Geometric and Resource-Theoretic Characterisation of Non-Stabiliserness in Quantum Algorithms

Tom Krueger*
Technical University of Applied Sciences Regensburg and FI CODE, Universität der Bundeswehr München

Wolfgang Mauerer†
Technical University of Applied Sciences Regensburg and Siemens AG, Foundational Technologies

(Dated: May 12, 2026)

While there is strong evidence for advantages of quantum over classical computation, the repertoire of computational primitives with proven or conjectured quantum advantage remains limited. A big challenge of quantum algorithmic design is a still incomplete understanding of the sources of quantum computational power. Advancing towards systematic quantum advantage calls for a better understanding of the efficient use of non-classical resources like non-stabiliser states.

We present an approach to track non-classical contributions in the form of non-stabiliserness across various algorithms by pairing resource theory of non-stabiliser entropies with the geometry of quantum state evolution, and introduce permutation agnostic distance measures that reveal and quantify non-stabiliser effects previously hidden by a subset of Clifford operations. We find different efficiency in the use of non-stabiliserness for structured and unstructured variational approaches, and show that greater freedom for classical optimisation in quantum-classical methods increases unnecessary non-stabiliser consumption. Our results open new means of analysing the efficient utilisation of quantum resources, and contribute towards the targeted construction of algorithmic quantum advantage.

I. INTRODUCTION

Contrary to the general discussion of quantum versus classical computing that often treats these as separate computational models, quantum computing (QC) can also be seen as an extension to the classical computational model that adds new primitives and resources. 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 in proposed algorithms 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 computation. 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; the amount of entanglement that manifests in computations is a prime candidate. It has been studied extensively from the early times of quantum computing [4] into the present, drawing substantial interest during the last few decades [5–8]. Entanglement is a distinct, perhaps the most non-classical aspect 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, at least when using the appropriate amount [10–14]. Trying to pinpoint exactly where and how such non-classical advantage is exploited necessitates more fine-grained insights. In particular, it is well known that not all forms of entanglement are equal (or: equally useful) [15]. 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 [16]. 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 (examples of non-STAB entangled states include W-states with three or more qubits) [17]. Circuits required for their preparation are believed to be classically hard to simulate [16, 18–20].

Stabiliser-Rényi-Entropies (SRE) have been recently introduced to entropically measure *non-stabiliserness*, also referred as *magic*, of quantum states [21]. We adopt SRE as measure 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

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

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 then provide intuition for qubit permutation invariant structures throughout the state evolution, from which we derive an extension of our tools to work under invariance of the qubit order, and show how this enables us to reveal previously 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. We conclude and discuss the potential of combining resource theoretic tools with geometric perspectives in Sections VI and VII.

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 [18], 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 [16], 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 even stabiliser circuits are able to harness some quantum effects, yet they can be efficiently simulated. This means the threshold to achieve quantum speedups must be somewhere behind the obvious first line drawn between classical and non-classical physics. For Clifford circuits, universality can be recovered by aninjection process [19] where *magic* states—non-stabiliser ancillas—serve as consumable resources restoring universality by interfering with the Clifford part of the circuit. Consequently, reaching a quantum advantage must be affected by magic state consumption. This motivated the the recent development of a resource theory for non-stabiliserness [22], producing a diverse set of non-stabiliserness measures including stabiliser rank [23], stabiliser fidelity [24], or stabiliser nullity [25]. Following up on these results, more abstract characterisations of measures like non-stabiliser monotones have been defined [26]. The most relevant for our work are stabiliser Rényi entropies (SRE) as introduced in [21]. Stabiliser Rényi entropies are also known to be monotones for non-stabiliserness resource theory [27]. While most measures of non-stabiliserness are hard to compute, SRE can be efficiently determined for low entanglement systes can be represented as matrix product states [28]. Further, SREs can also be determined empirically through measurements [29].

III. NON-STABILISERNESS

Before defining a measure for non-stabiliserness, it seems pertinent to take a brief moment for discussing the term *non-stabiliserness*. We already mentioned that STAB is given by the orbit of 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 $\mathcal{C}_n = \left\{V \in \mathcal{U}_{2^n} : V\,\mathcal{P}_\mathrm{n}\,V^\dagger = \mathcal{P}_\mathrm{n}\right\}$ is the normaliser of the Pauli group $\mathcal{P}_\mathrm{n} = \langle X, Y, Z\rangle^n$, 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.

Definition 1 (see Ref. [21]). *The α-Stabilizer-Rényi-Entropy is defined as:*

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

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

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 $\sum_{P\in\mathcal{P}_\mathrm{n}/\langle\pm\mathrm{i}\,\mathbb{1}_\mathrm{n}\rangle}\Xi_P(|\psi\rangle) = 1$. Thus, $\{\Xi_P\}_{P\in\mathcal{P}_\mathrm{n}/\langle\pm\mathrm{i}\,\mathbb{1}_\mathrm{n}\rangle}$ induces a probability distribution on a state $|\psi\rangle$. Here we also immediately see why we used the factor group of $\mathcal{P}_\mathrm{n}$, 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 $|\mathcal{P}_\mathrm{n}/\langle\pm\mathrm{i}\,\mathbb{1}_\mathrm{n}\rangle| = 2^n$ in contrast to the perhaps expected $4^n = |\mathcal{P}_\mathrm{n}|$. 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.

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 $P_i, P_j \in \mathcal{P}_\mathrm{n}/\langle\pm\mathrm{i}\,\mathbb{1}_\mathrm{n}\rangle$. In fact, $\mathcal{C}$ is isomorphic to the group of permutations in a sense

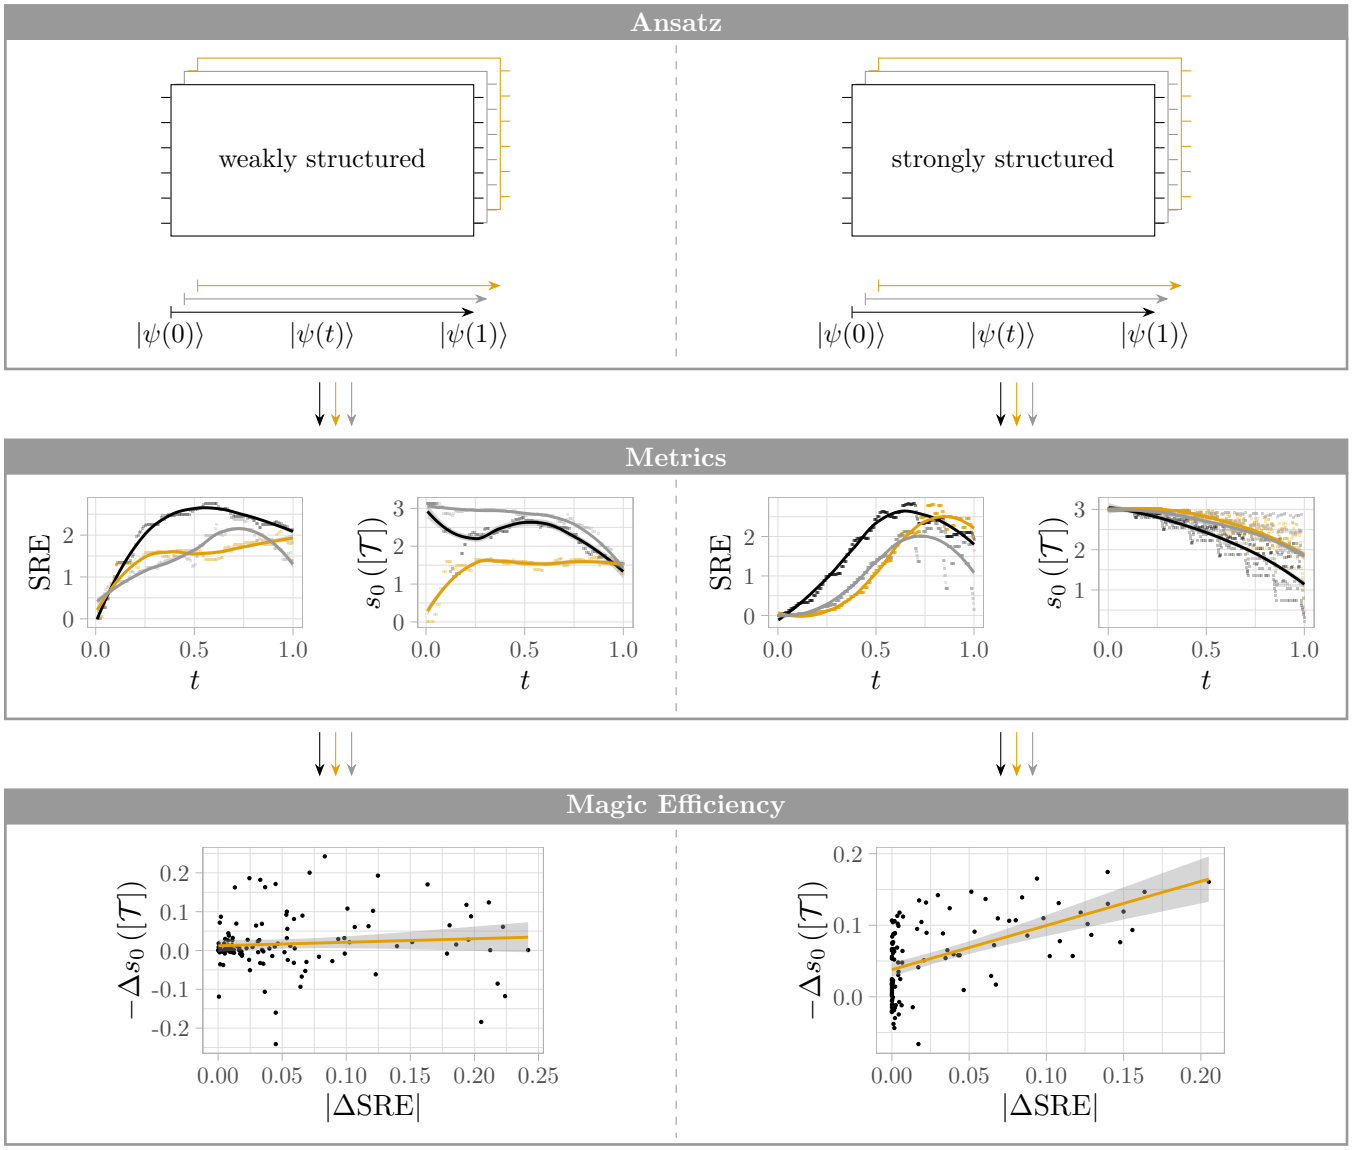

FIG. 1. Overview of our main methodology and contributions. In this paper we compare the state evolution $|\psi(t)\rangle$ (as $0 \overset{t}{\mapsto} 1$) throughout the circuit of a weakly and a strongly structured ansatz. Along this state evolution we compute two metrics. The first being the Stabiliser-Rényi-Entropy (SRE) and the second a geodesic distance from the target space invariant under qubit permutations. As can be seen in the middle pane, both metrics behave much more uniform in the strongly structured state evolution compared to the weakly structured one. In a final step (bottom pane) we compare the magic consumption $|\Delta\,\mathrm{SRE}|$ with the resulting geometric approach to the target $-\Delta s_0([\mathcal{T}])$. Doing so we can show that the strongly structured state evolution has a higher magic efficiency than the weakly structured one.

that $U^\dagger P_j U : P_j \mapsto P_{\pi(j)}$. Therefore, $\{\Xi_{P_j}(|\psi\rangle)\} = \{\Xi_{P_{\pi(j)}}(|0\rangle)\}$ and consequently $\sum_{P\in\mathcal{P}/\langle\pm\mathrm{i}\,\mathbb{1}\rangle}\Xi_P^\alpha(|\psi\rangle) = \sum_{P\in\mathcal{P}/\langle\pm\mathrm{i}\,\mathbb{1}\rangle}\Xi_P^\alpha(|0\rangle)$. Note that, $\langle a|\mathrm{Z}|a\rangle|_{a=0,1} = \pm 1$, $\langle a|\mathbb{1}|a\rangle|_{a=0,1} = 1$ and $\langle a|\mathrm{X},\mathrm{Y}|a\rangle|_{a=0,1} = 0$, which leads to the conclusion that

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

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

Now we will investigate the other direction, for which 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 $S \subset \mathcal{P}_\mathrm{n}$ such that $|S| = 2^n$ and $A|\psi\rangle = |\psi\rangle$ for all $A \in S$. Now let's assume $\mathrm{SRE}_\alpha(|\psi\rangle) = 0$ for some arbitrary state $|\psi\rangle$. Written out, that gives

us $(1-\alpha)^{-1}\log\sum_{P\in\mathcal{P}/\langle\pm\mathrm{i}\mathbb{1}\rangle}\Xi_P^\alpha(|\psi\rangle)-\log 2^n=0$ or rewritten $\log\sum_{P\in\mathcal{P}/\langle\pm\mathrm{i}\mathbb{1}\rangle}\Xi_P^\alpha(|\psi\rangle)=\log 2^{n(1-\alpha)}$. Thus, $\log\sum_{P\in\mathcal{P}/\langle\pm\mathrm{i}\mathbb{1}\rangle}2^{-n\alpha}\langle\psi|P|\psi\rangle^{2\alpha}=2^{n-n\alpha}$. From this we can derive a condition on $|\psi\rangle$ for $\mathrm{SRE}_\alpha(|\psi\rangle)$ 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 $P_i\in\mathcal{P}_\mathrm{n}/\langle\pm\mathrm{i}\mathbb{1}_\mathrm{n}\rangle$. Due to f being a constant function, we have $\frac{\mathrm{d}}{\mathrm{d}\alpha}f=0$. Additionally, we know that all $a_i\geq 0$ and therefore $\frac{\mathrm{d}}{\mathrm{d}\alpha}\langle\psi|P|\psi\rangle^{2\alpha}=0$ for all $P\in\mathcal{P}_\mathrm{n}/\langle\pm\mathrm{i}\mathbb{1}_\mathrm{n}\rangle$. From, $\frac{\mathrm{d}}{\mathrm{d}\alpha}\langle\psi|P|\psi\rangle^{2\alpha}=2\langle\psi|P|\psi\rangle^{2\alpha}\log\langle\psi|P|\psi\rangle=0$ we conclude, that $\langle\psi|P|\psi\rangle\in\{0,1\}$ for all $P\in\mathcal{P}_\mathrm{n}/\langle\pm\mathrm{i}\mathbb{1}_\mathrm{n}\rangle$. 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 $S\subset\mathcal{P}_\mathrm{n}/\langle\pm\mathrm{i}\mathbb{1}_\mathrm{n}\rangle$ such that $|S|=2^n$ and $A|\psi\rangle=|\psi\rangle$ for all $A\in S$; showing that $|\psi\rangle\in\mathrm{STAB}$. We arrive at the following key characteristic of the SRE measure proven in its introductory paper [21].

Theorem 1 (originally proven in [21])**.** *A state $|\psi\rangle$ is in STAB if and only if $SRE_\alpha(|\psi\rangle)=0$.*

We refer readers to the original paper for the original proof and a more in-depth discussion [21]. Also note following observation, which makes SRE such an interesting measure for identifying non-classical computation steps in quantum circuits.

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 $\mathrm{SRE}_\alpha(U|\psi\rangle)=\mathrm{SRE}_\alpha(|\psi\rangle)$. □

IV. GEOMETRIC PERSPECTIVE

If we take a random Haar sampled state $|\psi\rangle\sim\mathrm{Haar}$. Then the expected SRE is $\mathbb{E}_{|\psi\rangle\sim\mathrm{Haar}}(\mathrm{SRE}_\alpha(|\psi\rangle))\in\mathcal{O}(n)$ for all $\alpha\geq 2$, with overwhelming probability [30]. Additionally, SREs are linear upper bounded by $\mathrm{SRE}_\alpha(|\psi\rangle)\leq\log(2^n)\in\mathcal{O}(n)$ [21] (technically more precise, stabiliser entropy scales linearly with concentration bounds). This means, although intermediate states with SRE ≥ 0 are linked with and even necessary for quantum advantage, their occurrence is nothing special and has to be expected. Consequently, this raises the question if observed non-stabiliserness contributes to the computation, or if it is merely a by-product of a suboptimal choice of unitary propagators. To make this notion more precise, we need to clarify the meaning of *contributing to the computation* mean: Every computational process can be interpreted as a state evolution that drives a specific initial state $|\psi_0\rangle$ to a target state $|\psi_T\rangle$ that encodes a problem solution (or a superposition 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 [31], Anandan and Aharonov presented exactly this geometric perspective in conjunction with the concept of *geodesic efficiency* $\mu_\mathrm{gd}=s_0/s$ 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 computational problem, for instance 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 $|\psi_\mathrm{T}\rangle$, a quantum algorithm has the freedom to reach any state within the *target space*, which is the subspace of $\mathcal{H}^{\otimes\mathrm{n}}$ 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. In contrast to Hamiltonians usually found in combinatorial optimisation such two-level Hamiltonians encode solution spaces of decision problems, where the solutions are not weighted and the task is to find any valid solution. 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 $c:\mathbb{F}_2^\mathrm{n}\to\mathbb{F}_2$ the solution verifier of a problem with a finite set of classical solutions $T=\{t\in\mathbb{F}_2^\mathrm{n}:c(t)=1\}$. We then extend $c(\cdot)$ to quantum states $|\psi\rangle=\sum_{b\in\mathbb{F}_2^\mathrm{n}}\alpha_b|\mathrm{b}\rangle$ as a linear function on the basis states $\{|b\rangle\}_{b\in\mathbb{F}_2^\mathrm{n}}$:*

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

Further, we define a quantum target space

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

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

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

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

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

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

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

Proof. Let $\mathcal{B}_\mathcal{T} = \{|\mathrm{t}\rangle : t \in T\}$ the basis of $\mathcal{T}$. We then expand $\mathcal{B}_\mathcal{T}$ with $\mathcal{B}_{\overline{\mathcal{T}}} = \{|\mathrm{b}\rangle : b \in \mathbb{F}_2^{\mathrm{n}} \setminus T\}$ such that $\mathcal{B}_\mathcal{T} \cup \mathcal{B}_{\overline{\mathcal{T}}}$ forms a basis of $\mathcal{H}^{\otimes \mathrm{n}}$. Given that expanded basis, we can write every state $|\psi\rangle \in \mathcal{H}^{\otimes \mathrm{n}}$ 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 $|\mathrm{t}\rangle \in \mathcal{T}$ we have $\sum_{i=1}^{|\mathcal{T}|} |\tau_i|^2 = 1$. Now, let $\mathrm{H_c}$ be a problem Hamiltonian as defined in Definition 3, then $\mathrm{H_c} = \sum_{|\mathrm{t}\rangle \in \mathcal{B}_\mathcal{T}} |t\rangle\langle t|$ 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 $\langle \mathrm{H_c} \rangle = \langle \psi | \mathrm{H_c} | \psi \rangle$ is exactly the overlap between $|\psi\rangle$ and its projection onto the target space $\mathrm{H_c}|\psi\rangle$, which satisfies

$$\langle \mathrm{H_c} \rangle = \max_{|\mathrm{t}\rangle \in \mathcal{T}} |\langle \psi | t \rangle|$$

Now we use that $s_0(|\mathrm{t}\rangle) = 2 \arccos |\langle \psi | t \rangle|$ [32]. Now due to the monotonicity of arccos in $[0, 1]$ we can pull out the max from $2 \arccos\langle \mathrm{H_c} \rangle$ to end up at Eq. (8). □

B. Permutations

In a typical quantum circuit, qubits are sequentially numbered. This numbering implies an unsubstantiated sense of order of qubits in the quantum circuit and state vector picture. 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 $\sigma \in S_n$. Yet, geometrically the state vectors of $|0010\rangle$ and $|0001\rangle$ are orthogonal. In this section we will extend the geometric distance measure to address this dilemma. From a computational standpoint, looking at a concrete problem instance there is no order of variables by naming, there are relations between variables[1]. Variables are then being mapped to qubits and inter variable connections usually result in qubit interactions *i.e.* multi qubit gates. 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.

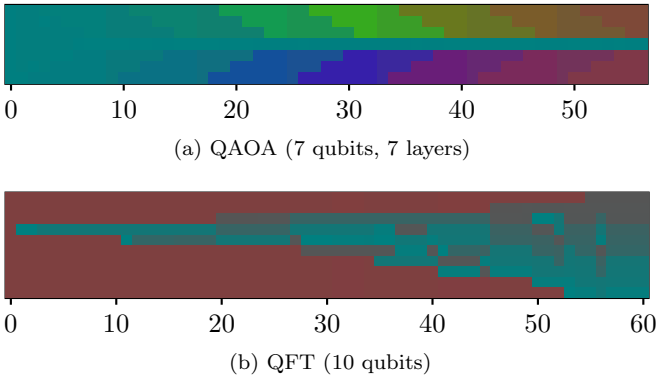

(a) QAOA (7 qubits, 7 layers)

(b) QFT (10 qubits)

FIG. 2. 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 $\mathrm{hsv}(\langle P_0 \rangle, \langle P_+ \rangle, \langle P_{+i} \rangle)$. Within a vertical slice, they are sorted according to the hsv tuple. This reveals stark differences in the evolution of qubit permutation invariant structures throughout the state evolutions driven by QAOA and QFT circuits.

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 2 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

[1] Which then could enforce an order not necessarily present in the naming of variables

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.

Our colour spectrum representation of states served as a visual intuition for qubit permutation invariant similarities between states. As an intermediate step, we will now reformulate this idea in a more mathematical way before defining a rigorous permutation invariant distance measure. Consider $\mathrm{hsv}(\langle P_0\rangle, \langle P_+\rangle, \langle P_{+i}\rangle)$ as a point in a colour space, then the color spectrum state representation is a set of points in this space, where each point corresponds to the vector of Pauli expectations $\langle P_0\rangle, \langle P_+\rangle, \langle P_{+i}\rangle$ for a specific qubit of the state. Given two states, we get two clouds of n points in the colour space. The similarity of both clouds can be computed by matching points of both clouds in pairs. The pairwise matching cost will be the euclidean distance between Pauli expectation vectors of both paired states. Finally, we get the similarity measure by taking the mean pairtwise matching cost of the minimal matching. Informally speaking, we calculated how far the points of both clouds drifted apart on average. Mathematically, we define a function $f : \mathbb{N} \times \mathcal{H}^{\otimes \mathrm{n}} \to \mathbb{R}^3$ such that

$$f(j, |\psi\rangle) = (\mathrm{tr}(\rho_j P_0), \mathrm{tr}(\rho_j P_+), \mathrm{tr}(\rho_j P_{+i})), \tag{9}$$

with ρ_j being the reduced density matrix of the j-th qubit. Using f we now define a state map $C : \mathcal{H}^{\otimes \mathrm{n}} \to \mathbb{R}^{n\times 3}$:

$$C(|\psi\rangle) = [f(i, |\psi\rangle)]_{j=1}^{n} \tag{10}$$

On this state representation we calculate a distance $d_{\mathrm{MMC}}(|\psi_1\rangle, |\psi_2\rangle)$ which is the mean of the minimal bipartite row matching between $C(|\psi_1\rangle)$ and $C(|\psi_2\rangle)$ with matching costs $\alpha_{j,l} = \|f(l, |\psi_2\rangle) - f(j, |\psi_1\rangle)\|$.

Figure 3 shows how this neatly bridges between the colour spectrum intuition and the rigorous permutation invariant distance we will define now. As a first step, we define a permutation operator capturing the notions discussed above.

Definition 4. *Given a permutation $\sigma \in S_n$ and $|\mathrm{b}\rangle \in \mathcal{B}^n$ where $|\mathrm{b}\rangle = |b_1\rangle \otimes \cdots \otimes |b_n\rangle$ with $|b_i\rangle \in \{|0\rangle, |1\rangle\}$, then we define $\hat{\sigma} \in U(\mathcal{H}^{\otimes \mathrm{n}})$ as*

$$\hat{\sigma}\,|\mathrm{b}\rangle = \left|b_{\sigma(1)}\right\rangle \otimes \cdots \otimes \left|b_{\sigma(n)}\right\rangle \tag{11}$$

and

$$\hat{\sigma} \sum_{b\in\mathcal{B}^n} \alpha_b\,|\mathrm{b}\rangle = \sum_{b\in\mathcal{B}^n} \alpha_b \hat{\sigma}\,|\mathrm{b}\rangle\,. \tag{12}$$

Note that in Definition 4 the inverse operator $\hat{\sigma}^\dagger \in U(\mathcal{H}^{\otimes \mathrm{n}})$ 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, $\hat{\sigma} \in \mathcal{H}^{\otimes \mathrm{n}}$ 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. □

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 if and only if there exist a permutation operator $\hat{\sigma}$ as defined in Definition 4, such that $|\psi_r\rangle = \hat{\sigma}|\psi_l\rangle$. Then*

$$[|\psi\rangle] = \{\hat{\sigma}|\psi\rangle : \forall \hat{\sigma}\} \tag{13}$$

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 $\mathcal{H}^{\otimes \mathrm{n}}$, we then extend this notion by defining

$$[\mathcal{T}] = \bigcup_{|\mathrm{t}\rangle\in\mathcal{T}} [|\mathrm{t}\rangle] \tag{14}$$

From Theorem 3 it also immediately follows that $\mathrm{SRE}_\alpha(|\psi\rangle) = \mathrm{SRE}_\alpha([|\psi\rangle])$. 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*

$$s_0([|\phi\rangle]) = \min_{|\phi'\rangle\in[|\phi\rangle]} s_0(|\phi'\rangle) \tag{15}$$

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

$$s_0([\mathcal{T}]) = \min_{|\mathrm{t}\rangle\in\mathcal{T}} s_0([|\mathrm{t}\rangle]) \tag{16}$$

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*

$$s_0([\mathcal{T}]) = s_0(|\psi\rangle, [\mathcal{T}]) = s_0([|\psi\rangle], \mathcal{T}) \tag{17}$$

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 $\min_{\hat{\sigma}} s_0(\hat{\sigma}\,|\psi\rangle, \mathcal{T}) = \min_{\hat{\sigma}} 2\arccos\left\langle \psi \middle| \hat{\sigma}^\dagger \mathrm{H_c}\, \hat{\sigma} \middle| \psi \right\rangle$. As we are minimising over the whole group of all permutation operators we can also

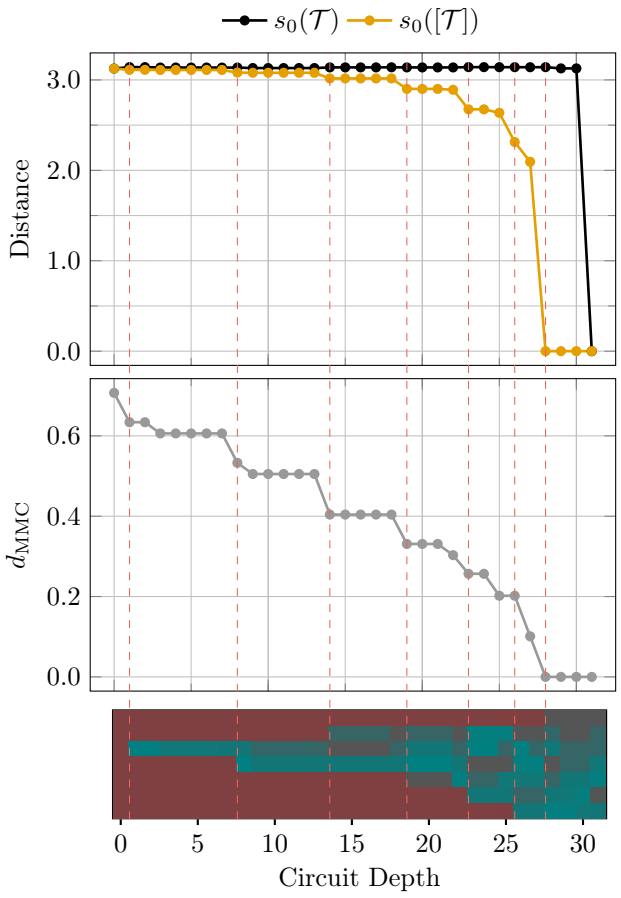

FIG. 3. Minimal geodesic distance for increasing circuit depths. Non-Clifford computational progress can be observed prior to the final qubit order reversal when all target space permutations $s_0([\mathcal{T}])$ (ochre) are considered. As such effects are not visible in the direct distance to the target 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).
The permutation invariant distance $s_0([\mathcal{T}])$ is also much more inline with the structural chances observed in the color spectrum representation of the state evolution under the QFT circuit. See Fig. 2 for a detailed description of the color spectrum representation.
Also, the mean minimal matching cost distance (d_{MMC}) to the final circuit state aligns with both $s_0([\mathcal{T}])$ and the structural changes of the color spectrum representation.

minimise over all complex conjugate operators instead $\min_{\hat{\sigma}^\dagger} 2\arccos\left\langle\psi\middle|\hat{\sigma}\,\mathrm{H_c}\,\hat{\sigma}^\dagger\middle|\psi\right\rangle$. By the canonical definition of $\mathrm{H_c}$ we have $\hat{\sigma}\,\mathrm{H_c}\,\hat{\sigma}^\dagger = \sum_{t\in T}\hat{\sigma}|t\rangle\langle t|\,\hat{\sigma}^\dagger$. Recall that $\{|\mathrm{t}\rangle : t \in T\}$ is the basis of the corresponding quantum target space $\mathcal{T}$. This means, by applying $\hat{\sigma}\,\mathrm{H_c}\,\hat{\sigma}^\dagger$ 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

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

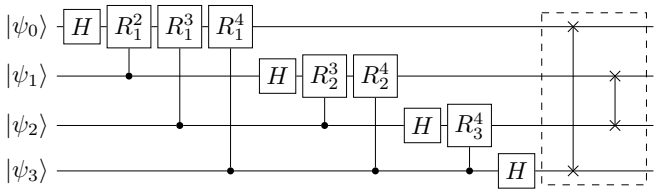

FIG. 4. QFT circuit with four qubits. The dashed box marks the qubit order inversion block of swap gates. Non-stabiliser computations take place before this block, but their computational influence on the geodesic distance is masked by the final qubit reordering.

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 speed-ups [33, 34], thus making it an interesting quantum primitive to study [35]. If one intends to use a distance measure to quantify computational progress it must be applicable to analyse the QFT as well. The interesting parts of the QFT circuit take part before the qubits are reordered in a final step (marked section in Fig. 4). 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 cannot be the case, a the QFT algorithm enables exponential speed-ups. Therefore valuable computational progress has to be made before, taking possible target space permutations into account to reveal such effects (see Fig. 3).

V. NUMERICAL SIMULATIONS

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 state 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.

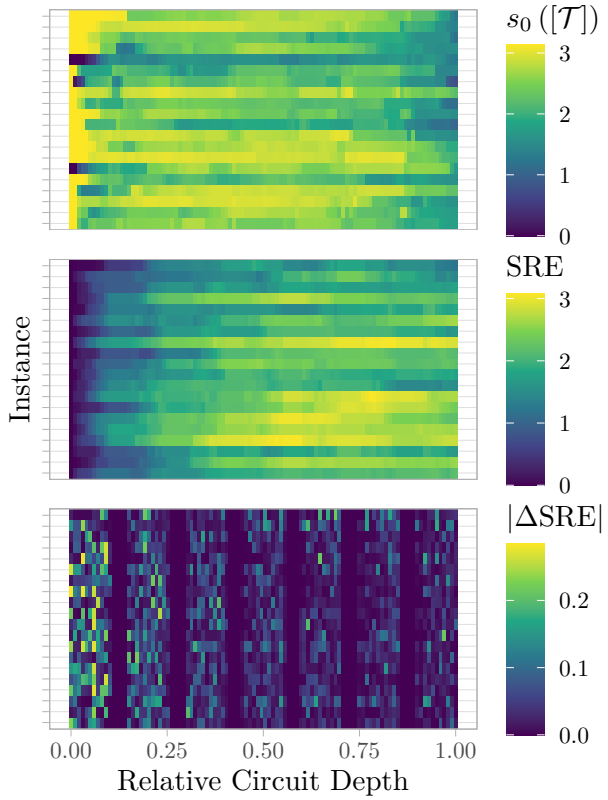

(a) Properties of intermediate states (x-axis) evolving under an **weakly structured** ansatz for different instances (y-axis). *Top:* The state approaches the target space with notable erraticity, as evidenced by irregular fluctuations in the geodesic distance from $[\mathcal{T}]$. While non-stabiliserness (*middle*) appears to evolve comparatively smooth, the variation in resource consumption, as quantified by |ΔSRE| (*bottom*), continues to display a markedly irregular behaviour.

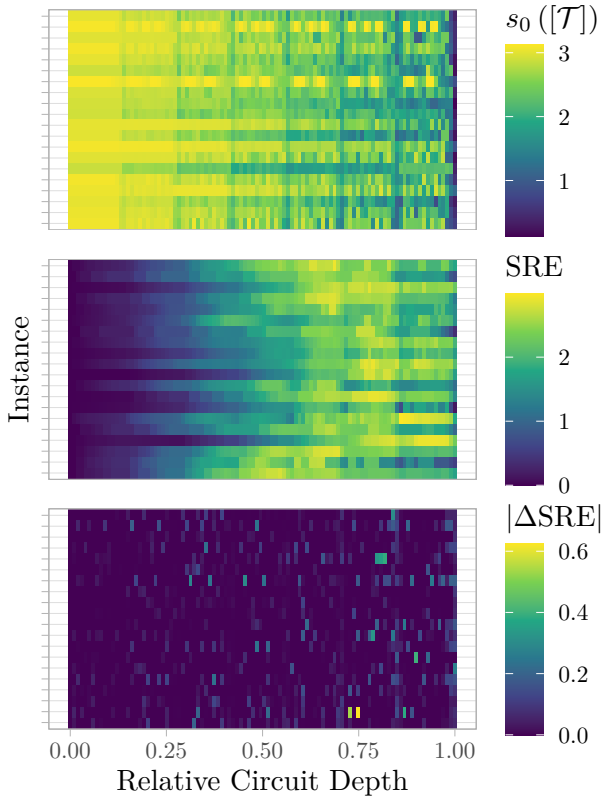

(b) Properties of intermediate states (x-axis) evolving under a **strongly structured** ansatz for different instances (y-axis). *Top:* The state approaches the target space $[\mathcal{T}]$ smoothly. The accumulation of non-stabiliserness (*middle*) follows a structured trajectory and reaches its apex 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. 5. Comparison of intermediate geodesic distances $s_0([\mathcal{T}])$, non-stabiliserness SRE and non-stabiliser consumption $|\Delta\,\mathrm{SRE}|$ between unstructured (Fig. 5a) and structured (Fig. 5b) state evolution. As the circuits vary in gate count, the depicted values are linearly interpolated to 100 steps to achieve uniform value distributions over the relative circuit depth accross all circuits.

A. Problem Structure in State Evolution Ansätze

When choosing an ansatz one can decide between different levels of expressiveness. Whilst a more expressive and thus more general ansatz seems to be the obvious choice, they usually need a high number of free parameters. Alternatively, information about the problem can be used to construct more specific ansätze. An example for a general ansatz would be a hardware efficient variational quantum eigensolver. In the case of QAOA on the other hand a problem hamiltonian is driving the unitary evolution of problem layers alternating with mixing layers. Here, the induced problem structure restricts the state evolution of the ansatz but therefore guides the with less variational parameters needed.

In this work, we will use the expressibility [36, 37] of an ansatz to quantify circuit induces structure on the resulting state evolution. Following [37] we define the expressibility of a circuit by the Kullback-Leibler divergence ($\mathrm{D_{KL}}$) between the fidelity distributions of pairwise sampled states from the ansatz's parametrised output distribution and Haar randomly sampled state pairs.

Definition 7. *Let $C(\theta)$ be a parametrised ansatz, with an angle parametrisation $\theta \in [0,2)^k$ for some $k \in \mathbb{N}$. Given two random states $|\psi\rangle$ and $|\phi\rangle$ the fidelity $F = |\langle\psi|\phi\rangle|^2$ is a random variable. Then, $P_C(F)$ is the probability distribution of fidelities $F(\theta,\varphi) = |\langle\psi(\theta)|\psi(\varphi)\rangle|^2$ with $|\psi(\theta)\rangle = C(\theta)|0\rangle$, $|\psi(\varphi)\rangle = C(\varphi)|0\rangle$ for angle parametrizations θ,φ sampled uniformly at random. Further, $P_{\mathrm{Haar}}(F)$ is the probability distribution of fidelities $F = |\langle\psi|\phi\rangle|^2$ with $|\psi\rangle$ and ϕ randomly sampled from the Haar distribution. Then, we define the expressibility of*

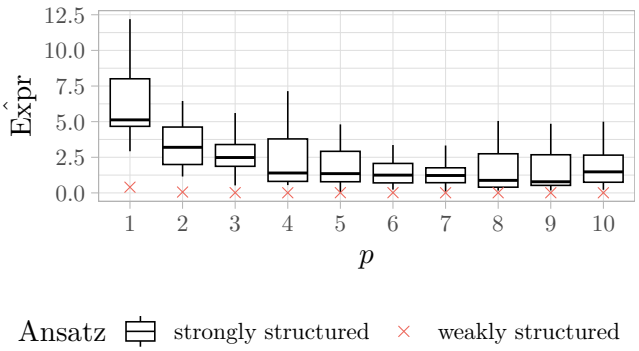

FIG. 6. Expressibility of a strongly structured ansatz (QAOA) and a weakly structured ansatz (VQE) with various layer depths p. In the strongly structured case, the ansatz is constructed with instance information. The distribution of expressibilities for 20 random instances per p is shown as box plots. The weakly structured VQE ansats is instance independent and therefore is depicted as a single data point. The weakly structured ansatz achieves close to full expressibility (*i.e.* Expr(0)) while the strongly structured one converges above Expr(0)

C by:

$$\mathrm{Expr}(C) = \mathrm{D}_{\mathrm{KL}}(P_C(F) \,||\, P_{\mathrm{Haar}}(F))$$

If an ansatz C has an expressibility $\mathrm{Expr}(C) \approx 0$, then its output distribution $\{C(\theta)\,|0\rangle\}_\theta$ with regards to a random angle parametrisation θ is close to the Haar distribution. Otherwise, the structure of C induces a bias to the output state distribution forcing it to diverge from the Haar distribution. In the picture of state evolutions we consider $|0\rangle \mapsto C(\theta)\,|0\rangle$ to be either a *weakly structured state evolution* if $\mathrm{Expr}\, C \approx 0$ or a *strongly structured state evolution* if $\mathrm{Expr}\, C \gg 0$.

Like [37], we will estimate $\mathrm{Expr}(C)$ by uniformly sampling pairs of parametrisations $\theta, \varphi \in [0, 2)^k$ and computing the corresponding fidelity $F(\theta, \varphi) = |\langle\psi(\theta)|\psi(\varphi)\rangle|^2$, where $|\psi(\cdot)\rangle = C(\cdot)\,|0\rangle$. In this picture, F is a random variable and we can estimate $P_C(F)$ by the histogram of the sampled $F(\theta, \varphi)$ realisations of F. The histogram of the F-distribution for random Haar pairs, can be analytically computed [36]. This gives us the estimator:

$$\hat{\mathrm{Expr}}(C) = \mathrm{D}_{\mathrm{KL}}(\hat{P}_C(F) \,||\, P_{\mathrm{Haar}}(F)) \tag{18}$$

We will now show two exemplary ansätze. One being weakly and the other strongly structured according to our expressibility based argument.

B. Weakly Structured State Evolution Ansatz

A highly expressible (*i.e.* weakly structured) state evolution ansatz requires 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 all problems and problem instances, otherwise problem or instance information would potentially induce circuit structures. Concrete instance or problem specific structure only gets introduced during the cost function optimisation process and is contained in the final optimised circuit parametrisation $\theta \in [0, 2\pi)$. These considerations motivate to investigate as an exemplary ansatz a hardware efficient variational quantum eigensolver (VQE). 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. 7.

Figure 6 shows the estimated expressibility $\hat{\mathrm{Expr}}$ of the hardware efficient VQE ansatz with different numbers of layers p as described above. To compute $\hat{\mathrm{Expr}}$ we sampled 1000 random fidelities $F(\theta, \varphi)$ and estimated P_C and P_{Haar} using histograms with 100 bins of uniform bin with. One can observe that the expressiblity of this ansatz converges close to 0, meaning that it does describe a weakly structured state evolution.

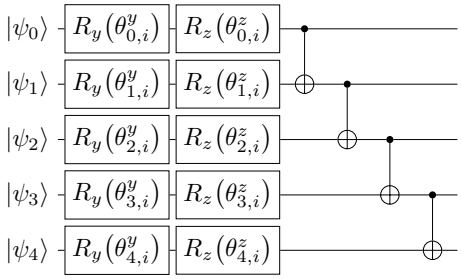

FIG. 7. The i-th layer of the hardware efficient ansatz used for unstructured state evolution.

C. Strongly Structured State Evolution Ansatz

In contrast to weakly structured 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 key parameters like the expected target space overlap of such structured state evolution methods [38]. This shows, that the structural infusion significantly impacts the ansatz even before instance specific cost function optimisation techniques are applied. As a representative for strongly structured state evolution, we chose a standard QAOA ansatz where the driving problem Hamiltonian is the problem Hamiltonian H_c defined above. This equals the construction by Krueger and Mauerer [38].

For the sake of completeness we will shortly recap the construction presented in [38]. Additionally, as it was not

required in the original paper, we will also provide a gate level decomposition of the mixer and problem unitaries. The i-th layer of the described QAOA ansatz is given by

$$U_L(\beta_i, \gamma_i) = e^{-i\beta_i \mathrm{H_x}} e^{-i\gamma_i \mathrm{H_c}}, \tag{19}$$

where $\mathrm{H_x} = \sum_{j=1}^{n} \sigma_j^x$. Here σ_j^x is the Pauli-X operator on the j-th qubit. The resulting mixing layer $e^{-i\beta_i \mathrm{H_x}}$ can thus be decomposed into single qubit rotating gates

$$e^{-i\beta_i \mathrm{H_x}} = \prod_{j=1}^{n} R_x^j(\beta_i),$$

where R_x^j is a x-rotation applied to the j-th qubit. To decompose $e^{-i\gamma_i \mathrm{H_c}}$, we need to understand its effect. Being a projector onto the target space $\mathcal{T}$, the constraint Hamiltonian is diagonal $\mathrm{H_c} = \mathrm{diag}(c(x_0), c(x_1), \cdots, c(x_{2^n-1}))$ with $x_i \in \mathbb{F}_2^n$ being the binary representation of i and $c(x_i) \in \mathbb{F}_2$ the constraint function which is only 1 if x_i is a valid solution. Thus, $e^{-i\gamma_i \mathrm{H_c}}$ is a phase gate which maps $|\mathrm{t}\rangle \mapsto e^{-i\gamma_i} |\mathrm{t}\rangle$ iff $t \in T$. To decompose $e^{-i\gamma_i \mathrm{H_c}}$ we construct a subcircuit $U_C(\gamma) = \prod_{t \in T} P(t, \gamma)$ where $P(t,\gamma)|x\rangle = e^{-i\gamma \mathbf{1}_{x=t}}|x\rangle$, with $\mathbf{1}_{x=t} \in \mathbb{F}_2$ being the indicator function which equals 1 iff $x = t$. Further, $P(t, \gamma)$ can be implemented by applying a X gate to each qubit $i \leq n-1$ if the i-th bit t_i in t equals 0. After that a $n-1$ multi qubit controlled phase gate $C_{n-1}P(\gamma)$ is applied to the n-th qubit. Finally, the X gates are being reversed, resulting in:

$$P(t, \gamma) = \prod_{i=1}^{n-1} X_i^{1-t_i} C_{n-1} P(\gamma) \prod_{i=1}^{n-1} X_i^{1-t_i} \tag{20}$$

Note that the X gates on the n-th qubit can be omitted as the phase gate only applies to the 1 state of the qubit. Figure 8 shows the construction of the $P(t, \gamma)$ gate. In summary, the p-layer QAOA ansatz used in our simulations is constructed as follows:

$$U_L(\beta_p, \gamma_p) \cdots U_L(\beta_2, \gamma_2) U_L(\beta_1, \gamma_1), \tag{21}$$

U_L as described by Eq. (19) with its decomposition as described above.

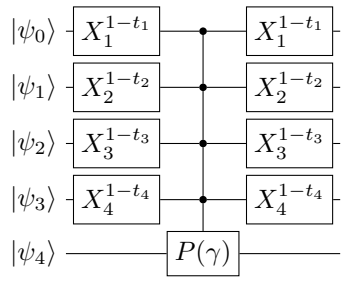

FIG. 8. The decomposition of the $P(t, \gamma)$ gate.

Just as in the weakly structured case we utilize the expressibility to argue that this ansatz is strongly structured. The first argument we make is, that its expressiblity depends on the specific problem instance. This shows, that instance information is used to induce structure into the circuit. Further, Fig. 6 shows that even with more layers the median expressiblity converges at a higher value than the weakly structured ansatz. As the box plots show, it also converges signifficantly above 0 ($\hat{\mathrm{Expr}} \gg 0$). Therfore, we argue that the ansatz represents strongly structured state evolution ansatzes.

D. 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 $\mathrm{H_c} := \mathbb{1} - \prod_{i=1}^{m} |t_i\rangle\langle t_i|$ satisfies Eq. (7). See Ref. [38] for more details.

E. Setup of Simulations

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. For our simulations we used the QuTiP library [39].

F. Results

Comparing the state evolution of strongly structured and weakly structured circuits we notice that the former apart from local fluctuations generally approaches the target space $[\mathcal{T}]$ in a more direct path, smoothly reducing the geodesic distance. In contrast, the weakly structured evolution seems to more erratically move through the state space, witnessed by bigger deviations away from the target space while passing through the circuit. The differences become apparent when comparing the top plots

of Figs. 5a and 5b. Although only exemplary, we also refer to the metric plots of Fig. 1, where the aforementioned difference becomes especially apparent. We now further analyse how both state evolutions approached the target space on a step by step basis. For this, we calculate the delta of s_0 before and after each step. Our interest is focused on non-classical computation steps. Therefore, we filter out all steps where no magic consumption took place. To account for numerical errors we filter for states where $|\Delta\,\mathrm{SRE}| > \epsilon$, with $\epsilon = 1 \times 10^{-5}$. Additionally, we discard outliers with a z-score > 3. The z-score is defined by

$$z(x) = \frac{|x - \mu|}{\sigma}, \tag{22}$$

where μ is the distribution mean and σ the standard deviation. Figure 10 shows that the distribution of the $\Delta s_0([\mathcal{T}])$ is symmetrically centred around zero, in the weakly structured case. For the strongly structured evolution, on the other hand, we observe that the distribution of Δs_0 values is skewed towards the regime below zero (Table I shows more detailed numbers). This indicates that speaking on a per-step basis the strongly structured ansatz more efficiently approaches the target, exhibiting a higher magic efficiency. For the weakly structured 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 (see Fig. 11) of bigger value seem to suggest that the weakly structured approach is able to reach further in larger individual steps, reaching the target faster if utilised efficiently.

Structuredness	Q1	Q2	Q3	$\Delta s_0 < -\epsilon$	$\Delta s_0 > \epsilon$
strongly	−0.0413	0.0000	0.0000	35.2%	6.01%
weakly	−0.0001	0.0000	0.0000	25.6%	22.8%

TABLE I. 25% (Q1), 50% (Q2), and 75% (Q3) quartiles of the $\Delta s_0([\mathcal{T}])$ distributions for strongly structured and weakly structured state evolution, Outliers with z-score > 3 were discarded. The last two columns depict fractions of steps that decrease ($\Delta s_0 < -\epsilon$) or increasing ($\Delta s_0 > \epsilon$) target distance, where $\epsilon = 1 \times 10^{-5}$. On the whole, the distribution of Δs_0 for strongly structured state evolutions exhibits a pronounced skewness towards negative values, in marked contrast to the weakly structured scenario, where it is mostly symmetrically centered about zero.

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. 5a and 5b (bottom), one sees that, the weakly structured ansatz has a significantly denser distribution of magic consuming evolution steps compared to

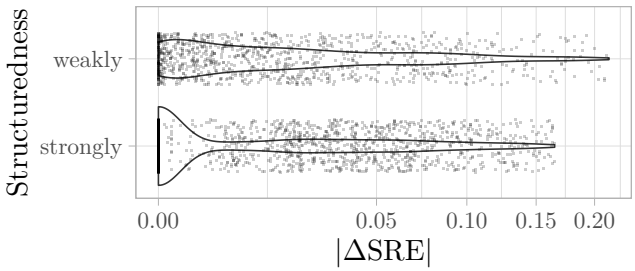

FIG. 9. The distribution of magic consumption $|\Delta\,\mathrm{SRE}|$ of computational steps that had no effect on the geometric target distance $s_0([\mathcal{T}])$. Outliers with z-score > 3 were discarded. The $|\Delta\,\mathrm{SRE}|$ axis is square root scaled

the strongly structured ansatz, whereas the latter shows a higher peak magic consumption throughout its evolution steps. This begs the question whether this shows a higher magic efficiency of the structured state evolution, especially taking the more uniform target appraoch of them compared to weakly structured evolutions. To explore this question we will compare each geometric step $\Delta s_0([\mathcal{T}])$ with its corresponding magic consumption $|\Delta\,\mathrm{SRE}|$. Figure 10 suggests a bimodal distribution off $\Delta s_0([\mathcal{T}])$ values differentiating between two cases: actual geometric steps ($\Delta s_0 > \epsilon$) and computational steps where the state did neither decrease nor increase its geometric distance to the target $\mathcal{T}$ ($|\Delta s_0| < \epsilon$), both up to a precision of ϵ. During the strongly structured state evolution only 23.5% of the computational steps had no effect on the geometric distance to $\mathcal{T}$ whereas during the weakly structured case 42.2% of the steps had no effect on s_0. Looking further into it, we see that the distance invariant computational steps on average consumed more magic in the weakly structured state evolution compared to the strongly structured one (see Fig. 9). For the second mode of computational steps with $|\Delta s_0([\mathcal{T}])| > \epsilon$, there is a clear positive correlation between step-wise geodesic distance reductions to the target space and magic consumption for the strongly structured state evolution. In contrast to that observation, there is no such correlation for the weakly structured case, as shown in Fig. 12. This further substantiates the hypothesis that the strongly structured ansatz utilises non-stabiliser resources more efficiently.

VI. DISCUSSION

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

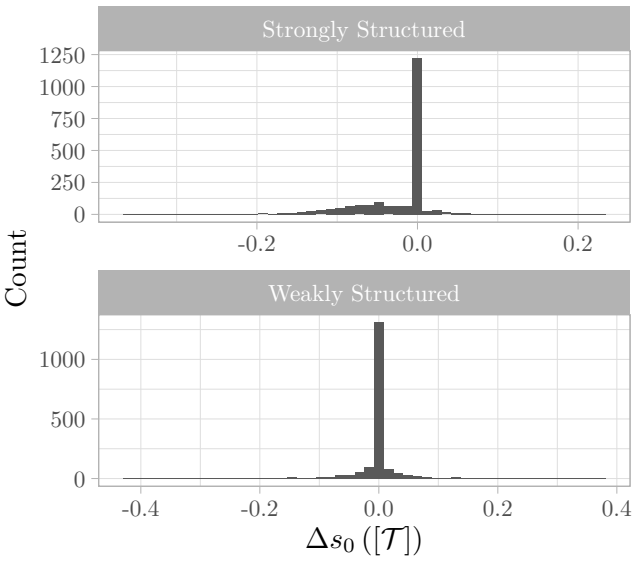

FIG. 10. Distribution of increments and decrements in distance to the target space ($x \sim \pm\Delta s_0([\mathcal{T}])$) of computational steps with magic a consumption $|\Delta\,\mathrm{SRE}| > 1 \times 10^{-5}$ and Δs_0 outliers (z-score > 3) removed. *Top:* Heavy skew toward decrements is observed for the structured ansatz, and most distance changes are negative. *Bottom:* The distribution of distance changes is approximately centered around zero for the unstructured ansatz, discounting a small number of outliers on the negative side.

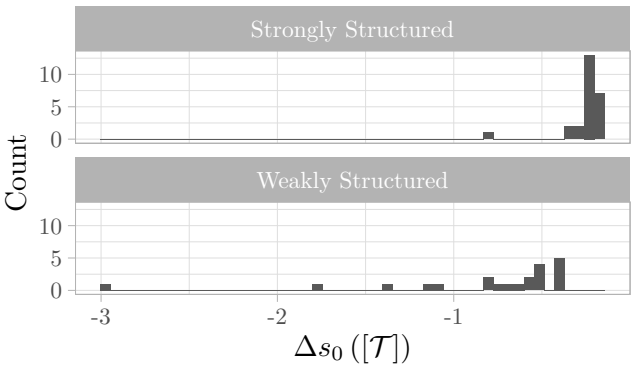

FIG. 11. Distribution of negative $\Delta s_0([\mathcal{T}]) < -\epsilon$ outliers (z-score > 3) discarded in Fig. 10 and Table I.

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 strongly structured and weakly structured state evolutions.

The different distributions of geodesic distance changes in addition with a less dense distribution of magic consuming evolution steps suggest a higher geodesic and magic efficiency for the strongly structured evolution.

Our analysis a bimodal distribution of geodesic step sizes for weakly and strongly structured state evolution. One mode captures all gates that have no effect on the geometric target distance. Here magic consumption can be considered unproductive. The weakly structured state evolution consumes more magic in these steps compared to the strongly structured one. The second mode captures all gates that do influence the geometric target distance, either positively or negatively. By combining resource theoretic Stabiliser-Rényi-Entropy and geometric geodesic distance measures, we were able to show that for these gates the strongly structured ansatz is significantly more efficient in the consumption of non-stabiliser resources than the weakly structured ansatz. On a methodical level, this demonstrated the potential of our combination of methodologies. In summary, the structured evolution is less wasteful in magic consumption with gates that don't affect the geometric target distance and additionally shows a higher magic efficiency when actually geometrically moving the state towards the target. Overall the structured state evolution is significantly more magic efficient.

It is also interesting to note that the magic consumption of the structured state evolution we investigated seems to match the observations of [40] quite well, where magic buildup peeked roughly in the middle of the evolution before it got reduced while closing in onto the target state. This is interesting as in [40] the authors investigate QAOA in the context of combinatorial optimisation with low degeneracy in the ground state (*i.e.* target space in our terms). The authors observed how the state has to pass a, in their terms, *magic barrier* in order to reach the non-degenerated (stabiliser) state. We argue that in a case where the initial and target state both are low in magic, this shape is necessary to achieve quantum speedups. This becomes clear under our perspective, where indeed not magic in itself is the quantum resource associated with quantum computations (and thus speedups) but rather the *magic consumption* $|\Delta\,\mathrm{SRE}|$. Thus, the *magic barrier* observed in [40] can be seen as the accumulated magic consumption during the state evolution as both steps increasing and decreasing magic are inherently non-classical and can both be efficient computational steps. In fact, evolving a stabiliser initial state to a stabiliser target state through a flat magic landscape would hint at a primarily classic computation not touching any quantum potential. In a variational scheme this would mean: moving the hard computational effort to the classical optimisation loop.

In contrast to [40] our construction captures decision problem that quite naturally have structured solution spaces leading to entangled non-stabiliser high magic targets spaces. Just thing of valid solutions of SAT formulas which have inherent dependencies between singe bit

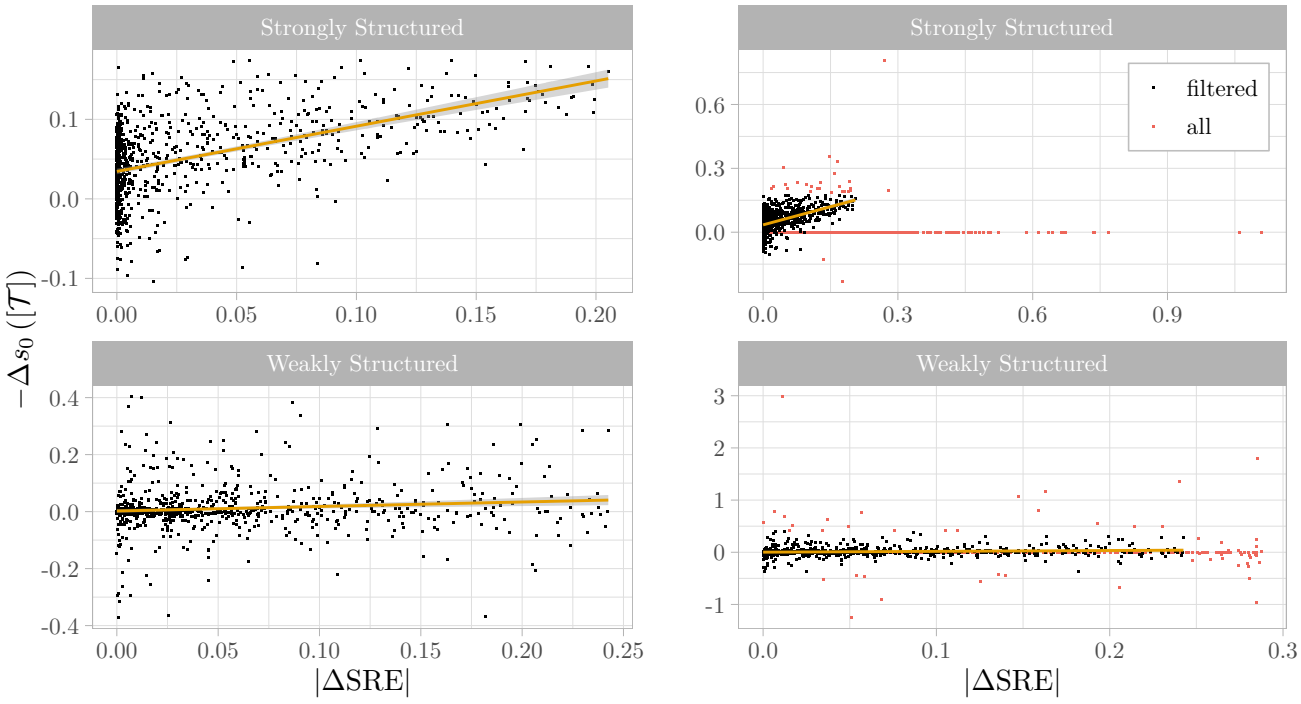

(a) A close up of the filtered data points.

(b) The filtered data points in the context of all data points.

FIG. 12. Correlation of magic consumption $|\Delta\,\mathrm{SRE}|$ and geometric target distance variation $\Delta s_0([\mathcal{T}])$ for all steps with $\Delta s_0 > \epsilon$: Strongly structured state evolution (top) shows clear correlation between magic 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 weakly structured ansatz (bottom). Outliers with z-score > 3 in $|\Delta\,\mathrm{SRE}|$ or $\Delta s_0([\mathcal{T}])$ were discarded.

flips cascading to other bits. Or in a more pathological case, we could also force our structured ansatz to prepare the $|w\rangle$ state, which has magic of $\mathrm{SRE}(|w\rangle) = 3\log_2(n) - \log_2(7n-6)$ [41]. Despite this stark difference we also witnessed a *magic barrier* like effect overreaching the final magic. This raises the question if such a *magic barrier* shape is a sign of efficient quantum resource consumption in a circuit. Further, one could analyse how the distribution between magic consumption during the state evolution and the final state magic differs weather the quantum speedup is routed more in the state evolution itself or the sampling from a high magic state.

Another interesting aspect probably worth a follow up investigation is the roll of the parameter space. Intuition would suggest that the encoded problem structure of strongly structured state evolution allows for a reduced number of well-placed parametrised operators. In the more general weakly structured case a higher density of parameterised gates is necessary to cover different problems and instances. On an per instance level perhaps only an unknown subset of gates would suffice, allowing the evolution to unproductively meander through the state space as the optimizer in a way *turning the wrong knobs*. Our framework could be utilized for instance to investigate the importance of individual gates in weakly structured ansätze to come up with a more ansätze for certain computational problems.

Recently, magic measures where proposed that capture non-local non-stabiliserness of bipartite systems [42, 43]. It would be interesting to deploy our framework with non-local magic measures. However, the impact of disregarding local magic to computational progress needs to be studied in more detail to infer sound conclusions. In general, our framework is flexible enough to facilitate arbitrary quantum state measures including entanglement measures.

VII. CONCLUSION & 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 $U|\psi_1\rangle = |\psi_2\rangle$. Although Clifford circuits are classically simulable, finding minimal Clifford decompositions is not always efficient. 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.

Appendix A: Reproducibility

A reproduction package [44] that contains all code required to run the simulations in this paper is available at [45][2]. All simulations are done on premise with seeded random number generators. No cloud services were used. Thus, given the availability of the used software packages, our simulations are fully reproducible.

ACKNOWLEDGMENTS

This work was supported by the German Federal Ministry of Research, Technology and Space (BMFTR), funding program *quantum technologies—from basic research to market*, grant number 13N17387. WM acknowledges support by the High-Tech Agenda Bavaria.

[1] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, Noisy intermediate-scale quantum algorithms, Reviews of Modern Physics **94**, 10.1103/revmodphys.94.015004 (2022).

[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] T. Rohe, D. Schuman, J. Nüßlein, L. Sünkel, J. Stein, and C. Linnhoff-Popien, The questionable influence of entanglement in quantum optimisation algorithms (2025), arXiv:2407.17204 [quant-ph].

[11] A. C. Nakhl, T. Quella, and M. Usman, Calibrating the role of entanglement in variational quantum circuits, Phys. Rev. A **109**, 032413 (2024).

[12] P. Díez-Valle, D. Porras, and J. J. García-Ripoll, Quantum variational optimization: The role of entanglement and problem hardness, Phys. Rev. A **104**, 062426 (2021).

[13] A. J. C. Woitzik, P. K. Barkoutsos, F. Wudarski, A. Buchleitner, and I. Tavernelli, Entanglement production and convergence properties of the variational quantum eigensolver, Phys. Rev. A **102**, 042402 (2020).

[14] D. Gross, S. T. Flammia, and J. Eisert, Most quantum states are too entangled to be useful as computational resources, Phys. Rev. Lett. **102**, 190501 (2009).

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

[16] 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.

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

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

[19] S. Bravyi and A. Kitaev, Universal quantum computation with ideal clifford gates and noisy ancillas, Physical

[2] A DOI save version will be made public with a camera ready version of the manuscript

Review A **71**, 10.1103/physreva.71.022316 (2005).
[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] L. Leone, S. F. Oliviero, and A. Hamma, Stabilizer rényi entropy, Physical Review Letters **128**, 10.1103/physrevlett.128.050402 (2022).
[22] E. Chitambar and G. Gour, Quantum resource theories, Reviews of Modern Physics **91**, 10.1103/revmodphys.91.025001 (2019).
[23] S. Bravyi, G. Smith, and J. A. Smolin, Trading classical and quantum computational resources, Physical Review X **6**, 10.1103/physrevx.6.021043 (2016).
[24] 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).
[25] 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).
[26] T. Haug and L. Piroli, Stabilizer entropies and nonstabilizerness monotones, Quantum **7**, 1092 (2023).
[27] L. Leone and L. Bittel, Stabilizer entropies are monotones for magic-state resource theory, Physical Review A **110**, 10.1103/physreva.110.l040403 (2024).
[28] 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).
[29] 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).
[30] 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).
[31] J. Anandan and Y. Aharonov, Geometry of quantum evolution, Physical Review Letters **65**, 1697–1700 (1990).
[32] C. Cafaro, E. Clements, and A. Alanazi, Aspects of complexity in quantum evolutions on the bloch sphere (2025).
[33] A. Y. Kitaev, Quantum measurements and the abelian stabilizer problem (1995).
[34] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, SIAM Journal on Computing **26**, 1484–1509 (1997).
[35] N. Linden and R. de Wolf, Average-Case Verification of the Quantum Fourier Transform Enables Worst-Case Phase Estimation, Quantum **6**, 872 (2022).
[36] G. Ilário Correr, I. Medina, P. C. Azado, A. Drinko, and D. O. Soares-Pinto, Characterizing randomness in parameterized quantum circuits through expressibility and average entanglement, Quantum Science and Technology **10**, 015008 (2024).
[37] S. Sim, P. D. Johnson, and A. Aspuru-Guzik, Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms, Advanced Quantum Technologies **2**, 10.1002/qute.201900070 (2019).
[38] T. Krüger and W. Mauerer, Out of the Loop: Structural Approximation of Optimisation Landscapes and non-Iterative Quantum Optimisation, Quantum **9**, 1903 (2025).
[39] N. Lambert, E. Giguère, P. Menczel, B. Li, P. Hopf, G. Suárez, M. Gali, J. Lishman, R. Gadhvi, R. Agarwal, A. Galicia, N. Shammah, P. Nation, J. R. Johansson, S. Ahmed, S. Cross, A. Pitchford, and F. Nori, Qutip 5: The quantum toolbox in python, arXiv preprint arXiv:2412.04705 (2024), arXiv:2412.04705 [quant-ph].
[40] C. Capecci, G. C. Santra, A. Bottarelli, E. Tirrito, and P. Hauke, Role of nonstabilizerness in quantum optimization (2025).
[41] J. Odavić, T. Haug, G. Torre, A. Hamma, F. Franchini, and S. M. Giampaolo, Complexity of frustration: A new source of non-local non-stabilizerness, SciPost Physics **15**, 10.21468/scipostphys.15.4.131 (2023).
[42] D. A. Korbany, M. J. Gullans, and L. Piroli, Long-range nonstabilizerness and phases of matter, Phys. Rev. Lett. **135**, 160404 (2025).
[43] D. Qian and J. Wang, Quantum nonlocal nonstabilizerness, Phys. Rev. A **111**, 052443 (2025).
[44] 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).
[45] W. M. Tom Krueger, Reproduction package, `https://github.com/lfd/quantum_dark_magic`.