BREAK THE MOULD

Silicon chips? So last century. They can't get any smaller, and they only think in straight lines. What comes next, asks *Andrew Smith*. A weird world where algae, mushrooms and microbes do our computing for us

Until the mid-20th century, there were only two certainties in life: death and taxes. Then out of the blue came a third — a descriptive "law" that anticipated the motion of science and society, and defined an era: Moore's Law. Moore's Law began as a prediction by a computer scientist named Gordon Moore in 1965 and was named after him retrospectively. For the foreseeable future, he decreed, rapid advances in computing technology were likely to see the silicon-based microprocessors at the heart of our machines double in speed and capability every 18 months.

Moore was right, which is how you come to have more computing power in your pocket — and quite possibly at home in your washing machine — than Nasa could muster on their way to the moon in 1969. Moore's Law has been so constant that most of us take it for granted. Try to imagine a world in which your next smartphone *doesn't* allow you to do, experience, store far more, far faster, than your present one. Chances are you can't.

Start trying. For three decades, experts have warned that Moore's law would eventually break down and the evolution of silicon-processing would first slow, then stop, leaving our current technologies stranded. The day has arrived: Moore's Law is collapsing. Engineers at ARM, the Cambridge firm whose chips control most of the world's smartphones, have just created a functioning microcircuit that can sit on the end of a hair. Silicon chips have shrunk about as small as they can go — to the atomic level. Any smaller and they become unstable.

Were the rapid advances of the silicon era a historical anomaly? Or — a weirder thought — could it be that the technologies driving it were transitional? Dig into these questions and you'll find something remarkable and hidden, in the form of a cadre of pan-global boffins working at the outer reaches of science who've been preparing for this day. Because, as with most things, the computing system we have is not the only one we might have had or could have. For some years, these visionary scientists have been theorising extraordinary systems that would compute with biological organisms, chemicals or light; would utilise microscopic nanotubes or the mind-blowing properties of quantum mechanics — the fabric of the universe itself — to create machines of such potential as to make the silicon years seem a stroll. Until recently most of the theories have sounded like the most outlandish science fiction. Now, in the push to find silicon's heirs, some of these machines are being built and tested.

I first stumbled across this idea in a meeting room at Intel's Silicon Valley headquarters about 12 years ago, while discussing the prospects for artificial intelligence with Justin Rattner, the chip-maker's affable head of microprocessor research. He mentioned it almost in passing, with 2020 being his estimate for the end of Moore's reign. "But what happens then?" I'd blurted, then peeled my chin off the floor as he shrugged: "Oh, people are experimenting with things like tanks of algae…" This drink his assistant had brought me — was it Kool-Aid? How on earth would algae compute?

"Well, if you think about it, a biological cell is a kind of switch." Meaning that it's either alive or dead: on or off. Like us. But also like the transistor switches printed in microscopic layers within a silicon chip. What a weird thought. Life as information processing.

Other avenues, he went on to say, included fungi, which might be grown and made to simulate neural networks — in effect, little mushroomy brains. And these would have one big advantage over silicon chips, namely that where the circuits printed on silicon work in two dimensions, these would operate in three dimensions, like our brains, allowing exponentially more connections. It's all a long way off, he said, but people were realising that nature already computes in ways we struggle to understand, using carbon, the building block of life. Or as ARM's Mike Muller puts it: "Organic-based lifeforms are clearly very power-efficient computing engines, and utilising that technology makes perfect sense. It's got to be the obvious replacement for — or more probably addition to — what we've got."

So you're saying that nature has given us the tools and techniques we need to replace or augment silicon? "Yeah. It's like, 'Catch up guys!' The question now is how to engineer it, program it, configure it." The only possible response to which is: *wow*.

Bizarre as it sounds, single-cell organisms and chemicals are already better than supercomputers at solving some problems. At the University of the West of England's International Centre for Unconventional Computing in Bristol, Ben de Lacy Costello,

THAT SHRINKING FEELING HOW THE SILICON ERA CAME AND WENT

1943
The world's first programmable electronic digital computer, Colossus, is built at Bletchley Park based on the theories of the codebreaker Alan Turing

1958
Integrated circuits on chips of geranium and silicon appear, courtesy of Jack Kilby at Texas Instruments and Robert Noyce at Fairchild Semiconductor. Noyce goes on to found Intel

1969
The Apollo Guidance Computer on board the command module and lunar module of Apollo 11 was more basic than many modern toasters, with around 64kb of memory

1974
The build-it-yourself MITS Altair 8800 kit computer, sold by mail order through magazine adverts, inspires Steve Jobs and Steve Wosniak to found Apple

1977
Jobs and Wosniak launch the home-computer revolution with Apple II, the first personal computer with colour graphics and a keyboard

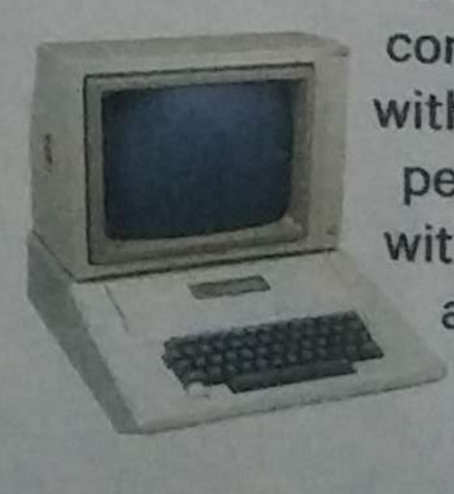

a chemist-turned computer scientist with a passing interest in philosophy, shows me why.

Our starting point is something called a Voronoi diagram, which scientists use to model complex systems in which "spheres of influence" are important — examples being galaxy clusters, the growth patterns of crystals or tumour cells, and competition for food among animal species. Alternatively, you might want to know how best to paint a pattern on your floor using two differently coloured tins of paint. These problems are extremely difficult to solve for our silicon-based computers, because they can only arrive at the best answer by trying all possible solutions one after the other, in a linear way. The principles behind this were theorised by the brilliant British wartime codebreaker Alan Turing, and have been in use ever since. But for really complex problems with myriad variables, this process could take centuries. And here's the point: nature, dealing as it does in complex, dynamic systems, tends to process information not in sequence, but in *parallel*.

To demonstrate what he means, De Lacy Costello drops two differently coloured chemicals into a petri dish and, over time, they form themselves into beautiful tessellations, finding the most efficient way to fill the surface area according to their innate programming: billions of tiny agents working independently and yet, somehow, in concert.

This is parallel processing, the holy grail of computing. The logical extension of Costello's work would be a shallow pool of chemicals whose elements we could manipulate in order to solve specific sorts of problems (like painting your floor efficiently) — a "reaction/diffusion computer", as first proposed by Costello's colleague Andy Adamatzky, professor of unconventional computing at UWE, back in 1992. If that sounds weird, ⋙

PREVIOUS PAGES: SLIME MOULD IMAGE BY HEATHER BARNETT. RIGHT: COURTESY OF DR BEN DE LACY COSTELLO AND DR ISHRAT JAHAN, UWE. BOTTOM: GETTY; APPLE

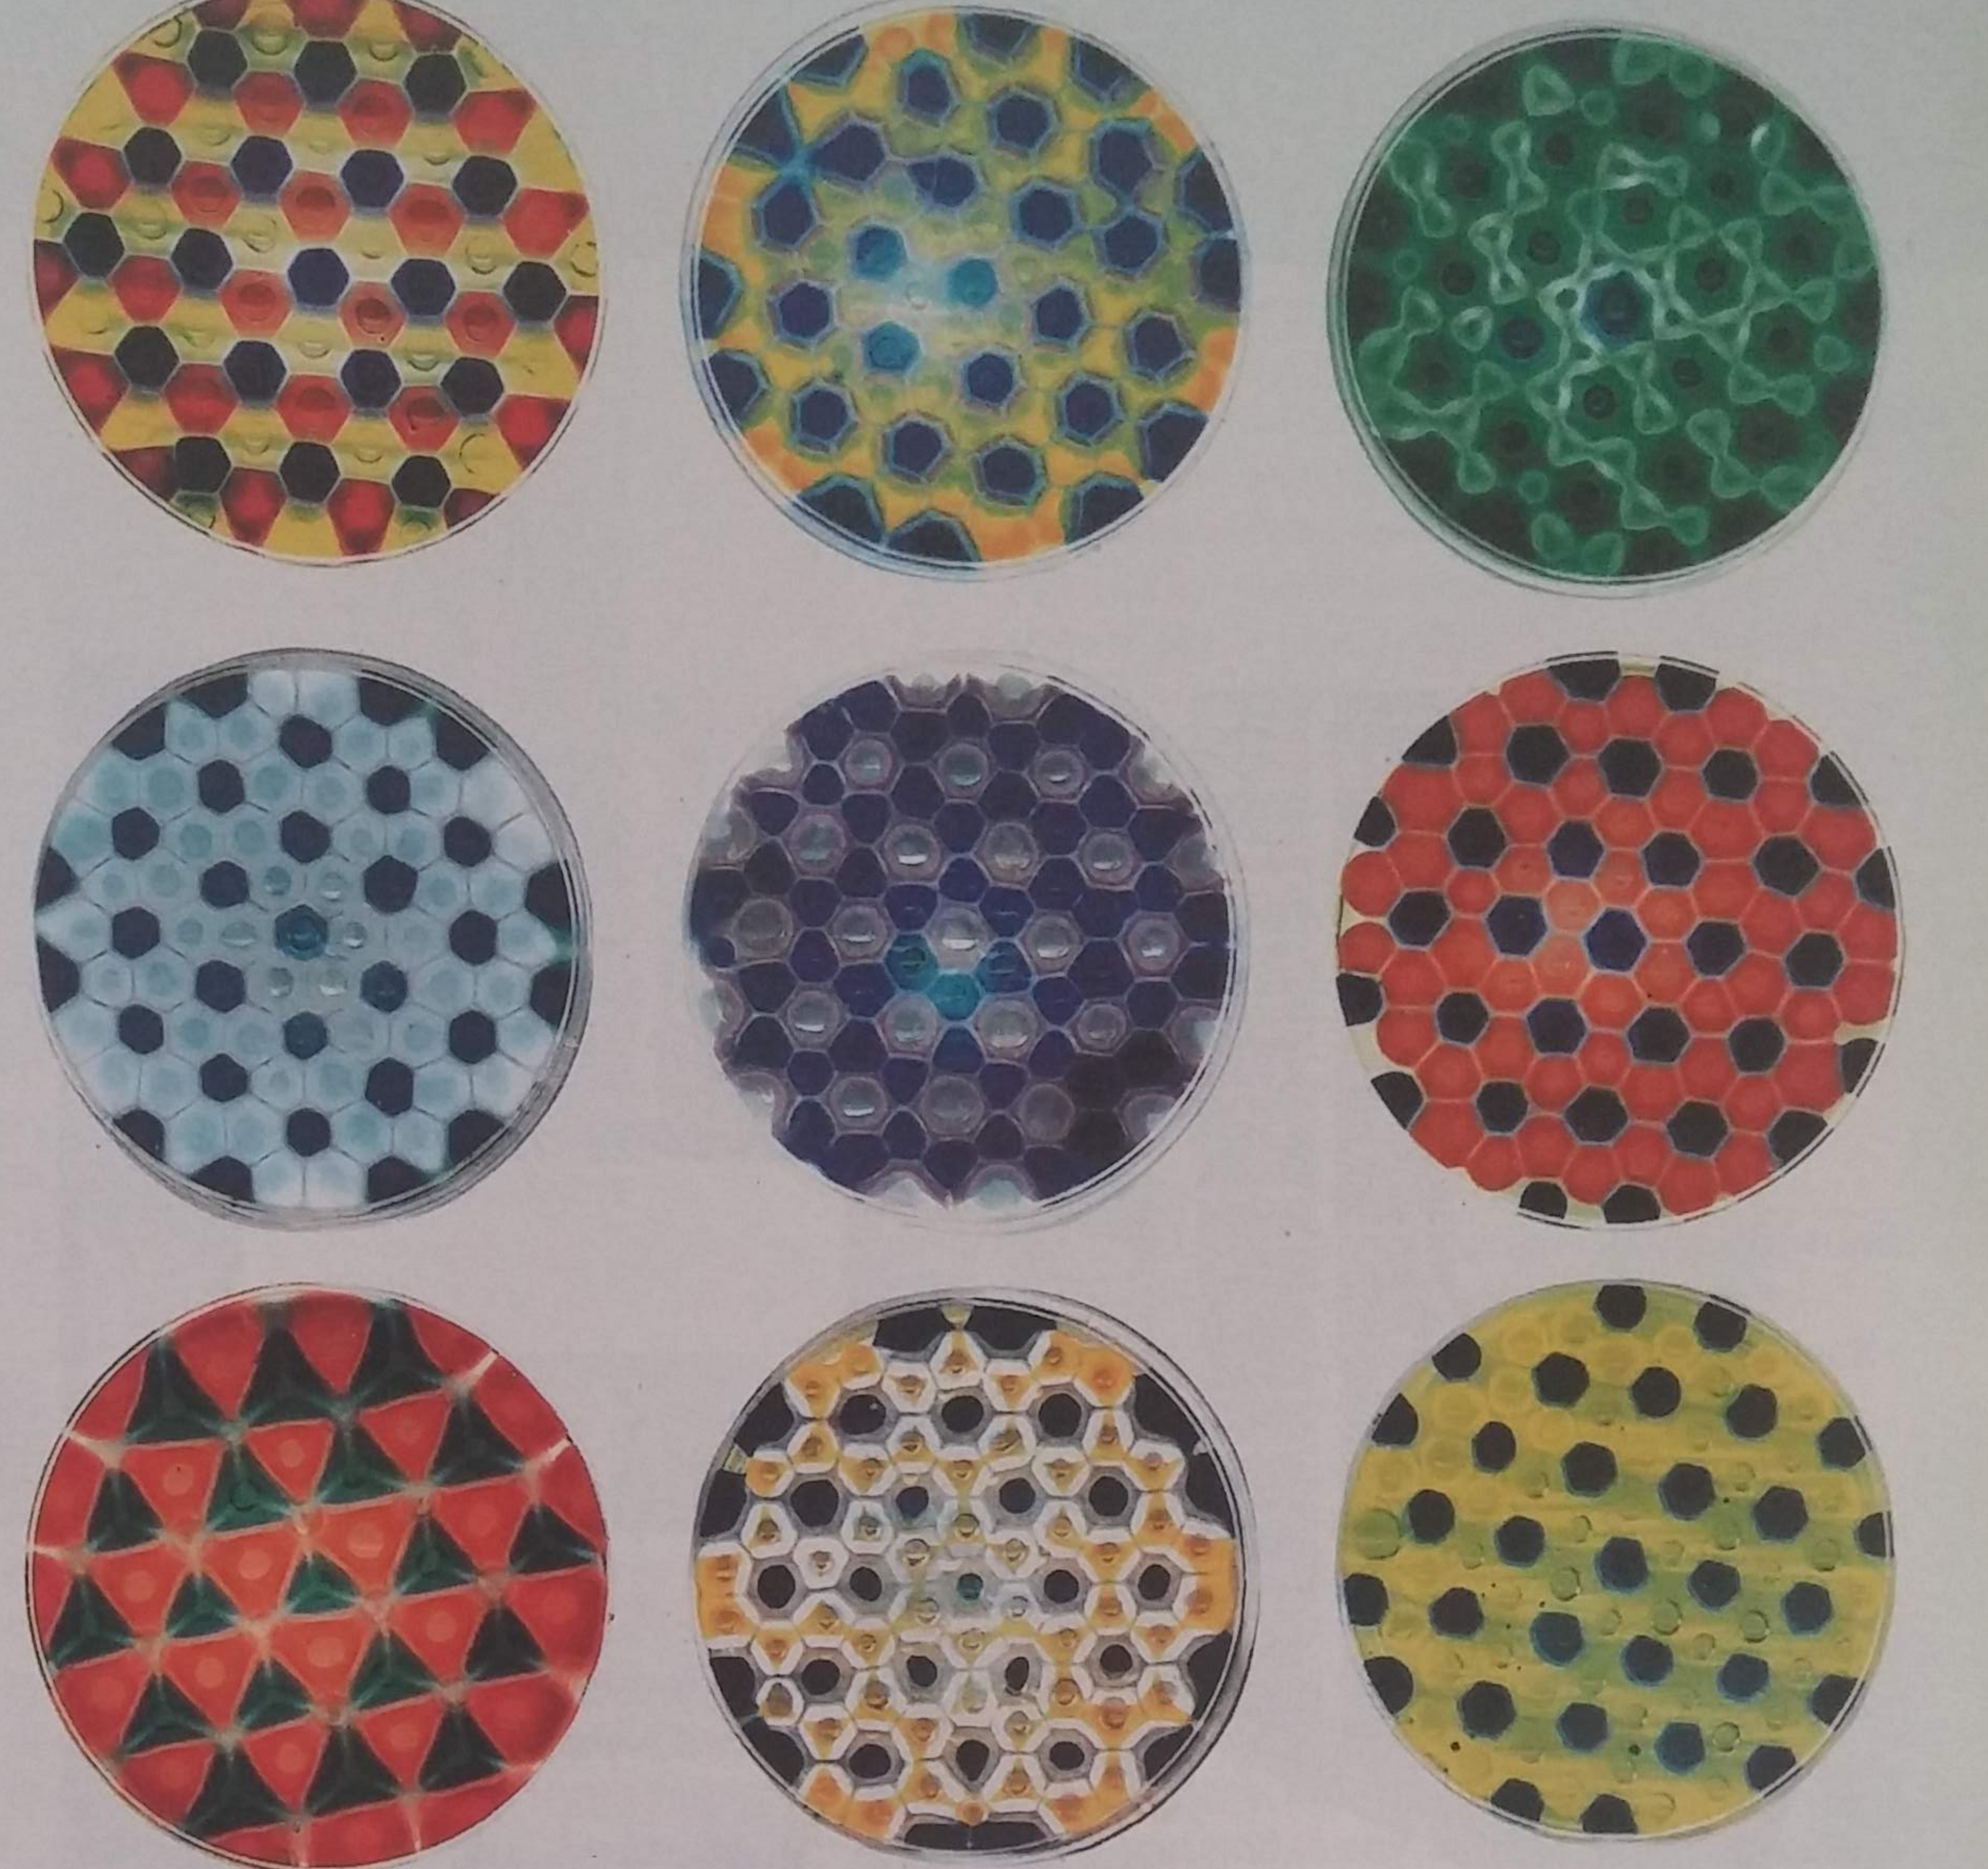

2015
Above: parallel processing in action — different coloured chemicals form patterns in a petri dish. Scientists are learning to harness these natural processes to compute with

1993
Intel launches its Pentium microprocessor, allowing computers to execute several instructions at the same time and support graphics and music

2007
Apple launches the iPhone, the first truly mobile popular computer

2010
D-Wave launch the world's first quantum computer — a step towards the holy grail of parallel processing

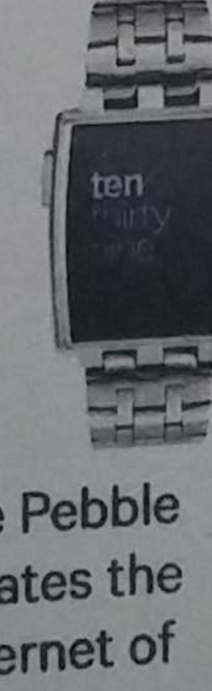

2013
The launch of the Pebble smartwatch indicates the arrival of the "internet of things", where computing becomes invisible and ubiquitous

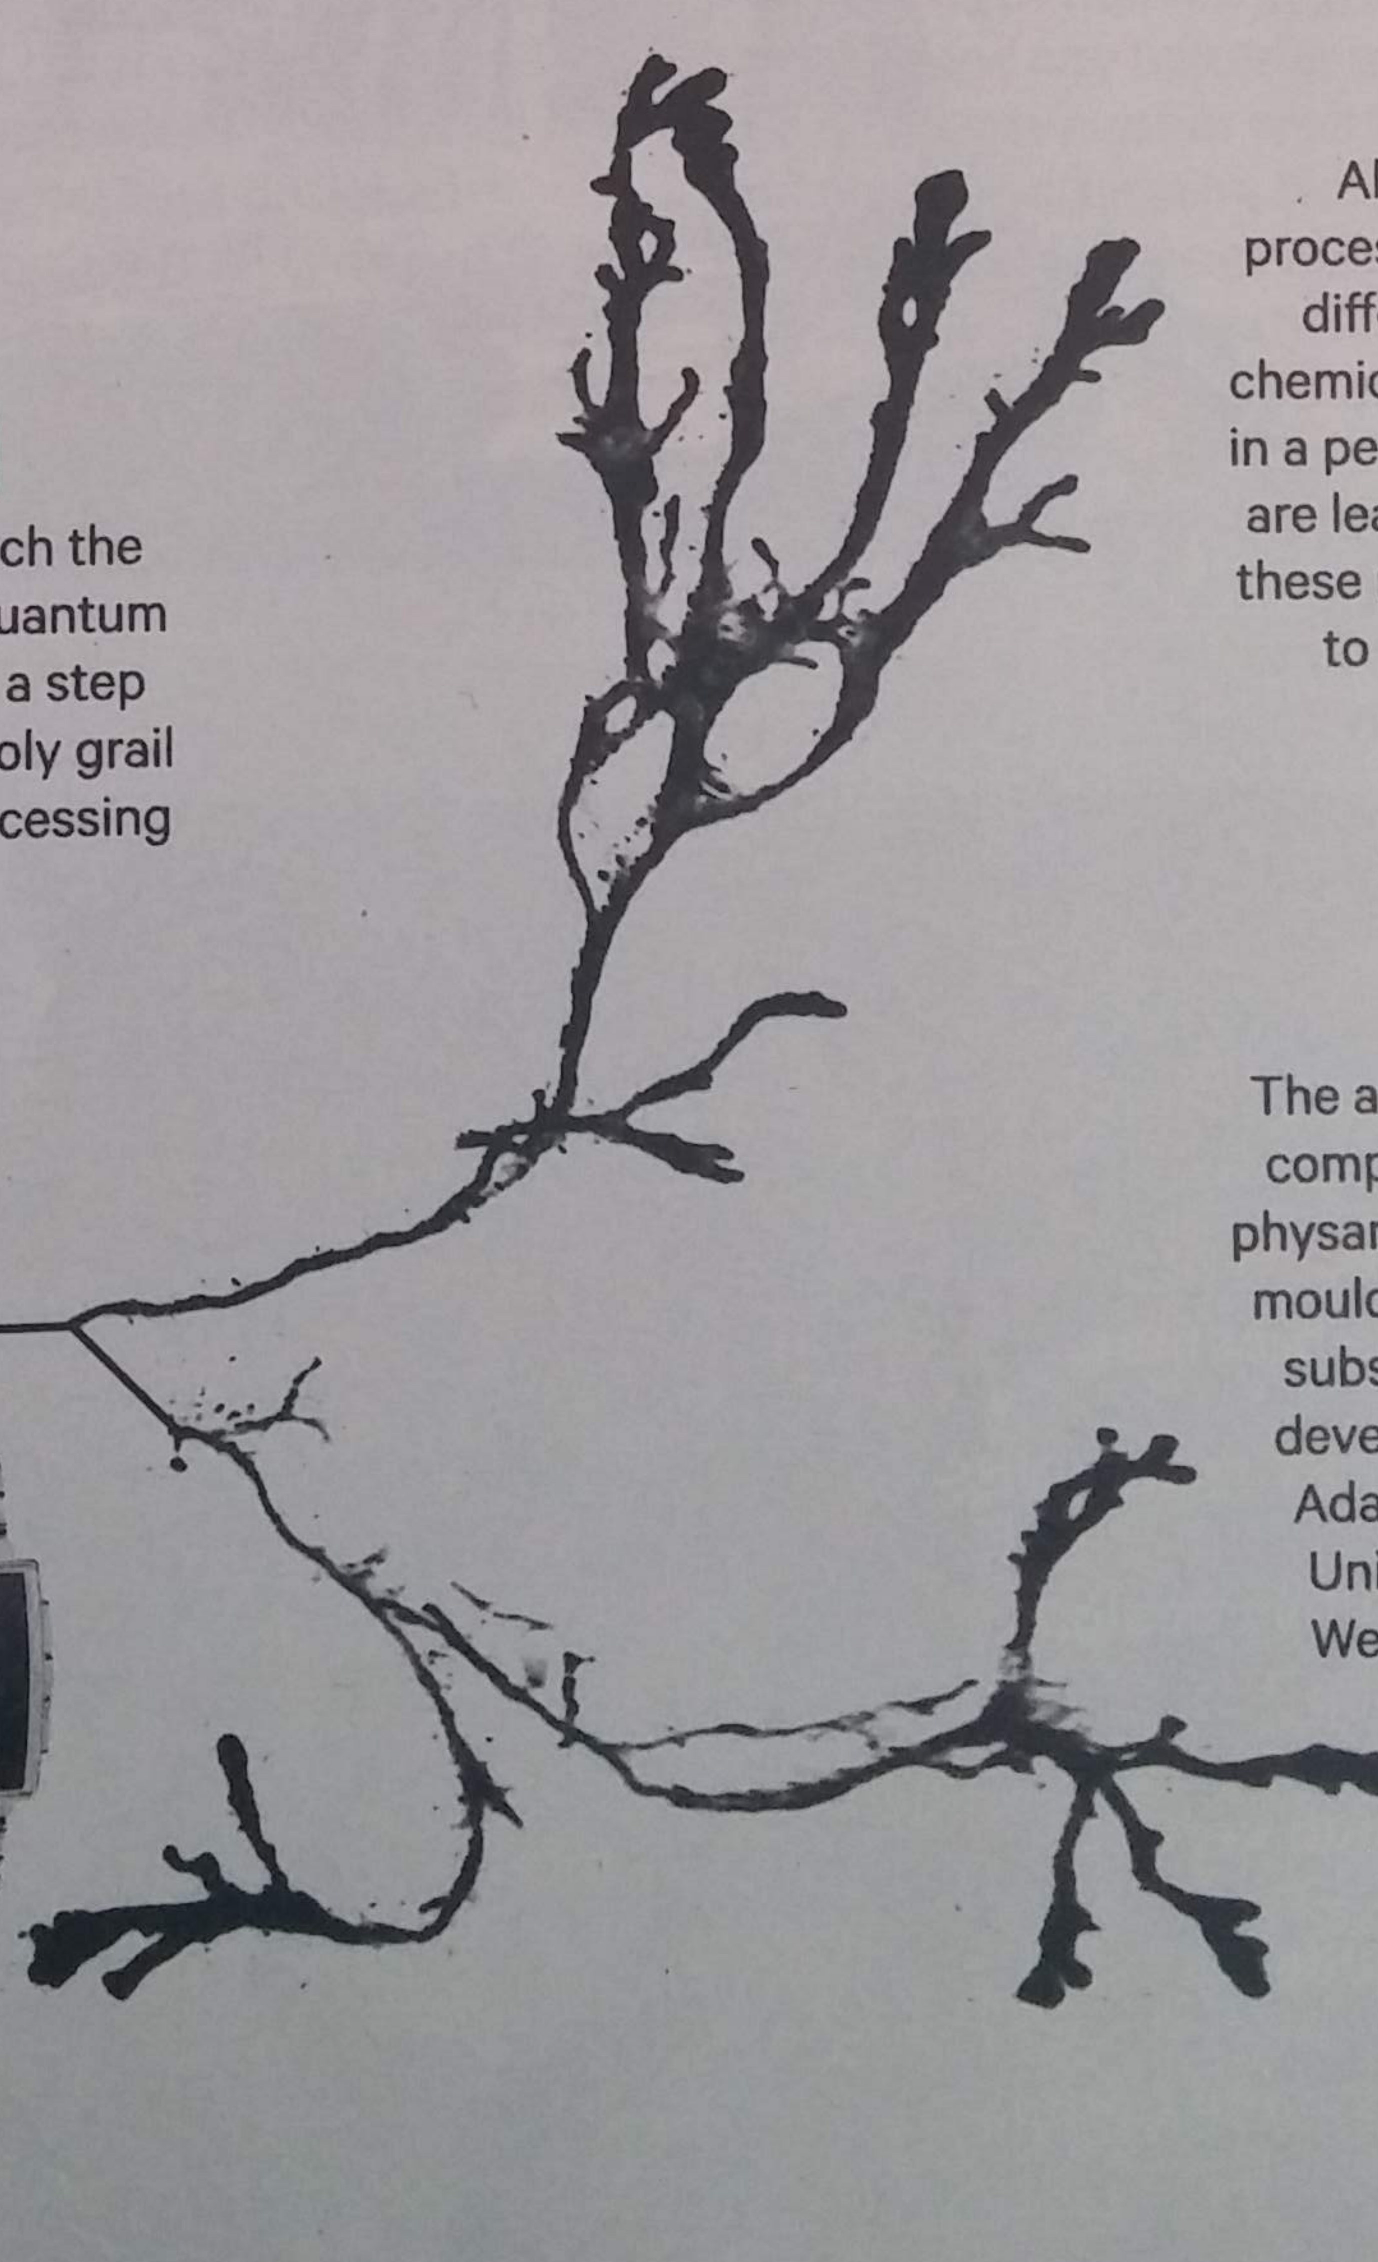

2014
The age of biological computing nears as physarum goo — slime mould — is used as a substrate in a chip developed by Andy Adamatzky of the University of the West of England

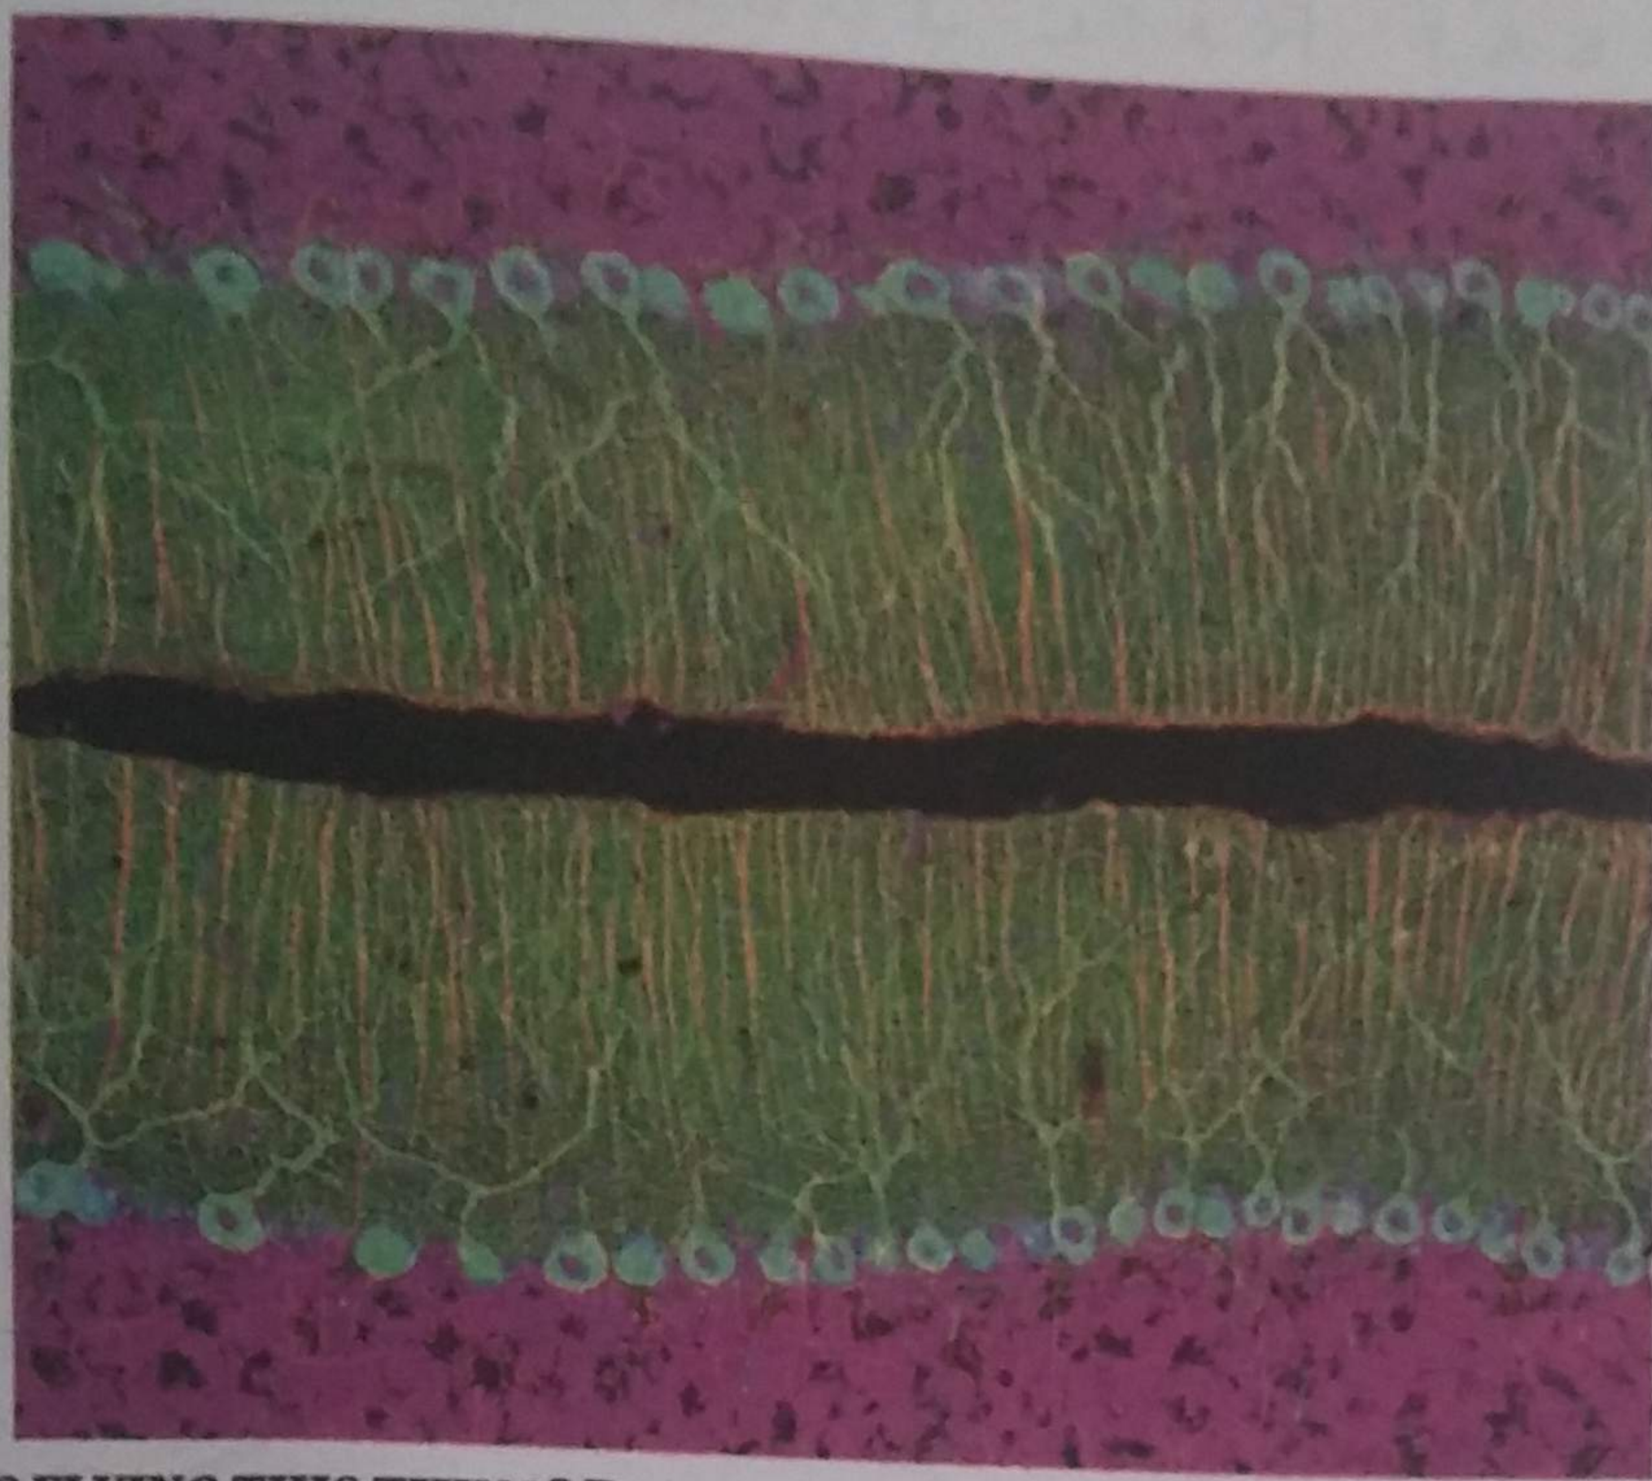

WHO'S FLYING THIS THING? Rat neurons (above) in a hybrid bio-electronic device have been trained to control a simple flight simulator

talk to Adamatzky about his current preferred computing "substrate" — *physarum goo*.

The goo looks like the slime mould you find on rotting trees in the forest (because it is) and is perhaps the most unlikely computing substrate since the Babel Fish, the ear-dwelling translator from The Hitchhiker's Guide to the Galaxy. A large, single-celled organism, *physarum polycephalum* can be controlled using light, vibration, salt and food (it's anyone's for an oat flake) and has proven to be as effective as conventional computers at solving certain problems — most notably in calculating the most efficient route through a series of points (as with a complex travel itinerary). And in terms of environmental impact, machines don't get any greener than this. Or sort of slithery yellow, to be precise.

Don't tell Nigel Farage, but this stuff is being taken seriously in Europe and research money is arriving. At the University of York, Dr Julian Miller channels the pan-European Nascence project, working on an entirely new kind of computational device that utilises blobs of tiny, randomly ordered carbon tubes with walls one atom thick and a range of unusual properties, called nanotubes.

Miller brings the issues surrounding Moore's Law into vivid focus. Turing-based computing systems try to solve problems by converting information into symbols — streams of ones and zeroes — then working through myriad possible solutions at the speed of light, which generates heat. They require ruinous amounts of energy to keep cool — Facebook has even built a data centre in the Arctic Circle to counter this. And as Miller speaks, I'm shocked to find myself articulating a thought I would previously have regarded as insane: namely that digital computation is kind of... dumb.

He smiles. "Well, from the point of view of humans trying to design computational circuits, it doesn't seem dumb at all, because we can make transistors in huge densities and make them perform very, very fast. But viewed from the perspective of evolution, it does seem very dumb. Your brain doesn't compute symbolically in this way. I think we can now use Turing machines to create technologies that will make Turing machines redundant."

Miller calls his current approach to novel computing "evolution-in-materio". At base, it is a recognition that Darwinian evolution has furnished us with tools, ideas, systems far richer than anything we have designed ourselves, and that rather than trying to fight and control them, we should harness them, effectively hitching a ride with nature.

An important feature of the chips he is developing is that, in terms of physics, we don't know precisely what goes on inside the nanotube blob — only that we can use its natural properties to generate algorithms with which to compute. But according to Miller — and this strikes me as a properly radical scientific idea — *it doesn't matter*.

"When we make things, we tend to make them to be stable and predictable and controllable. But nature doesn't do that. A bumblebee flies by creating vortices: it can only fly by creating non-linear effects — there's no smooth airflow. Evolution found out that there were physical effects there that could be exploited, so these insects utilise them. But if human beings tried to do that, we would say, 'Well, we need to know exactly how to model all those non-linear effects, and if we can't do that, forget it!'"

The approaches discussed so far represent the tip of a large and growing iceberg (in fact, look below the iceberg's waterline and you are likely to find a computer scientist attaching electrodes). When asked, UWE's Andy Adamatzky, one of the key figures in this burgeoning field, sends me an arm-long list with names like cryptic crossword clues: "bio-molecular computing, cellular automata VLSIs, in vitro computation, DNA computing, bacteria-based controllers for autonomous robots, reversible computing..." and so on. "The field is volatile, its boundaries are ill-defined and changing almost daily," he tells me, to which another practitioner adds, "It's like the Wild West right now."

At the extreme, Dr Mike Stannett of Sheffield University points me in the direction of Hungarian colleagues who conceive a "hyper-Turing" computer based on the black hole at the centre of our galaxy (operated by Obi-Wan Kenobi, one must hope).

Especially exciting to many is the burgeoning sphere of synthetic biology, and when you speak to Martyn Amos, professor of novel computing at Manchester Metropolitan University and author of a thrilling book called Genesis Machines, it's easy to see why.

"If you said to anybody in nanotechnology, 'Look, I've got this self-replicating robot and it doesn't need any batteries or external power, and it can sense other entities of its own class and engage with them, form consortia, combine and exchange information, and even self-destruct if necessary...' they would sell their grandmother for a look at one. But of course we have them! They're bacteria and they're yeast and they're mammalian cells. And we're beginning to understand how they're put together, and we're beginning to develop a circuit diagram which will allow us to replace components and persuade them to do what we want them to do. That's what synthetic biology is about."

These living cells can even be merged with traditional silicon machines to form truly hybrid bio-electronic devices. Amos and his team are working on one in which cultured rat neurons have been trained to control a simple flight simulator, so fasten your seatbelts tightly. As a bonus, the quest to adapt and make use of microbes, cells, DNA in this way is leading us to a better understanding of their often ingenious workings, says Amos.

"Did you know that bacteria have a form of sex? They exchange DNA molecules through physical contact, by squirting little plasmids into one another, and this is a major mechanism behind microbial resistance to drugs: it effectively allows bugs to 'broadcast' genes that protect them from antibiotics — like teenagers sending a tweet — and these genes can then be taken up and retweeted ⋙

"Did you know bacteria have sex? They squirt plasmids into one another. How about we use bacterial conjugation to send packets of data?"

GETTY

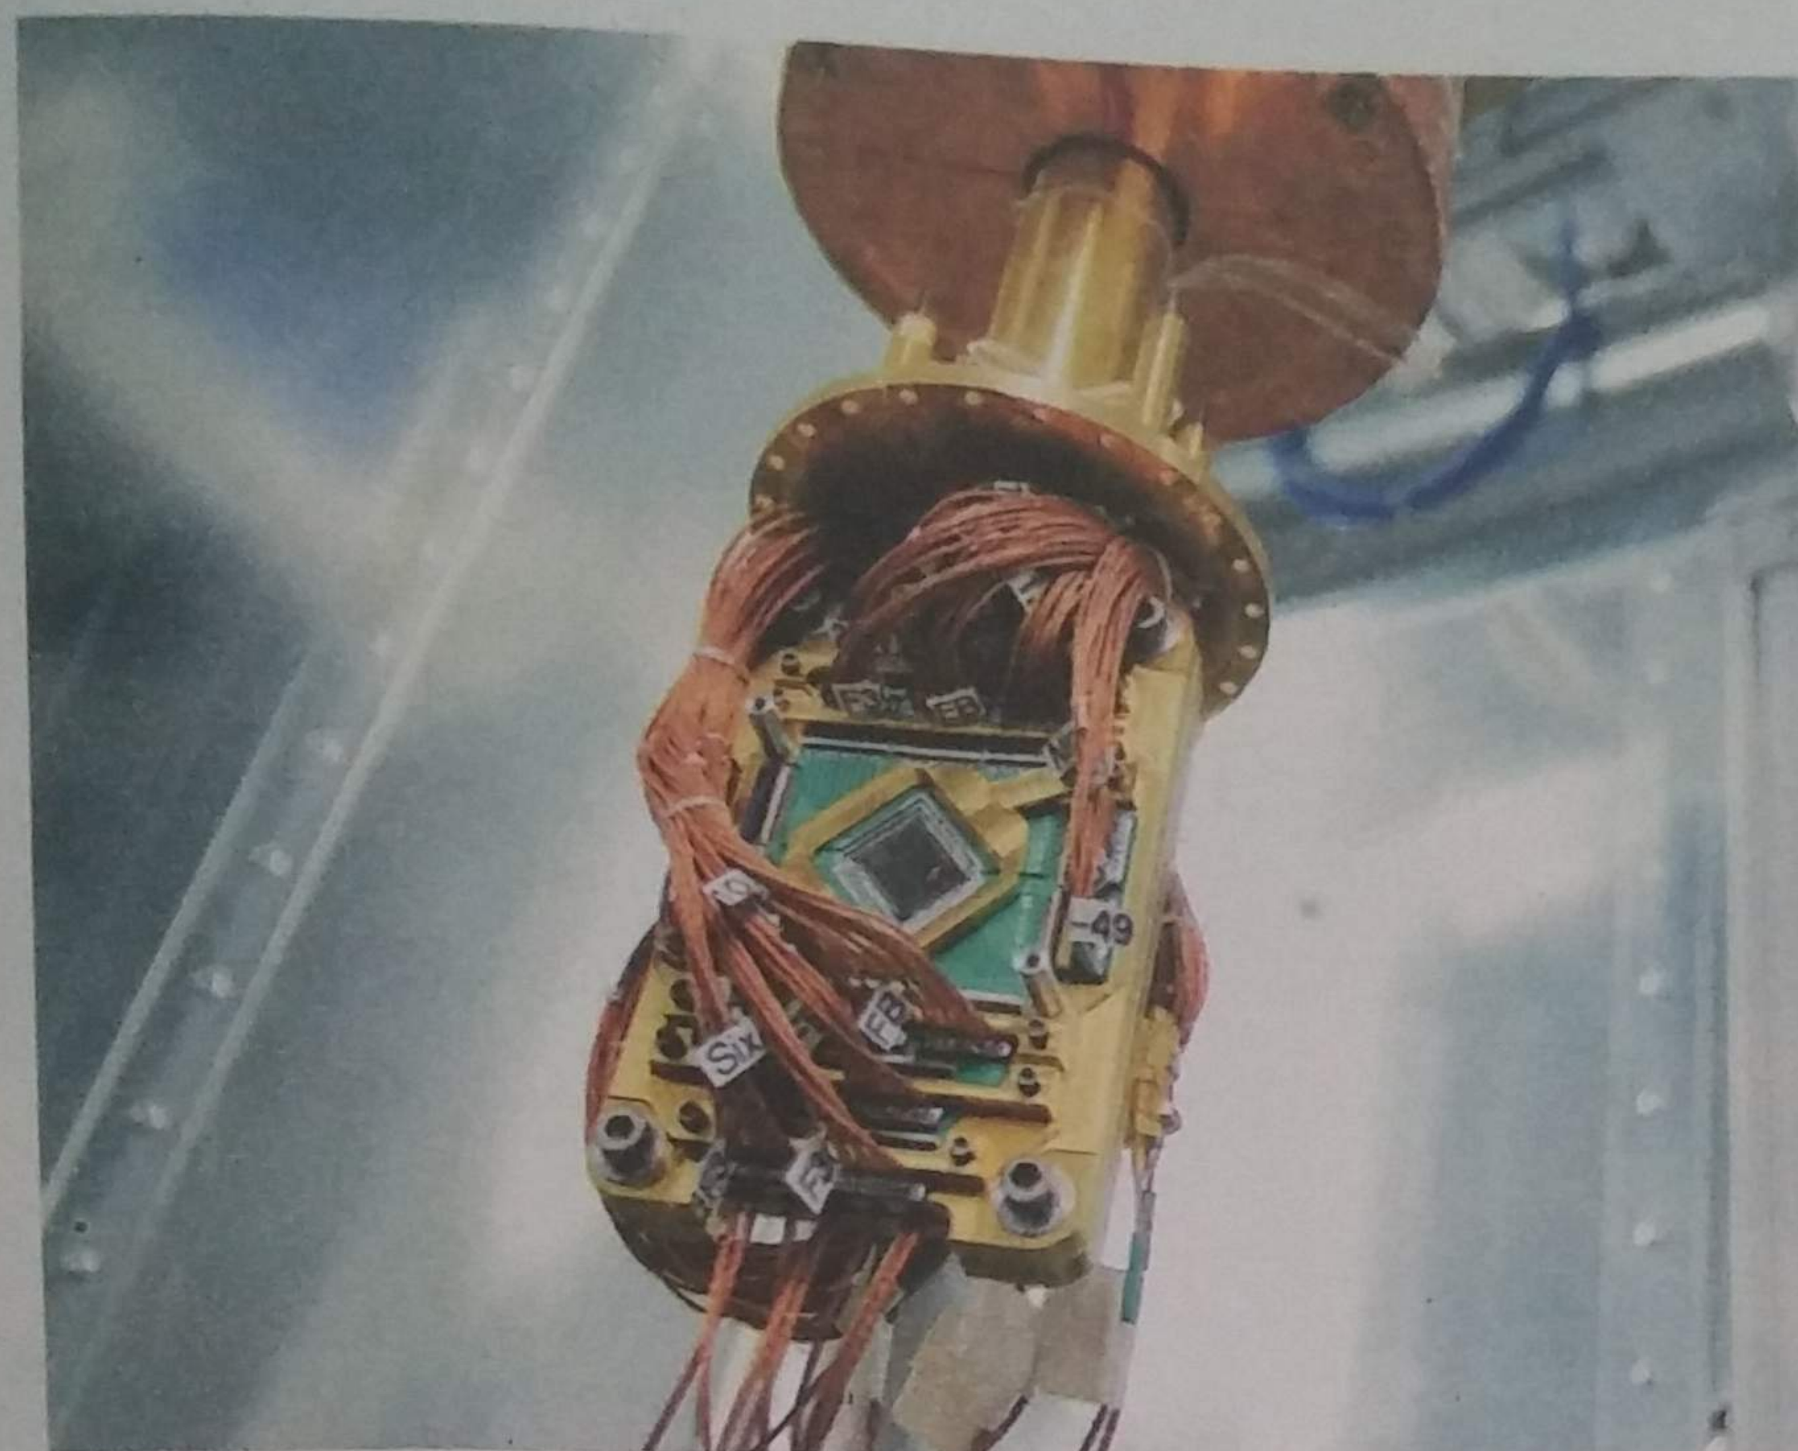

QUANTUM LEAP The D-Wave Two allows electrons to whizz around a loop in two directions at once — effectively summoning two parallel universes

by other microbes. Now our colleagues are saying, 'How about we use these mechanisms to engineer communication between cells? Using bacterial conjugation to send packets of data, the way our computers do.'" So instead of ones and zeroes, you would have As, Gs, Cs and Ts — the symbolic components of DNA. The difficulty with biological systems, Amos notes, is that small changes in environment can affect them profoundly, because biology is complex.

Amos recently co-edited a book of short stories called Beta Life, in which scientists were paired with writers including Frank Cottrell-Boyce, Lucy Caldwell and Toby Litt. One of the narratives features self-growing skyscrapers, a notion that seems wildly sci-fi until you talk to Rachel Armstrong, professor of experimental architecture at Newcastle University. She hasn't grown any buildings yet, but draws on unconventional computing to try to explore a less environmentally destructive "living architecture".

Whereas traditional architecture takes natural materials and forces them into alien contexts (eg, turning clay into bricks), wasting much energy in the process, living architecture aims to adapt or create new materials that intrinsically do what we want them to — for instance, self-repair or grow into hard structures, or extract carbon dioxide from the atmosphere. Rather than try to save Venice by replacing its wooden stilts, asks Armstrong, why not grow a limestone reef under it?

She believes the advent of the internet has brought about a "new nature: the millennial nature, which has a technological character". "It's very different to the 20th-century idea of nature, which was anti-human, or seen as something we're apart from and need to either conserve or drain as natural capital."

We can't go on for ever with a model of economic development that throws billions of tons of carbon into the atmosphere as cities, countries, industries expand, nor the gargantuan energy demands of the data farms that process our online life. Unconventional computing and the self-organising materials it points to — programmed to "grow" structures the way plants or crystals do — suggest a less destructive path forward.

For the D-Wave processor to work, the temperature must drop to 10 millikelvin — the lowest temperature in the universe

Perhaps the freakiest-sounding of all speculative, non-classical computer technologies is quantum computing — a technique first proposed by the physicist Richard Feynman, one of the giants of 20th-century science. At its core are the eerie laws of quantum mechanics, which Einstein, despite having birthed them, famously refused to believe. Of particular interest is the principle of "superposition", which holds that a physical system at the quantum level — an electron, say — can exist in more than one state at the same time; can be simultaneously in Russell Brand's beard and at the far end of the cosmos. In addition, "entanglement" describes the way two or more particles can be — how? Somehow! — connected over great distance, so that a change in the state of one affects the other, for reasons we're not even close to comprehending.

Weird, right? Now get this. A quantum computer uses the principle of superposition to enable the "bits" — the ones and noughts we know from classical computing — to adopt the state of one, nought, *or both at the same time*. What's more, these "qubits" have the capacity to process information in both states simultaneously, creating vast branching trees of parallel processing. (And if you can't visualise this, don't worry: no less an authority than Feynman said: "If you think you understand quantum mechanics, you don't understand quantum mechanics.")

At the turn of the millennium, I would have bet good money on never seeing a quantum machine in my lifetime, but in 2010 a Canadian company announced they'd built the first commercially available one — D-Wave One — and followed it up in 2012 with the 300,000-times more powerful D-Wave Two.

Down the line from Vancouver, Jeremy Hilton, D-Wave's head of processor development, describes the contortions necessary to bring it off. The qubits on a D-Wave processor are made of a metal called niobium, he tells me, formed into a loop. To use these in the way of standard processors, they are cooled to -264C, but to encourage quantum effects they must drop to 10 millikelvin, which is 150 times colder than interstellar space and 30,000 times colder than room temperature; thought to be the lowest sustained temperature in the universe. Once this is done, a magnetic vacuum must be created, until the current in the loop defies the laws of standard physics by travelling clockwise and counterclockwise at once. Now the electrons are in two states, the quantum chip having, by one interpretation, summoned two parallel universes. "So it gets a little crazy," says Hilton, with frankly heroic restraint.

The D-Wave Two has a footprint of 8ft by 12ft, stands 10ft tall and is reminiscent of Turing machines from the 1960s, being at a similar stage of development. In all likelihood, this size will reduce rapidly, but Google, Nasa and Lockheed have already bought in, at a price D-Wave won't discuss, but which has been estimated at $15m — a snip, if it works. At present, D-Wave Two is claimed to be on a par with the most powerful classical supercomputers, in conventional terms, but speed is not the issue right now: more important is the type of calculations quantum computers can address. Reportedly, Google are using theirs in pursuit of machine learning, a first step towards real artificial intelligence. That said, lest we forget that even the laws of quantum mechanics pale before the Law of Sod, one of the near-impossible problems quantum computers will solve in seconds is how to identify our internet security keys — meaning it really will be out with the old and in with the new ■

Andrew Smith, @wiresmith on Twitter, is the author of Totally Wired: the Wild Rise and Crazy Fall of the First Dotcom Dream (Simon & Schuster)

Is anybody out there?
How Nasa will use D-Wave to look for alien life
On tablet, or at *thesundaytimes.co.uk/unconventionalcomputing*