



THE NEXT STEP
IN WIRELESS TECHNOLOGY
COULD PROVE THE TIPPING POINT
FOR MULTICORE EMBEDDED PROCESSING.

4G wireless: EVOLUTION OR WATERSHED IN SOC ARCHITECTURES?

BY RON WILSON • EXECUTIVE EDITOR

The almost-mythological fourth generation of wireless service—4G—could be the fountainhead of an entirely new way of thinking about SOC (system-on-chip) architectures. Or it could drive a simple evolution of today's baseband wireless ICs. It could lead to entirely new kinds of mobile services for consumer clients. Or it could simply handle your e-mail attachments better. It could be a massive engineering challenge struggling into reality in 2015. Or it could happen in a couple of years.

To understand the impact that 4G is likely to have upon SOC design, you have to dig a little bit into just what people mean by the term, understand some of the computing challenges involved in supporting the service, and hear from some system architects on how they are approaching these challenges.

Many of the differences of opinion about the impact of 4G come from a single source: the lack of a clear definition. "You have to start with definitions," warns Bill Krenik, Texas Instruments' chief technology officer for wireless, "because all the controversy and

confusion surrounding the term have left it meaning very little."

Many people, Krenik says, think of 4G in terms of a new world of ubiquitous wireless connectivity—really anywhere, all the time—and in terms of the interactive, location-based, and media-rich services that such connectivity can support. Imagine walking down the street in an unfamiliar town, holding up your handset, and seeing it continuously display a real-time moving image of the street in front of you with map data, labels on buildings and sites of interest, paths to possible destinations, and locations of persons in your address file as

you walk. Or imagine that same handset turning the streets of the town into a multiplayer video game, complete with avatars for other players, 3-D images of aliens and weapons, and realistic rendering of the damage resulting from the virtual battle.

Others—who have to implement the underlying systems—often see 4G in more concrete terms. "Within TI, we don't try to have a dogmatic definition of what 4G is," Krenik explains. "Instead, we refer to the actual technologies by name: HSPA+ [high-speed packet access plus], WiMax, LTE [long-term evolution]. Until 3G Americas comes up with a standard, everything else is just opinion," he continues, referring to an organization whose mission is to promote the deployment of GSM (Global System for Mobile) communications and its evolution to 3G.

Still other engineers take a more quantitative view. Paralleling 3GPP's (Third Generation Partnership Project's) approach in defining LTE, these engineers frame 4G as "100-Mbps-peak throughput for mobile devices and 1-

Gbps peak for nomadic devices such as notebook computers,” says Alan Brown, a senior radio-product manager at Nokia Siemens Networks. Each of these perspectives leads to a different set of expectations for the baseband SOC that will implement the 4G handset.

EVOLVING A BASEBAND SOC

Start with the simplest set of expectations—those that LTE envisions—that the mobile device will somehow achieve a downlink peak data rate of at least 100 Mbps. “This [situation] leads to a baseband that is functionally no different from what we use today for UMTS [Universal Mobile Telecommunications System],” says Freescale Semiconductor Vice President and Senior Fellow Ken Hansen. Blocks include hardware accelerators for sample-rate functions, a CPU core to execute the MAC (media-access controller), a security engine, and a host interface.

At sample rate, data coming from the radio goes through analog-to-digital conversion, through some front-end digital processing, and into an FFT (fast-Fourier-transform) engine that separates the

AT A GLANCE

Fourth-generation wireless service, or 4G, means different things to different people.

Designers can achieve halfway measures with current chip architectures.

Cost and energy constraints will eventually force some architectural innovation.

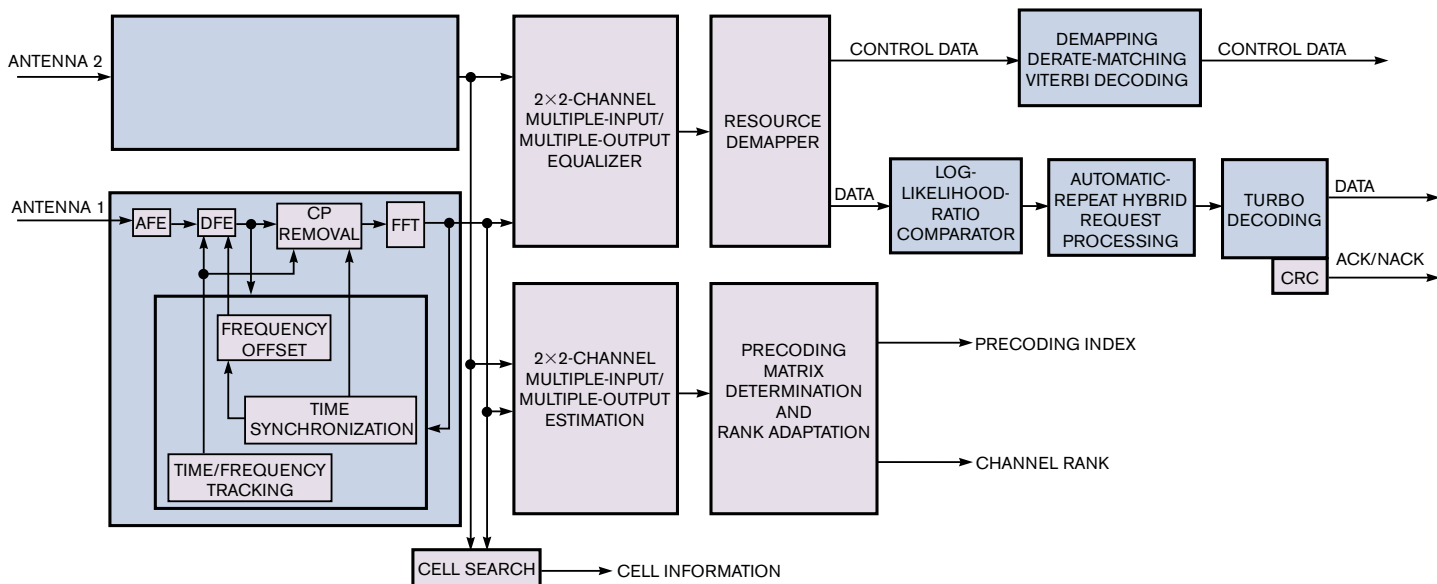
The 4G that application developers dream of may require the use of radical architectures.

OFDM (orthogonal-frequency-division-multiplexed) signal into its many constituent frequency bands. The frequency-domain signal then goes through further digital conditioning and into a detector—not unlike the read channel in a disk drive—that decodes the 64 QAM (quadrature-amplitude-modulation) signal on each of the carriers, producing a symbol from each active carrier. The symbols go through turbo decoding for decompression.

The difference between 3G and 4G in this architecture is a difference of

quantity, not kind. “In 3G, we extract about 1 bps per hertz of bandwidth,” points out Peter Carson, senior director of product management at Qualcomm CDMA Technologies. “To achieve 100-Mbps throughput, a 4G baseband would have to do significantly better than that: at least 3 or 4 bps per hertz, over a much wider band.”

In practice, this situation means many more carrier frequencies spread over a 20-MHz channel, compared, for instance, with the 5-MHz channel that UMTS 900 uses. It may also mean using multiple antennas in a MIMO (multiple-input/multiple-output) configuration. Today, MIMO configurations most often see use in channel equalization: You find a way to combine the signals from the two antennas to get the best possible reception. But 4G has something else in mind: using beam-forming algorithms to in effect make each pair of a base-station antenna and a receiver antenna into a separate channel, thus multiplying the effective bandwidth. “With multiple receivers, research has demonstrated, you can get about 1.75 times the data rate with two antennas,” Hansen says.



NOTES:
 INTERIOR OF ANTENNA 2 IS THE SAME AS THAT FOR ANTENNA 1.
 ACK/NACK=ACKNOWLEDGE/NO ACKNOWLEDGE.
 AFE=ANALOG FRONT END.
 CP=CYCLE PREFIX.
 CRC=CYCLIC-REDUNDANCY CHECK.
 DFE=DIGITAL FRONT END.
 FFT=FAST FOURIER TRANSFORM.

Figure 1 Infineon's vision of a downlink receiver for LTE shows the complexity of dealing with even two-antenna MIMO.

All of this capability requires silicon. The higher sample rate and wider channel mean a bigger, more power-hungry ADC and a faster, wider FFT engine. But the big hit comes from the need to provide for a 100-Mbps-peak throughput, which means faster symbol-rate processors, a lot more memory, and a faster processor for the MAC. “We are looking at 10 times the data rate coming into the MAC, with one-tenth the allowable latency on some transactions,” Hansen says. “But for power reasons, the MAC hardware has to run at a frequency much lower than the bit rate. This [problem] is interesting.”

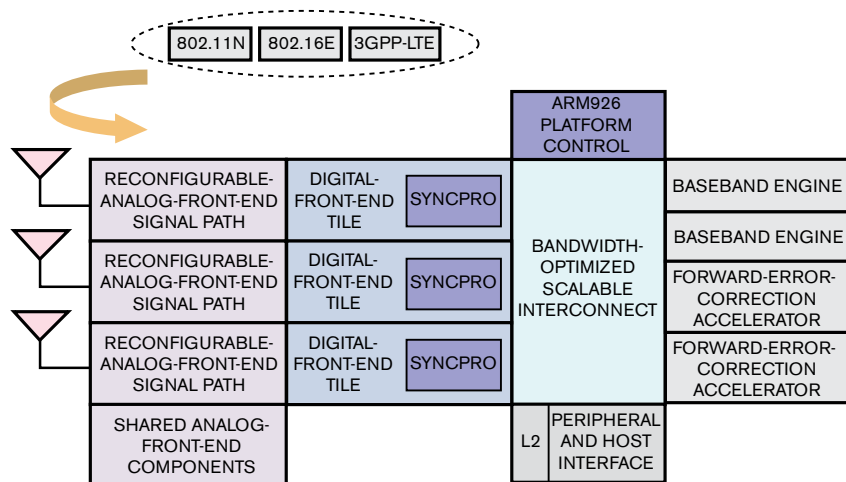
Qualcomm’s Carson agrees. “Peak data rate turns directly into die size. One thing architects will have to ask themselves is whether the specified peak data rate and the die size it requires will be justified in the average data rate the network will actually deliver.”

Given sufficient insensitivity to chip cost, the baseband architecture for this rate can be evolutionary. Carson says that Qualcomm’s current Snapdragon architecture is still perfectly manageable extended to 30- to 40-Mbps-peak data rates. That speed doesn’t meet the LTE specification, but LTE will come later—some call it late-term evolution—possibly allowing time for 32-nm CMOS to once again bail out the architects.

NONEVOLUTIONARY DESIGN

One of the first challenges to evolutionary architecture will come from MIMO. “MIMO is used to improve the quality of the wireless link,” explains Thuyen Le, PhD, of the Feature Phone Business Unit, Communication Business Group of Infineon Technologies AG. “One idea is to use it for transmitter and receiver diversity to combat fading. The other idea is to exploit fading for spatial multiplexing, which then allows transmitting independent data streams over the multiple transmitting antennas at the same time, hence increasing the user data rate. That [idea], however, depends on how well-conditioned the channel matrix is. So, my take is that MIMO is necessary to achieve high data rates, in light of both ideas.”

When air-interface designers shift from using a pair of receiving antennas to improve channel equalization to actually



NOTE:
3GPP-LTE=THIRD-GENERATION PARTNERSHIP PROJECT-LONG-TERM EVOLUTION.

Figure 2 IMEC researchers envision an array of on-the-fly reconfigurable resources shifting instantly between air interfaces and channel conditions as bandwidth needs change.

creating multiple channels through spatial-division multiplexing, the amount of duplicate hardware in the radio rises dramatically. Each antenna needs its own analog front end and digital front end, and the radio also needs either replication or increased throughput for much of the digital baseband (Figure 1). This requirement in itself does not mandate architectural innovation—just more of the same—but then there is the matter of power.

A limiting factor in any 4G architecture is that the radio must handle 10 times the peak data rate at a fraction of current energy consumption to make room for the dramatic increase in application-level processing energy. Will Strauss, president of research company Forward Concepts, estimates that a 4G handset will eventually require 100 times the computing power of a current 3G offering. “Everyone’s great hope is 32-nm processes,” Strauss observes, “but the reality is that energy consumption isn’t going down that much with new processes. What you gain in dynamic power you give back in leakage power. It may come down to finding novel architectures and power-management schemes or carrying a battery on your back for the handset.”

Another factor drives consideration of novel architectures, as well. It is the previously mentioned disparity between simply stating a peak data rate, as the

LTE specification does, and envisioning a new manner of using mobile devices, as do many of the visionaries who are evangelizing 4G to investors.

IMAGINING THE FUTURE

“It’s true there is no clear definition of 4G,” says Liesbet Van der Perre, science director at IMEC (Interuniversity Microelectronics Center). “But I believe we should be talking of a heterogeneous network supporting much higher mobility and data rates than are currently possible. Today, if you are truly mobile, you will see less than 2 Mbps, but 4G should mean 10 to 20 Mbps of real throughput. At least 10 Mbits sustained—not peak—is essential for good video, for instance. One of the disappointments of 3G is that it could not deliver the sustained data rate for good video.”

Van der Perre and other researchers describe an environment far more dynamic than anything that today’s wireless networks can realize. “Today, a handset-silicon vendor faces something like 30 air interfaces, multiple noncontiguous channels, and many very dif-

MORE AT EDN.COM

[Go to www.edn.com/071011df](http://www.edn.com/071011df)
and click on Feedback Loop to post
a comment on this article.

FOR MORE INFORMATION

Forward Concepts

www.fwdconcepts.com

Freescale Semiconductor

www.freescale.com

Infineon Technologies AG

www.infineon.com

IMEC (Interuniversity Microelectronics Center)

www.imec.be

Nokia Siemens Networks

www.nokiasiemens.com

Qualcomm CDMA Technologies

www.cdmatech.com

Texas Instruments

www.ti.com

3G Americas

www.3gamericas.org

3GPP (Third Generation Partnership Project)

www.3gpp.org

ferent services running simultaneously,” she observes. But the fact that one phone from one vendor supports only a small subset of this cacophony simplifies much of this complexity.

In the future, to ensure both sufficient sustained bandwidth—think of that real-time video aligned to the moving handset’s location and orientation—and sufficient energy efficiency—always choosing only just enough bandwidth and coding strength for the current task mix—a mobile device may be in continuous negotiation with a number of vendors, using a number of air interfaces from a number of base-station sites all at once (Figure 2). Bursts of data, video streams, control information, and a return channel from the keyboard and camera may all be traveling over different services and may switch in real time. For instance, holding the camera still allows the motion compensation in H.264 to drastically reduce the bit rate necessary to link the camera to a game-server farm. This action thus allows the radio controller to select an air interface with a lower bit rate.

In this world view, using today’s hardware-processing pipelines with dedicated blocks is an intermediate option, Van der Perre says. She sees modular, heterogeneous clusters of similar processors that you can specialize and a configurable interconnect network that can allow real-time dynamic processor configuration and task mapping. Aggressive energy-management techniques, including rapid voltage-frequency scaling,

moderately fine-grained power-gating of idle units, and agile shifting of algorithms between software and hardware, become possible in such an architecture. Indeed, this approach may be the only way to meet the energy-efficiency demands of the true 4G terminal, 32-nm CMOS notwithstanding.

All of these elements are taking shape at IMEC in various research projects, which perhaps explains Van der Perre’s world view. But it is far from an isolated point of view, at least in private. Companies publicly state their dedication to their pipeline-based hardware architectures, but one well-placed industry source claims that there are deeply embedded, heavily funded research teams at a number of major silicon suppliers exploring large multicore architectures for the 4G problem.

4G MAY BE THE FORGE ON WHICH DESIGNERS BEAT AN ENTIRELY NEW STYLE OF REAL-TIME EMBEDDED PROCESSING INTO SHAPE.

One major challenge with most large, multicore architectures is not an issue here: Much of the workload in high-bit-rate baseband processing is what the industry calls embarrassingly parallel. It’s not hard to spread around the tasks by simply dividing up the data. But the system-control, dynamic-load-balancing, and—perhaps most important—energy-management tasks are new, complex, and vital to the success of the design. In this respect, 4G may in fact not be evolutionary but rather the forge on which designers beat an entirely new style of real-time embedded processing into shape. **EDN**

You can reach
Executive Editor
Ron Wilson at
1-408-345-4427 and
ronald.wilson@reedbusiness.com.

