiver (or software radio), thereby reducing the cost, size, complexity and power consumption of a base station (perhaps as much as fivefold). Moreover, with the emergence of powerful DSP processors, base stations can now support a multilingual variety of air/modulation schemes and protocols.

By Andy McCann and Brad Brannon

With the 3G rollout occurring in earnest, operators are trying to unify the market and add capacity over a fixed spectrum. For operators, business hasn't grown as fast as they wanted and what growth has occurred has generated less profit than expected due to intense market competition—phone number portability being the latest barrier to customer captivity to disappear.

Operators want to roll out 3G networks that give customers a wider range of services, applications and capacity, providing the basis for increasing minutes-of-use-based revenue streams. But little financial capital exists to do so. One of the most increasingly competitive places

for operators to apply price pressure is on infrastructure equipment OEMs.

Mobile operators want a network that can add capacity, features and services easily because the dynamics of the market make it impossible to get a fix on what features and services customers will demand. This makes flexibility through software upgradeability one of the most important capabilities of 3G base stations in considering cost-to-operate and cost-to-own figures.

SDR moves to center stage

The subject of discussion for several decades, software-defined radio

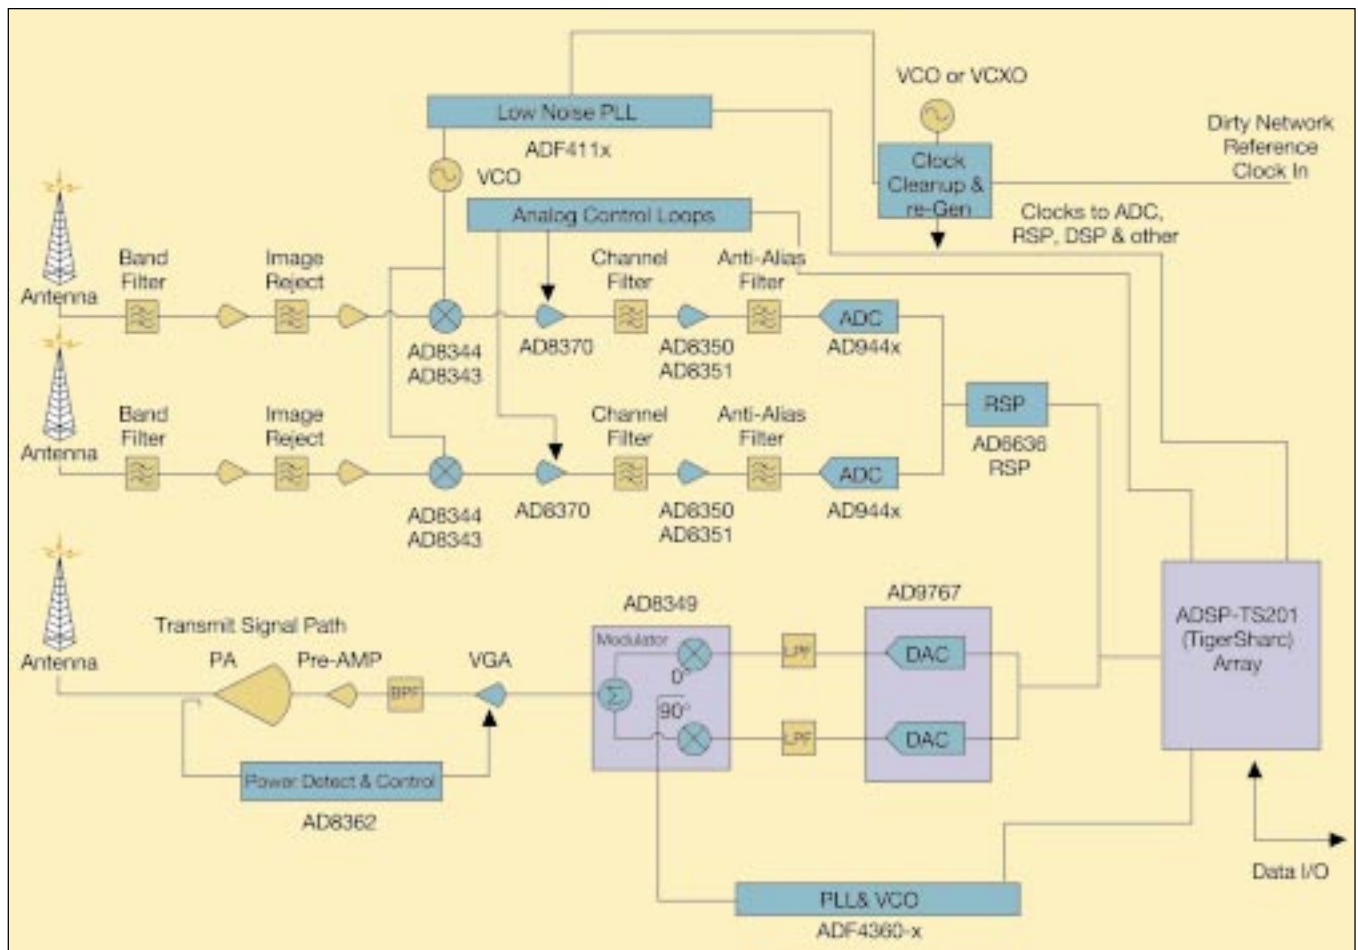


Figure 1. Options for high-performance ADCs are plenty, giving designers multiple ways to optimize a wideband receiver signal chain.

(SDR) has seemed a distant objective because of high costs and technical problems. These limitations have been evident in 3G W-CDMA base stations, where manufacturers have had to employ special-purpose electronic hardware and software to implement the core baseband section of the base station. High costs, low flexibility and restricted operational performance have stood in the way of a viable SDR base station for 3G networks. Until technology exists to support software-only development and deployment, less elegant and more expensive ways to fill the need will be utilized.

Developments in digital signal processing (DSP) and data conversion are enabling the commercialization of a software-based wideband receiver (or software radio), reducing the cost, size, complexity and power consumption of a base station (perhaps as much as fivefold). Powerful processors, such as Analog Devices' TigerSHARC family of DSPs, can support a multilingual variety of air/modulation schemes and protocols (WCDMA, cdma2000, TD-SCDMA, 802.16, 802.20, GSM and others). All the processing is done in software, so it is possible to load new protocols, upgrades and repaired code into the base station as they are developed.

After years in the concept stage, the all-software base station is technologically and economically feasible and is set to have a dramatic impact on the economics of 3G networks and services. The business case boils down to the advantages of removing the need to worry about obsolescence, interoperability or being trapped into out-dated standards—improving the economics of 3G for operators and manufacturers.

Delayed rollout of 3G

The economics of 3G are compelling: 3G networks deliver more

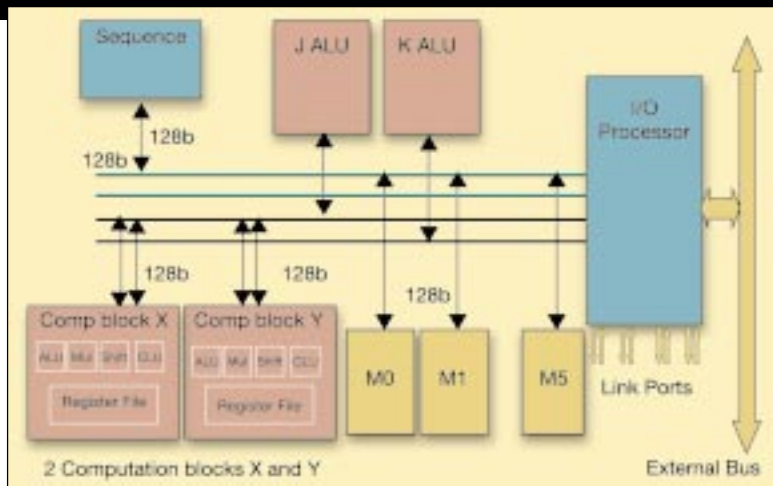


Figure 2. Internal architecture of the TigerSHARC.

capacity, but with costs that are lower than 2G (a voice channel of WCDMA costs about half the price of GSM). But the economics and the pressure to deploy 3G have not spurred a rapid rollout. One reason is the complexity of the technology and the resources required for a launch. Another reason is the pace of standards evolution, and the need within each upgrade for operators to differentiate their service offering through performance and functionality upgrades.

The availability of multiple air interface standards poses the same interoperability dilemma for equipment manufacturers and network operators as they are faced with 2G. In addition, the bridge connecting deployed second-generation systems to 3G systems is cluttered by legal, political, economic and cultural challenges. As a result, the

flexibility of baseband processing solutions is becoming a central issue among infrastructure equipment manufacturers.

Without the emergence of a single, flexible and scalable baseband processing platform that can be used for the competing 3G modes, infrastructure equipment OEMs are faced with escalating engineering costs, significant time-to-market delays and high market risks.

The rate of change and need for upgrading in 3G will be high. A high-profile example is High Speed Downlink Packet Access (HSDPA) in release 5 of WCDMA. HSDPA allows for 14 Mbps downloads from the base station to the mobile unit. HSDPA is commercially significant, but technically difficult. Carriers have little choice but to deploy now and face the need to replace systems in a year to get this feature, or risk a competitor stealing the market with a market-ready product.

An upgradeable base station product based on a processor like TigerSHARC overcomes these barriers for operators and OEMs. Software-defined base stations are not just an option, but a necessary decision for operators—the same platform transforms into a different base station, or one platform can support several versions at one time, removing the risk and fear of obsolescence, enabling carriers to upgrade to newer versions, and supporting multiple multilingual air-interfaces.

Base station requirements

A base station design requires radio frequency (RF) and power amplifier expertise, as well as a high-speed baseband, executing specific DSP operations and complex control protocols. The processing budget for 3G is estimated to be about 500 times greater than for GSM. WCDMA demands fast signal processing, and complex control algorithms such as searching, multipath tracking, and finger assignment, which are better suited for implementation in software.

In conventional cellular base stations, each channel has a receiver tuned to a specific band. Each of these receivers demand a degree of power, size and expense, and with up to 60 channels per cell, the expense of a base station compounds rapidly.

The first challenge is speed—the base station must operate across a wide frequency band. More challenging is the need for dynamic range. In the conventional approach, each radio only deals with a narrow band by filtering out interfering signals, the single carrier receiver concentrates on the desired signal, adjusting gain to optimize signal-to-noise performance and extracting a weak signal from a noisy background. However, with a wideband multicarrier receiver that is no longer possible; no signals can be filtered out in the radio, because they are all required (digital filtering is done after the digitizer). Worse still, there will be a range of signals—very strong ones from nearby mobiles, and very weak ones buried in noise.

As a result the receiver must have an extremely wide dynamic range for enough sensitivity to accurately recover the weak signals without their being swamped by the louder ones. It must also be extremely linear; any distortion or harmonics will generate images of strong signals—ghosts, indistinguishable from genuine signals.

There are two key specifications for the receivers that relate to ADC dynamic range. The first is spurious-free dynamic range (SFDR)—how much of the range below the converter full scale is clean and usable, usually specified with a single tone test but applies equally to the response of the converter to multiple complex signals. In the case of the 800 MHz cellular band, a 3G channel must co-exist with a narrowband signal (AMPS, IS-136 or GSM) that may be 87 dB larger at an offset of less than 1 MHz. While this selectivity is challenging in a traditional receiver with extensive analog filtering, a wideband SDR receiver must deal with the full brunt of the blocker through the entire signal chain. This results in the entire chain having to provide exceptional spurious response. An approximation of the required spurious response can be determined by understanding that the desired signal is dominated by the noise floor of the receiver including the noise figure (NF). With a NF of 4 dB, this places the antenna referred noise at -170 dBm/Hz (-174 dBm/Hz + 4 dB). If a 3G channel (CDMA2000) bandwidth of 1.25 MHz is

assumed, the total noise in the channel of interest is found to be -109 dBm by integrating the noise in this band. At the same time, the desired signal power is -117 dBm (-178 dBm/Hz when spread over 1.25 MHz). Therefore, the expected blocker would be -30 dBm, referenced to the antenna. If a spurious due to this blocker fell on our desired channel, as long as the contribution is at the same level as the noise floor, performance is acceptable. Therefore, this is a spurious response of 79 dBc (-30 dBm minus -109 dBm). However, as with any receiver design, this is the sum performance of the entire signal chain (antenna to baseband processing). Therefore, this must be distributed between each of the distorting components. If the ADC is being considered, it is often desirable for the ADC SFDR and intermodulation performance to be 5 dB to 15 dB below this, translating to between 85 dB and 95 dB SFDR. A similar analysis of W-CDMA will indicate an SFDR requirement for the ADC of between 80 dB and 90 dB.

The other figure of merit is signal-to-noise ratio (SNR). From the example, the noise floor referred to the antenna is -170 dBm including the noise figure of the receiver. Depending on how large the required blockers may be, the conversion gain of the receiver may be as high as 45 dB or as low as 30 dB, making the noise floor presented to the ADC anywhere from -140 dBm/Hz to -125 dBm/Hz. The gain is determined by how large the blockers are expected to be based on the requirements of the standard and which band to be deployed in. In any receiver design, it is important that the ADC noise floor reside well below that of the overall receiver to prevent non-linearities of the ADC from contributing to the characteristics of the receiver. This is due to the fact that DNL errors within an ADC cause the converter noise floor to deviate from the desired Gaussian distribution white noise spectrum. Therefore, to the degree that the ADC contributes to overall noise, these anomalies will affect overall receiver performance. In a practical system, these effects will cause problems ranging from channel power estimation errors to unpredictable changes in the noise floor as a function of input signal level. Therefore, the key is to keep ADC contributions to overall noise to a minimum. This involves keeping the converter noise at least 10 dB below the front end referred thermal noise resulting in required noise spectral densities from -150 dBm/Hz to -135 dBm/Hz. To convert this to an ADC SNR, this noise density must be integrated over the Nyquist band of the converter. Common converter samples rates are 61.44 MSPS, 92.16 MSPS and 122.88 MSPS. At 92.16 MSPS, this results in SNR's that range between 78 dB and 63 dB.

The tradeoff is due to what conversion gain the designer decides to use. In a traditional single channel receiver, the designer can trade off gain and filtering to achieve higher sensitivity and selectivity. However, in a wideband receiver, filtering is not an option and, therefore, higher gain means a lower ADC SNR could be used but at the expense of a higher SFDR requirement. Therefore, a large number of solutions to the wideband multicarrier receiver exist. Products such as the AD9235, AD9430, AD9244, AD6645 and many other 12-bit and 14-bit converters operating at low and high IF frequencies give the designer many different ways to optimize the signal chain (Figure 1).

Architecture issues

While they are software-defined, traditional DSP implementations have not been able to deliver sufficient processing power to economically support this approach. Some previous architectures have used a hybrid approach, combining whatever performance could be mustered by a previous-generation DSP with the hardware power of FPGAs or ASICs. But requirements differ from protocol to protocol. For example, the need for beam-forming and joint detection in Time Division Synchronous CDMA (TD-SCDMA) requires additional chip-rate resources, while the addition of HSDPA to the WCDMA standard makes the scheduler and MAC more complex.

Such heterogeneous multidevice architectures are inefficient because they must support the worst-case loading in each area indepen-

dently even though no one application would require such provisions. Heterogeneous hybrid architectures compound development and deployment headaches and the industry has desired a single, cost-effective processing element that could perform without the need for adjunct hardware support—a single homogeneous environment that can address the variety of tasks, move resources between them when required, and simplify aspects ranging from BOM inventory to in-field support.

DSP-powered SDR

Flexibility and cost, capital and long-term, are the advantages of an SDR base station. These are the key differentiators for base station manufacturers looking to produce products that stand out in price and performance, and which can be adapted easily as 3G standards evolve. They are also important to 3G network operators who need low-cost solutions that offer the best performance and the flexibility to upgrade or modify the system. Fully programmable SDR-based base stations are the only option for operators who want to survive in the 3G era.

The flexibility of a TigerSHARC-based digital stage means that the base station can be reprogrammed to work with new standards. While the concept of SDR based on an instruction set processor has been understood for at least 20 years, technical and cost challenges have made its feasibility elusive.

The TigerSHARC-enabled SDR baseband approach is highly scalable and can be adapted to implement base stations ranging from compact picocell systems, through to full-scale macrocell solutions. Such a baseband implementation is configurable and can be matched to whatever RF interface is required for a particular base station application. The system also provides a general-purpose Layer 1 interface suitable for connection to base station control and data-processing functions.

In order to execute the chip-rate and the symbol-rate processing capabilities required for baseband functionality, the DSP architecture must perform the high value, complex functions associated with symbol-rate processing as well as the cumbersome, low value, high-speed operations found within the chip-rate processing.

As a very long instruction word (VLIW) device, the TigerSHARC DSP permits multiple instructions per line, thereby reducing 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 co-processors, TigerSHARC DSPs give developers and mobile operators the flexibility, scalability and interoperability needed to address today's market dynamics. These characteristics and the associated benefits make the TigerSHARC-enabled SDR approach the optimal baseband processing solution.

Architectural features

There are four attributes of a DSP architecture that can enable software baseband solutions. These are core processing, internal memory bandwidth, memory integration and multiprocessing interfaces.

The most obvious arithmetic attribute of CDMA baseband processing is the complex correlation operations inherent in the rake receiver. Narrow bandwidth user data is translated into a wideband signal by correlating the user data with a two-bit complex code sequence. The receiver must recover time-delayed copies of this wideband signal for each user processed, resulting in a very high rate of complex correlation operations. Close coupling of these low-level operations with intelligent processing is required in the rake receiver. This has led to highly specialized silicon solutions, including applications-specific circuits of high complexity and cost.

The TigerSHARC solves this problem with a small set of instruction set enhancements. These instructions include powerful and flexible correlation operations that approach the processing capability of ASIC implementations. One example is the XCORRS instruction. The XCORRS instruction can despread eight samples of

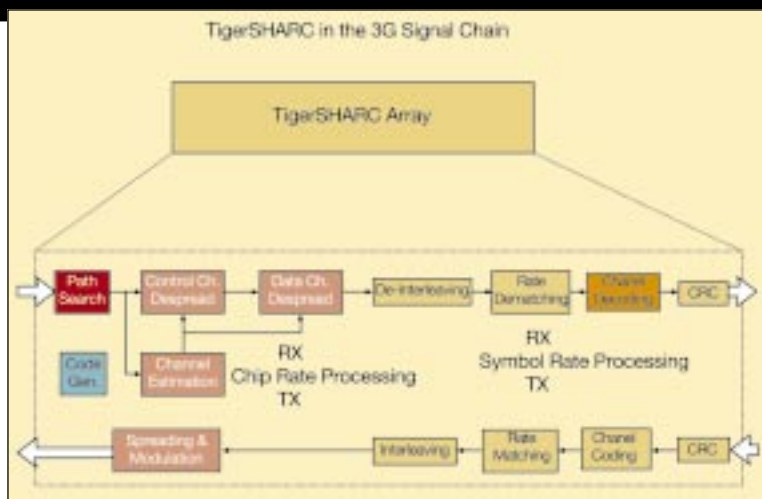


Figure 3. TigerSHARC in the 3G signal chain.

16-bit complex inputs with 16 separate codes sequences simultaneously. This amounts to an aggregate compute rate of 256 complex multiply-accumulates (MACs) per instruction cycle. At 600 MHz for the ADSP-TS201, this is more than 150 GMACs per second. A general-purpose DSP is capable of one or two complex multiply-accumulates per cycle, so a ADSP-TS201 chip-rate performance is as much as two orders of magnitude greater than a traditional DSP, even if that DSP is clocked at 1 GHz or more. A description of how the ADSP-TS201 instruction set enhancements are used in common baseband transceiver functions such as random access channel, path searcher, interpolation, control channel and data channel despreading, and channel decoder can be found in ¹.

All the core processing performance in the world is useless without a very high performance memory system to support it. The ADSP-TS201 is unique among general-purpose DSPs in that it has four 128-bit-wide internal buses connecting the core to the main memory. These buses run at the full rate of the core, resulting in an aggregate throughput of more than 38 Gbytes per second as shown in Figure 2.

The main memory of the ADSP-TS201 consists of 24 Mbits of embedded DRAM, three times the internal memory of the next closest competitor. This large memory allows for flexible software baseband architectures, particularly because it allows the ability to store antenna data buffers on-chip. An elegant and efficient use of processing and memory resources is described in ². In this scheme, DSP MHz are managed according to the traffic mix of voice and data users, saving processing resources that are wasted in traditional fixed partition baseband architectures.

HSDPA, the primary new feature of release 5 of the WCDMA telecommunication specification, brings its own challenges on the memory system. HSDPA is a packet-based channel allowing high downlink throughput rates from the base station to the mobile unit. Packet bursts are possible up to a peak rate of 14 Mbits per second. Copies of the transmitted data must be stored until a receipt message is returned since the data may need to be resent. The large internal memory of the ADSP-TS201 can hold a copy of the data being sent (that may need retransmission), in addition to buffering the next data packet. Traditional architectures based on application-specific processor and external memory solutions may require complex and specialized solutions to this problem.

From a technology perspective, embedded DRAM technology solves three challenges of memory systems needed to support software baseband architectures. The bit density of embedded DRAM is nominally six times better than SRAM, allowing for economical integration of more memory than SRAM-based silicon solutions. The memory size is important to enable storage of the antenna data. Second, the memory leakage performance of embedded DRAM is considerably better than SRAM. Embedded DRAM leakage in 0.13u

at high temperature is negligible, while leakage can be a considerable contributor to power dissipation in small geometry SRAM solutions. Last, the soft-error-rate performance of embedded DRAM is orders of magnitude better than SRAM. Soft errors are non-destructive errors in the memory bit cell due to cosmic rays or other particle contamination. In a large network such as a cellular system, soft errors can be a significant contributor to system down time.

The multiprocessing features of a DSP are an important element that enables a software baseband solution. The complexity of the 3G baseband operations preclude single chip system architectures of any type. However, efficient and simple interfaces allow an array of ADSP-TS201 devices to be viewed as a generic pool of memory and processing resources for baseband tasks (Figure 3). The power of this architecture is that it is flexible to allow multiple communications standards to be supported with a

single homogeneous system architecture, and scalable to allow OEMs to supply various portfolio products (e.g., macrocell, microcell, picocell) with the same software. Furthermore, as Moore's Law allows for more powerful DSPs, the existing software investment can be leveraged while reducing cost and adding features.

SDR destiny

Operators are demanding software feature-based business models and cost-of-ownership scenarios that are difficult to achieve with hardware centric, heterogenous solutions. Processors of the TigerSHARC class are making SDR solutions an economic reality, and base station OEMs are responding accordingly. The balanced combination of core processing power, internal memory bandwidth, memory size, and multiprocessor features allow for highly flexible and scalable software baseband architectures. These software baseband architectures are economical relative to traditional heterogeneous approaches, and are allowing OEMs and operators to employ new business models in a highly competitive market. RFD

1. M. Lopez, R. Rifaat, and Q. Zhang, "DSP for Basestations—The TigerSHARC" Software Defined Radio: Baseband Technologies for 3G Handsets and Basestations, pp. 177–200. John Wiley & Sons Inc. 2003, Edited by W. Tuttlebee.

2. K. Lange, G. Blanke, and R. Rifaat, "A Software Solution for Chip Rate Processing in CDMA Wireless Infrastructure," IEEE Communications Magazine, 2002, 40 (Feb.), 163-67.

ABOUT THE AUTHORS

Brad Brannon is a systems applications engineer, high speed converters. He works in the High Speed Converters Group of Analog Devices in Greensboro, NC. During his 20 years at ADI, he has worked in a number of different engineering roles, most recently as a systems application engineer. He has authored a number of articles on data converters and receiver architectures.

Andrew McCann is marketing manager, DSP Division, Analog Devices. He received his BSEE and MBA from Cornell University. McCann joined ADI in 1995 in corporate marketing. He has been in the DSP Division serving in a marketing role since 1996. Most recently he has been responsible for DSP marketing strategy in wireless infrastructure.