TELECOM // WIRELESS

FEATURE

Radios With Micromachined Resonators

Future wireless designs will replace electronics with precision mechanical components

Illustration: Harry Campbell

BY CLARK T.-C. NGUYEN // DECEMBER 2009

We do love our cellphones. And we hate them, too, of course—when they drop a call, go dead in the middle of a conversation, or simply fail to work in another country. Soon we'll probably be complaining about other things—perhaps that our handsets can't receive satellite TV broadcasts or last more than a week on a single charge.

You might guess that better microelectronics will soon provide higher data rates, lower power consumption, and greater flexibility in the types of communication that our handsets can manage. To some extent, that's true. But transistor advances alone will probably not be enough. The Moore's Law world of regularly doubling transistor densities has brought us cheap PCs that outperform the multimillion-dollar mainframes of 30 years ago, but those incredible shrinking transistors might not do much to eliminate dropped calls. In this respect, the most significant improvements may, in fact, come from what seems a bizarre source: better *mechanical* components.

The idea of adding a bunch of moving parts to a radio handset conjures up images of cellphones outfitted steampunk style with brass levers and steel gears. This, of course, is not what I mean. Rather, I'm suggesting that tomorrow's designs will benefit from advances in the kinds of mechanical devices already found in cellphones and other wireless equipment. If you don't believe your phone contains such things, open up the back and take a look. You'll see a battery and integrated circuits—and also such things as thin-film bulk acoustic resonators (FBARs), surface acoustic-wave (SAW) resonators, and quartz crystals. These components, which convert electrical signals to mechanical displacements, do the work that electronics struggle with—for example, selecting a narrow band of radio frequencies and removing interfering signals from all the energy captured by the antenna, and synthesizing extremely stable oscillating waveforms, which are needed to process the incoming radio-frequency signals.

This vision of wireless gear evolving to include more and better mechanical devices of this sort is very different from the approach some radio engineers are now pursuing. They seek to eliminate analog filters and use digital circuits to handle everything—interfering signals and all—using software to do all the filtering. The problem with building such a fully software-defined radio is that the ultrafast analog-to-digital converters that could deal with interfering signals in such a setup are not yet available. Worse, when they do become available, they will probably use too much power to be practical for battery-powered handsets.

Doing more of a radio receiver's up-front work mechanically has many advantages over relying solely on electronics and software. For example, it eases the demands on the electronics used for further signal processing—the additional filtering, analog-to-digital conversion, and so forth—and that in turn saves on the power consumed in that circuitry. This strategy would allow a portable radio receiver to monitor a wide swath of the spectrum at all times without swiftly burning through its batteries. That will be important in the much-anticipated world of cognitive radio, in which our handsets become agile enough to exploit frequencies that are fair game only when a higher-priority user hasn't claimed them.

It's likely that better mechanical components, and the cognitive-radio techniques they enable, will usher in the next wave of mobile telephony by giving our cellphones access to much more spectrum. These phones will operate in multiple bands, provide greater data throughput, and minimize if not eliminate the need for wireless providers to drop our calls because traffic exceeds capacity. Consumers will love the result, even if they don't know anything about the high-tech mechanics that may soon make it possible.

How can mechanical devices outperform electronic ones? One reason is that they generally consume no battery power. Another has to do with the quality factor of the resonating components, a quantity that physicists and engineers denote with the letter Q. The higher the Q, the more selective the resonator will be in responding only to a narrow range of frequencies.

Like any good radio receiver, the one in a cellphone requires resonators with Qs greater than 1000. Resonant electrical circuits, typically built with capacitors and inductors, have great difficulty achieving values that high. Indeed, the inductors in conventional integrated circuits are dismal, generally yielding circuits with Qs of less than 10. Vibrating mechanical resonators, on the other hand, can easily provide values in the required range.

Unfortunately, the need for such resonators makes handsets more costly to manufacture. If tomorrow's cellphones were to use many more of these mechanical components, the expense of including them could well dominate the cost of handsets. And the large sizes typical of these mechanical components could be a problem, too, although some are already quite small. Makers of FBARs, for example, use micromachining to construct on-chip gigahertz-frequency resonators with dimensions of about 200 micrometers.

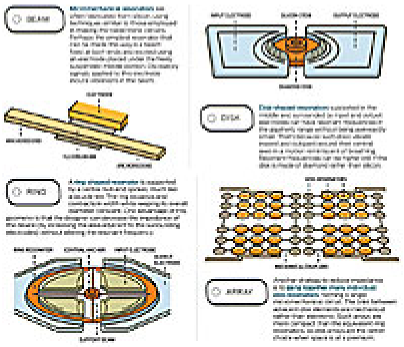

Illustrations: Harry Campbell

Micromechanical resonators
Click on image for larger view.

But small size is not the only requirement. Consumers will soon demand handsets that can handle all their wireless communication needs—Wi-Fi, Bluetooth, hookups with wireless sensor networks, cellular calls, broadcast television, even satellite links. And to do all that practically, designers will be compelled to put resonators of many different frequencies on a single chip. Unfortunately, the frequency of an FBAR resonator is set by the thickness of its constituent film, which means, for example, that 20 different deposition procedures would be needed to fabricate 20 different filters on the same chip. Doing so would likely end up being more expensive than just buying 20 individual FBAR devices.

Here's where the latest wizardry of microelectromechanical systems (MEMS) can save the day. A single layer of silicon, for example, can provide many different resonant frequencies if it is patterned so that parts vibrate from side to side, in directions parallel to the plane of the device. The designer just needs to create features with the correct lateral dimensions.

Easier said than done, of course. To get a sense of the challenge, consider a more familiar object whose lateral dimensions govern its resonance frequency: a steel guitar string, which spans about 64 centimeters. If tuned to middle A, it will have a fundamental resonance frequency of 110 hertz. The act of plucking a guitar string is capable of exciting any frequency, but this string will mechanically select just the one A note—and will do so with a Q of about 350. (That, by the way, is 50 times as good as a typical on-chip electrical circuit made of inductors and capacitors.)

Selecting a particular frequency is exactly what the filters of a radio do. Of course, they oscillate much faster than a guitar string vibrates, commonly hundreds of millions of times a second. To achieve such rates, you'd have to shrink a guitar string down to less than 10 µm and construct it out of a stiffer material, such as silicon. The result would be a tiny, flexible beam. You couldn't pluck such an object with a pick, of course, but you could easily excite it with an electric field. Such a micromechanical oscillator can be made to resonate with Qs in excess of 10 000.

As you might guess, shrinking things by another factor of 10 yields resonant frequencies in the gigahertz range, which are needed to reach the higher bands used for wireless communications today. However, as with nanometer-size electronics, some thorny engineering issues arise with mechanical resonators this small. For one, it's hard to control the resonance frequency precisely when the dimensions are so tiny. And even if you could do that consistently, you'd probably find that what you'd built couldn't handle much power.

Fortunately, there are ways to achieve gigahertz-range resonant frequencies without having to reduce dimensions to nanometer scale. My colleagues and I at the University of California, Berkeley, have, for example, fabricated some very useful gigahertz-frequency MEMS devices that measure several micrometers across. They can be tuned to the desired frequency relatively easily and are large enough to handle the power levels found in receiving circuitry. We've made them by fashioning the moving parts in the shape of thin disks, which resonate by expanding and contracting radially ever so slightly, rather than flexing like a beam.

If you make such a disk out of diamond instead of silicon, it will be stiffer and will consequently resonate at frequencies that can easily exceed 1 gigahertz. And if you arrange things so that this disk is supported only at its center—a point that doesn't move during the in-and-out oscillation—using a slender stem of silicon, the quality factor for this resonator can be stunningly high. We've built ones that measure 10 µm across, run at 1.5 GHz, and have *Q*s of more than 10 000—even higher when the air between the disk and the surrounding electrodes is removed. Versions of such disks running at 500 megahertz give Qs greater than 50 000.

This approach not only achieves the kinds of frequencies you'd want without having to build nanoscale objects, it also eliminates the need to use a vacuum to achieve high *Q*, reducing manufacturing costs. And because the resonant frequency of such a disk is roughly inversely proportional to its radius, even higher frequencies—such as Wi-Fi's 2.4-GHz, 3.6-GHz, and 5-GHz bands—with similar *Q*s should be possible simply by reducing the size.

But there's a limit to how much you can shrink things before the complications of working at the nanoscale start to emerge. Thankfully, we've discovered ways to dodge this problem. And happily enough, they don't require the use of diamond, which can be expensive to manufacture. In fact, all-silicon resonators can perform just as well, if not better, when constructed with somewhat more complicated geometries.

My colleagues and I have had success, for example, with resonators that take the form of a ring attached to a central support with four spokes. The ring expands and contracts in width while its average diameter remains fixed. (Imagine a bicycle wheel with just its rubber tire expanding and contracting slightly.) Our resonators can vibrate in this way at very high frequencies and, if properly designed, without losing much energy through their "spokes."

We have constructed one such ring that resonates at 1.46GHz and has a *Q* of 15 248, the current world record for an on-chip resonator operating above 1 GHz at room temperature. In a cellphone, that would translate into a filter with a pass band about 100 kilohertz wide—much more selective than the 35-MHz filters now found in cellphones. Indeed, it's narrow enough to remove all interfering signals and pass just a single communications channel. Because the processing circuitry that follows wouldn't have to deal with large-amplitude interfering signals, it could operate at lower power levels. And this basic design should work for much higher frequencies as well.

Ring resonators have other advantages too. Unlike what happens with a disk, it's easy for the designer to specify the electrical impedance of the device without changing its resonant frequency. This then allows the impedance to be matched to the circuits attached to the resonator, which is important for the same reason that it's important between, say, a stereo amplifier and its speakers: Without a good impedance match, power isn't transferred efficiently.

As nice as rings are, it might, in fact, be advantageous to gang several disks together in a mechanical circuit. A set of such disks will accomplish the necessary impedance matching and be physically smaller than an equivalent ring resonator. So a disk array may be a better choice when space is at a premium.

Some researchers, including Albert P. Pisano, my colleague at Berkeley, and Gianluca Piazza, of the University of Pennsylvania, are looking at another way to achieve the desired impedance and to handle high power levels. They are using piezoelectric MEMS resonators similar to FBARs, but with frequencies that can be controlled by adjusting certain lateral dimensions. The devices they've built should be sufficient even for the kinds of power levels found in transmitter circuitry, which are always much higher than what's encountered on the receiving end. The problem with piezoelectric resonators is that their *Q* values have so far been limited to about 3000.

In addition to working with disks and rings, I and other researchers around the world have experimented with MEMS resonators of other types: beams, squares, and combinations of these shapes. Lots of geometries are possible, of course, and no one will be surprised if something entirely new eventually proves even more capable than anything that's been built so far.

Some component vendors—for example, SiTime Corp. of Sunnyvale, Calif., and Discera, a company I founded in 2001, based in San Jose, Calif.—are currently marketing MEMS resonators for use in precision oscillators. These add to the growing number of applications where MEMS devices are turning up: principally in accelerometers, pressure sensors, gyroscopes, ultrasonic transducers, and microphones. The FBAR filters found in today's cellphones are also MEMS devices, albeit ones with lower *Q* values. I expect to see companies gearing up to apply more advanced MEMS technology to the construction of high-*Q* filters.

Suffice it to say that oscillators built with these mechanical resonators are far superior to their electrical counterparts. But these are not the only virtues of this technology. The best thing about these mechanical marvels is that they can do much more than just oscillate. If you're clever, you can transform a MEMS resonator into a complete radio receiver stage—one that can take an incoming RF signal and amplify it, down-convert its frequency, and filter the result—all with just one minuscule, passive mechanical device. This may seem like magic, so let me explain in more detail how this micromechanical prestidigitation works.

Whether built as a disk, a ring, or something else entirely, the vibrating mechanical part of these new MEMS resonators isn't placed in physical contact with its input or output terminals. Rather, it's coupled to the input and output signals by means of an electric charge placed on it. Because of the force between electric charges, the moving piece begins vibrating when an oscillating electrical signal is applied to the nearby input electrode. And similarly, the vibration of the electrically charged resonator induces an oscillatory signal on the adjacent output electrode.

To generate the required electric charge, you simply apply a bias voltage using a third terminal attached to the oscillating part of the resonator. Setting that bias voltage to zero turns the resonator off, effectively opening a switch between the input and output terminals. (That's better than putting a transistorized switch in the signal path, which is what you'd have to do with an electrical filter, because such switches degrade the signal.)

The bias voltage can, in fact, do more than just turn the resonator on or off: It can also amplify the incoming signal, down-convert its frequency, and filter it, as I mentioned. The trick is to apply an *oscillating* bias instead of a DC voltage.

If you think about what happens when the bias voltage varies with time, the magic begins to make sense. When the bias voltage is close to zero, the output will of course be very muted. When the bias is significant, the output will be large if the input is large and small if the input is small. In other words, by changing the bias, you can modulate whatever signal you apply to the input.

You could, for example, use this mechanism to shift your voice up to radio frequencies by attaching an RF oscillator to the input and using the voice signal to adjust the bias voltage. A MEMS resonator can be used equally well to shift frequencies in the other direction and down-convert the RF picked up on a radio's antenna: Just direct the captured RF to the input electrode of the resonator and hook up a suitable RF oscillator to the bias terminal. The MEMS device will automatically filter the output, too, passing only signals that are close to its resonant frequency. If the down-converted and filtered output of the resonator is too weak, increase the amplitude of the oscillating bias voltage; if it's too strong, use a smaller bias. As you can see, it's easy to control the operation of such an all-in-one MEMS radio stage.

Just as integrated circuits overtook discrete transistors decades ago, collections of MEMS devices—integrated micromechanical circuits—may eventually become common in wireless handsets. Designers could, for example, combine a large number of MEMS resonators to create a bank of elements each capable of selecting a single communications channel instead of a broad band containing a confusion of signals picked up by the antenna. Providing

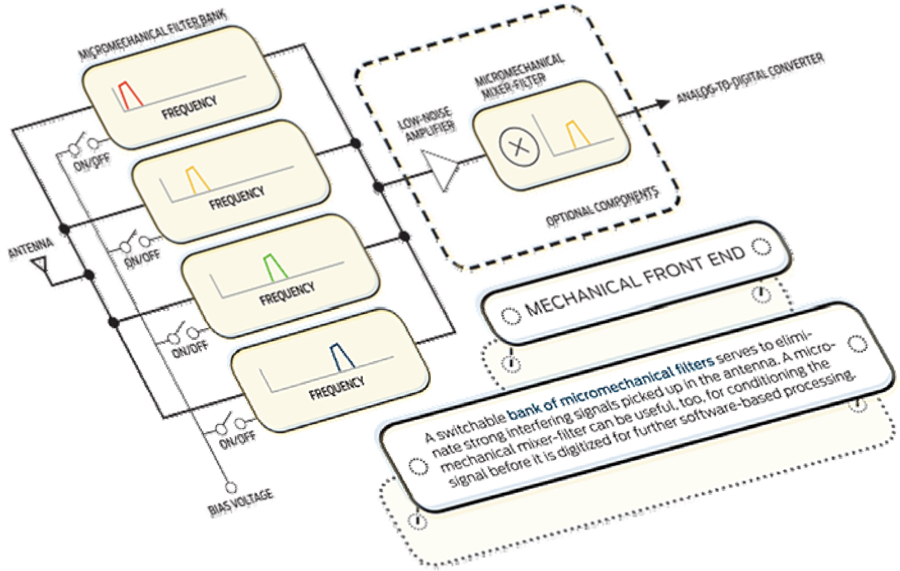

MECHANICAL FRONT END

A switchable **bank of micromechanical filters** serves to eliminate strong interfering signals picked up in the antenna. A micromechanical mixer-filter can be useful, too, for conditioning the signal before it is digitized for further software-based processing.

Click on image for a larger view.

a radio with filters that could separate out hundreds or even thousands of individual channels has been unthinkable before now. With MEMS, such separation becomes possible.

This strategy is very different from what goes on in the front end of a typical radio set, which uses tunable electronic circuitry to select which of the many frequencies captured by the antenna is sent on to subsequent amplifying and processing stages. It's also different from what's proposed for software-defined radios, which is to separate the signals on different channels computationally.

Doing channel selection mechanically would, like the computational approach, allow a smart radio receiver to monitor many different channels simultaneously, allowing it to identify in real time which slices of spectrum are free. The advantage of doing this with a bank of MEMS resonators is that they would draw very little power while doing their job, whereas the purely computational approach, even when it becomes feasible, seems destined to leave our handsets too often with dead batteries.

Using MEMS devices to replace typical front-end radio circuitry is very much still in the active research phase, but virtually all the handset makers are funding work in this area, as is the Defense Advanced Research Projects Agency. So some very sophisticated mechanical radio stages will very likely make the leap from lab bench to marketplace within the next few years.

These complex micromechanical devices will not only be integrated together; they'll most likely be put on the very same chip with electronic circuitry. This combination is possible because the fabrication processes used to create these tiny moving parts—depositing various materials in thin layers and patterning them in complex ways—is so similar to what's done in making complementary metal-oxide semiconductor (CMOS) circuits.

For example, Analog Devices, in Norwood, Mass., currently makes MEMS accelerometers and gyroscopes with mechanical and electronic components together on the same chip. This is awkward to do, though. The problem is that the construction of MEMS components requires high temperatures that would damage the copper or aluminum traces used in the electronics. So you have to lay down the mechanical parts first before putting down metal. Because most foundries won't work with wafers that are anything but pristine silicon, this arrangement prevents MEMS-device manufacturers from outsourcing the construction of the electronics and taking advantage of industry-wide advances in CMOS fabrication.

Analog Devices mixes MEMS and CMOS fabrication on a single wafer, first doing some of the steps needed for the transistors, then some for the MEMS, then returning to the transistors, and so forth. It would be better to arrange things so that all the CMOS circuitry is created first with the MEMS on top. Some researchers have done that by bonding a MEMS device to a CMOS chip using tiny wires. My colleagues and I have lately been exploring a different strategy: building the resonators out of metal, which can be laid down at temperatures low enough to avoid ruining the underlying electronics.

Whatever system eventually wins out, it's a good bet that the highly integrated chips that go into radios will slowly evolve from purely electronic devices into ones that are a complex mixture of electronic and mechanical components. The handset you carry a few years from now might not look very different from the one you have at the moment, but if you're technically savvy, you'll appreciate that it performs much better, thanks to internal mechanical parts. If you appreciate that fact enough, you might even be tempted to dress up your cellphone on the outside with a few brass levers or steel gears.

This article originally appeared in print as "Mechanical Radio."

About the Author

Clark T.-C. Nguyen argues in ”Mechanical Radio” that machining micromechanical parts into our electronics devices will lower their power consumption and make them more robust. It’s about time for micromechanical circuits to have their day, he says: ”You can see it starting up.” Nguyen, who has pioneered advances in MEMS communication technology through his start-up company, Discera, and as a program manager in DARPA’s Microsystems Technology Office, is a professor of electrical engineering and computer sciences at the University of California, Berkeley, where he received his B.S., M.S., and Ph.D. degrees.

most recent comments

ROBERT A. JOHNSON 12.11.2009
Clark, keep up the good work. Maybe you can comment on the frequency shift with temperature question that Vojak mentions. As you know, loss is low for FBAR and SAW filters and the TC of frequency in these two filters is low but how about the other filters you describe? Another possible problem we talked about a few years ago is intermodulation distortion. Thank you for a fine paper.--Bob.

C. FULLER 12.10.2009
This is some good stuff. MEMS filters are going to have to become a lot less lossy before they'll compete effectively with current filter technology. I don't understand why they are so lossy..

HENERY BACIN 12.10.2009
Mechanical filters (vibrating arm type) have been around since the 40's highly used in telecom and radios into the 80s. They have several problems, too high Q means long acquition time, high thermal drift of Fo, insertion loss, matching difficulity, and high cost. What is your thermal drift of Fo, in %, from -40 C to +100 C ? .

BRUCE VOJAK 12.10.2009
I didn't find any mention of the frequency sensitivity of these devices to changes in ambient temperature. When I last explored this topic over a decade ago at Motorola, that was perhaps the greatest barrier to MEMS technology being used in communications applications. MEMS devices were orders of magnitude more sensitive than quartz to temperature variation, rendering them fascinating, but useless for practical application at the time. Where does this stand today? Hopefully this has been / can be overcome..

CLIFF E PEERY 12.10.2009
very interesting. Awesome!.

JOHN KOTROSA 12.10.2009
Awesome article! Dr. Clark T.-C. Nguyen has reinvented radio! As others have commented on mechanical resonators being nothing new, I would argue that these aren't your granddad's resonators! This is a fascinating world of new possibilities - Great article and well written! Congratulations on some fabulous discoveries. I can't wait to see how this will change how we engineer the next gen of tiny low power receivers. I LOVE the innovative mixing and tuning techniques proposed here! So much for my career in DSP and software defined radio!.

MIKE MARCHYWKA 12.10.2009
I haven't look at this area lately beyond some passing interest in carbon MEM's (and I was surprised to see diamond mentioned after reading beyond first page and this got me interested) or cell phone accelerometers but I guess you could make digitally programmable "mechanical" devices that aren't much different from a trumpet, or guitar as you mention, with a few more valves. So, I wouldn't compare these to crystal resonantors that are generally assumed to "tune" only due to temperature changes (LOL). Monolithic inductors of course aren't obviously beneficial but just as people tried to convert paper systems into computer equivalents, I'm sure monolithic inductor approaches will be considered. Eventually of course you have to wonder what the future is for the lumped parameter circuit in which you can point to individual components characterised by a single parameter like L or C ? Resonantors have at least two, f and Q. LOL. re poster below, many FM IF's iirc were 10.7 mhz. .

DAN HICKS 12.07.2009
This is certainly not the first time that mechanical resonators have been used in radio applications. I can remember back in the 60s reading (likely in the long-defunct "Electronics World" magazine) of ceramic resonant filters that were being touted to replace the tuned intermediate frequency transformers in superhet AM receivers (455KHz, if I recall correctly). I suspect that the idea never caught on because FM was becoming popular, and FM required a

much higher IF frequency. I also vaguely recall that mechanical resonators in the form of resonant medal bars were used in some early CW radio sets. (But that was before my time.) Dan Hicks Byron MN.

WILL SOMMERVILLE **12.02.2009**
Thank you for an excellent article. Clark Nguyen is the undisputed leader in this field and I love his work. I wrote my undergrad thesis, titled Integrated MEMS Single Chip AM Radio Receiver at Sandia National Labs and referenced many of Clark's papers. I agree with his statements in this article and I am really glad that it is getting distribution. I was personally surprised when at Sandia, our first MEMS resonators that we connected to electronics fired up on the first try and filtered radio signals with a Q of 200 without any special vacuum packaging or treatment. This technology is proven and fairly mature - not at all as futuristic as it might sound. Again, thank you for the excellent article and I hope to see more MEMS in communications applications soon!.