Technical description of peaked circuits

qBitTensor Labs

1 Background

In recent years, much interest has developed toward quantum benchmarking, wherein one develops quantum-computational problems to suitably test the capabilities of a given quantum computer or simulator. Indeed, this is a preeminent goal of all so-called "quantum supremacy" experiments: The overarching aim is to describe a series of operations that can be performed by a quantum computer or simulator such that, if complete with results that satisfy a given bound in an appropriate metric, one can know with high confidence that the computer or simulator is capturing the "quantum-ness" of the operations.

The peaked circuit is one recently proposed idea, which is defined as a quantum circuit that takes a fixed input state to an output distribution over the computational space that contains a significant concentration of probability mass in a particular arbitrary computational basis state. Thus, successful completion of a peaked circuit benchmark is entailed by (a) execution of the peaked circuit, and (b) correct determination of the basis state on which the output distribution is concentrated. In a proper construction, the peaked circuit is qualitatively close to a completely randomized circuit and requires the executor to faithfully model the coherence and entanglement present in the quantum state (more on this below).

2 Generation and properties of peaked quantum circuits

Here, we take a quantum circuit on N qubits as a unitary operator acting on the appropriate Hilbert space describing the N-fold tensor product of concatenated two-level quantum systems. Any such circuit may be described by a complex-valued matrix $U \in \mathbb{C}^{2^N \times 2^N}$ with the property that U preserves the standard induced 2-norm of the underlying Hilbert space and has unit determinant. In particular, adapting Def. 1.1 in arXiv:.2404.14493, we may say that a given circuit C is δ, q-peaked if

$$\max_{s \in \{0,1\}^N} |\langle s|C|q\rangle|^2 \geq \delta \tag{1}$$

That is, the maximum probability mass accrued by any output state $|s\rangle$ for fixed input state $|q\rangle$ through the action of C is bounded from below by a real constant $\delta \in (0, 1]$.

Given fixed $|q\rangle$, one then may ask how to obtain a peaked circuit. Of course, one can obtain a $1, q$ -peaked circuit trivially by simply generating an arbitrary series of gates, appending the inverses of those gates in reverse order, and applying as many single-qubit NOT gates as it takes to transform $|q\rangle$ into a fixed target state $|s_0\rangle$. However, this structure is easily detected simply by analyzing the circuit classically, defeating the purpose of the benchmark.

Instead, we start by generating an approximately Haar-random state; that is, for input state $|q\rangle$ fixed to the N-qubit all-zero state $|0\rangle^{\otimes N}$, we sample a random circuit U_{r} of depth τ_{r}

$$U_{\mathrm{r}} \equiv \mathcal{T}\left\{\prod_{t=0}^{\tau_{\mathrm{r}}-1} \bigotimes_{k=0}^{\lfloor \frac{N-\bar{t}}{2} \rfloor} u^{(d,k)}_{2k+\bar{t},2k+\bar{t}+1}\right\} \tag{2}$$

where $\bar{x}$ denotes the reduction of x modulo 2 and $\mathcal{T}$ the time-ordering superoperator. Each $u_{q,q'}^{(t,k)} \in \mathrm{SU}(4)$ at depth t on qubits $q < q'$ is drawn uniformly under the Haar measure and applied in a brickwork arrangement to produce

$$|\psi_{\mathrm{r}}\rangle \equiv U_{\mathrm{r}}|0\rangle^{\otimes N}. \tag{3}$$

We then feed $|\psi_{\mathrm{r}}\rangle$ as input into another circuit Υ_{p} of depth τ_{p},

$$\Upsilon_{\mathrm{p}} \equiv \mathcal{T}\left\{\prod_{t=\tau_{\mathrm{r}}}^{\tau_{\mathrm{r}}+\tau_{\mathrm{p}}-1}\bigotimes_{j=0}^{\lfloor\frac{N-\bar{t}}{2}\rfloor} v_{2k+\bar{t},2k+\bar{t}+1}^{(t,k)}\right\} \tag{4}$$

with each $v_{q,q'}^{(t,j)} \in \mathrm{SU}(4)$ applied similarly in a brickwork arrangement but chosen specifically to peak $|\psi_{\mathrm{r}}\rangle$ to at least δ, i.e.

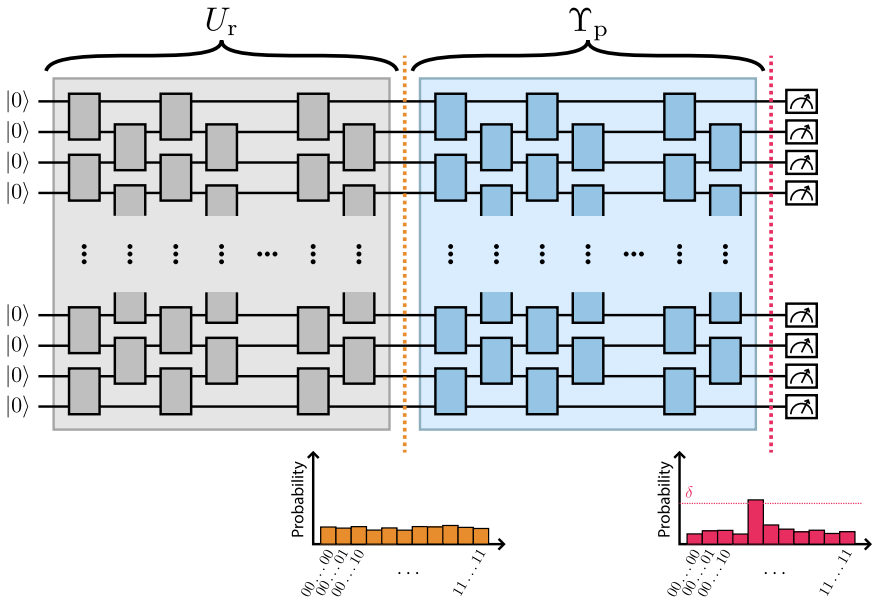

Figure 1: Peaking action of the circuit $C = \Upsilon_{\mathrm{p}} U_{\mathrm{r}}$. An approximately Haar-random state (orange) is prepared by the random brickwork circuit U_{r}, and passed through the peaking circuit Υ_{p}. Just before the final measurement layer (red), the state's probability distribution contains a peak with probability at least δ.

Thus, the benchmarking task we provide is simple but demanding:

Problem 1. (Peaked circuits)
Given a $\delta, 0$-peaked circuit C on N qubits, find the output state $|s_0\rangle, s_0 \in \{0,1\}^N$ such that

$$s_0 = \underset{s\in\{0,1\}^N}{\arg\max}\ |\langle s|C|0\rangle^{\otimes N}|^2 \tag{5}$$

3 Entanglement in random brickwork circuits

Problem 1, as with almost all problems in quantum computing, is subject to a fundamental exponential bound: By applying brickwork circuits comprising arbitrary two-qubit unitaries, we generate a significant amount of long-range entanglement in the quantum state, the information-theoretic content of which becomes an unavoidable obstacle to simulation. Simultaneously, we have fine-tuned the parameters of our circuits to avoid complete saturation of the well-known upper bound on entanglement–this is what allows them to even be simulated in the first place.

More specifically, the entanglement of a given state $|\psi\rangle$ is quantified using the entropy of entanglement, typically using the family of entropies known as the Rényi entropies of orders α across a given bipartition of the state, $A \otimes B = \mathcal{H} \ni \psi$:

$$S_\alpha^A(|\psi\rangle) = \frac{1}{1-\alpha} \log \operatorname{tr} \rho_A^\alpha \tag{6}$$

where ρ_A denotes the partial trace of the density matrix of $|\psi\rangle$ over B,

$$\rho_A = \operatorname{tr}_B |\psi\rangle\langle\psi| \equiv \sum_{b \in \text{ basis}(B)} \langle b|\psi\rangle\langle\psi|b\rangle. \tag{7}$$

Usually, one chooses to work with the von Neumann entropy of the state, which the Rényi entropies reduce to in the $\alpha \to 1$ limit. The von Neumann entropy can be expressed more naturally as

$$S^A(|\psi\rangle) \equiv \operatorname{tr}[\rho_A \log \rho_A]. \tag{8}$$

Notably, the entropy of entanglement in a fixed state $|\psi\rangle$ is a function of two variables. The first is the state itself, and the second is the exact bipartition chosen to evaluate S^A. As a function of the latter, the entropy reaches a maximal value for A and B chosen as symmetric halves of the total Hilbert space. Further, there also exist states of maximal entropy, saturating a global maximum

$$S_{|\psi\rangle}^{\max} = \frac{N}{2} \log 2 + \mathcal{O}(1) \tag{9}$$

in units of nats (1 bit $= \log 2 \approx 0.693$ nat). A state of maximal entanglement entropy can be prepared by means of a brickwork circuit comprising random two-qubit unitaries (as above), and the entropy in the state accrues entanglement as it progresses through the circuit as approximately this function of circuit depth t:

$$S_{|\psi\rangle}(t) = S_{|\psi\rangle}^{\max}\left(1 - e^{-t/N \log 2}\right). \tag{10}$$

From a practical perspective, this entanglement entropy is a fundamental barrier to computation because it represents irreducible information encoded in the state. Although algorithms exist to factor a state into smaller parts (see also so-called matrix product states), one cannot do so past a certain limit without significantly decreasing the fidelity of the computation: This manifests as growth of the irreducible part of the quantum state as it passes through successive layers of our circuits, and forms half the basis for the hardness of our circuits as a whole (see below for the other half).

In our design, the initial randomizing circuits U_r are constructed to generate significant amounts of entanglement entropy within the state, but *not* to saturate the entanglement bound $S^{\max}$ for larger states. Critically, this allows the clever miner room to find minimal representations of the circuit or state that will allow more efficient simulation; alternatively, one may choose to judiciously discard some information in the state as another route to feasible computation.

4 Peaking and difficulty scaling

Here we describe a finely tuned set of scaling functions that parameterize our peaked circuits in terms of a single "difficulty" value $0 \leq d \leq 5$. There are three important respects in which our brickwork peaked circuits can vary, which are

1. the number of qubits N;
2. the depths τ_{r} and τ_{p} of the randomizing and peaking circuits;
3. the extent to which the output state is peaked, defined as the ratio

$$k \equiv \frac{p_0}{p'} \tag{11}$$

of the target state s_0's probability p_0 relative to the next-highest basis state probability p'. The number of qubits N determines the dimension of the Hilbert space we work in, which ultimately sets the upper bound on how much information can be encoded (and, hence, must be represented in a faithful simulation) in the quantum state at a given point in the execution of the circuit. Said dimension is fundamentally exponential in N, so we construct a logarithmic scaling:

$$N = \lfloor 10 \log_2(d + 3.9) + 12 \rfloor \tag{12}$$

which admits a sub-exponential dependence on d in real circuit execution difficulty.

The depth τ_{r} of the randomizing circuit determines the extent to which this informational upper bound is saturated (see Section 3). The size of the quantum state increases exponentially with the entanglement entropy in the state, making it more computationally challenging to simulate accurately. However, most real-world circuits do not completely saturate the entanglement bound; hence we set

$$\tau_{\mathrm{r}} = \left\lfloor \frac{N}{2} \right\rfloor \tag{13}$$

specifically to reach around half the maximum entanglement entropy for the number of qubits. This allows room for clever miners to optimize their chosen simulation methods, rather than be forced into the hardest possible simulation regime.

Finally, we set the "peaking ratio" k using the peak probability p_0 as the tunable parameter

$$p_0 = \frac{10^{0.38d+2.102}}{2^N}. \tag{14}$$

That is, we set the absolute peaking probability to be a multiple $10^{0.38d+2.102}$ times the appropriate probability for a completely uniform probability over N qubits. In testing, we noted that, while this peaking is always attainable with $\tau_{\mathrm{r}} = \lfloor N/5 \rfloor$, the peaking procedure tends to cause non-uniform modifications to the other basis states' probabilities. Indeed, we choose this specific scaling for p_0 to keep k approximately in the 2–4 range, still ensuring that circuits are solvable, which has ramifications for sampling-based (as opposed to memory-intensive state vector-based) approaches to solution. We leave the determination of how many shots to take as a challenge to the miner.