Reversible Computing

Michael P. Frank, as a graduate at MIT, designed one of the first universal DNA computers, and discovered that it had to be reversible for fundamental thermochemical reasons. Since his graduation from MIT in 1999, Mike has been continuing his reversible computing research as an Assistant professor at the University of Florida's Department of Computer and Information Science and Engineering. *Digit* caught up with him to speak about reversible computing.

"Our civilisation will depend on exploring the problems that reversible computing poses"

What is reversible computing?

Reversible computing—revcomp for short—is the application of principles of recycling to computing. When a computer performs a logical operation, the unwanted bits after the operation are thrown away, and dissipated as heat. There is a limit on how far we can keep doing this. If we want computer speeds to keep on increasing, we need to 'uncompute' the unwanted bits.

Reversible computing means computing using a physical mechanism that is thermodynamically reversible—one that ballistically coasts along through its computation with low friction. Thus, it will dissipate only a small fraction of a bit's energy to heat with each logic operation.

If a computation is thermodynamically reversible, it would need to be logically reversible as well. Performing any desired computation is still possible, despite this constraint. But it strongly affects digital logic designs. In the long term, it will impact instruction sets and high-level programming languages as well. Eventually, these will also have to be reversible to provide optimal efficiency.

Why do we need reversible computing?

With device sizes fast approaching atomic-scale limits, more energy efficiency is essential. Energy efficiency fundamentally affects the speed of circuits—such as nanocircuits—and therefore the speed of most computing applications. Ballistic circuits that conserve information, by uncomputing bits instead of throwing them away, will soon offer the only physically possible way to keep improving energy efficiency.

What is a ballistic circuit?

A ballistic process is a process in which a system proceeds forward under its own momentum, with only a small fraction of the free energy in the system being lost to heat. For example, when an ICBM (Inter-Continental Ballistic Missile) makes a flight between continents, it 'coasts' (movement of a vehicle with hardly any generation of heat) for most of its journey, with only a small portion of its potential and kinetic energy being lost to friction during its trajectory through the atmosphere.

How soon do you think revcomp is going to become important?

It will become important when the chip-making industry runs out of other tricks for reducing the energy cost of logic operations. Throughout the history of the industry, it has been reducing the energy cost of logic operations by reducing bit energies; but, there are some absolutely irrevocable limits to this process. These limits will be reached within the next three decades. This has to do with the fact that the density of all of the physical information that is present in materials at normal temperatures and pressures is finite, and moreover, it is quite limited—at most, a small handful of bits per atom. We can never store a logical bit using less than 1 physical bit, and the irreversible transformation of a physical bit results in heat energy that must be expelled into the environment. In room-temperature environments, this results in a strict maximum of about 300 billion billion irreversible bit-operations that can be performed per Joule of energy consumed. This may sound like a lot, but we are only about a factor of 1,00,000 below this point. If past trends continue, we would expect to reach this point in only another 25 years, or so.

Why do we need to bother about it now?

To make reversible computing practical, there are many difficult engineering research problems that remain to be solved, such as the design of extremely high-quality oscillators. If we don't start working hard on developing solutions for these now, there is a significant danger that solutions will not be ready when traditional approaches run out of steam. If this happens, the development of computer technology will stall—no longer will we be able to expect that every year we will be able to get a more powerful processor with lower power consumption, which can be utilised for new applications not previously possible.

This can have a significant braking effect on the growth of the entire world economy, since productivity growth in many industries today relies largely on the increasingly extensive application of information technology. The rate at which IT can penetrate into new application areas will drastically slow down, if the performance-per-unit-power-consumption of digital technology stops improving.

> "...we risk hitting a technological plateau within the next 30 years."

If we do see a reversible computing device in as soon as 25 years, what will it look like? If not a desktop or a laptop, what will it be?
I believe that desktop and laptop—or maybe wearable, or even implanted—personal computers, will be a major potential area for application of revcomp technology. For example, powerful 'personal assistant' and 'intelligent agent' type applications might be built, using sophisticated artificial intelligence techniques. These algorithms require extremely high-performance computing capabilities. The more performance you have, the 'smarter' these techniques can become—with essentially no limit. Reversible computing is ideally suited for running these artificial intelligence algorithms, because they can all leverage a very high degree of massive parallelism.

If this is available, these algorithms do not even require extraordinarily high sequential performance. For comparison, individual neurons in the human brain, fire at frequencies of only about 1 KHz. Because there are on the order of 1 billion billion synapses packed into your skull, this translates to a total rate of about 1 quintillion (10^{18}) synpatic operations per second. This is about a thousand times greater than today's desktop computers. So, we can expect that achieving capabilities comparable to the human brain will require at least a thousand times greater performance than any desktop computer has today, as well as sophisticated AI algorithms. However, the example of the brain teaches us that intelligence does not necessarily require the individual devices to operate at extremely high frequencies.

The goal of AI, thus, fits perfectly what reversible computing offers, which is a much larger overall rate of operations within a given power budget. This will be achieved through massive parallelism of large numbers of extremely energy-efficient micro- or nano-scale devices. These can individually operate at only a moderate speed, which is required in order to maintain their high energy efficiency.

How realistic is the 25-year timeframe you spoke about for reversible computing to become mainstream?
I do not know for certain whether the semiconductor industry will continue to find it economical to continue taking conventional computer technology forward for as long as 25 years at the same rate that the technology has been progressing for the last 40 years or so. It is very difficult to say, because many problems to further progress loom close ahead, yet many ingenious ideas have been proposed that may push the limits a bit farther out. However, the point is that at whatever time the conventional approaches run out of steam and stop improving so quickly, reversible computing will then be the only way to rapidly introduce further reductions in power dissipated per unit of performance. If the traditional approaches hit a practical limit and stall in considerably less than 25 years, then so much the better for reversible computing!

You spoke about micro- or nano-scale devices, and massive parallelism. Where is the intersection between nanotechnology and reversible computing?
The fundamental limits on the energy efficiency of conventional semiconductor-based computing also apply identically to all possible nanotechnologies. As a result, nanotechnology, if not augmented by reversible computing, can only inch us a bit closer to the limits of irreversible computing than semiconductors might be able to get, but it cannot take us beyond these limits. The only way that nanotechnology can possibly circumvent energy dissipation limits, and keep computer performance improving beyond the next two or three decades, is if it wholeheartedly and thoroughly adopts reversible computing principles.

In the 1980s and 1990s, K. Eric Drexler and Ralph Merkle, the pioneers of nanotechnology, discussed the need for reversible computing in nanotechnology, and designed a wide variety of nanoscale mechanical and electronic reversible logic mechanisms.

If reversible computing is so fundamental and important, how come we've never heard about it until now?
The lack of attention to reversible computing has been due to a stew of misconceptions about revcomp that have been floating around. For instance, myths circulate that reversible computing has been proven impossible, or that it violates some law of thermodynamics, or that it is fundamentally impractical. None of these myths are true—no valid proof of any of these claims exists anywhere in the literature. I know, because I have spent years searching for one—since I have no desire to waste my career pursuing an impossible goal! There have been a number of attempted proofs of the impossibility of reversible computing, but all of them contain either fatally-flawed assumptions, or logical fallacles.

It is not an exaggeration to say that the entire long-term future of computing technology, and indeed of our entire civilisation, will depend on exploring the problems that revcomp poses. We need to start seriously tackling these problems today, or we risk hitting a technological plateau within the next 30 years. And, if the industry ever becomes accustomed to a state of stagnation rather than of growth, who knows how long it will take us start things up again? For us to ignore reversible computing today seems a far more dangerous strategy than for us to pursue it with vigour.

RAM MOHAN RAO
ram_mohan@thinkdigit.com