

Isoholonomic inequality and tight implementations of holonomic quantum gates

Ole Sönnernborn ^{*}

*Department of Mathematics and Computer Science, [Karlstad University](#), 651 88 Karlstad, Sweden
and Department of Physics, [Stockholm University](#), 106 91 Stockholm, Sweden*



(Received 18 December 2024; accepted 1 May 2025; published 16 May 2025)

In holonomic quantum computation, quantum logic gates are realized by cyclic parallel transport of the computational space. The resulting quantum gate corresponds to the holonomy associated with the closed path traced by the computational space. The isoholonomic inequality for gates establishes a fundamental lower bound on the path length of such cyclic transports, which depends only on the spectrum of the holonomy, that is, the eigenvalues of the implemented quantum gate. The isoholonomic inequality also gives rise to an estimate of the minimum time required to execute a holonomic quantum gate, underscoring the central role of the inequality in quantum computation. In this paper, we show that when the codimension of the computational space is sufficiently large, any quantum gate can be implemented using a parallel transporting Hamiltonian in a way that saturates the isoholonomic inequality and the corresponding time estimate. We call such implementations tight. The treatment presented here is constructive and lays the foundation for the development of efficient and optimal implementation strategies in holonomic quantum computation.

DOI: [10.1103/PhysRevA.111.052430](https://doi.org/10.1103/PhysRevA.111.052430)

I. INTRODUCTION

Nonadiabatic parallel transport has been proposed as an efficient technique for implementing quantum logic gates that combines the robustness of adiabatic holonomic quantum computation with enhanced resilience to noise [1–5].

A key result in holonomic quantum computation is the isoholonomic inequality, which establishes a fundamental lower bound on the path length necessary for a cyclic transformation of the computational space to realize a holonomic gate. This inequality also provides an estimate of the minimum time required to execute a holonomic quantum gate [6]. In this paper, we first review the principles of nonadiabatic holonomic quantum computation and formulate the isoholonomic inequality and its associated execution time estimate. We then show that if the codimension of the computational space is at least as large as its dimension, any quantum gate can be tightly implemented using a parallel transporting Hamiltonian. Here, “tightly” signifies that the isoholonomic inequality and the execution time estimate are saturated. The proof is constructive and explicit, making it a suitable starting point for the design of tight implementation schemes for holonomic quantum gates.

II. PARALLEL TRANSPORT AND HOLONOMIC GATES

In nonadiabatic holonomic computation, gates are implemented by cyclic parallel transport of the computational space

within an ambient Hilbert space. To accurately describe this procedure we need to introduce the Stiefel-Grassmann principal bundle.

A. The Stiefel-Grassmann bundle

Let $\mathcal{H}$ be a Hilbert space and n be a positive integer less than the dimension of $\mathcal{H}$. The Grassmann manifold $\mathcal{G}(n; \mathcal{H})$ is the manifold of all n -dimensional subspaces of $\mathcal{H}$. We provide this manifold with the smooth and Riemannian structures induced by its canonical identification with the space of orthogonal projection operators of rank n on $\mathcal{H}$, equipped with the Hilbert-Schmidt Riemannian metric.

Assume the computational space $\mathcal{R}$ is a subspace of $\mathcal{H}$ of dimension n , and thus is an element of $\mathcal{G}(n; \mathcal{H})$. A cyclic transformation of $\mathcal{R}$ corresponds to a curve in $\mathcal{G}(n; \mathcal{H})$ starting and ending at $\mathcal{R}$.¹ In this paper we focus exclusively on cyclic transformations generated by Hamiltonians, that is, closed curves in $\mathcal{G}(n; \mathcal{H})$ of the form $\mathcal{R}_t = U_t(\mathcal{R})$, where U_t is the time evolution operator of a given Hamiltonian.

An n -frame in $\mathcal{H}$ is a sequence of n orthonormal vectors in $\mathcal{H}$, and the Stiefel manifold $\mathcal{V}(n; \mathcal{H})$ is the manifold of all such frames. The linear span of an n -frame is an element of $\mathcal{G}(n; \mathcal{H})$, and the surjective map

$$\mathcal{V}(n; \mathcal{H}) \ni V \mapsto \text{span } V \in \mathcal{G}(n; \mathcal{H}) \quad (1)$$

is the projection of the Stiefel-Grassmann principal bundle.

The symmetry group of the Stiefel-Grassmann bundle is the group $U(n)$ of $n \times n$ unitary matrices, acting on $\mathcal{V}(n; \mathcal{H})$

^{*}Contact author: ole.sonnerborn@kau.se

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. Funded by [Bibsam](#).

¹In this paper we use the word *curve* to refer to a piecewise smooth one-parameter family of objects such as vectors, operators, or subspaces. We call the parameter time, denote it by t , and assume that it varies between 0 and some unspecified finite value τ .

from the right. To describe this action more explicitly we represent the elements of $\mathcal{V}(n; \mathcal{H})$ as formal row matrices,

$$V = (|v_1\rangle |v_2\rangle \cdots |v_n\rangle). \quad (2)$$

The action of a unitary matrix U on V is then given by the matrix multiplication $V \mapsto VU$. Since this multiplication replaces each component of V by a linear combination of the components of V and thereby preserves its span, the action of the symmetry group preserves the fibers of the Stiefel-Grassmann bundle.

Remark 1. In this paper, all frames are assumed to be of length n unless explicitly stated otherwise. Therefore, from now on, we will simply write *frame* and omit the reference to the number of components.

B. Parallel transport

Parallel transport ensures that the gates implemented are the result of purely geometric effects. To rigorously define the concept of parallel transport, we must first define what is meant by a curve of frames being horizontal.

We say that a tangent vector $\dot{V}$ at the frame V is horizontal if the velocity vectors that make up $\dot{V}$ are perpendicular to the span of V . In short, $\dot{V}$ is horizontal if $V^\dagger \dot{V} = 0$. Furthermore, we say that a curve of frames V_t is horizontal if its velocity vector is always horizontal,

$$V_t^\dagger \dot{V}_t = 0. \quad (3)$$

A fundamental result from the theory of principal fiber bundles ensures that, for every curve $\mathcal{R}_t$ in $\mathcal{G}(n; \mathcal{H})$ and every frame V spanning $\mathcal{R}_0$, there exists a unique horizontal curve of frames V_t with $\text{span } V_t = \mathcal{R}_t$ and $V_0 = V$. We call V_t the horizontal lift of $\mathcal{R}_t$ starting at V .

Suppose $\mathcal{R}_t = U_t(\mathcal{R})$, where U_t is the evolution operator of a Hamiltonian H_t . For any frame V spanning $\mathcal{R}$, we can define a lift V_t of $\mathcal{R}_t$ as $V_t = U_t V$, where U_t acts componentwise on V . This lift can be horizontal or not. If it is horizontal, that is, if condition (3) is satisfied at all times, we say that H_t is parallel transporting. The property of being parallel transporting does not depend on the choice of frame V but is intrinsic to H_t itself. By writing $|v_{k;t}\rangle$ for the k th component of V_t , the parallel transport condition can be written as

$$\langle v_{k;t} | H_t | v_{l;t} \rangle = 0 \quad (k, l = 1, 2, \dots, n). \quad (4)$$

C. Holonomic gates

Suppose $\mathcal{R}_t$ is a closed curve at $\mathcal{R}$, and V_t is an arbitrary horizontal lift of $\mathcal{R}_t$. Then the two frames V_0 and V_τ span $\mathcal{R}$. As a result, the unitary operator

$$\Gamma[\mathcal{R}_t] = V_\tau V_0^\dagger = \sum_{k=1}^n |v_{k;\tau}\rangle \langle v_{k;0}| \quad (5)$$

preserves $\mathcal{R}$. This operator is independent of the choice of the horizontal lift and is called the holonomy of $\mathcal{R}_t$. We define a holonomic gate to be any unitary transformation of the computational space $\mathcal{R}$ that is implemented as the holonomy of a parallel transport of $\mathcal{R}$ around a loop. If the parallel transport is unitary, $\mathcal{R}_t = U_t(\mathcal{R})$, then the restriction of the final time unitary U_τ to $\mathcal{R}$ is equal to the holonomy of $\mathcal{R}_t$

and is thus a holonomic gate. This is how holonomic gates are implemented in practice.

Remark 2. If the parallel transport condition is satisfied, $\mathcal{R}$ is transported in such a way that no time-local rotation occurs within $\mathcal{R}$. Nevertheless, any set of cotransported vectors in $\mathcal{R}$ may end up rotated relative to its initial configuration. The rotation is described by the holonomy.

D. Projective holonomic gates

Two gates differing only by a phase factor transform states in an identical way and are therefore computationally interchangeable. Consequently, one can more generally define a gate on $\mathcal{R}$ to be a projective unitary operator on $\mathcal{R}$ or, equivalently, a holomorphic isometry of the projective space of $\mathcal{R}$. Reference [7] developed a gauge theory for nonadiabatic holonomic quantum computation with projective gates. The projective parallel transport condition is given by

$$\langle v_{k;t} | H_t | v_{l;t} \rangle = \epsilon_t \delta_{kl} \quad (k, l = 1, 2, \dots, n), \quad (6)$$

where ϵ_t is any real-valued function. We recover the conventional parallel transport condition by requiring that $\epsilon_t = 0$.

III. THE ISOHOLONOMIC INEQUALITY

The isoholonomic bound of a gate Γ on an n -dimensional subspace of $\mathcal{H}$ is the quantity

$$L(\Gamma) = \sqrt{\sum_{j=1}^n \theta_j (2\pi - \theta_j)}, \quad (7)$$

where $\theta_1, \theta_2, \dots, \theta_n$ are the phases of the eigenvalues of Γ in the interval $[0, 2\pi)$. The isoholonomic inequality for gates states that the length of a closed curve $\mathcal{R}_t$ in $\mathcal{G}(n; \mathcal{H})$ is lower bounded by the isoholonomic bound of its holonomy:

$$\text{Length}[\mathcal{R}_t] \geq L(\Gamma[\mathcal{R}_t]). \quad (8)$$

The length of $\mathcal{R}_t$ is defined as

$$\text{Length}[\mathcal{R}_t] = \int_0^\tau dt \sqrt{\frac{1}{2} \text{tr}(\dot{P}_t^2)}, \quad (9)$$

where P_t is the orthogonal projection onto $\mathcal{R}_t$. The isoholonomic inequality for gates was derived in Ref. [6].

From the isoholonomic inequality one can derive a quantum speed limit for the execution time of holonomic gates: Suppose $\mathcal{R}_t$ is generated by the Hamiltonian H_t . Then the length of $\mathcal{R}_t$ can be expressed as

$$\text{Length}[\mathcal{R}_t] = \tau \langle \langle \sqrt{I(H_t; \mathcal{R}_t)} \rangle \rangle. \quad (10)$$

The factor in double angle brackets is the average of the square root of the skewness measure

$$I(H_t; \mathcal{R}_t) = -\frac{1}{2} \text{tr}([H_t, P_t]^2) \quad (11)$$

over the evolution time interval. The skewness measure quantifies both (the square of) the instantaneous speed with which $\mathcal{R}_t$ traverses Hilbert space and the extent to which $\mathcal{R}_t$ is not invariant under H_t ; see Ref. [8]. It follows from Eqs. (8) and (10) that if H_t is configured to parallel transport the computational subspace along the closed curve $\mathcal{R}_t$, thereby executing

a holonomic gate Γ , then the execution time is not less than the isoholonomic bound of Γ divided by the mean value of the square root of the skewness measure:

$$\tau \geq \frac{L(\Gamma)}{\langle \sqrt{I(H_t; \mathcal{R}_t)} \rangle}. \quad (12)$$

A. The projective isoholonomic inequality

There is also a projective version of the isoholonomic inequality which states that the length of any closed curve $\mathcal{R}_t$ in $\mathcal{G}(n; \mathcal{H})$ is bounded from below by the isoholonomic bound of its projective holonomy:

$$\text{Length}[\mathcal{R}_t] \geq L(\bar{\Gamma}[\mathcal{R}_t]). \quad (13)$$

The isoholonomic bound of a projective gate $\bar{\Gamma}$ is defined as the minimum of the isoholonomic bounds of the gates representing $\bar{\Gamma}$:

$$L(\bar{\Gamma}) = \min\{L(\Gamma) : \Gamma \in \bar{\Gamma}\}. \quad (14)$$

It was shown in Ref. [7] that if Γ is an arbitrary representative of $\bar{\Gamma}$, $\theta_1, \theta_2, \dots, \theta_n$ are the phases in $[0, 2\pi)$ of the eigenvalues of Γ , and $\theta_0 = 0$, then

$$L(\bar{\Gamma}) = \min_{0 \leq k \leq n} \left\{ \sum_{j=1}^n |\theta_j - \theta_k| (2\pi - |\theta_j - \theta_k|) \right\}. \quad (15)$$

The isoholonomic inequality for projective gates gives rise to a quantum speed limit for the execution time of projective gates similar to that in Eq. (12).

B. Review of the derivation of the isoholonomic inequality

The isoholonomic inequality for gates (8) extends and generalizes the isoholonomic inequality for states derived in Ref. [9] (see also the Supplemental Material in Ref. [6]). The latter inequality asserts that if ρ_t represents a closed curve of pure states, then its Fubini-Study length is bounded from below as follows:

$$\text{Length}[\rho_t] \geq \sqrt{\theta(2\pi - \theta)}. \quad (16)$$

The quantity θ is the Aharonov-Anandan geometric phase of ρ_t in the interval $[0, 2\pi)$.²

Let V_t , with components $|v_{k;t}\rangle$, be a horizontal lift of $\mathcal{R}_t$ initiated at an orthonormal basis for the holonomy of $\mathcal{R}_t$:

$$\Gamma[\mathcal{R}_t]|v_{k;0}\rangle = e^{i\theta_k}|v_{k;0}\rangle. \quad (17)$$

Each curve of pure states $\rho_{k;t} = |v_{k;t}\rangle\langle v_{k;t}|$ is then closed with the geometric phase θ_k . Furthermore it holds that

$$\text{Length}[\mathcal{R}_t]^2 \geq \sum_{k=1}^n \text{Length}[\rho_{k;t}]^2. \quad (18)$$

See Ref. [6] for details. The isoholonomic inequality for gates is obtained by applying the isoholonomic inequality for states to each of the terms on the right.

²The holonomy of a closed curve of pure states is an element in $U(1)$, and the Aharonov-Anandan geometric phase is defined as the argument of this holonomy modulo 2π ; see [10]. The phase θ is the unique representative of this argument in the interval $[0, 2\pi)$.

IV. TIGHT IMPLEMENTATIONS

In this section we show that the isoholonomic inequality is tight when the codimension of the computational subspace is at least as large as its dimension. More precisely, we show that under these conditions it is possible (at least in principle) to implement any quantum gate on the computational space using a parallel transporting Hamiltonian that transports the computational space along a cyclic path whose length saturates the isoholonomic inequality and hence the quantum speed limit (12). We call such an implementation a tight implementation.

A. Tight parallel generation of a geometric phase

The proof of the isoholonomic inequality suggests that we first focus on tight parallel transports of a state that generates a given geometric phase. Let $|v\rangle$ be a normalized vector in $\mathcal{H}$ and $\mathcal{E}$ be any two-dimensional subspace of $\mathcal{H}$ containing $|v\rangle$. Inspired by Ref. [11], we show that there exists a Hamiltonian H_t which acts trivially on the orthogonal complement of $\mathcal{E}$ and parallel transports $|v\rangle$ along a curve $|v_t\rangle$ with the following properties:

(1) $|v_t\rangle$ remains within $\mathcal{E}$, and for an arbitrarily chosen time $\tau > 0$ we have that $|v_\tau\rangle = e^{i\theta}|v\rangle$, where θ is any pre-determined phase in the interval $[0, 2\pi)$.

(2) $\rho_t = |v_t\rangle\langle v_t|$ is periodic with period τ and has the constant Fubini-Study speed $(2\pi\theta - \theta^2)^{1/2}/\tau$.

Since H_t is parallel transporting, that is, $\langle v_t|H_t|v_t\rangle = 0$, each loop of ρ_t acquires the geometric phase θ . Moreover, the speed of ρ_t is calibrated so that each loop saturates the isoholonomic inequality for states. The ability to confine the dynamics to a given two-dimensional subspace will be crucial in the following sections.

To construct a Hamiltonian that satisfies the above criteria let $|\epsilon_0\rangle, |\epsilon_1\rangle$ be an orthonormal basis for $\mathcal{E}$ such that

$$2\pi|\langle\epsilon_1|v\rangle|^2 = \theta, \quad (19)$$

and define the triple of Pauli-like operators σ as follows:

$$\sigma_1 = |\epsilon_0\rangle\langle\epsilon_1| + |\epsilon_1\rangle\langle\epsilon_0|, \quad (20a)$$

$$\sigma_2 = i(|\epsilon_0\rangle\langle\epsilon_1| - |\epsilon_1\rangle\langle\epsilon_0|), \quad (20b)$$

$$\sigma_3 = |\epsilon_1\rangle\langle\epsilon_1| - |\epsilon_0\rangle\langle\epsilon_0|. \quad (20c)$$

The Bloch vector of $\rho = |v\rangle\langle v|$ is the unit-length 3-vector $\mathbf{r}$ with components $r_k = \langle v|\sigma_k|v\rangle$. Define A and H as

$$A = \mathbf{a} \cdot \sigma, \quad \mathbf{a} = (0, 0, \pi/\tau), \quad (21a)$$

$$H = \Omega \cdot \sigma, \quad \Omega = \mathbf{a} - (\mathbf{a} \cdot \mathbf{r})\mathbf{r}. \quad (21b)$$

The claim is that the Hamiltonian

$$H_t = e^{-itA} H e^{itA} \quad (22)$$

meets the required criteria.

Let U_t be the time evolution operator associated with H_t , and set $|v_t\rangle = U_t|v\rangle$. The time evolution operator can be decomposed as $U_t = e^{-itA} e^{-it(H-A)}$, and the condition (21b) ensures that ρ commutes with $H - A$:

$$[H - A, \rho] = 2i((\Omega - \mathbf{a}) \times \mathbf{r}) \cdot \sigma = 0. \quad (23)$$

As a result, ρ evolves as if A were the Hamiltonian:

$$\rho_t = U_t \rho U_t^\dagger = e^{-itA} \rho e^{itA}. \quad (24)$$

Condition (21b) also ensures that H_t is parallel transporting:

$$\langle v_t | H_t | v_t \rangle = \text{tr}(H_t \rho_t) = \text{tr}(H \rho) = \boldsymbol{\Omega} \cdot \mathbf{r} = 0. \quad (25)$$

From Eqs. (21a) and (24) it follows that ρ_t is periodic with the period $\pi/|\mathbf{a}| = \tau$. Furthermore, each loop acquires the Aharonov-Anandan geometric phase θ :

$$\pi(1 + \langle v | \sigma_3 | v \rangle) = 2\pi |\langle \epsilon_1 | v \rangle|^2 = \theta. \quad (26)$$

Here we have applied a well-known formula for the geometric phase of a qubit whose dynamics is governed by a time-independent Hamiltonian (see, e.g., [12]), which is effectively the case here according to Eq. (24).

By construction, H_t acts trivially on the orthogonal complement of $\mathcal{E}$, ensuring that $|v_t\rangle$ remains in $\mathcal{E}$. Since H_t is parallel transporting, $|v_t\rangle$ is Aharonov-Anandan horizontal. We also have that $|v_\tau\rangle = e^{i\theta}|v\rangle$ since ρ_t completes a loop with the geometric phase θ in the time τ .

It remains to calculate the Fubini-Study speed of ρ_t . The square of the speed is

$$\frac{1}{2} \text{tr}(\dot{\rho}_t^2) = -\frac{1}{2} \text{tr}([A, \rho]^2) = |\mathbf{a} \times \mathbf{r}|^2. \quad (27)$$

Furthermore,

$$\begin{aligned} \tau^2 |\mathbf{a} \times \mathbf{r}|^2 &= \tau^2 |\mathbf{a}|^2 (1 - \langle v | \sigma_3 | v \rangle^2) \\ &= \pi^2 (1 - (2|\langle v | \epsilon_1 \rangle|^2 - 1)^2) \\ &= 2\pi |\langle v | \epsilon_1 \rangle|^2 (2\pi - 2\pi |\langle v | \epsilon_1 \rangle|^2) \\ &= \theta(2\pi - \theta). \end{aligned} \quad (28)$$

We conclude that ρ_t has a constant speed given by

$$\sqrt{\frac{1}{2} \text{tr}(\dot{\rho}_t^2)} = |\mathbf{a} \times \mathbf{r}| = \frac{1}{\tau} \sqrt{\theta(2\pi - \theta)}. \quad (29)$$

Figure 1 illustrates the perpendicular configuration of the Bloch vector $\mathbf{r}$ and the Rabi vector $\boldsymbol{\Omega}$ representing the state and the Hamiltonian at the time $t = 0$. As time passes, the Bloch and Rabi vectors of the state and the Hamiltonian rotate with the same angular speed, $2|\mathbf{a}|$, around the axis spanned by $\mathbf{a}$, and thus remain perpendicular.

Remark 3. If $\theta = \pi$, then $H_t = A$. This is the only nonzero geometric phase of a qubit system that can be implemented by a time-independent parallel transporting Hamiltonian. For all other nonzero phases, the Hamiltonian must be time dependent.

B. Tight implementation of a holonomic gate that has only one eigenvalue with a non-zero phase

Section V provides schemes for the tight implementation of the $\pi/8$, Hadamard, and CNOT gates. Common to these is the property that they can be tightly implemented with a Hamiltonian that acts nontrivially on only one eigenvector of the gate to be implemented. As an application of the previous section, we outline an approach for this case. The general case will be treated in detail in the next section.




FIG. 1. The Bloch vector $\mathbf{r}$ of the state and the Rabi vector $\boldsymbol{\Omega}$ of the Hamiltonian are perpendicular at $t = 0$. These vectors rotate with time around the axis spanned by the vector $\mathbf{a}$, representing the operator A . The rotation is illustrated by cones. The red cone is traced out by the Bloch vector of the evolving state, and the blue cone is traced out by the Rabi vector of the time-varying Hamiltonian. The Bloch and Rabi vectors rotate with the same angular speed and therefore remain perpendicular.

Suppose the gate to be implemented, Γ , has only one eigenvalue whose phase θ is not zero. Let $|v\rangle$ be a corresponding normalized eigenvector. To implement Γ tightly, the Hamiltonian only has to act nontrivially on the linear span of this vector and a vector $|w\rangle$ in the orthogonal complement of the computational space. We follow the strategy of the previous section and set

$$|\epsilon_0\rangle = \sqrt{1 - \frac{\theta}{2\pi}} |v\rangle - \sqrt{\frac{\theta}{2\pi}} |w\rangle, \quad (30a)$$

$$|\epsilon_1\rangle = \sqrt{\frac{\theta}{2\pi}} |v\rangle + \sqrt{1 - \frac{\theta}{2\pi}} |w\rangle. \quad (30b)$$

The overlap condition (19) is then satisfied. We define Pauli-like operators as in Eqs. (20a)–(20c), and we define A and H as in Eqs. (21a) and (21b), with the Bloch vector being that of $|v\rangle\langle v|$. The Hamiltonian is then given by

$$H_t = e^{-itA} H e^{itA}. \quad (31)$$

This Hamiltonian parallel transports $|v\rangle$ to $e^{i\theta}|v\rangle$ within the span of $|v\rangle$ and $|w\rangle$, while keeping the orthogonal complement of $|v\rangle$ in the computational space fixed. The Hamiltonian thus parallel transports the computational space in a closed curve with holonomy Γ . The length of this curve is the same as the Fubini-Study length of the evolution curve of $|v\rangle\langle v|$, which is $(2\pi\theta - \theta^2)^{1/2}$. The length thus matches the isoholonomic bound of Γ , which means that the implementation is tight.

C. Tight implementation of a holonomic gate

Assume the codimension of the computational space is at least equal to its dimension. One can then apply a tight transformation to each eigenstate of the target gate without simultaneously coupling it to other eigenstates. In the

following, we show that this selective approach allows a tight holonomic implementation of any desired gate.

Consider a target gate Γ acting on the computational space $\mathcal{R}$, which we aim to implement tightly using a parallel transporting Hamiltonian H_t . Let $\theta_1, \theta_2, \dots, \theta_n$ be the phases of the eigenvalues of Γ , chosen within the interval $[0, 2\pi)$, and let V be any frame of corresponding eigenvectors for Γ ,

$$V = (|v_1\rangle |v_2\rangle \cdots |v_n\rangle), \quad \Gamma|v_k\rangle = e^{i\theta_k}|v_k\rangle. \quad (32)$$

Furthermore, let $|w_1\rangle, |w_2\rangle, \dots, |w_n\rangle$ be an orthonormal set of vectors in the orthogonal complement of $\mathcal{R}$. As established in the previous section, for each eigenvector $|v_k\rangle$ there exists a Hamiltonian $H_{k;t}$ that evolves $|v_k\rangle$ along a trajectory $|v_{k;t}\rangle$ within the two-dimensional subspace $\mathcal{E}_k$ spanned by $|v_k\rangle$ and $|w_k\rangle$. The Hamiltonian is configured so that the projection $\rho_{k;t} = |v_{k;t}\rangle\langle v_{k;t}|$ traces a closed path over a fixed time τ , thereby accumulating the geometric phase θ_k , and the Hamiltonian leaves the orthogonal complement of $\mathcal{E}_k$ unaffected. We define the full Hamiltonian as

$$H_t = \sum_{k=1}^n H_{k;t}. \quad (33)$$

Let U_t be the time evolution operator of H_t . The curve of frames $V_t = U_t V$ has the $|v_{k;t}\rangle$ s as components and is a horizontal lift of $\mathcal{R}_t = \text{span } V_t$. This follows from H_t parallel transporting the eigenstates of Γ within mutually orthogonal subspaces of $\mathcal{H}$. Also, at time $t = \tau$, the subspace $\mathcal{R}_t$ returns to $\mathcal{R}$ because $|v_{k;\tau}\rangle = e^{i\theta_k}|v_k\rangle$, implying that V_τ spans the same space as V . The holonomy of $\mathcal{R}_t$ is given by

$$U_\tau|_{\mathcal{R}} = V_\tau V^\dagger = \sum_{k=1}^n e^{i\theta_k}|v_k\rangle\langle v_k| = \Gamma. \quad (34)$$

To determine the length of $\mathcal{R}_t$, we start by calculating the skewness of H_t relative to $\mathcal{R}_t$:

$$\begin{aligned} I(H_t; \mathcal{R}_t) &= -\frac{1}{2} \text{tr}([H_t, V_t V_t^\dagger]^2) \\ &= \sum_{k=1}^m -\frac{1}{2} \text{tr}([H_{k;t}, |v_{k;t}\rangle\langle v_{k;t}|]^2) \\ &= \sum_{k=1}^m \frac{1}{2} \text{tr}(\hat{\rho}_{k;t}^2) \\ &= \frac{1}{\tau^2} \sum_{k=1}^m \theta_k (2\pi - \theta_k). \end{aligned} \quad (35)$$

In the second line we have used the fact that H_t preserves the mutually orthogonal subspaces $\mathcal{E}_k$. We find that the length of $\mathcal{R}_t$ corresponds to the isoholonomic bound of Γ :

$$\begin{aligned} \text{Length}[\mathcal{R}_t]^2 &= \tau^2 I(H_t; \mathcal{R}_t) \\ &= \sum_{k=1}^n \theta_k (2\pi - \theta_k) \\ &= L(\Gamma)^2. \end{aligned} \quad (36)$$

This shows that the implementation of Γ is tight.

V. A COMPLETE SET OF QUBIT GATES

A complete set of qubit gates allows the approximation of any quantum gate with arbitrary precision on a computational space associated with a register of qubits. The single-qubit $\pi/8$ and Hadamard gates, together with the two-qubit CNOT gate, form a complete set [13]. Here we provide tight holonomic implementation schemes for these gates.

A. The $\pi/8$ gate

Let $|0\rangle, |1\rangle$ be the computational basis of a qubit. The $\pi/8$ gate T , which adds $\pi/4$ to the relative phase between these basis vectors, is the gate represented by the matrix

$$T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}. \quad (37)$$

This gate has eigenvalues 1 and $e^{i\pi/4}$, with phases 0 and $\pi/4$. Consequently, the isoholonomic bound of the $\pi/8$ gate is

$$L(T) = \sqrt{\frac{\pi}{4} \left(2\pi - \frac{\pi}{4} \right)} = \frac{\pi\sqrt{7}}{4}. \quad (38)$$

To implement T holonomically and tightly in the time τ , we need to find a Hamiltonian that tightly parallel transports $|1\rangle$ to $e^{i\pi/4}|1\rangle$ in the time τ while keeping $|0\rangle$ fixed. Guided by the approach described in Sec. IV B, we fix a unit vector $|2\rangle$ in the orthogonal complement of the computational space and introduce the orthonormal vectors

$$|\epsilon_0\rangle = \sqrt{\frac{7}{8}}|1\rangle - \sqrt{\frac{1}{8}}|2\rangle, \quad (39a)$$

$$|\epsilon_1\rangle = \sqrt{\frac{1}{8}}|1\rangle + \sqrt{\frac{7}{8}}|2\rangle. \quad (39b)$$

The choice of coefficients is such that Eq. (19) is satisfied. We define A and H as in Eqs. (21a) and (21b). Then $H_t = e^{-itA} H e^{itA}$ tightly implements T in the time τ . The operators A and H are represented by the matrices

$$A = \frac{\pi}{4\tau} \begin{pmatrix} -3 & \sqrt{7} \\ \sqrt{7} & 3 \end{pmatrix}, \quad (40a)$$

$$H = \frac{\pi}{4\tau} \begin{pmatrix} 0 & \sqrt{7} \\ \sqrt{7} & 0 \end{pmatrix}, \quad (40b)$$

relative to the frame $|1\rangle, |2\rangle$.

Remark 4. Sometimes the $\pi/8$ gate is defined as

$$T' = \begin{pmatrix} e^{-i\pi/8} & 0 \\ 0 & e^{i\pi/8} \end{pmatrix}. \quad (41)$$

The phases in $[0, 2\pi)$ of the eigenvalues of this matrix are $\pi/8$ and $15\pi/8$. Thus, the isoholonomic bound of T' is

$$L(T') = \frac{\pi}{4} \sqrt{\frac{15}{2}}, \quad (42)$$

which is greater than the isoholonomic bound of T . Both T and T' represent the same projective gate, but one must transport the computational space along a longer curve to implement T'

holonomically compared to T. It follows from Eq. (15) that the isoholonomic bound of the projective gate represented by T and T' is the same as that of T.

B. The Hadamard gate

The Hadamard gate is the qubit gate represented by

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad (43)$$

relative to a computational basis $|0\rangle, |1\rangle$. The Hadamard gate has eigenvalues 1 and -1 , with phases 0 and π , implying that the isoholonomic bound of the Hadamard gate is

$$L(H) = \sqrt{\pi(2\pi - \pi)} = \pi. \quad (44)$$

An eigenvector of H with eigenvalue -1 is

$$|v\rangle = \sin\left(\frac{\pi}{8}\right)|0\rangle - \cos\left(\frac{\pi}{8}\right)|1\rangle. \quad (45)$$

We fix a unit vector $|2\rangle$ in the orthogonal complement of the computational space and define

$$|\epsilon_0\rangle = \frac{1}{\sqrt{2}}(|v\rangle - |2\rangle), \quad (46a)$$

$$|\epsilon_1\rangle = \frac{1}{\sqrt{2}}(|v\rangle + |2\rangle). \quad (46b)$$

Then the overlap condition (19) is satisfied. Since the geometric phase to be implemented is π , we set $H_t = A$ as per Remark 3. With respect to the frame $|0\rangle, |1\rangle, |2\rangle$,

$$A = \frac{\pi}{\tau} \begin{pmatrix} 0 & 0 & \sin(\pi/8) \\ 0 & 0 & -\cos(\pi/8) \\ \sin(\pi/8) & -\cos(\pi/8) & 0 \end{pmatrix} \quad (47)$$

Remark 5. The projective gate represented by H has the same isoholonomic bound as H.

C. The CNOT gate

Suppose the computational space is spanned by the two-qubit computational basis $|00\rangle, |01\rangle, |10\rangle, |11\rangle$. The CNOT

gate is the entangling two-qubit gate

$$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}. \quad (48)$$

The CNOT gate has eigenvalues 1 and -1 , the former with multiplicity 3. Thus, the isoholonomic bound of CNOT is

$$L(\text{CNOT}) = \sqrt{\pi(2\pi - \pi)} = \pi. \quad (49)$$

An eigenvector with eigenvalue -1 is

$$|v\rangle = \frac{1}{\sqrt{2}}(|10\rangle - |11\rangle). \quad (50)$$

In complete analogy to the case of the Hadamard gate, CNOT can be tightly implemented with a time-independent Hamiltonian $H_t = A$, which in this case is given by the matrix

$$A = \frac{\pi}{2\tau} \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & -\sqrt{2} \\ 0 & -\sqrt{2} & 0 \end{pmatrix} \quad (51)$$

relative to $|10\rangle, |11\rangle, |2\rangle$. As before, $|2\rangle$ is an arbitrary normalized vector in the orthogonal complement of the computational space.

VI. SUMMARY

The isoholonomic inequality establishes a lower bound on the path length required to parallel transport the computational space for implementing a quantum gate holonomically. This inequality is a cornerstone of holonomic quantum computation, as it not only imposes fundamental constraints but also provides an estimate of the minimum time needed to execute a holonomic gate. As such, it is crucial for designing fast and efficient schemes for holonomic quantum computation.

In this paper, we have demonstrated that the isoholonomic inequality and its associated execution time estimate are both sharp and saturable. Specifically, we have proven that any quantum gate can be implemented using a Hamiltonian that parallel transports the computational space along a loop that exactly saturates these bounds. Our presentation, being both constructive and explicit, paves the way for the development of optimal and efficient strategies for implementing holonomic quantum gates.

DATA AVAILABILITY

No data were created or analyzed in this study.

-
- [1] P. Zanardi and M. Rasetti, Holonomic quantum computation, *Phys. Lett. A* **264**, 94 (1999).
 [2] E. Sjöqvist, D. M. Tong, L. M. Andersson, B. Hessmo, M. Johansson, and K. Singh, Non-adiabatic holonomic quantum computation, *New J. Phys.* **14**, 103035 (2012).
 [3] G. F. Xu, J. Zhang, D. M. Tong, E. Sjöqvist, and L. C. Kwek, Nonadiabatic holonomic quantum computation in decoherence-free subspaces, *Phys. Rev. Lett.* **109**, 170501 (2012).

- [4] E. Sjöqvist, V. Azimi Mousolou, and C. M. Canali, Conceptual aspects of geometric quantum computation, *Quantum Inf. Process.* **15**, 3995 (2016).
 [5] J. Zhang, T. H. Kyaw, S. Filipp, L.-C. Kwek, E. Sjöqvist, and D. Tong, Geometric and holonomic quantum computation, *Phys. Rep.* **1027**, 1 (2023).
 [6] O. Sönnernborn, Estimate of the time required to perform a nonadiabatic holonomic quantum computation, *Phys. Rev. A* **109**, 062433 (2024).

- [7] O. Sönnernborn, preceding paper, Parallel transport in rotating frames and projective holonomic quantum computation, *Phys. Rev. A* **111**, 052429 (2025).
- [8] S. Luo and Y. Sun, Skew information revisited: Its variants and a comparison of them, *Theor. Math. Phys.* **202**, 104 (2020).
- [9] N. Hörnedal and O. Sönnernborn, Tight lower bounds on the time it takes to generate a geometric phase, *Phys. Scr.* **98**, 105108 (2023).
- [10] Y. Aharonov and J. Anandan, Phase change during a cyclic quantum evolution, *Phys. Rev. Lett.* **58**, 1593 (1987).
- [11] N. Hörnedal and O. Sönnernborn, Closed systems refuting quantum-speed-limit hypotheses, *Phys. Rev. A* **108**, 052421 (2023).
- [12] E. Sjöqvist, Geometric phases in quantum information, *Int. J. Quantum Chem.* **115**, 1311 (2015).
- [13] P. O. Boykin, T. Mor, M. Pulver, V. Roychowdhury, and F. Vatan, A new universal and fault-tolerant quantum basis, *Inf. Process. Lett.* **75**, 101 (2000).