Quantum Dark Magic
Efficiency of Intermediate Non-Stabiliserness

Tom Krueger[*]
Technical University of Applied Sciences Regensburg

Wolfgang Mauerer[†]
Technical University of Applied Sciences Regensburg and
Siemens AG, Foundational Technologies
(Dated: July 11, 2025)

Computational advantages of quantum systems over classical computers have been established in principle since decades. However, only a handful of algorithmic primitives with proven, expected or suspected quantum advantage over the best possible (or known) classical techniques are available. Despite substantial progress in limitations of de-quantisation and insights on quantum-classical separation, it remains an open challenge to systematically construct algorithms that exhibit quantum advantage. This can, at least partly, be attributed to a still incomplete understanding of the origins of quantum computational power.

Intermediate non-stabiliserness, the traversal of states outside the Clifford orbit, indicates non-classical behaviour of a quantum circuit. Therefore, non-stabiliserness is strictly necessary to reach quantum advantage. On the other hand, "the more, the better" is not a good measure for non-classical behaviour, as high non-stabiliserness is the norm for quantum states, demonstrated by the expected non-stabiliserness of a random Haar sampled state, which is with overwhelming probability close to its upper bound. To progress towards quantum advantage, it therefore seems pertinent to understand the efficient use of the non-classical, yet abundantly available resource of non-stabiliser states at the right rate and places in the quantum computation.

In this paper, we present an approach towards improving the required understanding by tracking the behaviour of non-stabiliserness across various algorithms with known of suspected quantum advantages. In particular, we pair results of resource theoretic work around non-stabiliser entropies with geometric considerations about how direct a quantum state evolution approaches the solution space. Using our techniques we are able to unveil different efficiencies in the use of non-stabiliserness comparing structured and unstructured variational state evolution. This leaves us to hypothesise that greater degrees of freedom to the classical optimisation step in such methods introduces the risk of unnecessary non-stabiliser consumption, which becomes increasingly costly when transitioning from the NISQ era into the early era of error correction.

I. INTRODUCTION

Contrary to the general discussion of quantum versus classical computing, which often treats these as separate computational models, quantum computing (QC) extends the classical computational model instead of replacing it. Quantum computations can (and for many suggested approaches also do, particularly for any variational ansatz) contain classical parts [1–3], which shifts the question of separating the two models to a more nuanced approach of identifying inherently quantum parts in computations. While possible speed-ups over purely classical approaches must obviously originate from quantum parts of a computation, not every quantum sub-computation necessarily needs to positively contribute to overall solution finding. Identifying reasons for and structure of quantum speed-ups is a crucial question to improve the understanding of chances and limitations of quantum approaches. In this paper, we address this question from a novel point of view by using geometrical distance arguments within a solution space.

Several measures for quantumness have been established; entanglement is a prime candidate that not only originates from the very beginnings of quantum mechanics [4], but has also drawn substantial interest during the last few decades [5–8]. Entanglement is a distinct, non-classical feature of quantum mechanics, and is considered one of the fundamental resources of QC [9]. However, its effect on computational power is not easy to characterise from a computer science point of view. It is generally acknowledged and understood that entanglement plays a fundamental role in many quantum algorithms and protocols. Trying to pinpoint exactly where and how such non-classical advantage is exploited necessitates more fine-grained insights. In particular, it is well known by now that not all forms of entanglement are equal (or: equally useful) [10]. Even maximally entangled states like the seminal GHZ state can be prepared by Clifford circuits; it is known that these can be efficiently simulated by a classical computer [11]. States within the orbit of the Clifford group are called stabiliser states (STAB). Conversely, states outside of STAB are referred to as non-stabiliser state (Example of non-STAB entan-

---

[*] tom.krueger@othr.de
[†] wolfgang.mauerer@othr.de

gled states include W-states with three or more qubits) [12]. Circuits required for their preparation are believed to be classically hard to simulate.

Stabiliser-Rényi-Entropies (SRE) have been recently introduced to entropically measure *non-stabiliserness*, also referred as *magic*, of quantum states [13]. In this work, we adopt SRE as measures of intermediate states to locate how and where non-classical effects appear during the execution of contemporary quantum algorithms.

The structure of the paper is as follows: In Section II, we review history and significance of non-stabiliser resource theory and measures, particular stabiliser Rényi entropies. Followed by Section III where we provide an introduction into the relevant definitions and characteristics of stabiliser Rényi entropies. We then present a geometric perspective in Section IV, and show how to calculate geodesic distances to target spaces by the means of taking the expectation value of a special problem Hamiltonian. We extend this approach to invariance of the qubit order, and show how this enables us to reveal previously easily overlooked non-stabiliser effects. After that, we put our theoretic framework to use in Section V to analyse the differences of intermediate non-stabiliser consumption in structured and unstructured state evolutions. By combining the quantum resource theoretic SRE measures with a geometric perspective, we are able to qualify the efficiency of non-stabiliser consumption. We observed a significantly higher efficiency for the structured evolution than for the unstructured case. Rounding things up, we conclude our results and discuss the potential of combining resource theoretic tools with geometric perspectives in Sections VI and VII.

Notice that a reproduction package [14] that contains all code that is required to run the experiments in this paper is available at www.TODO.de (link in PDF).

II. RELATED WORK

Given the significance of non-stabiliser effects in quantum computing, it is no surprise that understanding their properties and effects has been considered in numerous contexts. In 1997, Gottesman presented the *stabiliser formalism* in his PhD thesis, laying the foundations of quantum error correction protocols. This formalism already covers many quantum specific phenomena like GHZ entanglement. The seminal Gottesman-Knill theorem [11], presented shortly after, showed a quantum-classical separation by stating that every stabiliser circuit can be efficiently simulated by a classical computer. A stabiliser circuit is restricted to using gates from the Clifford group. An interesting conflict arises as stabiliser circuits are able to harness some quantum effects, yet an efficient simulate requires that the separation of classical and quantum complexity lies somewhere behind the obvious first line drawn between classical and non-classical physics. Colloquially speaking: There seams to be a threshold of non-classical effects needed to achieve

non-classical speedups. For Clifford circuits, universality can be recovered by a magic state injection process [16]: *Magic* states are provided on an ancilla register, letting them interfere with the Clifford part of the circuit. These magic states are non-STAB states. This turns STAB states into resources that are consumed by the magic injection procedure. Consequently, in the key to reaching a quantum advantage must be affected by magic state consumption. This motivated the the recent development of a resource theory for non-stabiliserness [17], producing a diverse set of non-stabiliserness measures including stabiliser rank [18], stabiliser fidelity [19], or stabiliser nullity [20]. Following up on these results, more abstract characterisations of measures like non-stabiliser monotones have been defined [21]. The most relevant for our work are stabiliser Rényi entropies (SRE) [13] as introduced by Leone, Oliviero, and Hamma. Stabiliser Rényi entropies are also known to be monotones for non-stabiliserness resource theory [22]. Most measures of non-stabiliserness are hard to compute. The SRE on the other hand is known to be efficiently computable for matrix product state [23]. Further, SREs can also be determined imperially through measurements [24].

III. NON-STABILISERNES

Before defining a measure for non-stabeliserness, it seems partinent to take a brief moment for discussing the term *non-stabeliserness*. We already mentioned that STAB is given by the orbit if the Clifford group (recall that the orbit of an element x in a group G is given by $G(x) := \{gx \in G : g \in G\}$). The Clifford group $\mathcal{C}_n = \{V \in \mathcal{U}_{2^n} : V\mathcal{P}_n V^\dagger = \mathcal{P}_n\}$ is the normaliser of the Pauli group $\mathcal{P}_n = \langle X, Y, Z \rangle^n$, where X, Y, Z are the Pauli operators and n denotes the number of qubits. In the following, we drop suffix n if the number of qubits can be deduced from the context, or to describe systems of arbitrary (but finite) size.

An entropy function can be defined as follows:

Definition 1 (see Ref. [13])**.**

$$SRE_\alpha(|\psi\rangle) = \frac{1}{1-\alpha} \log \sum_{P \in \mathcal{P}_n /\langle \pm i\, \mathbb{1}_n \rangle} \Xi_P^\alpha(|\psi\rangle) - \log 2^n \tag{1}$$

$$\Xi_P(|\psi\rangle) = \frac{1}{2^n} \langle\psi|P|\psi\rangle^2 \tag{2}$$

Definition 1 may require some explanation to establish an intuitive understanding. Let us start from Eq. (2). Note that $\Xi_P(|\psi\rangle) \leq 1$ and $\sum_{P \in \mathcal{P}_n /\langle \pm i\, \mathbb{1}_n \rangle} \Xi_P(|\psi\rangle) = 1$. Thus, $\{\Xi_P\}_{P \in \mathcal{P}_n /\langle \pm i\, \mathbb{1}_n \rangle}$ induces a probability distribution on a state $|\psi\rangle$. Here we also immediately see why we used the factor group of $\mathcal{P}_n$, ignoring the scalar unit factors ± 1 and $\pm i$. Due to the applied square in Eq. (2), they would only double up in the sum, not contributing any valuable information to the distribution. This also

explains the normalisation factor of 2^{-n} corresponding to $|\mathcal{P}_n / \langle \pm i \mathbb{1}_n \rangle| = 2^n$ in contrast to the perhaps expected $4^n = |\mathcal{P}_n|$. We also see, with $\{\Xi_P\}$ being a probability distribution, Eq. (1) simply defines a family of Rényi entropies offset by $\log 2^n$.

To demonstrate and provide intuition about how Stabiliser-Rényi-Entropies work, it is worth looking into their main characteristics, to understand how SREs characterise stabiliser states. For this, we will revisit the property which is most important for our work, namely that non-STAB states are characterised by a non-zero SRE. For didactic reasons and the sake of completeness we also provide a proof. The interested reader may be referred to the original paper for a more in depth discussion [13].

Theorem 1 (see [13]). *A state $|\psi\rangle$ is in STAB if and only if $SRE_\alpha(|\psi\rangle) = 0$.*

Proof. Let $|\psi\rangle \in$ STAB be some not further specified stabiliser state. Then $|\psi\rangle$ is in the Clifford orbit of $|0\rangle$, meaning that there exist a $U \in \mathcal{C}$ such that $U |0\rangle = |\psi\rangle$. Due to the Clifford group stabilising the Pauli group, we have $U^\dagger P_j U = P_j$ for $P_i, P_j \in \mathcal{P}_n / \langle \pm i \mathbb{1}_n \rangle$. In fact, $\mathcal{C}$ is isomorphic to the group of permutations in a sense that $U^\dagger \cdot U : P_j \mapsto P_{\pi(j)}$. Therefore, $\{\Xi_{P_j}(|\psi\rangle)\} = \{\Xi_{P_{\pi(j)}}(|0\rangle)\}$ and consequently $\sum_{P \in \mathcal{P}/\langle \pm i \mathbb{1} \rangle} \Xi_P^\alpha(|\psi\rangle) = \sum_{P \in \mathcal{P}/\langle \pm i \mathbb{1} \rangle} \Xi_P^\alpha(|0\rangle)$. Note that, $\langle a|Z|a\rangle|_{a=0,1} = \pm 1$, $\langle a|\mathbb{1}|a\rangle|_{a=0,1} = 1$ and $\langle a|X,Y|a\rangle|_{a=0,1} = 0$, which leads to the conclusion that

$$\langle 0|P|0\rangle = \begin{cases} 0 & \text{if} \quad \exists i : \sigma_i \in \{X, Y\} \\ 1 & \text{otherwise} \end{cases} \tag{3}$$

for all $P = \sigma_i \otimes \cdots \otimes \sigma_n \in \mathcal{P}_n / \langle \pm i \mathbb{1}_n \rangle$. There are 2^n many $P \in \mathcal{P}_n / \langle \pm i \mathbb{1}_n \rangle$ such that $\langle 0|P|0\rangle = 1$. As a result $\sum_{P \in \mathcal{P}_n / \langle \pm i \mathbb{1}_n \rangle} \Xi_P^\alpha(|\psi\rangle) = 2^n 2^{-n\alpha} = 2^{n(1-\alpha)}$. Here is where the offset of $\log(2^n)$ in Eq. (1) comes into play, as $\text{SRE}_\alpha(|\psi\rangle) = (1-\alpha)^{-1} \log 2^{n(1-\alpha)} - \log 2^n = \log 2^n - \log 2^n$. We conclude that $\text{SRE}_\alpha(|\psi\rangle) = 0$ for all $|\psi\rangle \in$ STAB.

To prove the other direction, we will use an alternative characterisation of stabiliser states, which is that $|\psi\rangle$ is in STAB if and only if there exists a subset $S \subset \mathcal{P}_n$ such that $|S| = 2^n$ and $A|\psi\rangle = |\psi\rangle$ for all $A \in S$. Now let's assume $\text{SRE}_\alpha(|\psi\rangle) = 0$ for some arbitrary state $|\psi\rangle$. Written out, that gives us $(1-\alpha)^{-1} \log \sum_{P \in \mathcal{P}/\langle \pm i \mathbb{1} \rangle} \Xi_P^\alpha(|\psi\rangle) - \log 2^n = 0$ or rewritten $\log \sum_{P \in \mathcal{P}/\langle \pm i \mathbb{1} \rangle} \Xi_P^\alpha(|\psi\rangle) = \log 2^{n(1-\alpha)}$. Thus, $\log \sum_{P \in \mathcal{P}/\langle \pm i \mathbb{1} \rangle} 2^{-n\alpha} \langle \psi|P|\psi\rangle^{2\alpha} = 2^{n-n\alpha}$. From this we can derive a condition on $|\psi\rangle$ for $\text{SRE}_\alpha(|\psi\rangle)$ to equal to 0:

$$f(\alpha) = 2^n \tag{4}$$

where $f(\alpha) = a_1(\alpha) + \cdots + a_{2^n}(\alpha)$ with $a_i = \langle \psi|P_i|\psi\rangle^{2\alpha}$ and $P_i \in \mathcal{P}_n / \langle \pm i \mathbb{1}_n \rangle$. Due to f being a constant function, we have $\frac{d}{d\alpha} f = 0$. Additionally, we know that all $a_i \geq 0$ and therefore $\frac{d}{d\alpha} \langle \psi|P|\psi\rangle^{2\alpha} = 0$

for all $P \in \mathcal{P}_n / \langle \pm i \mathbb{1}_n \rangle$. From, $\frac{d}{d\alpha} \langle \psi|P|\psi\rangle^{2\alpha} = 2\langle \psi|P|\psi\rangle^{2\alpha} \log \langle \psi|P|\psi\rangle = 0$ we conclude, that $\langle \psi|P|\psi\rangle \in \{0,1\}$ for all $P \in \mathcal{P}_n / \langle \pm i \mathbb{1}_n \rangle$. Note that $\langle \psi|P|\psi\rangle = 1$ only if $P|\psi\rangle = |\psi\rangle$ and $f(\alpha) = 2^n$. Thus, there exists a subset $S \subset \mathcal{P}_n / \langle \pm i \mathbb{1}_n \rangle$ such that $|S| = 2^n$ and $A|\psi\rangle = |\psi\rangle$ for all $A \in S$; showing that $|\psi\rangle \in$ STAB. $\quad\square$

Corollary 1. *Stabiliser-Rényi-Entropies are invariant under Clifford operations.*

Proof. This follows directly from the isomorphism between the Clifford group and permutations. Let $|\psi\rangle$ be an arbitrary state and $U \in \mathcal{C}$, then we get that $\{\Xi_{P_j}(U|\psi\rangle)\} = \{\Xi_{P_{\pi(j)}}(|\psi\rangle)\}$ and thus $\text{SRE}_\alpha(U|\psi\rangle) = \text{SRE}_\alpha(|\psi\rangle)$. $\quad\square$

IV. GEOMETRIC PERSPECTIVE

If we take a random Haar sampled state $|\psi\rangle \sim$ Haar. Then it's expected SRE is $\mathbb{E}_{|\psi\rangle \sim \text{Haar}}(\text{SRE}_\alpha(|\psi\rangle)) \in \mathcal{O}(n)$ for all $\alpha \geq 2$, with overwhelming probability [25]. Additionally, SREs are linear upper bounded by $\text{SRE}_\alpha(|\psi\rangle) \leq \log(2^n) \in \mathcal{O}(n)$ [13]. This means, although intermediate states with SRE ≥ 0 are linked with and even necessary for quantum advantages, their occurrence is nothing special and has to be expected. Consequently, this raises the question whether the observed non-stabiliserness has a contributing factor to the computation or if it is merely a byproduct of a suboptimal choice of unitary propagators. So, what does *contributing to the computation* mean? Every computation can be interpreted as a state evolution starting at a specific initial state $|\psi_0\rangle$ to a target state $|\psi_T\rangle$ encoding a problem solution or a superposition of thereof. Geometrically speaking, such a state evolution resembles a rotation of the state vector. The whole circuit represents one singular unitary, which in turn corresponds to a direct rotation from the initial to the final state around the rotational axis defined by said unitary. This only applies from the most top-level view and discard the actual realisation of the circuit's unitary given by a concrete partitioning into quantum gates and their correct sequencing. The circuits gate level realisation induces a path of the resulting state evolution, which is most likely diverging from the shortest path at some point. In geometric terms, the shortest, most direct path of this state evolution would be characterised by the geodesic from the initial state to the target. In [26] Anandan and Aharonov presented exactly this geometric perspective in conjunction with the concept of *geodesic efficiency* $\mu_{\text{gd}} = s_0/s$ of a state evolution where s_0 is the geodesic distance and s the actual distance travelled. If we want to specify the distance to a specific state $|\phi\rangle$, we write $s_0(|\phi\rangle)$ and $s(|phi\rangle)$ and if the initial state is not clear from the context we write $s(|\psi\rangle, |\phi\rangle)$ and $s_0(|\psi\rangle, |\phi\rangle)$

A. Problem Hamiltonian

Usually there is more than one unique solution to a computation problem, e.g. all binary variable assignments satisfying a propositional satisfiability problem. This adds variability to the geometric perspective discussed above. Instead of rotating our initial state to a specific target state $|\psi_T\rangle$, a quantum algorithm has the freedom to reach any state within the *target space*, which is the subspace of $\mathcal{H}^{\otimes n}$ that contains all superpositions of quantum states encoding problem solutions. Thus, the geodesic distance s_0 from above needs to be reinterpreted to be the shortest geodesic distance to one of the states in the target space. In the following, we will address this by first defining the target space based on an indicator function of problem solutions and a two-level problem Hamiltonian projecting on said target space. We then show, that the expected value of this problem Hamiltonian corresponds to the scaled inverse geodesic distance from the initial state to the target space.

Definition 2. *Let be* $c : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$ *the solution verifier of a problem with a finite set of classical solutions* $T = \{t \in \mathbb{F}_2^n : c(t) = 1\}$. *We then linear extent* $c(\cdot)$ *to quantum states* $|\psi\rangle = \sum_{b \in \mathbb{F}_2^n} \alpha_b |b\rangle$:

$$c(|\psi\rangle) = \sum_{b \in \mathbb{F}_2^n} |\alpha_b|^2 c(b) \tag{5}$$

Further, we define a quantum target space

$$\mathcal{T} = \{|t\rangle : c(|t\rangle) = 1\} \subset \mathcal{H}^{\otimes n} \tag{6}$$

Remark 1. *Note that* $\mathcal{T}$ *is indeed a complete subspace of* $\mathcal{H}^{\otimes n}$, *spanned by* $\{|b\rangle : b \in \mathbb{F}_2^n, c(b) = 1\}$. *Thus,* $\mathcal{T}$ *has a dimension of* $|T|$.

Definition 3. *Based on* $c(|\psi\rangle)$, *we define a 2-level problem Hamiltonian* H_c *by the condition that*

$$\langle H_c \rangle = c(|\psi\rangle) \tag{7}$$

We quickly see that H_c is a projector onto $\mathcal{T}$ and can explicitly defined by $H_c = \sum_{t \in T} |t\rangle\langle t|$.

Theorem 2. *Given a target space* $\mathcal{T}$ *and the corresponding problem Hamiltonian* H_c *according to Definitions 2 and 3, we have*

$$s_0(\mathcal{T}) := \min_{|t\rangle \in \mathcal{T}} s_0(|t\rangle) = 2 \arccos\langle H_c \rangle \tag{8}$$

Proof. Let $\mathcal{B}_\mathcal{T} = \{|t\rangle : t \in T\}$ the basis of $\mathcal{T}$. We then expand $\mathcal{B}_\mathcal{T}$ with $\mathcal{B}_{\overline{\mathcal{T}}} = \{|b\rangle : b \in \mathbb{F}_2^n \setminus T\}$ such that $\mathcal{B}_\mathcal{T} \cup \mathcal{B}_{\overline{\mathcal{T}}}$ forms a basis of $\mathcal{H}^{\otimes n}$. Given that expanded basis, we can write every state $|\psi\rangle \in \mathcal{H}^{\otimes n}$ as $|\psi\rangle = \sum_{i=1}^{|T|} \tau_i |t_i\rangle + \sum_{i=1}^{n-|T|} \beta_i |b_i\rangle$, with $\sum_{i=1}^{|T|} |\tau_i|^2 + \sum_{i=1}^{n-|T|} |\beta_i|^2 = 1$ and for all $|t\rangle \in \mathcal{T}$ we have $\sum_{i=1}^{|\mathcal{T}|} |\tau_i|^2 = 1$. Now, let H_c be

a problem Hamiltonian as defined in Definition 3, then $H_c = \sum_{|t\rangle \in \mathcal{B}_\mathcal{T}} |t\rangle\langle t|$ and

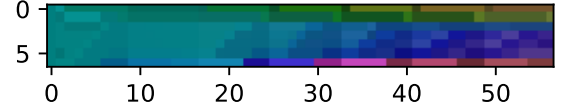$$0 \leq \langle H_c \rangle = \sum_{i=1}^{|T|} |\tau_i|^2 \leq 1$$

Now, let's take an arbitrary state $|\psi\rangle$, then $\langle H_c \rangle = \langle\psi|H_c|\psi\rangle$ is exactly the overlap between $|\psi\rangle$ and its projection onto the target space $H_c |\psi\rangle$, which satisfies
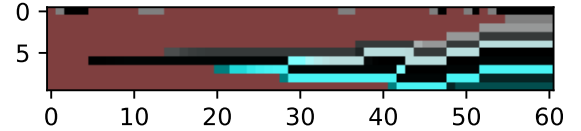
$$\langle H_c \rangle = \max_{|t\rangle \in \mathcal{T}} |\langle\psi|t\rangle|$$

Now we use that $s_0(|t\rangle) = 2 \arccos|\langle\psi|t\rangle|$ [27]. Now due to the monotonicity of arccos in $[0,1]$ we can pull out the max from $2\arccos\langle H_c \rangle$ to end up at Eq. (8). $\square$

B. Permutations

(a) QAOA (7 qubits, 7 layers)

(b) QFT (10 qubits)

FIG. 1. Evolution of the colour representation of the state in quantum circuits. Every vertical slice at $x = i$ represents the colour spectrum of the state after the ith gate. The reduced one qubit density matrices are mapped to a hue-saturation-value colour with hsv($\langle P_0\rangle$, $\langle P_+\rangle$, $\langle P_{+i}\rangle$). Within a vertical slice, they are sorted according to the hsv tupel.

In a typical quantum circuit, qubits are sequentially numbered. This numbering implies an unsubstantiated sense of locality or neighbourhood. Indeed, it is actually completely arbitrary and nothing more of a naming convention. Qubit q_i and q_j could also be remapped $q_{\sigma(i)}$ and $q_{\sigma(j)}$ for some permutation $\sigma \in S_n$. The actual relevance of locally comes from the concrete problem instance [28], which imposes relations between variables, which in turn are mapped to qubits. Those connections could better be represented as a graph, which in turn is isomorphic under vertex permutations. On this graph, we assign to each vertex a colour based on properties of its linked qubit. By keeping track of the necessary permutations, we can, at each intermediate time, determine the permutation order of qubits based on their assigned colour.

One could thus visualise a quantum state evolution as the change of a colour spectrum through time. Solving a classical problem, we are basically interested in the measurement probabilities of all qubits and the resulting bit-string, hopefully encoding a possible solution to the problem. We therefore exemplary map $\langle P_0 \rangle$, $\langle P_+ \rangle$ and $\langle P_{+i} \rangle$ to the hue, saturation and value component of an HSV colour. Here $\langle P_0 \rangle$, $\langle P_+ \rangle$ and $\langle P_{+i} \rangle$ are the probabilities of the reduced mixed system of said qubit being in the state $|0\rangle$, $|+\rangle$ and $|+i\rangle$. Now the, qubits can be ordered according to their hue. Figure 1 demonstrates how this representation, which is qubit permutation invariant, still reveals highly specific structures of quantum state evolutions. As we will show below, the ordering does not alter non-stabiliserness, as it can be performed by an efficient Clifford circuit. Therefore, it can be ignored regarding our analysis of non-stabiliserness resource consumption. Additionally, introducing, at the worst case, one ordering and reordering before and after each computational step does also not change the complexity theoretic characterisation of the circuit, as it, given the presumption of a polynomial sized initial circuit, only adds a polynomial amount of permutation circuits which themselves also only have a polynomial complexity. Therefore, questions regarding the link between non-stabiliser consumption and quantum advantages can be investigated with frameworks factoring out permutational degrees of freedom.

We now will bring the intuition of a shift on the colour spectrum to a concrete mathematical representation. After that we have to formulate a permutation robust version of our geometric measure. As a first step, we define a permutation operator capturing the notions discussed above.

Definition 4. *Given a permutation $\sigma \in S_n$ and $|b\rangle \in \mathcal{B}^n$ where $|b\rangle = |b_1\rangle \otimes \cdots \otimes |b_n\rangle$ with $|b_i\rangle \in \{|0\rangle, |1\rangle\}$, then*

$$\hat{\sigma} |b\rangle = \left|b_{\sigma(1)}\right\rangle \otimes \cdots \otimes \left|b_{\sigma(n)}\right\rangle \qquad (9)$$

and

$$\hat{\sigma} \sum_{b \in \mathcal{B}^n} \alpha_b |b\rangle = \sum_{b \in \mathcal{B}^n} \alpha_b \hat{\sigma} |b\rangle . \qquad (10)$$

Note that in Definition 4 the inverse operator $\hat{\sigma}^\dagger \in \mathcal{H}^{\otimes n}$ corresponds to the inverse permutation $\sigma^{-1} \in S_n$. Next, we have to show that the SRE measure is invariant under such permutation operators.

Theorem 3. *Let $\hat{\sigma}$ be a permutation operator as defined in Definition 4, then*

$$SRE_\alpha(\hat{\sigma} |\psi\rangle) = SRE_\alpha(|\psi\rangle)$$

Proof. Every permutation $\sigma \in S_n$ can be decomposed into a sequence of 2-cycles, which can be realised by a single swap gate. Thus, $\hat{\sigma} \in \mathcal{H}^{\otimes n}$ can be realised by a sequence of swap gates, which are Clifford operations. Since, SRE_α is invariant under Clifford operations, it also is for all permutation operators constructed as defined in Definition 4. $\square$

Now, after we have formalised the idea of invariance under permutation on the operational side, we will do the same for the objects of interest. We do this by subsuming all states equal under permutation into equivalence classes and then extend this to the target space itself.

Definition 5. *We define an equivalence relation $|\psi_l\rangle \sim |\psi_r\rangle$ which is satisfied iff there exist a permutation operator $\hat{\sigma}$ as defined in Definition 4, such that $|\psi_r\rangle = \hat{\sigma}|\psi_l\rangle$. Then*

$$[|\psi\rangle] = \{\hat{\sigma}|\psi\rangle : \forall \hat{\sigma}\} \qquad (11)$$

is the corresponding equivalence class of $|\psi\rangle$ under $\sim$ and further $SRE([|\psi\rangle]) = SRE(|\psi\rangle)$. Let $\mathcal{T}$ be a subspace of $\mathcal{H}^{\otimes n}$, we then extend this notion by defining

$$[\mathcal{T}] = \bigcup_{|t\rangle \in \mathcal{T}} [|t\rangle] \qquad (12)$$

From Theorem 3 it also immediately follows that $SRE_\alpha(|\psi\rangle) = SRE_\alpha([|\psi\rangle])$. For the geodesic distance, we need to extend the definition to equivalence classes.

Definition 6. *Let $[|\phi\rangle]$ be a equivalence class of states then*

$$s_0([|\phi\rangle]) = \min_{|\phi'\rangle \in [|\phi\rangle]} s_0(|\phi'\rangle) \qquad (13)$$

For $[\mathcal{T}]$ we extend s_0 in a similar fashion to

$$s_0([\mathcal{T}]) = \min_{|t\rangle \in \mathcal{T}} s_0([|t\rangle]) \qquad (14)$$

Determining the distance $[\mathcal{T}]$ requires tracing all permutations of all possible solution states, which can be a bit tricky. Lucky, we can show that the distance from $|\psi\rangle$ to $\mathcal{T}$ is equal to the distance from $[|\psi\rangle]$ to $\mathcal{T}$.

Theorem 4. *Given a target space permutation equivalence class $[\mathcal{T}]$ as defined in Definition 5, it holds that*

$$s_0([\mathcal{T}]) = s_0(|\psi\rangle, [\mathcal{T}]) = s_0([|\psi\rangle], \mathcal{T}) \qquad (15)$$

Proof. By definition, we have that $s_0([|\psi\rangle], \mathcal{T}) = \min_{|\psi'\rangle \in [|\psi\rangle]} s_0(|\psi'\rangle, \mathcal{T})$ which equals $\min_{\hat{\sigma}} s_0(\hat{\sigma}|\psi\rangle, \mathcal{T}) = \min_{\hat{\sigma}} 2 \arccos \left\langle \psi \left| \hat{\sigma}^\dagger H_c \hat{\sigma} \right| \psi \right\rangle$. As we are minimising over the whole group of all permutation operators we can also minimise over all complex conjugate operators instead $\min_{\hat{\sigma}^\dagger} 2 \arccos \left\langle \psi \left| \hat{\sigma} H_c \hat{\sigma}^\dagger \right| \psi \right\rangle$. By the canonical definition of H_c we have $\hat{\sigma} H_c \hat{\sigma}^\dagger = \sum_{t \in T} \hat{\sigma}|t\rangle\langle t| \hat{\sigma}^\dagger$. Recall that $\{|t\rangle : t \in T\}$ is the basis of the corresponding quantum target space $\mathcal{T}$. This means, by applying $\hat{\sigma} H_c \hat{\sigma}^\dagger$ we are performing a basis transformation on the target space, measuring the expected probability of $|\psi\rangle$ being in he permuted target space. By minimising over all permutations we get $s_0(|\psi\rangle, [\mathcal{T}])$, thus in conclusion $s_0(|\psi\rangle, [\mathcal{T}]) = s_0([|\psi\rangle], \mathcal{T})$. $\square$
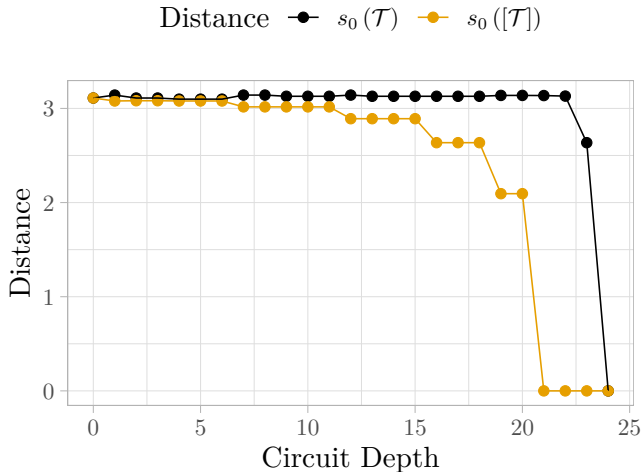
FIG. 2. Minimal geodesic distance for increasing circuit depths. Non-Clifford computational progress can be observed prior to the final qubit order reversal when all target space permutations $s_0([\mathcal{T}])$ (ochre) are considered. As such effects are not visible in the direct distance to the target space $s_0(\mathcal{T})$ (black) that neglects permutations, this demonstrates how potential non-stabiliser effects can be masked by non-Clifford-agnostic measures (lines are used to guide the eye and have no significance).
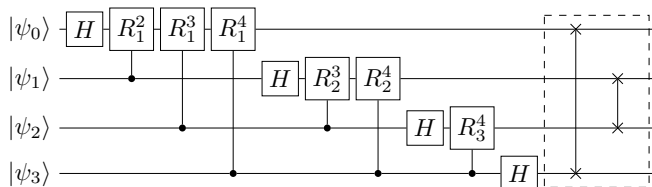

FIG. 3. QFT circuit with four qubits. The dashed box marks the qubit order inversion block of swap gates. Non-stabiliser computations take place before this block, but their computational influence on the geodesic distance is masked by the final qubit reordering.

The quantum Fourier transform (QFT) is a good example to demonstrate this effect, and additionally shows how to calculate the distance of the closest target space permutation. As of today, the QFT is regarded as *the* seminal primitive contributing to quantum advantage, finding application in a wide range of quantum algorithms reaching proven quantum speedups such as the famous integer factorisation by Shor [29] . The interesting parts of the QFT circuit take part before the qubits are reordered in a final step (marked section in Fig. 3). This is problematic when looking at distance measures based on state to space overlaps like the geodesic distance. The block of swap gates implementing the reordering is entirely Clifford, yet looking at the geodesic instance $s_0(\mathcal{T})$ one could be under the impression that all the computational progress takes place in this section of the circuit.

This, can not be the case, as we know that the QFT algorithm provides an exponential speedup. Therefore valuable computational progress has to be made before, taking possible target space permutations into account reveals such effects (see Fig. 2).

V. EXPERIMENTS

General state evolution algorithms usually are quite high level from an algorithmic standpoint. The logical structure of problem instances usually is encoded in a Hamiltonian either driving the sate evolution like in quantum annealing and its gate based counterparts (e.g., QAOA) or serving as a cost function representation expressing the solution quality, which then can be used to optimise free parameters of a quantum circuit. In both cases, the problem structure is quite removed from the description of the algorithmic dynamics. This divide between descriptive dynamics and problem structures introduces a high level of abstraction masking the actual dynamics.

State evolution can be broadly grouped into *structured* and *unstructured* techniques. The former introduce little restrictions on the circuit logic, and leave more freedom to the optimisation step. The latter directly impose the problem structure onto the circuit, which significantly reduces the number of free parameters.

A. Problem Description

We now want to showcase our methods introduced above to reveal actual differences in the evolution of structured and unstructured state evolution techniques. As an exemplary problem, we chose the seminal NP complete problem of boolean satisfiability (SAT), more precisely the problem of finding a satisfying variable assignment of a 3-CNF boolean formula $F : \mathbb{F}_2^n \to \mathbb{F}_2$. Let's define a problem Hamiltonian satisfying Definition 3. We start by defining the classical solution space T where $t = t_1 t_2 \cdots t_n \in T$ iff $c(t) := F(t) = 1$. Then target space shall be defined as $\mathcal{T} = \{\bigotimes_{i=1}^n |t_i\rangle : t_1 t_2 \cdots t_n \in T\}$. Note that F is a 3-CNF boolean formula, therefore $F = \prod_{i=1}^m f_i$ with $f_i : \mathbb{F}_2^3 \to \mathbb{F}_2$ are disjunctions. This means, every f_i has one unique unsatisfying assignment $\overline{t_i}$. Now it is easy to see that the Hamiltonian $H_c := \mathbb{1} - \prod_{i=1}^m |t_i\rangle\langle t_i|$ satisfies Eq. (7). See Ref. [30] for more details.

B. Unstructured State Evolution

In an unstructured state evolution ansatz a a generic circuit template leaving maximal flexibility to be adjusted later in an optimisation step minimising a cost function, which is minimal if the final state is in the target space $\mathcal{T}$. The initial circuit ansatz is the same for
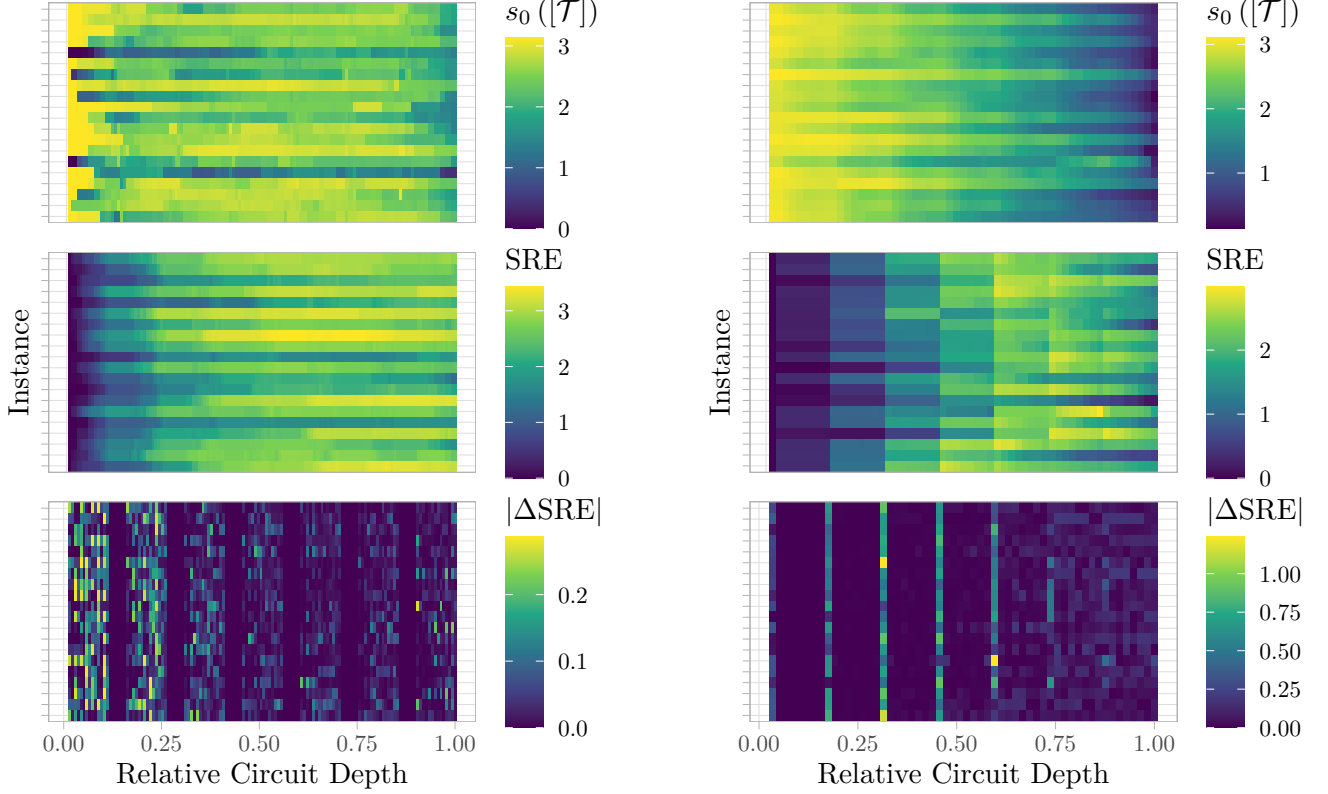
(a) Properties of intermediate states (x-axis) evolving under an **unstructured** ansatz for different instances (y-axis). *Top:* The state approaches the target space with notable erraticity, as evidenced by irregular fluctuations in the geodesic distance from $[\mathcal{T}]$. While non-stabiliserness (*middle*) appears to evolve comparatively smooth, the variation in resource consumption, as quantified by $|\Delta\text{SRE}|$ (*bottom*), continues to display a markedly irregular behaviour.

(b) Properties of intermediate states (x-axis) evolving under a **structured** QAOA ansatz for different instances (y-axis). *Top:* The state approaches the target space $[\mathcal{T}]$ smoothly. The accumulation of non-stabiliserness (*middle*) follows a structured trajectory and reaches its apex at about 75% relative circuit depth, after which it gradually diminishes. This behaviour is mirrored in the patterns of non-stabiliser resource consumption (*bottom*).

FIG. 4. Comparison of intermediate geodesic distances $s_0([\mathcal{T}])$, non-stabiliserness SRE and non-stabiliser consumption $|\Delta\,\text{SRE}|$ between unstructured (Fig. 4a) and structured (Fig. 4b) state evolution.

all problems and problem instances. Concrete instance or problem specific structure only gets introduced during the cost function optimisation process. As an exemplary ansatz we investigated a hardware efficient variational quantum eigensolver. We chose a layered architecture where one layer exists of a stack of $R_y(\theta_i^y)$ gates applied to each qubit i followed by a similar stack of $R_z(\theta_i^z)$ gates and a ladder of cnot gates to provide entanglement. For a full circuit for the layer structure see Fig. 5.

C. Structured State Evolution

In contrast to unstructured state evolution techniques, in the structured case the ansatz already gets infused with instance structures. One can show that problem structures extrapolated from common instance structures are sufficient to successfully approximate the state evolu-
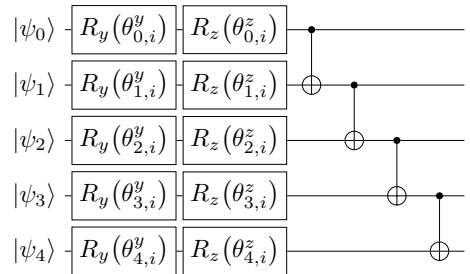

FIG. 5. The i-th layer of the hardware efficient ansatz used for unstructured state evolution.

tion of such methods [30]. This shows, that the structural infusion significantly impacts the ansatz even before instance specific cost function optimisation techniques are applied. As a representative for structured state evolution, we chose a standard QAOA ansatz where the driv-

ing problem Hamiltonian is the problem Hamiltonian H_c defined above. This equals the construction presented in [30].

D.   Experimental Setup

For each ansatz we solved 20 SAT instances with the circuits spanning $n = 7$ qubits and $p = 7$ layers. Every instance was randomly sampled with a clause to variable ratio of $|C|/|V| = 3$, which generates SAT instances that are constrained enough to be at the start of the easy to hard phase transition. At the same time those instances are still not too hard to solve such that we can expect the state evolutions to get fairly close to the target space, assuring that we witness a state space traversal travelling a significant part of the distance necessary to successfully solve the problem.

E.   Results and Comparison

Comparing the state evolution of structured and unstructured circuits we notice that the former approaches the target space $[\mathcal{T}]$ in a direct path, smoothly reducing the geodedic distance with each step. In contrast, the unstructured evolution seems to erratically jump through the state space, witnessed by jumps in the geodesic distance $s_0([\mathcal{T}])$ while passing through the circuit. The differences become apparent when comparing the top plots of Figs. 4a and 4b. We now further analyse how both state evolutions apprached the target space on a step by step basis. For this, we calculate the delta of s_0 before and after each step. Fig. 6 shows that the distribution of $\Delta s_0([\mathcal{T}])$ is symmetrically centred around zero, ignoring a few outliers of big jumps of negative Δs_0. For the structured evolution, on the other hand, we observe that the distribution of Δs_0 values is skewed towards the regime below zero, and the majority of values is less than zero (Table I shows more detailed numbers). This indicates that speaking on a per-step basis the structured ansatz more efficiently approaches the target. For the unstructured ansatz the majority of steps seem to move towards or away from the target with equal probability, de-facto cancelling each other out on the macroscopic level. That being said, negative Δs_0 outliers of bigger value seem to suggest that the unstructured approach is able to reach further in larger individual steps, reaching the target faster if utilised efficiently.

Another important aspect of efficiency is the consumption of non-stabiliserness. As already mentioned above SRE is invariant under Clifford gates. In conclusion, a change in the SRE of the intermediate state being evolved indicates the use of a non-Clifford operation. Therefore, we will use the absolute step SRE difference $|\Delta \mathrm{SRE}|$ as an indicator of non-stabiliserness consumption, which is inherently linked to costly operations. Comparing Figs. 4a and 4b (bottom), one sees that, similar to the

Ansatz	Q1	Q2	Q3	$\Delta s_0 < 0$	$\Delta s_0 > 0$
Structured	-0.0792	-0.0377	0.0000	76.7%	16.6%
Unstructured	-0.0021	0.0000	0.0010	32.3%	33.7%

TABLE I. 25% (Q1), 50% (Q2), and 75% (Q3) quartiles of the $\Delta s_0([\mathcal{T}])$ distributions for structured and unstructured state evolution. The last two columns depict fractions of steps that decrease ($\Delta s_0 < 0$) or increasing ($\Delta s_0 > 0$) target distance. On the whole, the distribution of structured Δs_0 exhibits a pronounced skewness towards negative values, in marked contrast to the unstructured scenario, where it is mostly symmetrically centred about zero.

geodesic distance $s_0([\mathcal{T}])$ (top) the unstructured ansatz is also more erratic than its structured counterpart, when it comes to non-stabiliserness consumption. This begs the question whether there is a connection between both observations. In fact, there is a positive correlation between step-wise geodesic distance reductions to the target space and non-stabiliserness consumption for the structured state evolution. In contrast to that observation, there is no such correlation for the unstructured case. This further substantiates the hypothesis that the structured ansatz utilises non-stabiliser resources more efficiently.
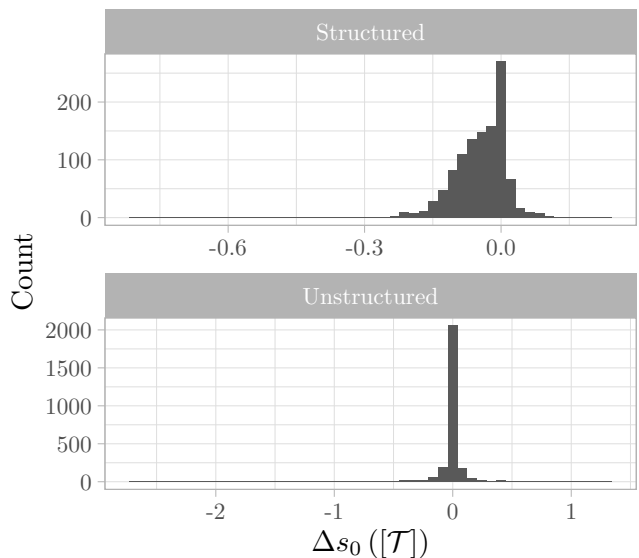

FIG. 6.   Distribution of increments and decrements in distance to the target space ($x \sim \pm\Delta s_0([\mathcal{T}])$). *Top:* Heavy skew toward decrements is observed for the structured ansatz, and most distance changes are negative. *Bottom:* The distribution of distance changes is approximately centred around zero for the unstructured ansatz, discounting a small number of outliers on the negative side.
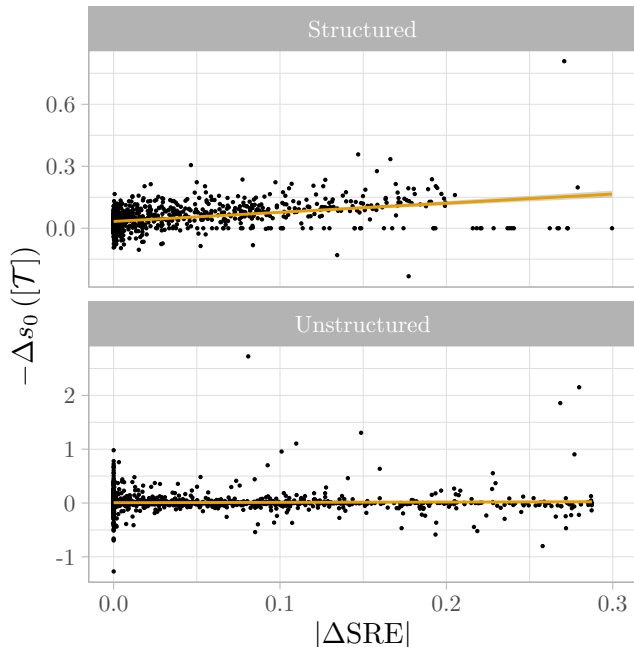
FIG. 7. Structured state evolution (top) shows clear correlation between non-stabiliser consumption $|\Delta \mathrm{SRE}|$ and steps reducing the geodesic distance to the target space $-\Delta s_0([\mathcal{T}])$. In contrast, we cannot observe a similar correlation for the unstructured ansatz (bottom). Data are restricted to $|\Delta \mathrm{SRE}| < 0.3$ to filter outliers (this retains approximately 98% of the original data points).

VI. CONCLUSION

We have extended the concept of geodesic distance measures in state evolutions targeting a specific state to evolutions targeting a more complex target space $\mathcal{T}$. We showed how the geodesic distance to $\mathcal{T}$ can be derived from the expected value of a Hamiltonian satisfying Eq. (7). This Hamiltonian based definition fits well into widely used frameworks of Hamiltonian cost function encodings. It also further allows for establishing empirical measurement based setups that integrating nicely with existing toolkits of quantum computing practitioners.

We further provided a qubit order agnostic version of the geodesic framework by introducing equivalence classes of states that are equal under permutation. Considering the quantum Fourier transform as a use-case, we demonstrated how our approach can cut through Clifford layers, and thus unveil previously hidden computational progress in the circuit. We then applied the developed methods to comparatively analyse of structured and unstructured state evolutions. The different distributions of geodesic distance changes suggest a higher geodesic efficiency for the structured evolution.

By combining resource theoretic Stabiliser-Rényi-Entropy and geometric geodesic distance measures, we

where able to show that the structured ansatz is significantly more efficient in the consumption of non-stabiliser resources than the unstructured ansatz. On a methodical level, this demonstrated the potential of our combination of methodologies.

VII. DISCUSSION & PERSPECTIVE

We believe our methodology opens new means of analysing non-stabiliser effects and the efficient utilisation of non-stabiliser resources in quantum circuits. A nuanced understanding of such effects seems crucial for advancing the systematic development of quantum algorithms, particularly with regard to realising quantum speed-ups in a well-principled manner. Furthermore, we anticipate that our results will become increasingly pertinent as the field transitions into the era of early fault-tolerant quantum computing: In such regimes, non-stabiliser operations pose significantly greater challenges for error correction compared to stabiliser operations. Consequently, the use of this resource must be optimised, and we believe that our analytical framework offers a valuable instrument in progressing towards this objective.

We showed how permutation agnostic distance measures can reveal internal non-stabiliser effects previously hidden by a subset of Clifford operations. Our construction based on permutation operators $\hat{\sigma}$ could be extended to accept general Clifford operators in the sense that two states $|\psi_1\rangle \sim |\psi_2\rangle$ iff there exists a $U \in \mathcal{C}$ such that $U|\psi_1\rangle = |\psi_2\rangle$. Even though Clifford circuits can be efficiently simulated by classical systems, is is not necessarily possible to construct them efficiently. Thus, such an extension would need to impose some complexity theoretic bounds on U to avoid grouping states that can only be reached by overly powerful oracles.

We showed that the combination of resource theoretic and geometric tools offers a mean to qualify resource consumptions by efficiency. We see a potential to embed quantum resource theoretic measures like Stabiliser-Rényi-Entropies into a proper differential geometric framework. This is a second promising avenue for improvement that would allow us to analyse resources consumed by state evolutions following different paths over the projective state manifold.

ACKNOWLEDGMENTS

[1] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational quantum algorithms, Nature Reviews Physics **3**, 625–644 (2021).

[2] S. Thelen, H. Safi, and W. Mauerer, Approximating under the influence of quantum noise and compute power, in *Proceedings of WIHPQC@IEEE QCE* (2024).

[3] M. Franz, T. Winker, S. Groppe, and W. Mauerer, Hype or heuristic? quantum reinforcement learning for join order optimisation, in *Proceedings of the IEEE International Conference on Quantum Computing and Engineering* (2024).

[4] R. Jozsa, Entanglement and quantum computation, arXiv preprint quant-ph/9707034 (1997).

[5] J. Preskill, Quantum computing and the entanglement frontier, arXiv preprint arXiv:1203.5813 (2012).

[6] R. Jozsa and N. Linden, On the role of entanglement in quantum-computational speed-up, Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences **459**, 2011 (2003).

[7] A. Datta and G. Vidal, Role of entanglement and correlations in mixed-state quantum computation, Physical Review A—Atomic, Molecular, and Optical Physics **75**, 042310 (2007).

[8] W. Mauerer, *On colours, keys, and correlations: multimode parametric downconversion in the photon number basis*, Ph.D. thesis, Erlangen, Nürnberg, Univ., Diss. (2009).

[9] A. Khrennikov, Roots of quantum computing supremacy: superposition, entanglement, or complementarity?, The European Physical Journal Special Topics **230**, 1053 (2021).

[10] J. Odavić, M. Viscardi, and A. Hamma, Stabilizer entropy in non-integrable quantum evolutions, arXiv preprint arXiv:2412.10228 (2024).

[11] D. Gottesman, The Heisenberg representation of quantum computers, in *22nd International Colloquium on Group Theoretical Methods in Physics* (1998) pp. 32–43, arXiv:quant-ph/9807006.

[12] H. J. García, I. L. Markov, and A. W. Cross, On the geometry of stabilizer states, Quantum Info. Comput. **14**, 683–720 (2014).

[13] L. Leone, S. F. Oliviero, and A. Hamma, Stabilizer rényi entropy, Physical Review Letters **128**, 10.1103/physrevlett.128.050402 (2022).

[14] W. Mauerer and S. Scherzinger, 1-2-3 reproducibility for quantum software experiments, in *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)* (2022).

[15] D. Gottesman, *Stabilizer codes and quantum error correction* (California Institute of Technology, 1997).

[16] S. Bravyi and A. Kitaev, Universal quantum computation with ideal clifford gates and noisy ancillas, Physical Review A **71**, 10.1103/physreva.71.022316 (2005).

[17] E. Chitambar and G. Gour, Quantum resource theories, Reviews of Modern Physics **91**, 10.1103/revmodphys.91.025001 (2019).

[18] S. Bravyi, G. Smith, and J. A. Smolin, Trading classical and quantum computational resources, Physical Review X **6**, 10.1103/physrevx.6.021043 (2016).

[19] S. Bravyi, D. Browne, P. Calpin, E. Campbell, D. Gosset, and M. Howard, Simulation of quantum circuits by low-rank stabilizer decompositions, Quantum **3**, 181 (2019).

[20] M. Beverland, E. Campbell, M. Howard, and V. Kliuchnikov, Lower bounds on the non-clifford resources for quantum computations, Quantum Science and Technology **5**, 035009 (2020).

[21] T. Haug and L. Piroli, Stabilizer entropies and nonstabilizerness monotones, Quantum **7**, 1092 (2023).

[22] L. Leone and L. Bittel, Stabilizer entropies are monotones for magic-state resource theory, Physical Review A **110**, 10.1103/physreva.110.l040403 (2024).

[23] S. F. E. Oliviero, L. Leone, and A. Hamma, Magic-state resource theory for the ground state of the transverse-field ising model, Physical Review A **106**, 10.1103/physreva.106.042426 (2022).

[24] S. F. E. Oliviero, L. Leone, A. Hamma, and S. Lloyd, Measuring magic on a quantum processor, npj Quantum Information **8**, 10.1038/s41534-022-00666-5 (2022).

[25] A. Gu, L. Leone, S. Ghosh, J. Eisert, S. F. Yelin, and Y. Quek, Pseudomagic quantum states, Physical Review Letters **132**, 10.1103/physrevlett.132.210602 (2024).

[26] J. Anandan and Y. Aharonov, Geometry of quantum evolution, Physical Review Letters **65**, 1697–1700 (1990).

[27] C. Cafaro, E. Clements, and A. Alanazi, Aspects of complexity in quantum evolutions on the bloch sphere (2025).

[28] An additional establishment of locality and neighbourhood comes from restrictions imposed by the coupling graph of the concrete quantum processor executing the circuit. This consideration is beyond the more abstract arguments presented in this paper.

[29] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, SIAM Journal on Computing **26**, 1484–1509 (1997).

[30] T. Krüger and W. Mauerer, Out of the loop: Structural approximation of optimisation landscapes and non-iterative quantum optimisation (2024).