Distributed Integrated Circuits: Wideband Communications for the 21st Century

by Ali Hajimiri

Global communications have rendered our world a smaller, yet more interesting place, making it possible to exchange visions, ideas, goals, dreams—and PoKéMoN cards—across our small planet. Modern communications systems, such as the internet and portable wireless systems, have added new dimensions to an already complex world. They make us aware of our similarities and differences and give us an opportunity to communicate with people we have never met from places we have never been. The fusion of education with communication is already bringing about new levels of awareness, accompanied by creative upheavals in all aspects of modern life.

However, the ever-increasing demand for more connectivity inevitably increases the complexity of such systems. Integrated systems and circuits continue to play a central role in the evolution of component design. Silicon-based integrated-circuit technologies (particularly complementary metal oxide semiconductor, or CMOS) are the only technologies to date capable of providing a very large number (over a

Ali Hajimiri

million) of *reliable* active (e.g., transistors) and passive (e.g., interconnect) devices. Further, they are relatively inexpensive to incorporate into mass-market products.

The realization of revolutionary ideas in communications depends heavily on the performance of the integrated electronic circuits used to implement them. Let's consider some well-established theoretical background for a moment. The maximum number of bits (1s and 0s) that can be trans-

mitted per second (i.e., *bit rate*) determines the speed of a digital communications system. C.E. Shannon, the founder of modern information theory, proved that the maximum achievable bit rate of a digital communications system increases linearly with the available range of frequencies (i.e., *channel bandwidth*) and logarithmically with the signal-to-noise ratio. Thus, three critical parameters, namely, *bandwidth, signal power,* and *noise*, are the most important parameters in determining the performance of any given communications system.

One of the more common methods of increasing the bandwidth, and hence the bit rate, of any given system is to migrate to higher operating frequencies. The maximum speed of operation in electrical systems is determined by the performance of both active and passive devices. While in modern integrated-circuit technologies the single-transitor maximum frequency of operation can be quite high, actual circuits rarely operate anywhere near these frequencies.[1] This provides further motivation to pursue alternative approaches to allevi-

---

[1]   A transistor in a given process technology is usually characterized by its unity-gain frequency shown as f_T. This is the frequency at which the current gain (the ratio of the output current to input current) of a transistor drops to unity. While the unity-gain frequency of a transistor provides an approximate measure to compare transistors in different technologies, the circuits built using these transistors scarcely operate close to the f_T and usually operate at frequencies 4 to 100 times smaller depending on the complexity of their function. There are two main reasons for this behavior. First, many systems rely on closed-loop operation based on negative feedback to perform a given function independent of the parameter variations. An open-loop gain much higher than one is thus required for the negative feedback to be effective. This higher gain can be only achieved by operating the transistors at a lower frequency than the f_T to provide the desired gain. Second, integrated passive devices necessary in most of the high-speed analog circuits have their own frequency limitations due to parasitic components that can become design bottlenecks. The combination of these two effects significantly lowers the maximum frequency of reliable operation in most conventional circuit building blocks

ate bandwidth limitations, particularly in silicon-based systems which, despite their reliability, suffer from low transistor speed, poor passive performance, and high noise compared with other technologies.

The complex and strong interrelations between constraints in modern communications systems have forced us to reinvestigate our approach to system design. "Divide and conquer" has been the principle used to solve many scientific and engineering problems. Over many years, we have devised systematic ways to divide a design objective into a collection of smaller projects and tasks defined at multiple levels of abstraction artificially created to render the problem more tractable. While this divide-and-conquer process has been rather successful in streamlining innovation, it is a double-edged sword, as some of the most interesting possibilities fall in the boundary between different disciplines and thus hide from the narrow field of view available at each level. Thus, approaching the problem across multiple levels of abstraction seems to be the most promising way to find solutions not easily seen when one confines the search space to one level.

Distributed circuit and system design is a multi-level approach allowing more integral co-design of the building blocks at the circuit and device levels. This approach can be used to greatly alleviate the frequency, noise, and energy efficiency limitations of conventional circuits. Unlike conventional circuits, which often consist of a sin-
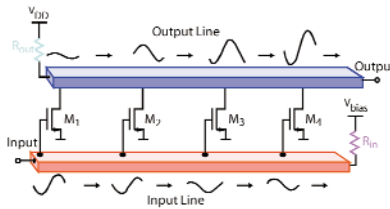

Figure 1. A distributed amplifier consisting of two transmission lines and multiple transistors providing gain through multiple signal paths that amplify the forward traveling wave. Each transistor adds power in phase to the signal at each tap point on the output line. Each pathway provides some gain and therefore the whole amplifier is capable of providing a higher gain-bandwidth product than a conventional amplifier.

gle signal path, distributed integrated systems and circuits rely on multiple parallel paths operating in harmony to achieve an objective. However, this multiple signal-path feature often results in strong electromagnetic couplings between circuit components, which makes it necessary to perform the analysis and the design of distributed circuits across multiple levels, a task not crucial when using the "divide and conquer" approach.

This concept can be best seen through the distributed amplifier (originally suggested by Percival and first implemented by Ginzton) sketched in Figure 1. This amplifier consists of two transmission lines on the input and the output, and multiple transistors providing gain through multiple signal paths. The forward (to the right in the figure) wave on the input line is amplified by each transistor. The incident wave on the output line travels forward in synchronization with the

traveling wave on the input line. Each transistor adds power in phase to the signal at each tap point on the output line. The forward traveling wave on the gate line and the backward (traveling to the left) wave on the drain line are absorbed by terminations matched to the loaded characteristic impedance of the input line, R_{in}, and output line, R_{out}, respectively, to avoid reflections.

In a distributed amplifier, one tries to avoid a "weakest-link" situation by providing multiple, equally strong (or equally weak) parallel paths for the signal. In the absence of passive loss, additional gain can be achieved without a significant reduction in the bandwidth by addition of extra transistor segments. This is the direct result of multiple signal paths in the circuit. The extended bandwidth of the distributed amplifier comes at the price of a larger time delay between its input and output, as there is a trade-off between the bandwidth and delay in an amplifier. Alternatively, one can think of this approach as a method of absorbing the parasitic capacitances of the transistors into the transmission lines and making them a part of the passive network.

A t Caltech, one of our most exciting breakthroughs has been in the area of silicon-based distributed circuits for communication systems; we have achieved unprecedented performance for communication blocks and systems.

In particular, we have used the concept of distributed systems to demonstrate an extremely high-speed voltage-controlled oscillator using a low-performance CMOS technology with small cut-off frequencies for the active and passive components (see Figure 2). This oscillator uses the delay introduced by the distributed amplifier to sustain electrical oscillation by continuous amplification of the signal around a loop. The oscillation frequency is determined by the round-trip time delay, i.e., the time it takes the wave to travel through the transmission lines and get amplified by the transistors.

Tunability is an essential feature for such distributed voltage-controlled oscillators (DVCOs), and thus it is necessary to devise a method to control the oscillation frequency. The oscillation frequency is inversely proportional to the total delay and hence the total length of the transmission lines. This property leads to a frequency tuning approach based on changing the effective length of the transmission lines. Frequency control can be achieved by introducing shortcuts in the signal path. This concept can be seen using the racetrack analogy of Figure 2a. Here the signals traveling on the input and output lines are analogous to two runners on two tracks running side-by-side to be able to pass a torch at all times. The time it takes them to complete a lap (oscillation period) can be changed by introducing symmetrical shortcuts for both of them and controlling what percentage of the time they go through each one.

This concept has been successfully demonstrated in the distributed voltage-controlled oscillator of Figure 2b where alternative signal paths have been introduced to change the electrical length seen by the traveling wave.

Another component we have devised is the distributed active transformer (DAT) power amplifier. The design of a fully integrated silicon-based power amplifier with high output power, efficiency, and gain has been one of the unsolved major challenges in today's pursuit of a single-chip integrated communication systems. Although several advances

Figure 2a. The racetrack analogy.

Figure 2b. A die photo of the 10-GHz distributed voltage-controlled oscillator using 0.35-μm CMOS transistors with a tuning range of 12% and phase noise of –114dBc/Hz at 1-MHz offset. In addition to higher frequency of oscillation, the DVCO provides better frequency stability.

have been made in this direction, a watt-level, truly fully integrated CMOS power amplifier has not been demonstrated using the traditional power-amplifier design techniques.

T wo main obstacles in the design of a fully integrated power amplifier are the low breakdown voltages of transistors and the high losses of passive components. The low breakdown voltage limits the voltage swing at the output node, which in turn lowers the produced output power. The high passive loss reduces the amplifier's power efficiency by dissipating the generated power in the signal path. These problems are exacerbated in most commonly used CMOS process technologies, as the MOS transistor's minimum feature size is continuously scaled down for faster operation, resulting in lower substrate resistivity and smaller breakdown voltages.

Our DAT power amplifier uses the distributed approach to perform impedance transformation and power combining simultaneously to achieve a large output power while maintaining acceptable power efficiency. It overcomes the low breakdown voltage of short-channel MOS transistors and alleviates the substrate loss problems by providing the power gain through multiple similar stages and signal paths.

Figure 3a shows the essential features of the DAT, which consists of multiple distributed push-pull circuits in a polygonal geometry.
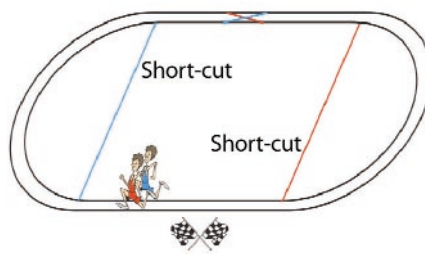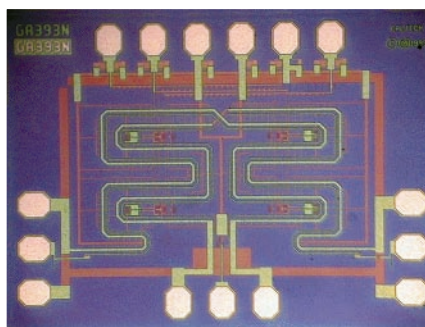
Each side of the square is a single amplifier consisting of a transmission line, two transistors, and input matching lines. This particular positioning of the push-pull amplifiers makes it possible to use a wide metal line as the drain inductor to provide natural low-resistance paths for the dc and ac currents to flow.

The four transmission lines are used as the primary circuit of a magnetically coupled active transformer. The output power of these four push-pull amplifiers is combined in series and matches their small drain impedance to the load. These four push-pull amplifiers, driven by alternating phases, generate a uniform circular current at the fundamental frequency around the square, resulting in a strong magnetic flux through it. A one-turn metal loop inside the square is used to harness this alternating magnetic flux and acts as the transformer secondary loop. This is where multiple signal paths converge. Using the DAT, a fully integrated watt-level power amplifier was demonstrated in a standard CMOS process technology for the first time, as shown in Figure 3b. The distributed nature of the DAT structure reduces the sensitivity of the power amplifier's efficiency to

the substrate power losses while providing a large overall output power using low-breakdown-voltage MOS transistors. The strong electromagnetic coupling between

Figure 3a. The essential features of our distributed active transformer (DAT).

Figure 3b. A die photo of the 0.18-μm CMOS DAT generating 3.6 W of power at 1.9 GHz into a 50-Ω load using a 1.8-V power supply, while achieving a power added efficiency of 51%. It is fully self-contained and uses no external components. The distributed nature of the DAT structure reduces the sensitivity of the power amplifier's efficiency to the substrate power losses while providing a large overall output power using low-breakdown-voltage MOS transistors.

multiple signal paths in a DAT necessitates an analysis and design approach spanning architecture, circuits, device physics, and electromagnetics.

These examples demonstrate some of the basic concepts of distributed integrated circuit design. The combination of multiple distributed signal paths working in harmony and a design approach covering several levels of abstraction allow us to achieve higher frequencies of operation, higher power and efficiency, while creating more robust systems.

Bringing this state-of-the-art technology into the commercial realm, substituting easily mass produced silicon-based circuits for the traditional GaAs-based circuits in use today in everything from cell phones to communications satellites, will further the revolution in communications systems that defines our modern era. E N G

Ali Hajimiri is Assistant Professor of Electrical Engineering.
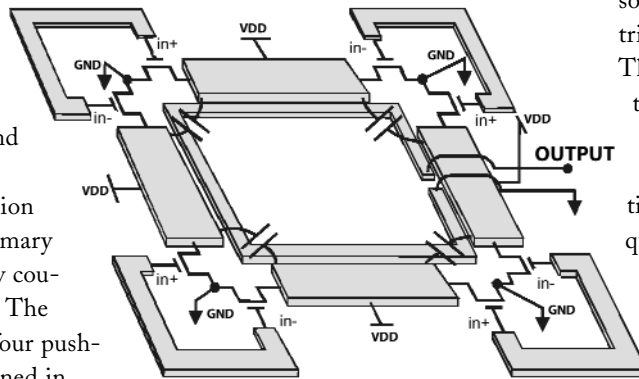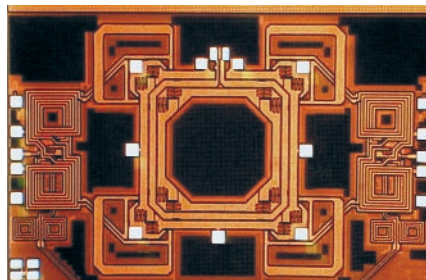
There is more on Professor Hajimiri at
http://www.chic.caltech.edu