Low-rank geometry of two-qubit gates

Llorenç Balada Gaggioli *
LAAS-CNRS, Université de Toulouse, Toulouse, France
and Faculty of Electrical Engineering, Czech Technical University in Prague, Prague, Czech Republic

(Received 16 April 2026; accepted 27 October 2026; published 5 November 2026)

We present a framework based on the determinantal geometry of two-qubit gates. Combining the Weyl chamber representation with operator Schmidt theory, we interpret gate synthesis as a distance problem to determinantal varieties. This gives an operational geometry to the Weyl chamber, quantifying nonlocal complexity. We show that the $\sqrt{i\text{SWAP}}$ gate is the closest perfect entangler to the variety of local operations, and that no perfect entangler can be approximated by a local gate with average gate fidelity above 78.3%. The three different determinantal costs form a synthesis-adapted coordinate system that encodes nonlocal complexity and generally reconstruct the Weyl coordinates in the reduced sector.

DOI: 10.1103/nrx9-9p92

I. INTRODUCTION

Two-qubit gates are the fundamental nonlocal resource of quantum computation. Combining them with arbitrary one-qubit rotations is sufficient to generate universal quantum circuits, and moreover, any two-qubit entangling operation is universal up to local operations [1–4]. But universality alone does not make all entangling gates equivalent in terms of synthesis cost, interaction requirements, and entanglement structure. This motivates the search for descriptions that isolate the intrinsic nonlocal content and organize all two-qubit operations in a canonical and operationally meaningful geometry [5,6].

To isolate the genuinely nonlocal content of two-qubit gates, we have to separate the local degrees of freedom from one-qubit operations, leading to a classification of gates up to local equivalence [7]. This compresses the full space of two-qubit unitaries into a compact geometric representation, based on the Cartan decomposition [8–10]. In this picture, subsets of gates like the perfect entanglers have a clear geometric interpretation, making two-qubit nonlocality not only classifiable but also geometrically explicit.

Within this geometric framework, several notions of two-qubit nonlocality have been developed. Complete sets of local invariants characterize when gates are locally equivalent, while the Weyl-chamber picture identifies specific subsets like the perfect entanglers. State-based quantities like entanglement power and entangling capability were introduced to quantify how much entanglement gates can generate on separable states [9,11]. On the other hand, operator-level approaches based on the entanglement of quantum evolution [12], strength measures of unitary dynamics [13], and operator Schmidt decomposition [14–18] emphasized that the nonlocality of two-qubit gates is not

*Contact author: llorenc.balada.gaggioli@fel.cvut.cz

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

fully determined by their action on states. Together, these works show that local classification, entangling capacity, and operator complexity are closely related but genuinely distinct aspects of two-qubit gate nonlocality.

Despite the success of existing geometric and entanglement-based descriptions, a synthesis-oriented geometry of two-qubit gates is still lacking. Classification, entangling power, and implementation cost capture different aspects of nonlocality, but they do not directly quantify how far a certain gate is from structurally simpler gates. Here, we address this problem by combining the Weyl-chamber picture with operator Schmidt theory and determinantal geometry. Two-qubit synthesis is reformulated as a distance problem to low-rank operator families, giving an operational notion of nonlocal complexity to the Weyl chamber. This approach yields closed-form operator Schmidt data from Weyl coordinates, identifies the $\sqrt{i\text{SWAP}}$ gate as the closest perfect entangler to locality, gives fidelity bounds under constrained architectures, and provides a set of determinantal coordinates that encode synthesis complexity almost injectively.

II. FRAMEWORK

A. Operator Schmidt decomposition

Let $\mathcal{H}_A$ and $\mathcal{H}_B$ be two-dimensional Hilbert spaces. Any operator U acting on $\mathcal{H}_A \otimes \mathcal{H}_B$ has an operator Schmidt decomposition

$$U = \sum_{j=0}^{r-1} s_j A_j \otimes B_j, \tag{1}$$

where r is the Schmidt rank, $\{A_j\}, \{B_j\}$ are orthonormal operator sets, and $s_j \geqslant 0$ are the Schmidt coefficients satisfying $\sum_{j=0}^{r-1} s_j^2 = \|U\|_F^2 = 4$, for $U \in SU(4)$. If $r = 1$, we have a purely local two-qubit gate, while if $r > 1$ then there is some nonlocality. Therefore, the Schmidt rank determines how many independent product operators are required to represent a gate U.

In order to compute the Schmidt coefficients, we can use the so-called realignment map $R(U)$ satisfying $R_{(i_A,j_A),(i_B,j_B)} = U_{(i_A,i_B),(j_A,j_B)}$. The Schmidt coefficients of U are then given by the singular values of $R(U)$, $s_j = \sigma_j(R(U))$.

B. Determinantal geometry

For any matrix M, we call the set

$$\mathcal{D}_k = \{M\text{: } \operatorname{rank}(R(M)) \leqslant k\} \tag{2}$$

the k-determinantal variety, defined by the vanishing of all $(k+1) \times (k+1)$ minors.

We can now define the determinantal cost as the distance to the determinantal varieties by letting $U^\star$ be a target gate. Then, we can write the distance as

$$d_k(U^\star)^2 = \min_{V \in \mathcal{D}_k} \|U^\star - V\|_F^2 = \sum_{j \geqslant k} s_j^2, \tag{3}$$

where $s_0 \geqslant s_1 \geqslant s_2 \geqslant s_3$. This is given by the Eckart-Young theorem [19].

Importantly, $\mathcal{D}_k$ contains general matrices and is not restricted to physically realizable unitary operators. We therefore denote by $\mathcal{U}_k = \mathcal{D}_k \cap U(4)$ the unitary subset of the variety. For example, $\mathcal{U}_1$ is the set of local unitary operators.

If a specific architecture is limited to the synthesis of rank $\leqslant k$ gates, then $d_k(U^\star)$ is the minimal Frobenius norm error to generate $U^\star$ inside the full variety $\mathcal{D}_k$. The determinantal cost then provides a lower bound on the cost restricted to the unitary subset $\mathcal{U}_k$ and a fundamental bound on the achievable performance under gate locality constraints.

III. THE WEYL CHAMBER

The Schmidt coefficients and ranks are invariant under local unitary transformations. Let $U \in SU(4)$; two gates U, V are said to be locally equivalent if there exist single-qubit unitaries $K_i \in SU(2)$ such that $V = (K_1 \otimes K_2)U(K_3 \otimes K_4)$. If $U = \sum_{j=0}^{r-1} s_j A_j \otimes B_j$, then $V = \sum_{j=0}^{r-1} s_j (K_1 A_j K_3) \otimes (K_2 B_j K_4)$.

This hints we should move to the space of locally equivalent two-qubit unitaries. Local gates do not generate entanglement, so classifying unitaries up to local equivalence allows for the study of the truly nonlocal content of the gates. The Cartan decomposition [10] states that every two-qubit gate is locally equivalent to

$$U = \exp\left[-\frac{i}{2}(c_1 X \otimes X + c_2 Y \otimes Y + c_3 Z \otimes Z)\right].$$

The full *Weyl chamber* is the tetrahedron with vertices $(0,0,0)$, $(\pi,0,0)$, $(\frac{\pi}{2},\frac{\pi}{2},0)$, $(\frac{\pi}{2},\frac{\pi}{2},\frac{\pi}{2})$. This representation reduces the 15-parameter original manifold into a three-dimensional nonlocal space.

For the analytical results below, we work in the reduced sector $0 \leqslant c_3 \leqslant c_2 \leqslant c_1 \leqslant \frac{\pi}{2}$. The complementary half of the conventional Weyl chamber is mapped to this reduced sector by the reflection $c_1 \mapsto \pi - c_1$. Under this reflection, the squared operator Schmidt coefficients are permuted, and therefore the ordered Schmidt spectrum and all determinantal costs are unchanged. For visualization, we use the conventional extended representation of the Weyl chamber, in which the c_1 axis is displayed over $[0, \pi]$.

The relation between the Weyl coordinates of two-qubit gates and their operator-Schmidt spectrum is known [14,18]; here, we recall the result in our parameters.

Theorem 1. Let $U = \exp[-\frac{i}{2}(c_1 XX + c_2 YY + c_3 ZZ)]$ and define $C_j = \cos(\frac{c_j}{2})$, $S_j = \sin(\frac{c_j}{2})$. Then, the operator Schmidt coefficients are

$$\begin{aligned}
s_0 &= 2\sqrt{(C_1C_2C_3)^2 + (S_1S_2S_3)^2},\\
s_1 &= 2\sqrt{(C_1S_2S_3)^2 + (S_1C_2C_3)^2},\\
s_2 &= 2\sqrt{(S_1C_2S_3)^2 + (C_1S_2C_3)^2},\\
s_3 &= 2\sqrt{(S_1S_2C_3)^2 + (C_1C_2S_3)^2}.
\end{aligned}$$

Moreover, through the reduced Weyl chamber, $s_0 \geqslant s_1 \geqslant s_2 \geqslant s_3$ and $\sum_{i=0}^{3} s_i^2 = 4$. The proof of these closed-form expressions for the Schmidt coefficients is provided in Appendix A.

Example 1 (The Heisenberg diagonal). Let us consider the section of the Weyl chamber where $(c_1, c_2, c_3) = (\alpha, \alpha, \alpha)$, which includes gates such as the identity, $\sqrt{\text{SWAP}}$, and SWAP. The Schmidt spectrum is

$$s_0 = \sqrt{1 + 3\cos^2(\alpha)}, \quad s_1 = s_2 = s_3 = \sin(\alpha),$$

which is ordered for s_0 the largest. The Frobenius distance to the rank-1 variety is then the sum of squares of the three smallest Schmidt coefficients, which in this case give

$$d_1(\alpha)^2 = \sum_{i=1}^{3} s_i^2 = 3\sin^2(\alpha).$$

When $\alpha = 0$, we have $U = I$, so a local gate, as we see from $d_1(0)^2 = 0$. For $\alpha = \pi/4$, U is locally equivalent to the $\sqrt{\text{SWAP}}$ gate, with $d_1^2 = 3/2$. And for $\alpha = \pi/2$, U is locally equivalent to SWAP, with $d_1^2 = 3$, which is the maximum along this line. This provides a closed-form expression for a lower bound on the synthesis cost under a rank-1 constraint.

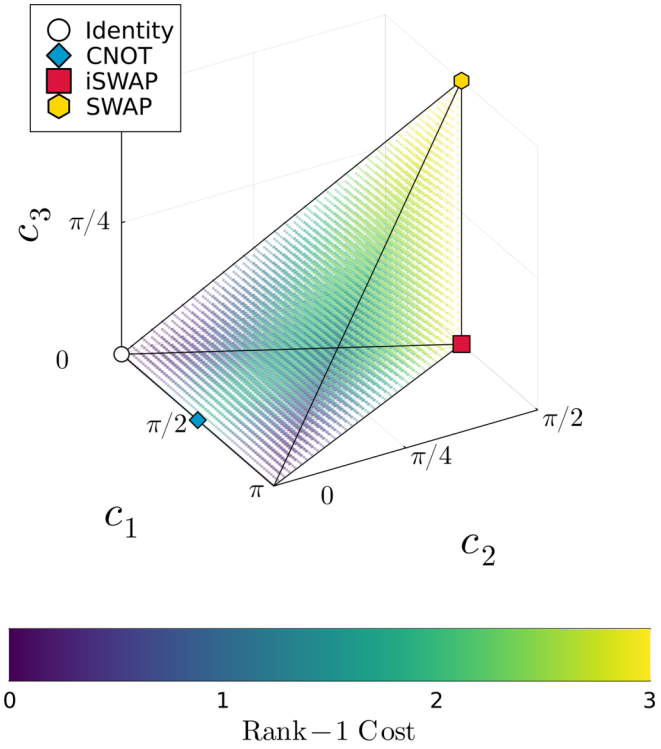


FIG. 1. The Weyl chamber colored by the rank-1 determinantal cost.

We can consider the cost to the different varieties for the entirety of the Weyl chamber, as we only have to compute the Schmidt coefficients (see Figs. 1 and 2). While the Weyl chamber provides a coordinate system (c_1, c_2, c_3) that classifies the two-qubit gates in their local equivalence classes, we now consider the underlying geometry that can add additional crucial information to the two-qubit gate synthesis.

IV. COST OF PERFECT ENTANGLERS

The Weyl chamber has a section that is occupied fully by the so-called perfect entangling gates. These are gates that can produce a maximally entangled state from an unentangled one [10]. The perfect entangler region in the reduced Weyl chamber is bounded by the following planes:

$$c_1 + c_2 \geqslant \frac{\pi}{2}, \quad c_2 + c_3 \leqslant \frac{\pi}{2}.$$

All the gates in this region share the same property of being perfect entanglers, but we want to answer the question of how costly it is to synthesize them when we are constrained by the rank of the possible gates we can use. We solve the optimization problem

$$\min_{c_1,c_2,c_3} \quad d_k(U(c_1, c_2, c_3)) \text{ subject to } U \text{ a perfect entangler.} \tag{4}$$

Theorem 2. Among all perfect entanglers in the reduced Weyl chamber, the distance to the rank-1 variety satisfies

$$d_1^2 \geqslant \tfrac{5}{2} - \sqrt{2}, \tag{5}$$

with equality if and only if $(c_1, c_2, c_3) = (\frac{\pi}{4}, \frac{\pi}{4}, 0)$. This point corresponds to the $\sqrt{i\text{SWAP}}$ gate.

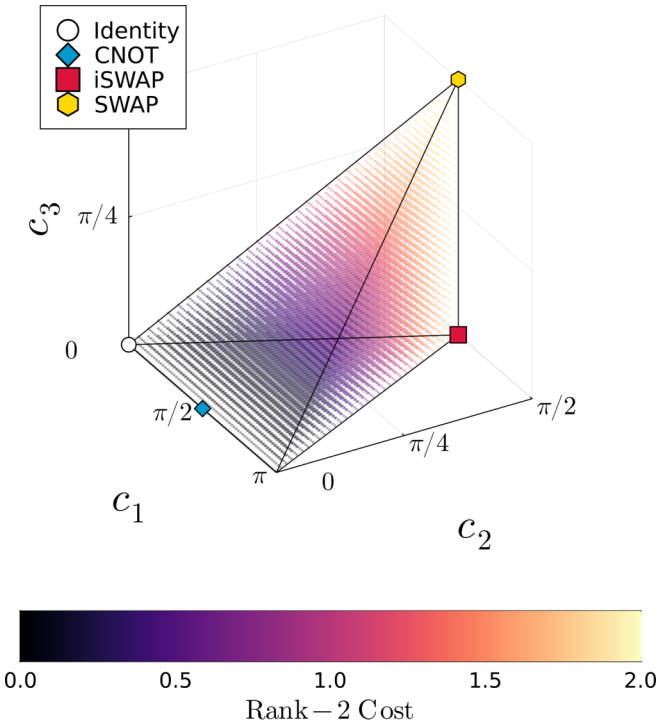


FIG. 2. The Weyl chamber colored by the rank-2 determinantal cost.

We prove the theorem in Appendix B. This result sets a strictly positive gap between the rank-1 variety and the set of perfect entanglers. The $\sqrt{i\text{SWAP}}$ lies then at the minimal cost, and is therefore the least nonlocal perfect entangler under the operator-Schmidt geometry.

A. Different norms, different costs

Up to this point, we have been using the Frobenius norm to define the distance. However, we can pick any unitarily invariant norm as the Eckart-Young theorem implies that the distance under any Schatten p-norm is fully determined by the Schmidt spectrum. For this, we now define the distance in realignment space, for $1 \leqslant p < \infty$, as

$$d_{1,p}(U) = \min_{V \in \mathcal{D}_1} \|R(U) - R(V)\|_{S_p} = \left(s_1^p + s_2^p + s_3^p\right)^{1/p} \tag{6}$$

and $d_{1,\infty}(U) = s_1$. For $p = 2$, realignment preserves the Frobenius norm, and therefore $d_{1,2}(U) = d_1(U)$. For $p \neq 2$, however, these are costs in realignment space, which changes the cost interpretation we had for $p = 2$, but we can still have a notion of distance that could be interesting to study. For $p = 2$, the normalization $\sum_{j=0}^{3} s_j^2 = 4$ implies $d_{1,2}(U)^2 = 4 - s_0^2$, so minimizing the cost is equivalent to maximizing s_0. This argument does not extend to $p \neq 2$, since in general $\sum_j s_j^p$ is not constant over the Weyl chamber.

If we let $p = 1$, we have the trace norm. In this case, $d_{1,1} = s_1 + s_2 + s_3$, and the optimal perfect entangler that minimizes the distance to the rank-1 variety is the CNOT, as we can see in Fig. 3 and show in Appendix B.

In the limit $p \to \infty$, the cost is governed by the largest remaining Schmidt coefficient, and the distance is minimized in the $(c_1, c_2, c_3) = (\frac{\pi}{4}, \frac{\pi}{4}, \alpha)$ line. The gates $\sqrt{i\text{SWAP}}$ and $\sqrt{\text{SWAP}}$ are part of this line.

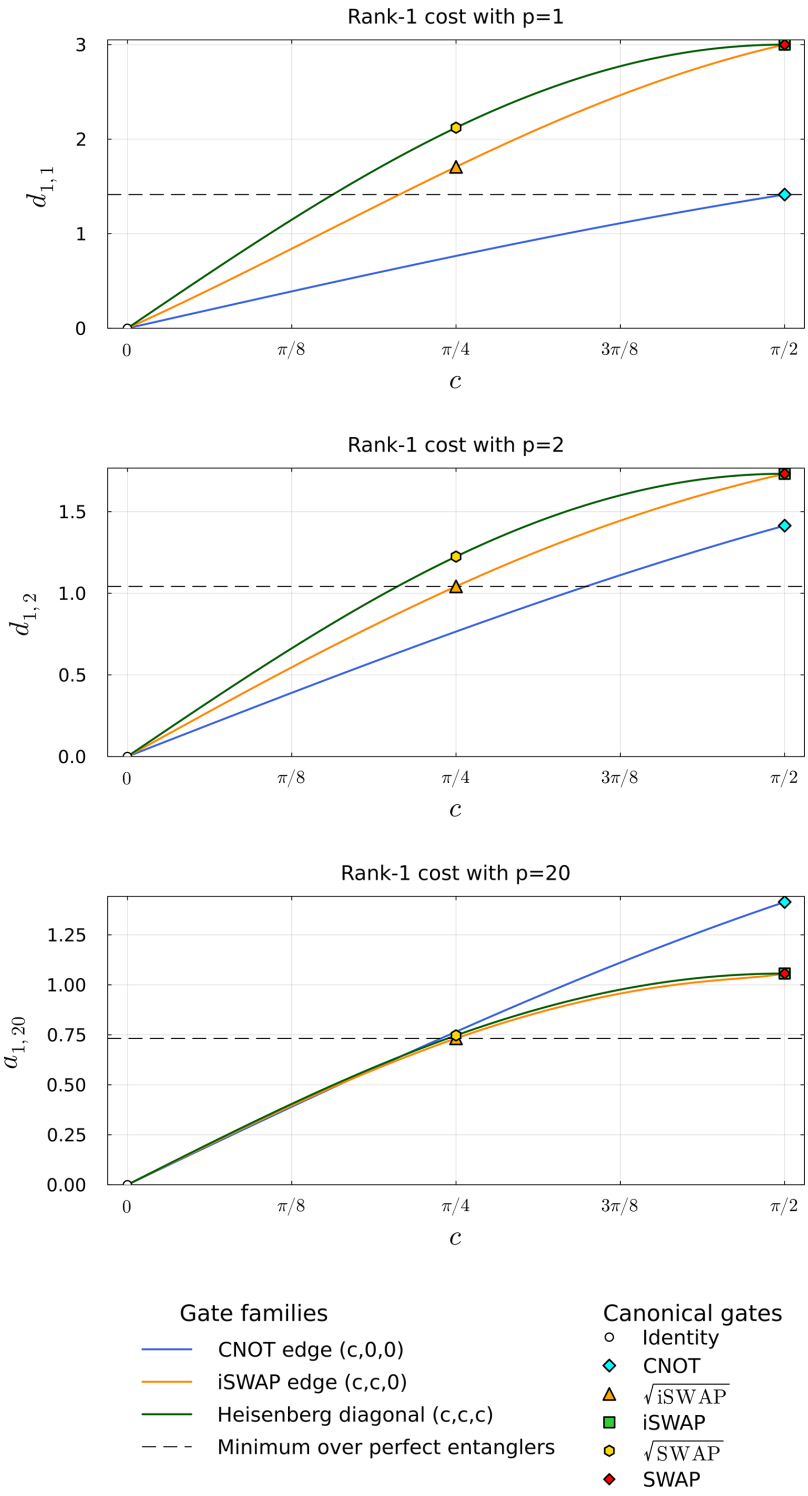


FIG. 3. Rank-1 distance $d_{1,p}$ along representative one-parameter families of the Weyl chamber for three Schatten norms.

These results show that the notion of a "closest" perfect entangler depends on how the residual Schmidt tail is weighted. The Schatten 1-norm penalizes the tail additively and therefore favors sparse operator-Schmidt structure, the Frobenius norm gives a quadratic rms-type penalty, and the $p \to \infty$ limit is governed by the largest remaining Schmidt coefficient. In concrete hardware settings, these different norms may become relevant when the physical control or noise model induces a corresponding cost functional. The Frobenius norm is particularly interesting because it leads directly to fidelity bounds.

B. Fidelity bounds

Having bounds on the Frobenius norm means that we can get bounds on the synthesis fidelity of a target gate under rank constraints. With this we know how well we can generate a specific gate if the given physical architecture has constraints of rank. For unitaries $U, V \in U(d)$, we write the average gate fidelity [20] as

$$F_{\text{avg}}(U, V) = \frac{|\text{Tr}(U^\dagger V)|^2 + d}{d(d+1)}.$$

Theorem 3. No perfect entangler can be approximated by a local gate with average gate fidelity exceeding 78.3%. The bound is tight at the $\sqrt{i\text{SWAP}}$ Weyl point.

Proof. For a local unitary $V = A \otimes B$, with $A, B \in U(2)$, define $\widetilde{A} = A/\sqrt{2}$ and $\widetilde{B} = B/\sqrt{2}$, so that $\|\widetilde{A}\|_F = \|\widetilde{B}\|_F = 1$ and $V = 2\widetilde{A} \otimes \widetilde{B}$. By the variational characterization of the largest operator Schmidt coefficient, Cauchy-Schwarz and orthonormality give $|\text{Tr}[U^\dagger(\widetilde{A} \otimes \widetilde{B})]| \leqslant s_0$. Hence, for any local unitary V, $|\text{Tr}(U^\dagger V)| \leqslant 2s_0$. For Weyl gates, the Schmidt operators are normalized Pauli operators, which are proportional to unitaries, so this bound can be saturated. Therefore, $\max_{V \in \mathcal{U}_1} |\text{Tr}(U^\dagger V)| = 2s_0$.

From Theorem 2, $d_1(U)^2 = 4 - s_0^2 \geqslant \frac{5}{2} - \sqrt{2}$, and hence $s_0^2 \leqslant \frac{3+2\sqrt{2}}{2}$. Therefore, for $d = 4$, $F_{\text{avg}}(U, V) \leqslant \frac{4s_0^2+4}{20} \leqslant \frac{5+2\sqrt{2}}{10} \approx 0.7828$. The bound is attained at $(c_1, c_2, c_3) = (\pi/4, \pi/4, 0)$, corresponding to the $\sqrt{\text{ISWAP}}$ gate. ■

We can also obtain fidelity bounds for the synthesis of specific gates under fixed operator Schmidt-rank constraints. For example, consider the SWAP gate, whose Schmidt spectrum is $s_0 = s_1 = s_2 = s_3 = 1$. Its rank-2 determinantal cost is therefore $d_2(\text{SWAP})^2 = s_2^2 + s_3^2 = 2$. Since $\mathcal{U}_2 \subseteq \mathcal{D}_2$, for any $V \in \mathcal{U}_2$ we have $\min_\phi \|\text{SWAP} - e^{i\phi}V\|_F^2 \geqslant d_2(\text{SWAP})^2 = 2$. Using $\min_\phi \|U - e^{i\phi}V\|_F^2 = 2d - 2|\text{Tr}(U^\dagger V)|$, with $d = 4$, it follows that $|\text{Tr}(\text{SWAP}^\dagger V)| \leqslant 4 - \frac{d_2(\text{SWAP})^2}{2} = 3$. Hence, $F_{\text{avg}}(\text{SWAP}, V) \leqslant \frac{3^2+4}{20} = 0.65$. Thus, 0.65 is an upper bound on the average gate fidelity achievable when approximating SWAP by a unitary gate of operator Schmidt rank at most 2.

V. DETERMINANTAL COORDINATES

The quantities $d_k(U)^2$ are not only approximation costs to the determinantal varieties $\mathcal{D}_k$. They define a set of coordinates that encode information about the synthesis complexity of the nonlocal two-qubit landscape. We let $x := d_1(U)^2$, $y := d_2(U)^2$ for simplicity.

Theorem 4. Let $U(c_1, c_2, c_3)$ be a two-qubit Weyl gate; for every such gate, the image of the reduced Weyl chamber under $U \mapsto (x, y)$ is the strip

$$\mathcal{R} = \left\{(x, y) \in \mathbb{R}^2 : 0 \leqslant y \leqslant 2,\ \tfrac{3}{2}y \leqslant x \leqslant 2 + \tfrac{1}{2}y\right\}. \tag{7}$$

The perfect entangler region lies in the region

$$\mathcal{R}_{\text{PE}} = \left\{(x, y) \in \mathbb{R}^2 : 0 \leqslant y \leqslant 2,\ L_{\text{PE}}(y) \leqslant x \leqslant 2 + \tfrac{1}{2}y\right\}, \tag{8}$$

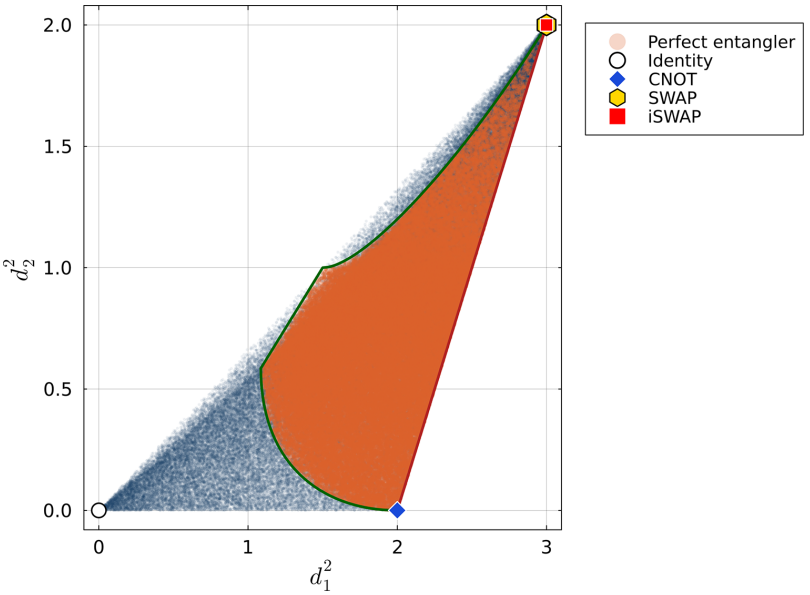


FIG. 4. Image of the Weyl chamber in the determinantal plane. The orange subregion is the perfect entangler set.

where

$$L_{\mathrm{PE}}(y) = \begin{cases} 2 + \frac{y}{2} - \frac{(4-y)\sqrt{y(4-y)}}{4}, & 0 \leqslant y \leqslant 2-\sqrt{2}, \\ y + \frac{1}{2}, & 2-\sqrt{2} \leqslant y \leqslant 1, \\ y + \frac{1}{2} + \frac{1}{2}\sqrt{(y-1)(3-y)}, & 1 \leqslant y \leqslant 2. \end{cases}$$

We prove this theorem in Appendix C. Within the reduced sector, the lower bound of the full slice is attained exactly at the Heisenberg diagonal $c_1 = c_2 = c_3$, and the upper bound is attained at the face $c_1 = \frac{\pi}{2}$. Note that at the face $c_1 = \frac{\pi}{2}$, the map from the Weyl chamber is not injective, as we see in Fig. 4, where the SWAP and iSWAP gates coincide at the same point. This means that this determinantal strip is useful to consider the map of the perfect entanglers, but not to classify them.

This tells us that the determinantal coordinates do not occupy the full plane, but once we fix one, the other one is restricted to a smaller set of values. The perfect-entangler (PE) region indicates that the ability to maximally entangle states does not determine by itself the operator synthesis complexity of the gate. Gates sharing the perfect-entangler property might lie in very different areas of the strip.

A. Complexity: A spectral witness

It is known that to implement two-qubit operations we need at most three CNOT gates, and the exact synthesis problem admits explicit universal and minimal constructions [6,21–24]. With the determinantal coordinates, we can build a spectral witness of CNOT complexity coming directly from the Schmidt spectrum of the gate.

Theorem 5. All unitaries lying on the determinantal point $(x, y) = (0, 0)$ are local gates; one-CNOT gates are locally equivalent to CNOT and map to the point $(x, y) = (2, 0)$. The set of gates implementable with at most two CNOT gates lies in the Weyl face $c_3 = 0$, with image

$$\mathcal{R}_2 = \left\{(x, y) : 0 \leqslant y \leqslant 2,\ \ 2y - \tfrac{1}{4}y^2 \leqslant x \leqslant 2 + \tfrac{1}{2}y\right\}.$$

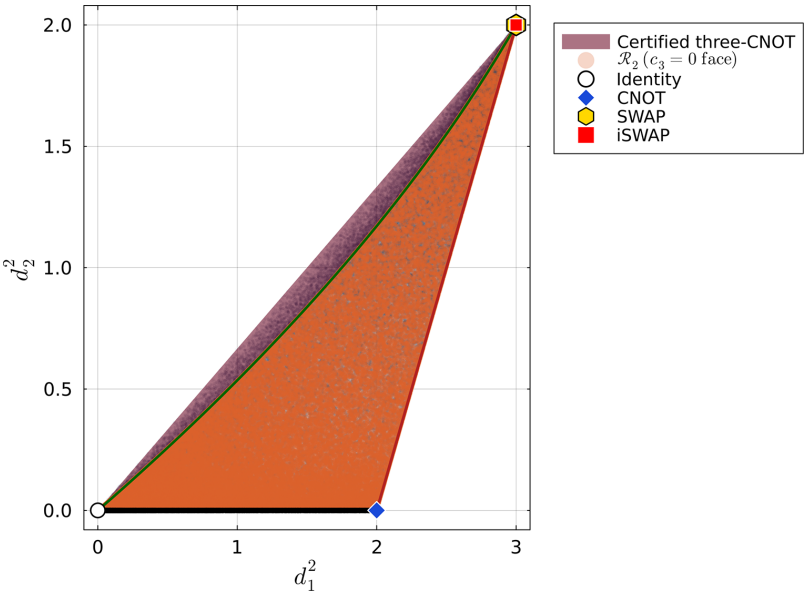


FIG. 5. CNOT complexity regions in the determinantal plane.

Therefore, every gate with $x < 2y - \frac{1}{4}y^2$ requires three CNOT operations to be synthesized.

We show this in Appendix C. The complexity regions of the plane are illustrated in Fig. 5, showing that the coordinates $(d_1(U)^2, d_2(U)^2)$ not only quantify nonlocality but also separate the levels of synthesis complexity in a geometrical way.

As the map $(c_1, c_2, c_3) \mapsto (x, y)$ is not injective, membership in $\mathcal{R}_2$ does not certify that a gate has two-CNOT complexity. However, the other direction is useful, points lying outside of $\mathcal{R}_2$ require three CNOTs. Therefore, (x, y) acts as a spectral witness rather than a classifier of CNOT complexity.

B. Completeness: The determinantal chamber

If we let $z := d_3(U)^2$, we have a three-coordinate system (x, y, z). The first two coordinates already provide a witness of synthesis complexity, but now with z we can complete this picture into a three-dimensional synthesis-adapted chamber.

Theorem 6. The map $(c_1, c_2, c_3) \mapsto (x, y, z)$ is injective on the reduced Weyl chamber for $c_1 < \frac{\pi}{2}$, so each Weyl point is uniquely determined by (x, y, z). On the face $c_1 = \frac{\pi}{2}$, the map is no longer injective, and two points can have the same determinantal data.

The image of the reduced Weyl chamber with $c_1 < \pi/2$ under the determinantal coordinates is

$$\mathcal{D} = \left\{ (x, y, z) \in \mathbb{R}^3 : \begin{array}{l} 0 \leqslant 2z \leqslant y \leqslant x, \\ x < z + 2, \\ x + z \geqslant 2y, \\ y^2 - y(x + z) + 4z \geqslant 0 \end{array} \right\}.$$

Its closure, obtained by including the face $c_1 = \pi/2$, is given by replacing $x < z + 2$ with $x \leqslant z + 2$.

The proof of this result is provided in Appendix C. The determinantal chamber is a synthesis analog of the Weyl chamber as it quantifies the distance to low-rank operator structure and, away from the singular face $c_1 = \frac{\pi}{2}$, the coordinates uniquely determine a Weyl point in the reduced sector. This geometry does not merely describe the landscape of two-qubit gates, but it organizes it in operational terms.

VI. CONCLUSION

In this work, we use determinantal geometry to turn the Weyl chamber from a descriptive classification of the two-qubit operations to a synthesis-oriented geometric landscape. The distances to low-rank varieties provide fundamental lower bounds on implementation error under rank constraints, singling out gates like the $\sqrt{i\text{SWAP}}$ as the least nonlocal perfect entangler. This geometry also allows for the study of the CNOT complexity of the gates, showing how the determinantal distances are not only a way to study synthesis cost but a natural language to describe the compilation and control of two-qubit operations under implementation constraints.

Within the framework presented in this paper, many more things can be studied. The same geometry can be used to study how nonlocal cost grows under a given Hamiltonian, how determinantal cost can serve as a predictor of compilation depth and robustness, and how anisotropy selects which gate families are naturally reachable. For the interaction Hamiltonian $H = \frac{1}{2}(J_x XX + J_y YY + J_z ZZ)$, free evolution follows a trajectory $(c_1, c_2, c_3) = t(J_x, J_y, J_z)$ in the reduced Weyl chamber. An Ising interaction generates the $(c, 0, 0)$ family, an XY interaction generates the $(c, c, 0)$ family, and the isotropic Heisenberg interaction generates the diagonal (c, c, c). In this sense, the determinantal picture is not only a static description of two-qubit gates, but a flexible framework for relating algebraic rank structure, physical interactions, and gate-synthesis complexity.

Future work should investigate in more depth how the determinantal costs can have an active role in hardware-aware gate design, as the norm dependence of the cost suggests that the geometry can be adapted to different physical platforms. The next major conceptual challenge is to study whether an analogous determinantal geometry arises beyond two qubits. For bipartite qudit systems, the operator-Schmidt and determinantal constructions apply directly, but the main obstacle is the loss of the simple three-dimensional Weyl chamber parametrization. For multipartite systems, an additional difficulty arises, there is no single Schmidt decomposition, so one must choose bipartitions, flattenings, or higher-order rank notions.

ACKNOWLEDGMENTS

I am grateful to Jiří Vala and Jakub Mareček for insightful discussions regarding two-qubit operations. The author acknowledges the use of AI for assistance with brainstorming ideas, mathematical development, coding, and drafting the manuscript. The final content, analysis, and conclusions remain the sole responsibility of the author. This work has been supported by European Union's HORIZON-MSCA-2023-DN-JD program under the Horizon Europe (HORIZON) Marie Skłodowska-Curie Actions, Grant Agreement No. 101120296 (TENORS).

DATA AVAILABILITY

Code and data are available at the GitHub repository [25].

APPENDIX A: PROOF OF SCHMIDT SPECTRUM

We can write a two-qubit gate via the Cartan decomposition as $U = e^{-\frac{i}{2}c_1 XX} e^{-\frac{i}{2}c_2 YY} e^{-\frac{i}{2}c_3 ZZ}$. If we let $C_j = \cos(\frac{c_j}{2})$, $S_j = \sin(\frac{c_j}{2})$, then we have

$$U = (C_1 I - iS_1 XX)(C_2 I - iS_2 YY)(C_3 I - iS_3 ZZ).$$

Recall, under usual product, $X^2 = Y^2 = Z^2 = I$, $XY = iZ$, $YZ = iX$, and $ZX = iY$. Therefore, we can write

$$U = a_0 II + a_1 XX + a_2 YY + a_3 ZZ,$$

with

$$a_0 = C_1C_2C_3 - iS_1S_2S_3,$$
$$a_1 = C_1S_2S_3 - iS_1C_2C_3,$$
$$a_2 = S_1C_2S_3 - iC_1S_2C_3,$$
$$a_3 = S_1S_2C_3 - iC_1C_2S_3.$$

To write the unitary from orthonormal bases, we let

$$U = (2a_0)\frac{II}{2} + (2a_1)\frac{XX}{2} + (2a_2)\frac{YY}{2} + (2a_3)\frac{ZZ}{2},$$

which yields $s_i = |2a_i|$.

Now define

$$A = \cos c_1 \cos c_2, \quad B = \cos c_1 \cos c_3, \quad C = \cos c_2 \cos c_3.$$

The squared Schmidt coefficients can be equivalently written as

$$s_0^2 = 1 + A + B + C, \quad s_1^2 = 1 - A - B + C,$$
$$s_2^2 = 1 - A + B - C, \quad s_3^2 = 1 + A - B - C.$$

In the reduced Weyl chamber $0 \leqslant c_3 \leqslant c_2 \leqslant c_1 \leqslant \pi/2$, so $0 \leqslant A \leqslant B \leqslant C$. Since $s_j \geqslant 0$, this proves

$$s_0 \geqslant s_1 \geqslant s_2 \geqslant s_3 \geqslant 0.$$

From unitarity, normalization follows

$$\sum_{j=0}^{3} s_j^2 = \|U\|_F^2 = \mathrm{Tr}(U^\dagger U) = 4.$$

APPENDIX B: PROOFS OF GEOMETRIC THRESHOLDS FOR PERFECT ENTANGLERS

Let us write the Frobenius distance to the rank-1 determinantal variety as

$$d_1^2 = \sum_{i \geqslant 1} s_i^2 = 4 - s_0^2,$$

such that minimizing the distance is equivalent to maximizing s_0^2.

Lemma B1. The maximizer of s_0^2 over the perfect-entangler region of the reduced sector has $c_3 = 0$.

Proof. Let $\theta_j = c_j/2$. From Theorem 1,

$$\frac{s_0^2}{4} = (\cos\theta_1 \cos\theta_2 \cos\theta_3)^2 + (\sin\theta_1 \sin\theta_2 \sin\theta_3)^2.$$

For fixed c_1, c_2, let $t = \cos^2\theta_3$. Since $0 \leqslant \theta_3 \leqslant \theta_2$, we have $t \in [\cos^2\theta_2, 1] \subseteq [1/2, 1]$, and

$$\frac{s_0^2}{4} = (\cos\theta_1 \cos\theta_2)^2 t + (\sin\theta_1 \sin\theta_2)^2 (1 - t).$$

Since $0 \leqslant \theta_1, \theta_2 \leqslant \pi/4$, $\cos\theta_1 \cos\theta_2 \geqslant \sin\theta_1 \sin\theta_2$, so the expression is maximized at $t = 1$, corresponding to $c_3 = 0$. ∎

If we let $c_3 = 0$, then the relevant constraint for perfect entanglers is $c_1 + c_2 \geqslant \frac{\pi}{2}$. The product $\cos\theta_1 \cos\theta_2$ is maximized when $c_1 + c_2$ is minimized, so at $c_1 + c_2 = \frac{\pi}{2}$.

Lemma B2. On the line $c_1 + c_2 = \frac{\pi}{2}$, the product $\cos\theta_1 \cos\theta_2$ is maximized at $c_1 = c_2 = \frac{\pi}{4}$.

Proof. Write $\cos\theta_1 \cos\theta_2 = \frac{1}{2}(\cos(\theta_1 + \theta_2) + \cos(\theta_1 - \theta_2))$, and note $\theta_1 + \theta_2 = \frac{1}{2}(c_1 + c_2) = \frac{\pi}{4}$, so we want to maximize $\cos(\theta_1 - \theta_2)$, which happens at $\theta_1 = \theta_2$. Therefore, we conclude $\theta_1 = \theta_2 = \frac{\pi}{8}$. ■

At this point, $s_0^2 = 4\cos^4(\frac{\pi}{8}) = \frac{3+2\sqrt{2}}{2}$, and hence $d_1^2 = 4 - s_0^2 = \frac{5}{2} - \sqrt{2}$.

Now, we consider different norms, starting with $p = 1$. Let $\theta_j = c_j/2$, $C_j = \cos\theta_j$, and $S_j = \sin\theta_j$. From the Schmidt coefficients in Theorem 1,

$$d_{1,1} = s_1 + s_2 + s_3 \geqslant 2S_1C_2C_3 + 2C_1S_2C_3 + 2C_1C_2S_3 = 2C_3\sin(\theta_1 + \theta_2) + 2C_1C_2S_3.$$

Since $\sin(\theta_1 + \theta_2)/(C_1C_2) = \tan\theta_1 + \tan\theta_2 \leqslant 2$ and $S_3/(1 - C_3) = \cot(\theta_3/2) \geqslant \cot(\pi/8) > 2$, we have $C_1C_2S_3 \geqslant (1 - C_3)\sin(\theta_1 + \theta_2)$, and therefore

$$d_{1,1} \geqslant 2\sin(\theta_1 + \theta_2).$$

For a perfect entangler, $c_1 + c_2 \geqslant \pi/2$, so $\theta_1 + \theta_2 \geqslant \pi/4$, while $\theta_1 + \theta_2 \leqslant \pi/2$, and hence $d_{1,1} \geqslant 2\sin(\pi/4) = \sqrt{2}$. Equality requires $c_3 = 0$ and $S_1S_2 = 0$ together with $c_1 + c_2 = \pi/2$; using $c_1 \geqslant c_2$, this gives $(c_1, c_2, c_3) = (\pi/2, 0, 0)$, corresponding to the CNOT gate. Thus, CNOT is the perfect entangler minimizing the $p = 1$ realignment-space distance.

For $p = \infty$, we have $d_{1,\infty} = s_1$, so we minimize s_1 over the perfect-entangler region. Let $u = \cos c_1$, $v = \cos c_2$, and $w = \cos c_3$, so that $s_1^2 = 1 - uv - uw + vw$. If $c_2 \geqslant \pi/4$, then for fixed v, w this expression is minimized by maximizing u, and the Weyl ordering gives $u = v$, i.e., $c_1 = c_2$. Hence, $s_1^2 = 1 - v^2 = \sin^2 c_2 \geqslant 1/2$, with equality for $c_1 = c_2 = \pi/4$, independently of $c_3 \in [0, \pi/4]$. If $c_2 \leqslant \pi/4$, the perfect-entangler condition $c_1 + c_2 \geqslant \pi/2$ gives $u \leqslant \sin c_2$, and since $v - u \geqslant 0$, $s_1^2 = 1 - uv + w(v - u)$ is minimized by taking $w = v$ and $u = \sin c_2$, resulting in $s_1^2 \geqslant 1 + \cos^2 c_2 - 2\sin c_2 \cos c_2 \geqslant 1/2$, where the last expression is minimized at $c_2 = \pi/4$. Therefore $d_{1,\infty} \geqslant 1/\sqrt{2}$, with equality for the family $(c_1, c_2, c_3) = (\pi/4, \pi/4, c_3)$, $0 \leqslant c_3 \leqslant \pi/4$, which contains the $\sqrt{i\text{SWAP}}$ and $\sqrt{\text{SWAP}}$ gates.

APPENDIX C: PROOFS OF DETERMINANTAL COORDINATES

Here, we prove the three theorems of Sec. V. For this, we define $a = \cos c_1 \cos c_2$, $b = \cos c_1 \cos c_3$, $c = \cos c_2 \cos c_3$. With these, note that $s_0^2 = 1 + a + b + c$, $s_1^2 = 1 - a - b + c$, $s_2^2 = 1 - a + b - c$, $s_3^2 = 1 + a - b - c$, and therefore $x = d_1^2 = 3 - a - b - c$, $y = d_2^2 = 2 - 2c$, $z = d_3^2 = 1 + a - b - c$.

Proof of Theorem 4. Recall that $x = 3 - a - b - c$ and $y = 2 - 2c$, where $a = \cos c_1 \cos c_2$, $b = \cos c_1 \cos c_3$, and $c = \cos c_2 \cos c_3$. Fixing y therefore fixes $c = 1 - y/2$, and we have $x = 3 - c - \cos c_1(\cos c_2 + \cos c_3)$. Let $u = \cos c_1$, $v = \cos c_2$, and $w = \cos c_3$, so that $u \leqslant v \leqslant w$ and $vw = c$. Since $u(v + w) \leqslant v(v + w) = v^2 + c \leqslant 2c$, the lower boundary of the reduced Weyl chamber is attained when $u = v = w$, i.e., $c_1 = c_2 = c_3$, giving $x = 3 - 3c = \frac{3}{2}y$. The upper boundary is obtained for $u = 0$, i.e., $c_1 = \frac{\pi}{2}$, giving $x = 3 - c = 2 + \frac{y}{2}$. For the perfect-entangler region, we additionally impose $c_1 + c_2 \geqslant \frac{\pi}{2}$ and $c_2 + c_3 \leqslant \frac{\pi}{2}$, which imply $u \leqslant \sqrt{1 - v^2}$ and $w \geqslant \sqrt{1 - v^2}$. For fixed v and w, the quantity $u(v + w)$ is maximized by taking $u = \min\{v, \sqrt{1 - v^2}\}$. The active constraints then change according to the value of $c = 1 - y/2$, producing three regimes. For $\frac{1}{\sqrt{2}} \leqslant c \leqslant 1$, equivalently $0 \leqslant y \leqslant 2 - \sqrt{2}$, one has $v \geqslant 1/\sqrt{2}$ and hence $u = \sqrt{1 - v^2}$. The quantity $\sqrt{1 - v^2}(v + c/v)$ is maximized at the smallest allowed value $v = c$, corresponding to $w = 1$, so $c_3 = 0$ and $c_1 + c_2 = \frac{\pi}{2}$. Substituting $c = 1 - y/2$ gives $x = 2 + \frac{y}{2} - \frac{(4-y)\sqrt{y(4-y)}}{4}$. For $\frac{1}{2} \leqslant c \leqslant \frac{1}{\sqrt{2}}$, equivalently $2 - \sqrt{2} \leqslant y \leqslant 1$, the maximum occurs at $v = 1/\sqrt{2}$, so $c_1 = c_2 = \frac{\pi}{4}$ while c_3 varies, giving $x = y + \frac{1}{2}$. Finally, for $0 \leqslant c \leqslant \frac{1}{2}$, equivalently $1 \leqslant y \leqslant 2$, we have $v \leqslant 1/\sqrt{2}$ and therefore the maximizing choice is $u = v$. The largest allowed v is obtained by saturating

$w = \sqrt{1-v^2}$, corresponding to $c_1 = c_2$ and $c_2 + c_3 = \frac{\pi}{2}$. Since $c = v\sqrt{1-v^2}$, we obtain $v^2 = \frac{1-\sqrt{1-4c^2}}{2} = \frac{1-\sqrt{(y-1)(3-y)}}{2}$. Substituting into $x = 1 + y - v^2$ gives $x = y + \frac{1}{2} + \frac{1}{2}\sqrt{(y-1)(3-y)}$. Therefore, the lower boundary of the perfect-entangler region is $L_{\mathrm{PE}}(y)$, with the three branches stated in Theorem 4, while the upper boundary remains $x = 2 + \frac{y}{2}$, attained on the face $c_1 = \frac{\pi}{2}$.

Proof of Theorem 5. Known optimal CNOT synthesis results characterize the CNOT count directly in the Weyl chamber. A circuit containing one CNOT and arbitrary local operations has the form K_1 CNOT K_2, so locally equivalent to CNOT, with Weyl coordinates $(\pi/2, 0, 0)$. A two-CNOT circuit can be written as K_1 CNOT K_2 CNOT K_3, the resulting local-equivalence classes are precisely the Weyl face $c_3 = 0$ [6]. Since three CNOTs are enough to synthesize any arbitrary two-qubit unitary, points with $c_3 > 0$ require three CNOTs. We now map these subsets into determinantal coordinates.

On the face $c_3 = 0$, we have $y = 2(1 - \cos c_2)$ and $x = 3 - \cos c_2 - \cos c_1(1 + \cos c_2)$. For a fixed c_2, the ordering $c_2 \leqslant c_1 \leqslant \pi/2$ implies $0 \leqslant \cos c_1 \leqslant \cos c_2$. Fixing $\cos c_1 = 0$ gives the upper boundary $x = 2 + y/2$, and $\cos c_1 = \cos c_2$ gives the lower boundary $x = 2y - y^2/4$. Hence, the image of the $c_3 = 0$ face is $\mathcal{R}_2$, and any point to the left of this region cannot be synthesized with two-CNOT gates and therefore requires three CNOTs.

Proof of Theorem 6. The triple (x, y, z) already determines the Schmidt spectrum, as $s_3^2 = z$, $s_2^2 = y - z$, $s_1^2 = x - y$, $s_0^2 = 4 - x$. From these, we get the products a, b, c that, away from the $c_1 = \frac{\pi}{2}$ face, allow us to recover the Weyl coordinates

$$\cos c_1 = \sqrt{\frac{ab}{c}}, \quad \cos c_2 = \sqrt{\frac{ac}{b}}, \quad \cos c_3 = \sqrt{\frac{bc}{a}}.$$

For $c_1 < \frac{\pi}{2}$, the map is then injective. However, in the face $c_1 = \frac{\pi}{2}$ we have $a = b = 0$, and the data only depend on $c = \cos c_2 \cos c_3$, so different Weyl points collapse to the same determinantal coordinates.

We now obtain the description of the determinantal chamber $\mathcal{D}$ by translating the elementary constraints on a, b, c into constraints on x, y, z. The ordering of the Schmidt coefficients gives $0 \leqslant 2z \leqslant y \leqslant x$. Moreover, $x - z = 2 - 2a$, with $a = \cos c_1 \cos c_2$. Since $a > 0$ for $c_1 < \pi/2$, we obtain $x < z + 2$. Equality is recovered only on the face $c_1 = \pi/2$, where $a = 0$.

We also have

$$x + z - 2y = 2(c - b) = 2\cos c_3(\cos c_2 - \cos c_1) \geqslant 0,$$

as $c_1 \geqslant c_2$ in the Weyl chamber, so $x + z \geqslant 2y$.

The nonlinear constraint follows directly from the Weyl-coordinate parametrization. For $c_1 < \pi/2$, we have $a > 0$, and $\frac{bc}{a} = \cos^2 c_3 \leqslant 1$, so that $bc \leqslant a$. In terms of the determinantal coordinates,

$$a = \frac{2 - x + z}{2}, \quad b = \frac{2 - x + y - z}{2}, \quad c = \frac{2 - y}{2}.$$

Therefore,

$$4(a - bc) = y^2 - y(x + z) + 4z \geqslant 0,$$

which gives the nonlinear inequality in the definition of $\mathcal{D}$. On the face $c_1 = \pi/2$, the same condition is recovered by continuity.

Conversely, suppose (x, y, z) satisfies the inequalities defining $\mathcal{D}$. From $x < z + 2$, $2z \leqslant y$, and $x + z \geqslant 2y$, we obtain $0 < a \leqslant b \leqslant c$, while $y^2 - y(x + z) + 4z \geqslant 0$ is exactly $bc \leqslant a$. These inequalities imply $0 \leqslant y \leqslant 2$, hence $c \leqslant 1$. Define $p = \sqrt{ab/c}$, $q = \sqrt{ac/b}$, and $r = \sqrt{bc/a}$. Then, $0 < p \leqslant q \leqslant r \leqslant 1$, with $pq = a$, $pr = b$, and $qr = c$. Therefore, setting $c_1 = \arccos p$, $c_2 = \arccos q$, and $c_3 = \arccos r$ gives $0 \leqslant c_3 \leqslant c_2 \leqslant c_1 < \pi/2$ and reproduces the original (x, y, z), showing that every point satisfying the inequalities lies in the image.

[1] D. P. DiVincenzo, Two-bit gates are universal for quantum computation, Phys. Rev. A **51**, 1015 (1995).
[2] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, Phys. Rev. A **52**, 3457 (1995).
[3] J. L. Dodd, M. A. Nielsen, M. J. Bremner, and R. T. Thew, Universal quantum computation and simulation using any entangling Hamiltonian and local unitaries, Phys. Rev. A **65**, 040301 (2002).
[4] S. Lloyd, Almost any quantum logic gate is universal, Phys. Rev. Lett. **75**, 346 (1995).
[5] K. Hammerer, G. Vidal, and J. I. Cirac, Characterization of nonlocal gates, Phys. Rev. A **66**, 062321 (2002).
[6] G. Vidal and C. M. Dawson, Universal quantum circuit for two-qubit transformations with three controlled-NOT gates, Phys. Rev. A **69**, 010301 (2004).
[7] Y. Makhlin, Nonlocal properties of two-qubit gates and mixed states, and the optimization of quantum computations, Quantum Inf. Proc. **1**, 243 (2002).
[8] N. Khaneja and S. J. Glaser, Cartan decomposition of $SU(2^n)$ and control of spin systems, Chem. Phys. **267**, 11 (2001).
[9] B. Kraus and J. I. Cirac, Optimal creation of entanglement using a two-qubit gate, Phys. Rev. A **63**, 062309 (2001).
[10] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley, Geometric theory of nonlocal two-qubit operations, Phys. Rev. A **67**, 042313 (2003).
[11] P. Zanardi, C. Zalka, and L. Faoro, Entangling power of quantum evolutions, Phys. Rev. A **62**, 030301 (2000).
[12] P. Zanardi, Entanglement of quantum evolutions, Phys. Rev. A **63**, 040304 (2001).
[13] M. A. Nielsen, C. M. Dawson, J. L. Dodd, A. Gilchrist, D. Mortimer, T. J. Osborne, M. J. Bremner, A. W. Harrow, and A. Hines, Quantum dynamics as a physical resource, Phys. Rev. A **67**, 052301 (2003).
[14] S. Balakrishnan and R. Sankaranarayanan, Operator-Schmidt decomposition and the geometrical edges of two-qubit gates, Quantum Inf. Proc. **10**, 449 (2011).
[15] S. Balakrishnan and R. Sankaranarayanan, Measures of operator entanglement of two-qubit gates, Phys. Rev. A **83**, 062320 (2011).
[16] S. Balakrishnan and M. Lakshmanan, Chaining property for two-qubit operator entanglement measures, Eur. Phys. J. Plus **129**, 231 (2014).
[17] J. E. Tyson, Operator-Schmidt decompositions and the Fourier transform, with applications to the operator-Schmidt numbers of unitaries, J. Phys. A: Math. Gen. **36**, 10101 (2003).
[18] M. Musz, M. Kuś, and K. Życzkowski, Unitary quantum gates, perfect entanglers, and unistochastic maps, Phys. Rev. A **87**, 022111 (2013).
[19] G. Golub, A. Hoffman, and G. Stewart, A generalization of the Eckart-Young-Mirsky matrix approximation theorem, Linear Algebra Appl. **88–89**, 317 (1987).
[20] L. H. Pedersen, N. M. Møller, and K. Mølmer, Fidelity of quantum operations, Phys. Lett. A **367**, 47 (2007).
[21] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley, Exact two-qubit universal quantum circuit, Phys. Rev. Lett. **91**, 027903 (2003).
[22] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley, Minimum construction of two-qubit quantum operations, Phys. Rev. Lett. **93**, 020502 (2004).
[23] F. Vatan and C. Williams, Optimal quantum circuits for general two-qubit gates, Phys. Rev. A **69**, 032315 (2004).
[24] V. V. Shende, I. L. Markov, and S. S. Bullock, Minimal universal two-qubit controlled-NOT-based circuits, Phys. Rev. A **69**, 062321 (2004).
[25] https://github.com/llorebaga/LR2qubit.