Trends in low power handset Software Defined Radios
Author: Tanuj Raja, Vice President
Dr. John Glossner, EVP & CTO

Sandbridge Technologies, Inc.
1 North Lexington Ave, 10th Floor
White Plains, NY, 10601
914-287-8500

INTRODUCTION

Most SDR designs are concentrating on product implementations for BTS, RF front ends (both for handsets and BTS) and military applications. Relatively few are working on the development of commercial technologies which enable SDR Tier 2 level programmability (e.g. complete physical layer implementations in software). Technical challenges for low power handset SDR implementations include: programmability, power consumption, and cost (size). These are significant technological barriers to overcome, especially since the need for higher data rates, increasing implementation complexity of communication systems (OFDM, MIMO, etc) continues to grow. In addition, mobile terminal development cycles can be long and unpredictable.

Meanwhile the mobile terminal market continues to evolve. The current generation mobile terminals are primarily designed to address a specific geographic or a user market without a possibility of being used as platforms for different models or markets (except for different form factors). In addition, new connectivity investments by the wireless carriers into standards such as WiMAX, HSDPA and even ubiquitous WiFi require mobile terminals to be versatile, have the ability to connect to multiple services and provide continuous concurrent usage. These terminals are complex, complicated to design and difficult to bring through Inter Operability Testing (IOT).

Moreover video, gaming, and digital broadcast are emerging as "must have" technologies in mobile terminals. With VGA LCDs and the desire to have built-in video capabilities for multicasting, viewing digital broadcasts, the traditional applications processor does not have the capability to address the plethora of different multimedia codecs such as JPEG, MP3, H264, DVB.H, ISDB-T, etc. Additionally, GPS and Bluetooth are still required.

For truly convergent devices the following use model emerges:

- Ability to connect to the highest data rate available (LAN) with 802.16e, 802.11 a, b, g, n, HSDPA etc.
- Ability to connect to the Wide Area Network for lower data rates and voice with standards such as GSM/GPRS, CDMA1X, CDMA EV-DO, WCDMA, TD-SCDMA etc
- Broadcast capabilities with DVB.H, ISDB-T, DMB (Satellite and Terrestrial), DAB, FM, XM, etc.
- Audio and Video capabilities with MP3, JPEG, H264, MPEG2-4
- And GPS, Bluetooth etc.

While not all of these standards are required simultaneously for one single mobile terminal - for a truly multimode, multifunction convergent device support of at least one or possibly two LAN standards is needed, at least one WAN standard, at least one on PAN standard, and a combination of broadcast, audio, video capabilities depending upon the wireless carrier service requirements. The challenge is to successfully manage development of this complex, multifunction, multimode mobile terminal.

From the carrier perspective, additional new services are important to retain customers and to provide an enhanced user experience while creating new revenue streams. This has clearly been a challenge. ARPU (Annual Revenue per User) for wireless carriers has been on a

steady decline. In the US, some carriers have deployed initial data services with plans to provide large area networks, but these services are still not generating significant revenue. The penetration levels in European and Asian countries are high and new revenue streams (from technologies such as 3G) have been slow to materialize due to lack of sufficient and user friendly mobile terminals. In emerging markets such as China and India carriers hope to capitalize on lack of pre-existing infrastructure and leapfrogging towards 4G technologies.

True convergence of multimedia, cellular, location and connectivity technologies is expensive, time consuming, and complex at all levels of development. Moreover, the standards themselves have failed to converge which has led to multiple market segments. For mobile terminal OEM's to maintain and/or increase market share any handset development effort must include dozens of combinations of communications systems and multimedia functionality. This requires the handset companies to support multiple platforms and multiple hardware solutions from multiple technology suppliers.

HARDWARE BASED APPROACH

Traditional hardware based approaches for multimode, multifunction terminals are not cost effective. In order to develop these mobile terminals *at least three* different processors are required. Current semiconductor solutions do not offer a single ASIC solution, which prevents the handset OEM's to reduce cost and complexity. Even though there are significant concurrency requirements for the above mentioned standards, a mobile terminal is still a single user interface device. Depending on the usage, at least one or two processors in the hardware based solution stay idle. This is not an optimum solution especially if a smaller form factors are desired. Even more important is the cost associated with multiple ASIC's. Assuming a single hardware based SoC was offered combining all the different standards (resulting in large die size) for these terminals, the solution would still not be cost effective because of the increased die size resulting in increased cost.

Moreover, current hardware based solutions require 18 months of development for commercial SoC's and another 9 to 12 months for successful mobile terminal development with no ability to change, update or modify functions. In the case of the multimode baseband modems, there is no ability to quickly verify updates or add new functionality because any modification requires a re-spin of the SOC. This is a costly and time consuming proposition. Also, wireless carriers and handset OEM's have to anticipate end user requirements and standards evolution up to three years in advance.

For OEM's the long turn around time, increased cost and nearly stagnant market not only erodes margins but with decreased R&D productivity increases the overall development cost. Clearly, this is not a productive solution. In fact the entire wireless supply chain from end users to component suppliers in a traditional hardware based solution requires an unacceptable turn around time for new features, updates and functionality.

SDR BASED APPROACH

SDR enabling technologies clearly have an advantage over the traditional hardware based SoC's; it provides programmability, reduced time to market and reduced overall cost by reusing ASIC hardware for different communication and multimedia standards.

The benefits to end users include:
- Mobile Terminal Independence with the ability to "choose" desired feature sets.
- Global connectivity with the ability to roam across operators
- Future Scalability and a longer lifetime of the handset.

Wireless carriers realize expanded growth from the ability to differentiate their offerings from competitors, and constantly adding new services and increasing their ARPU. Also, the possibility exists to reduce churn since most users switch carriers to upgrade to new features. This provides wireless carriers with the ability to provide over the air (OTA) software upgrades.

Most importantly, by employing an SDR approach, handset OEM's can develop

terminals quickly, provide a rich set of features, and address multiple markets with multiple communications combinations all with a single platform. This design methodology reduces the R&D effort for the handset developers by providing the ability to define the physical layer in software which can then be reused in other combinations for different terminals. The system level validation and mobile terminal prototype verification can start the moment optimized C code for the physical layer and the associated protocol stacks are available. In fact, if desired, the prototype development platform and the physical layer programming can be done in parallel, enabling simultaneous verification, type approval testing, and field testing. By conservative estimates the development time can be reduced by 30%-40%. In addition, once the engineering team is familiar with the platform, the complexity of dealing with a new platform for each mobile terminal is reduced.

Given the programmability, the cost and the time to market advantage, for SDR based solution to be truly effective: power consumption for the processors need to be near parity with traditional approaches. Power dissipation constraints have historically prohibited the use of SDR based baseband and communication processors. Further differentiating a Tier-2 SDR solution, multiple solutions exist including: Processor SDR and. Reconfigurable Defined Radio (RDR). We define processor SDR as an instruction set processor with a software development methodology. It is dynamically reconfigurable in nanoseconds and provides word level parallel processing. In contrast Reconfigurable Defined Radio (RDR) is usually an FPGA based bit level processor with a hardware development methodology providing dynamic reconfigurability in milliseconds. Given the goal of fast dynamic reconfiguration times and real-time requirements, a processor SDR solution is desirable.

Further considerations for the SDR processor are needed to ensure that low latency requirements are met. An SDR processor can either be developed as a DSP or a general purpose processor. However, real-time constraints highlight differences in the designs: \

DSP	General Purpose
Execution Predictability • Required to guarantee real-time constraints	Fast But Non-predictable • Dynamic Instruction Issue • Non-deterministic caches
1 cycle MAC	Multi cycle MAC
0-overhead Loop Buffer	Branch Prediction
Complex Instructions • Multiple Operations Issued	RISC Superscalar Instructions • Multiple Instructions Issued
Harvard Memory Architecture • Multiple memory access	Von Neumann Architecture • Split Cache has similar benefit
Operate on Vector Stream Data	Caches Assume Locality
Data-independent Execution	Data-dependent Execution • Dependent upon operands
Fractional Arithmetic	Integer Arithmetic
Pipeline Non-interlocked • Shallow Pipeline (3-5 stage)	Pipeline Typically Interlocked • Deep Pipeline (5+ stage)
Delayed Branch	Multi cycle Branch

Historically, only general purpose processors have provided programmability from high level languages. DSP's on the other hand have provided real-time parallel execution capabilities. SDR processors should be both programmable and execute in real-time to provide optimized communication and multimedia systems in software

SANDBRIDGE SOLUTION

Based on the Sandblaster® DSP, a highly parallel, multithreaded architecture, easily programmable in C language, the SB3011 processor solves the power consumption

conundrum by utilizing power optimization techniques at core and chip level

The Sandblaster® DSP, designed for parallel, real time communication system processing is optimized for low power consumption by utilizing the following techniques

Instruction set architecture
- Instruction organization simplifies decoding
- Vector operations
  - less decode for same work
- Compound operations with simple decode
  - E.g. load with update
- Sleep/Nap/Doze modes

Circuits
- Low voltage operation
- Minimum dimension transistors
- Limited use of low threshold transistors
- Glitch minimization
- Complex logic flip-flops

Multi-threading
- Turn off threads
- Slower latency memories

IC Implementation
- Register files are 1 read + 1 read/write port
- All memories are single ported
- Data memories are multi-banked, at most 2 banks active per cycle
- Only 1 way of I-cache active
- Jumps resolve before directory lookup
- Directories accessed before cache
- Only 1 stall (memory stall)
- Minimal control signals
- Easy clock gating/functional unit shutdown
- Separate decode stage

By optimizing a single Sandblaster® DSP core (~300K), all instances of the core on the single SoC are optimized, resulting in a highly power efficient design for the processing engine part of the SoC. Each Sandblaster® DSP core runs at 600-800 MHz

As described in the figure 1, SB3011 incorporates four instances of the Sandblaster® DSP to provide 9.6 to 12.8 GMACs sustainable performance to enable real time SDR applications with an integrated ARM9 and widely used smart phone peripherals with 1.5MB on chip memory.
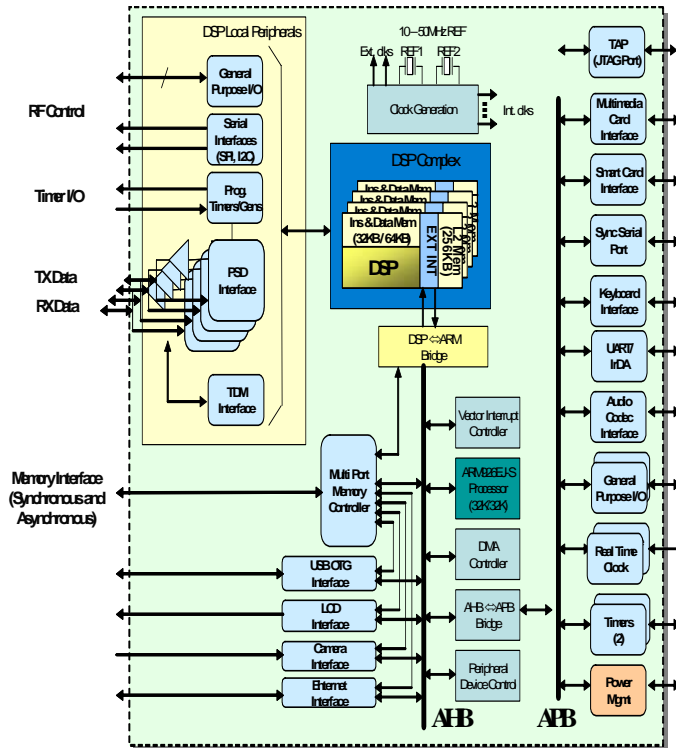

Figure 1: Conceptual view of SB3011 SoC

At the chip level, granular power control for all major units and sub-systems such as DSPs, IO Node/Ring/L2, PLLs, ARM and Peripherals is implemented. An embedded device power management unit (DPMU) incorporates the following features
- Minimal Logic, always powered
- Clocked with external reference
- Accessible to all processors
- Provides Unit "Power-down" and "Wake-up" sequencing
- Provides access to Clock/Frequency generators
- Controls Internal Power Gates
- Communicates with external Power Management Unit (over I2C) for Power Enable/Disable or Voltage Scaling
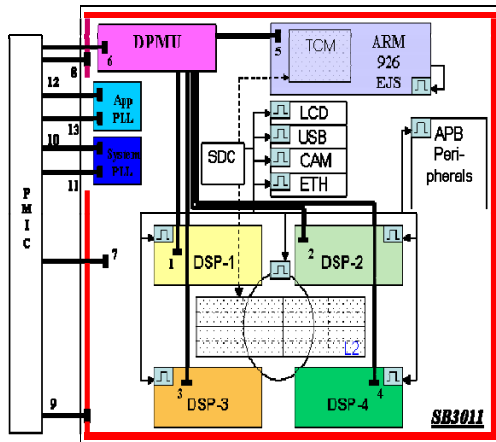
Figure 2: Description of the DPMU   implementation on the SB3011 SoC

The net impact of the core level and chip level power optimization, coupled with architectural decisions for the Sandblaster® DSP is a low power implementation of SDR based baseband and multimedia processor. Following are the measured results on the actual hardware for each Sandblaster® DSP running at 600 MHZ at 0.9V

Target Application	SB3011 (mw) Core w/ L1 Instances 32K Icache 64K DMem	SB3011 (mw) Core w/o L1 Instances
GSM/GPRS, Class 14 (needs one Core)	45	40
WCDMA @ 384 KB (entire application	180	162

needs 3 cores)

SUMMARY

Communication technologies have become more complex, and mobile terminals require multiple applications for high data rate services, multimedia services. In addition, different combinations for different target markets are required Traditional hardware based solutions cannot sustain the time-to-market requirements and cannot offer the flexibility required for combining multiple communication and multimedia services. Moreover, multiple ASIC's are required which increase cost. SDR based solutions offer the ability to quickly modify, update, and enhance the user experience at a lower cost. This is possible through implementation of low power DSP based architecture while incorporating techniques for real time execution of the communication protocols.

Sandbridge's SB3011 bridges the gap of low power design and necessary computing performance requirements. It provides four instances of the Sandblatser® DSP on one SoC and includes the necessary peripherals for a complete mobile terminal design. This solution offers complete baseband and required multimedia processing is software. The software development is done through a high level (C programming) language, thus enabling future reuse and the ability to generate multiple mobile terminals on a single platform. Future implementation of Sandblaster® technology will provide capability for all 4G requirements

A Note about the Authors

 Mr. Tanuj Raja is Vice President of Business Development at Sandbridge Technologies, Inc. With more twelve years of industry experience, Tanuj joined Sandbridge Technologies in February, 2002. As a member of Sandbridge's senior management team, Tanuj is responsible for developing and defining the company's global business strategy. Prior to Sandbridge, Tanuj spent four years in the Wireless Design group at Tality, Inc. (Subsidiary of Cadence Design Systems), Inc. where he held various positions such as Director of Business Development, Worldwide Business Manager for 3G Technologies, and 3G Program Manager. He also has five years of software development experience at Cadence. He has authored and co-authored numerous journal and conference papers. Tanuj has extensive

business development experience in Asia Pacific and Europe; he has successfully developed wireless business opportunities with many Asian, European US companies.

Dr. John Glossner is co-founder, CTO, and Executive Vice President at Sandbridge Technologies. Prior to co-founding Sandbridge, John managed the Advanced DSP Technology group, Broadband Transmission Systems group, and was Access Aggregation Business Development manager at IBM's T.J. Watson Research Center. Prior to IBM, John managed the software effort in Lucent/Motorola's Starcore DSP design center. John received a Ph.D. in Computer Architecture from TU Delft in the Netherlands for his work on a Multithreaded Java processor with DSP capability. He also received an M.S. degree in Engineering Management and an M.S. degree in Electrical Engineering from The National Technological University. John holds a B.S. degree in Electrical Engineering from Penn State. John is a Senior Member of the IEEE and has more than 80 publications and 20 issued patents.