

Lightcone bounds for quantum circuit mapping via uncomplexity

Steinberg, Matthew; Bandić, Medina; Szkudlarek, Sacha; Almudever, Carmen G.; Sarkar, Aritra; Feld, Sebastian

DOI

[10.1038/s41534-024-00909-7](https://doi.org/10.1038/s41534-024-00909-7)

Publication date

2024

Document Version

Final published version

Published in

NPJ Quantum Information

Citation (APA)

Steinberg, M., Bandić, M., Szkudlarek, S., Almudever, C. G., Sarkar, A., & Feld, S. (2024). Lightcone bounds for quantum circuit mapping via uncomplexity. *NPJ Quantum Information*, 10(1), Article 113. <https://doi.org/10.1038/s41534-024-00909-7>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.




Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

<https://doi.org/10.1038/s41534-024-00909-7>

Lightcone bounds for quantum circuit mapping via uncomplexity



Matthew Steinberg^{1,2,4}  , Medina Bandić^{1,2,4}, Sacha Szkuclarek¹, Carmen G. Almudever³, Aritra Sarkar^{1,2} & Sebastian Feld^{1,2} 


Efficiently mapping quantum circuits onto hardware is integral for the quantum compilation process, wherein a circuit is modified in accordance with a quantum processor's connectivity. Many techniques currently exist for solving this problem, wherein SWAP-gate overhead is usually prioritized as a cost metric. We reconstitute quantum circuit mapping using tools from quantum information theory, showing that a lower bound, which we dub the lightcone bound, emerges for a circuit executed on hardware. We also develop an initial placement algorithm based on graph similarity search, aiding us in optimally placing circuit qubits onto a device. 600 realistic benchmarks using the IBM Qiskit compiler and a brute-force method are then tested against the lightcone bound, with results unambiguously verifying the veracity of the bound, while permitting trustworthy estimations of minimal overhead in near-term realizations of quantum algorithms. This work constitutes the first use of quantum circuit uncomplexity to practically-relevant quantum computing.

The promise of quantum technology extends to many areas of modern theoretical physics, computer science and cryptography, among others¹. In spite of much success over the past 30 years, current-generation quantum technology is characterized by noisy, intermediate-scale devices that are severely limited not only by the depth and size of the quantum circuits that can be executed, but also by the qubit connectivity of such devices². Such processors have allowed for the first generation of quantum-technology demonstrations, ranging from experimental realizations of hybrid quantum-classical optimization techniques^{3,4} to resource-intensive algorithms such as fault-tolerant quantum error-correction codes (QECCs)^{5–7}.

With such promise as is forecasted for quantum technology, *full-stack* design approaches have emerged, in order to delegate resources efficiently and to ensure high success rates for a given quantum circuit, realized on a quantum processor^{8–10}. As such, one of the cardinal issues to emerge for practical quantum computing is that of *quantum compilation*, which can be broadly defined as the various engineering-level steps required to translate and prepare a quantum circuit for execution on a quantum processor¹¹. Central to quantum compilation is the *quantum circuit-mapping problem* (QCMP), which concerns the assignment and rearrangement of qubits from an algorithm to a processor as a quantum circuit is executed, in order to guarantee high fidelity of the resulting state^{12,13}. It is known that the quantum circuit-mapping problem is NP-complete^{14,15}, and has been likened to the traveling salesman problem (TSP) on a torus¹⁶. The QCMP is also related to token swapping¹⁷. Many competing approaches have been proposed for

solving the QCMP, with all of the state-of-the-art strategies trading accuracy for speed, among other considerations^{18–34}. However, to our knowledge, no work has attempted to formulate the QCMP from a standpoint grounded in theoretical physics and quantum information theory. The motivation for such an endeavor is twofold. First, since the QCMP is a physical process, such a description can provide new insights and perspectives on how best to solve it. Second, by providing a fundamental description of the QCMP, we lay the groundwork for uniting various contemporary approaches towards a solution, and show how they compare to each other in a self-consistent framework. In short, a physics-motivated description of the QCMP offers consensus for current and future solution strategies, and how best to compare them.

Underpinning the advances in quantum technology, *quantum information theory* seeks to quantify the achievable limits of information processing on a fundamental mathematical basis using quantum physics^{1,35,36}. While much progress is already notable, many fields outside of the immediate scope of quantum information theory have benefited from incorporating quantum-information-theoretic interpretations to outstanding research problems, including theoretical physics^{37,38}, network science (which studies the behavior of complex networks from the standpoint of statistical mechanics and graph theory)^{39–41}, among many others. Bearing in mind such potential, we apply the machinery of quantum information theory, in particular, *quantum circuit complexity*^{42,43}, *spectral graph entropy*^{39,44–46}, and the *quantum operations* formalism^{35,47–49} to the QCMP, in

¹Quantum Computing Division, QuTech, Delft University of Technology, Delft, the Netherlands. ²Department of Quantum & Computer Engineering, Delft University of Technology, Delft, the Netherlands. ³Computer Engineering Department, Technical University of Valencia, Valencia, Spain. ⁴These authors contributed equally: Matthew Steinberg, Medina Bandić.  e-mail: m.a.steinberg@tudelft.nl

order to describe the problem of preparing certain quantum states on a quantum processor whose qubit connectivity is restricted. As a quantum circuit itself describes a sequence of unitary transformations under which a quantum state transforms, addressing such quantum operations under the guise of a processor’s connectivity is not only reasonable using quantum information theory, but also embodies a natural extension of quantum information theory to the setting of practical quantum computing.

Several recent proposals have sought to establish links between graph theory and the QCMP^{28,29,50–52}. Since then, it has become commonplace in the literature to consider an *interaction graph* (IG), in which edges represent the necessary two-qubit interactions for implementing a quantum circuit, and a *coupling graph* (CG), whose edges determine allowed two-qubit interactions between neighboring subsystems on the processor⁵³. In many of these proposals, the SWAP-gate count required to realize a quantum algorithm on a given quantum processor is considered to be a typical metric for the objective function of a mapping strategy²².

In this vein, we strengthen this connection by initiating a study of the QCMP from the theoretical standpoint, using graph theory and network science as a foundation. More concretely, we propose a special case of the QCMP in which all two-qubit interactions of a given IG can be compressed into a single time slice of the quantum circuit; this simplification can be likened to a sort of “lightcone” path through a configuration space, which we explain in detail in Section III. Starting from this point, we translate the IG and CG into density matrices, and calculate their *thermodynamic path length* in the configuration space of density matrices, given certain allowed superoperator transformations. These allowed superoperator transformations consist of a combination of *doubly-stochastic quantum operations*^{49,54} to permute vertices of the CG, and Bell measurements on the IG, in order to sequentially and methodically minimize the path length over the configuration-space geodesic. Using recent results in *quantum circuit complexity*^{42,43} and methods from *quantum information geometry*^{55–57}, we carefully show that entropic divergence measures can be used in order to minimize the distance between density matrices describing the IG and CG along the configuration-space geodesic of allowed quantum operations, and that a minimal SWAP-gate count can be ascertained using this method. This minimal SWAP-gate count is shown to coincide with the *quantum circuit uncomplexity*^{42,43}; as such, we name this lower bound the *SWAP uncomplexity*. As this lower bound does not take into account the traditionally-used gate-dependency graph of the IG, the SWAP uncomplexity represents a lightcone solution to the QCMP, in which infinite parallelization of two-qubit gates is possible.

In addition, we develop a novel algorithm for the *qubit assignment* (or initial placement) of qubits from the IG to the CG, based on a subgraph isomorphism and graph similarity search^{58–60}, which has applications for *multi-programming* on a quantum device⁶¹ and may be of independent interest. In this case, however, it serves as a necessary step in our formulation and further enables our approach by constraining the coupling graph to match the size of an IG, which is one of the method’s crucial requirements.

This algorithm also facilitates a calculation that we use to compute a maximal SWAP-gate count.

Together with the formalism introduced, a combined approach is constructed that searches for the best qubit assignment in terms of the *graph-edit distance* (GED), and then calculates the SWAP uncomplexity, given an IG/CG pair as inputs. We test the resultant algorithm against IBM’s Qiskit compiler, finding that in all cases, the SWAP-gate count as calculated by the SWAP uncomplexity algorithm is never surpassed, in full agreement with our formalism. This lower bound is of great importance, as such constraints can help with the prediction of compilation performance, as well as for making design choices relevant in application-specific mapping strategies and quantum devices⁸.

Results

Quantum Circuit Mapping

Generally, quantum algorithms and their associated circuit-level descriptions are developed without considering the architecture-specific limitations of particular devices, i.e., they are developed in an *architecture-free* manner. For example, the *elementary gate set* (or primitives) for a particular device may differ significantly from what has been indicated at the level of a generic circuit description; as such, several actions must be performed in order to translate the quantum algorithm into a circuit that a quantum device can actually execute. Another example can be seen in the physical connectivity properties of a quantum processor, which must be considered to ensure that the necessary qubit-qubit interactions of the circuit can be performed on the device. Although certain exceptions may exist (in which several of the aforementioned steps may not necessarily be carried out), these procedures are collectively known as *quantum circuit mapping*¹¹.

The task of quantum circuit mapping itself is usually divided into several steps, which typify the process: A) *elementary gate-set decomposition*, which involves the translation of a circuit to a native gate set utilized by a quantum processor; B) *scheduling*, which concerns the formation of a logical time ordering for algorithm execution, and includes considerations for parallelism of operations and for the shortening of circuit depth; C) *qubit assignment*, which relates to the initial assignment of qubits from an algorithm to the physical qubits on a quantum architecture; and D) *qubit routing*, which examines the increase in gate overhead as extra operations are inserted into the algorithm as a function of physically moving qubits around the processor, such that the required two-qubit operations are realizable¹¹. Typically, operations such as SWAP gates are utilized in order to adapt the circuit to hardware; these amount to classical permutation operations on a product state, but other approaches exist as well^{32,62,63}.

As a simple example, consider the quantum circuit-mapping procedure in Fig. 1. The circuit on the left is decomposed into IG form (a)–(b), wherein we do not consider single-qubit gates for simplicity, and we assume that the two-qubit gates shown in the circuit diagrams are taken as general two-qubit operations. Upon comparing the IG with the available qubit-qubit interactions afforded by a quantum device, it is apparent that the




Fig. 1 | An example of the QCMP as a sequence of steps needed to assign qubits from an algorithm to a quantum device. The two-qubit gates in the two circuit diagrams are used to represent general two-qubit unitary operations (with the exception of SWAP gates); here, we do not consider single-qubit gates, and the two-qubit interactions shown in (a) and (d) are taken to be general two-qubit unitary operations. **a** The quantum circuit is transformed into an *interaction graph*, as shown

in **b**. Next, it is compared with the connectivity properties of the *coupling graph* (**c**). As no *graph isomorphism* (i.e., no exact matching between the vertices of the IG and CG, which upholds all of the edge relations of both) exists between the IG and CG, one can compensate for the lack of connectivity by introducing SWAP operations to the circuit in order to realize the circuit. **d** These operations degrade the fidelity of the final output state.

connectivity available on the device and the connectivity required by the algorithm differ (the two-qubit interaction represented by the edge $e_{d_0 d_1}$ is not possible, as shown in (b)–(c)). As such, this discrepancy can be compensated for by adding detrimental SWAP operations to the circuit’s initial assignment (d). This example provides an illustration of *quantum circuit mapping* based on an exact graph matching between IG and CG. Implicit to this example was the assumption that exactly such a mapping between the vertices of the IG and CG exists, which preserves all of the edge relationships of the IG; this is known as a *graph isomorphism*, and will be discussed in more detail in the Supplementary Material.

In this paper, we formulate and solve a special case of the QCMP. In contrast to Fig. 1, our formalism exhibits three main simplifications. First, as is typical in the QCMP, we do not consider single-gate interactions in the formalism that we present starting in Section II B. Second, we assume the existence of a noiseless quantum device, as such an ideal scenario precisely allows for the emergence of a lower bound. Third, we assume that no further gate simplifications via *quantum circuit-synthesis techniques* can be further leveraged^{64–66}. Finally, in real quantum hardware, typically we only consider that one multi-qubit gate can be performed on a given hardware qubit during a given moment; this necessitates the division of the quantum circuit into time slices. Instead, we consider the scenario in which all two-qubit gate interactions of the circuit can be performed within a single unit time slice, i.e. the causal structure of our circuits are taken to be indefinite^{67,68}, as essentially infinite parallelization of two-qubit gate operations can be implemented. We devote more detail to these concepts in Section III.

Graphs as density matrices

In quantum physics, the most general manner of describing a quantum state involves the use of *density matrices*^{1,69}. A density matrix ρ is a Hermitian, positive semidefinite matrix, whose trace is equal to unity. A system ρ is termed *pure* if and only if the bound $\text{Tr}[\rho^2] \leq 1$ is saturated. The density matrix admits a spectral decomposition as

$$\rho = \sum_j p_j |\psi_j\rangle\langle\psi_j|, \tag{1}$$

for an orthonormal basis $\{|\psi_j\rangle\}$, where p_j are non-negative eigenvalues summing to 1.

In this work, as in ref. 39,40, we make use of the concept of a density matrix to describe a complex network (i.e. a graph with many edges and vertices, and assumed topological structure⁷⁰), by defining a matrix from a network, which fulfills the mathematical properties of a density matrix. One such candidate was previously shown in³⁹ to be promising for adhering to the property of *subadditivity* for the VNE; this equilibrium *Gibbs state* is defined as

$$\rho_L = \frac{e^{-\beta L}}{Z}, \tag{2}$$

where ρ_L , $e^{(\cdot)}$, β , and Z represent: the density matrix of graph Laplacian L ; the matrix exponential; the inverse temperature (or diffusion time⁴⁴); and the partition function, which is defined as $Z = \text{Tr}[e^{-\beta L}]$, respectively. We define the graph Laplacian as $L := D - A$, following^{39,71}. Throughout the text, we will refer to the graph Laplacians for the IG and CG using the notation L_{IG} , L_{CG} , respectively, and ρ_{IG} , σ_{CG} to refer to the corresponding IG and CG density-matrix forms, respectively. Additionally, we refer to edges in a graph-theoretic context as a line connecting two vertices; in the density-matrix formalism, we will make reference to this instead with the term *subsystem interactions*.

Using these objects to describe complex networks is advantageous for several reasons. Firstly, although it is known that the graph Laplacian is uniquely determined up to vertex-numbering assignments⁷², the *eigenvalue spectrum* of the graph Laplacian does

not allow for unique identification of a graph. For example, two graphs can be *cospectral*, i.e. possessing the same eigenvalue spectrum, but with different connectivity⁷². As such, the approach we detail in this work is motivated by the fact that *entropic divergence measures* allow for a unique differentiation between two quantum states ρ and σ . Secondly, the VNE is *permutation-invariant*, i.e., the VNE is invariant under a reordering of subsystems. For example, suppose we have a state vector of five subsystems a, b, c, d, and e. If two such subsystem orderings give rise to density matrices $\eta = \sum_{abcde \in Z_2} |abcde\rangle\langle abcde|$ and $\xi = \sum_{abcde \in Z_2} |baced\rangle\langle baced|$, it can be shown that the equality $S(\eta) = S(\xi)$ holds⁴⁹. Lastly, as discussed in ref. 39,40, previous attempts to calculate the classical entropy of a complex network fail, as these measures are dependent on a probability distribution resultant from a specific network descriptor. In contrast, the quantum approach we utilize does not depend on a specific network descriptor, but rather the entire network, rescaled and normalized as a Gibbs state.

Distance Measures in Quantum Information Theory

The task of distinguishing two quantum states is in general a highly non-trivial problem in quantum physics, with many interpretations useful for distinct scenarios^{1,35,36,73}. However, at the core of these distance measures lies a central object known as the *quantum Fisher information*⁵⁷, and is typically calculated as

$$G_{ij} = \sum_{i,j=0}^{d-1} \frac{\text{Re}\left(\langle\lambda_i|\partial_i\rho|\lambda_j\rangle\langle\lambda_j|\partial_j\rho|\lambda_i\rangle\right)}{\lambda_i - \lambda_j}, \tag{3}$$

where $\partial_i = \partial\rho/\partial i$ for some density matrix parameterized by a vector $\theta = \{\theta_1 \dots \theta_m\}$, and we write i, j as shorthand for the parameters $\theta_i, \theta_j, \lambda_i, \lambda_j$ represent the eigenvalues associated with $\partial_i\rho$ and $\partial_j\rho$.

The quantum Fisher information is a fundamental object in quantum information theory, allowing for the derivation of an extended family of statistical inference measures that distinguish between parameterized quantum states in different settings^{55,74}. In particular, it is known that the quantum Fisher information is closely related to the *Bures distance* \mathcal{B}_{ij} ^{1,36,55}, as well as to the *quantum relative entropy* as

$$G_{ij} = 4\mathcal{B}_{ij} = 8\left(1 - \sqrt{\mathcal{F}(\rho_i, \rho_j)}\right) \approx \sqrt{2S(\rho_i||\rho_j)}, \tag{4}$$

where $\mathcal{F}(\rho_i, \rho_j) = (\text{Tr}[\sqrt{\sqrt{\rho_i}\rho_j\sqrt{\rho_i}}])^2$ represents the *fidelity function*, $S(\rho_i||\rho_j) = S(\rho_i) + S(\rho_j)$ is the *quantum relative entropy* (QRE), and $S(\rho_i)$ is the *Von Neumann entropy* (VNE).

From the quantum relative entropy, one can immediately define a similar divergence measure which will be useful to the present work, the *quantum Jensen-Shannon divergence*. It is defined, using the quantum relative entropy, as

$$\mathcal{D}_{qJSD}(\rho_i||\rho_j) = \frac{1}{2}\left[S\left(\rho_i||\frac{\rho_i + \rho_j}{2}\right) + S\left(\rho_j||\frac{\rho_i + \rho_j}{2}\right)\right]. \tag{5}$$

It is also well-known that the quantum Fisher information (and by extension, the quantum Jensen-Shannon divergence) is closely related to the *quantum Wasserstein distance*^{56,73,75}.

Quantum information theory & entropic divergence measures

The task of actually distinguishing two quantum states ρ_i and ρ_j can be accomplished through the use of the *entropic divergence measures*^{1,36,45,46,49,55,74,76,77}. In particular, we employ the *quantum*

Jensen-Shannon divergence (qJSD), which is defined as

$$\mathcal{D}_{\text{qJSD}}(\rho_i || \rho_j) = \mathcal{S}\left(\frac{\rho_i + \rho_j}{2}\right) - \frac{1}{2}\left(\mathcal{S}(\rho_i) + \mathcal{S}(\rho_j)\right). \quad (6)$$

Here we defined the Von Neumann entropy^{1,49} as

$$\mathcal{S}(\rho_i) = -\text{Tr}(\rho_i \log \rho_i), \quad (7)$$

where all logarithms are of natural base, and we utilize the convention $0 \log_2 0 := 0$.

One may ask why we chose to utilize the qJSD, and not other quantum entropic measures, such as the mutual information or the quantum relative entropy^{1,74}. Our deference to the qJSD is due to several useful properties (partially originating from the VNE), but arguably the most important one originates from the square root of the qJSD lies in a metric space $\mathcal{D}(x, y)$ for two objects x, y that we wish to distinguish. A metric space is endowed with the properties of:

- **Distance:** Let x, y, z be the elements inside a set X , then the function $\mathcal{D} : X \times X \rightarrow \mathbb{R}$ upholds $\mathcal{D}(x, y) \geq 0$, with the case of $\mathcal{D} = 0$ if $x = y$.
- **Symmetry:** The function $\mathcal{D}(x, y)$ also obeys $\mathcal{D}(x, y) = \mathcal{D}(y, x)$.
- **Adherence to the Triangle inequality:** lastly, $\mathcal{D}(x, y) + \mathcal{D}(x, z) \geq \mathcal{D}(y, z)$.

If these conditions are all upheld, we say that $\mathcal{D}(\cdot, \cdot)$ is a metric space⁴⁶. More specifically, the qJSD defines a bounded metric space of the form

$$0 \leq \sqrt{\mathcal{D}_{\text{qJSD}}(\rho_i || \rho_j)} \leq 1, \quad (8)$$

with a value of 0 signifying that $\rho_i = \rho_j$ and a value of 1 used for the case of $\rho_i \perp \rho_j$ ^{76,78}. As we are comparing the density matrices related to the IG and CG of a quantum circuit and processor, it is imperative to understand the closeness of one to the other, using some bounded distance measure. As a contrasting incentive, consider measuring the quantum relative entropy of two orthogonal states; in this case, the divergence is unbounded and gives $\mathcal{S}(\rho_i || \rho_j) \rightarrow \infty$ ¹. In the practical setting of the QCMP, such a measure is therefore not useful and does not convey the necessary distance information.

In addition to the metric space property, the qJSD is *symmetric*. This property is concomitant to the previous property related to metric spaces, but we address it here separately. Symmetry means that the qJSD obeys the relation

$$\mathcal{D}_{\text{qJSD}}(\rho_i || \rho_j) = \mathcal{D}_{\text{qJSD}}(\rho_j || \rho_i). \quad (9)$$

For the QCMP, we observe that this relation is desirable, as we wish for the notion of distance between two density matrices to stay the same, regardless of whether one is derived from L_{IG} or L_{CG} . As we will see in Section II E, it is in fact this distance quantity that we relate to the *quantum circuit uncomplexity*^{42,43}. Additionally, the concept of symmetry is paramount, as it permits us to directly relate the qJSD back to the *quantum Fisher information*; indeed, it was shown using the quantum Fisher information that the qJSD exactly calculates the *thermodynamic path length* between two equilibrium quantum states, and lower bounds their divergence on a Riemannian manifold^{55,56,74,75,79–81}.

Finally, we note that the qJSD is *non-increasing* under the action of a CP map⁷⁶, which can be formally stated as

$$\mathcal{D}_{\text{qJSD}}(\rho_i || \rho_j) \leq \mathcal{D}_{\text{qJSD}}(\Lambda(\rho_i) || \Lambda(\rho_j)), \quad (10)$$

where $\Lambda(\cdot)$ represents the superoperator of a *quantum operation*. The most general form of a quantum operation can be written in several

representations; in this work, we will concentrate on the Kraus representation (also known as the *operator-sum* representation), stated as

$$\Lambda(\cdot) := \sum_i E_i \cdot E_i^\dagger, \quad (11)$$

where E_i is the i^{th} term in the sum of operators, and $\Lambda(\cdot)$ is taken to be a general quantum operation superoperator, constrained to the *completely positive* (CP) condition^{1,35,36}.

We also introduce here the class of *doubly-stochastic* (DS) quantum channels, with the term *quantum channel* distinguishing from *quantum operation* in that, in addition to the CP constraint, we additionally impose *trace preservation* (TP)¹. In this work, we will refer to CPTP maps using $\Phi(\cdot)$. Moreover, doubly-stochastic quantum channels are *unital*, meaning that the fixed point of the channel upholds the equality $\Phi(\mathbb{I}_n) = \mathbb{I}_n$ ^{47,49}. In defining the class of doubly-stochastic quantum channels, we use the fact that any Kraus operator can be factorized, as all systems of Kraus operators implementing a quantum operation are related by a unitary transformation. A particular decomposition can be defined as

$$E_i = \sum_j \sqrt{\theta_j} P_j. \quad (12)$$

Here $P_j \in \mathbb{P}_n$ refers to permutation matrices from the set of $n \times n$ permutation matrices, and θ_j refers to a probability distribution⁴⁹ (we have also omitted the indices i on the right-hand side for clarity). The existence of this class of convex decomposition comes from the *Birkhoff-Von-Neumann Theorem*^{36,49,54} for which it is known that such a decomposition can be found in polynomial time⁸².

Lastly, we present *projective measurements* for density matrices constructed from graph Laplacians. Following the treatment of ref. 41, we define a set of orthogonal projectors Π_k such that $\sum_k \mathcal{M}_k = \mathbb{I}_n$. The post-measurement state of a general density matrix is then

$$\mathcal{M}(\rho) = \frac{\mathcal{M}_k \rho \mathcal{M}_k}{\text{Tr}[\mathcal{M}_k \rho]}, \quad (13)$$

where $\text{Tr}[\mathcal{M}_k \rho]$ represents the probability of the k^{th} measurement outcome. Note that projective measurements are known as a specific example of a CP map^{1,35,36}. In Section II E, we shall use projective measurements to erase subsystem interactions from the density-matrix form of the IG, ρ_{IG} , as well as for selecting appropriate subsystem permutations of the CG density matrix σ_{CG} .

Thermodynamic path length & (un)complexity

As mentioned in Section II C, the quantum Fisher information defines an entire family of statistical distance measures, from which we have taken particular interest in the family of entropic divergences. However, it is still not clear how to connect this to the more profound notion of *thermodynamic path length*. In order to provide an answer, let us start from quantum thermodynamics^{75,80,81}: the distance between any two quantum states (in continuum spacetime) can always be described as the average work extracted over some path through configuration space ξ :

$$W = \int_{\xi} dt \text{Tr} [\dot{H}_t \rho_t], \quad (14)$$

where we consider $H_t = \sum_i \lambda_i X_i$ as a time-dependent Hamiltonian with time-dependent, experimentally controllable parameters λ_i and time-independent observables X_i , and ρ_t is the time-evolved density matrix from $t \in [0, \tau]$, following the work of^{75,80,81}.

We also generally know how ρ_t evolves in spacetime via the Lindblad master equation^{80,81}, which is given by

$$\frac{d\rho}{dt} = i/\hbar[H, \rho(t)] + \mathcal{L}(\rho(t)), \tag{15}$$

where $\mathcal{L}(\rho(t)) = \sum_j \gamma_j (L_j \rho L_j^\dagger - 1/2(L_j^\dagger L_j \rho + \rho L_j^\dagger L_j))$, and $\rho_t, \rho(t)$ represent the time-evolved density matrix at time t , as well as the time-dependence of the density matrix on t , respectively. L_j are known as *jump operators*, and describe the channel that the quantum system is subjected to as it interacts with external environmental degrees of freedom⁸¹. γ_j are known as the *decoherence rates*.

In the case of the QCMP, this machinery is not needed, as we simply wish to understand the optimal case of quantum circuit mapping. That is to say, we wish to consider a noiseless quantum processor, capable of infinite parallelization (as we described above and in the manuscript’s discussion section). In that case, there are no environmental factors nor decoherence to consider, and we can examine our problem from the standpoint of a closed quantum evolution; therefore, we set $\gamma_j = 0$ and recover the original Von Neumann equation. As is expected for a closed system of pure quantum states, the dynamics now depend only on the Hamiltonian.

In order to calculate thermodynamic path length⁸¹, considers the amount of work dissipated into the environment due to restricted thermodynamic transformations on Gibbs states. One can directly find this from Equation (15), by optimizing the geodesic equations and accompanying Christoffel symbols for λ_i ⁸¹, ending with

$$W_{\text{diss}} = 1/\beta \int_{\xi} dt \dot{\lambda}_i(\mathbf{G}_{ij}) \dot{\lambda}_j, \tag{16}$$

Where β here represents the inverse temperature related to the Gibbs state, as is standard. From W_{diss} , we can formally define thermodynamic path length as

$$W_{\text{path}} = 1/\beta \int_{\xi} dt \sqrt{\dot{\lambda}_i(\mathbf{G}_{ij}) \dot{\lambda}_j}. \tag{17}$$

In these previous two equations, we recognize the quantum Fisher information \mathbf{G}_{ij} in the integral kernel. As we discussed previously in Section II C, it is known that thermodynamic path length and the Jensen-Shannon divergence compute the lower-bound distance between quantum states, and that this distance constitutes a geodesic in a configuration space of allowed transformations between equilibrium states^{56,57,73,75,79–81}. Lastly, we know that geodesics not only represent the shortest paths on a Riemannian manifold, but also that the distance between any two infinitesimally small intervals on the geodesic are locally the shortest path as well (i.e. the distance function we have defined must monotonically decrease along the thermodynamic path).

Before progressing, there are a few further points to mention. Firstly, when two CG vertices are adjacent to one another such that a corresponding IG two-qubit edge (gate) may be performed, we say that this edge is executable, and therefore should not factor further into the shortest-path calculation. Therefore, we compensate for this by performing *Bell measurements* on these edges, an operation already shown in⁴¹ to erase edges of simple graphs. We use the measurement scheme in quantum-operation form, as discussed at the end of Section II D. Secondly, in order to permute vertices on the CG, we make extensive utilization of the *doubly-stochastic* quantum channel forms which are also described in Section II D. However, we must make a slight modification due to the practical considerations of the QCMP. As we are limited to performing only nearest-neighbor SWAP gates on the CG, this signifies that, for the Kraus operators $E_i = \sum_j \sqrt{\theta_j} P_j$, we have that $P_j \in \mathbb{P}_n(\text{CG}) \subset \mathbb{P}_n$, where $\mathbb{P}_n(\text{CG})$ is the subgroup of all nearest-neighbor permutations available on the CG at a given time instant.

Thirdly and lastly, the operations of a quantum circuit take place over the discrete configuration space of $SU(2^k)$, with $k = 2, 3$ in most contexts, representing the number of qubits participating in a given gate⁴². Although it

is the case that $SU(2^k)$ is a Lie group and is therefore continuous, the permitted operations of the QCMP lie strictly within discrete configuration space, as we are restricted to only the set of Bell measurements $\{\mathcal{M}_e\}_{e \in E_{\text{IG}}}$ on the IG, and doubly-stochastic quantum operations $\{\Lambda(\cdot)|E_i \in \mathbb{P}_n(\text{CG})\}$ on the CG, both of which are represented by discrete simple graphs.

We can then discretize the integral over $\xi \in \{\{\mathcal{M}_e\}_{e \in E_{\text{IG}}}, \Lambda(\cdot)\}$ by considering infinitesimally small time translations $t + \Delta t$ with $\Delta < 1$, such that

$$\rho(t + \Delta t) \approx \mathcal{E}(\rho(t)), \tag{18}$$

where $\mathcal{E}(\cdot)$ is the action of a quantum channel acting on ρ . If we then perform this action m times, then we have

$$\rho(t + m(\Delta t)) \approx \mathcal{E}_m \circ \dots \circ \mathcal{E}_1(\rho(t)). \tag{19}$$

Keeping all of these points in mind, we can re-write the integral from Equation (16) in discretized form as

$$W_{\text{diss}} = \sum_m \Pi^l \left[\mathcal{M}_i^l \left[\arg \min_{\mathcal{G}_j} \left[\Pi^m \left[\Lambda_j^m(\mathbf{G}_{ij}) \right] \right] \right] \right], \tag{20}$$

where $\Pi^l[\mathcal{M}_i^l(\cdot)] = \mathcal{M}_i^l \circ \dots \circ \mathcal{M}_i^l(\cdot)$, i.e. the sequence of measurements executed on the IG when two-qubit gates are possible on the CG. Additionally, $\Pi^m(\Lambda_j^m(\cdot)) = \Lambda_j^m \circ \dots \circ \Lambda_j^m(\cdot)$, i.e. the sequence of SWAP-gate permutations undertaken in order to move qubits on the CG such that IG two-qubit gates can be performed. We sum over all of the permutations performed as we are erasing IG edges. Finally, as we recognize from^{56,73,79} and Equation (4) that the quantum Jensen-Shannon divergence is directed related to the quantum Fisher information matrix, we can directly substitute and obtain the form

$$W_{\text{diss}} = \sum_m \Pi^l \left[\mathcal{M}_i^l \left[\arg \min_{\mathcal{D}_{ij}^{\text{qIS}}} \left[\Pi^m \left[\Lambda_j^m(\mathcal{D}_{ij}^{\text{qIS}}(\rho_{\text{IG}} || \rho_{\text{CG}})) \right] \right] \right] \right], \tag{21}$$

where $\mathcal{D}_{ij}^{\text{qIS}}(\rho_{\text{IG}} || \rho_{\text{CG}}) = \mathcal{S}(\frac{\rho_{\text{IG}} + \rho_{\text{CG}}}{2}) - \frac{1}{2}(\mathcal{S}(\rho_{\text{IG}}) + \mathcal{S}(\rho_{\text{CG}}))$, and the subscripts i, j denote quantum operations on the IG and CG, respectively. We have also absorbed the terms λ_i and λ_j into the description of their respective quantum channels, as these terms represent time-dependent, externally-controllable parameters in the first place. The equation above can be likened to the process of *parallel transport* on a Riemannian manifold^{56,73}, and preserves the structure of the metric, as well as the geodesic form.

As we sum over all of the m permutations performed, we eventually erase all of the edges of the IG, resulting in an effective distance between the original CG and the maximally mixed state (now the erased IG):

$$\begin{aligned} W_{\text{diss}} &= ||\mathbb{I}_{\text{IG}} - \Pi^m[\Lambda_j^m(\rho_{\text{CG}})]||_{\mathcal{O} \in \{\{\mathcal{M}_e\}_{e \in E_{\text{IG}}}, \Lambda(\cdot)\}} \\ &= \mathcal{C}(\mathbb{I}_{\text{IG}}) - \mathcal{C}(\rho_{\text{CG}}) = \mathbf{U}_{\text{SWAP}}, \end{aligned} \tag{22}$$

where $||\cdot||_{\mathcal{O}}$ is a distance measure subject to the restrictions on transformations \mathcal{O} for transporting $\mathcal{D}_{ij}^{\text{qIS}}(\rho_{\text{IG}} || \rho_{\text{CG}})$ along the geodesic. It is obvious from the lower half equalities of Equation (22) that our equation exactly coincides with the form of *quantum circuit uncomplexity* given by⁴² and expounded upon in⁴³. Additionally, Equation (22) is directly related to the *quantum Wasserstein distance*, another known distance measure for calculating the shortest path between two quantum states in terms of number of gates within some restricted set of allowed transformations^{56,73}.

Benchmark Results

In this section, we describe the numerical results obtained from comparing the SWAP uncomplexity against IBM’s Qiskit compiler^{83*}, as well as against

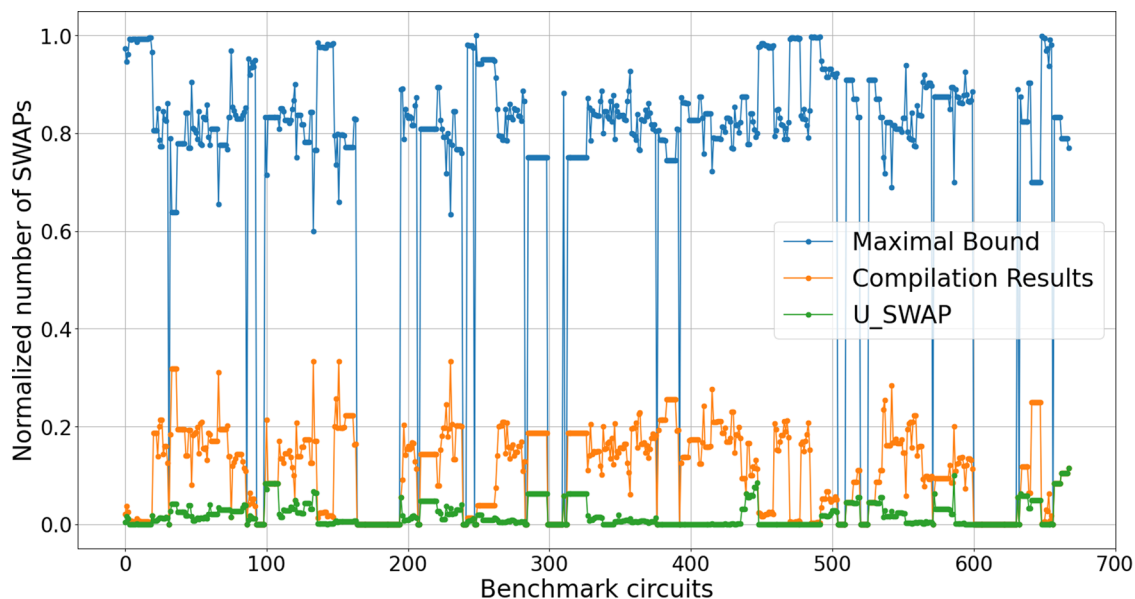


Fig. 2 | Simulation results for various IG / CG benchmark pairs. The horizontal axis enumerates each benchmark circuit tested (sorted by the number of two-qubit gates), and the vertical axis describes the normalized number of SWAPs, due to very high maximal bounds (the SWAP-bound values of each benchmark are divided by their sum). The results are color-coded as follows: the SWAP uncomplexity of

Section II E (green); the Qiskit compiler with default options Sabre router and circuit optimization level 1⁸³ (orange); and the maximal maximal bound calculated as in the Supplementary Material (blue). In every IG/CG pair, the bound calculated captures the SWAP uncomplexity that is either approachable or unattainable by the Qiskit compiler, thus empirically demonstrating our formulation.

a brute-force approach⁸⁴. These experiments were carried out for two main reasons. Firstly, we wish to subject the SWAP uncomplexity formalism and algorithm to a concrete, rigorous sanity check; after all, if the SWAP uncomplexity algorithm does in fact solve for the minimal SWAP-gate count, then the bounds we calculate should not be surpassed by any known compilation or brute-force optimization method in existence. By such logic, a compiler should be able to attain but not find fewer SWAP gates for an arbitrary IG / CG pair. In order to perform this empirical check, we chose to run our algorithm (Section IV A) against the Qiskit compiler, since it is considered to be the state-of-the-art approach at the moment. Secondly, to the best of our knowledge, there is scant literature on bounding required SWAP gates for an IG / CG pairing; at the moment, the latest work we are aware of addresses only up to quantum circuits of 6 qubits via a brute-force optimization algorithm^{84,85}. In contrast, our simulation results demonstrate scalability that greatly exceeds this brute-force optimization technique⁸⁴, as we achieved results for circuits of up to 16 qubits.

57 benchmark circuits were selected from the *qbench* suite⁵³. These benchmarks cover a range of 3 to 20 qubits and represent a wide spectrum of possible IG connectivities (47 different connectivities) encountered in quantum algorithms. More details about the selected benchmarks can be found in the Supplementary Material. As for the CGs, we chose connectivity graphs from a set of 16 in-use quantum devices, ranging from 5 to 72 qubits. The specific details of these devices are provided in the Supplementary Material. As some of the benchmarks are too large to be run on some of the smaller processors from our list, in total we devised 675 simulation experiments with the Qiskit compiler⁸³. In these simulations, we utilized Qiskit’s transpiler with the default circuit-optimization setting.

The results of our simulations are shown in Figs. 2-4. In Fig. 2, we display the normalized number of SWAP gates found by: the SWAP uncomplexity from Section II E (shown in green); the Qiskit compiler (denoted in orange); and the maximum SWAP-gate bound (depicted in blue). We observe clearly that the SWAP uncomplexity can be reached but never surpassed by the Qiskit compiler for select benchmark trials. As expected, the Qiskit compiler significantly outperforms the maximum SWAP-gate count calculated. Figure 3 also showcases the relation of the two-qubit gate count of the circuits (before compilation) and the bounds.

The results in this figure, however, do not only depend on the circuit complexity, but also on the coupling graph.

In order to more thoroughly scrutinize our results, we have included the relative graph-theoretic edge complexity for the benchmark circuits and have depicted them in Fig. 3. Here, results are plotted for only one device, the Google Bristlecone device, and for a circuit size of up to six qubits. Two subgraphs are shown, relating the normalized number of SWAP gates to two different measures: in a) the relation between the bounds and the number of two-qubit gates and in b) the relation between the bounds and the IG size of the circuit, shown as nodes-edges pairs with correspondingly small IG figures as guides for the reader. It is evident that, while the number of IG nodes or qubits has the biggest influence on the results, the number of edges and gates is also important. We refer the reader to the Supplementary Material in order to locate the circuit corresponding to points labeled on either of the horizontal axes of Fig. 3.

The non-triviality of the maximal and minimal SWAP-gate counts becomes evident in Fig. 4, where we present a covariance matrix with correlation coefficients ranging as $[-1, 1]$, with 0 indicating no correlation⁸⁶. This matrix compares the results that we obtained throughout the simulation; in particular, we compare the effective correlation between the SWAP uncomplexity; the Qiskit compiler SWAP calculation results; and the maximal bound as calculated in the Supplementary Material. Notably, both of our bounds (i.e., those obtained from our minimal bound with the SWAP uncomplexity algorithm, as well as the maximum SWAP-gate count) exhibit a substantial positive correlation (34% and 73%, respectively) with the actual results obtained from the compiler. The correlations here exemplify the non-triviality of the bounds; in other words, the SWAP uncomplexity and maximal bounds grow proportionally with the actual compilation results. The results at best coincide with each other, meaning that the lower bound equals the actual SWAP-gate count from the presence of a graph isomorphism; in this case, the SWAP uncomplexity, Qiskit result, and the maximal bound all obtain the same amount (which is zero if a graph isomorphism is present). These checks provide not only hard evidence for the usability of our methods, but additionally serve as a crucial sanity test that was passed for the algorithmic realization of the SWAP uncomplexity.

Although not shown in Fig. 4, it is also worth observing the considerable impact of the initial placement on the resulting bounds; this

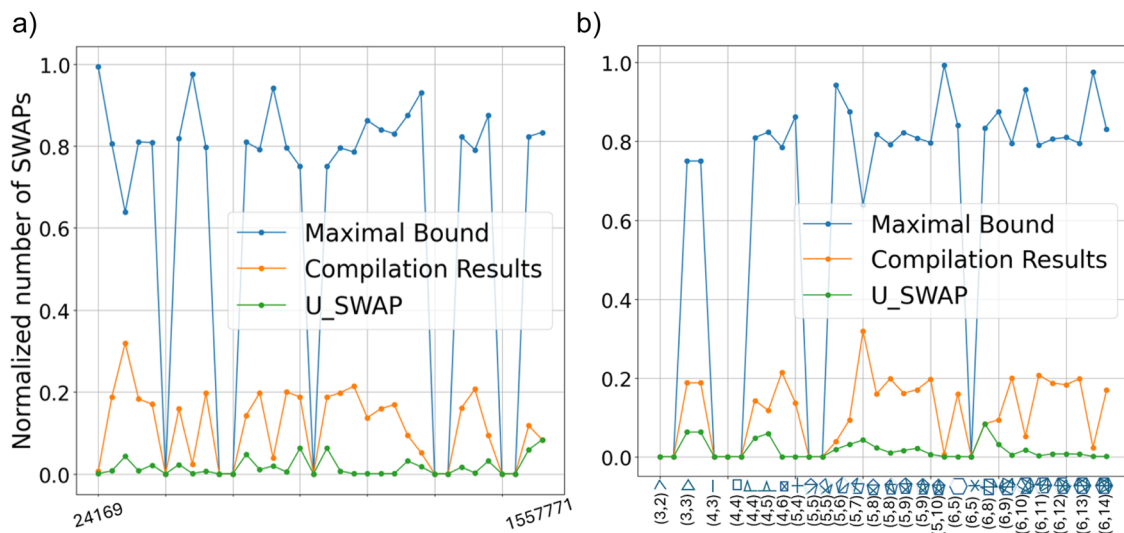


Fig. 3 | Subset of simulation results from Fig. 2 where only one CG is shown (Google Bristlecone). Benchmark circuits are sorted by a) number of two-qubit gates; and b) IG complexity (number of nodes and edges of IGs). The respective benchmarks with their respective nodes-edges count are detailed in the Supplementary Material.

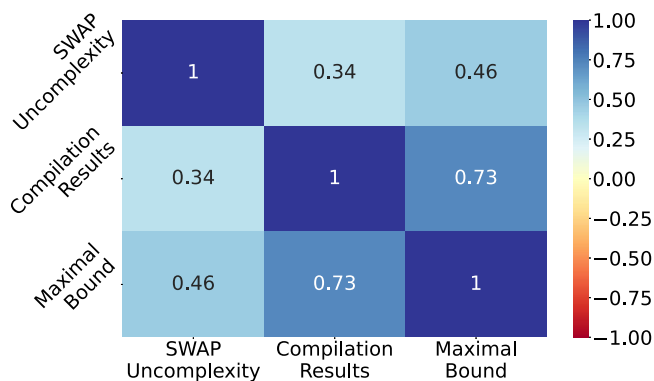


Fig. 4 | The Pearson correlation matrix⁸⁶ between the critical parameters measured for our benchmark investigation. The values range between - 1 and 1 for negative and positive correlation, respectively. When one of the parameters changes, the other one changes in the same direction. In this figure, we observe a high positive correlation between all the selected parameters where the Pearson correlation coefficient ranges between 0.34 and 1.0.

particularly depends on the GED and the number of missing edges in the chosen CG partition compared to the IG. We therefore calculated the correlation coefficients of these two parameters (as well as the case when compared to our retrieved bounds), resulting in correlations of 79% and 61% for the SWAP uncomplexity and maximal bounds, respectively. Using the same initial placement for the Qiskit compiler resulted in a 45% correlation with the parameters related to initial placement.

The qubit-assignment strategy was initially tested with 729 benchmarks, showing a success rate of 92.6%. The remaining 7.4% of the benchmarks could not be finished due to insufficient computing resources. Recognizing the limited scalability of the approach (up to 16 circuit qubits), we developed a more relaxed method for complete graphs, mentioned in Section IV B. Indeed, the scalability of our exact algorithm already exceeded that of the exact state-of-the-art algorithms, which struggled beyond 6 qubits⁵¹. Furthermore, our initial placement encountered no difficulties in exploring a vast search space. It successfully executed circuits on all tested devices, extending up to a size of 72 physical qubits in our case.

Lastly, Fig. 5 shows a comparison between our SWAP uncomplexity algorithm (Section IV A) and the brute-force optimization results from⁸⁴. In this approach, the authors utilize an optimizer, which essentially tries every




Fig. 5 | The Pearson correlation matrix⁸⁶ between the critical parameters measured for our benchmark investigation. The values range between - 1 and 1 for negative and positive correlation, respectively. When one of the parameters changes, the other one changes in the same direction. In this figure, we observe a high positive correlation between all the selected parameters where the Pearson correlation coefficient ranges between 0.34 and 1.0.

permutation of SWAP placements possible while respecting the gate-dependency graph and weighted IG of the original quantum circuit. In all cases, we see clearly that the brute-force algorithm only achieves but never surpasses the SWAP uncomplexity bound.

Discussion

It is known that the QCMP is NP-complete¹⁴. As such, we have made three simplifications in order to derive the SWAP uncomplexity. Firstly, we do not consider single-qubit interactions, as it is known that such gates do not heavily affect calculated success rates¹³. Secondly, we have removed all two-qubit interaction noise from the CG; this should come as no surprise, as we are mainly interested in finding a lower bound for the number of SWAP gates required, and such a lower bound mandates the existence of a hypothetically noiseless quantum processor. Thirdly, we consider the limit in

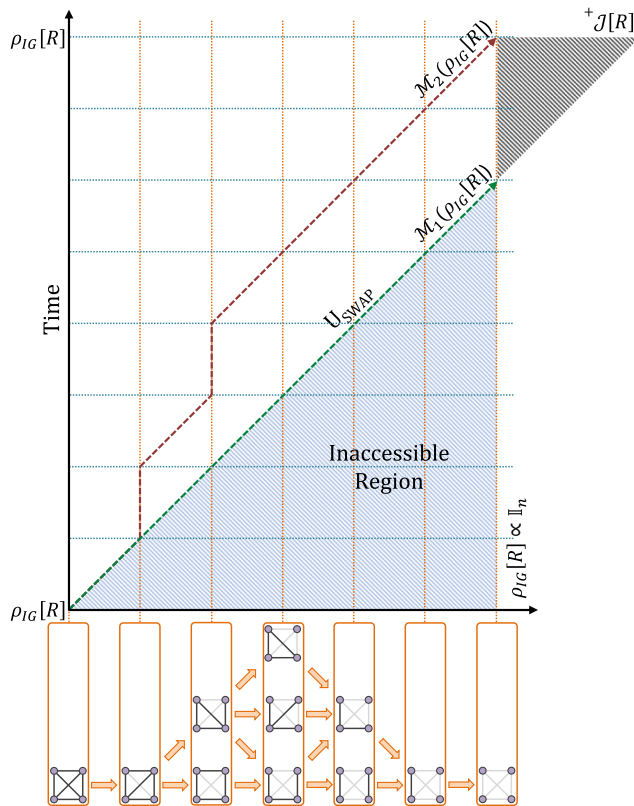


Fig. 6 | A Penrose diagram representing the quantum circuit complexity for the evolution of a quantum state, which is related to a quantum circuit ρ_{IG} (for the sake of simplicity, we choose the state ρ_{IG} to be derived from the K_4 graph, but one can choose other examples). It is possible to generate different quantum complexities by adding different amounts and orderings of SWAP gates to the circuit as we approach the event horizon of a black hole (shown as a shaded triangle). Here, the connectivity limitations of the CG capture the role of the background spacetime geometry⁸⁷, as both determine the ease by which certain operations can be performed. The SWAP uncomplexity, as it ignores the effects of time ordering and the amount of qubit-qubit interaction present, can be associated with the null lightlike geodesic $^+ \mathcal{J}[R]$, shown as a green dashed line. Below this arrow, the bottom-right part of the diagram represents possible states that are inaccessible to us, given the restrictions of the CG, as well as the operations available to us (SWAP gates, in our case).

which gate dependencies for the IG are not considered; implicitly, we assume the existence of not only a noiseless quantum processor, but additionally one that can perform all two-qubit gate interactions required in one unit time slice, i.e. one whose two-qubit gate operations are infinitely parallelizable.

These concepts can be related using *Penrose diagrams*^{87,88}, as shown in Fig. 6. A Penrose diagram typically shows the causal structure of events unfolding in a *spacetime geometry*⁸⁷. In Fig. 6, the horizontal axis refers to purely spatial evolutions, which in our case are shown by potential SWAP erasures. Each of the sets outlined in sky blue represent elements of the same total number of edges, but different spectral properties. Each graph within a given set represents a unique spectral signature which can be shared by multiple four-node subgraphs. The vertical axis depicts the evolution of time, and is known as the *null time geodesic*, under which the set of *trivial time-ordered gate operations* (i.e. idling, which in our simplified picture, is noiseless) evolve ρ_{IG} from point R to the same later state. Here, the trivial minimal SWAP-gate count for the QCMP is represented, under which no SWAP gates are ever applied in order to adapt the quantum circuit to the device and its connectivity restrictions; in effect, the state is left to freely evolve for an infinite amount of time, with no regards to gate operations.

Conversely, the red dashed line represents possible spacetime evolutions arising from the application of distinct instances of time-ordered quantum operations $\mathcal{M}_1(\cdot), \mathcal{M}_2(\cdot)$ (which in this case are insertions of SWAP gates and subsequent erasures from the IG). In this way, every trajectory on the diagram can be associated with a given SWAP uncomplexity from a sequence of quantum operations. The possible endpoints of the quantum circuit are shown along the shaded triangle in dark gray, which represents $\rho_{IG} = \mathbb{I}_n$, i.e. the state of maximal circuit complexity, a maximally mixed state.

Furthermore, erasure transformations on the original quantum circuit proceed according to restrictions dictated by the background geometry (which in our present case is analogous to the CG connectivity). Evolution commences at a spacetime point R . The green dashed line traces out the *lightlike null geodesic* (i.e. future lightcone) $^+ \mathcal{J}[R]$. This geodesic signifies the SWAP uncomplexity, which is the path that the state takes under the minimal set of causally indefinite operations such that we approach the maximally mixed state in minimal time. As we solve for the SWAP uncomplexity, without consideration of time ordering and as dictated by the thermodynamic path length calculable via the Fisher information^{75,79–81}, we can interpret our bound as a sort of *lightcone evolution* of our initial state ρ_{IG} towards the event horizon of a black hole (shown as the shaded triangle). Previous work has already alluded to the concept of optimization over thermodynamic distance^{81,89}; as such, our results point to a natural and reasonable extension of this trend for the quantum circuit mapping problem.

As we touched upon earlier, it is possible to interpret the shaded triangle in Fig. 6 as the event horizon of a black hole. Consider a benign black hole scenario in which the black hole itself can only erase information from a density matrix in accordance with only certain SWAP gates from some constrained architecture (i.e., the black hole itself is described with respect to a background geometry, which constrains which operations can be performed). As black holes are known to be the fastest information scramblers in nature⁹⁰, the QCMP can be viewed through the lens of a scrambling process, yielding the most-efficient method to maximally mix the information of the IG’s density matrix for a given β . This process exemplifies the traits of quantum circuit uncomplexity and we have shown that the quantity U_{SWAP} can in fact be optimized for using our technique.

Taking stock, we would then expect that any realistic quantum compiler which takes into account time ordering and finite qubit-qubit interactions per unit of time slicing to be limited by the lightlike null geodesic. Indeed, surpassing the lightlike null geodesic would introduce operations, which are not inside of the lightcone, giving access to the uncomputable region to the bottom-right. Such trajectories could be made possible using a larger set of routing resources, such as teleportation^{62,63}. As a concrete counterexample, consider a hypothetical quantum compiler, which could surpass U_{SWAP} . One of the main assumptions that we utilize in our formulation above is that possible interactions can occur if and only if a subsystem interaction between qubits exists in the density matrix picture. Consequently, moving outside the region embellished by the future lightlike geodesic corresponds to new operations which must be taken into account. One simple example lies in *teleportation-based* quantum circuit mapping, which can be used to swap CG qubits, which do not share a physical subsystem interaction, and can allow for a smaller quantum circuit complexity⁶². From the standpoint of our formalism, this difference would correspond to allowing for non-nearest-neighbor permutation matrices to arise in the doubly-stochastic quantum channel described in Equation (12). One may suspect that architectures in the future may benefit from such on-chip teleportation procedures, as work has shown that speedups exist over classical SWAP methods for exchanging distant qubits^{32,62,63}, albeit with larger circuit and entanglement overhead.

The QCMP itself has been described using several approaches from computer science, many of which have allowed for the development of entirely new strategies for solving the problem. Our contribution here serves a different purpose. At the level of theoretical physics, as well as quantum information theory, we have solved a simplified subproblem of the QCMP,

which we dub the “lightcone bound” to the QCMP. We have shown that solving for the lightcone SWAP uncomplexity bound is optimal, in the sense that it implicitly defines the shortest path through a configuration space of restricted gate operations. At precisely the SWAP uncomplexity limit, it is expected that: the quantum device is noiseless; the two-qubit gate interactions given by the device can be performed with an indefinite causal order; and that any number of two-qubit interactions can be performed in parallel within one unit time slice. Therefore, we surmise that no quantum compiler in existence can violate this lower bound, allowing for a fundamental means of comparison between differing strategies for solving the real-world QCMP. We have provided, to the best of our knowledge, the first instance of a solvable lower bound for SWAP-gate count in the context of quantum circuit compilation. The SWAP uncomplexity was derived using tools from graph theory, quantum information theory, quantum circuit complexity, and information geometry. In addition to the use case discussed in this work, potential applications of uncomplexity for quantum machine learning are discussed in⁹¹. This work also represents the first application of quantum circuit uncomplexity to the realm of practical quantum information processing.

Of independent interest may be the qubit-assignment algorithm which was designed to aid in the calculation of the SWAP uncomplexity. This algorithm, grounded as a graph similarity search, inspects distinct n -qubit partitions of a given CG, and returns the most-similar resultant to the IG provided. Employing this method has enabled us to map circuits with up to 16 qubits onto devices with up to 72 physical qubits. For larger circuits, we devised an alternative approach. The initial placement precedes the minimal SWAP-gate count solution, which is further utilized for routing and minimal bound calculation. Additionally, we calculated a maximal bound by leveraging known classical graph metrics; both of these novel structures provide additional tools of interest outside of the scope of this work.

We would now like to draw attention to several open problems regarding our work, as well as several future possible directions for research:

1. *Improvements to the subgraph similarity search algorithm.* In this work, although our qubit-assignment algorithm outperforms the current state of the art solver⁸⁵, we were still limited by the scalability of the qubit-assignment algorithm constructed in Section IV B. However, once a suitable qubit assignment is set, the calculation of U_{SWAP} can be completed in $\sim \mathcal{O}(\dim(\rho_{\text{CG}})^4)$ timesteps, as per the Birkhoff-Von-Neumann algorithm⁸². A future research goal could involve making further scalability improvements to the qubit-assignment algorithm, or by considering more advanced methods of routing, such as those using *ancillary qubits*²⁵.
2. *Searching for the optimal β value.* In our work, we have taken a somewhat naive approach to optimizing for β ; however, because of the similarity to a phase diagram, one may be able to use concepts from condensed matter theory^{92–94} in order to devise a suitable gradient-based optimization method.
3. *Analytical expression for tightness of the SWAP uncomplexity to the brute-force solution.* We have given empirical evidence for tightness, but it still remains to define an analytical expression for how similar in general our solution is, compared to the brute-force solution proposed in⁸⁴, and how tightness scales as the size of the quantum circuit to be mapped increases in both register and depth.
4. *Extension to incorporate bridge gates, teleportation-based quantum circuit mapping, and shuttling.* There are other methods commonly in use, in addition to the SWAP gate, for conforming a quantum circuit to hardware. Our approach is extendable for the Bridge gates mentioned in^{95–99}, as well as the quantum teleportation-based protocols of^{32,62,63,100} and shuttling-based approaches for spin-qubit architectures^{101,102}, trapped-ion architectures^{103,104}, and neutral-atom devices¹⁰⁵.
5. *Extension for quantum error correction codes, in particular syndrome extraction circuits.* It is well-known that various *fault-tolerance* protocols are required in order to ensure that quantum error correction codes function up to their full code distance^{106–117}. As our bound constitutes a non-trivial resource requirement, it may be useful

to adapt fault-tolerance protocols further to the setting of quantum compilation, in which an error correction code is adapted to a device not specifically designed for a particular code family^{118,119}.

6. *Extension for entanglement/qubit routing in quantum communications networks and modular architectures.* Several other extensions may be possible as well, including those allowing for bounds on the QCMP for modular scenarios¹²⁰ as well as for *entanglement distribution* in noisy quantum networks¹²¹.

Finally, we remark that the problem of assessing similarities between two complex networks is a problem spanning many disciplines. Indeed, our work follows recent trends of utilizing quantum information theory and statistical mechanics to study complex networks^{39,40,44,122}. As the task of comparing the distance between graphs appears in many different areas of science^{70,123}, we expect the implications of our work to stretch beyond the realm of quantum information science.

Methods

Algorithmic implementation

The pseudocode for calculating the SWAP uncomplexity is shown in Algorithm 1. The algorithm proceeds similarly to the mathematical derivation detailed in Section II E. Firstly, the density matrices ρ_{IG} and ρ_{CG} are provided as inputs, and U_{SWAP} is set to zero. Next, we immediately calculate the qJSD in order to check if an isomorphism exists between ρ_{IG} and ρ_{CG} . If none exists, then we first remove all of the edges of the IG, which directly match up with edges of the CG, obtaining the density matrix $\bar{\rho}_{\text{IG}}$. We then set a SWAP-gate counter m to zero. Afterwards, a `FOR` loop begins with the eventual goal to erase all of the IG’s edges; the process by which this happens begins with the calculation of a first qJSD $\text{qj} \text{sd}_1$ after the i^{th} and m^{th} actions of measurement and doubly-stochastic quantum channels on their respective density matrices. Additionally, we calculate a second qJSD with an extra permutation applied to the CG. The optimal choice of this particular permutation requires a worst-case search over all edges of CG for each iteration. In practice, decomposing a doubly-stochastic quantum channel will result in the superposition of several possible permutation matrices^{49,54}; in order to make a hard decision, we choose to apply the permutation matrix with the maximal θ_j value, as shown in Equation (12). The reason for choosing the maximal θ_j lies in the fact that performing the most-likely permutation matrix at every iteration step of the algorithm allows us to follow and stay on the geodesic at every time step^{79,81}. After applying the permutation matrix, we compute the second qJSD $\text{qj} \text{sd}_2$; if it is found that $\text{qj} \text{sd}_2 < \text{qj} \text{sd}_1$, then we simply add one to the SWAP counter and the same process of comparing subsequent qJSDs continues until $\text{qj} \text{sd}_2 \geq \text{qj} \text{sd}_1$. Upon arriving here, we first check to see if the current iteration of $\mathcal{M}^i(\bar{\rho}_{\text{IG}})$ is equivalent to the identity matrix \mathbb{I}_{IG} ; in this case, the algorithm is complete and we break out of the `FOR` loop, returning the number m associated to the SWAP uncomplexity U_{SWAP} . If $\mathcal{M}^i(\bar{\rho}_{\text{IG}}) \neq \mathbb{I}_{\text{IG}}$, then we continue by performing the next subsequent measurement, $\mathcal{M}^{i+1}(\bar{\rho}_{\text{IG}})$, associated with whichever edges are currently matched up between the IG and the CG.

At this point, the erasure of a remaining subsystem interaction implies that the qJSD will again increase, as we know that the VNE under a CP map always increases¹²⁴. We must then perform the commensurate doubly-stochastic quantum channel operation(s) again and select the appropriate θ_j -valued permutation such that the divergence decreases to its minimal value once more. The algorithm terminates upon the successful erasure of all subsystem interactions in $\bar{\rho}_{\text{IG}}$, leaving a maximally mixed state.

One may ordinarily surmise that the runtime complexity of Algorithm 1 is quite high; after all, inside the `FOR` loop lies several seemingly difficult optimization problems. However, due to the *Birkhoff-Von-Neumann algorithm*, decomposition of any doubly-stochastic quantum channel is guaranteed in polynomial timesteps^{49,82}. Taking stock, we conclude that the algorithm’s runtime complexity is bounded by the number of edges in the IG, multiplied by the number of edges in the CG queried by an optimizer to determine the maximum- θ_j permutation for estimating $\text{qj} \text{sd}_2$, using Λ^{m+1} .

Since $|\bar{E}_{IG}| \leq |\bar{E}_{CG}|$, the worst case complexity is quadratic in the number of edges in a complete graph of size CG, or more succinctly, $O((\dim(\rho_{CG})(\dim(\rho_{CG}) - 1)/2)^2) = O(\dim(\rho_{CG})^4)$. However, for the optimization loop, since we need to consider edges within the connectivity constraints of quantum processors, these graphs are typically planar instead of all-to-all connected (i.e. as in a complete graph), with much more benign runtime expectation for the pragmatic use case.

Algorithm 1. Pseudocode for an algorithmic optimization of the SWAP Uncomplexity U_{SWAP} .

```

Input:  $\rho_{IG}, \rho_{CG}$ 
Initial_Qubit_Assignment  $\leftarrow$  Qubit_Assignment( $\rho_{IG}, \rho_{CG}$ )
 $U_{SWAP} \leftarrow 0$ 
Output:  $U_{SWAP}$ 
Assert:  $\mathcal{D}_{qjs}(\rho_{IG} || \rho_{CG}) == 0$ 

if  $\mathcal{D}_{qjs}(\rho_{IG} || \rho_{CG}) \neq 0$  then
   $\bar{\rho}_{IG} \leftarrow$  Remove_Trivial_Edges( $\rho_{IG}$ )
   $m \leftarrow 0$ 
  for  $i \in [0, |\bar{E}_{IG}|]$  do
     $qj\text{sd}_1 \leftarrow \mathcal{D}_{qjs}(\mathcal{M}^i(\bar{\rho}_{IG}) || \Lambda^m(\rho_{CG}))$ 
     $qj\text{sd}_2 \leftarrow \mathcal{D}_{qjs}(\mathcal{M}^i(\bar{\rho}_{IG}) || \Lambda^{m+1}(\rho_{CG}))$ 
    if  $qj\text{sd}_2 < qj\text{sd}_1$  then
       $m \leftarrow m + 1$ 
      continue
    else
      if  $\mathcal{M}^i(\bar{\rho}_{IG}) \equiv \mathbb{I}_{IG}$  then
        break
      else
         $\mathcal{M}^i(\bar{\rho}_{IG}) \leftarrow \mathcal{M}^{i+1}(\bar{\rho}_{IG})$ 
      end if
    end if
  end for
end if
return  $U_{SWAP} \leftarrow m = U_{SWAP}$ 

```

Qubit Assignment

As mentioned previously in Section IIA, the initial stage of the QCMP is known as *qubit assignment* (also known as initial placement, qubit allocation, or initial mapping)^{51,52}. This procedure plays a pivotal role in quantum circuit execution¹². In our proposal for calculating the SWAP uncomplexity in the QCMP, we also must assign qubits from the IG to the CG initially in an optimal way, as this influences how many SWAP gates will be utilized.

In⁵¹, the concept of qubit assignment was introduced as a search for a subgraph isomorphism for an IG/CG pair. To our knowledge, this technique has not yet seen widespread implementation in practical qubit-assignment techniques, despite its potential. Instead, most existing approaches focus on alternative factors such as sequential gate flow in the circuit or the number of interactions between qubits as in^{20,125,126}. Nevertheless, some work has explored the subgraph-isomorphism concept for the QCMP^{28,52,127,128}.

Building upon the foundation of the well-known VF2 algorithm^{28,129}, our approach to qubit-assignment searches for an exact location on the quantum device where our circuit can run without requiring additional gates. If a solution is feasible, we are left with an optimal assignment. In cases where a solution is not possible, we conduct a graph similarity search. This process involves the GED calculation and comparison of the IG to all distinct subgraphs of the same size within the CG, which opens up alternative assignment possibilities. In this fashion, we condense the search space for alternative solutions, while also highlighting the potential utility of our approach for *multi-programming* applications (i.e. executing multiple circuit in parallel on a quantum device)⁶¹.

Let $|V_{IG}|$ be the number of qubits in the IG, and $|V_{CG}|$ be the number of physical qubits on the IG and CG, respectively. Our qubit-assignment

process consists of the following steps, which are described diagrammatically in the Supplementary Material:

1. *Preprocessing:*
 - (a) Select a quantum algorithm described as a quantum circuit and extract its IG $G_{IG}(E_{IG}, V_{IG})$, where $|E_{IG}|$ represents the number of edges in the IG.
 - (b) Choose a quantum device to execute the circuit on and extract its CG, represented as graph $G_{CG}(E_{CG}, V_{CG})$, where $|E_{CG}|$ stands for the number of edges in the CG.
 - (c) In order to increase the efficiency of steps later on, and reduce the search space, we find all distinct subgraphs of size $|V_{IG}|$ within graph G_{CG} .
2. *Subgraph isomorphism using VF2 and subgraph similarity search:*
 - (a) Use the VF2 algorithm to check if a subgraph isomorphism exists between graphs G_{IG} and G_{CG} .
 - i. If a subgraph isomorphism is found, we immediately determine the location within the CG for qubit assignment.
 - ii. When a subgraph isomorphism does not exist, we utilize the *graph-edit distance* (GED) to identify structurally most similar subgraph of the CG when compared to the IG. During this process, we compare IGs only to distinct subgraphs of a CG derived from Item 1c.
 - (b) Assign the IG to the CG in accordance with the result from the previous step.
3. *Calculating the maximal SWAP-gate count* as it depends on the qubit assignment. We describe the computation of this bound in more detail in the Supplementary Material.

Data availability

The software developed for this project is available at <https://github.com/QML-Group/QCMP-complexity-bound>.

Code availability

The software developed for this project is available at <https://github.com/QML-Group/QCMP-complexity-bound>.

Received: 22 February 2024; Accepted: 20 October 2024;

Published online: 09 November 2024

References

1. Nielsen, M. & Chuang, I. Quantum Computation and Quantum Information. ISBN 978-1-107-00217-3 (Cambridge University Press, Cambridge, UK, 2010).
2. Preskill, J. Quantum computing in the nisq era and beyond. *Quantum* **2**, 79 (2018).
3. Hempel, C. et al. Quantum chemistry calculations on a trapped-ion quantum simulator. *Phys. Rev. X* **8**, 031022 (2018).
4. O'Malley, P. J. J. et al. Scalable quantum simulation of molecular energies. *Phys. Rev. X* **6**, 031007 (2016).
5. Cory, D. G. et al. Experimental quantum error correction. *Phys. Rev. Lett.* **81**, 2152–2155 (1998).
6. Chiaverini, J. et al. Realization of quantum error correction. *Nature* **432**, 602–605 (2004).
7. Terhal, B. M. Quantum error correction for quantum memories. *Rev. Mod. Phys.* **87**, 307 (2015).
8. Murali, P. et al. Full-stack, real-system quantum computer studies: Architectural comparisons and design insights. In *Proceedings of the 46th International Symposium on Computer Architecture*, 527–540 (2019).
9. Bandic, M., Feld, S. & Almudever, C. G. Full-stack quantum computing systems in the nisq era: algorithm-driven and hardware-aware compilation techniques. In *2022 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 1–6 (IEEE, 2022).
10. Bertels, K. et al. Quantum computer architecture toward full-stack quantum accelerators. *IEEE Trans. Quantum Eng.* **1**, 1–17 (2020).

11. Bandic, M., Zarein, H., Alarcon, E. & Almudever, C. G. On structured design space exploration for mapping of quantum algorithms. In *2020 XXXV conference on design of circuits and integrated systems (DCIS)*, 1–6 (IEEE, 2020).
12. Li, G., Ding, Y. & Xie, Y. Tackling the qubit mapping problem for nisq-era quantum devices. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '19*, 1001–1014 (Association for Computing Machinery, New York, NY, USA, 2019). <https://doi.org/10.1145/3297858.3304023>.
13. Siraichi, M. Y., Santos, V. F. d., Collange, C. & Pereira, F. M. Q. Qubit allocation. In *Proceedings of the 2018 International Symposium on Code Generation and Optimization, CGO 2018*, 113–125 (Association for Computing Machinery, New York, NY, USA, 2018). <https://doi.org/10.1145/3168822>.
14. Cowtan, A. et al. On the qubit routing problem. In *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019)* (Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019).
15. Alon, N., Chung, F. R. & Graham, R. L. Routing permutations on graphs via matchings. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of Computing*, 583–591 (1993).
16. Paler, A., Zulehner, A. & Wille, R. Nisq circuit compilation is the travelling salesman problem on a torus. *Quant. Sci. Technol.* **6**, 025016 (2021).
17. Miltzow, T. et al. Approximation and hardness for token swapping. *arXiv preprint arXiv:1602.05150* (2016).
18. Wagner, F., Bärmann, A., Liers, F. & Weissenböck, M. Improving quantum computation by optimized qubit routing. *Journal of Optimization Theory and Applications* 1–34 (2023).
19. Murali, P., Baker, J. M., Javadi-Abhari, A., Chong, F. T. & Martonosi, M. Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers. In *International Conference on Architectural Support for Programming Languages and Operating Systems*, 1015–1029 (2019).
20. Tannu, S. S. & Qureshi, M. K. Not all qubits are created equal: A case for variability-aware policies for NISQ-era quantum computers. In *International Conference on Architectural Support for Programming Languages and Operating Systems*, 987–999 (2019).
21. Li, G., Ding, Y. & Xie, Y. Towards efficient superconducting quantum processor architecture design. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, 1031–1045 (2020).
22. Zulehner, A., Paler, A. & Wille, R. An efficient methodology for mapping quantum circuits to the IBM QX architectures. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* (2018).
23. Venturelli, D. et al. Quantum circuit compilation: An emerging application for automated reasoning. In *Scheduling and Planning Applications Workshop* <https://openreview.net/forum?id=S1eEBO3nFE> (2019).
24. Lao, L. et al. Mapping of lattice surgery-based quantum circuits on surface code architectures. *Quant. Sci. Technol.* **4**, 015005 (2019).
25. Lao, L., Van Someren, H., Ashraf, I. & Almudever, C. G. Timing and resource-aware mapping of quantum circuits to superconducting processors. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **41**, 359–371 (2019).
26. Herbert, S. & Sengupta, A. Using reinforcement learning to find efficient qubit routing policies for deployment in near-term quantum computers. *arXiv:1812.11619* (2018).
27. Lye, A., Wille, R. & Drechsler, R. Determining the minimal number of swap gates for multi-dimensional nearest neighbor quantum circuits. In *Asia and South Pacific Design Automation Conference*, 178–183 (2015).
28. Li, S., Zhou, X. & Feng, Y. Qubit mapping based on subgraph isomorphism and filtered depth-limited search. *IEEE Trans. Comput.* **70**, 1777–1788 (2020).
29. Biuki, A., Mohammadzadeh, N., Wille, R. & Sargaran, S. Exact mapping of quantum circuit partitions to building blocks of the saqip architecture. In *2022 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, 402–405 (IEEE, 2022).
30. Molavi, A. et al. Qubit mapping and routing via maxsat. In *2022 55th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 1078–1091 (IEEE, 2022).
31. Moro, L., Paris, M. G., Restelli, M. & Prati, E. Quantum compiling by deep reinforcement learning. *Commun. Phys.* **4**, 178 (2021).
32. Devulapalli, D., Schoute, E., Bapat, A., Childs, A. M. & Gorshkov, A. V. Quantum routing with teleportation 2204.04185 (2022).
33. Upadhyay, S., Saki, A. A., Topaloglu, R. O. & Ghosh, S. A shuttle-efficient qubit mapper for trapped-ion quantum computers. In *Proceedings of the Great Lakes Symposium on VLSI 2022*, 305–308 (2022).
34. Nottingham, N. et al. Decomposing and routing quantum circuits under constraints for neutral atom architectures. *arXiv preprint arXiv:2307.14996* (2023).
35. Wilde, M. M. *Quantum Information Theory* (Cambridge University Press, 2017), 2 edn.
36. Watrous, J. *The theory of quantum information* (Cambridge University Press, 2018).
37. Di Meglio, A. et al. Quantum computing for high-energy physics: State of the art and challenges. *PRX Quantum* **5**, 037001 (2024).
38. Harlow, D. Jerusalem lectures on black holes and quantum information. *Rev. Mod. Phys.* **88**, 015002 (2016).
39. De Domenico, M. & Biamonte, J. Spectral entropies as information-theoretic tools for complex network comparison. *Physical Review X* **6** <https://doi.org/10.1103/PhysRevX.6.041062> (2016).
40. Faccin, M. & Biamonte, J. Complex networks from classical to quantum. *Nature Communications Physics* **2** <https://doi.org/10.1038/s42005-019-0152-6> (2019).
41. Braunstein, S. L., Ghosh, S. & Severini, S. The laplacian of a graph as a density matrix: A basic combinatorial approach to separability of mixed states. *Ann. Combinatorics* **10**, 291–317 (2006).
42. Brown, A. R. & Susskind, L. Second law of quantum complexity. *Phys. Rev. D* **97**, 086015 (2018).
43. Yunger Halpern, N. et al. Resource theory of quantum uncomplexity. *Phys. Rev. A* **106**, 062417 (2022).
44. Glos, A., Krawiec, A. & Pawela, Ł. Asymptotic entropy of the gibbs state of complex networks. *Sci. Rep.* **11**, 1–9 (2021).
45. Briët, J. & Harremoës, P. Properties of classical and quantum jensen-shannon divergence. *Phys. Rev. A* **79**, 052311 (2009).
46. Lamberti, P. W., Majtey, A. P., Borrás, A., Casas, M. & Plastino, A. Metric character of the quantum jensen-shannon divergence. *Phys. Rev. A* **77**, 052311 (2008).
47. Wolf, M. Quantum channels and operations: A guided tour <https://www-m5.ma.tum.de/foswiki/pub/M5/Allgemeines/MichaelWolf/QChannelLecture.pdf> (2012).
48. Christandl, M. Quantum information theory lecture notes (2018).
49. Nielsen, M. An introduction to majorization and its applications to quantum mechanics: Lecture notes <https://michaelnielsen.org/blog/talks/2002/maj/book.ps> (2002).
50. Steinberg, M. A., Feld, S., Almudever, C. G., Marthaler, M. & Reiner, J.-M. Topological-graph dependencies and scaling properties of a heuristic qubit-assignment algorithm. *IEEE Trans. Quantum Eng.* **3**, 1–14 (2022).
51. Siraichi, M. Y., Santos, V. F. d., Collange, C. & Pereira, F. M. Q. Qubit allocation. In *Proceedings of the 2018 International Symposium on Code Generation and Optimization*, 113–125 (2018).
52. Siraichi, M. Y., Santos, V. F. D., Collange, C. & Pereira, F. M. Q. Qubit allocation as a combination of subgraph isomorphism and token swapping. *Proc. ACM Program. Lang.* **3**, 1–29 (2019).
53. Bandic, M., Almudever, C. G. & Feld, S. Interaction graph-based characterization of quantum benchmarks for improving quantum circuit mapping techniques. *Quantum Mach. Intell.* **5**, 40 (2023).

54. Sagawa, T. Entropy, divergence, and majorization in classical and quantum thermodynamics 2007.09974 (2020).
55. Bengtsson, I. & Życzkowski, K. *Geometry of quantum states: an introduction to quantum entanglement* (Cambridge university press, 2017).
56. Datta, N. & Rouzé, C. Relating relative entropy, optimal transport and fisher information: a quantum hwi inequality. In *Annales Henri Poincaré*, vol. 21, 2115–2150 (Springer, 2020).
57. Liu, J., Yuan, H., Lu, X.-M. & Wang, X. Quantum fisher information matrix and multiparameter estimation. *J. Phys. A: Math. Theor.* **53**, 023001 (2020).
58. Zager, L. A. & Verghese, G. C. Graph similarity scoring and matching. *Appl. Math. Lett.* **21**, 86–94 (2008).
59. Koutra, D., Parikh, A., Ramdas, A. & Xiang, J. Algorithms for graph similarity and subgraph matching. In *Proc. Ecol. inference conf*, vol. 17 (Citeseer, 2011).
60. Samanvi, K. & Sivasadan, N. Subgraph similarity search in large graphs. *arXiv preprint arXiv:1512.05256* (2015).
61. Niu, S. & Todri-Sanial, A. Multi-programming cross platform benchmarking for quantum computing hardware. *arXiv preprint arXiv:2206.03144* (2022).
62. Bapat, A., Childs, A. M., Gorshkov, A. V. & Schoute, E. Advantages and limitations of quantum routing. *PRX Quantum* **4**, 010313 (2023).
63. Hillmich, S., Zulehner, A. & Wille, R. Exploiting quantum teleportation in quantum circuit mapping. In *Proceedings of the 26th Asia and South Pacific Design Automation Conference, ASPDAC '21*, 792–797 (Association for Computing Machinery, New York, NY, USA, 2021). <https://doi.org/10.1145/3394885.3431604>.
64. Ge, Y. et al. Quantum circuit synthesis and compilation optimization: Overview and prospects. *arXiv preprint arXiv:2407.00736* (2024).
65. Maslov, D., Dueck, G. W., Miller, D. M. & Negrevergne, C. Quantum circuit simplification and level compaction. *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* **27**, 436–444 (2008).
66. Nash, B., Gheorghiu, V. & Mosca, M. Quantum circuit optimizations for nisq architectures. *Quantum Sci. Technol.* **5**, 025010 (2020).
67. Brukner, Č. Quantum causality. *Nat. Phys.* **10**, 259–263 (2014).
68. Goswami, K. & Romero, J. Experiments on quantum causality. *AVS Quantum Science* **2** (2020).
69. Sakurai, J. J. & Commins, E. D. *Modern quantum mechanics*, revised edition (1995).
70. Strogatz, S. H. Exploring complex networks. *nature* **410**, 268–276 (2001).
71. Green, F. Review of handbook of graph theory, combinatorial optimization, and algorithms. *SIGACT N.* **50**, 6–11 (2019).
72. Godsil, C. & Royle, G. F. *Algebraic graph theory*, vol. 207 (Springer Science & Business Media, 2001).
73. Li, L., Bu, K., Koh, D. E., Jaffe, A. & Lloyd, S. Wasserstein complexity of quantum circuits. *arXiv preprint arXiv:2208.06306* (2022).
74. Jarzyna, M. & Kołodyński, J. Geometric approach to quantum statistical inference. *IEEE J. Sel. Areas Inf. Theory* **1**, 367–386 (2020).
75. Van Vu, T. & Saito, K. Thermodynamic unification of optimal transport: Thermodynamic uncertainty relation, minimum dissipation, and thermodynamic speed limits. *Phys. Rev. X* **13**, 011013 (2023).
76. Majtey, A. P., Lamberti, P. W. & Prato, D. P. Jensen-shannon divergence as a measure of distinguishability between mixed quantum states. *Phys. Rev. A* **72**, 052310 (2005).
77. De Domenico, M., Nicosia, V., Arenas, A. & Latora, V. Structural reducibility of multilayer networks. *Nat. Commun.* **6**, 6864 (2015).
78. Rossi, L., Torsello, A. & Hancock, E. R. Attributed graph similarity from the quantum jensen-shannon divergence. In *Similarity-Based Pattern Recognition: Second International Workshop, SIMBAD 2013, York, UK, July 3-5, 2013. Proceedings 2*, 204–218 (Springer, 2013).
79. Crooks, G. E. Measuring thermodynamic length. *Phys. Rev. Lett.* **99**, 100602 (2007).
80. Scandi, M. & Perarnau-Llobet, M. Thermodynamic length in open quantum systems. *Quantum* **3**, 197 (2019).
81. Abiuso, P., Miller, H. J., Perarnau-Llobet, M. & Scandi, M. Geometric optimisation of quantum thermodynamic processes. *Entropy* **22**, 1076 (2020).
82. Johnson, D. M., Dulmage, A. & Mendelsohn, N. On an algorithm of g. birkhoff concerning doubly stochastic matrices. *Can. Math. Bull.* **3**, 237–242 (1960).
83. Qiskit contributors. Qiskit: An open-source framework for quantum computing (2023).
84. Lye, A., Wille, R. & Drechsler, R. Determining the minimal number of swap gates for multi-dimensional nearest neighbor quantum circuits. In *The 20th Asia and South Pacific Design Automation Conference*, 178–183 (2015).
85. Shafaei, A., Saeedi, M. & Pedram, M. Qubit placement to minimize communication overhead in 2d quantum architectures. In *2014 19th Asia and South Pacific Design Automation Conference (ASP-DAC)*, 495–500 (2014).
86. Freedman, D., Pisani, R. & Purves, R. *Statistics* (international student edition). *Pisani, R. Purves, 4th edn. WW Norton & Company, New York* (2007).
87. Carroll, S. M. *Spacetime and geometry* (Cambridge University Press, 2019).
88. Penrose, R. Asymptotic properties of fields and space-times. *Phys. Rev. Lett.* **10**, 66–68 (1963).
89. Ito, S. Stochastic thermodynamic interpretation of information geometry. *Phys. Rev. Lett.* **121**, 030605 (2018).
90. Sekino, Y. & Susskind, L. Fast scramblers. *J. High. Energy Phys.* **2008**, 065 (2008).
91. Sarkar, A. *Applications of Quantum Computation and Algorithmic Information: for Causal Modeling in Genomics and Reinforcement Learning*. Phd thesis, Delft University of Technology (2022).
92. Goldenfeld, N. *Lectures on phase transitions and the renormalization group* (CRC Press, 2018).
93. Brightwell, G. R. & Winkler, P. Graph homomorphisms and phase transitions. *J. Combinatorial Theory, Ser. B* **77**, 221–262 (1999).
94. Dorogovtsev, S. N., Goltsev, A. V. & Mendes, J. F. F. Critical phenomena in complex networks. *Rev. Mod. Phys.* **80**, 1275–1335 (2008).
95. Niu, S., Suau, A., Staffelbach, G. & Todri-Sanial, A. A hardware-aware heuristic for the qubit mapping problem in the nisq era. *IEEE Trans. Quantum Eng.* **1**, 1–14 (2020).
96. Itoko, T., Raymond, R., Imamichi, T. & Matsuo, A. Optimization of quantum circuit mapping using gate transformation and commutation. *Integration* **70**, 43–50 (2020).
97. Park, S., Kim, D., Kweon, M., Sim, J.-Y. & Kang, S. A fast and scalable qubit-mapping method for noisy intermediate-scale quantum computers. In *Proceedings of the 59th ACM/IEEE Design Automation Conference, DAC '22*, 13–18 (Association for Computing Machinery, New York, NY, USA, 2022). <https://doi.org/10.1145/3489517.3530402>.
98. Liu, J. et al. Tackling the qubit mapping problem with permutation-aware synthesis. *arXiv preprint arXiv:2305.02939* (2023).
99. Park, S., Kim, D., Sim, J.-Y. & Kang, S. Mcqa: Multi-constraint qubit allocation for near-ftqc device. In *Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design, ICCAD '22* (Association for Computing Machinery, New York, NY, USA, 2022). <https://doi.org/10.1145/3508352.3549462>.
100. Hillmich, S., Zulehner, A. & Wille, R. Exploiting quantum teleportation in quantum circuit mapping. In *2021 26th Asia and South Pacific Design Automation Conference (ASP-DAC)*, 792–797 (IEEE, 2021).
101. Paraskevopoulos, N., Almudever, C. G. & Feld, S. besnake: A routing algorithm for scalable spin-qubit architectures 2403.16090 (2024).
102. Paraskevopoulos, N., Sebastiano, F., Almudever, C. G. & Feld, S. Spinq: Compilation strategies for scalable spin-qubit architectures. *ACM Trans. Quantum Comput.* **5**, 1–36 (2023).

103. Kreppel, F. et al. Quantum circuit compiler for a shuttling-based trapped-ion quantum computer. *Quantum* **7**, 1176 (2023).
104. Sivaram, S. et al. t|ket>: a retargetable compiler for NISQ devices. *Quantum Sci. Technol.* **6**, 014003 (2020).
105. Patel, T., Silver, D. & Tiwari, D. Geysler: a compilation framework for quantum computing with neutral atoms. In *Proceedings of the 49th Annual International Symposium on Computer Architecture*, 383–395 (2022).
106. Preskill, J. Reliable quantum computers. *Proc. R. Soc. Lond. Ser. A: Math., Phys. Eng. Sci.* **454**, 385–410 (1998).
107. Bhatnagar, D., Steinberg, M., Elkouss, D., Almudever, C. G. & Feld, S. Low-depth flag-style syndrome extraction for small quantum error-correction codes. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 01, 63–69 (2023).
108. Gottesman, D. An introduction to quantum error correction and fault-tolerant quantum computation. *Quantum Inf. Sci. its contributions Math., Proc. Symposia Appl. Math.* **68**, 13–58 (2010).
109. Gottesman, D. Theory of fault-tolerant quantum computation. *Phys. Rev. A* **57**, 127–137 (1998).
110. Chamberland, C. & Beverland, M. E. Flag fault-tolerant error correction with arbitrary distance codes. *Quantum* **2**, 53 (2018).
111. Steane, A. M. Overhead and noise threshold of fault-tolerant quantum error correction. *Phys. Rev. A* **68**, 042322 (2003).
112. Cross, A. W., DiVincenzo, D. P. & Terhal, B. M. A comparative code study for quantum fault-tolerance. *arXiv preprint arXiv:0711.1556* (2007).
113. Campbell, E. T., Terhal, B. M. & Vuillot, C. Roads towards fault-tolerant universal quantum computation. *Nature* **549**, 172–179 (2017).
114. Aliferis, P., Gottesman, D. & Preskill, J. Quantum accuracy threshold for concatenated distance-3 codes. *arXiv preprint quant-ph/0504218* (2005).
115. Aharonov, D. & Ben-Or, M. Fault-tolerant quantum computation with constant error. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, 176–188 (1997).
116. Knill, E. Quantum computing with realistically noisy devices. *Nature* **434**, 39–44 (2005).
117. Chao, R. & Reichardt, B. W. Quantum error correction with only two extra qubits. *Phys. Rev. Lett.* **121**, 050502 (2018).
118. Lidar, D. A. & Brun, T. A. *Quantum error correction* (Cambridge university press, 2013).
119. Lao, L. & Almudever, C. G. Fault-tolerant quantum error correction on near-term quantum processors using flag and bridge qubits. *Phys. Rev. A* **101**, 032333 (2020).
120. Bandic, M. et al. Mapping Quantum Circuits to Modular Architectures with QUBO. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)* 790–801 (IEEE, 2023). <https://doi.org/10.1109/QCE57702.2023.00094>.
121. Pant, M. et al. Routing entanglement in the quantum internet. *npj Quantum Inf.* **5**, 25 (2019).
122. Nicolini, C., Vlasov, V. & Bifone, A. Thermodynamics of network model fitting with spectral entropies. *Phys. Rev. E* **98**, 022322 (2018).
123. Albert, R. & Barabási, A.-L. Statistical mechanics of complex networks. *Rev. Mod. Phys.* **74**, 47–97 (2002).
124. Benatti, F. & Narnhofer, H. Entropy behaviour under completely positive maps. *Lett. Math. Phys.* **15**, 325–334 (1988).
125. Bahreini, T. & Mohammadzadeh, N. An MINLP model for scheduling and placement of quantum circuits with a heuristic solution approach. *J. Emerging Technol. Comput.* **12**, 29 (2015).
126. Li, G. et al. Tackling the qubit mapping problem for NISQ-era quantum devices. In *International Conference on Architectural Support for Programming Languages and Operating Systems*, 1001–1014 (2019).
127. Jiang, H., Deng, Y. & Xu, M. Quantum circuit transformation based on tabu search. *arXiv preprint arXiv:2104.05214* (2021).
128. Peham, T., Burgholzer, L. & Wille, R. On optimal subarchitectures for quantum circuit mapping. *ACM Trans. Quantum Comput.* **4**, 1–20 (2023).
129. Cordella, L. P., Foggia, P., Sansone, C. & Vento, M. A (sub) graph isomorphism algorithm for matching large graphs. *IEEE Trans. Pattern Anal. Mach. Intell.* **26**, 1367–1372 (2004).

Acknowledgements

We thank Kenneth Goodenough, Hans van Someren, Pablo le Henaff, David Elkouss, and Tariq Bontekoe for insightful discussions and useful manuscript feedback. MS, MB, and SF are grateful for financial support from the Intel corporation. CGA also acknowledges support from the Spanish Ministry of Science, Innovation and Universities through the Beatriz Galindo program 2020 (BG20-00023) and the European ERDF under grant PID2021-123627OB-C51. Also from the QuantERA grant EQUIP with the grant number PCI2022-133004, funded by Agencia Estatal de Investigación, Ministerio de Ciencia e Innovación, Gobierno de España, MCIN/AEI/10.13039/501100011033, and by the European Union “NextGenerationEU/PRTR. AS acknowledges funding from the Dutch Research Council (NWO) through the project “QuTech Part III Application-based research” (project no. 601.QT.001 Part III-C-NISQ).

Author contributions

M.S., A.S., and S.S. developed the theoretical framework and formalism. M.B., S.S., and A.S. developed the numerical algorithm based on the framework and implemented all numerical simulations. M.B. developed the subgraph isomorphism qubit-assignment algorithm. Part of this work was conducted during the master thesis for S.S., and was supervised by M.S., M.B., and S.F. C.A. and S.F. supervised and coordinated the project, and provided guidance during the writing process.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41534-024-00909-7>.

Correspondence and requests for materials should be addressed to Matthew Steinberg.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2024