Machines learn from biology

Since the 1940s, computers have been employed as predictable, deterministic machines, carrying out tasks by running continuously through sequences of operations. Of course, every operation, simple or complex, requires energy to drive currents, alter voltages or otherwise change the states of digital computing elements. Computing requires energy, as Rolf Landauer famously argued in 1961. Only because devices could take their energy supplies along — in the form of compact, rechargeable batteries — did we see the recent revolution in mobile computing.

Another element of the computing paradigm dominant for the past 70 years is that computers should do their assigned tasks with high ‘quality of service’, as computer scientists put it. The device should reliably run through the expected set of steps. That a device should deliver good computations while using as little energy as possible is a secondary requirement. Service first; energy preservation second.

This computing paradigm isn’t coming to an end any time soon. The United States, China, Japan and the European Union are all striving to reach so-called exascale computing in the next few years, with machines able to execute more than 10^{18} floating point operations per second. These machines will open new possibilities in artificial intelligence, weather and climate studies, and seismic analysis, as well as being put to use in national security.

But a competing paradigm may soon be just as important, and could invert the traditional preference of service over energy. Technologists have been talking for a couple of decades about the coming Internet of Things, in which compact computing devices will spill out into the world and embed themselves in the physical fabric around us, inhabiting not only our products — from toasters and handbags to articles of clothing — but also the environment, including our bodies, the atmosphere, forests and oceans. It’s happening now.

In January, the CES 2020 trade show offered a hyped celebration of the future, at least as seen by tech companies. Among the newest products revealed was a toothbrush enabled with machine-learning-based artificial intelligence. One might well question the necessity of that, or of computing and wireless communication elements tucked inside in mattresses, electrical sockets, and roof tiles. But there are real benefits to be gained in other areas, in motoring and managing buildings to save energy,

for example, or in delivering real-time information to doctors on their patients’ blood pressure, weight and other physiological parameters during extended treatments.

Our future microscopic computing devices will come to live on their own, unattended and unaided by us, essentially as independent entities

This trend faces some serious obstacles linked to the unavoidable need for energy. The number of connected devices is projected to be over 20 billion in the next year or two, and 70 billion by 2025. Until now, our computing devices have remained within the scope of our oversight and maintenance. Rechargeable batteries can be recharged, and replaced when necessary. That will increasingly become not only uneconomical, but practically impossible due to their sheer number. As an alternative, engineers will shift to the pursuit of devices able to harvest their own energy, although it will require a profound shift in computing strategy. In particular, future computing could shift to emphasizing energy over performance, only providing the maximum quality of service over certain intervals. As recently described by Sivert Sliper and colleagues (Slipert, S., Cetinkaya, O., Weddell, A. S., Al-Hashimi, B. & Merrett, G. V. *Phil. Trans. R. Soc. A* <https://doi.org/dmkd>; 2019), engineers see three progressive steps in moving toward this vision.

In step one, engineers will try to improve the lifetime of batteries by giving them some help from elements able to act as a buffer between energy supply and demand. For example, microscopic capacitors could store energy from the environment in a time of surplus, and give it back later. Such ‘energy-neutral’ devices would maintain high quality of service, although doing so would impose costs and an increase in physical size. These devices would also eventually expire due to limits to the number of charge cycles or ageing.

A step further would pursue so-called power-neutral operation, in which devices would adapt their power consumption to match the currently available harvested power, thereby reducing or even eliminating the need for energy storage. This would make the devices simpler and smaller, but at a

further cost to performance. Without energy storage, the volatile memory on which so much computation depends would be lost, so any running application would halt if harvested power fell too low, and would need to start again from the beginning whenever power returns. This kind of operation would have only niche applications, as it could not support any computations running over long periods of time.

The most ambitious goal is ‘intermittent operation’, which fully embraces how energy harvesting depends on the environment. Here the idea is to accept that power failures will be frequent and unpredictable, and design operations to proceed anyway. Long-running applications would be carried out not continuously, but with sporadic interruptions, with many intervals in which the device finds insufficient energy and goes dormant, waiting to commence again on the return of energy availability. In principle, such devices could attain maintenance-free operation and very long lifetimes at low cost.

Realizing this vision is a deep challenge. One likely application is in sensors which record events and report them to a central server. One key problem is that intermittent devices have no way to keep track of time during periods of inactivity, and timing is crucial to support functions such as wireless communications. Some research is trying to measure intervals of downtime by the discharge of a capacitor, but these efforts fall far short of the required accuracy. Another problem is designing devices to avoid “Sisyphian tasks” — tasks that would repeatedly use up all the available energy, stop before completion, and then senselessly restart again when more energy is available, only to fail again.

This vision of computing moves in the direction of biology, where some organisms can go dormant in the right conditions. Viruses routinely do this, inserting their DNA into a host’s genome until restored to activity later when their replication is more viable. Similarly, seeds or spores can remain dormant in the environment for years, protected and using no energy, before being reactivated.

But perhaps this is not surprising, as our future microscopic computing devices will come to live on their own, unattended and unaided by us, essentially as independent entities. Not quite organisms, but certainly on their way to becoming like them. □

Mark Buchanan

Published online: 6 March 2020
<https://doi.org/10.1038/s41567-020-0830-z>