Viewpoint

Christof Teuscher and Moshe Sipper

Hypercomputation: Hype or Computation?

Can we physically implement a hypercomputer? So far no one has.

In 1936, Alan Turing laid the theoretical groundwork for modern computing science (along with others, including Alonzo Church and Emil Post), by defining what would later become known as a universal Turing machine (UTM). This typewriter-like device is able to perform complex computations by employing a suitably programmed head that can move along, read from, and write over an infinite memory tape, divided into discrete squares. Indeed, according to the Church-Turing (CT) thesis—a touchstone of modern computing theory—a UTM can carry out any effective computation, or, in other words, it can simulate any other machine capable of performing a well-defined computational procedure (care should be taken when invoking the CT thesis as there exist many different formulations, and it has often been misinterpreted). But does a UTM really capture the essence of any and all forms of computing? Or, are there hypercomputers: super-Turing machines capable of going beyond the Turing limit?

Hypercomputation, which was the focus of a workshop in London [4], has a decades-long history. Turing himself presented a model of an abstract hypercomputer—the O-machine—in his 1939 doctoral thesis. An O-machine is a UTM augmented by an oracle (or "black box") performing a computation which a UTM operating in finite time cannot compute [1]. Turing gave no indication, however, on how such an oracle might be implemented.

More recently, Hava Siegelmann presented a hypercomputer model based on analog recurrent neural networks (ARNN) [5]. The ARNN is a finite network of nodes (called neurons) and connections (called synapses), wherein the synaptic weight associated with each connection (determining the coupling strength of the two end-nodes) is a real (analog) value. Siegelmann showed that ARNNs are more powerful than the Turing-machine model in that they can perform computations provably uncomputable by a UTM. When the real synaptic weights are replaced by rational numbers, the network's computational power is reduced to that of a Turing machine. Mike Stannett presented yet another hypercomputer model [6] and showed that it can solve the classical halting problem that no Turing machine is able to solve.

All hypercomputers presented to date are purely theoretical, and we may well ask whether they can actually be implemented (or, perhaps we should say pseudo-implemented; infinite memories being in short supply, a personal hypercomputer would be no more an implementation of a hypercomputer model, than a personal computer is of a UTM). So far, no one has physically implemented a hypercomputer and the strong version of the Church-Turing thesis—stating that no realizable physical device can be more powerful than a Turing machine—has not been refuted. Neither Church nor Turing, however, had demonstrated the impossibility of hypercomputers.

The practicality of physical hypercomputation has, in fact, been questioned by several researchers. Most hypercomputer models involve analog computation with infinite precision or try to compute the infinite in finite time. Given the lessons of quantum mechanics, it seems nature would not tolerate our building infinitely precise machines. Moreover, the effect of noise on analog computation presents yet

LISA HANEY

COMMUNICATIONS OF THE ACM August 2002/Vol. 45, No. 8 **23**

another obstacle to the implementation of hyper-machines. As for quantum computation—wherein physical phenomena at the quantum level are directly employed to build more powerful computers (at least in theory)—it is too early to tell whether this domain holds any promise where hypercomputation is concerned. To date, one form of a universal quantum computer has been proved to be UTM-equivalent [4]. This is not to say that quantum computers cannot compute much faster than classical ones, but that they cannot compute that which is not computable by classical, digital computers (the latter of which characterizes hypercomputers).

Assume, nonetheless, you've managed to overcome these implementation obstacles and (pseudo) implement a hypercomputer, to what use could you now put it? Perhaps you could build a "brain"; one of the major application domains of hypercomputation being envisaged is that of machine intelligence. This raises the question of whether the human brain itself is a super-Turing machine, an issue over which opinions diverge widely. Most computational brain models presented to date are less powerful than a UTM. On the other hand, the modest success of classical AI has urged researchers such as Stannett to speculate that "if biological systems really do implement analog or quantum computation, or perhaps some mixture of the two, it is highly likely that they are provably more powerful computationally than Turing machines" [6]. This statement implies that truly intelligent behavior cannot be implemented on standard digital machines, an opinion shared by Roger Penrose, who believes mechanical intelligence is impossible since purely physical processes are uncomputable.

A Turing machine is a closed system that does not accept input while operating, whereas the brain continually receives input from the environment. Based on this observation, Jack Copeland has proposed the coupled Turing machine which is connected to the environment via one or more input channels [2]. However, any coupled machine with a finite input stream can be simulated by a UTM since the data can be written on the machine's tape before it begins operation (note that a machine with a finite life span handles a finite amount of input data). Consider the following imaginary experiment: we've managed to

record all of a human's inputs and outputs, and received and emitted over a lifetime of interaction; would a UTM now be able to map the inputs to the outputs? In other words, can a human's lifelong behavior be described by a Turing-machine computable function? Copeland wrote that "it would—or should—be one of the great astonishments of science if the activity of Mother Nature were never to stray beyond the bounds of Turing-machine computability" [3].

The field of hypercomputation still encompasses several open questions of a fundamental nature, including:

- Can we (physically) build a hypercomputer?
- Can super-Turing machines pave the way toward AI? Are super-Turing machines necessary for the creation of AI?
- Does the brain perform functions not computable by a Turing machine? Is thinking more than computing? Does thinking amount to hypercomputing?

So, hype or computation? At this juncture, it seems the jury is still out—but the trial promises to be riveting. ▣

REFERENCE
1. Copeland, B.J. and Proudfoot, D. Alan Turing's forgotten ideas in computer science. *Scientific American 280*, 4 (Apr. 1999), 77–81.
2. Copeland, B.J. and Sylvan, R. Beyond the universal Turing machine. *Australasian J. Philosophy 77, 1* (Mar. 1999), 46–66.
3. Deutsch, D. Quantum theory, the Church-Turing principle of the Universal Quantum Computer. In *Proceedings of the Royal Society of London* (1985), 97–117.
4. Hypercomputation workshop. University College, London (May 24, 2000); www.AlanTuring.net/turing_archive/conference/hyper/hypercom.html.
5. Siegelmann H.T. Computation beyond the Turing limit. *Science 268* (Apr. 1995), 545–548.
6. Stannett, M. X-machines and the halting problem: Building a super-Turing machine. *Formal Aspects of Computing 2, 4* (1990), 331–341.

CHRISTOF TEUSCHER (christof@teuscher.ch) is a researcher in the Logic Systems Laboratory of the Swiss Federal Institute of Technology in Lausanne.
MOSHE SIPPER (sipper@cs.bgu.ac.il) is an associate professor in the Department of Computer Science, Ben-Gurion University, Israel.