Greg Nilsen

Prof. Ellen Walker

Hiram College

VLSI Design

10 May 2000

CMOS and VLSI for Wireless Communication

The complimentary metal-oxide semiconductor, or CMOS, is the mechanism used in the transistors that are found in many microchips. By controlling the flow of electrons in a material (typically silicon) a transistor, combining positive and negative gates in order to minimize power usage, is created. This technology has allowed for the rapid development of the very large-scale integration (VLSI) processors that are in the majority of the computers of today, as well as many other electronic devices. The technology has evolved over the years, and now eyes are looking towards wireless communications. Three areas of work are most prominent in this area – digital phone and mobile data networks, wireless local area networks, and transistor optimization.

In 1970, at a time when bipolar transistors were thought to be the wave of the future, Intel developed a NMOS chip that contained 2,000 transistors. Maurice Wilkes, 1967 ACM Turing Award winner, states that this may “mark the beginning of very large-scale integration” (16). The technology then switched to CMOS where the power dissipation was much smaller than in either NMOS or bipolar technologies (Hendrich). With the use of the CMOS innovation, the most advanced chips created today have as many as 9.3 million transistors on them. However, there are still some features of the bipolar technology that are better for current applications than the CMOS form. This has lead some researchers to work towards developing an integration of the two called BiCMOS. Each of these forms has become a player in the wireless communication field of today.

Digital phone and mobile networks are radio-frequency (RF) networks that are primarily used for voice communications, as well as some limited data utilities. They are designed to handle many users over a wide area at relatively low data rates. These networks are typically used for phone systems, such as the PCS Network developed by Sprint. Though some signal format work is necessary, most of the concentration is on the actual communications units. The goals of these units are “minimum power consumption for battery autonomy, minimum silicon area for maximum functional integration per die to obtain a small, low cost pocket size radio telephone” (Haspeslagh, 71).

There are four main components that go into the construction of a radiotelephone transceiver. The first of these is the low noise radio receiver. In this unit, “one or two cascaded mixer stages bring the radio signal down from RF to base band with sufficient gain to further process the analog receive signal in noisy standard CMOS technology” (Haspeslagh, 71). The CMOS is also used in order to provide for a very low power consumption measure. The unit signal then goes through a layer of CMOS filtering in order to properly process the base band signal.

Depending on the arrangement of the different processes and the techniques used in the receiver, one of two different types can be created. A heterodyne receiver has a local oscillator frequency that separates itself more easily from the other bands on the receiver. However, the filters needed to separate these frequencies have encountered many difficulties in their integration with rest of the unit. The other type, a homodyne receiver, lends itself to easy filtering with CMOS filters. However this receiver does cause an offset, but this problem has been overcome using an offset subtraction CMOS filter (Haspeslagh, 72-4).

The next portion of the transceiver is the voice codec. Using an interface already in place with the existing analog interface found in regular telephones, a simple converter into the binary signal is used. Currently this process and the rest of the processing within the phone are contained on separate chips within the telephone unit. However, “nothing but the silicon area per die prevents us from integrating the voice interface on a single chip with the base band CMOS radio transceiver" (Haspeslagh, 74).

The third of the four portions of the transceiver is the low noise radio transmitter. This unit picks out the noises from the stream and filters out which ones need to be sent based on how 'strong' they are. The part also has noise constraints in order to prevent the phones from jamming one another and the receivers of the base stations for distant users (Haspeslagh, 74). The problem with this unit is that the power amplifier is inefficient and expensive. Chip designers are trying to alleviate this problem with new designs in silicon and gallium-arsenic (Haspeslagh, 75).

The final basic portion of the transceiver is the local oscillator synthesizer. These oscillators currently rely on "temperature controlled reference crystal oscillators" to provide the local oscillator signal (Haspeslagh, 75). However, this does not seem to be the most accurate method, and CMOS is being looked into to provide a replacement.

From the descriptions of the parts of these radiotelephone transceivers, CMOS is an essential part of their design. Even in portions where CMOS use is limited, it is being looked into to provide an improved product. The fact that CMOS transistors use almost no power when not needed keeps the power consumption of the processor low (Rosch). Those reasons, as well as the decreasing silicon area per die, are main advantages of CMOS in the area of digital phone and mobile networks. The future in this area seems to belong to CMOS as researchers have seen that they are only scratching the surface with its capabilities. However, the successor to CMOS, BiCMOS may take over the field in the future because of its combination of high speed and memory capabilities.

Wireless local area networks, or WLANs, are an extension of the existing physically connected LAN architecture. Like the digital phone and mobile networks, the wireless LANs run on radio-frequency schemes. However, these networks are designed to have fewer users in a local area, transmitting at high data rates. Ideally these wireless LANs should be seamlessly integrated with their existing wired counterparts, with matching speeds and hybrid connections.

Bongjin Jung and Wayne Burleson tell us:

> In contrast to wireline communication, the physical bandwidth of RF wireless communication systems is relatively limited and is unlikely to grow significantly in the future. Hence it is advantageous to increase the effective bandwidth of communication channels at the expense of complex processing at both the sending and

> receiving entities…to improve the performance of wireless local area networks (27).

An existing LAN has a fairly broad range of bit transfer rates, ranging from 10 Mbps in Ethernet to 100 Mbps in FDDI. The frequency bandwidth is limited to begin with, but wireless LANs require the use of a spread spectrum, which leaves the actual data bandwidth being less than one tenth of the original. For this reason, efficient bandwidth use is a must in order to obtain even 10 Mbps speeds over the wireless LAN.

One method of making bandwidth use more efficient is loss-less data compression (Burleson, 27). As the name implies, this is a method of transmitting data without losing any information. This is done because in the case that some of the transmitted data is lost, then bandwidth is wasted. This is unlike the digital phone and mobile networks that have a high tendency to send excess information in order to assure that enough information reaches its destination. That is due to the wide area nature of those types of networks.

In order to process this loss-less data compression, real-time compression and decompression is necessary (Burleson, 27). This would be done to minimize the network delay. If the network delay were to become too high, then performance would be affected and data loss would become imminent as the receiver ran out of space to hold incoming transmissions.

Unlike wired LANs, whose bandwidth seems to be almost limitless, the bandwidth of wireless LANs is the bottleneck for high-speed data connections. However, with the use of VLSI architecture, this bottleneck becomes much less of a problem. Jung and Burleson present a low-power VLSI architecture that achieves a high compression rate, which they state is “well above the data rate required by any current or foreseeable future wireless LAN” (38). In fact, they have been able to achieve an average compression rate of 50 Mbps, which stacks up very well against existing wired LAN systems (Burleson, 27).

While looking to improve the capabilities of the wireless LAN, these methods also look for cost- and power-conscious ways to improve the technology. But then, when has it been good to take a step forward in one direction and one backward in another direction?

Another area of focus in wireless communication is on transistor optimization. Though this area is not unique to the wireless domain, it is extremely important because of all the power that researchers and companies are trying to put into very small, usually handheld, devices. Since many devices contain several small VLSI chips, it would be advantageous to optimize the transistor so they could all be fit onto one VLSI chip. The discipline works to minimize the area each transistor occupies, as well as the power consumption while minimizing the delay on the chip.

The transistor optimization is done for several reasons. One reason is that by minimizing power consumption, it allows devices to be used for more extended periods. The optimization can also increase the efficiency of the chip. If the transistors are arranged in a way that optimizes them and lowers the delay of the chip, then the chip becomes much more efficient at processing data. By putting more power into less space, the market is potentially opened up for more devices that become increasingly powerful.

Transistor optimization is done using technology-mapping algorithms. These algorithms rearrange the layout of the transistors on the VLSI chip in ways that try to minimize the area occupied and power consumed while ensuring that the delay does not exceed an upper limit (Prasad, 281). According to Prasad and Roy, “the existing algorithms do not take power consumption into account and/or require too long a run-time to be used on the large circuits of today” (281).

That statement came from the duo in 1996. That same year they introduced a new algorithm for transistor reordering. This method “is used to select an ordering of series-connected transistors found in CMOS gates to achieve lower power consumption” (Prasad, 281). Where most algorithms are based on iterative improvement, theirs is based on dynamic programming for computing the longest path in the circuit along with circuit optimization using gate-input-reordering. The combination of these two methods results in a more efficient and more effective algorithm for transistor reordering (Prasad, 280-2).

Currently, the field of wireless communications is a very busy and volatile one. There are many areas of research going in many different directions, but they still are working together to improve one another. Someday there will probably be an integrated phone and network system in place that utilizes the transistor optimizations. However, today is not that day and the different areas are still working towards more specialized goals, and it is very likely that they will be recombined in the future. This gives the potential of having an integrated system, which functions better than what exists today, without all the wires. Wherever the technology does lead our society, it is not difficult to see the promise of the wireless communications field.

The use of the CMOS and VLSI technologies in these areas is abundant. These tools have allowed for the rapid development in the field, and will continue to do so in one form or another, whether it is CMOS, BiCMOS or some other new technology that will no doubt be based on the existing techniques in some way.

Still, the use of these technologies is very important to the field and should continue to be far into the future.

There is almost no limit to the possibilities that can be seen in the wireless realm. And in all that, CMOS and VLSI have been key players. Thanks to the use of these technologies, the wireless field will be able to complete its diamond of evolution – taking the thing apart to find which parts need to be improved, and then taking those improvements and putting them all back together.

Works Cited

Networks through VLSI data compression”. Wireless Networks Volume 4. 1998: 27-39. ACM Digital Library. On-line. ACM. 23 Apr. 2000.

Haspeslagh, D., J. Sevenhans, and J. Wenin. “Wireless telecom silicon integration: analog design for radio, baseband and speech spectrum”. Wireless Networks Volume 4. 1998: 71-77. ACM Digital Library. On-line. ACM. 23 Apr. 2000.

Hendrich. “CMOS technology demonstration” (22 Apr. 1999): n.pag. On-line. Internet. 30 Apr. 2000. Avaliable WWW: http://tech-www.informatik.uni-hamburg.de/applets/cmos/cmosdemo.html.

Prasad, S.C. and K. Roy. “Transistor Reordering for Power Minimization Under Delay Constraint”. ACM Transactions on Design Automation of Electronic Systems Apr. 1996: 280-300. ACM Digital Library. On-line. ACM. 23 Apr. 2000.

Rosch, Winn L. The Winn L. Rosch Hardware Bible (Third Edition). 1994. On-line. Whatis.com. 30 Apr. 2000. Available WWW: http://www.whatis.com/cmos.htm

Wilkes, Maurice V. “The Rise of the VLSI Processor”. Communications of the ACM Dec. 1990: 16-17, 44. ACM Digital Library. On-line. ACM. 23 Apr. 2000.