

1 Chip, 2x the Bandwidth



Today's radio technology has almost nothing in common with its counterpart from a century ago. Countless advances have produced gear that is fantastically smaller, more reliable, and more efficient in its use of power and bandwidth. And yet, one enormous limitation remains from long ago: Radios still can't transmit and receive on the same frequency at the same time. • Such a capability, called full duplex, would be a great advance for wireless. It would in one fell swoop double a network's capacity, the physical ability to carry data. At a time when unused radio-frequency spectrum has already been mined to near extinction and demands for data are increasing—and with emerging 5G networks targeting a 1,000-fold increase in data flows—full-duplex wireless has become a holy grail in the search for a way to ease the spectrum crunch. • Now, at last, we and others have demonstrated that a full-duplex wireless system can be practical and reliable. Research in our own labs at Columbia University under the FlexCoN project and in Europe under the DUPLO initiative has demonstrated full-duplex operation within CMOS integrated circuits, the kind that are ubiquitous in today's computing and communications gear.



**A ONE-ANTENNA CHIP THAT CAN SEND AND RECEIVE SIMULTANEOUSLY
COULD DOUBLE THE DATA CAPACITY OF PHONE NETWORKS**



By Harish
Krishnaswamy
& Gil Zussman

This work follows the first demonstrations of this technique, which were accomplished only a few years ago in separate projects at Rice and Stanford universities using benchtop laboratory equipment. The Stanford research has since been spun off into a startup company, Kumu Networks, that is using discrete components to bring full-duplex capability to base stations and infrastructure, where cost and size constraints are more relaxed than in handsets.

In the wired world, duplex circuits are old hat. Early pre-electronic telephone handsets were able to use a single channel to send and receive simultaneously by isolating the earpiece from the microphone using a hybrid transformer circuit; that way, the outgoing and returning signals could each pass on a pair of twisted wires without interfering with each other.

In the wireless realm, the idea dates to the 1970s, when the Plessey Groundsat system gave soldiers full-duplex radio over channels within the 30- to 76-megahertz VHF band. But in those days the feat could be managed only where money was no object and where it was possible to put some distance between the transmitting and receiving antennas. In today's military systems, simultaneous transmit and receive capability on the same frequency channel is being pursued using photonics techniques to isolate the receiver from the transmitter.

But getting full duplex into civilian applications involving compact units, for example in cellular communications and Wi-Fi, is harder because the transmission signal creates at the receiver a tremendous self-interference, or echo. Because that echo can be anywhere from a billion to a trillion times as strong as the signal that needs to be detected, the system can be made to work only by canceling the echo very, very accurately.

Military systems pursue full-duplex capability by accepting more weight, bulk, and cost than their civilian counterparts could tolerate. To make the technique feasible for consumer products like smartphones, researchers

had to push the state of the art in such disciplines as antennas, circuit design, and algorithms.

That's why today's wireless networks are only half duplex. Transmitters and receivers either transmit and receive in different time slots (which is called time-division duplexing, or TDD) or at the same time but at different frequencies (frequency-division duplexing, or FDD). Because time or frequency resources are being used only part of the time, such



the task justice. Remember, the echo is a billion to a trillion times as loud as the signal you're straining to hear, so you have to cancel reverberations extremely accurately. That means performing cancellation across several domains: radio frequency, analog, digital, and even within the antenna interface, as we will describe later. The cancellation in each domain must be coordinated with the cancellation in all the others.

In addition, the transmitter signal is not entirely known. As the transmitter signal leaks into the receiver, it undergoes frequency distortion through the antenna interface. Further complicating matters is the fact that the transmitter signal can reflect off nearby objects and follow several paths that arrive at the receiver at different times. Therefore, to achieve near-perfect cancellation—in which a constructed replica of the echo signal is subtracted from the echo so that no more than, say, one part per billion of the echo remains—the wireless self-interference channel from the transmitter to the receiver must be determined and reproduced very accurately.

The interfering signals are of course received in analog circuits, and in an ideal world they'd be immediately handed off to digital circuits, which can process them with far greater flexibility. But in full duplex the echo is so powerful that the analog circuitry would hand off a badly distorted version of the echo. So we have to do some cancellation in the analog side first.

One method is called time-domain cancellation. Let's say the interference signal from the transmitter arrives at the receiver along with time-delayed copies of itself that have been reflected off nearby objects. Here, synchronizing the cancellation signal with the interference requires slowing down the cancellation signal by sending it through prepared pathways, perhaps centimeters long, that mimic the delays experienced by the interference signal. The problem is that an integrated circuit is itself much smaller than a centimeter square, so such a time-domain cancellation scheme couldn't physically fit onto the chip.

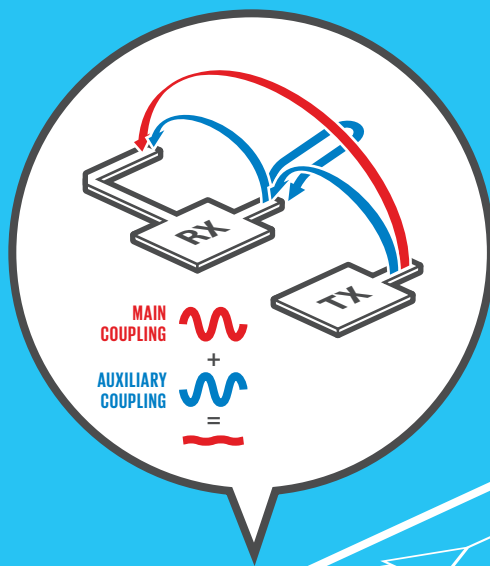
Therefore, one of us (Krishnaswamy), along with his Ph.D. student Jin Zhou



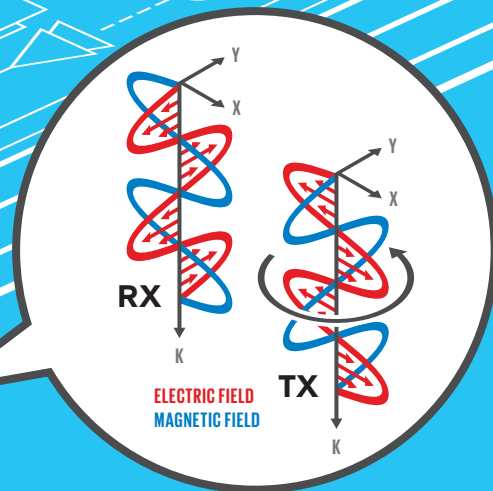
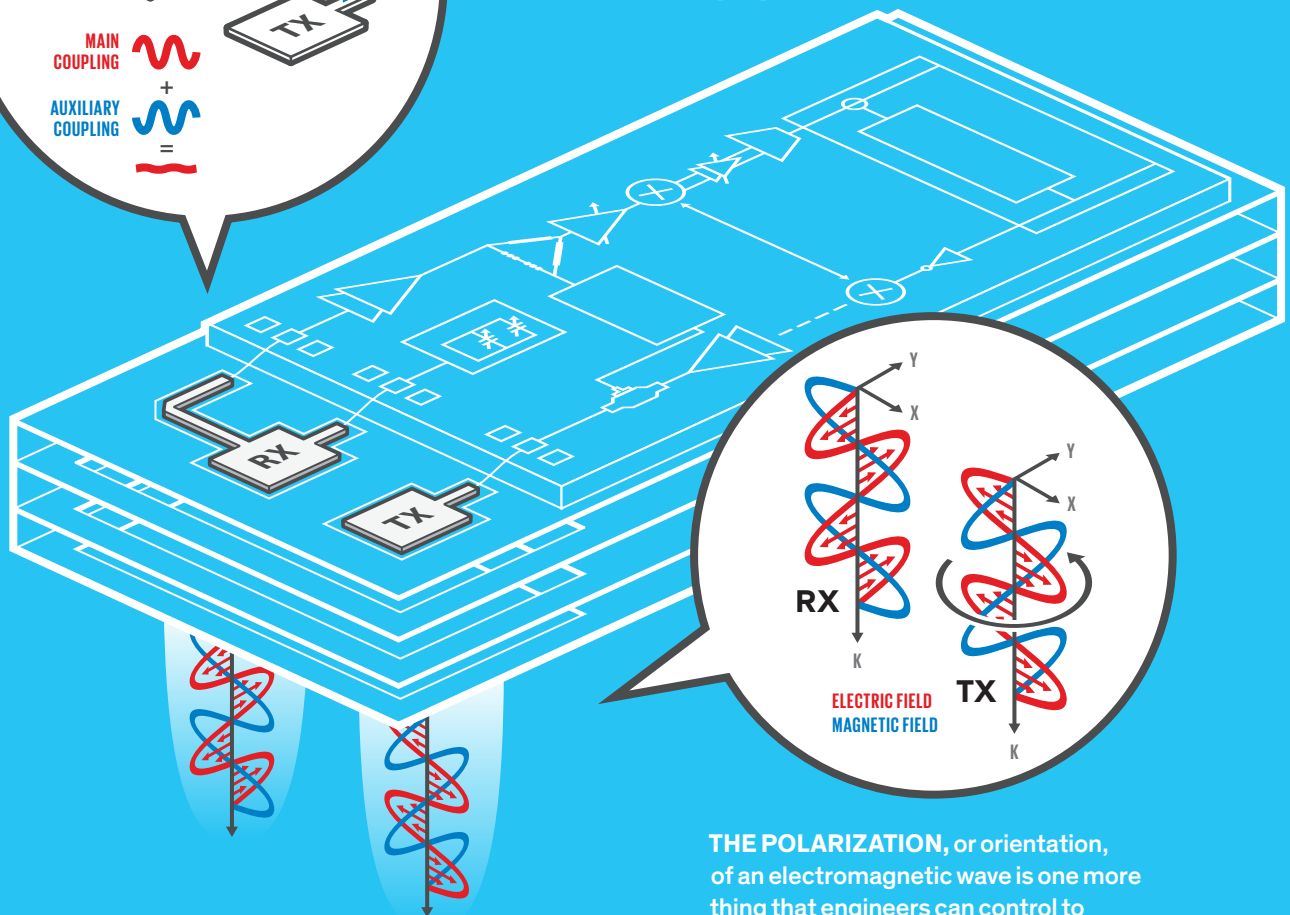
DO THE TWIST: This IC “twists” the polarity of an electromagnetic wave, making it the first nonmagnetic, on-chip circulator [chip, above; commercial package, below].

networks achieve only half the basic network capacity that is possible in an ideal full-duplex network.

To go from half to full duplex requires solving a basic problem: A wireless device has trouble transmitting and receiving at the same time for the same reason you can't hear a whisper while screaming at the top of your lungs: The interference is louder than the signal. Canceling that interference at the receiver involves subtracting the known transmitter signal. But that pithy description hardly does



THE TRANSCEIVER taps a portion of the transmitted signal [blue], conditions it, and then couples it into the receiver so that it cancels out the main interference [red].



THE POLARIZATION, or orientation, of an electromagnetic wave is one more thing that engineers can control to prevent transmitted and received signals from interfering with each other. This 60-gigahertz full-duplex transceiver developed at Columbia University uses this polarization technique to separate the transmitted from the received signals.

Polarization-Based Duplexing

proposed another approach—frequency-domain equalization. It works something like the graphic equalizer in a stereo system, which corrects sound by adjusting the power of audio signals in particular bands of frequencies. To divide the signal into multiple frequency bands that can be individually manipulated, we use filters that each have a very sharp frequency response (or high quality factor);

they take the incoming signals and let only a very narrow range of frequencies through. We use a number of filters centered at different frequencies spread out across the signal's full spectrum.

Historically, it had been impossible to build sharp radio-frequency filters on a chip. We did it in nanoscale CMOS chips by using a circuit-design trick, called the N-path filter. A conventional filter uses

inductors and capacitors, but inductors are notoriously hard to put on a chip. The N-path filter instead uses switches, which is to say transistors—the essence of integrated-circuit technology.

Other groups had already put N-path filters on a chip, but we were the first to apply N-path filters for frequency-domain equalization. The result was

echo cancellation across a very wide band of frequencies for a full-duplex wireless radio.

In our system, a bank of N-path filters taps a small portion of the transmitter signal. Next it divides that RF signal into two frequency bands (though more than two is also quite feasible). Then it conditions the signal in each of those bands to mimic the self-interference that's arriving at the receiver.

This multiband approach divides the bandwidth into bite-size chunks, a divide-and-conquer strategy that makes it easier to condition each chunk of bandwidth—that is, to adjust it for power and for phase. Circuitry performs the conditioning for each band according to the weights

nology. Our full-duplex receiver can operate at any frequency between 0.8 and 1.4 gigahertz, and the RF self-interference canceler suppresses the transmitter interference, for a variety of antenna types, over a bandwidth that is about 10 times as great as what you can get with existing, conventional cancellation techniques. We achieved this 10x performance advantage with just two N-path filters in the bank. That's good enough to make it compatible with many advanced wireless standards, including LTE and Wi-Fi. More filters would enable even wider cancellation bandwidths.


Another advantage of our frequency-based cancellation scheme is its compatibility with existing wireless systems,

tunable duplexers, but such duplexers typically are less effective in isolating the transmitter from the receiver than their fixed-frequency counterparts. Consequently, the receiver is particularly prone to transmitter self-interference. And that's where self-interference cancellation comes in.

The entry point for self-interference is right at the antenna, and it would be wonderful to suppress the interference there, before it has a chance to leak into the receiver. The main challenges are to keep the antenna compact—say, for use in a cellphone—and make sure that the self-interference doesn't come back every time the electromagnetic environment changes. In other words, we need the antenna to be smart.

Such a smart antenna can manipulate the radio wave's obvious electronic characteristics—amplitude, phase, and frequency—but also the extra dimension: wave polarization. A radio wave is really two fields joined at the hip, one electrical, the other magnetic—hence the word “electromagnetic.” Each field oscillates at a given frequency, and the oscillation of the electric field induces the magnetic field, and vice versa. The two fields are perpendicular, and the way the pair of them are oriented in space is called their polarization. Electromagnetic waves of different polarization can pass through each other without interference.

Krishnaswamy and his Ph.D. student Tolga Dinc were able to use polarization for duplexing within a pair of compact antennas (configured for 4.6 GHz), one for the transmitter and one for the receiver. We were able to place them right next to each other because the waves that were coming and going to the two antennas were orthogonally polarized with respect to each other, which effectively isolated them. But though this isolation minimizes self-interference, it does not eliminate it entirely. That's why we also installed a port in the receiving antenna that's copolarized with the transmitting antenna. The port samples a small portion of the transmitted signal, conditions the signal through a filter, and then passes it on to the receiver port. Result: near-perfect cancellation. Because the filter



ONE SMALL PACKAGE: This is the world's first full-duplex transceiver on a chip. It cancels interference with the familiar technique of frequency-domain equalization.

assigned to that band. Again, to use the audio analogy, it's like dialing up the bass and dialing down the treble, and doing whatever else it takes to get the output to match the input signal.

The next step is to automate this weighting process so that the output changes accordingly as the environment changes. Of course, these environmental changes are dynamic, fluctuating from second to second, so the process has to be automated. Though we have shown some initial and promising demonstrations of such automation, there is more work to be done here.

We designed a prototype of the receiver-cum-canceler and fabricated it using 65-nanometer CMOS tech-

which use what's called multiband frequency-division duplexing. As mentioned earlier, FDD is a half-duplex strategy in which the transmitter and the receiver operate at the same time but at different frequencies. It requires duplexer filters to separate the transmitter and the receiver signal at the common antenna. Because these off-chip duplexer filters cannot be tuned, today's smartphones use a separate duplexer filter to support each of the FDD bands—and 4G LTE supports 25 bands, therefore requiring 25 separate duplexer filters!

You can reduce the bulk and cost of the radio component of a cellphone by replacing those filters with just a few

ONE CHIP, TWICE THE BANDWIDTH

CONTINUED FROM PAGE 42

can be programmed in the field, it can reconfigure its cancellation so that it meets the demands of a changing electromagnetic environment.

Our prototype achieved an isolation of 50 decibels—that's five orders of magnitude, or 100,000-fold—over a 300-MHz bandwidth. This level of isolation represents a 1,000-fold improvement over the isolation that is achieved without the polarization-based cancellation. Even when we put a strongly reflecting metallic plate near the antenna to heighten the self-interference, we were still able to reconfigure the system and fully restore the cancellation.

Our method of coordinating antenna design with cancellation techniques is easy to adapt to higher frequencies—30 GHz and above. In that part of the spectrum, the wavelengths are measured in mere

millimeters, so the antennas that receive and transmit them are small as well.

These high frequencies are particularly promising for next-generation communication networks because they offer significantly wider bandwidths than radio frequency. We have implemented a 60-GHz full-duplex transceiver IC that employs both our reconfigurable polarization-based antenna cancellation and RF and digital cancellation. The device achieves self-interference suppression of nearly 80 dB (or 100 million-fold) over 1-GHz bandwidth. It thus enabled the world's first millimeter-wave full-duplex link over a distance of almost a meter. That's a respectable distance for millimeter-wave links, which are being considered for various short-range applications, such as wireless USB connections.

Our system can work in both Wi-Fi and cellular networks. The cellular application is the harder one because every link must be carefully coordinated by the

base station, which might assign one person to one frequency and another person to another, timing everything perfectly. But Wi-Fi is a free-for-all, with none of these constraints. Also, cellphones transmit over a kilometer or more, but Wi-Fi reaches only a few score meters, producing signals that are fewer in number and much lower in power. All this makes it easier to perform echo cancellation.

One of the authors (Zussman) and his student Jelena Marašević were able to analyze the benefits from full duplex running on real chips (not merely in an idealized scenario, as others had done). They found that you can't always assume that the cancellation is perfect: The received signal needs to be a lot stronger than the weak echo that remains even after you've done all your cancellations.

In the meantime, several problems will have to be solved before we can declare wireless full duplex fully achieved. First, to take full advantage of the compactness of our duplex-on-a-chip, we have to construct what's called



Messe München

Connecting Global Competence



Planet e: Where the future begins.

Electronics of tomorrow. Today.

Tickets & Registration: electronica.de/en/tickets

World's Leading Trade Fair for Electronic Components,
Systems and Applications
Messe München | November 8–11, 2016 | electronica.de

Contact: German American Chamber of Commerce, Inc.
Phone +1 646 437 1013 | dbartels@munich-tradefairs.com

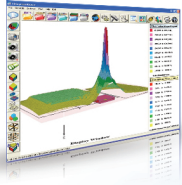


electronica 2016

inside tomorrow

THE INDUSTRY STANDARD SOFTWARE PACKAGES IN GROUNDING AND ELECTROMAGNETIC INTERFERENCE

Are now more powerful with the addition of new standard components such as arbitrary transformers and cables.



Classical Equipotential Grounding

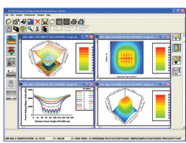
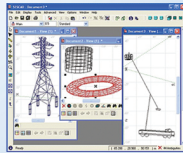
Multiple grounding systems having any shape, in simple or complex soils, including any number of layers or heterogeneous soil volumes.

- AutoGroundDesign
- AutoGrid Pro
- AutoGround & MultiGround

Complex Frequency Domain Grounding, Interference Analysis & Environmental Impact Assessment

Perform fast, yet complex and accurate, Interference analyses on pipelines, railways, etc. Examine electromagnetic impacts to the environment.

- SESTLC
- Right-Of-Way
- MultiLines & SES-Enviro
- MultiGroundZ



Frequency & Time Domain Arbitrary Network Analysis

Upgrade to the full power of MultiFields or CDEGS to tackle any electromagnetic problem involving aboveground and buried conductor networks including lightning or switching surge analysis.



World Leader in Grounding & EMI
www.sestech.com

800-668-3737

450-622-5000

a circulator—a device that shares a single antenna between the transmitter and the receiver. That's not easy because such a circulator must be nonreciprocal—that is, it must process a signal coming in differently from the way it processes a signal going out. Only in that way could a single antenna serve as a transmitter and a receiver simultaneously.

But treating a signal differently depending on which way it's moving violates the basic physical precept of Lorentz reciprocity. Fortunately, this precept applies only to *most* materials and systems. Ferrite materials are one exception, and indeed researchers have used them to build nonreciprocal circulators for years.

The ferrite materials work by twisting an electromagnetic wave clockwise with respect to the axis of motion when it's going one way, counterclockwise when it's going the other way. But because ferrite circulators don't fit on a chip, we came up with another way: We twist the wave with switches—transistors, in other words.

This past April, Krishnaswamy and his Ph.D. student Negar Reiskarimian did manage to build a nonreciprocal circulator that uses transistors to mimic the behavior of ferrites. This was the first such unit ever built on a chip, and we integrated it with a single-chip, full-duplex, echo-canceling receiver. The result: single-antenna full duplex.

A second challenge is to extend our method of canceling self-interference to transceivers that have multiple inputs and multiple outputs (MIMO). Such transceivers are common in base stations, and they can transmit over multiple parallel streams, greatly increasing the data rate. Unfortunately, self-interference plagues every possible transmitter-receiver pair, and if you try to implement the filters to handle every possible pair, the complexity will scale quadratically with the number of MIMO elements. This is a real problem, and though we have some ideas on how to proceed, there is a lot of work ahead.

But even now, we expect to be able to use full duplex on a chip to improve key aspects of existing wireless systems. One of the first applications, probably within two years, should come in short-range wireless links and in systems, such as Wi-Fi, in which the received signal-to-noise ratio is generally high. Getting the technology into cellphones, with their more powerful self-interference signals, will probably take about five years. Another application that will benefit is fixed, point-to-point microwave and millimeter-wave backhaul and relays, which form the backbone of our telecommunications network.

Many major players in the wireless industry have expressed interest in our work, and some, such as Qualcomm, have even funded us. We have also received funding from the U.S. Defense Advanced Research Projects Agency and the U.S. National Science Foundation.

Five years down the line, we expect to be using full duplex for small cellular networks. We are confident that the day will finally come when all our wireless devices will speak and listen at one time, on one channel, through one antenna, within one chip. ■

POST YOUR COMMENTS at <http://spectrum.ieee.org/fullduplex0716>



TEMASEK RESEARCH FELLOWSHIP (TRF)

A globally connected cosmopolitan city, Singapore provides a supportive environment for a vibrant research culture. Its universities Nanyang Technological University (NTU), National University of Singapore (NUS) and Singapore University of Technology and Design (SUTD) invite outstanding young researchers to apply for the prestigious TRF awards.

Under the TRF scheme, selected young researchers with a PhD degree have an opportunity to conduct and lead defence-related research. It offers:

- A 3-year research grant of up to S\$1 million commensurate with the scope of work, with an option to extend for another 3 years
- Postdoctoral or tenure-track appointment (eligibility for tenure-track will be determined by the university)
- Attractive and competitive remuneration

Fellows may lead, conduct research and publish in these areas:

- Adaptive Camouflage Techniques and Technologies
- Cyber Security for Autonomous Systems
- Perception under Adverse Conditions for UGV Navigation

For more information and application procedure, please visit:

NTU—http://www3.ntu.edu.sg/trf/index_trf.html

NUS—<http://www.nus.edu.sg/dpr/InfoForResearchers/trf.html>

SUTD—<http://temasek-labs.sutd.edu.sg/funding-opportunities/temasek-research-fellowship-trf>

Closing date: 23 September 2016 (Friday)

Shortlisted candidates will be invited to Singapore to present their research plans, meet local researchers and identify potential collaborators in February/ March 2017.