Los robots en la sociedad del futuro

Ariel Palazzesi

Investigadores españoles han realizado un estudio sobre el impacto que tendrán los robots en la sociedad del futuro. Los resultados son inquietantes: según sus descubrimientos para el año 2020 los robots serán tan “inteligentes” y su interacción con los humanos será tan grande que existirá un desequilibrio tecnológico enorme entre quienes posean o no una estas herramientas.

Hemos hablado hasta el cansancio sobre el papel que jugarán (o no) los robots en los conflictos bélicos del futuro. Pero, ¿cómo cambiaran la vida de cada uno de nosotros, en nuestro ámbito laboral o social? Afortunadamente, un equipo de investigadores españoles, liderados por Antonio López Peláez, se ha planteado esta cuestión, llegando a conclusiones sorprendentes sobre el impacto social que tendrá la robótica en las próximas décadas.

Antonio López Peláez es un profesor de Sociología de la UNED, que ha entrevistado a expertos en robótica de todo el mundo para obtener un pronóstico de cómo cambiarán nuestra vida diaria los robots. Según la opinión de lo investigadores, en el año 2020 se producirá un punto de inflexión tecnológica, gracias al cual los robots “serán capaces de ver, actuar, hablar, dominar el lenguaje natural y ser más inteligentes. Entonces nuestra relación con ellos será más constante y más cercana”, dice López Peláez. Los autómatas dejarán de ser máquinas sofisticadas que llaman nuestra atención en exposiciones o series de TV para convertirse en herramientas cotidianas que nos ayudarán en las tareas más comunes.

Según el investigador, los robots androides que construiremos a partir de ese año, contarán con funciones y niveles de inteligencia tales que se convertirán en compañeros para la especia humana. De hecho, y tal como lo plantea el profesor Kurzweil desde hace años, esta singularidad hará que la inteligencia de nuestras máquinas sea equiparable a la nuestra. En este contexto, resultan obvias las diferencias que tendrán entre sí las sociedades que cuenten con estas máquinas y aquellas que no las posean.

La clave para que se produzca este punto de inflexión se encuentra en los avances que se han producido (y que se puede especular que se producirán en breve) en áreas tales como el reconocimiento de voz, el sentido del tacto robótico, la inteligencia artificial, la nanotecnologia, la antropología robótica y, cómo no, la capacidad de los robots de superar el famoso Test de Turing.

Robots con estas características cambiarán nuestro futuro. Suponiendo que evitemos su uso como maquinas de matar, podremos tenerlos en casa para ayudarnos con las tareas de limpieza o incluso con la educación de nuestros hijos. También tendrán trabajo en las granjas, cosechando el cereal y hasta ordeñando las vacas. En las fábricas, un robot con

estas cualidades físicas e intelectuales será mucho más eficiente que los que usamos hoy para, por ejemplo, montar automóviles. Serán más flexibles y capaces de solucionar problemas que aparezcan en las cadenas de montajes. Y todo eso trabajando en tres turnos, las 24 horas del día. De hacerse realidad las predicciones de Antonio, la incorporación de robots evitará la exposición de los obreros a ambientes peligrosos, estresantes o poco saludables, eliminado en fantasma de los riesgos laborales.

Nanomáquinas y cyborgs

Pero no todos los robots serán así de grandes. De hecho, es posible que la mayor ayuda provenga de sus hermanos más pequeños, aquellos que se construirán gracias a los avances de la nanotecnologia. Estos pequeños ingenios, con tamaños micrométricos, podrán hasta ingresar a nuestro organismo y realizar “reparaciones” en nuestras arterias, venas y órganos internos. ¿Tienes una arteria tapada por culpa del colesterol? No necesitarás cirugía, solo un nanorobot que se desplace por su interior y quite la obstrucción.

Durante años se ha especulado con multitud de mecanismos construidos a escala nanométrica. Motores, pistones, trozos de circuitos y casi cualquier cosa que te imagines puede ser construido a escala molecular. Por fin, y luego de millones de horas hombre de investigación, estamos en condiciones de comenzar a fabricar cosas útiles con esta tecnología.

Toda la tecnología que hará posible la existencia de esos maravillosos robots también podrá utilizarse para sustituir nuestras partes defectuosas. Efectivamente, no hay razones para que, disponiendo de brazos robóticos dotados de manos capaces de reconocer superficies mediante el tacto, o cámaras que ven mejor que un ojo, no las utilicemos como piezas de reemplazo en nuestros cuerpos.

El famoso punto de inflexión, o singularidad, propuesto por Kurzweil parece finalmente estar a la vuelta de la esquina. Solo queda determinar su momento exacto, y las consecuencias que ocasionará tanto para los agraciados que se encuentren dentro de ella, como para aquellos que se queden de la nueva brecha tecnológica.

BIOMEDICAL / BIONICS

FEATURE

Good-bye, Wheelchair, Hello Exoskeleton

This year, the Ekso Bionics exoskeleton for paraplegics hits the market

By ELIZA STRICKLAND / JANUARY 2012

In a warehouse that looks like a cross between a mad inventor's garage and a climbing gym, a pair of mechanical legs hangs from the ceiling on ropes. With the quiet whir of four motors, one in each hip and knee, the legs take a step, then another and another. This is an exoskeleton walking suit, and it is taking the hundreds of thousands of steps that regulators demand to prove that it's no mere toy but a reliable medical device, one that just might change the lives of people who thought they'd never again rise from a wheelchair.

The Berkeley, Calif., warehouse is the home of Ekso Bionics (formerly known as Berkeley Bionics), a young company that's about to step out onto the world stage. Early this year the company will begin selling its Ekso suit to rehab clinics in the United States and Europe, to allow patients with spinal cord injuries to train with the device under a doctor's supervision. By the middle of 2012, the company plans to have a model for at-home physical therapy.

TOP TECH 2012 Special Report

When you don the Ekso, you are essentially strapping yourself to a sophisticated robot. It supports its own 20-kilogram weight via the skeletal legs and footrests and takes care of the calculations needed for each step. Your job is to balance your upper body, shifting your weight as you plant a walking stick on the right; your physical therapist will then use a remote control to signal the left leg to step forward. In a later model the walking sticks will have motion sensors that communicate with the legs, allowing the user to take complete control.

"We took the idea of the external skeleton, and we added nerves in the form of sensors and motors that represent your muscles and computers that represent your brain," says Eythor Bender, CEO of Ekso Bionics.

The company began its evolution in 2005 with the ExoHiker, an exoskeleton that allows able-bodied people to carry 90 kg (about 200 pounds) with minimal exertion. The company's engineers at first thought it would take 5 kilowatts to power such an exoskeleton, which would have meant bulky batteries and motors. The breakthrough was a redistribution of weight that reduced the power requirements by three orders of magnitude. A later system, the load-carrying HULC (Human Universal Load Carrier), was licensed to Lockheed Martin Corp. for military development in 2009, and Ekso Bionics' engineers began looking for a new direction. Their energy-efficient devices, they realized, left them with a "power budget" that could be spent on moving the user's legs. That's when paraplegic people became the company's target customers.

Photo: Gabriela Hasbun

HELLO, EKSO User Tamara Mena, who was paralyzed in 2005, gleefully puts her exoskeleton walking suit through its paces. *Click to enlarge.*

A few other companies around the world are bringing out exoskeletons for people with disabilities, but Ekso Bionics' push in 2012 may give it a market advantage. Ten top U.S. rehab clinics have already signed up for the first batch of production units.

One of the first devices will go to Mount Sinai Hospital, in New York City, where Kristjan T. Ragnarsson, chairman of the department of rehabilitation medicine, has treated spinal cord patients for 40 years. His patients' priorities have never changed. "The first thing they want to know is whether they will walk again," says Ragnarsson. "As their physician, I always have to address that question."

Over the years he has told his patients about the latest inventions, from stiff air-filled garments to devices that electrically stimulate the muscles, but all these contraptions proved too difficult for the patients to operate. "They were completely exhausted after just a few steps," he says.

Ragnarsson thinks the Ekso can succeed where so many others failed, because the powered device does most of the labor for the patient. "I'm optimistic, actually, that this will work," he says. "I think my patients will be able to stand up and take a few steps and face the next person directly on!"

In a rehab room at Mount Sinai Hospital in November, Robert Woo, an architect, is justifying his doctor's optimism. His high-tech wheelchair, which his kids have plastered with stickers, sits in the corner. That wheelchair has been his only source of mobility in the four years since a construction crane dropped 6 metric tons of metal on his job-site trailer, crushing his spine.

But today Woo stands tall. He leans forward on the two walking sticks as physical therapists and Ekso technicians buzz around him. Standing in the Ekso feels like floating, Woo explains, and he's still getting used to it. "I'm always looking down, wondering, How are my feet doing?" he says with a smile. "I have to build up confidence in myself and confidence in the device." Then a therapist calls out, "Step!" and punches a button on her controller. With the creak of plastic joints and the whir of motors, Woo is walking.

His progress is slow and jerky, and on one step he leans too far forward and starts to lose his balance before the therapists grab his arms and steady him. But most of the 308 steps Woo takes during the hour-long session are remarkably fluid, with a heel-to-toe gait. Achieving that natural stride was a major technological challenge, as it requires a sophisticated transition of control among the four servomotors that supply the torque to the knee and hip joints.

The session was only Woo's third time using the device, but he had already graduated from a wheeled walker to the walking sticks. He'll have to wait a bit before he makes any more progress: The Ekso was at Mount Sinai for only a week's tryout. Ragnarsson expects the hospital to buy a device of its own this year for the steep price of about US $100 000; that price should fall as production increases.

The company expects to test its physical therapy model soon on patients with other diagnoses, like multiple sclerosis and stroke. And by 2014, it plans to release a personal model that can be used not just for rehab but for everyday living. Woo says he looks forward to a day when he can stand in the Ekso and cook dinner in his kitchen or take a walk in the park with his children.

Before that day comes, though, Ekso Bionics must get approval from the U.S. Food and Drug Administration by proving both the device's safety and its benefits. And its engineers must work out how the user will initiate more complicated movements, like climbing stairs and sitting down.

Photo: Cathy Clarke/Mount Sinai Medical Center

STEPPING OUT At Mount Sinai Hospital, in New York City, Robert Woo uses the Ekso to walk again. *Click image to enlarge.*

Bender, Ekso Bionics' CEO, is confident that controlling Eksos will come to feel utterly natural for the people who rely on them for mobility. "People talk about robots taking over, especially in factories, and taking away our jobs. But the way we see it is, if you can't beat them, join them!" says Bender. "We are joining with the robots. And this has the ability to make us stronger, more productive, and to improve our quality of life."

This article originally appeared in print as "Good-bye, Wheelchair."

Uhfrp p hqg <9#shrsd#uhfrp p hqg#wkl1#Eh#wkh#luvw#ri#rxu#uhqgv1

GREEN TECH / ADVANCED CARS

FEATURE

Plug-in Vehicles Proliferate

Full battery-electric and hybrid-electric plug-in vehicles will start hitting the streets in significant numbers later this year

By DAVID SCHNEIDER / JANUARY 2012

Illustration: Paul Tebbott

Do you own a plug-in car? Do you even know anybody who does? Probably not. But that might very well change this year. Electric cars aren't new, of course. Mitsubishi's all-electric i-MiEV has been available in several countries since 2010. The US $35 200 Nissan Leaf, another pure electric vehicle, has been selling in Japan and the United States for more than a year. The Chevy Volt is also available in America, as is the Tesla Roadster, a $109 000 toy for the rich.

And yet, despite the intense press coverage and feverish anticipation, relatively few of these cars have found their way into garages and driveways. In 2012, however, a big influx of plug-in cars will help these vehicles start shedding their novelty image and in some places may even justify the construction of public charging stations [see "State of Charge," in this issue].

By the end of 2012, for example, Toyota will introduce no fewer than three new plug-in models, and Ford and Volvo will both have their first-ever plug−in electrics in showrooms. Nissan and General Motors will ramp up production of the Leaf and the Volt as well. Nissan aims to produce about 4000 Leafs per month at its plant in Oppama, Japan. Capacity will rise by another 150 000 cars annually when modifications to Nissan's assembly plant in Smyrna, Tenn., are completed later this year, the company says. And Nissan projects that further production increases in 2013 will bring the total to about a quarter million cars a year, which is on par with the number of Volkswagens of all types sold in the United States. If Nissan achieves that kind of volume, the Leaf will have truly broken out of its niche, becoming the Prius of pure electrics.

Chevy's Volt, which costs a little more than $39 000, will also be seen in far greater numbers. Production of the Volt, which last year totaled about 5000 cars, is slated to rise to 60 000 in 2012, with some being sold outside North America.

Even Tesla is trying to go mainstream. This past October, the California carmaker showcased beta versions of an all-electric sedan, the Model S. At $57 400, it ain't cheap, but it's much less pricey than the Roadster.

To scale up its operations, Tesla acquired the New United Motor Manufacturing plant, in Fremont, Calif., where General Motors and Toyota formerly built cars. Tesla projects that later this year, Model S vehicles will be made there at a rate of 20 000 per year.

Much of the money to buy that factory and get it tooled up came from the initial public offering of Tesla stock, which raised $226 million for the company in June of 2010. Just weeks before the IPO, Tesla and Toyota agreed to collaborate on the development of electric vehicles. The two companies began by working on an all-electric version of Toyota's

RAV4, an SUV. Commercial versions of that vehicle will be produced later this year alongside conventional RAV4s at Toyota's assembly plant in Woodstock, Ont., Canada.

This is not the first time the RAV4 has been electrified. Between 1997 and 2003, Toyota produced an earlier battery-powered version of this car, some 1500 of which were sold or leased in California. The later ones included a 27-kilowatt-hour nickel metal hydride battery pack, which provided a range of approximately 160 kilometers.

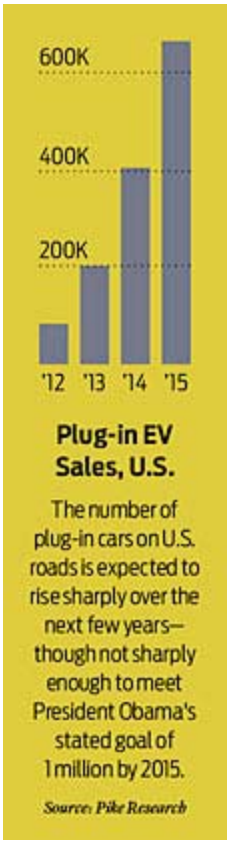

The electric RAV4 design that Toyota and Tesla have worked out will have a similar range but much better acceleration and top speed. That's because it will incorporate Tesla-made drivetrains, including the same kind of battery modules [PDF] found in the Roadster and the Model S. These modules are unique among electric cars in that they contain thousands of thumb-size lithium-ion cells—the kind you might find inside your laptop.

Although other car manufacturers are now also using lithium−ion batteries to power their plug−in vehicles, they construct battery packs using a smaller number of large-format cells. The Leaf's 24-kWh battery, for example, is composed of 192 cells. And the electric version of the Ford Focus, which debuts this year, uses a 23-kWh pack of just 98 cells.

Toyota's decision to use Tesla's many-thousand-cell battery modules for its electric RAV4 is a puzzle. Toyota will be selling two other plug-in vehicles in 2012—an all-electric version of the diminutive Scion iQ and a plug-in version of the company's well-known Prius hybrid—and their battery packs will have the more standard configuration. Why go with something so different for the new RAV4 EV?

Sheldon Brown, executive program manager at the Toyota Technical Center in Saline, Mich., explains the reason for the dueling battery technologies. "We see this as an opportunity to challenge our current way of thinking and to reevaluate our own development processes, ensuring that we are not complacent," he says.

You can make a strong case that the likeliest to succeed is the plug-in Prius, which will go on sale in Europe, Japan, and the United States early this year for $32 000. The reason? It breaks from established technology *the least*.

The car is essentially a clone of the company's third-generation Prius hybrid. But in place of the latter's 1.3−kWh nickel metal hydride battery, the plug-in model contains a 4.4−kWh lithium-ion battery. This battery can propel the car for about 24 km without help from the car's gasoline engine. That limited electric range leaves some plug-in proponents underwhelmed—the Chevy Volt, by contrast, squeezes about 60 km from its battery.

But Mike Ferry, transportation program manager at the California Center for Sustainable Energy, in San Diego, likes the new Prius. He and his coworkers took part in the demonstration and testing program Toyota mounted last year using preproduction versions of this car. Ferry let many of his colleagues use these plug−ins to see how they'd fare. "We have people who live 6 or 7 miles away—all urban driving. They were able to go all week without the gas engine coming on," he says. Most couldn't do that, but on average, his group burned less than 3 liters per 100 km (between 80 and 85 miles per gallon), almost doubling the efficiency of the fuel-sipping regular Prius.

Toyota expects to be able to produce 50 000 plug-in Prius cars annually. While that's a modest number compared with sales of the conventional Prius, which exceed 400 000 cars a year, it's yet another sign that plug-in vehicles are finally going places.

This article originally appeared in print as "Plug-ins Proliferate."

COMPUTING / SOFTWARE

FEATURE

Windows 8: A Redo From Redmond

The 2012 release of Windows 8 will show whether Microsoft can surf the tablet tidal wave

By STEVEN CHERRY / JANUARY 2012

Photo: Alex Gallardo/Reuters

WIN-WIN? Microsoft's Steven Sinofsky gives attendees an early look at Windows 8 at a September conference.

TOP TECH 2012 Special Report

The first time Micro-soft was caught off guard, Bill Gates managed to turn things around by sending the entire company his famous e-mail, "The Internet Tidal Wave." The current challenge—the drift away from PCs in favor of mobile devices—may well be greater.

The move began in 2007, with the introduction of Apple's iPhone, and it was kicked into high gear with the advent of the iPad in 2010 and an Android equivalent in 2011. Suddenly, Microsoft's business was based on yesterday's platforms.

"Their products are fine on notebooks. But on phones and tablets, they're in bad shape," says Michael Silver, an analyst at Gartner, a technology research firm.

Yesterday's platforms are coming to be represented by an ever-narrowing slice of the computing pie. Last year people bought twice as many smartphones as notebooks, and tablets are also set to exceed notebooks by 2016, according to Gartner. Microsoft's quick solution to its mobile-computing problem was Windows Phone 7, but in the third quarter of 2012 the company is expected to launch its more considered answer: Windows 8. The future of the company may depend on it.

"Windows 8 has to be really good," says Michael Cherry, lead analyst for operating systems at Directions on Microsoft, in Kirkland, Wash. (and no relation to the author). "They can't have another Vista," the ill-regarded system that preceded Windows 7. "Especially when you're late to market, you have to leapfrog," he says.

The early indications—which mainly come from a prebeta version of the software that Microsoft handed out at its Build conference, in September—are mixed.

"They have a lot of work ahead of them," Cherry says. "The version they gave out at Build is good for developers, to give them some idea of the changes, to test programs, and so on. But performance is very slow. And they only had code for Intel processors, not ARM." That last point may be critical: Today's smartphones and tablets all use ARM-based microprocessors.

Microsoft's hardest task might be the interface, says Cherry. "It's going to have to do what no operating system has yet done: seamlessly bridge the divide between smartphones and PCs."

James Ashley, a design architect at the Web-development firm Razorfish, in Seattle, thinks Microsoft may have met the challenge. Indeed, he says he was "blown away" by what he saw of Windows 8 at Build. There, Microsoft showed a tablet in a docking station that also had an external desktop monitor. The user can touch the tablet's screen or use a keyboard and mouse to work the external screen.

"They have what looks like two different operating systems side by side. And the part that took everybody by surprise was that you're switching back and forth between [them] casually," says Ashley. That kind of integration of the mobile and desktop experiences might be just the leapfrogging jump that Microsoft needs.

"A lot of people want to use a tablet in the office now," he says. "What Microsoft has done is push the story a little bit - further—we can use a tablet with a docking station and a big screen, so we can have a traditional desktop experience, but at the same time we can just pull the tab let out and take it to a meeting, or home."

The greater threat to Microsoft might be Amazon. Its inexpensive Android-based Kindle Fire, and the ecosystem of movies, books, music, apps, browsing, and shopping that surrounds it, may tempt an unexpectedly large number of consumers away from the PC entirely. Says Cherry, "What if people say, 'Gee, instead of a new laptop, we'll get three Kindles and everyone can have one'? That might be something that no amount of PC-tablet-smartphone integration can overcome."

This article originally appeared in print as "A Redo from Redmond?"

Uhfrp p hqg 56#shrsdn#uhfrp p hqg#wklv1#Eh#wkh#luvw#ri#|rxu#uhqgv1

GREEN TECH / CONSERVATION

FEATURE

LED Bulbs for Less

In 2012, there will finally be a first-rate LED bulb you can afford

By RICHARD STEVENSON / JANUARY 2012

Photo: Philips

BRILLIANT IDEA The Philips L Prize–winning LED bulb draws one-sixth the power of an incandescent and lasts at least 25 times as long.

The passing of Edison's bulb has already been decreed, and which of the two alternatives will replace it is at last becoming clear. It will be the LED.

The success of the light-emitting diode means curtains for the compact fluorescent light (CFL). This clunky, mercury-ridden, hard-to-dim, excessively white device has just two things going for it: It's more efficient than Edison's bulb and, right now, cheaper than the LED-based alternative.

But the LED's quality is rising and its price is dropping—fast. Even now you can pick up a 40-watt-equivalent LED bulb with an appealingly warm hue for just US $9.97. By the end of 2012, a 60–W cousin could be available for about the same price, and within a few years for much less than that.

A glimpse of what's to come appeared this past August, when an LED lightbulb from Philips Lighting North America won the U.S. Department of Energy's $10 million Bright Tomorrow Lighting Prize, better known as the L Prize.

"Our L Prize bulb is essentially the Ferrari of lighting," boasts Todd Manegold, director of LED lamps marketing at Philips. "It does everything that any lightbulb could ever or should ever want to do."

In the beginning—perhaps in the first half of this year—the bulb will sell at a Ferrari price, perhaps $50 apiece. But the pressure to cut that premium will grow as early adopters buy up other brands of LED bulbs that trade efficiency for lower prices. And indeed, those prices are falling fast.

Philips had to design the bulb to hit a slew of engineering targets set by the DOE. First, the bulb had to put out at least 900 lumens—as much light as a 60-W incandescent bulb, the most common kind in the United States. Then it had to last for 25 000 hours, which is roughly 25 times as long as a standard incandescent. And it had to draw less than 10 W.

Philips built its bulb around the Luxeon Rebel, an LED radically different from those that backlit the keypads and displays of early handsets. In a conventional LED, most of the light bounces around within a stack of semiconductor layers, which have a far higher refractive index than air, so only a small proportion of the light goes out in the proper direction. In the Luxeon Rebel, a metallic mirror on the bottom of the chip keeps light from leaking out the wrong way, and the roughening of the top surface allows more of the properly directed light out of the box. As a result, the bulb needs just 9.7 W to yield 910 lumens, whereas an incandescent's 60 W yields only 800 lumens.

The L Prize bulb also features omnidirectional emission, a hallmark of the incandescent bulb. Simply putting a handful of white LEDs into a glass bulb will not lead to uniform illumination, because each device produces a beam. It would be like trying to illuminate a room with a dozen flashlights.

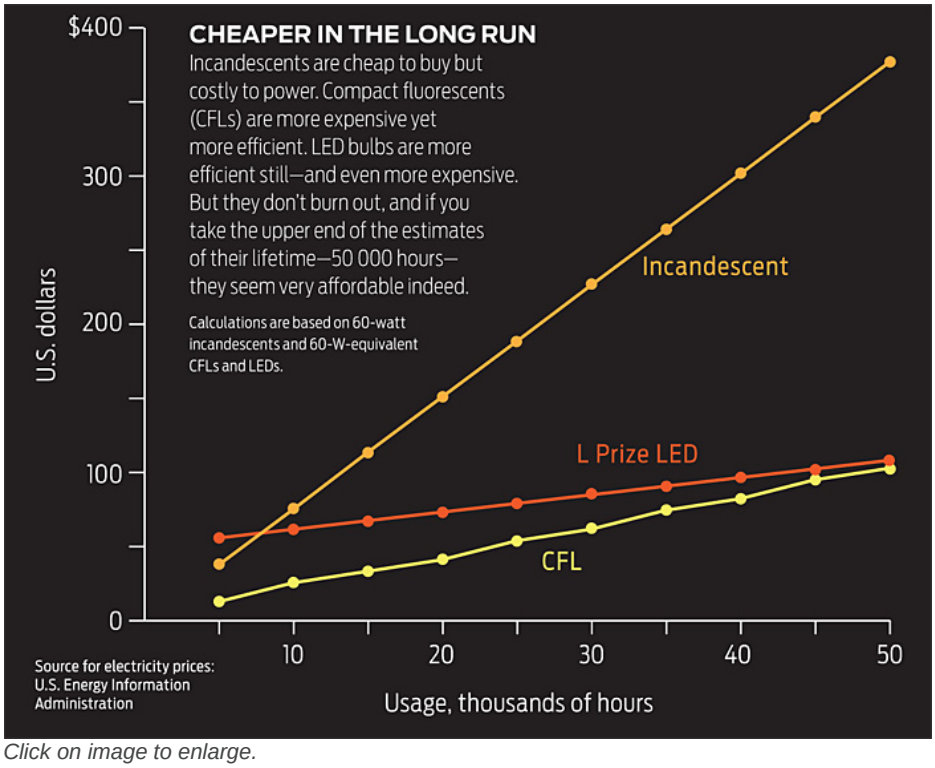

Click on image to enlarge.

The problem comes down to how white LEDs are constructed. Traditionally, white LEDs are made by applying a yellow phosphor to a blue-emitting chip of gallium nitride and indium gallium nitride. That way, the mixing of the yellow emitted by the phosphor and the blue that passes through it combine to create white light, but only from the coated face of the chip. Philips gets around this problem by putting the phosphor close to the outside of the bulb and illuminating it fairly evenly with a battalion of 18 LEDs. "When their light hits the phosphors, it is diffused into a pattern that gives you a clean, uniform look," explains Manegold.

The resulting glow is reasonably similar to what you'd get from an incandescent bulb. Says Manegold: "In our headquarters for lamps, we put an LED like the L Prize, an incandescent, a halogen, and a CFL behind lamp shades and make people play the guessing game." During the past few months he has played this game with many guests, and none of them could tell the difference between the LED bulb and the incandescent.

Philips's bulb also had to work well when shaken about and subjected to temperature extremes, high levels of humidity, and distortions in supply voltage. "We compared [them] to CFL lamps—the better quality ones. All the LEDs made it through, and none of the CFLs made it through," says James Brodrick, solid-state lighting program manager at the DOE.

So if you can't afford the very best LED bulb, what can you expect from the cheaper products that are already out there? One option is Home Depot's EcoSmart range, which is made by the Lighting Science Group Corp. of Satellite Beach, Fla. For $9.97 you can pick up a 40-W equivalent, and by this summer there should be a 60-W version retailing for less than $15. What's more, Lighting Science Group says it will equal the efficiency of the Philips L Prize bulb.

By year's end you can expect a further drop in price. The powerful LEDs used for lighting will be 27 percent less expensive, says the United Kingdom–based market analysis firm IMS Research, and that drop in cost will make a big impact on the price of the bulb.

One reason why it's getting cheaper to make LED chips is that producers are migrating to larger wafers. However, the most important reason is simple economics: The LED market is awash with devices following a production ramp-up in China, which is making great efforts to cut energy consumption, says market analyst Ross Young of IMS Research.

The Chinese government is offering a wide range of incentives to any firm willing to try its hand at manufacturing. Until recently these incentives included a subsidy covering up to three-quarters of the purchase price of the semiconductor fabrication tool needed to grow the nitride films that form the heart of LEDs, a multiwafer MOCVD (metal organic chemical vapor deposition) reactor. No wonder Chinese purchases of these tools have gone through the roof. In 2010 and 2011, shipments hit about 800, quadruple the number in previous years, and about 70 percent of the market.

XV#′5333
Frvw#ri#ulggiqj

Right now, though, most Chinese firms are just producing little glowworms for Christmas lights and the flashing heels of children's sneakers. These outfits have yet to master the "black art" of

d#urrp #ri#wkh p hufxu|#urp d#eurnhq#FIO

epitaxy and, in particular, to learn how to reduce the number of light-quenching defects in nitride films. These defects stem from the 16 percent difference in the spacing of atoms between the films and the sapphire substrates on which they are deposited. Only after these tricks are learned will the Chinese firms produce better material. After they do, they'll have to develop light-extracting LED architectures like that of the Luxeon Rebel to make devices bright enough for use in lightbulbs.

In the meantime, China's flooding of the low-end market with cheap, low-power LEDs is having an effect on the high end, according to Young. When the LED oversupply began, in the second half of 2010, it drove other LED makers out of their bread-and-butter market in display backlighting—and into general lighting. The domino effect has thus created a surplus in that market, too. China's efforts continue, and so the prices of all LEDs will fall throughout 2012.

Will this price decline immediately make the LED lightbulb a worthwhile purchase? The answer depends on what you're screwing into your sockets today. If you still use incandescents, switching to an LED bulb makes a lot of sense. That's because it costs about $180 to run a 60-W incandescent for 25 000 hours; in comparison, the electricity bill for an LED bulb producing the same amount of light is between $25 and $50. Besides, that 25 000 hours of lighting will require 25 incandescents as opposed to just one LED bulb.

However, if you now use CFLs—as is the case in most of Europe—it will be harder to decide whether 2012 is the time to switch to LEDs. CFLs are still much cheaper than LED bulbs, and this easily makes up for their 10 000-hour lifetime.

In any case, this year you will for the first time be able to afford an LED bulb that's clearly superior to a CFL. It will give off a nice warm glow, work with your dimmer switch, use energy frugally and when you finally replace it after 15 years, you can just throw it in the dustbin.

By then, you'll struggle to recollect what the letters "CFL" ever stood for.

This article originally appeared in print as "LEDs for Less."

About the Author

Richard Stevenson is himself an expert on the raw materials that go into such devices: compound semiconductors. He got a Ph.D. at the University of Cambridge studying these materials. Then he went into industry and made them. Now, as a freelance journalist based in Wales, he writes about them. (His previous feature for *Spectrum*, however, was on good old silicon: "A Driver's Sixth Sense," October 2011.)

Uhfrp p hqg 46#shrsd#uhfrp p hqg#wklv1#Eh#wkh#luvw#ri#|rxu#uhqgv1

AEROSPACE / ASTROPHYSICS

FEATURE

Single Blue Planet Seeks Same

In 2012, a new exoplanet hunter will look for worlds like our own

By RACHEL COURTLAND / JANUARY 2012

Photo: Gianni Tessicini/INAF-FGG

NEW HORIZONS Italy's Telescopio Nazionale Galileo on La Palma, one of the Canary Islands, will host a new planet-finding instrument. *Click on image to enlarge.*

TOP TECH 2012 Special Report

Francesco Pepe leans over a large black chest in the basement of the Geneva Observatory. He lifts the lid to reveal a length of rubbery fiber-optic cable coiled along the bottom and surrounded by steep cliffs of foam. "I think they may have gone a little overboard with the packaging," Pepe says.

Dwarfed by its container, the cable doesn't look that impressive, but it will soon be a key component in the hunt for planets like our own. In just a few months, astronomers will use the fiber to feed starlight into a new detector—the High Accuracy Radial velocity Planet Searcher–North (HARPS-N), an ultra-precise spectrograph that is being assembled nearby, not far from a field of grazing horses.

In the coming months, Pepe, an astronomer, and his colleagues will take apart the go-cart–size instrument, box up the pieces, and put them together again in a room near Italy's Telescopio Nazionale Galileo on La Palma, one of the Canary Islands.

When HARPS-N begins observations in April, astronomers expect it will be the most precise planet hunter in the northern hemisphere. They also reckon it will be particularly well placed to help bag the most coveted extrasolar quarry of all—planets small enough to have rocky surfaces and cool enough to have liquid water. The -reason? HARPS-N will not be working alone. It will view a particular patch of the sky that has been the singular focus of NASA's Kepler telescope for almost three years.

Kepler, which trails behind Earth in a solar orbit all its own, is on the hunt for the subtlest of flickers—periodic dips in stellar brightness caused by planets that pass, or "transit," in front of their host stars and briefly block about 0.01 percent of the light. And by all accounts, the spacecraft's dogged focus on a single region of the sky has paid off nicely. In December, Kepler's astronomers announced the discovery of a planet orbiting far enough from its sun-like host to be able to boast liquid water.

With a little luck and enough time, Kepler will find many more such planets, some quite similar in size to Earth.

Photo: Rachel Courtland

PRECISION WORK Astronomer Francesco Pepe stands beside the partially assembled HARPS-N spectrograph.

To confirm its discoveries and learn as much as possible about these planets, Kepler needs help on the ground. Unless the space telescope spots a close-knit system of planets that influence one another's orbits, Kepler's measurements can reveal only the size of a planet, not its density. To distinguish a gas giant like Neptune from a rocky super-Earth, astronomers need a tool capable of measuring mass precisely—a spectrograph.

A spectrograph uses a diffraction grating to separate incoming starlight so that different wavelengths hit different parts of a charge-coupled device (CCD). By a trick of inference, astronomers can use changes to surrounding pixels to pinpoint—with subpixel precision—the shift in any particular peak or trough in a star's spectrum. These shifts can occur if the star is moving toward or away from Earth. If a detector shows a star wiggling back and forth, and if these wiggles occur with regularity, the signal is a good indication that there is a planet in orbit, tugging on its host.

HARPS-N should be able to detect changes in the position of spectral lines to within 15 nanometers, which is about one-thousandth the width of a single pixel on its 6-centimeter-wide, 16-megapixel CCD. That's enough to pace the to-and-fro motion of a star to within 1 meter per second—about walking speed—from hundreds of light-years away.

Even that's not quite good enough to spot something like Earth from afar, says David Latham, a member of the HARPS-N team at the Harvard-Smithsonian Center for Astrophysics in Cambridge, Mass. But he says HARPS-N will still be more than twice as precise as the Kepler team's main follow-up instrument, the High Resolution Echelle Spectrometer, which sits atop Mauna Kea, Hawaii at the W.M. Keck Observatory.

As a result, HARPS-N and Kepler together will be able to offer portraits of planets that are on the edge of habitability. Chances are good, team members say, that HARPS-N will be able to see rocky planets that are roughly a few times the mass of Earth and on orbits that lie within the ”Goldilocks” zone, where temperatures are just right for life. ”We're not going all the way to characterizing true Earth twins, but we are pushing in that direction,” says Latham, who helped to get the HARPS-N project started. ”It's a big step.”

The detector's main components are nearly identical to those of its progenitor, the HARPS detector at the European Southern Observatory's high-altitude La Silla Observatory, in Chile. But HARPS-N will also get a few improvements that could help boost its sensitivity.

The telescope's new fiber is one of the key changes. The original HARPS instrument uses a circular fiber to feed starlight into the spectrograph. But the position of the light that exits the fiber can shift depending on how a star is positioned in the telescope, and that can register as movement of spectral lines on the detector. ”If [the lines] move because of differing illumination, that is catastrophic for us,” Pepe says. To remedy the situation, HARPS-N will use a fiber with an octagonal cross section. The team hopes that the eight facets will scatter and homogenize the light so that it will always exit the fiber evenly.

A few months after installation, HARPS-N will play guinea pig for another potential improvement, a new kind of light source called a laser frequency comb, which could be used to calibrate the detector. Today many astronomical spectrographs are calibrated using the spectral lines emitted by thorium argon lamps, which have irregular lines that vary widely in brightness and spacing and shine on only a small fraction of the CCD cells in a detector. A laser comb instead uses brief pulses of laser light to create millions of evenly spaced lines over a wide range of wavelengths, which can be tuned to calibrate every pixel in a spectrograph's CCD. A comb calibrator can also be made very stable by matching key frequencies with those of an atomic clock. The HARPS-N team will test one incarnation of the laser comb, which uses a titanium sapphire laser, while others are working on fiber laser-based systems.

In the next few years, astronomers hope these combs could help improve the precision of stellar spectrographs by a factor of 10, if not more. If there are no confounding factors—not least the natural fluctuations in a star's atmosphere—that should be precise enough to detect the planet of our dreams: a true Earth twin.

This article originally appeared in print as "In Search of Alien Worlds."

About the Author

For more about the author, see the Back Story, "A Planet Like Our Own."

Uhfrp p hqg Eh#wkh#iluw#ri#rxu#uhqgv#wr#uhfrp p hqg#wkv1

TELECOM / WIRELESS

FEATURE

Fantastic 4G

Hundreds of telecoms will invest in 4G LTE networks in 2012

By MARISA PLUMB / JANUARY 2012

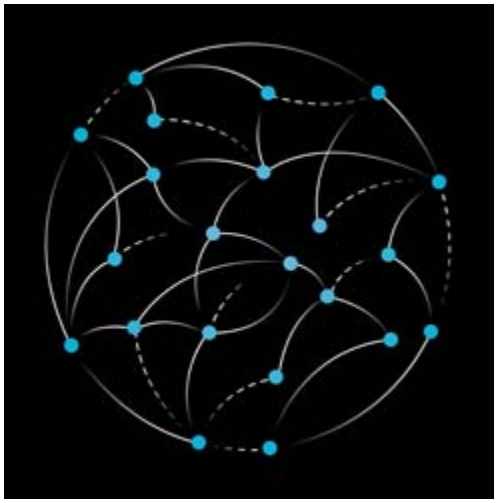

Illustration: Paul Tebbott

It's 5:00 in the afternoon. Do you know how much data your smartphone apps are sucking out of the ether?

It's probably at least 10 megabytes per hour, and it may be as much as 115 MB/h, according to a recent study by the British firm Virgin Media Business. In other words, depending on what you're using it for, your phone or tablet might be consuming—or more likely struggling but failing to consume—the data equivalent of about 100 medium-size books every hour.

If you're on a 3G network, you have our sympathies. But cheer up: 4G is coming to a cellphone tower near you, if it hasn't already. With a conformity rare in technology, let alone in perennially fractious telephony, the world's wireless operators are falling in line behind a 4G standard called LTE, which stands for Long Term Evolution. According to the market research firm iSuppli Corp., the number of 4G LTE subscribers will grow by 400 percent this year, and about 10 percent of global wireless subscribers will have LTE connections by 2015.

Today's LTE networks deliver data download rates about 10 times those of 3G while making more efficient use of the radio spectrum. Basically, 4G LTE can keep up, if just barely, with the soaring data demands of the fast-growing ranks of ever more sophisticated smartphones.

Most smartphones can already accommodate data at rates that exceed the networks' capabilities, says Paul Kapustka of Sidecut Reports, a service that analyzes carrier technologies. Given the chance, broadband phones will choke their wireless networks to death. Telecom equipment makers such as Alcatel-Lucent and Ericsson have convinced network operators around the world that LTE is the technology that will keep their networks moving fast.

The Global Mobile Suppliers Association reports that 4G LTE networks are being built by 280 operators in 90 countries, and dozens of these networks will light up the airwaves in 2012. It's easy to see why. Increasingly, people are using their phones and tablets to stream music, movies, and television shows; check in with their social networks; place two-way video calls; and even have virtual visits with their doctors. All those bandwidth hogs are more reliable and fun to use with LTE. Indeed, within a year of the Swedish telecom TeliaSonora taking its 4G LTE network live, the average users' monthly download total rose to about 15 gigabytes from about 3 GB. LTE is fast enough to support such features as high-definition video. So Phil Solis, research director of mobile networks at ABI Research, predicts a sharp rise this year in users of video call services like Skype, FaceTime, and Fring.

LTE networks are not all created equal, but early rollouts are more than validating backers' claims of a tenfold speed increase over 3G. In Australia, a 3G connection averages just under 2 megabits per second, says Mike Wright, executive director of Telstra Corp. networks and access technologies. When Telstra, the country's largest telecom, launched a 4G network in 2011, average download speeds rose to 20 Mb/s. In Dallas, between April and October of 2011, Verizon launched a new LTE network and saw a jump from around 0.75 Mb/s to between 11 and 14 Mb/s, according to data analysts at RootMetrics. That's about double the speed most Americans get through their cable and DSL broadband services.

It's good news for the millions of people who primarily use their cellular connections to access the Web. And the good news doesn't stop there. Mobile connections at 4G speeds are supercharging older platforms as well: Forty percent of today's 4G consumer devices are routers that bring the Internet to one or more computers, and 25 percent are dongles for laptop computers.

LTE and 4G have not always been one and the same. LTE was once one of two standards marketed as 4G. However, decisions in the past year have effectively edged out LTE's competitor, WiMax. The U.S. operator Sprint, which initially deployed WiMax, will build out a separate LTE network this year. Even though the company's WiMax service will be around for some time, nobody expects it to be a long-term solution. LTE's triumph is all the more remarkable considering that WiMax had a head start.

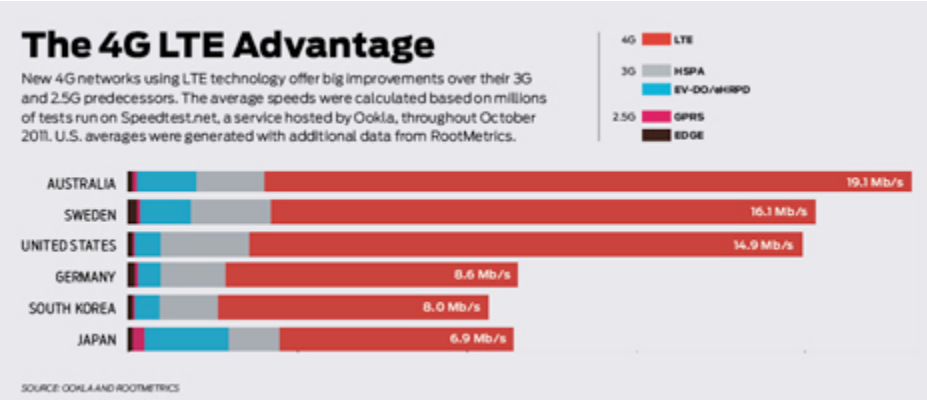

Click on image for the full graphic view.

LTE and WiMax both qualify as 4G networks because they each use an all-Internet Protocol scheme to streamline their architectures and boost data rates. In the future, they'll be capable of delivering data, including digitized voice, in packets, just like on the Internet. For now, LTE is capable of handling only data. Voice calls are handled on an operator's 3G or 2G network or by applications that use a different layer of the network to deliver voice over IP (VoIP), like Skype. But upgrades that will start this year and continue through 2014 will pave the way for always handling voice exactly like data.

In a typical IP network, packets are assigned labels that indicate how they fit together, so that they can be sent over any available channel and reassembled when they reach their destination. To access a website, a smartphone on an LTE network connects to a base station, or an "evolved node" in LTE parlance. This station is typically a 3G base station upgraded with more processing power to handle packet-switched data and radio technology for LTE's different bands of spectrum. For now, an individual LTE base station works on one spectral band, but there will be an upgrade to allow for multiband evolved nodes in the next few years.

43{
Lqfuhdvh#lq#wkh#
gdwd#grz qordg#
udwh#ru#OWH#7J#
frp sduhg#wr#6J

Packetized data then streams between the node and smartphone over the widest available of several channels within a spectral band. The evolved node then connects to what's called an evolved packet core. The EPC, which processes data and sends and receives video images, Web pages, and soon, voice, can connect to all kinds of other networks.

Because 4G uses IP, like the Internet itself, the data's trip from the base station to the Internet or other networks can use any means—copper, fiber, terrestrial microwave, or satellite. But before it can do that, the data stream must first be tweaked by another layer of technology called the IP Multimedia Subsystem. IMS helps optimize bandwidth for each application, for instance, by guaranteeing a certain maximum latency for VoIP traffic.

The LTE scheme is a lot simpler than 3G, which requires a temporary but dedicated connection between a device and a base station to transmit information. That connection is divided into time slots, with each device occupying its own slot, even when there's silence during the conversation.

Although WiMax handles voice just as it does any other data, the carriers that tried it out are now switching to LTE. So why is LTE emerging as the clear favorite? One reason is capacity. The LTE specification promises peak download rates of 100 Mb/s. WiMax could match that speed in the future, but early versions fell short and don't use a carrier's available spectrum as efficiently as LTE. The latest version of WiMax is more efficient, however, so the performance gap between the two technologies is likely to narrow.

According to Solis of ABI Research and Kapustka of Sidecut Reports, one big reason LTE will remain ahead is a campaign against WiMax by the companies with the biggest stakes in LTE technologies. Qualcomm, for example, owns many patents for LTE technology, while one of its biggest rivals, Intel, has been the driving force behind WiMax. Alcatel-Lucent and Ericsson have also invested a lot in LTE equipment and were highly motivated to push LTE in influential U.S. markets, says Kapustka.

It is worth noting that neither technology's maximum download speed is a reality yet. Clearwire Corp., which provides Sprint with the most extensive WiMax network in the United States, delivers average speeds under 7 Mb/s. So far, commercial LTE networks have posted much higher average speeds. In some countries, like Australia and Sweden, average LTE download rates are more than 20 Mb/s in urban areas.

All-IP networks are inevitable, and operators would like to install them everywhere overnight. But the upgrade will be gradual, and 2G and 3G networks will be around for years to come. It's not that upgrading a 3G base station to LTE is always a terribly demanding or disruptive process, says Todd Rowley, vice president of 4G technologies at Sprint. It's just that during the deployment, the entire coverage network and its devices need to work seamlessly with both LTE and the existing 3G networks.

Because 4G and 3G can't occupy the same bands, there's been a scramble for spectrum. Carriers will have to either buy new spectrum from other companies or at auction, or they'll have to repurpose bands from less-populated 2G networks.

It's been easier to manage such spectrum swaps in some countries than in others. Australia, for example, seems to have handled things with minimal disruption: Telecom giant Telstra was able to repurpose its 2G band, at 1800 megahertz, for its 4G network because its base of 2G subscribers had dwindled significantly, says Telstra's Wright. Even better, the 1800-MHz band is the LTE standard for Europe and other countries, so radio gear designed for it is readily available.

5;3
Qhwz run#
rshudwruv#
lqyhvwlqj#lq
7J#OWH
qhwz runv

Elsewhere, the transition hasn't been so smooth, because telecoms had to horse-trade to get spectrum for 4G.

"LTE is enabling more capacity for data services, but it's also increasing the range of scenarios that a device needs to work in," says Bill Davidson, senior vice president of global marketing and investor relations at Qualcomm. "Different countries have auctioned off the 700-MHz band, then the 2.5-gigahertz band, then 800 and 900 bands of spectrum over the last few years, and that's made it more complex for us." To cover all the bases, Qualcomm's latest high-end handset chips can accommodate up to 40 bands. And Ericsson, which makes 4G base stations, sells equipment for at least 15 different bands.

These challenges aren't stopping anyone. But telecoms have to get their new systems up in a hurry—even as LTE networks are being rolled out in major urban centers, LTE's successor is already being planned: LTE-Advanced promises faster data rates, wider channel bandwidths, more advanced antennas, and support for a larger number of low-power "picocells" to increase coverage. It will offer another boost in efficiency that's just as irresistible as was LTE's.

"It's the drive for spectral efficiency that will always motivate operators to use the newest technologies," says Terry McCabe, chief technology officer at Mavenir Systems, which provides VoIP for LTE networks. And so networks, even those meant for the long term, will continue to evolve.

Uhfrp p hqg 65#shrsdu#uhfrp p hqg#wkl v1#Eh#wkh#luvw#ri#rxu#uhqgv1

AEROSPACE / AVIATION

FEATURE

Private Spaceflight: Up, Up, and Away

This year, commercial spaceflight will really take off

By JAMES OBERG / JANUARY 2012

Photo: Virgin Galactic

FUN RIDE Virgin Galactic's *SpaceShipTwo*, attached to its carrier, *WhiteKnightTwo*, should be making test flights this year.

The 20th of February, 2012, will be the 50th anniversary of the first U.S. orbital manned spaceflight. To mark the occasion, retired pilot Craig Russell had an over-the-top idea: Reenact astronaut John Glenn's mission, but do it with private funding and off-the-shelf technologies.

Ultimately, a lack of funding killed Russell's dream, but don't lose heart. Truth is, if you've got a more practical reason for putting a person in space, there's never been a better time to try. Over the last decade, a broad advance in the commercial availability of aerospace technologies has allowed small private entities to attempt feats that once had been the monopoly of major governments.

TOP TECH 2012 Special Report

In 2012 privately funded human spaceflight will advance from promises and one-off stunts to serious flight-testing of spaceships. Governments will be the biggest customers, with unmanned systems possibly docking with the International Space Station (ISS) this year and perhaps eventually taking the place of the retired U.S. space shuttles. Meanwhile, spacecraft designed to give well-heeled tourists a thrill will be firing up their rockets, letting their passengers enjoy a few minutes of weightlessness, and gliding in for landings.

Indeed, this could be the year that spaceflight moves beyond the 1960s inspirational phrase "man in space" toward a more inclusive one: "*Any* man or woman in space."

In spaceflight, as in many other fields, there's an advantage to being first. Virgin Galactic has that advantage. Its *SpaceShipTwo* is the roomier follow-up to the craft that won the Ansari X-Prize in 2004 for crossing the legal boundary of space (100 kilometers). The British company has already sold nearly 500 tickets at US $200 000 apiece, and it opened a spaceport this past October in New Mexico.

In 2012, "our hope, our plan, is to do powered flight tests, and if things go well, we have a shot of getting into space," says CEO George T. Whitesides. Using the large *WhiteKnightTwo* aircraft as a carrier, the rocket plane has already begun unpowered drop tests from its development base in Mojave, Calif. And it recently concluded ground-based engine burns that were nearly long enough to put a vehicle into space. Full-duration burns will last between 60 and 70 seconds during an ascent that involves a little less than 4 *g*'s of acceleration.

Actual test flights will be made in small steps, Whitesides says. Turning the engine on and off will be a first step, followed by a 10-second burn, a 20-second burn, and so forth. When asked how many test flights there would be before the first tourists fly, he laughed, saying, "As many as you need."

The other tourist spaceship that is expected to get airborne in 2012 is the *Lynx* rocket plane from Xcor Aerospace, in Mojave. Its concept is less grandiose than *SpaceShipTwo*'s, with a single passenger sitting in a cockpit next to a single pilot and with initial flights to altitudes of just over 60 kilometers. Company spokesman Mike Massee says the rocket plane's propulsion system was designed for speedy testing and will eventually allow for two flights per day using a single ground crew.

The Year in Space, Almost

Company	Spaceship	2012 Goals
Virgin Galactic	*SpaceShipTwo*	Powered test flight
Xcor Aerospace	*Lynx*	First test flights
Blue Origin	*New Shepard*	Unknown
SpaceX	*Dragon*	Two resupply trips to the International Space Station
Sierra Nevada Corp.	*Dream Chaser*	Drop tests
Excalibur Almaz	A variant of a Soviet spacecraft	Design and safety reviews

You can also expect to hear more about Amazon.com founder Jeff Bezos's secretive spaceflight company, Blue Origin, based in Kent, Wash. Little was known about its spacecraft, *New Shepard*, which takes off and lands vertically, until an unmanned supersonic test flight this past summer ended in an emergency engine shutdown and the crash-and-burn loss of the vehicle. Another craft from the company is expected to fly this year and at some point could carry one or more test pilots. While it's unknown what altitude pilots will be aiming for, breaking the 100-km barrier is the likely goal for selling seats to passengers.

Up and down isn't the only human spaceflight game in town. The more important work, from a commercial and scientific perspective, is orbital flight. And that arena, too, could pass real milestones this year. Some suborbital tourist firms have dreams of eventual orbital flight, but others are going straight for it.

In terms of spacecraft engineering, the step from short suborbital missions to orbital flight is about two orders of magnitude of energy, vehicle stress, and general complexity. Private ticket prices reflect that: Initial suborbital *SpaceShipTwo* prices start at $200 000, while space tourist trips to the ISS in a Soyuz have ranged from $20 million to $40 million.

NASA hopes that at least one corporate team will succeed in developing an Earth-to-orbit "space taxi" in the next few years, so it can stop paying Russia $65 million per seat on Soyuz vehicles. If the price is low enough, that same operator could someday also sell tickets to other orbital destinations, should there ever be any—and some are under consideration. But at recent hearings in Washington, D.C., some U.S. congressmen expressed serious skepticism about whether a private market for such flights would ever evolve while the ticket prices remain in the multimillion-dollar range.

That's why SpaceX's *Dragon* flights this year are so important. The firm, based in Hawthorne, Calif., plans two unmanned resupply runs to the ISS. If successful, these cargo flights will help establish a track record that could eventually lead NASA to approve a SpaceX craft for manned flights. Before that happens, though, SpaceX needs to finish developing and testing its launch escape system.

Two other orbital projects in the works include a shuttle-shaped craft built by Sierra Nevada Corp. that repurposes a surplus Soviet-era design, and a modernized version of the Apollo capsule from Boeing. Sierra Nevada's *Dream Chaser* will be tested in the second half of the year. It will be carried to altitude and dropped by Virgin Galactic's *WhiteKnightTwo*, and it will then glide under autopilot and land itself. Future flights could be flown by test pilots, but Sierra Nevada has not been forthcoming. These tests,

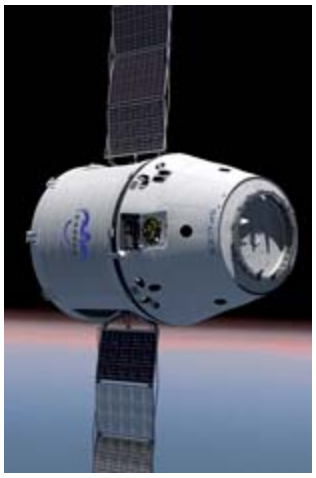

Illustration: SpaceX

DRAGON, FLY SpaceX's Dragon capsule will see testing this year.

say experts, should help dispel persistent concerns among some spaceflight professionals about the design's controllability. Boeing's craft, the CST-100 (for Crew Space Transportation), is further from flight-testing.

Little is known about the spacecraft from Excalibur Almaz, except that it's based on a Soviet manned vehicle, with modernized innards, and is launched by a commercial booster. Company official Arthur Dula states that it "could provide a crew vehicle two years earlier than the current NASA plans." The company, headquartered on the Isle of Man, recently inked a development agreement with NASA and will spend 2012 on design and safety reviews, not flight-testing.

Harsh reality will eventually sort out the winners from the losers, but for now the field of contestants is satisfyingly broad, as they exploit combinations of classic, proven ideas with bold new innovations. Successes and failures are likely to be spectacular, as the competing projects attempt to push open a wider way for human access to space. Expect a sky full of excitement this year.

This article originally appeared in print as "Up, Up, and Away."

About the Author

James Oberg, worked as an aerospace engineer at NASA for 22 years. He switched to journalism in the late 1990s and now makes his living reporting on space for such outlets as *Popular Science*, NBC News, and of course, *Spectrum*.

Uhfrp p hqg 8#shrsd#uhfrp p hqg#wkl1#Eh#wkh#luvw#ri#|rxu#uhqgv1

COMPUTING / HARDWARE

FEATURE

China’s Homegrown Supercomputers

In 2012, China’s chips will power the Dawning 6000

By JOSEPH CALAMIA / JANUARY 2012

Photo: Xu Suhui/Xinhua/Landov

GREAT WALL The Sunway BlueLight supercomputer is the first deployed using processors designed in China.

In late October 2011, the Sunway BlueLight MPP made headlines as China's first high-performance computer to harness the power of a homegrown chip, the ShenWei SW1600. And the Dawning 6000, scheduled to come on line in December 2011, will use another indigenous processor, the Godson-3B. These are supercomputers that China can truly call its own.

At press time, engineers were busy optimizing the Dawning 6000's ability to run Linpack—the benchmark software library used to rank computers in the Top500 list. The Sunway BlueLight, a petaflops-level supercomputer, has already been put through its Linpack paces and claimed 14th place in the November Top500 ranking. But don't be fooled: Neither machine is a speed demon. Consider them rather as steps toward technological independence.

TOP TECH 2012 Special Report

"The Dawning 6000 is really trying to master the tricks of this domain so that the Chinese have the ability to develop their own chips, their own IT from the ground up," says IEEE Fellow Tarek El-Ghazawi, a professor of electrical and computer engineering at George Washington University and a codirector of the NSF Center for High-Performance Reconfigurable Computing. Given that today's exotic supercomputer components are tomorrow's quotidian hardware for personal computers, El-Ghazawi predicts that these research projects will prove a boon for Chinese commercial chips, which he expects to become widespread in China's marketplace in around 10 years. "Then, in the next 20 years," he says, "they may be selling chips to the world, including the U.S."

That would be a swapping of roles. The Dawning 5000 line, released in 2008, relied on U.S.-made AMD Opteron CPUs. And even China's Tianhe-1A, for a few months the world's top-ranked supercomputer, owed a good part of its 2.57-petaflops performance to Western chips—a total of 7168 Nvidia Tesla GPUs complemented by 14 396 Intel Xeon CPUs.

"The Tianhe was opportunistic," El-Ghazawi says. "They looked at the top-performing chips out there and applied them. With Dawning, from the ground up, they are building a machine with careful consideration to each level of the architecture—chip, node, and system—with the requirements of the software in the back of their minds."

The Tianhe-1A did not sacrifice all innovation in the race for the top. The machine was also celebrated for its indigenous interconnect system, the channels for shuttling information between computer nodes. The interconnect system, called

Arch, was developed by China's National University of Defense Technology. Capable of 160 gigabytes per second, Arch had greater bandwidth than commercially available alternatives, such as InfiniBand.

"If you're developing your own supercomputer, you would have to build both your own processors and interconnect to connect them together," says Jack Dongarra, a professor of electrical engineering and computer science at the University of Tennessee who helps to compile the Top500 ranking. "I would guess that the Chinese would want to move toward a system that they have developed themselves....They want to be in a position where they can develop an industry that can generate computers for China and the rest of the world rather than relying on Western components."

Different supercomputers represent different strategies, argues David K. Kahaner, founding director of the Asian Technology Information Program, headquartered in Albuquerque. For example, the Tianhe-1A, still the fastest machine in China, and the Sunway BlueLight have roots in defense research, while the Dawning 6000 might be thought of as an academic research supercomputer. "China is a big country with a tremendous number of capable people, and they are striking out in a number of directions," Kahaner says. "Competition is good for everybody." Right now, he sees the BlueLight as the most indigenous, noting its use of both homegrown chips and a unique water-cooling system.

Photo: Xu Suhui/Xinhua/Landov

HOMESPUN LOGIC The ShenWei SW1600, which was designed in China, powers the Sunway BlueLight supercomputer.

None of the machines has completely broken away from Western influences. For example, the Dawning 6000's Godson-3B processors, from the Chinese Academy of Sciences, appear to use a Western instruction set—the CPU's overarching architecture as it relates to programming. The MIPS instruction set, which Godson uses, is more commonly found in the microprocessors inhabiting television set-top boxes. Mark Pittman, MIPS Technologies' vice president for Asia and Pacific sales, says the Chinese Academy of Sciences was one of the first groups to target the MIPS instruction set for high-performance computing. Early in the project, the Godson researchers used the instruction set without a license, but Pittman notes that this issue has been resolved, adding that since 2009 his company has directly licensed MIPS to the academy's Institute of Computing Technology.

"Leading researchers in China feel it is better to innovate a microprocessor with an existing instruction set rather than take a long time to develop an instruction set and then innovate a processor based on that," Pittman says. Developing a new instruction set would require porting existing operating systems, programs, and drivers. "All the things that run on MIPS would have to be re-created for a new instruction set, and even in China that's an expensive proposition."

The BlueLight's ShenWei SW1600, made by the National Research Center of Parallel Computer Engineering and Technology, is rumored to have taken one step further toward independence. Though earlier ShenWei chips used a modified Alpha instruction set, designers claim this processor uses an architecture of their own design, Kahaner says.

Both machines may also still use Western interconnect systems. Kahaner confirms that the BlueLight uses a modified InfiniBand interconnect system. El-Ghazawi says a prototype node of the Dawning 6000 also used a modified InfiniBand network, combining it with a specialized network for performing frequent tasks more efficiently. "In the end, we may see some new, efficient kind of Chinese switching that adheres to the InfiniBand standards," he says, adding that the Dawning's interconnect system is "very forward looking."

The designers of the processors were also prescient in the stress they've put on conserving power. "The race to exaflop computing will be a race to energy efficiency," says Steve Scott, now chief technology officer of Nvidia's Tesla unit and formerly a senior vice president of Cray, a pioneering supercomputer company.

The Godson-3B, capable of 128 billion flops using just 40 watts, claims almost double the peak power efficiency of some U.S. competitors. At press

Chinese Supercomputers

Supercomputer	Developer	Debut in Top500	Processors
Tianhe-1A	National University of Defense Technology	November 2010	Nvidia Tesla GPUs and Intel Xeon CPUs with custom interconnect
Sunway BlueLight MPP	National Research Center of Parallel Computer Engineering and Technology	November 2011	ShenWei SW1600 CPUs with an InfiniBand QDR interconnect
Dawning 6000	Chinese Academy of Sciences and Dawning Information Industry	June 2012 (expected)	Godson-3B CPUs

time, however, the Godson's energy efficiency had yet to be tested using a standard benchmark like Linpack. And, Dongarra points out, the processor is only one part of an energy budget that also includes interconnects and memory. The BlueLight's complete system also turned heads with its efficiency. The system can perform 741 megaflops per watt, compared to 636 Mflops/W for the Tianhe -1A, Dongarra says.

Researchers at the Chinese Academy of Sciences are already upping the efficiency of the next class of processors, the Godson-3C. According to one of the chip's architects, Yunji Chen, the 3C will have an even higher performance-to-power ratio, mainly because the processor will be built using a 32-nanometer fabrication process as opposed to the 3B's older, 65-nm process and because the new chip will feature an improved three-level cache memory.

Even more energy could be saved by moving from CPUs to GPUs, as other high-performance computers have done. Such graphics processing chips—now used in general computing—can do simple operations on great gobs of data in parallel, rather than in a more one-at-a-time fashion as a CPU core does. That parallelism economizes on energy. Nvidia's Scott says that an Intel Westmere CPU takes about 1.7 nanojoules per operation at peak performance, while an Nvidia Fermi GPU takes less than a seventh of that. Though China has produced some homegrown midlevel GPUs, the Chinese Academy of Sciences appears to be focusing its efforts on the Godson line of CPUs.

But favoring CPUs may just be another part of China’s strategy, says El-Ghazawi. He notes that it isn’t nearly as easy to program for GPUs, and Chinese supercomputer makers are looking not merely for speed records but for market share. “Although they are a latecomer,” he says, “they are really hitting the ground running.”

About the Author

Joseph Calamia is a freelance writer based in New Haven, Conn. Despite the geographic disadvantage, he was the natural choice to write "China's Homegrown Supercomputers", because he'd reported on the development of the key microprocessors involved less than a year ago. A frequent contributor to *IEEE Spectrum* and an alumnus of the MIT Graduate Program in Science Writing, Calamia has also written for *Discover* and *Popular Mechanics*.

Uhfrp p hqg Eh#wkh#iluw#ri#|rxu#uhqgv#wr#uhfrp p hqg#wkh1