IAN FYFE

Francisco Monteiro looks at the remarkable achievements in error-free digital communications

WE ARE the first generation that is able to contact friends on the other side of the world, from anywhere, at any time. Whether in the living room or in the middle of a park, we can use a tiny laptop, apparently connected to nothing except the air we breathe, to chat with friends on a webcam whilst a missed TV show streams in another window. We take this for granted, yet the development of error-free, wireless transmission is one of the most astonishing intellectual achievements of modern science.

Most of us know that any piece of music, painting or text can be represented by a combination of just two symbols, known as binary digits or bits (for simplicity, we call them *zeros* and *ones*). And we know that we want lots of them coming to us in a short time. But marketing tells us to ask for higher speeds and this is misleading. Data can be received more quickly if more bits are transmitted per second, but the bits themselves do not travel any faster. So what marketing should tell us is to ask for a higher *bit rate*, not a higher speed.

These digital signals, in the form of *zeros* and *ones*, must be detected and decoded against corrupting background noise. For example, temperature causes random movement of electrons in receivers, which disrupts the signal. Error-free transmission of binary digits under such conditions is not easy. Some *ones* may be mistaken for *zeros* and vice versa, and errors increase with faster bit rates. The challenge is to maximise the bit rate whilst minimising errors.

Is there a limit to the bit rate we can achieve, whilst keeping the link free from errors? This was one of the questions that Claude Shannon asked in his 1948 seminal paper *A Mathematical Theory of Communication*. Shannon formulated the concept of a channel's information capacity; the maximum achievable rate of error-free data transfer in a given channel (the Shannon limit). He showed that if we transmit below the capacity of a channel, some code should exist that would allow the correction of all the bits that have been corrupted. It is similar to a word processor suggesting corrections for misspelt words; more specifically the proficiency with which it identifies the most likely correct word.

> “Marketing tells us to ask for higher speeds, but this is misleading”

Some of the brightest mathematicians, engineers, and computer scientists devoted themselves to the problem of finding such a feasible error correction code. However, by 1993, even the best codes were still performing far from the capacity limit.

Then the unexpected happened. In a leading conference, a paper, claiming to have a feasible family of codes (dubbed turbo-codes) that operated near the Shannon limit, was presented by Claude Berrou and Alain Glavieux, two French engineering professors who where rather unknown at the time to the coding theory community.

“They got it wrong,” people mumbled

IAN FYFE

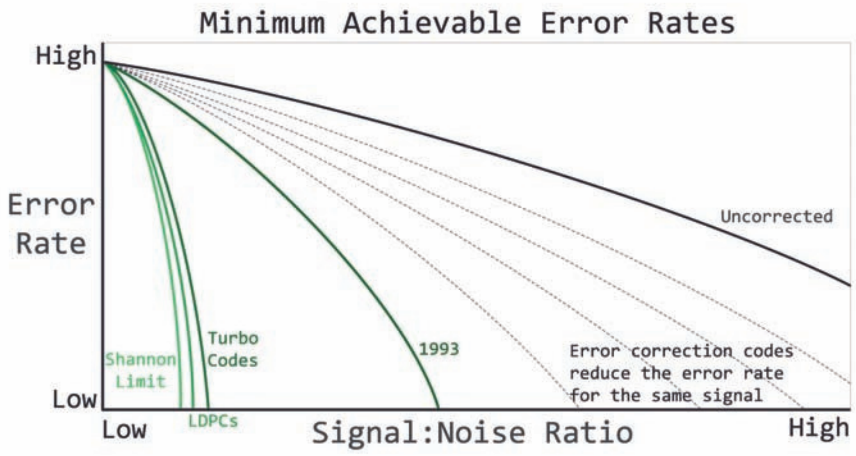

Progress in error correction codes allowed error rates to approach the Shannon limit

“By 1993 even the best codes were still performing far from the capacity limit”

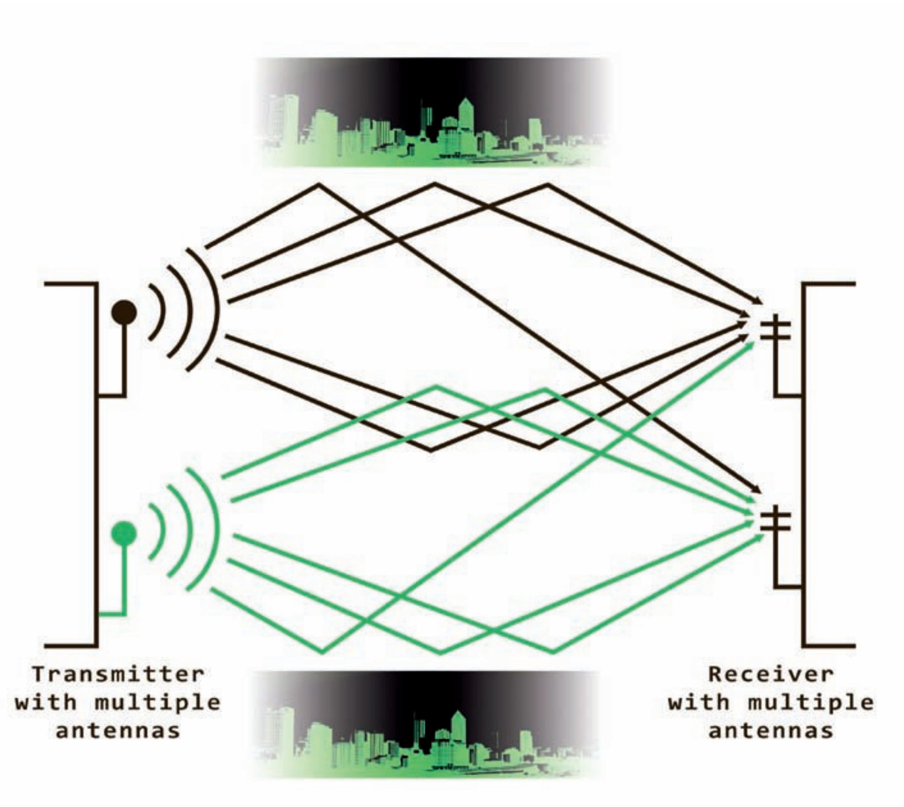

IAN FYFE

MIMO space-time processing takes advantage of multiple reflections that act to artificially create independent communication streams

at the end of the presentation, “They must have forgotten to divide by two somewhere!” Everybody rushed back to their labs and tried to replicate the results. They could not believe what they found: turbo-codes were performing just as claimed. However, it was unclear why they worked.

At around the same time, Cambridge Professor David MacKay, along with Radford Neal at the University of Toronto, was looking at the problem from a fresh perspective. In 1995, he devised codes operating even closer to the Shannon limit. For some time, his Low Density Parity Check Codes (LDPCs) made Cambridge the home of some computers that were running the best error-correcting codes in the world.

Interestingly, his research revealed that LDPCs had been devised by MIT professor Robert Gallager in his 1962 PhD thesis, but had been forgotten. This was probably because there was not enough computing power at the time to implement them, or because he did not include them in his textbook published in 1968.

Mackay’s papers triggered a boom of research and LDPCs were further refined by researchers in America and Switzerland. Currently, turbo-codes play a central role in the correct detection of the bits received by mobile broadband, and help to receive images from the probes on Mars. The patent-free LDPCs will take their place soon.

It had taken almost 50 years to reach the Shannon limit. But a further burst of research in the second half of the 1990s proved that the maximum possible bit rate within a fixed spectrum had not been reached. Shannon’s formula for typical electrical channels considered thermal noise only, not additional “perturbations” such as multiple reflections of the signal in the environment, as is the case in wireless communication. For many years, this type of “self-interference” was perceived as an additional obstacle for correct signal detection at the receiver.

However, it was later proven, mathematically and experimentally, that by considering space in addition to time when designing a code, the Shannon limit could be surpassed. With rather complex algebra and computing, we can artificially create several independent communication streams using the so called space-time coding on multiple-input multiple-output (MIMO) systems. In electronics, this translates to the use of multiple antennas on the outside and much more processing complexity on the inside.

“4G mobiles will reach rates of up to 1 gigabits per second”

The same MIMO principles are now being used to take advantage of the different reflecting paths that light waves can take inside optical fibres. Even in bundles of landline cables, the mutual interference can be used in a similar way.

Soon, 3.5G mobiles will provide gross bit rates of up to 100 megabits per second (Mbps) and, inside the house, the next Wi-Fi standard will provide up to 600 Mbps. Later, 4G mobiles will reach rates of up to 1 gigabits per second. To put that in perspective, in 2008 the average download speed in the UK was 4 Mbps. At 2 Mbps it takes 47 minutes to download a typical film; at 10 Mbps this is already down to twelve minutes, so at 600 Mbps it will literally take seconds.

These are the plans for the next decade, but a revolution has recently started in academic circles: network coding theory and collaborative networks have all users in a network helping all other users to sustain the error-free bit flow. At this stage, the capacities for such networks are unknown, and a new Shannon is needed.

Francisco Monteiro is a PhD student in The Computer Laboratory and in the Engineering Department