

Introduction

Quantum information theory is an exciting new development at the intersection of physics, mathematics and computer science. It is based on the empirical observation that quantum systems behave in a fundamentally different way from “classical” ones.

In the study of digital information we may profitably ignore many of the specific aspects of a system. We do not care whether our information is stored electrically, magnetically, optically, or even on paper punch cards, all that matters is that it can be interpreted as bits, and that the system is capable of manipulating those bits.

A similar situation obtains in the study of quantum information systems. We need not concern ourselves with whether quantum information is represented using photon polarization, nuclear spin, etc., only that our system is capable of storing and manipulating quantum bits, or “qubits”. This is not to trivialize the matter. To date, only a few specialized quantum information systems have been physically realized. Many scientific and engineering challenges remain, foremost among them, how to faithfully store collections of qubits in a way that avoids decoherence.

The great – and as yet unrealized – prize is the construction of a general purpose *quantum computer*. What could such a machine do? It would not be able to compute any (mathematical) functions that cannot be computed on a classical computer, such as the halting problem for Turing machines or the word problem for groups. The reason is that we can (very inefficiently) *simulate* a quantum computer on a classical one.

However, a quantum computer would potentially make many computable but infeasible functions feasible. For reasons that we will learn about shortly, quantum information systems have the property that they contain an inherent parallelism that can lead to a massive reduction of run-time resource utilization of programs. Such a *quantum speed-up* is what will potentially bring whole classes of problems into the realm of feasibility. This will open up new possibilities, and potentially, create new difficulties.

One example of a quantum speed-up is the task of searching an unsorted list for

a particular element. Classically, the best we can do is to examine each element in turn, so we would expect to find our target in $n/2$ time on average. However, there is a known quantum algorithm (due to Grover) that can do it in \sqrt{n} time. A machine that could execute such an algorithm would be invaluable to the likes of Google and Oracle.

Another, potentially even more disruptive example involves prime factorization. The best known algorithms for this task take time exponential in the size of the input. But there is a known quantum algorithm (due to Shor) that can do it in polynomial time. A machine that could execute such an algorithm would be invaluable to mathematicians – and to intelligence agencies, because the infeasibility of this, and related, problems underlies much of the cryptography that we rely on for protecting everything from national security secrets to our credit card numbers when shopping online. But in this case, when the universe closes a door, it opens a window: quantum information systems introduce the possibility of *quantum cryptography*, in which resistance to certain types of attack has the status of a natural law.

The problems of searching and factoring lie within a (classical) complexity class known as NP. Such problems may (as in the case of factoring) be hard to solve, but have the property that a candidate solution is easy to verify – for factoring, you just need to multiply. But there are many computational problems of interest that lie beyond even the complexity class NP. Salient among these are physical simulations. When we write a computer program to simulate a chemical or nuclear reaction, the only way we can know whether our simulation is any good is to compare its predictions to experimental results. For a classical computer, the simulation typically involves exponential overhead, so the behavior of non-trivial systems is both infeasible to compute and impossible to check without running the experiment anyway.

For physical systems for which quantum mechanics provides an accurate description (either because it turns out to be the fundamental theory of our reality or because whatever other effects there may be are negligible) we should be able to *efficiently* simulate their behavior using a quantum computer. The conjecture that *all* physical systems are inherently quantum, and can thus be efficiently simulated by a quantum computer is known as the quantum strong Church-Turing thesis.

Through decades of experimenting and theorizing, people have discovered mathematical structures which seem to accurately describe the behavior of quantum information systems; that is, they have formulated a *theory* of quantum information.

Basically, quantum information systems are represented by finite-dimensional vector spaces with an inner product operation, known as (finite-dimensional) *Hilbert spaces*. From the structure of Hilbert spaces, we are able to derive properties of quantum information systems. This presentation seems to be consistent with all known experimental observations. However it has some drawbacks: is

not modular, especially perspicuous, or for that matter, very well-motivated. We will incrementally come to understand some of the reasons for this, but to (perhaps over-) generalize, the core issue is that it is an analytic characterization.

Consider the apt analogy of geometry. There are two general approaches that can be taken. One is *analytic*. Historically this approach came much later, but for many applications it has proved wildly successful. In the analytic approach, properties of geometry emerge as consequences of the properties of another, ambient, system; e.g. of vector spaces with a choice of basis and inner-product over the real numbers. This approach is highly amenable to computation – indeed, most computer geometry-based systems take this approach – but the flip side of this fact is that many qualitative facts about geometry require a lot of computation.

In contrast, the synthetic approach to geometry is the one pioneered in antiquity by Euclid. Here the entire theory is reduced to a collection of basic axioms (e.g. Euclid’s famous five postulates) and a collection of rules by which new theorems may be deduced from old. This approach has the advantage of perspicuity: proofs tend to have a natural structure. It is also modular: famously, the parallel postulate may be swapped out for a variant, leading to spherical or hyperbolic geometry.

A *synthetic* approach to quantum information systems shares these advantages of modularity and perspicuity. Furthermore, it is amenable to the tools of a branch of math called *category theory*, and in particular, to description by diagrammatic *graphical languages*.