Uncomputability, Intractability and the Multiverse

Ruben Ruiz-Torrubiano

Abstract

If the many-worlds interpretation of quantum mechanics is correct, is there any way to harness the computational power of the multiverse to compute in practice intractable problems? We introduce the Computational Censorship Hypothesis to answer this question and explore its implications.

The Multiverse and Tractability

The *many-worlds interpretation* of Quantum Mechanics (QM) was first proposed by Hugh Everett in 1957 [1]. In Everett's interpretation, there is no collapse of the wave function, as it is the case in the Copenhagen interpretation. Instead of that, the wave function evolves deterministically governed by Schrödinger's equation alone. As a consequence, all possible outcomes of a given quantum measurement are actually realized on a parallel universe, branching from the universe the measurement was performed on.

We consider the question whether, if such parallel universes are real, they can be used to make computations which are otherwise unfeasible or intractable, which is to say in practice uncomputable. First, let us take a look at a particular, intractable problem.

The Traveling Salesman Problem (TSP) is a classic example of a problem that is easy to state, easy to verify and hard to solve. In its essence, the problem involves a list of cities and their pairwise distances. The objective is to find the optimal tour (i.e. the route that minimizes the total distance) for a salesman such that each city is only allowed to be visited once (and the salesman returns to the starting point). This problem is known to be NP-hard [2]: There is no known efficient algorithm for solving it. By "efficient" we mean an algorithm running in time polynomial on the number of cities given as input. In practice, the problem is solvable by means of approximation algorithms, which are not guaranteed to return the optimal solution, but a solution which is no worse than the optimal one by a given amount. Finding the optimal solution by brute-force (checking all possible tours and picking the best one) would typically require longer than several orders of magnitude of the age of the universe.

In theory research, computer scientists use models of computation which are used to abstract the main aspects of computation independently of physical computing devices. A Turing machine (TM) [3] is a mathematical abstraction of a digital computer which

has exactly the same computational capabilities (i.e. everything a digital computer can compute can also be computed by a TM – which is called *Turing completeness*) and can be used for theoretical algorithm analysis. In a nutshell, a TM is a mathematical object having a writing/reading head (that can be moved one position at a time to the left or to the right), an input tape, a work tape and an output tape, where only one symbol of a given alphabet Σ can be written on each cell of the tape (plus an "empty" symbol). You can think, for instance, of the binary alphabet $\Sigma = \{0, 1\}$ as the one used by modern computers. Then, given the state of the machine at any given time and an input symbol, a transition function δ returns the next state of the TM, the next symbol that will be written to the tape, and where the head is to be moved to the left or to the right (or remains unmoved). In general, there is an initial state and an accepting or a rejecting state. In a deterministic TM, the transition function is allowed to return exactly one transition: there is a deterministically unique next state according to the current state of the TM and the next symbol to be processed. On the contrary, a non-deterministic TM is allowed to return a set of states (that can be empty) which represents all possible computations that are feasible and lead to an accepting state. In other words, a non-deterministic TM is able to perform all necessary computations at once, picking only the optimal or successful one.

A helpful way of thinking about determinism and non-determinism in computation is to see a deterministic computation as a linear sequence of steps which lead to a result from an initial state. On the contrary, non-deterministic computations can be seen as decision trees involving all possible computation paths, as shown in Figure 1.
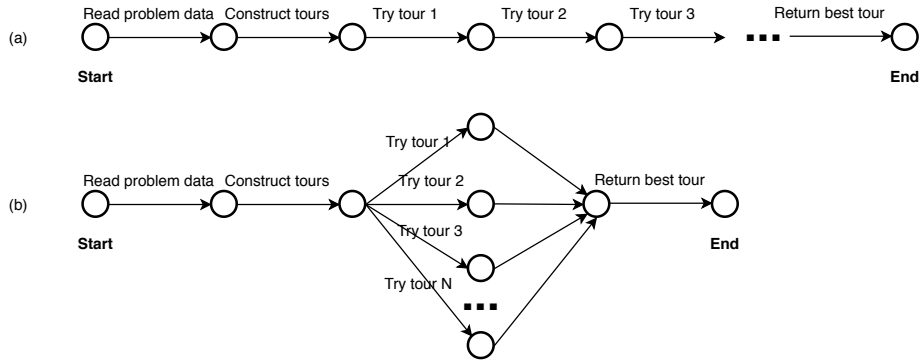

Figure 1: (a) Computing the best tour for a TSP instance deterministically by brute-force. (b) Non-deterministic computation for the same instance.

Of course, such trees can also be computed deterministically, since any deterministic TM can be used to emulate a non-deterministic one, just as you would do it with a pencil and a piece of paper: just enumerate all possible states and transitions, write them down, and select an optimal path. The caveat is, for this brute force approach to work you would actually need worst-case exponential time in the size of the input. So, in practice, you would not be able to compute it.

2

Non-deterministic computations

Let us explore what non-deterministic computations have to do with the many-worlds interpretation. We will construct a hypothetical device whose computational behavior can be modeled by a non-deterministic TM, and call this device *non-deterministic computational device* (NDCD). One possibility to implement such a NDCD would be to branch on a different parallel universe at any junction in the algorithm's decision tree. This branching could be implemented by attaching to each junction a physical quantum measurement, like measuring the spin of a photon. Consider the case when the computation is about to branch to a given set of states. The current state of our machine would be described by the Hamiltonian H of some quantum system. By construction, measuring the state of the system by our NDCD would result in computing each branch in a different parallel universe. In order to compute the final result, every accepting branch in each parallel universe would have to be checked and the best result according to some preference criterium (in an optimization problem, that would be maximizing or minimizing a given utility function or cost) would be returned. Figure 2 shows a schematic representation of the function of the machine.

How could this check procedure work? One possibility would be that an accepting state triggers the quantum measurement of an entangled particle. The particle's entanglement would then produce new branchings were the result of the computation is known. In particular, all parallel observers in all parallel universes would perceive the result of the measurement. So any from these parallel observers would get the same result, say accept or reject. An optimization problem could be solved by asking the question: *does a solution exist with a cost function value lower than K?* and then search for the minimum K (for instance, by applying a binary search method).

We argue that such a computation would be precluded by the Computational Censorship Hypothesis (CCH), which can be informally written as:

> **Hypothesis (Computational Censorship Hypothesis)** *"If Everett's many-worlds interpretation of quantum mechanics is correct, it is not feasible to harness the computational power of parallel universes in a particular universe".*

The main argumentation for the CCH is information theoretical in the sense that, if parallel universes exist in the Everett sense, there can be no information exchange between them. Any device that is able to exploit parallel universes for computational tasks would violate this principle, since to perform any computation it would be necessary to transfer information between parallel universes. Moreover, since new information means lower entropy, the direct implication would be an entropy decrease in all newly branched parallel universes.

One slightly more formal statement of the CCH would therefore be:

> **Hypothesis (Computational Censorship Hypothesis)** *"If Everett's many-worlds interpretation of quantum mechanics is correct, there can be no information exchange between parallel universes and, therefore, no computations involving parallel universes are allowed".*

Therefore, if the CCH is true, then it is impossible to construct any NDCD. If the multiverse exists, this impossibility extends to all possible parallel universes. That would mean that, as far as computability and predictability are concerned, and if the multiverse is a part of our physical reality, it cannot be used to perform computational tasks.

If, on the other hand, the CCH could be proven not to be true, then the consequences for uncomputability and unpredictability would be manifold. First, in practice uncomputable or intractable problems could experience an super-polynomial speedup. Specially tailored algorithms for the NDCD could be written to approach any known computationally hard problem. It could be even possible to write a general interpreter that, for a given problem and instance definition, could run a general algorithm launching a brute-force non-deterministic search.
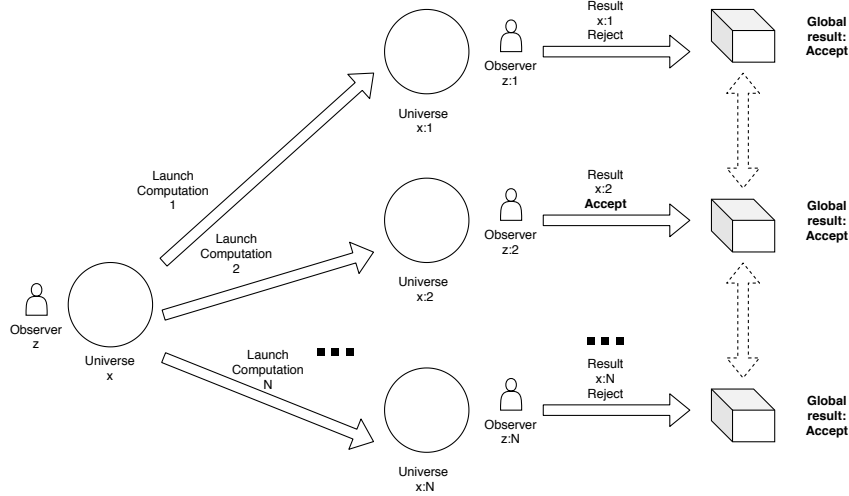

Figure 2: Schematic representation of the workings of a NDCD based on Everettian parallel universes. The computation is started in our current universe, then parallel computations are started as different possible measurement outcomes of a system described by a Hamiltonian H. An accepting state in a parallel universe $x : i$ branched from universe x triggers a measurement, which causes the accepting state in all parallel universes to materialize (global result of the computation).

What about making predictions about processes which are otherwise unfeasible to predict? Would a NDCD posess an advantage regarding the capacity of making

predictions? Predictability, or the lack thereof, has been raised as an issue with Everett's many-worlds interpretation [4]. If every possible realization of a given process actually happens in a parallel universe, then the multiverse as a theory is not falsifiable (but it may be a prediction of a falsifiable theory). It has been argued [5] that what we as observers perceive as randomness (or lack of predictability) is just a realization of a given outcome in one of infinitely many possible parallel universes. In this sense, the physical laws of quantum mechanics would represent a principal obstacle to using a NDCD to predict measurement outcomes in a given universe.

Recent studies have found evidence of quantum supremacy in a noisy intermediate-scale quantum (NISQ) device solving the problem of sampling the output of pseudo-random quantum circuits [6]. Can we interpret such results in terms of the CCH? If the CCH is not true, then qubit-based quantum computers can indeed perform computations in parallel Everettian universes, but using a different principle as the NDCD described earlier. In the ordinary quantum setting, qubits are set in a super-position of quantum states and computations are performed on those qubits, trying to avoid decoherence. The result of the computation is then measured and, with a high level of confidence, the result of the measurement is a solution to the problem at hand. So measurement is a key component of any quantum algorithm. Moreover, the Hamiltonian of the particular problem instance encodes the probabilistic amplitudes of observing the system in a given configuration. Both components are missing on NDCDs: these are able to actually perform all computations in parallel. It is an open problem to determine how non-determinism is related to quantum computing, as expressed by their respective complexity classes BQP and NP. So it is generally not known if there is some kind of equivalence between quantum computers and NDCDs.

Conclusion

We have explored the consequences of Everett's many-worlds interpretation of quantum mechanics for unpredictability and intractability. By harnessing the computational power of parallel universes, non-deterministic computing devices could be used to make in practice intractable problems computable. Using these non-deterministic devices, super-polynomial speedups on otherwise practically uncomputable problems could be achieved. We have proposed the Computational Censorship Hypothesis as a way of answering the question if such devices are possible, and provide a (negative) answer based on information theoretic considerations.

References

[1] Hugh Everett. *Theory of the universal wavefunction.* PhD thesis, Princeton University, 1957.

[2] The traveling salesman problem. In *Combinatorial optimization: algorithms and combinatorics*, volume 21. Springer, Berlin, Heidelberg, 2008.

[3] Alan Turing. On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, s2-42(1):230–265, 1936.

[4] Sabine Hossenfelder. *Lost in math: how beauty leads physics astray.* Basic Books, New York, 2018.

[5] Max Tegmark. *Our mathematical universe: my quest for the ultimate nature of reality.* Penguin Books, 2014.

[6] Frank Arute, Kunal Arya, Ryan Babbush, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.