Quantum Computers By Degrees

Our take from the debate between Gil Kalai and Aram Harrow

Volker Strassen has made many famous contributions to theoretical computer science, but one might not know that they include a limerick. Yet the following exchange in 1998 with Peter Shor and his wife prefaces Chapter 5 in the bellwether quantum computing textbook by Michael Nielsen and Isaac Chuang:

> If computers that you build are quantum,
> Then spies everywhere will all want 'em.
> Our codes will all fail,
> And they'll read our email,
> Till we get crypto that's quantum, and daunt 'em.
> – Jennifer and Peter Shor.
>
> To read our E-mail, how mean
> of the spies and their quantum machine;
> be comforted though,
> they do not yet know
> how to factorize twelve or fifteen. – Volker Strassen.

Today we, Dick and I, explain our own skepticism of quantum computing, and say what we have learned from the detailed debate between Gil Kalai and Aram Harrow.

The number 15 was factorized by IBM researchers three years later, in 2001, though it took runs by other groups to be fully convincing that Shor's factoring algorithm was being implemented in a faithful manner, with observable entanglements being created. The most recent repetition for 15 was published this year, and claims a 48% success rate. That 15 has remained the highest number for over a decade makes it reasonable for us to propagate Strassen's implied question:

> When will they factorize twenty-one? Or eighteen, or thirty?

Strassen also proved the best-known general lower bounds on arithmetical circuits for low-degree polynomials, in conjunction with his "close friend" Walter Baur among colleagues he salutes in verse. As I covered here, the bounds arise in terms of the *geometric degree* of an algebraic variety associated to the polynomial. I suspect something like Strassen's mechanism operates in the quantum realm, but governing *entanglement* rather than merely circuit size.

Ken: Shor-Sure but Grover Not Bowled Over

The thrust of my skepticism is that *the standard gate-counting measure of the quantum circuit model understates the effort required to operate the circuit.* Of course all quantum-computing skepticism asserts this; what's particular to mine is the nature of the overlooked cost factor. I am fairly confident that it is more than a constant factor, and also that it is at most polynomial. The latter actually makes me "Shor-sure" in Scott Aaronson's terms, i.e. a believer in the polynomial feasibility of Shor's algorithm. The main questions within my skepticism is whether the extra factor is linear or logarithmic, and if the latter, whether it is absorbed by the log factor overhead already present in the Quantum Fault Tolerance Theorem, or separate and paid up-front.

I was led to this by descriptions of Grover's algorithm as running "in sub-linear time," *viz.*, $O(\sqrt{N})$ overall work. With intuition admittedly of early 20th Century vintage, I think a procedure that polls N locations must expend N units of effort. I might not have trouble in cases like solving SAT where $N = 2^n$ and n physical qubits are used to search N assignments that are not really in N "locations."

Moreover, it is not clear whether "effort" should mean *time* or *work* (i.e., entropy change) or something in-between, noting that evolutions according to Schrödinger's equation take time but are reversible and dissipate no energy. One thing we can postulate, however, is that an effort measure $E(C)$ for a circuit C—quantum or otherwise—should satisfy the following axiom:

> If C is a disjoint union of circuits C_1 and C_2, with no connections between them, then $E(C) = E(C_1) + E(C_2)$.

The analogous notion of disjointness for quantum systems ψ_1 and ψ_2 is the product state $\psi_1 \otimes \psi_2$. An *entanglement measure* $\mathcal{E}$ is said to be (*strongly*) *additive* if $\mathcal{E}(\psi_1 \otimes \psi_2) = \mathcal{E}(\psi_1) + \mathcal{E}(\psi_2)$. A rider of such measures is that whenever ϕ is comprised of ϕ_1 and ϕ_2 that are mutually entangled, then

$$\mathcal{E}(\psi) > \mathcal{E}(\psi_1) + \mathcal{E}(\psi_2).$$

At bottom my skepticism asserts that a similar inequality holds for a natural measure $E(C)$ that lower-bounds the "true effort" to implement the quantum circuit C, whenever C is broken into entangled pieces. Let's contrast this with "classical" circuits.

Quantum Versus Boolean Circuits

In simplistic terms, here is the main complexity distinction from familiar Boolean circuits that I feel must be accounted for. A Boolean circuit C has a well-defined local value at every juncture. If you break C into pieces C_1, C_2 such that no wires go from C_2 to C_1, then you can call the values on each wire going from C_1 to C_2 as outputs of C_1 and inputs of C_2, writing the function f computed by C as the composition of f_1 computed by C_1 and f_2 computed by C_2. The point is that then the complexity of f equals the sum of the complexity of f_1 and that of f_2. The whole equals the sum of the parts because all values are local. This intuitively holds even under ways of breaking the circuit that have wires going both ways between pieces.

For quantum circuits this is not so because of entanglement. The simplest example is a two-qubit circuit consisting of an Hadamard and a controlled-NOT gate, on the basis input $|00\rangle$.

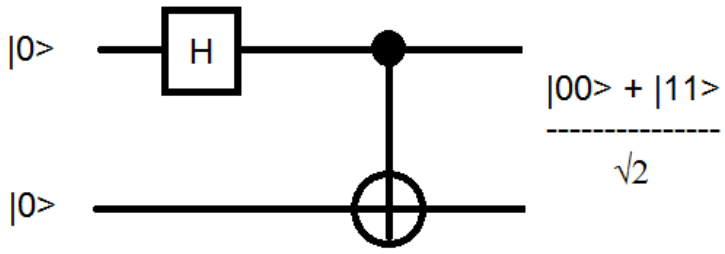

At the right-hand ends one speaks of two qubits, but neither qubit has an individual value. Attempting to define a local value for either qubit entails *tracing out* the other one, which leaves a classical random bit. Saying the values are two classical random bits is false in view of the entanglement. Hence the whole is not the sum of its local parts. But summing local steps is what we do in complexity measures. Thus I contend the true complexity picture must be something other than what the diagram leads us to believe.

The "quantum grille" I described here supplies a local variable for each juncture of the circuit, but on pain of involving computations that are generically #P-complete, well beyond the believed power of the circuits.

While there is consensus on quantifying entanglement between two systems, there are a wide unresolved variety of proposed multipartite entanglement measures, such as for n-qubit (pure) states in general. Hence there is no salient notion of entanglement produced

by circuits either. One subtlety is that mapping a non-entangled state to an entangled one is not a circuit invariant, even in the binary case. The above circuit applied to the product state $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)|0\rangle$ produces the product state $|00\rangle$. But finding a best measure of "entangling capacity" for a circuit is something we should try to do.

Two Ways to Break Quantum Circuits

To supply first the intuition for why Strassen's degree measure may be relevant, let us look more closely at how quantum circuits C might be pieced apart. Let's use notation:

- n for the number of qubits; n_0 for inputs if we need to distinguish them from ancilla qubits.
- $N = 2^n$; $N_0 = 2^{n_0}$ if needed.
- k for the maximum (or typical) arity of a gate—usually $k = 2$ or $k = 3$ so we can ignore this.
- s for the number of quantum gates—or alternately, the sum of 2^k over all gates.
- h for the number of Hadamard gates, or similar non-deterministic gates. (More formally, $h = \log_2 R^2$ where R is the product of normalizing constants in the gates.)
- d for the number of *levels* in the circuit, so $\frac{s}{n} \leq d \leq s$.

Flying by intuition not definition for now, let's consider how $E(C)$ might grow as we break down and re-assemble the circuit, in two different ways.

• *By qubits:* Break off each qubit line. It may be entangled with the rest of the circuit. Using the simple gate-counting measure to represent the "heft" of what the line is entangled with, we get s as the cost of each line, for an upper bound of sn on $E(C)$.

• *By levels:* Now we first argue that we can ignore the issue of entanglements between levels by thinking of each level as a discrete timestep and submerging it into the issue of maintaining the quantum state entering each timestep. Thus we estimate $E(C)$ by d times the effort needed to maintain an n-qubit state. We analogize this to the *description complexity* of (approximating) the n-qubit state, and take a special cue from graph states, which are important in quantum error-correcting codes.

In this 2005 paper by Caterina Mora and Hans Briegel, general n-qubit graph states are argued to have description complexity of order about n^2. This gives dn^2 for the whole again. The preparation complexity measure of Peter Høyer, Meidi Mhalla, and Simon Perdrix for graph states promises lower effort, but may be circular in this context as its its value is stated as the gate-counting size of a quantum circuit that does the preparation. However, the by-levels breakdown allows that the answer could be lower like $O(dn)$ or $O(dn \log n)$.

Information bounds may answer the intuition

There is a separate consideration that is also highlighted by graph states. An n-vertex directed graph G can encode $n(n-1) \approx n^2$ bits of information according to its edges and non-edges. However, the n-qubit graph state prepared from G does not allow extracting more than n bits of classical information, by Holevo's theorem. This may argue that the "heft" at each level is really only n after all, which over d levels yields dn, and leaves just the gate-count s when $d = \Theta(s/n)$.

Privately, Aram has told me of work on a paper that addresses fine-grained concerns of encoding graphs by qubits, and it is possible that this may go even further than Holevo's bound to blunt the complexity intuition here. However, for now we say that the "by-levels" breakdown argues that the effort measure $E(C)$ (at least for C involving graph states) is typified by $d \cdot n \cdot$ "some factor," where the factor may be constant, n, or something in-between. We now apply Strassen's measure to intuit where an "in-between" factor might come from.

Strassen's Measure

Strassen's measure $\mu(f)$ for a function $f(x_1, \ldots, x_n)$ is the **geometric degree** of the mapping

$$V = (y_1 - \frac{\partial f}{\partial x_1}, \ldots, y_n - \frac{\partial f}{\partial x_n}),$$

where the y_i are fresh variables. This is definable as the maximum finite value of $||V \cap A||$ for an n-dimensional affine linear subspace of $\mathbb{C}^{2n}$. Aided by Walter Baur on a lemma connecting the complexity of f to that of its partial derivatives, Strassen proved (see this for more details):

Theorem 0.1. *Every arithmetic cricuit computing f must have at least $m = \frac{1}{2}\log_2(\mu(f))$ multiplication gates.*

For example with $f = x_1^d + \cdots x_n^d$, we can define A by setting $y_i = 1$ (or rather, $y_i = n - 1$) for each i, so that $V \cap A$ consists of all n-tuples of $(d-1)$st complex roots of unity. Since there are $(d-1)^n$ such tuples, we have $\mu(f) \geq (d-1)^n$, and in fact this is tight because $(d-1)^n$ is the maximum possible geometric degree for a mapping V defined by equations of degree $d-1$ in n variables. Thus we have $m = \Omega(n \log d)$, which for $d = n^{O(1)}$ gives $\Omega(n \log n)$.

This is still the best-known lower bound for general low-degree polynomials such as the permanent, even though it holds for this simple function f. There are notions of "arithmetic degree" that take exponentially higher values, but analogues of Strassen's theorem either fail or are unknown for them. What further significance might these ideas of "degree" have?

Application to Quantum Circuits

My work introduced in this post associates to every quantum circuit C (whose gates satisfy a minimal "balance" condition) a polynomial P_C in variables $z_{i,j}$ for i ranging over qubits and j over the d levels. The polynomial P_C is a product of polynomials for each level, but this becomes a simpler product P'_C of polynomials p_g for each gate g upon performing substitutions that leave $n + h$ variables. Each P_g has degree 2^a where a is the arity of g. So for the overall arity bound k, P'_C has degree bounded by $2^k s$, which for bounded k we may just call s.

The maximum possible geometric degree is then $(n + h)^{2^k s}$, but there is reason to believe it is also tight for many natural and simple circuits, as for the simple polynomial f above. Taking logs gives us the following analogue of Strassen's circuit size lower bound:

$$m = 2^k s \log_2(n + h).$$

When $h = \Theta(s)$ and k is constant, this is $\sim s \log_2 s$. Thus if m is a lower bound on the hypothesized true-effort measure $E(C)$, then the effort has an extra non-constant factor on top of the gate-counting measure.

One plus point of defining $E(C) = \log_2 \mu(p)$ for some polynomial p, itself a product of terms, associated to C is that when the circuit is a disjoint union of C_1, C_2 we can expect the respective p_1, p_2 to have disjoint variable sets, and give $p = p_1 p_2$. Then we have

$$\mu(p_1 p_2) = \mu(p_1)\mu(p_2),$$

so $E(C) = E(C_1) + E(C_2)$ as desired. The question is whether P_C or P'_C (by invariance under projections or substitutions as described here they should be equivalent) is the "right" polynomial to serve as p. Since the whole framework is a discrete version of Richard Feynman's "sum-over-paths" paradigm, it is possible that a more-relevant polynomial to use as p has a variable for each path, whose value represents the contribution of that path. In that case, if the polynomial's degree is still s (notwithstanding $2^k s$), we would have

$$m = s \log_2(n + 2^h) \approx sh.$$

This would yield the case discussed above of the extra factor being linear rather than logarithmic. And this extra factor h would, I contend, show as an immediately felt scaling obstacle, even more than the $\log s$ factor.

In any event, the above supplies grounds for my contention that the gate-counting measure undershoots by a non-constant factor that is immediately felt for small n. It is not a "galactic" consideration like the amount of (classical) circuitry that an fit in a given volume being physically asymptotically quadratic rather than cubic owing to the Bekenstein bound. It also appears to be felt on top of, rather than subsumed by, the $O(\log s)$ factor from fault-tolerant encodings.

Lessons From the Debate

One strong lesson emphasized in every argument by Aram is that sustaining a skeptical position on quantum computing ultimately requires establishing a new property of physics. I had originally thought I was avoiding this by demurring only from the quantum circuit *model* and its simple counting measure. Now I appreciate that my position is willy-nilly asserting a property lurking in the geometry of quantum field theory, at least Feynman's version of it. I have asked about such a possibility for years, and have found hints in papers such as this and related ones by Michael Nielsen, this by Howard Brandt, and this by Joseph Landsberg, but have not yet struck gold. My attraction to Gil Kalai's skeptical papers was that they might supply a physical mechanism behind the abstract reasoning, but this is what Aram finds lacking.

Open Problems