Look Up, Look Down, Look All Around

By Douglas L. Smith

Along with GPS receivers and DVD players, some luxury cars these days come with autonomous cruise control, which measures the distance to the car in front of you and automatically eases up on the gas if you start to get too close. Some versions will even hit the brakes for you. But they don't come cheap—Mercedes-Benz's Distronic Plus, for example, which incorporates two onboard radar units that stare straight ahead to different distances, costs about three grand. "A radar system consists of thousands of individual parts," says Professor of Electrical Engineering Ali Hajimiri. "A lot of that complexity arises from the fact that you are trying to pull information from one module to another one at very high frequencies through interconnects. Once you put everything on the same chip, you can actually eliminate a lot of the complexity." Hajimiri has done just that, creating a complete radar system that fits on the head of a thumbtack. Since we can put millions of transistors on a chip these days, the unit contains all the supporting electronics needed to send, receive, and process the radar signal, and even to sweep the beam back and forth. The chip was made with standard industrial processes, and in mass production would cost about a buck each, Hajimiri estimates.

The key, he says, was "a different approach. You can't just take a traditional, module-based design and simply transplant all of the modules onto the same chip. All of the elements had to be designed for this purpose. Very few people thought that solid-state technology could be used at very high frequencies, but we found ways of doing so and then combined it with digital signal processing. Up to then, integrated phased-array radar had been one of the last bastions of analog signal processing."

From left: Grad students Jay Chen, Edward Keehr, Aydin Babakhani (MS '05), Yu-Jiu Wang (MS '06), professor Ali Hajimiri, grad students Juhwan Yoo (BS '06), Florian Bohn (BS '01), Jennifer Arroyo, Hua Wang, and research engineer Sagguen Jeon (MS '04, PhD '06) create a phased array in Millikan Pond.

GUIDED RIPPLES, FOCUSED POWER

That the system is a "phased-array" radar means that, instead of being a dish that tilts, in the simplest case it's a line of dipole antennas, each one like an old-fashioned portable radio antenna. Such an antenna radiates its signal in every direction, like the ripples from a rock dropped from a bridge over a pond. But say that bridge had a line of people, armed with sticks instead of rocks, leaning over the railing. If one person poked the water, it would make a set of ripples as before. But if the next person poked the water an instant later—slightly out of phase, in other words—two sets of ripples would form. The overlapping ripples would add up in some directions and cancel out in others, making a new set that would travel at an angle determined by that fractional hesitation between pokes. Now, if everybody did this sequentially from left to right, in a sort of inverted version of the wave you see at sporting events, you could get a sizeable set of swells going, and if you altered the poke interval you could actually steer the waves from one bank of the pond to the other. Similarly, in radar, a phased array's beam is steered electronically by adjusting the phase of the signal at each antenna. There are no moving parts. "I can do that very fast, on the order of a nanosecond," says Hajimiri. "Or even faster, as the technology improves. I can have the beam pointing somewhere completely different the next nanosecond. No mechanical part can move nearly as fast."

A line of antennas can steer the angle of a beam across a plane, but a two-dimensional array can be steered in every direction, sweeping out the hemisphere in front of its radiating surface. Three-dimensional arrays are even possible, says Hajimiri. "Two-dimensional arrays generally make nice, narrow beams when you look straight ahead. But when you look sideways past a certain point, the beam widens. A three-dimensional array is uniform all over the place. And the three-dimensional array would generate a lot more power, because of all the antennas and the signal generators attached to them."

The phased-array concept has been around for quite a while—some antiballistic missile early warning radars of the 1960s used it, and the U.S. Navy's Aegis shipboard system has been operational since the '80s. But, says Hajimiri, "they're mostly used in airborne radar, because as you can imagine, you don't want any moving parts in an airplane that's going to pull 9 and 10 g's. As soon as you turn, you need a new antenna." A typical radar array in a fighter jet, nose-mounted to scan the skies in front of the plane, is about a meter square, several centimeters thick, and contains several thousand tiny antennas.

A tiny antenna puts out a tiny signal, but lots and lots of tiny antennas, properly synchronized, put out power in proportion to the square of their numbers. Even a piddling four-antenna array produces a sixteenfold increase in power over a single one—in the desired direction. It's like the difference between an ordinary light bulb and a laser—a classroom laser pointer is five milliwatts, or 20,000 times less powerful than a 100-watt desk lamp, but you don't need to point it at your eye to know which beam is brighter. This collective boost is vital, Hajimiri explains, because "the smaller the antenna, the less power, and the more antennas you need. So there's a tradeoff."

These airborne radars currently use expensive compound semiconductors, such as gallium arsenide microwave monolithic integrated circuits, which we civilians have in our cell phones' power amplifiers. Each module has one antenna and all its supporting electronics—the transmitter's power amplifier, the phase-delay controller, the low-noise receiver, and the receiver's gain control. The formidable problem of synchronizing the phase delays is handled by separate, highly complex (and very expensive) modules. These radars broadcast in the microwave band of the spectrum—whose wavelengths, confusingly enough, are actually a few centimeters long—at frequencies of around 10 gigahertz (GHz), or 10 billion cycles per second. Trying to coordinate a set of fixed delays for, say, a radar that always looks down toward the ground at a 60-degree angle to track the terrain, is tricky enough. If you want the beam to sweep, the system has to calculate variable time adjustments finer than a fly's eyelash all across the array. The slightest jaggedness, and the wavefront dissolves in chaos, like the din at a family reunion where everyone's talking at once. Internal travel times become critical, says Hajimiri. "If you try to connect parts with cables or leads, they have to match to hundredths or thousandths of a centimeter, and their lengths shouldn't change with temperature or variations in the electronics."

The Soviet Union's MiG-31 "Foxhound," which entered service in 1983, was the world's first production aircraft with an electronically scanned phased-array radar. With a forward range of 200 kilometers, it could track 10 targets simultaneously and engage four.

Courtesy of the Defense Visual Information Center

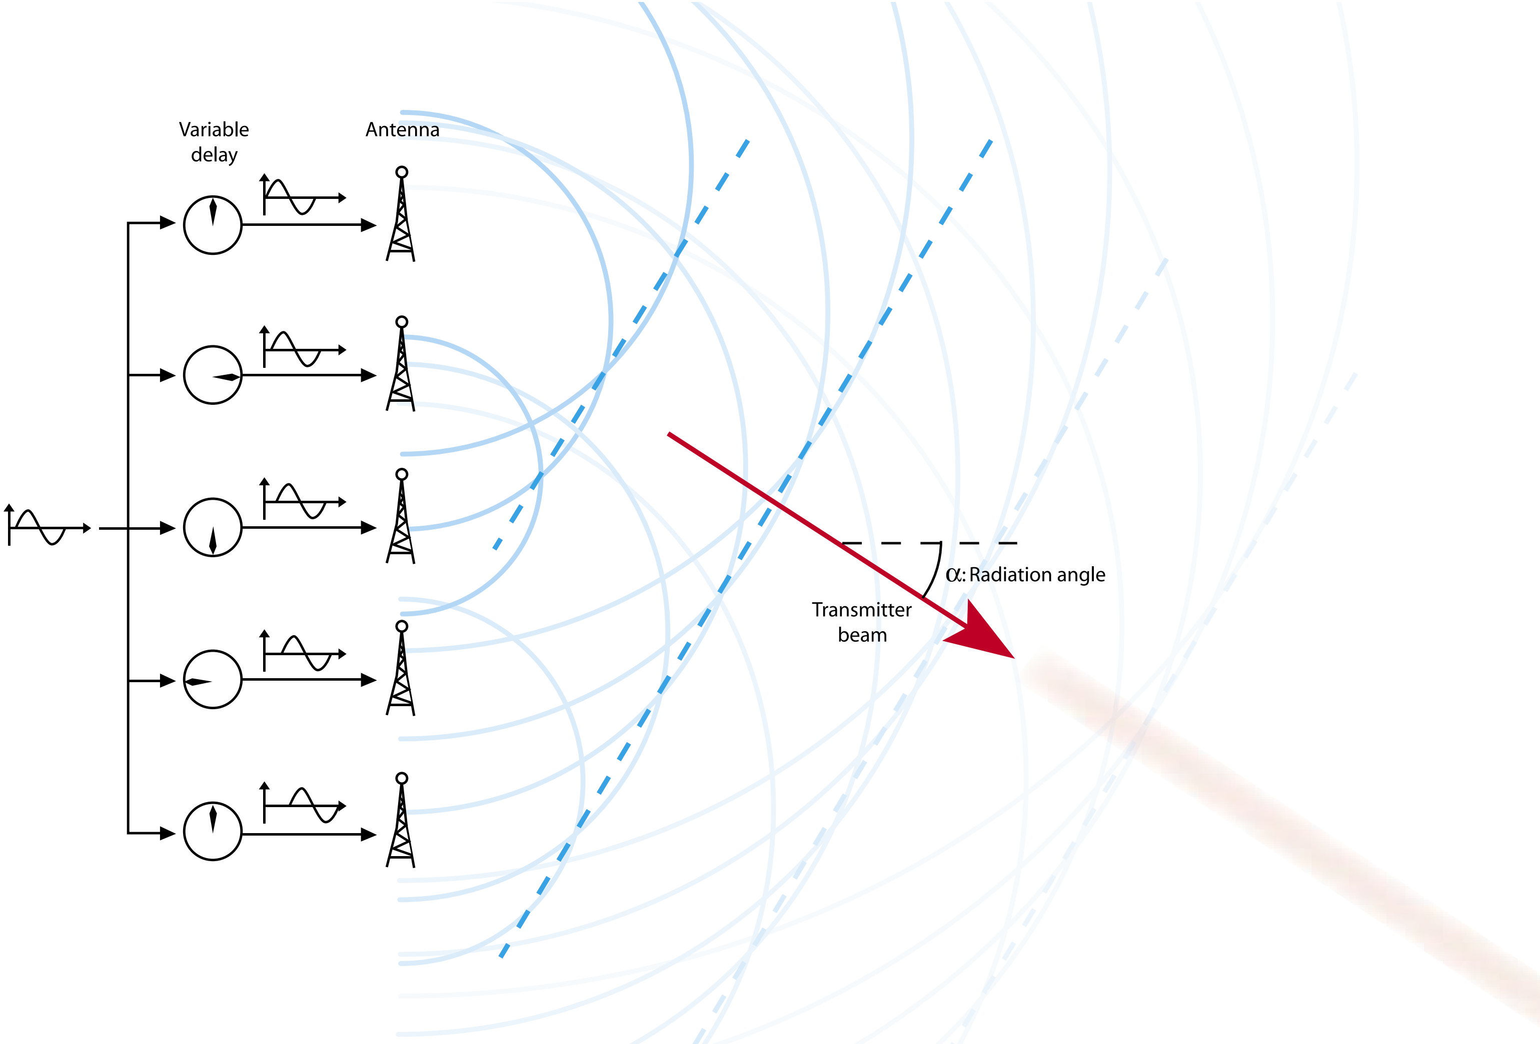

STEERING CLEAR OF BROADBAND CLUTTER

But the biggest civilian market for phased arrays is most likely going to be in broadband communications. The Federal Communications Commission recently opened up several bands, including the 24 and 60 GHz bands, for wireless communications systems. You might not think you could *get* any more networked, but just wait until your kitchen appliances need to talk to your Blackberry so that they can get your breakfast going super-early on the day of your big meeting. And when your microwave begins downloading Jon Stewart so that you can have something to stare at besides the instant oatmeal on the carousel (sure, you laugh now), we're gonna need a lot of bandwidth. Or consider the wireless office, with everybody's gadgets talking to each other all at once—as if we didn't have enough office clutter already.

The theoretical upper limit for information transfer at any given frequency increases with the broadcasting power. (Actually, to be accurate, it increases with the signal-to-noise ratio.) "At 24 GHz and 60 GHz we can get up to several gigabits per second on a few milliwatts of power over several tens of feet," says Hajimiri. "Your dial-up line is some 50 kilobits per second and your DSL is probably at best around one to two megabits per second. So this is another four orders of magnitude. If you can communicate at a couple of gigabits per second, you can transmit the entire contents of a DVD in about 10 seconds."

In such crowded airwaves, a steerable two-way communications beam has obvious advantages. Even a two-element phased array beats a single antenna, as anybody who has ever been to a loud party knows. Says Hajimiri, "If you're politely listening to a conversation which is getting quite boring, and there's a juicier conversation on the other side, you can tune into it while still nodding and looking like you're listening. The optimal antenna spacing in a phased array is half the wavelength, and for one kilohertz, which is kind of the middle of the audible frequency range, this works out to about 15 centimeters, which is more or less how far apart your ears are."

A phased-array cell phone would reduce the need for more ugly towers, and spare us from the feeble attempts to disguise them as trees. Each antenna in the array would pick up the incoming signal at a slightly different time. With the right delay at each antenna, all the signals coming from a certain angle would be in phase, amplifying themselves. But signals at the same frequency arriving from other directions would be out of phase, "and in fact they can cancel each other out if you design the delays right."

A phased-array transmitter (opposite page) takes an incoming signal (the sine wave at far left) and broadcasts it from each antenna in the array at a slightly different time. The sets of radiated waves from each antenna interfere with one another, creating a new set of waves that travels at an angle to the array. A phased-array receiver (below right) uses a similar delay to reassemble the original signal, amplifying it in the process.

Turning Silicon to Gold

"Silicon is truly today's alchemy," says Hajimiri. "It's a way of turning sand into money, quite literally. My philosophy is simple: if you can do anything in the digital domain, it should and will be done in the digital domain. And anything that can be done in silicon will be done in silicon. And most particularly, anything that can be done in CMOS should and will be done in CMOS." CMOS, for Complementary Metal Oxide Semiconductor, is an integrated-circuit manufacturing process that permits the fabrication of billions of transistors on a chip with a very high probability of all of them working. It is what makes today's PCs possible.

Doing everything on one chip gave the Hajimiri group a leg up on the thorniest problem—how to adjust the delay between antennas. The chip synchronizes the delays with a master clock circuit called a local oscillator, which can be built as a ring of amplifiers around which a pulse of voltage chases itself. Each lap takes one wavelength to execute, so by choosing the point in the loop where each antenna draws its time signal, you can steer the beam—in Hajimiri's case, by increments of 7.2 degrees. Says Hajimiri, "It's impossible to implement a local oscillator phase delay in a module-based architecture, due to the inevitable variations in the properties of the components and the off-chip interconnections among them."

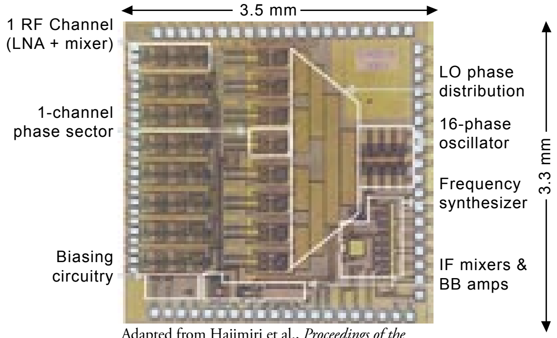

Adapted from Hajimiri et al., *Proceedings of the IEEE*, Vol. 93, No. 9, Sept. 2005, pp. 1637–1655.

The lab's first phased-array chip received signals at 24 GHz from eight off-chip antennas.

Hossein Hashemi (MS '01, PhD '04, now an assistant professor at USC) and Xiang Guan (MS '02, PhD '06) and Hajimiri created the lab's first successful device, built in 2003 and premiered at the annual International Solid-State Circuits Conference (ISSCC) in San Francisco in February 2004. This receive-only chip contained all the electronics needed to collect, amplify, and combine incoming signals in the correct phase—using the same local-oscillator concept, but in reverse. The

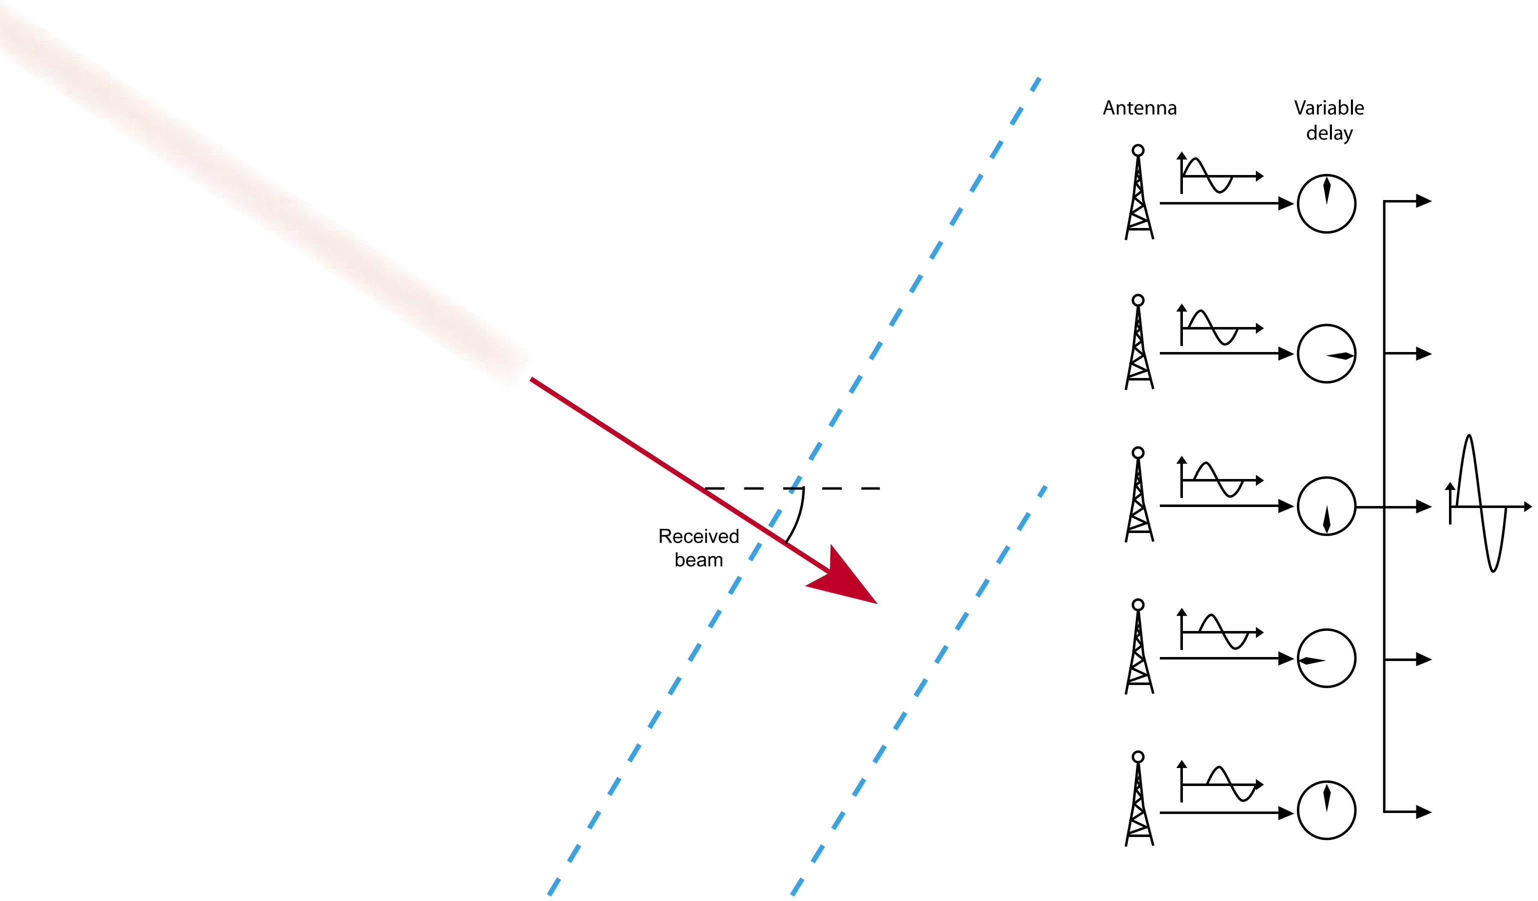

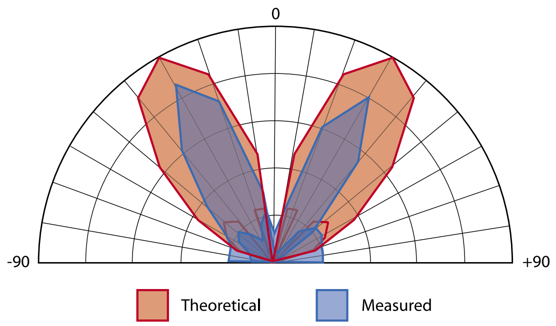

Even with only four antennas in use (the maximum the lab's testing equipment could accommodate), the receiver array's angular selectivity, shown here for a setting of either plus 30° or minus 30°, exceeded expectations. Data adapted from Hajimiri et al., *Proceedings of the IEEE*, Vol. 93, No. 9.

trio signed their work in the upper right-hand corner. "We tried to put in the Caltech logo," says Hajimiri, "but it was hard to lay out. You have to design chip elements as assemblies of squares, and there's software that combines them, but at the last minute the students said, 'We'd rather sleep after five days,' so they just put down Caltech." The plot above shows the array's theoretical and measured sensitivity to test sources placed at various angles. The chip had even better angular discrimination than predicted, as shown by the narrowness of the measured lobes.

Arun Natarajan (MS '03, PhD '07) and Abbas Komijani (PhD '05) and Hajimiri next built a four-element transmitter chip in 2004 that was unveiled at the 2005 ISSCC. The array gave a nice, tightly focused beam, as shown below. But a narrow beam counts for naught if you can't tell what's being sent. So the next test was to send actual data, in this case a random string of ones and zeroes, and see what came out. The results were plotted as an "eye diagram," so called because if you're looking at the output in an oscilloscope, the line should be at the top of the screen for a one and at the bottom for a zero. The middle should be blank, forming a wide-open eye. A squinting or closed eye reveals intermediate values, where the receiver will have to guess what was sent. At one gigabit per second, or actually 500 megabits per second per channel for two overlapping channels, says Hajimiri, "you can easily distinguish between all the ones and all the zeros. If I say, 'This is the cutoff line: anything below this is a zero, anything above it is a one,' I don't have to make any tough calls. This is considered a perfect eye—people who deal with eye diagrams are actually used to ones that are a lot less open."

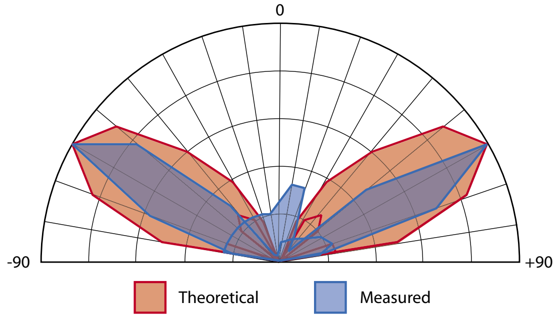

The lab's first transmitter chip, which also operated at 24 GHz, had four external antennas to match the testing apparatus. It, too, performed better than predicted. Data adapted from Hajimiri et al., *Proceedings of the IEEE*, Vol. 93, No. 9.

In both these designs, the antennas were still off-chip components. There were two reasons for this—the size of the antenna is usually proportional to the wavelength, and a 24 GHz dipole antenna would be about three centimenters long, or 10 times the length of the chip itself.

But the other problem was more fundamental—silicon makes a lousy antenna. It has a very high dielectric constant, which means that it literally soaks up the radiating electromagnetic field. It's also a semiconductor, which means that it drains the incoming electric field away before it ever reaches the antenna. For an on-chip antenna, with silicon on one side and air on the other, fully 95 percent of the power leaks into the silicon. The group spent a couple of years playing around with several possible fixes, none of which worked particularly well. "In the end we said, 'If we can't get rid of it, we'll make it a feature,'" laughs Hajimiri. "We'll just redesign the system so the chip radiates from the backside." They essentially put the chip facedown in its mounting and let the signal travel in the direction it wanted to go anyway. They even

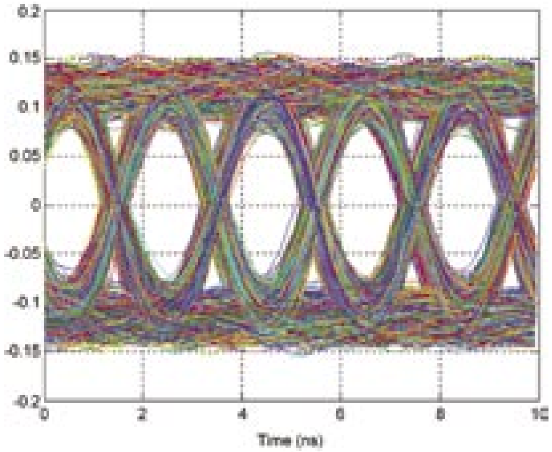

From Natarajan et al., *IEEE Journal of Solid-State Circuits*, Vol. 40, No. 12, Dec. 2005, pp. 2502–2514.

These colorful skeins are oscilloscope traces of digital data transmission from the 24 GHz chip at one gigabit per second. “Ones” are displayed along the top, and “zeroes” along the bottom; the crossovers happen when successive bits have opposite values.

turned this to their advantage by adding a silicon hemisphere to the back of the chip, now its front, to form a lens and improve the radiative properties of the antenna. The focusing effect gives the beam a longer range—up to 100 meters, says Hajimiri.

“In the end we said, ‘If we can’t get rid of it, we’ll make it a feature,’” laughs Hajimiri. “We’ll just redesign the system so the chip radiates from the backside.”

Anything that helps boost the power output is a bonus, because if you apply too large a voltage to silicon circuits, you’ll fry them, a phenomenon known in the trade as a low breakdown voltage. It doesn’t take much of a voltage difference to convert a zero into a one in your PC, which is good because the less power you consume the less cooling you need. But broadcasting is a whole ’nother ball game. Explains Hajimiri, “The low breakdown voltage limits how much power you can transmit without killing the transistors, because it limits how large a voltage swing you can have in the circuit, and that determines how much power you can generate. So we had to find a way to use a large number of transistors, each one of them generating a little bit of power, and then combine all that power somehow.”

But this had to wait until someone else had used all those transistors to make computers smart enough to help with the design. Radar frequencies are far above where other solid-state devices operate—Pentium chips, for example, run at a leisurely couple of gigahertz. Says Hajimiri, “Transistor performance has not been modeled very well at higher frequencies, and you’re basically prone to the ‘garbage in, garbage out’ principle. If you don’t know what you’re designing with, you can’t expect the product to be exactly like what you simulated.” So the group spent close to a year developing a very accurate three-dimensional model of how the electromagnetic field propagates through the volume of the chip.

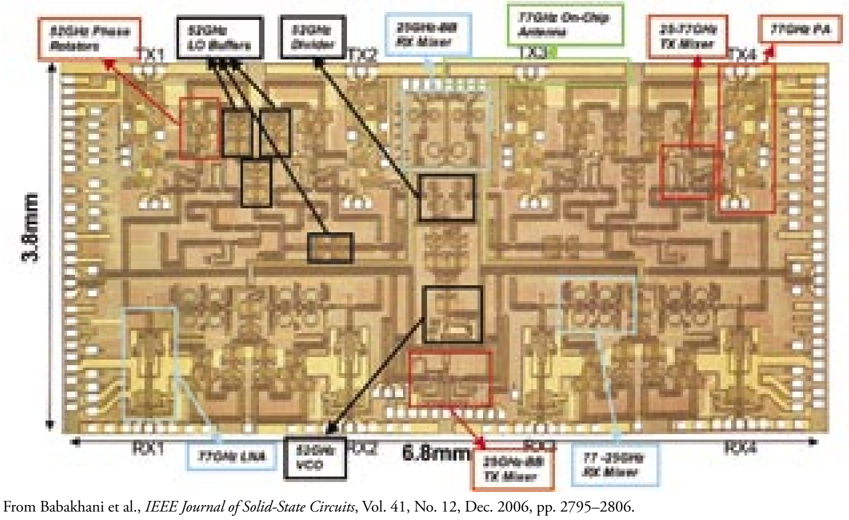

The 77 GHz model was the world’s first phased-array chip to have it all—transmitter, receiver, and antennas—on one slab of silicon. The aluminum antennas, 600 millionths of a meter long and 50 wide, are the four golden bars running along the chip’s top and bottom edges.

From Babakhani et al., *IEEE Journal of Solid-State Circuits*, Vol. 41, No. 12, Dec. 2006, pp. 2795–2806.

The model, which runs on multiprocessor PCs and can take several hours to execute, isn't perfect, either. Therefore Hajimiri's chips include self-correcting circuitry to account for performance variations in the individual transistors, as well as such external factors as temperature and humidity. "This is one thing I tell the students in my electronic design class every year—in integrated circuits, extra transistors are essentially free, so use as many as you want; use them any way you like. If it helps you, use them."

"Radio astronomy traditionally has been done with sparse, huge antennas. But if you can make them cheap, you can cover a very large area with a very large number of them—you could use an army of mice, or maybe ants, instead of an occasional elephant."

The group was now ready to put the transmitter, the receiver, and the antennas on the same chip. Natarajan and Aydin Babakhani (MS '05), Guan, Komijani, and Hajimiri spent 2005 working on a 77 GHz phased-array transceiver chip (consisting of about 15,000 transistors—peanuts compared to the tens of millions of transistors on a Pentium) that debuted, once again, at the ISSCC in 2006. The researchers used some of those transistors to simplify the design by having separate transmitter and receiver arrays—four antennas per—each with, again, all of their supporting circuitry. The chip has only two inputs: one to set the angle you want the beam steered to, and one for the data; ditto for the outputs. "That chip took about 10,000 man-hours of design time," says Hajimiri. "At 77 gigahertz, the antennas are small enough that we can put them on a chip. Putting a strip of metal on a chip is easy; putting a strip of metal on a chip that does the right thing is very, very hard. But by then we had figured out how to do it."

It's Both the Dessert Topping and the Floor Wax

With dirt-cheap send-and-receive units, you could put a whole bunch of them all over a car, and feed their outputs to a dashboard display that shows everything around you in 3-D, right next to the GPS screen. Better, to avoid sensory overload, "you could couple all those chips into a central system that does autonomous cruise control, self-parking, brake boosting, all of those kinds of features, in an integrated approach, instead of a patchwork of a little sensor here, a little sensor there, all doing different things, and not quite as well."

Lexus is touting the self-parking LS 460, on sale now. Parallel parking separates the wheat from the chaff in driver's ed, and some folks never truly master it. But in Lexus's TV spots, the driver pulls up next to a vacant space, pushes a button, and the car backs up, cuts the wheel, slips neatly into the slot, and then stops automatically. On the other hand, or perhaps the other foot, "brake boosting" is an electro-hydraulic system that keeps the brake lines as pressurized as possible for maximum stopping power when needed. The next generation of autonomous cruise control will tie into the booster controller so that if somebody cuts you off, the brakes will instantly clamp down hard at the slightest touch of the pedal. "And that's very important because most accidents—and I didn't know this until I started talking to the car companies—are caused by the fact that you don't apply the full force of the brake as soon as you see the problem, you just gradually increase it," Hajimiri says.

Which brings us to collision-avoidance systems, which are *not* coming soon to a dealer near you. On the most basic level, this could be a car telling a semi "DON'T CHANGE LANES! I'm right beside you!!" or two oncoming cars negotiating who is going to get out of the way based on their speeds, maneuverabilities, and surroundings. More

advanced systems could allow emergency vehicles such as ambulances to part the traffic in front of them. “If you have a portable device that has beam-forming capability, you can both use it as a sensor and a communication device at the same time—it’s both the dessert topping and the floor wax,” Hajimiri says. “Car-to-car communication is quite important if you want to flock cars—group them into a flock and have them drive together. They have to talk to each other constantly.” In fact, such a system would have to have all sorts of other sensors talking, too, about such things as tire pressure, for example—you wouldn’t want a car in the middle of the flock to run over a nail and suddenly get a flat.

Hajimiri is now trying to generalize this approach into broader applications, such as scalable arrays. Scalable means that the entire surface of a car, an airplane, or anything could be tiled with chips that act in unison. Maintaining phase synchronization between all of these chips is an enormous technical challenge, and things get even more interesting if the surface is curved, as the surfaces of airplanes and automobiles tend to be. “The calculations are complex but doable,” says Hajimiri. “They are somewhat similar to the calculations done for curved space-time, such as the differential geometric ones for general relativity.” And it’s well worth it, he adds, because of the tremendous power boost at the transmitter and the increased sensitivity of the receiver that results. James Buckwalter (BS ’99, PhD ’06, now an assistant professor at UC San Diego), Babakhani, and Hajimiri have developed a two-by-two scalable array that operates at 60 GHz.

Such arrays could crop up in all sorts of unexpected places, including radio astronomy. “Radio astronomy traditionally has been done with sparse, huge antennas. But if you can make them cheap, you can cover a very large area with a very large number of them—you could use an army of mice, or maybe ants, instead of an occasional elephant. You could have a square-kilometer array, or even larger, and you wouldn’t have to use fancy cryogenic systems because of the phenomenal combined gain from all those antennas.”

The group is also working on multiband, multibeam chips. Research engineer Sanggeun Jeon (MS ’04, PhD ’06), grad students Florian Bohn (BS ’01), Yu-Jiu Wang (MS ’06), and Hua Wang, and Hajimiri have built a phased-array receiver chip that can listen in on up to four beams at once, each at any frequency of your choice between six to 18 GHz.

Intel cofounder Gordon Moore (PhD ’54) wrote a visionary paper in the journal *Electronics* in 1965. In it, among other things, he first made the empirical observation that has since become known as Moore’s Law, which is usually quoted as saying that the number of transistors one can put on a chip doubles about every two years. Like Einstein’s postulation of the existence of gravity waves, every prediction Moore made has come true—except, until now, for one. The final sentences read, “It is difficult to predict at the present time just how extensive the invasion of the microwave area by integrated electronics will be. The successful realization of such items as phased-array antennas, for example, using a multiplicity of integrated microwave power sources, could completely revolutionize radar.” Beams Hajimiri, “I’m glad to tell you that we’ve done this 40 years after his prediction, and I’m glad Caltech did it. Gordon must be thrilled.” □

PICTURE CREDITS: 32 — Mercedes-Benz; 34-35, 36 — Doug Cummings; 38-39 — Bob Paz