Public-Key Pseudoentanglement and the Hardness of Learning Ground State Entanglement Structure

Adam Bouland
Department of Computer Science, Stanford University, CA, USA

Bill Fefferman
Department of Computer Science, University of Chicago, IL, USA

Soumik Ghosh
Department of Computer Science, University of Chicago, IL, USA

Tony Metger
ETH Zürich, Switzerland

Umesh Vazirani
Department of Electrical Engineering and Computer Sciences,
University of California, Berkeley, CA, USA

Chenyi Zhang
Department of Computer Science, Stanford University, CA, USA

Zixin Zhou
Department of Computer Science, Stanford University, CA, USA

Abstract

Given a local Hamiltonian, how difficult is it to determine the entanglement structure of its ground state? We show that this problem is computationally intractable even if one is only trying to decide if the ground state is volume-law vs near area-law entangled. We prove this by constructing strong forms of pseudoentanglement in a public-key setting, where the circuits used to prepare the states are public knowledge. In particular, we construct two families of quantum circuits which produce volume-law vs near area-law entangled states, but nonetheless the classical descriptions of the circuits are indistinguishable under the Learning with Errors (LWE) assumption. Indistinguishability of the circuits then allows us to translate our construction to Hamiltonians. Our work opens new directions in Hamiltonian complexity, for example whether it is difficult to learn certain phases of matter.

2012 ACM Subject Classification Theory of computation → Quantum computation theory; Theory of computation → Pseudorandomness and derandomization

Keywords and phrases Quantum computing, Quantum complexity theory, entanglement

Digital Object Identifier 10.4230/LIPIcs.CCC.2024.21

Related Version *Full Version*: `https://arxiv.org/abs/2311.12017`

Funding B.F. and S.G. acknowledge support from AFOSR (FA9550-21-1-0008). This material is based upon work partially supported by the National Science Foundation under Grant CCF-2044923 (CAREER) and by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers (Q-NEXT). This research was also supported in part by the National Science Foundation under Grant No. NSF PHY-1748958. A.B., B.F., C.Z., and Z.Z. were supported in part by the DOE QuantISED grant DE-SC0020360. A.B. and C.Z. were supported in part by the U.S. DOE Office of Science under Award Number DE-SC0020266. A.B. was supported in part by the AFOSR under grant FA9550-21-1-0392. C.Z. was supported in part by the Shoucheng Zhang graduate fellowship. T.M. acknowledges support from SNSF Grant No. 200021_188541, the ETH Zurich Quantum Center, and an ETH Doc.Mobility Fellowship. U.V. was supported in part by DOE NQISRC QSA grant FP00010905, NSF QLCI Grant No. 2016245, and MURI Grant FA9550-18-1-0161.

39th Computational Complexity Conference (CCC 2024).
Editor: Rahul Santhanam; Article No. 21; pp. 21:1–21:23
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Acknowledgements We thank Rotem Arnon-Friedman, Jordan Docter, Tudor Giurgica-Tiron, Andru Gheorghiu, Hsin-Yuan Huang, Vinod Vaikuntanathan, and Thomas Vidick for helpful discussions. We thank the Simons Institute for the Theory of Computing, where some of this work was conducted.

1 Introduction

The central problem in Hamiltonian complexity is to understand the structure of ground states of local Hamiltonians and the difficulty of learning properties of them, e.g. [20, 29, 24, 6, 5]. In this work we study the following question: given a local Hamiltonian H, how difficult is it to learn about the entanglement structure of its ground state? For example, given H, can you tell if its ground state is area-law or volume-law entangled? We call such questions about the qualitative features of the entanglement structure a *Learning Ground State Entanglement Structure* (LGSES) problem. This problem is related to both condensed matter physics – where ground state entanglement structure is a central object of study – and may also shed light on questions in quantum gravity regarding how entanglement could possibly be dual to other physical quantities [12, 18].

Whereas most effort in many-body physics has been directed towards the positive side of this question, i.e. finding conditions under which properties of the ground state can be efficiently learnt or computed (see e.g. [15, 24, 21]), here our goal is to prove hardness results for the LGSES problem from cryptographic assumptions. This explores the limitations that any algorithm for such problems must run up against. In other words, hardness results for the LGSES problem point to qualitative features of Hamiltonian ground states that are inherently computationally intractable to compute.

To prove computational hardness results for the LGSES problem, we will relate it to a notion called *pseudoentanglement*, a term recently introduced in [1]. Informally, pseudoentangled state ensembles consist of low-entanglement states that masquerade as high-entanglement states to computationally bounded observers; in other words, a pseudoentangled state only looks like a high-entanglement state to bounded observers, whereas its actual (information-theoretic) entanglement is low. The main result of [1] is that it is possible to create pseudoentangled states which hide vast differences in entanglement. More concretely, they construct two ensembles of quantum states, Ψ^{high} and Ψ^{low}, such that any state $|\psi\rangle \in \Psi^{\mathrm{high}}$ has high entanglement entropy across any bipartition of the qubits and states in Ψ^{low} have low entanglement, but any computationally bounded quantum algorithm that receives a state from $\Psi^{\mathrm{high}} \cup \Psi^{\mathrm{low}}$ cannot tell which kind of state it received.

This notion of pseudoentanglement is interesting in its own right and has been used for various applications in [1], but its relevance to Hamiltonian complexity is unclear. This is because the settings are inherently different: the notion of pseudoentanglement in [1] involves a *quantum* input, namely copies of the relevant quantum states, being given to the distinguisher. If we tried to translate this into a setting involving Hamiltonian ground states, we would end up in a model where we study the properties of Hamiltonian ground states given quantum copies of these ground states, but without knowing the actual Hamiltonian itself. In contrast, in Hamiltonian complexity we assume that we know a classical description of the Hamiltonian under consideration and would like to determine properties of its ground state.

Therefore, if we want to use some notion of pseudoentanglement to prove hardness results for the LGSES problem, we need to consider a model where the distinguisher is not just given quantum copies of the high- or low-entanglement states, but rather an efficient classical description of these states. This leads to a different kind of pseudoentanglement, which we call *public-key* pseudoentanglement since the description of the states (the “key”) can be made public. This notion is already implicit in earlier work by Gheorghiu and Hoban [18], who gave a construction of public-key pseudoentanglement that we discuss in more detail below.

▶ **Definition 1** (Public-key pseudoentanglement (implicit in [18])). *Two ensembles of n-qubit poly(n)-gate quantum circuits $\{C_k\}$, $\{C'_k\}$ are public-key pseudoentangled with entanglement gap $f(n)$ vs $g(n)$ if they are computationally indistinguishable to poly-time quantum algorithms, and yet with high probability over the ensembles, $C_k|0^n\rangle$ has entanglement $\geq f(n)$ and $C'_k|0^n\rangle$ has entanglement $\leq g(n)$ across one or more cuts of the system.*

The key difference between this definition and the one in [1] is that here a classical description of the circuit used to prepare the states is given as input to the distinguisher, whereas in [1] the distinguisher only receives copies of the output state of the circuit. A distinguisher can therefore not only prepare copies of the output state, but also analyse the classical description of the circuits directly to gain additional information, making it harder to “hide” the entanglement of pseudoentangled states. Hence, public-key pseudoentanglement is a much stronger notion than that in [1], which we will call *private-key* pseudoentanglement as the circuits are hidden.

There are generally two features we care about in a pseudoentanglement construction: the entanglement gap, which we want to be as large as possible in order to hide as much entanglement as possible, and the set of cuts across which this entanglement gap holds. Informally, a more general set of cuts (i.e. bipartitions of the qubits) across which the entanglement gap holds corresponds to hiding more qualitative information about the entanglement structure of the state; this is what we are most interested in for the LGSES problem.

The pioneering work of Gheorghiu and Hoban gave the first pseudoentanglement construction with entanglement gap n vs $n - \Omega(1)$ across a single cut based on the Learning With Errors (LWE) assumption [18]. Using this they showed that it is difficult to learn *fine-grained* properties of the ground state entanglement – namely if the ground state has entanglement n vs $n - \Omega(1)$ across a fixed cut. The basic idea is that if one passes the circuit to prepare pseudoentangled states through a modified Kitaev clock construction [23, 26], the ground state of the resulting Hamiltonian has the same entanglement properties as the output state of the circuit. We emphasize that the Kitaev clock Hamiltonian encodes the circuit used to prepare the state in plaintext. Consequently, this application necessarily requires a *public-key* construction if we aim to construct the hard instance via the Kitaev clock.

However, Gheorghiu and Hoban’s construction only suffices to prove hardness of detecting very small differences in the entanglement of the ground state across a single cut. In contrast, the LGSES problem asks about *qualitative* or *coarse-grained* features of the entanglement structure, as very fine-grained properties are usually not physically relevant. This raises the following question: is it possible to get a public-key pseudoentanglement construction that hides qualitatively different entanglement structures? This would lead to natural hardness statements for the LGSES problem.

1.1 Near area-law public-key pseudoentanglement

Our first result is to construct strong forms of public-key pseudoentanglement from LWE. In particular we show it is possible to hide 1D near area-law vs volume-law entanglement.

▶ **Theorem 2** (Volume vs area-law public-key pseudoentanglement (informal))**.** *Assuming subexponential-time hardness of LWE, there exist public-key pseudoentangled ensembles with volume-law vs near area-law entanglement when the qubits are arranged on a 1D line.*

That is, in one case the states have entanglement $\Omega(\min(k, n-k))$ across any division of the qubits into k vs $n-k$ qubits (volume-law), and in the other case the entanglement of any cut is $\leq |A|\text{polylog}(n)$ where $|A|$ is the area of the cut when the qubits are arranged on a 1D line (i.e. the number of times the cut crosses the 1D line). We call this *near area-law entangled*. This is optimal because if the polylog factor were changed to a log, then these states would be efficiently distinguishable from one another by standard Matrix Product State learning algorithms [15, 24]. Hiding such qualitatively different entanglement structures requires a completely different construction from [18]. We note that while in Theorem 2 we assume subexponential-time hardness of LWE, we also show that the standard LWE assumption implies a similar result.[1] We will discuss the application of our result to Hamiltonian complexity shortly, after we present its proof sketch.

Proof sketch for single-cut pseudoentanglement. Let us first consider a single cut partitioning the qubits into sets A and B. In this case, the most natural states to consider are of the form $\sum_{x\in\{0,1\}^n} |x\rangle_A |h(x)\rangle_B$ for some function h. If one chooses h to be injective, then this state has entanglement entropy n across this cut; if one chooses h to be 2^k-to-1, then the entanglement is $n-k$. [18] used exactly these states with trapdoor claw-free functions from [13], which are functions that are either injective or 2-to-1, but the two cases are computationally hard to distinguish (given a description of the function) assuming the hardness of LWE. There are some additional subtleties arising from the fact that the trapdoor claw-free functions from [13] do not output numbers, but probability distributions, and are only approximately 2-to-1. We refer to [18] for a detailed analysis of this construction.

To increase the entanglement gap, we need to make the many-to-1 functions *more compressing*. The functions in [18, 13] additionally have a trapdoor; however, we observe that for pseudoentanglement, we can dispense with the trapdoor and only need so-called *lossy functions*: these are functions that are either injective or 2^k-to-1, but again the two kinds of functions are hard to distinguish. Starting from this observation, it turns out to be possible to combine the construction from [18] with ideas from a recent randomness generation protocol [25] to achieve an entanglement gap of n vs n^δ for any $\delta > 0$.[2]

The challenge with this approach based on [18] is that it appears fundamentally restricted to a single cut. However, our goal is to obtain pseudoentangled states with near area-law vs volume-law entanglement structure, which requires low entanglement across *exponentially many cuts* simultaneously. This would require controlling the entanglement not just between regions A and B, but also within these regions. With the approach of [18] it seems difficult to appropriately modify the state in register A without jeopardising the entanglement across the cut between A and B.

For this reason, we need a different kind of state that allows us to control the entanglement across all cuts simultaneously. It turns out that a useful class to consider are binary phase states as in [22, 14, 1], i.e., states of the form $\sum_{x\in\{0,1\}^n} (-1)^{f(x)} |x\rangle$. Our approach will be

[1] In particular, the polylog correction factor to the area-law scaling is replaced by an n^ϵ correction, where ϵ can be any constant > 0.

[2] We note that independently from and simultaneously to our work, Gheorghiu and Hoban updated their results to include a construction of this form, which achieves the aforementioned gap of n^δ vs n across a fixed cut with n qubits in one set and poly(n) in the other.

as follows: we start from a phase state with high-entanglement across every cut. We will then modify this state (in a computationally undetectable way) to have low entanglement across some particular cut. This achieves essentially the same as the [18]-based construction above.[3] However, crucially our "modification procedure" is iterable: this means that we can perform essentially the same entanglement reduction operation across many cuts in sequence and end up with a state with low entanglement across every cut.

We first describe the construction and proof for a single cut. For simplicity, let us first consider the cut between the first and second $n/2$ qubits. One can easily compute that the reduced density matrix of the first half of the state $\sum_{x\in\{0,1\}^n}(-1)^{f(x)}|x\rangle$ is $\rho \propto TT^\top$, where $T_{ij} = (-1)^{f(i\|j)}$. Here, $\|$ denotes the concatenation of two strings, and $i, j \in \{0,1\}^{n/2}$ correspond to the first and second halves of the string x, i.e., T is the truth table of the function f written out in a matrix form. By a direct calculation, one can show an upper bound on the entanglement entropy of our phase state (i.e. the von Neumann entropy of ρ) in terms of the rank of T, and a lower bound on the entanglement entropy in terms of the Frobenius norm of $TT^\top$. Our strategy will therefore be to start from a "T-matrix" for a high-entanglement state, and then perform one of two modifications in a computationally indistinguishable way: either modify T to reduce its rank to get low-entanglement states, or modify T in a way that does *not* reduce $\left\|TT^\top\right\|_2$ too much so that the entanglement of the states remains high, where $\|\cdot\|_2$ denotes the Frobenius norm, as defined in Section 2.1. These modified T-matrices then correspond to modified phase states, and our goal is to hide which of the two procedures we performed, even when handing out a classical description for preparing the corresponding states.

In [1] the idea is to apply a private key cryptographic hash function to replace rows of the matrix T with copies of other rows to reduce its rank. That is, we pick a phase function $f(x)$ which yields high entanglement [14], and then "whittle down" its entanglement by replacing $f(x)$ with $f(h(i) \| j)$ where $i, j \in \{0,1\}^{n/2}$ and $h : \{0,1\}^{n/2} \to \{0,1\}^{n/2}$ is a 2^k-to-1 function. This reduces the number of distinct rows in T (each of which is now repeated many times) and correspondingly decreases the rank of T. As a result, this procedure lowers the entanglement of ρ by k. On the other hand, if we pick h to be a 1-to-1 function, we simply permute the rows of T and do not change $\left\|TT^\top\right\|_2$. To make this public-key, we need a way of applying a 2^k-to-1 or a 1-to-1 hash function to these phase states in such a way that it is hard to tell which one was applied, even when the state preparation circuit (and therefore the source code for h) is public.

At first sight, it may seem that we can use the same idea as above: take the lossy functions from [25] and use them to reduce the rank of T. However, the phase state construction is much less flexible than states of the form $\sum_x |x\rangle|h(x)\rangle$: for the latter, the codomain of h does not matter and we can use the functions h from [25], whose codomain are probability distributions over $\mathbb{Z}_q^m$ for $m \gg n$. In contrast, for phase states we need lossy functions $h : \{0,1\}^{n/2} \to \{0,1\}^{n/2}$, i.e. the codomain has to be the same as the domain. This is a somewhat unusual requirement from a cryptographic perspective and forces us to use a custom construction of "imperfect" lossy functions based on LWE. Concretely, we first show how to create lossy functions mapping $\{0,1\}^{n/2} \to \{0,1\}^{\mathrm{poly}(n)}$ which are *exactly* injective vs 2^k-to-1 (for sufficiently large k) with high probability. This is not yet what we need, as the codomain is exponentially larger than the domain. To fix this, we compose these

[3] One advantage of this construction even for a single cut is that it allows an entanglement gap of n vs $\mathrm{poly}\log n$ for an $O(n)$-qubit state, whereas in the [18]-based construction the B-register had to have $\mathrm{poly}(n)$ qubits for a comparable entanglement gap.

functions with pairwise independent hash functions which shrink the codomain size back to $2^{n/2}$. This can only make the 2^k-to-1 functions more compressing, which is to our benefit. However, this also introduces unwanted collisions for the injective functions which might break the high-entanglement case. To deal with this, we show that the repetition pattern this produces in the matrix T is sufficiently well-behaved that the corresponding states still have high entanglement. Intuitively, this holds because even though the "injective" functions or no longer actually injective, they are still pairwise independent, which ensures enough independence in the row repetition pattern of T to give a strong lower bound on the Frobenius norm $\left\|TT^\top\right\|_2$.

From single-cut to multi-cut pseudoentanglement. As we mentioned, the advantage of the above construction is that it can be extended to pseudoentanglement across all cuts simultaneously. We now explain how this extension works at a high level. A first observation is that to achieve pseudoentangled states with 1D near area law scaling, reducing the entanglement across all $n-1$ contiguous cuts of the line to $O(\text{polylog}(n))$ suffices by strong subadditivity. Therefore, a natural approach is to perform a 1D sweep of the line, reducing the entanglement of the contiguous cuts one at a time. For reasons that will become apparent below, we perform this sweep right-to-left. The final phase function is then a complicated composition of n independent lossy functions and hash functions.

To show that this construction indeed achieves pseudoentanglement across all cuts simultaneously, we need to worry about two issues: firstly, for the low entanglement states we need to ensure that performing the entanglement reduction operation for a cut towards the left of the line does not inadvertently increase the entanglement across earlier cuts to the right, so that the low entanglement across those earlier cuts is preserved. Secondly, for the high entanglement states we need to make sure that the fact that we are using "imperfect" injective functions does not decrease the entanglement too much even after applying these functions across every cut.

The first concern is relatively easy to deal with due to the relationships between the "T-matrices" for different cuts. To see this, consider a phase state $|\psi\rangle = \sum(-1)^{f(x)}|x\rangle$ on $n+m$ qubits on a line and let $T^{n|m} \in \{\pm 1\}^{2^n \times 2^m}$ be the T-matrix for the cut between the first n and last m qubits, i.e. $T^{n|m}_{ij} = (-1)^{f(i\|j)}$ for $i \in \{0,1\}^n, j \in \{0,1\}^m$. After performing the entanglement reduction operation, this matrix will only have some smaller number R of distinct rows, each repeated many times. In the next step of the sweep (recalling that we move right to left), we consider the cut between the first $n-1$ and last $m+1$ qubits. We denote the corresponding T-matrix by $T^{n-1|m+1} \in \{\pm 1\}^{2^{n-1} \times 2^{m+1}}$. From the definition of the T-matrix, one can see that the first row of $T^{n-1|m+1}$ simply consists of the first two rows of $T^{n|m}$ stacked side by side. More generally, the i-th row of $T^{n-1|m+1}$ is simply the horizontal concatenation of rows $2i-1$ and $2i$ from $T^{n|m}$. Suppose we now reduce the rank of $T^{n-1|m+1}$ by removing some rows and duplicating others. We then need to check that if we go back to the cut $n|m$, the resulting T-matrix (denoted $\tilde{T}^{n|m}$) still has rank at most R. This is the case since the rows of $\tilde{T}^{n|m}$ consist of the first and second parts of the rows of $T^{n-1|m+1}$; since the rank reduction only repeats, but does not modify, rows in $T^{n-1|m+1}$, every row in $\tilde{T}^{n|m}$ must have already appeared in $T^{n|m}$. As a result, the subsequent rank reduction for cut $n-1|m+1$ can only decrease, not increase, the rank of the T-matrix across $n|m$. It is not too hard to see that this argument generalises to any future rank reduction operation, not just the immediately subsequent one.

The second concern is more difficult to deal with. When applying this sweep with *approximately* injective functions, the entanglement is reduced slightly each time. The rightmost (first) cut in particular has its entanglement reduced n times, so even a tiny loss could kill the entire construction. Perhaps surprisingly, we show that this is not the case, and the entanglement losses do not compound too badly. We show that different rows have different probabilities of being hashed together due to the structure of the sweep, and a careful accounting of this process reveals that not much entanglement is lost, even in the first cut. The analysis is somewhat technical and we refer to Section 3.3 for details. Finally, we note that while the construction we described here does not produce pseudorandom states ensembles (i.e. the families of pseudoentangled states, without the public key, are not necessarily pseudorandom states [22]), we can make a simple modification to our construction to ensure that this is the case.

1.2 Hardness of learning ground state entanglement structure

Our second result is to show that this public-key, area vs volume-law pseudoentanglement construction enables new applications in quantum Hamiltonian complexity. Because our public-key pseudoentanglement construction can hide qualitative features of the entanglement structure, we can show natural results for the hardness of the Learning Ground State Entanglement Structure (LGSES) problem for broad differences in entanglement structure. As we mentioned above, the main idea for turning pseudoentanglement constructions into hard instances of LGSES is to use a circuit-to-Hamiltonian construction on the state preparation circuit for the pseudoentangled state. There are a variety of circuit-to-Hamiltonian constructions and using these on our pseudoentangled states yields a variety of hardness statements for LGSES. In this paper, we consider three different constructions: a Kitaev clock construction with a binary clock, a Kitaev clock construction with a unary clock, and a customised version of a geometrically local 2D construction [3]. As we discuss in Section 1.4, an interesting open problem is whether a custom circuit-to-Hamiltonian construction that is focused purely on preserving entanglement structure (rather than QMA-hardness) can produce hard instances of the LGSES problem for more physically natural Hamiltonians.

Using a Kitaev clock construction with a binary clock [23], we get the following result (see Theorem 26 for the formal statement).

▶ **Theorem 3** (informal). *Assuming subexponential-time hardness of LWE, LGSES is intractable when the input Hamiltonian is $O(\log n)$-local on n qubits, and the goal is to decide whether the ground state is volume-law or near area-law entangled for the qubits arranged on a 1D line.*

This result follows relatively straightforwardly from the standard Kitaev clock construction. There are only two issues that need to be addressed: firstly, the ground state of the Hamiltonian in the Kitaev clock construction is the history state of the circuit, not the output state. However, our pseudoentanglement construction only provides guarantees on the entanglement structure of the output state. This problem can be addressed using a "padding trick" [26]: we can simply pad the pseudoentanglement circuit with a large (polynomial) number of identity gates at the end. This will ensure that the history state has most weight on the output state. Using continuity properties of the von Neumann entropy, this implies that the history state has the desired entanglement structure, too. The second issue is that we have no control over the entanglement within the clock register of the Hamiltonian. However, this does not matter for the coarse-grained entanglement structure of the state: since the clock register only has logarithmically many qubits, discarding it only changes the entanglement by $O(\log n)$, which is irrelevant for our $O(\text{poly} \log n)$ vs $\Omega(n)$ entanglement gap.

The Hamiltonian in Theorem 3 does not achieve constant locality because the Hamiltonian terms acting on the binary clock register require locality $\log n$. By using a unary clock instead of a binary clock, we can make the Hamiltonians in Theorem 3 have constant locality [23]. This is also what was used in [18] to study a Hamiltonian version of their entropy difference problem. However, the clock register now has $\Theta(n)$ qubits, and because it has so many qubits, the analysis from Theorem 3 no longer yields the desired entanglement gap. However, if we trace out the clock register and measure entanglement of the remaining mixed state by any operational mixed-state entanglement measure, we show that we still recover a maximal entanglement gap across any cut. Intuitively, this is because due to the padding trick, after tracing out the clock register the remaining mixed state is close in trace distance to the (pure) output state of the pseudoentanglement circuit. We refer to Theorem 28 for the formal statement.

The main downside of the Kitaev clock construction is that the resulting Hamiltonian is not *geometrically* local, i.e. even though we imagine the qubits as arranged on a 1D line in order to talk about area and volume law entanglement, the Hamiltonian itself has no inherent 1D geometrical structure. In contrast, most physical Hamiltonians are geometrically local. To obtain hard instances of the LGSES problem for geometrically local Hamiltonians we can use a more sophisticated 2D clock construction, where we account for time across one of the spatial dimensions instead of needing to add extra clock qubits to the circuit. We first state the resulting hardness statement for LGSES informally and then briefly sketch the proof. We refer to Theorem 31 for the formal statement and Section 4.3 for details of the construction.

▶ **Theorem 4.** *Assuming subexponential-time hardness of LWE, LGSES is intractable when the input Hamiltonian is* geometrically local *on a 2D grid of $n \times \mathrm{poly}(n)$ qudits with constant local dimension $d = O(1)$, and the goal is to decide whether the ground state has entanglement scaling* $\mathrm{polylog}(n)$ *or n across horizontal cuts.*

At a high level, the construction of the 2D geometrically local case is similar to before: we take padded versions of our pseudoentanglement circuits and convert them to local Hamiltonians using a 2D circuit-to-Hamiltonian construction [3]. This circuit-to-Hamiltonian construction produces a geometrically local Hamiltonian by dispensing with an explicit clock register. As a result, the ground state also does not have a clock register and is instead of the form $\sum_t V_t |\psi_t\rangle$ (with normalization), where $|\psi_t\rangle$ is the state of the circuit after time step t and V_t are isometries such that $V_t^\dagger V_{t'} = 0$ for any $t \neq t'$. In other words, similarly to the Kitaev clock construction, the ground state of the Hamiltonian has the form of a history state, with different time steps encoded in mutually orthogonal states. Since we padded the circuit with identities, we can approximate this ground state by $\sum V_t |\psi_{\mathrm{out}}\rangle$ with $|\psi_{\mathrm{out}}\rangle$ the output state of our pseudoentanglement circuit.

The challenge in bounding the entropy of the reduced states of this "history state" is that the different time steps are encoded in different bases, specified by the isometries V_t. This is in contrast to the Kitaev construction, where the intermediate states are all encoded in the same basis. As a result, when we trace out part of the state $\sum V_t |\psi_{\mathrm{out}}\rangle$, we get a state that looks very different from just the reduced state of $|\psi_{\mathrm{out}}\rangle$. With the construction of [3], we do not know how to bound the entanglement of these reduced states.

We therefore need to modify the construction from [3] to gain better control over the entropy of these reduced states. We do this by increasing the local dimension of the qudits in order to better keep track of different steps of the circuit execution. With this modified construction, we can ensure that reduced states of different $V_t |\psi_{\mathrm{out}}\rangle$ corresponding to different phases of the circuit execution are, in a certain sense, "cutwise" orthogonal The overall reduced state is now a sum of different "orthogonal" reduced states, each corresponding to a different phase of the circuit execution. We can compute the entropy of each of these individual reduced states relatively easily from the entanglement properties of our pseudoentangled states. Using cutwise orthogonality allows us relate the entropy of the overall state to the entropies of the individual reduced states that we sum over. As a result, we can compute the entropy of the overall reduced state even though all the different time steps are encoded in different bases.

1.3 Related work

We have already given a detailed discussion of the work of Gheorghiu and Hoban [18], which introduced the idea of public key-pseudoentanglement and gave the first construction, and the work of Aaronson et al. [1], which coined the term pseudoentanglement and gave a private-key construction with maximal entanglement gap across any cut.

The main motivation in [18] was to provide a hardness result for the so-called (quantum) entropy difference problem: given two (quantum) circuits, decide whose output has more entropy when acting on a uniformly random input. If the circuit depth is polynomial, these problems are known to be complete for the complexity classes QSZK and SZK, respectively [19, 10]. Gheorghiu and Hoban showed that for constant-depth circuits with unbounded fan-out or logarithmic-depth circuits with constant fan-out, both the QED and ED problems are still at least as hard as breaking LWE. In their proof, the entropy difference between the high- and low-entropy circuits was a single bit. Our improved pseudoentanglement construction implies that both ED and QED remain LWE-hard with large entropy gaps.[4] This is similar in spirit to the classical result that SZK gaps can be amplified [28].

Recently, independent and complementary work of Arnon-Friedman, Brakerski, and Vidick [7] gave a new definition of pseudoentanglement. Their definition is private-key and is natural in the context of operational tasks in quantum Shannon theory. Consequently, they focus on operational mixed-state entanglement measures across a single cut and require their states to be efficiently preparable under LOCC. In contrast in our work we focus on creating *public-key* pseudoentanglement with different large-scale geometrical structures, which is driven by our applications in Hamiltonian complexity.

Finally, we discuss the relationship between the LGSES problem and existing algorithms for properties of ground states. While the physics literature on computing properties of ground states is too vast to survey here, we highlight two results closer to computer science. First, in [24] the authors provide a polynomial-time algorithm that, given a classical description of a one-dimensional geometrically local Hamiltonian with constant spectral gap, outputs an MPS description of the ground state. Therefore 1D constant gapped geometrically local Hamiltonians cannot "hide" anything about their ground state entanglement structure. We note that this algorithm cannot be used on the Hamiltonians we construct in this paper as they are neither 1D geometrically local nor have constant spectral gap. Therefore our results limit potential further improvements to their algorithm. Second, the recent result [21] considers the problem of distinguishing phases of matter given labelled examples of states

[4] This result only requires our single-cut pseudoentanglement construction, for which the depth can be made logarithmic as in [18]. In fact, as mentioned earlier, an independently updated version of [18] also achieves single-cut pseudoentanglement with a large gap, implying the same hardness result for the (Q)ED problem that we obtain from our construction, although their circuits have poly(n) output qubits for an entropy gap of n^δ vs n, whereas ours only have $O(n)$ output qubits, i.e. achieve a larger relative gap. We refer to [18] for a more detailed analysis.

in different phases. The authors show that if there is a constant spectral gap and the separation between the phases is sufficiently well-conditioned[5], then a classical algorithm can efficiently learn to distinguish between the phases using information from only few-body measurements. In condensed matter physics, different qualitative entanglement structures are often associated with different quantum phases of matter; therefore our result also limits potential further improvements to such algorithms, i.e. it is not possible to relax some of their assumptions e.g. to gapless phases. An interesting direction for future work is to make our pseudoentanglement Hamiltonians "more physical" to be closer to the assumptions of theses algorithmic settings. This would help to better delineate the boundary between tractability vs intractability of learning properties of ground states of local Hamiltonians.

1.4 Discussion and open questions

In this work, we have introduced and studied the Learning Ground State Entanglement Structure (LGSES) problem: given a classical description of a local Hamiltonian, determine qualitative properties of the entanglement of its ground state, e.g. whether it is area-law or volume-law entangled. To prove hardness results for this problem, we have related it to a notion that we call public-key pseudoentanglement: low-entanglement states that are computationally indistinguishable from high-entanglement states even when given the state preparation circuit. Our main technical contribution is to construct public-key pseudoentanglement with (near) area-law vs volume-law scaling assuming the hardness of LWE.

Psedoentanglement is a relatively new idea with many avenues for future work. We suggest three main directions: (i) improving pseudoentanglement constructions themselves, (ii) strengthening the link between pseudoentanglement and condensed matter physics, and (iii) applications of pseudoentanglement beyond Hamiltonian complexity. We briefly discuss each in turn.

(i) Our public-key pseudoentanglement construction achieves essentially optimal parameters, but its construction uses a fairly involved iterated entanglement reduction procedure. In contrast, the private-key construction from [1] is very simple and based on subset states. It would be desirable to have a similarly straightforward construction of public-key pseudoentanglement, too. Furthermore, as we suggested in our discussion of [7], one can extend our definition of public-key pseudoentanglement to include a trapdoor that allows for efficient distillation of the "hidden" entanglement. It is not obvious how to extend our construction to include this feature.

(ii) Our hardness results for the LGSES problem use Hamiltonians that differ from the Hamiltonians typically studied in condensed matter physics. For example, while we do prove a hardness result for the LGSES problem for 2D geometrically local Hamiltonians, this only holds for a certain set of cuts across the system. We expect that these results can be improved to be closer to the settings studied in condensed matter physics, such as to geometrically local Hamiltonians with more natural entanglement structures and larger spectral gaps,[6] which would have implications for the hardness of studying

[5] In particular, there must be a well-behaved function of few-body observables that separates the phases.

[6] Of course, we cannot hope to construct hard instances of the LGSES problem where both the area and volume law Hamiltonians are geometrically local and have constant spectral gap. This is simply because the area law (proven in 1D [20] and in 2D under extra conditions [4], but widely believed to hold generally) requires *any* such Hamiltonian to have area law entanglement. However, this does not rule out computationally indistinguishable families of Hamiltonians where the area law Hamiltonian has constant gap and the volume law Hamiltonian has inverse polynomial gap, since determining the spectral gap itself is computationally infeasible [16, 9].

quantum phases of matter. This may require developing new sorts of clock constructions where the only goal is to preserve entanglement structure of BQP computations rather than to encode more general QMA-complete problems.

(iii) While we have focused on applications in Hamiltonian complexity in this work, pseudoentanglement might be a useful tool for proving hardness results in other domains, too. For example, recent work [11, 7] has analysed the computational resources required to execute certain tasks from quantum Shannon theory, e.g. entanglement distillation. As observed in [7], proving hardness results for such problems is closely related to pseudoentanglement, and we hope that our construction of public-key pseudoentanglement will lead to additional and stronger hardness results in this direction.

Furthermore, public-key pseudoentanglement might also be interesting from a quantum cryptographic point of view, in particular its trapdoor-variant we suggested above. For example, recent work has focused on finding minimal assumptions for quantum cryptography (see e.g. [32] and references therein for an overview), and it would be interesting to explore how pseudoentanglement is related to these assumptions.

Finally, it is natural to ask if public-key pseudoentanglement might have applications in quantum gravity. The AdS/CFT correspondence postulates that gravitational theories are dual to quantum mechanical systems, and that the entanglement structure of the quantum system is related to the geometry of the gravitational system [27]. Our results show that it is difficult to estimate the entanglement of quantum states. In contrast, geometry seems to be easy to determine, which might provide an argument that this duality is exponentially hard to compute, as first suggested in [12]. Indeed this was part of the motivation for prior works of pseudoentanglement [18, 1, 7]. Our public-key extension might allow one to argue about hardness of different versions of the duality, e.g. the duality remains hard to compute even if given a parent Hamiltonian for the quantum state.

2 Preliminaries

2.1 Notation

We write $[n]$ for the set $\{1, \ldots, n\}$. For a bitstring $x \in \{0,1\}^n$, we denote the m most and least significant bits by $\mathrm{MSB}_m(x)$ and $\mathrm{LSB}_m(x)$, respectively. We denote the concatenation of strings by $x \parallel y$. For a set of indices $I \subset [n]$ and bitstrings $x \in \{0,1\}^{|I|}$, $y \in \{0,1\}^{n-|I|}$ we denote by $z = x \parallel_I y$ the string z that equals x in indices in I and y on indices in $[n] \setminus I$. We will occasionally think of a bitstring as a $\mathbb{Z}_2$-vector, in which case we denote it as $\vec{x}$.

For a matrix $A \in \mathbb{C}^{m \times n}$, we denote by $\|A\|_p = \mathrm{Tr}\big[(A^\dagger A)^{p/2}\big]^{1/p}$ its Schatten p-norm. The 1-norm is also called the trace norm, the 2-norm is the Frobenius norm (or Hilbert-Schmidt norm), and the ∞-norm the operator norm.

Quantum systems are denoted by capital letters A, B, etc. For a pure state $|\psi\rangle_{AB}$ or a mixed state ρ_{AB} on systems A and B, we denote the reduced states on system A by ψ_A and ρ_A, respectively. We write quantum circuits as $\mathsf{C} = U_T \cdot U_{T-1} \cdots U_1$, where U_i are elementary gates. This should be thought of as a list of gates, not simply a large unitary; in particular, inserting identity gates into the circuit does change the circuit (although of course it does not change the unitary implemented by the circuit).

2.2 Independent hash functions

▶ **Definition 5** (r-wise independent function family). *A function family $H = \{h_k : [N] \to [M]\}_{k \in \mathcal{K}}$ indexed by some set of keys $\mathcal{K}$ is r-wise independent if for all distinct $x_1, \ldots, x_r \in [N]$, the random variables $h_k(x_1), \ldots, h_k(x_r)$ (for $k \in \mathcal{K}$ chosen uniformly) are uniform i.i.d.*

The following is a standard result (see e.g. [30, Corollary 3.34]):

▶ **Lemma 6.** *For any $n, m, r \in \mathbb{N}$, there exists an r-wise independent function family $H_n = \{h_k : \mathbb{Z}_q^n \to \mathbb{Z}_q^m\}_{k \in \mathcal{K}}$ such that each $k \in \mathcal{K}$ has length $\mathrm{poly}(n, m, r, \log q)$ and given $k \in \mathcal{K}$, the function h_k can be evaluated in time $\mathrm{poly}(n, m, r, \log q)$.*

2.3 Entropies

We recall the basic definitions of quantum entropies. Throughout, we use the convention that $0 \log 0 = 0$.

▶ **Definition 7** (von Neumann entropy). *The von Neumann entropy of a quantum state ρ is defined as*

$$S(\rho) = -\mathrm{Tr}[\rho \log \rho] \ .$$

▶ **Definition 8** (Conditional von Neumann entropy). *The conditional von Neumann entropy of a quantum state ρ_{AB} is defined as*

$$S(\rho_{A|B}) = S(\rho_{AB}) - S(\rho_B).$$

▶ **Definition 9** (Binary entropy function). *The binary entropy function is defined as*

$$h(x) = -x \log x - (1 - x) \log(1 - x),$$

for $x \in [0, 1]$.

2.3.1 Continuity properties

▶ **Lemma 10** (Continuity of the von Neumann entropy [17, 8]). *Let ρ_{AB} and σ_{AB} be the density matrix of two n-qubit quantum states respectively, each partitioned into subsystems A and B, and let*

$$\frac{1}{2} ||\rho_{AB} - \sigma_{AB}||_1 \leq \epsilon.$$

Then,

$$|S(\rho_{AB}) - S(\sigma_{AB})| \leq \epsilon \cdot n + h(\epsilon),$$

where $h(\cdot)$ is the binary entropy function.

▶ **Lemma 11** (Continuity of the conditional von Neumann entropy [31]). *Let ρ_{AB} and σ_{AB} be the density matrix of two n-qubit quantum states respectively, each partitioned into subsystems A and B, and let*

$$\frac{1}{2} ||\rho_{AB} - \sigma_{AB}||_1 \leq \epsilon.$$

Then,

$$|S(\rho_{A|B}) - S(\sigma_{A|B})| \leq 2\epsilon \cdot \log |A| + (1 + \epsilon) \cdot h\left(\frac{\epsilon}{1 + \epsilon}\right),$$

where $|A|$ is the dimension of the Hilbert space for subsystem A and $h(\cdot)$ is the binary entropy function.

2.4 Entanglement measures

2.4.1 Pure state entanglement measure

For pure states on systems AB, the entanglement between A and B is quantified using the so-called entanglement entropy, which is simply the von Neumann entropy of the reduced state on either subsystem.

▶ **Definition 12** (Entanglement entropy). *For a* pure *state $|\psi\rangle_{AB}$, the entanglement entropy between systems A and B is defined as $S(\psi_A)$. Note that this is invariant under swapping A and B since for a pure state $|\psi\rangle_{AB}$, $S(\psi_A) = S(\psi_B)$.*

2.4.2 Entanglement entropy for phase states

▶ **Definition 13** (T-matrix associated with phase states). *Let $s : \{0,1\}^n \to \{0,1\}$. For an n-qubit phase state*

$$|\psi\rangle = \sum_x (-1)^{s(x)} |x\rangle$$

and a subset $X \subseteq [n]$, we define the "T-matrix" with respect to the cut X as a $\{\pm 1\}^{2^{|I|} \times 2^{n-|I|}}$-matrix with entries

$$T_{ij} = (-1)^{s(i \|_X j)} ,$$

where $\|_X$ is the "index string concatenation" defined in Section 2.1.

▶ **Lemma 14.** *Let $s : \{0,1\}^n \to \{0,1\}$ and $|\psi\rangle = \sum_x (-1)^{s(x)} |x\rangle$. Then for any cut $X \subseteq [n]$, the entanglement entropy of that cut is bounded by*

$$-\log\left(\left\|\frac{1}{2^n} T T^\top\right\|_2\right) \le S(\psi_X) \le \log \operatorname{rank}(T) ,$$

where T is the T-matrix of $|\psi\rangle$ across cut X.

3 Public-key pseudoentanglement: definition and construction

In this section, we define and construct public-key pseudoentanglement. Our construction uses a similar idea as in [1, Appendix A], which is to consider phase states whose phases have been manipulated in a particular way to create high or low entanglement. The "manipulation" of these phases is by means of applying a one-to-one or many-to-one function (see Section 3.2 for details). We therefore need to construct such *lossy functions* with the appropriate parameters, which we do in Section 3.1 based on the LWE assumption.

In Section 3.2, we then use these lossy functions to construct indistinguishable families of quantum states where states in one family have high entanglement and states in the other family have low entanglement across a *single fixed bipartition* of the qubits. In Section 3.3, we extend this construction to states that have high or low entanglement *for (almost) every cut on a 1Dimensional line*, i.e. we imagine the qubits of the state being arranged on a line and consider all bipartitions into left and right qubits. We show that under the subexponential-time LWE assumption, it is possible to construct pseudoentangled states of this form where either all cuts a have a linear or a polylogarithmic amount of entanglement, which is the largest possible separation, as discussed in Remark 19. In this sense our public-key pseudoentanglement construction is optimal. We will use this multi-cut construction in Section 4 to show that learning the ground state entanglement structure of (classically described) local Hamiltonians is computationally hard (under the LWE assumption).

3.1 Construction of lossy functions

We define the following rounding function for $\mathbb{Z}_q$ elements.

▶ **Definition 15.** *For $q = cp$ with $c \in \mathbb{N}$, divide $\mathbb{Z}_q$ into p consecutive bins. We define $\lfloor x \rceil_p \in \mathbb{Z}_p$ as the index of the bin in which x lies. For a vector $x \in \mathbb{Z}_q^m$, $\lfloor \vec{x} \rceil_p \in \mathbb{Z}_p^m$ is defined as the element-wise application of $\lfloor \cdot \rceil_p$.*

▶ **Definition 16** (Lossy function construction)**.** *Choose parameters $\ell(m), r(m) \leq \mathrm{poly}(m)$. Let $p = 2^4$, $q = 2^m$, and $\sigma = q/m^3$. Let $H_m = \{h_k : \mathbb{Z}_p^m \to \mathbb{Z}_2^m\}_{k \in \mathcal{K}_m^{\mathrm{hash}}}$ be the $r(m)$-wise independent function family from Lemma 6. We define two families of functions $f : \{0,1\}^m \to \{0,1\}^m$ indexed by key sets $\mathcal{K}_m^{\mathrm{inj}}, \mathcal{K}_m^{\mathrm{lossy}} \subset \mathbb{Z}_q^{m \times m} \times \mathcal{K}_m^{\mathrm{hash}}$ as follows:*

- *To sample a key from $\mathcal{K}_m^{\mathrm{inj}}$, denoted $k \leftarrow \mathcal{K}_m^{\mathrm{inj}}$, sample $A \leftarrow U_q^{m \times m}$ and $k^{\mathrm{hash}} \in \mathcal{K}_m^{\mathrm{hash}}$ uniformly. Set $k = (A, k^{\mathrm{hash}})$.*
- *To sample a key from $\mathcal{K}_m^{\mathrm{lossy}}$, denoted $k \leftarrow \mathcal{K}_m^{\mathrm{lossy}}$, sample $A \leftarrow L_{q,\ell,\sigma}^{m \times m}$ and $k^{\mathrm{hash}} \in \mathcal{K}_m^{\mathrm{hash}}$ uniformly. Set $k = (A, k^{\mathrm{hash}})$.*
- *For a key $k = (A, k^{\mathrm{hash}}) \in \mathcal{K}_m^{\mathrm{inj}} \cup \mathcal{K}_m^{\mathrm{lossy}}$, the function $f_k : \{0,1\}^m \to \{0,1\}^m$ is defined as*

$$f_k(\vec{x}) = h_{k^{\mathrm{hash}}}\left(\lfloor A \cdot \vec{x} \rceil_p\right) .$$

3.2 Public-key pseudoentanglement across a single cut

We will now use our lossy function construction from Section 3.1 to construct public-key pseudoentangled states for the middle cut that separates the left and right half of qubits. In Section 3.3 we will use the ideas from this section in an iterated way to construct pseudoentangled states for qubits arranged on a line that are pseudoentangled across every cut on the line. Strictly speaking, all the results in this section follow from the more general analysis in Section 3.3. We spell them out nonetheless because it may be easier for readers to first understand the single-cut construction in detail before moving on to Section 3.3.

We begin by defining single-cut public-key pseudoentanglement formally.

▶ **Definition 17** (Public-key pseudoentanglement across a single cut)**.** *A public-key pseudoentangled state ensemble with entanglement gap $(f(n), g(n))$ across cuts $X_n \subset [n]$ consists of two sequences of families of quantum states $\Psi_n^{\mathrm{low}} = \{|\psi_k\rangle\}_{k \in \mathcal{K}_n^{\mathrm{low}}}$ and $\Psi_n^{\mathrm{high}} = \{|\psi_k\rangle\}_{k \in \mathcal{K}_n^{\mathrm{high}}}$ indexed by key sets $\mathcal{K}_n^{\mathrm{low}}$ and $\mathcal{K}_n^{\mathrm{high}}$ respectively with the following properties:*

(i) *Every $|\psi_k\rangle \in \Psi_n^{\mathrm{low}} \cup \Psi_n^{\mathrm{high}}$ is an n-qubit state.*

(ii) *Every key $k \in \mathcal{K}_n^{\mathrm{low}} \cup \mathcal{K}_n^{\mathrm{high}}$ has length $\mathrm{poly}(n)$, and there exists an efficient sampling procedure that, given as input n and a label "high" or "low", outputs a key $k \in \mathcal{K}_n^{\mathrm{low}}$ or $k \in \mathcal{K}_n^{\mathrm{high}}$, respectively. We write $k \leftarrow \mathcal{K}_n^{\mathrm{low}}$ and $k \leftarrow \mathcal{K}_n^{\mathrm{high}}$ for keys sampled according to this procedure.*

(iii) *Given $k \in \mathcal{K}_n^{\mathrm{low}} \cup \mathcal{K}_n^{\mathrm{high}}$, the corresponding state $|\psi_k\rangle$ is efficiently preparable (without knowing whether $k \in \mathcal{K}^{\mathrm{low}}$ or $k \in \mathcal{K}^{\mathrm{high}}$). Formally, there exists a uniform polynomial-time circuit family $\{C_n\}$ such that C_n takes as input a key $k \in \mathcal{K}_n^{\mathrm{low}} \cup \mathcal{K}_n^{\mathrm{high}}$ and outputs a state negligibly close to $|\psi_k\rangle$.*

(iv) *The keys from $\mathcal{K}_n^{\mathrm{low}}$ and $\mathcal{K}_n^{\mathrm{high}}$ are computationally indistinguishable. Formally, for all $\mathrm{poly}(n)$-time quantum adversaries $\mathcal{A}$ that take as input a key $\mathcal{K}_n^{\mathrm{low}} \cup \mathcal{K}_n^{\mathrm{high}}$ and output a single bit:*

$$\left| \Pr_{k \leftarrow \mathcal{K}_n^{\mathrm{low}}}[\mathcal{A}(k) = 0] - \Pr_{k \leftarrow \mathcal{K}_n^{\mathrm{high}}}[\mathcal{A}(k) = 0] \right| = \mathrm{negl}(n) .$$

(v) *With overwhelming probability, across the cut X_n states in Ψ_n^{low} have entanglement entropy $\Theta(f(n))$ and states in Ψ_n^{high} have entanglement entropy $\Theta(g(n))$. Formally, there exist constants $0 < C_1 < C_2$ such that for all sufficiently large n,*

$$\Pr_{k \leftarrow \mathcal{K}_n^{\mathrm{low}}}[S((\psi_k)_{X_n}) \in [C_1 f(n), C_2 f(n)]] \geq 1 - \mathrm{negl}(n)\,,$$
$$\Pr_{k \leftarrow \mathcal{K}_n^{\mathrm{high}}}[S((\psi_k)_{X_n}) \in [C_1 g(n), C_2 g(n)]] \geq 1 - \mathrm{negl}(n)\,.$$

Here, $S((\psi_k)_{X_n})$ is the von Neumann entropy of the reduced state of $|\psi_k\rangle$ on the qubits in the set $X_n \subset [n]$.

▶ Remark 18. We will frequently abuse notation and use entanglement gaps of the form $(O(f(n)), \Omega(g(n)))$. By this, we mean that (across a specified cut X_n) Ψ_n^{low} has entanglement at most $O(f(n))$ and Ψ_n^{high} has entanglement at least $\Omega(g(n))$. Formally, this means that for this case Item v of Definition 17 has to be modified as follows:

$$\Pr_{k \leftarrow \mathcal{K}_n^{\mathrm{low}}}[S((\psi_k)_{X_n}) \leq O(f(n))] \geq 1 - \mathrm{negl}(n)\,,$$
$$\Pr_{k \leftarrow \mathcal{K}_n^{\mathrm{high}}}[S((\psi_k)_{X_n}) \geq \Omega(g(n))] \geq 1 - \mathrm{negl}(n)\,.$$

▶ Remark 19. A natural question is what the optimal entanglement gap for pseudoentangled states is. Clearly, the high-entanglement states can have entanglement at most $g(n) = O(n)$ across any cut, since the entanglement entropy is upper bounded by the number of qubits. For the low-entanglement states, one can show that the entanglement entropy needs to scale faster than $\log n$, i.e. $f(n) = \omega(\log n)$. Otherwise, one could distinguish the low-entanglement states from the high-entanglement states using a variant of the SWAP test. This was proven in [22] and [1, Appendix F] for private-key pseudoentangled states and the same proof applies to the public-key setting, too.

We now give a construction of single-cut pseudoentangled states based on our lossy function construction from Section 3.1. As we will show in Theorem 21, these states do indeed form pseudoentangled ensembles in the sense of Definition 17. Under the standard LWE assumption, we can achieve an entanglement gap of $(O(n^\delta), \Omega(n))$ for any $\delta > 0$, where n is the number of qubits and the cut divides the qubits into two equal halves; under the stronger subexponential-time LWE assumption we can achieve an entanglement gap of $(O(\mathrm{poly}\log n), \Omega(n))$, which is essentially optimal as noted in Remark 19.

For simplicity, for the rest of this subsection we always assume that n is even and write $m = n/2$. We will consider the cut that divides the qubits into two sets of size n; we could also consider more general cuts and treat them with the same technique, which we do in Section 3.3.

▶ **Definition 20.** *Fix a function $f(n)$. (This will be treated as a parameter of the construction.) Let $H_n = \{h_k : \{0,1\}^n \to \{0,1\}\}_{k \in \mathcal{K}_n^4}$ be a 4-wise independent family as given in Lemma 6. Instantiate the lossy functions from Section 3.1 with parameters $\ell(m) = \sqrt{f(2m)}$ and $r(m) = 2$. (Recall that $m := n/2$.)*

We first describe the sampling procedure for the keys $\mathcal{K}_n^{\mathrm{low}}$ and $\mathcal{K}_n^{\mathrm{high}}$.

(i) *To sample $k \in \mathcal{K}_n^{\mathrm{low}}$, sample $k^{\mathrm{rep}} \leftarrow \mathcal{K}_m^{\mathrm{lossy}}$ and $k^{\mathrm{fin}} \in \mathcal{K}_n^4$ uniformly. Set $k = (k^{\mathrm{rep}}, k^{\mathrm{fin}})$.*

(ii) *To sample $k \in \mathcal{K}_n^{\mathrm{high}}$, sample $k^{\mathrm{rep}} \leftarrow \mathcal{K}_m^{\mathrm{inj}}$ and $k^{\mathrm{fin}} \in \mathcal{K}_n^4$ uniformly. Set $k = (k^{\mathrm{rep}}, k^{\mathrm{fin}})$.*

For $k = (k^{\mathrm{rep}}, k^{\mathrm{fin}})$, define the labelling function *$r_k : \{0,1\}^n \to \{0,1\}^n$ by*

$$r_k(x) = (f_{k^{\mathrm{rep}}}(i) \parallel j) \qquad \text{with } i = \mathrm{MSB}_m(x), j = \mathrm{LSB}_m(x)\,.$$

We next define the function $s_k : \{0,1\}^n \to \{0,1\}$ as

$$s_k(x) = h_{k^{\mathrm{fin}}}(r_k(x))\,.$$

The states $|\psi_k\rangle$ are then given by

$$|\psi_k\rangle = \sum_{x\in\{0,1\}^n} (-1)^{s_k(x)}|x\rangle\,.$$

▶ **Theorem 21.**

(i) *Under the standard LWE assumption, for any function $f(n) = n^\delta$ for $\delta > 0$, the state families $\Psi_n^{\text{low}} = \{|\psi_k\rangle\}_{k\in\mathcal{K}_n^{\text{low}}}$ and $\Psi_n^{\text{high}} = \{|\psi_k\rangle\}_{k\in\mathcal{K}_n^{\text{high}}}$ from Definition 20 form a pseudoentangled state ensemble with entanglement gap $(O(f(n)), \Omega(n))$ across the cuts $X_n = [n/2]$.*

(ii) *Under the subexpoential-time LWE assumption, there exists a function $f(n) = \operatorname{poly}\log n$ such that the state families $\Psi_n^{\text{low}} = \{|\psi_k\rangle\}_{k\in\mathcal{K}_n^{\text{low}}}$ and $\Psi_n^{\text{high}} = \{|\psi_k\rangle\}_{k\in\mathcal{K}_n^{\text{high}}}$ from Definition 20 form a pseudoentangled state ensemble with entanglement gap $(O(f(n)), \Omega(n))$ across the cuts $X_n = [n/2]$.*

3.3 Area-law public-key pseudoentangled states on a 1D line

In this section we will give a (nearly) area-law public-key pseudoentangled states construction on a line based on the row repetition technique we introduced in Section 3.2. This means that we will construct public-key pseudoentangled states such that if we imagine the qubits arranged on a line, the entanglement gap is $\operatorname{poly}\log n$ vs n across all cuts that separate the qubits into left and right qubits on the line.

We begin by formally defining this multi-cut version of pseudoentanglement. The definition is almost identical to Definition 17, except that we now need to require an entanglement gap across all cuts on a line simultaneously. One slight subtlety is that if we consider cuts close to the end of the line, the entanglement will be low simply by virtue of the fact that there are only very few qubits on one side of the cut. Therefore, in the high-entanglement case, we need to require the entanglement to be at least g(distance from end of line) rather than simply $g(n)$. Furthermore, very close to the boundary (namely, $O(\log n)$ close), certain properties of our construction break down. Therefore, we only consider cuts that are at least $\omega(\log n)$ far from the boundary. Since we are primarily interested in large entanglement gaps of the form $(O(\operatorname{poly}\log n), \Omega(n))$, this small boundary region is of no particular interest to us. Nonetheless, it is possible to modify our construction to work for such small boundary regions too.

As in the single-cut case, we use entanglement gaps of the form $(O(f(n)), \Omega(g(n)))$ for pseudoentangled states where we only have an upper bound on the entanglement in the low-entanglement case and a lower bound in the high-entanglement case (see Remark 18 for details). To simplify the definition slightly, below we state the definition directly for this case; it is straightforward to adapt it to the case where one wants the exact scaling rather than one-sided bounds, but we will not need this for our results.

▶ **Definition 22** (Public-key pseudoentanglement across geometrically local cuts in 1D). *A public-key pseudoentangled state ensemble with entanglement gap $(O(f(n)), \Omega(g(n)))$ across geometrically local cuts on a 1D line consists of two sequences of families of quantum states $\Psi_n^{\text{low}} = \{|\psi_k\rangle\}_{k\in\mathcal{K}_n^{\text{low}}}$ and $\Psi_n^{\text{high}} = \{|\psi_k\rangle\}_{k\in\mathcal{K}_n^{\text{high}}}$ indexed by key sets $\mathcal{K}_n^{\text{low}}$ and $\mathcal{K}_n^{\text{high}}$ respectively that satisfy Items i–iv from Definition 17 and the following modified version of Item v from Definition 17:*

(v') *For any function $b(n) = \omega(\log n)$, with overwhelming probability, states in Ψ_n^{low} have entanglement entropy $O(f(n))$ and states in Ψ_n^{high} have entanglement entropy $\Omega(g($distance from end of line$))$ for all geometrically local cuts that are at least $b(n)$ far from the end of the line. Formally,*

$$\Pr_{k \leftarrow \mathcal{K}_n^{\text{low}}}\left[\forall c \in \{b(n), \ldots, n-b(n)\} : \ S((\psi_k)_{[c]}) \leq O(f(n))\right] \geq 1 - \text{negl}(n)\,,$$
$$\Pr_{k \leftarrow \mathcal{K}_n^{\text{high}}}\left[\forall c \in \{b(n), \ldots, n-b(n)\} : \ S((\psi_k)_{[c]}) \geq \Omega\Big(\min(g(c), g(n-c))\Big)\right] \geq 1 - \text{negl}(n)\,.$$

Here, $(\psi_k)_{[c]}$ is the reduced states of $|\psi_k\rangle$ on qubits $(1, \ldots, c)$.

▶ **Definition 23.** *Fix a function $f(n)$ (this will be treated as a parameter of the construction). Let $H_n = \{h_k : \{0,1\}^n \to \{0,1\}\}_{k \in \mathcal{K}_n^4}$ be a 4-wise independent family as given in Lemma 6. For $m \in \{f(n), f(n)+1, \ldots, n\}$, instantiate the m-bit lossy function from Section 3.1 with parameters $\ell(m) = \sqrt{f(n)}$ and $r(m) = 2$.*

We first describe the sampling procedure for the keys $\mathcal{K}_n^{\text{low}}$ and $\mathcal{K}_n^{\text{high}}$.

(i) *To sample $k \in \mathcal{K}_n^{\text{low}}$, for $m \in \{f(n), f(n)+1, \ldots, n\}$, independently sample $k_m^{\text{rep}} \leftarrow \mathcal{K}_m^{\text{lossy}}$ (see Definition 16) and $k^{\text{fin}} \in \mathcal{K}_n^4$ uniformly. Set $k = (k_{f(n)}^{\text{rep}}, k_{f(n)+1}^{\text{rep}}, \ldots, k_n^{\text{rep}}, k^{\text{fin}})$.*

(ii) *To sample $k \in \mathcal{K}_n^{\text{high}}$, for $m \in \{f(n), f(n)+1, \ldots, n\}$, independently sample $k_m^{\text{rep}} \leftarrow \mathcal{K}_m^{\text{inj}}$ and $k^{\text{fin}} \in \mathcal{K}_n^4$ uniformly. Set $k = (k_{f(n)}^{\text{rep}}, k_{f(n)+1}^{\text{rep}}, \ldots, k_n^{\text{rep}}, k^{\text{fin}})$.*

For $k = (k_{f(n)}^{\text{rep}}, k_{f(n)+1}^{\text{rep}}, \ldots, k_n^{\text{rep}}, k^{\text{fin}})$, define the labelling functions *$r_k^{f(n)}, \ldots, r_k^n : \{0,1\}^n \to \{0,1\}^n$ recursively by*

$$r_k^m(x) = \begin{cases} x & m = n+1, \\ r_k^{m+1}(f_{k_m^{\text{rep}}}(i) \parallel j) & f(n) \leq m \leq n, \text{MSB}_m(x) = i, \text{LSB}_{n-m}(x) = j, \end{cases}$$

where $\text{MSB}_m(x)$ is the first m bits of x, $\text{LSB}_m(x)$ is the last m bits of x, and $i \parallel j$ is the concatenation of bit strings i and j. For simplicity, we define $r_k(x) = r_k^{f(n)}(x)$.

With this notation, for $k = (k_{f(n)}^{\text{rep}}, k_{f(n)+1}^{\text{rep}}, \ldots, k_n^{\text{rep}}, k^{\text{fin}})$, we next define the function $s_k : \{0,1\}^n \to \{0,1\}$ as

$$s_k(x) = h_{k^{\text{fin}}}(r_k(x))\ ,$$

The states $|\psi_k\rangle$ are then given by

$$|\psi_k\rangle = \sum_{x \in \{0,1\}^n} (-1)^{s_k(x)} |x\rangle\,.$$

Our main result is that this construction satisfies the requirements from Definition 22 as summarised by the following theorem. In particular, we show that under the subexponential-time LWE assumption, our construction achieves an entanglement gap of $\text{poly}\log n$ vs n, which is essentially optimal by Remark 19. On a 1D line, this entanglement scaling corresponds to area-law (up to poly log factors) vs volume law entanglement, which is why we call this result area-law pseudoentangled states.

▶ **Theorem 24.**

(i) *Under the standard LWE assumption, for any function $f(n) = n^\delta$ for $\delta > 0$, the state families $\Psi_n^{\text{low}} = \{|\psi_k\rangle\}_{k \in \mathcal{K}_n^{\text{low}}}$ and $\Psi_n^{\text{high}} = \{|\psi_k\rangle\}_{k \in \mathcal{K}_n^{\text{high}}}$ from form a pseudoentangled state ensemble with entanglement gap $(O(f(n)), \Omega(n))$ across geometrically local cuts in 1D.*

(ii) *Under the subexponential-time LWE assumption, there exists a function $f(n) = \text{poly}\log n$ such that the state families $\Psi_n^{\text{low}} = \{|\psi_k\rangle\}_{k \in \mathcal{K}_n^{\text{low}}}$ and $\Psi_n^{\text{high}} = \{|\psi_k\rangle\}_{k \in \mathcal{K}_n^{\text{high}}}$ from Definition 23 form a pseudoentangled state ensemble with entanglement gap $(O(f(n)), \Omega(n))$ across geometrically local cuts in 1D.*

3.4 Area-law public-key pseudoentangled states on a 2D grid

We can easily generalise the 1D construction from Section 3.3 to a system of qubits arranged on a 2D grid. This is the same construction as in [1, Appendix D.5], so we only provide a short sketch. Let $|\psi_k\rangle$ be an n-qubit state from a pseudoentangled state ensemble. We can arrange the qubits of this state on an $\sqrt{n} \times \sqrt{n}$ grid as shown in Figure 1. Now consider a contiguous 2D subregion R of this $\sqrt{n} \times \sqrt{n}$ grid. Let $|R|$ be the size of R (i.e. the number of qubits in R) and $|\partial R|$ the size of the boundary of R. Unfolding the "snake", this region R corresponds to a (not necessarily geometrically local) cut in 1D.

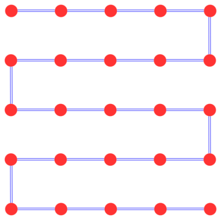

Figure 1 Arranging an n-qubit state on a $\sqrt{n} \times \sqrt{n}$ grid.

For the pseudoentangled state ensembles we constructed in Theorem 24, a high-entanglement state $|\psi_k\rangle$ has entanglement entropy $\Omega(|R|)$ for any (sufficiently large) cut R, even if the cut is not 1D geometrically local. This means that arranged on a 2D grid, the high-entanglement states from Theorem 24 exhibit volume law entanglement scaling.

Conversely, consider a low-entanglement state $|\psi_k\rangle$ from the construction in Theorem 24. From the geometry of Figure 1 it is easy to see that a region R corresponds to a 1D cut that divides the qubits into at most $O(|\partial R|)$ contiguous regions; this is because the boundary of the region R can cut the "snake" at most $O(|\partial R|)$ times. For each of these $O(|\partial R|)$ cuts in 1D, we know from Theorem 24 that $|\psi_k\rangle$ has entanglement entropy at most $O(\operatorname{poly}\log n)$ across that cut. Using subadditivity of entanglement entropy, it then follows that the entanglement entropy of the region R is at most $O(|\partial R| \cdot \operatorname{poly}\log n)$, which corresponds to area-law scaling in two dimensions (up to polylogarithmic factors).

4 Computational hardness of learning ground state entanglement structure

Our public-key pseudoentanglement constructions can be leveraged to construct Hamiltonians such that it is hard to learn the entanglement structure of their ground state. This is what we will discuss in the next sections. Specifically, we will study variants of the following problem, which we define somewhat informally.

▶ **Definition 25** (Learning Ground State Entanglement Structure (LGSES) problem). *Given a classical description of a k-local Hamiltonian H on n qubits with spectral gap at least $\frac{1}{poly(n)}$, decide whether the ground state of H has entanglement structure* A *or* B*? Here,* A *and* B *should be two qualitatively different, pre-specified entanglement structures, e.g. near are-law and volume law entanglement.*

We will see three different variants of this problem for three different types of local Hamiltonians and correspondingly three slightly different entanglement structures. We will progressively make our constructions more local – in some sense, more local corresponds to more physical Hamiltonians – but we will pay a slight price in terms of how straightforwardly the entanglement structures can be described.

(i) In Section 4.1, we will study the hardness of LGSES for $O(\log n)$-local Hamiltonians on n qubits arranged in a 1D line. The two entanglement structures to distinguish between will be $\operatorname{poly}\log n$ vs $\Omega(n)$ entanglement across geometrically local cuts in 1D. In other words, we are asked to distinguish 1D near area-law vs volume-law entanglement.

(ii) In Section 4.2, we will improve upon the locality of the Hamiltonian and study the hardness of LGSES for $O(1)$-local Hamiltonians on n qubits arranged in a 1D line. However, the entanglement structure will be slightly more complicated: we will consider the reduced states of ground states on a specific subsystem and ask whether this has 1D near area-law or volume-law entanglement structure for a mixed state measure of entanglement.

(iii) In Section 4.3, we will study the hardness of LGSES for 2-local Hamiltonians on a 2D grid of size $n \times \operatorname{poly}(n)$ and with constant local dimension, where all Hamiltonian terms are *geometrically local* (i.e. only nearest neighbors on the grid can interact). The two entanglement structures to distinguish will be entanglement entropy $O(\operatorname{poly}\log n)$ vs $\Omega(n)$ across horizontal cuts through the grid.

4.1 1D Hamiltonians with $\log n$-locality and pure states

In this section, we will show how to obtain two families of $\log n$-local Hamiltonians, one whose ground state has $\operatorname{poly}\log n$ entanglement scaling and the other whose ground state has $\Omega(n)$ entanglement scaling across geometrically local cuts in 1D, such that given the description of one of these Hamiltonians it is computationally hard to decide which family it belongs to.

We will start with the public-key pseudoentangled state constructions in Section 3.3, use the padded circuit-to-Hamiltonain construction, and then use the trace distance closeness property to show that the entanglement structures of the ground states of these Hamiltonians resemble the entanglement structure of the public-key pseudoentangled states.

▶ **Theorem 26.** *For every $n \in \mathbb{N}$, there exist two families $\mathcal{H}_n^{\mathrm{low}}$ and $\mathcal{H}_n^{\mathrm{high}}$ of $O(\log n)$-local Hamiltonians on $(n + O(\log n))$ qubits arranged on a 1D line with spectral gap $\Omega(1/\operatorname{poly}(n))$ and efficient procedures that sample (classical descriptions of) Hamiltonians from either family (denoted $H \leftarrow \mathcal{H}_n^{\mathrm{low}}$ and $H \leftarrow \mathcal{H}_n^{\mathrm{high}}$) with the following properties:*

(i) *Hamiltonians sampled according to $H \leftarrow \mathcal{H}_n^{\mathrm{low}}$ and $H \leftarrow \mathcal{H}_n^{\mathrm{high}}$ are computationally indistinguishable under the assumption that LWE is subexponentially hard.*

(ii) *With overwhelming probability, the ground states of Hamiltonians $H \leftarrow \mathcal{H}_n^{\mathrm{low}}$ have 1D near area-law entanglement and Hamiltonians $H \leftarrow \mathcal{H}_n^{\mathrm{low}}$ have 1D volume-law entanglement. Formally, this means that for geometrically local cuts in 1D of size $r = \omega(\log n)$, the ground states of the Hamiltonians have entanglement entropy $O(\operatorname{poly}\log n)$ or $\Omega(\min(r, n-r))$, respectively.*

▶ Remark 27. Under the standard LWE assumption instead of the subexponentially hardness assumption, Theorem 26 still holds, but with the smaller entanglement gap $O(n^{\delta})$ vs $\Omega(n)$ for any $\delta > 0$. This mirrors directly the statement in Theorem 24.

4.2 1D Hamiltonians with constant locality and mixed states

In this section, we will modify the construction in Section 4.1 with a unary clock to get constant locality. However, because the clock register will now have $\operatorname{poly}(n)$ qubits, we can no longer simply remove the clock qubits as we did in Theorem 26. Therefore, we will consider the entanglement structure of the reduced density matrices of the ground state with the clock register traced out. Using mixed state entanglement measures, we will show that one such density matrix will have high entanglement, and the other will have low entanglement.

As discussed in Section 2.4, there are many mixed state measures of entanglement. We will show that for any "natural" mixed state entanglement measure, the ground state of the our Hamiltonian (with the clock register traced out) has either high or low entanglement. We achieve this by giving an upper bound on the entanglement of formation of our low entanglement construction and a lower bound on the distillable entanglement of our high entanglement construction. This gives an entanglement gap for any natural entanglement measure. In fact, Hamiltonians constructed from our ensembles of pseudoentangled states achieve a maximally large gap.

▶ **Theorem 28.** *For every $n \in \mathbb{N}$, there exist two families $\mathcal{H}_n^{\text{low}}$ and $\mathcal{H}_n^{\text{high}}$ of $O(1)$-local Hamiltonians on $(n + poly(n))$ qubits arranged on a 1D line with spectral gap $\Omega(1/\operatorname{poly}(n))$ and efficient procedures that sample (classical descriptions of) Hamiltonians from either family (denoted $H \leftarrow \mathcal{H}_n^{\text{low}}$ and $H \leftarrow \mathcal{H}_n^{\text{high}}$) with the following properties:*

(i) *Hamiltonians sampled according to $H \leftarrow \mathcal{H}_n^{\text{low}}$ and $H \leftarrow \mathcal{H}_n^{\text{high}}$ are computationally indistinguishable under the assumption that LWE is subexponentially hard.*

(ii) *If we trace out poly(n) many qubits from the ground state of each Hamiltonian, the entanglement gap between the resultant quantum states in the high and low families is $\Omega(min(r, n-r))$ versus $\mathcal{O}(poly \log n)$, for a cut of size $(r, n-r)$, for any natural entanglement measure. With overwhelming probability, the reduced states on the first n qubits of the ground states of Hamiltonians $H \leftarrow \mathcal{H}_n^{\text{low}}$ have 1D near area-law entanglement and Hamiltonians $H \leftarrow \mathcal{H}_n^{\text{low}}$ have 1D volume-law entanglement with respect to any natural mixed state entanglement measure.*

▶ Remark 29. Just as in Remark 27, under the standard LWE assumption Theorem 28 still holds, but with the smaller entanglement gap $O(n^{\delta})$ vs $\Omega(n)$ for any $\delta > 0$.

4.3 2D Hamiltonians with geometric locality and pure states

In this section, we will show how to obtain two families of 2D Hamiltonians on a 2D grid of poly(n) qudits, one whose ground state has entanglement entropy of order n and the other whose ground state has entanglement entropy of order poly $\log n$, with respect to most horizontal cuts across the 2D grid. Thus, arguably, this gives us a relatively more complicated entanglement structure than the constructions in Theorem 26 and Theorem 28. However, we gain in geometric locality: the Hamiltonian only has nearest neighbor interactions on a 2D grid. Formally, 2D Hamiltonians are defined as follows.

▶ **Definition 30** (2D (local) Hamiltonian, [2])**.** *Let H be a Hermitian operator (interpreted as a Hamiltonian, giving the energy of some system). We say that H is an r-state Hamiltonian if it acts on r-state qudits (i.e. $d = r$). When $r = 2$, namely, when the qudits are qubits. We say that H is k-local if it can be written as*

$$H = \sum_i H_i,$$

where each H_i acts non-trivially on at most k qudits. Note that this term does not assume anything about the physical location of the qudits. We say that H is a 2D Hamiltonian if the qudits are arranged on a 2D grid and the terms H_i interact only pairs of nearest neighbor qudits. In particular, a 2D Hamiltonian is 2-local.

▶ **Theorem 31.** *For every $n \in \mathbb{N}$, there exist two families $\mathcal{H}_n^{\text{low}}$ and $\mathcal{H}_n^{\text{high}}$ of* geometrically 2D-local *Hamiltonians on $(n \times \operatorname{poly} n)$ qudits arranged in an $n \times \operatorname{poly}(n)$ grid with spectral gap $\Omega(1/\operatorname{poly}(n))$, and there are efficient procedures that sample (classical descriptions of) Hamiltonians from either family (denoted $H \leftarrow \mathcal{H}_n^{\text{low}}$ and $H \leftarrow \mathcal{H}_n^{\text{high}}$) with the following properties:*

(i) *Hamiltonians sampled according to $H \leftarrow \mathcal{H}_n^{\text{low}}$ and $H \leftarrow \mathcal{H}_n^{\text{high}}$ are computationally indistinguishable under the assumption that LWE is subexponentially hard.*

(ii) *With overwhelming probability, the ground states of Hamiltonians $H \leftarrow \mathcal{H}_n^{\text{low}}$ have* $\operatorname{poly}\log n$ *entanglement across horizontal cuts through the grid that are at least $\omega(\log n)$ far from the boundary, and Hamiltonians $H \leftarrow \mathcal{H}_n^{\text{high}}$ have $\Omega(n)$ entanglement across the same cuts.*

Overview. The main steps of our construction are as follows:

- First, we start with two n-qubit public key pseudoentangled states across multiple cuts, according to the construction in Section 3.3, and consider the circuits for preparing them. Suppose these circuits have $K = \operatorname{poly}(n)$ gates. Without loss of generality, we assume the circuit can be decomposed into $R = \operatorname{poly}(n)$ "rounds", each made up of exactly n nearest-neighbor interactions on qubits 1, (1,2), (2,3), etc. Any circuit can be transformed into this form by inserting a polynomial number of identity and swap gates. Hence, the circuit contains nR gates in total. As in Section 4.1, we pad the circuits with nM identity gates at the end for a sufficiently large $M = \operatorname{poly}(n)$.
- Then, we use a modified version of the 2D clock construction [3] to construct two families of 2D Hamiltonians such that the padded circuit is embedded into its ground state. That is, if $\mathsf{C} = U_{nT} \cdot U_{nT-1} \cdots U_1$ is the circuit with padding, where $T = M + R$ is the total number of rounds after padding, we construct a Hamiltonian H such that the ground state $|\psi_{\text{ground}}\rangle$, on an $n \times T$ grid of qudits, encodes the time evolution of the padded circuit.
- We show how, because of the padding, the entanglement structure of the 2D ground state across any horizontal cut resembles the entanglement structure of the state

$$|\psi_{\text{output}}\rangle = U_{nR} \cdot U_{nR-1} \cdots U_1 |0^n\rangle, \tag{4.1}$$

  across the same cut.
- By virtue of our pseudoentanglement construction, the state in Equation (4.1) either has high or low entanglement, whenever the cut has distance $\omega(\log n)$ to the boundary of the grid. Then, by a continuity argument, we show that the ground state $|\psi_{\text{ground}}\rangle$ also inherits the high or low entanglement property.

References

1 Scott Aaronson, Adam Bouland, Bill Fefferman, Soumik Ghosh, Umesh Vazirani, Chenyi Zhang, and Zixin Zhou. Quantum pseudoentanglement. *arXiv preprint v2*, 2023. `arXiv:2211.00747v2`.

2 Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe. The power of quantum systems on a line. *Communications in mathematical physics*, 287(1):41–65, 2009.

3 Dorit Aharonov, Wim Van Dam, Julia Kempe, Zeph Landau, Seth Lloyd, and Oded Regev. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM review*, 50(4):755–787, 2008.

4 Anurag Anshu, Itai Arad, and David Gosset. An area law for 2d frustration-free spin systems. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 12–18, 2022.

5 Anurag Anshu and Srinivasan Arunachalam. A survey on the complexity of learning quantum states. *arXiv preprint*, 2023. arXiv:2305.20069.

6 Itai Arad, Zeph Landau, Umesh Vazirani, and Thomas Vidick. Rigorous RG algorithms and area laws for low energy eigenstates in 1D. *Communications in Mathematical Physics*, 356:65–105, 2017.

7 Rotem Arnon-Friedman, Zvika Brakerski, and Thomas Vidick. Computational entanglement theory, 2023. arXiv:2310.02783.

8 Koenraad M R Audenaert. A sharp continuity estimate for the von Neumann entropy. *Journal of Physics A: Mathematical and Theoretical*, 40(28):8127–8136, June 2007. doi:10.1088/1751-8113/40/28/s18.

9 Johannes Bausch, Toby S Cubitt, Angelo Lucia, and David Perez-Garcia. Undecidability of the spectral gap in one dimension. *Physical Review X*, 10(3):031038, 2020.

10 Avraham Ben-Aroya, Oded Schwartz, and Amnon Ta-Shma. Quantum expanders: Motivation and constructions. In *2008 23rd Annual IEEE Conference on Computational Complexity*, pages 292–303. IEEE, 2008.

11 John Bostanci, Yuval Efron, Tony Metger, Alexander Poremba, Luowen Qian, and Henry Yuen. Unitary complexity and the uhlmann transformation problem. *arXiv preprint*, 2023. arXiv:2306.13073.

12 Adam Bouland, Bill Fefferman, and Umesh Vazirani. Computational pseudorandomness, the wormhole growth paradox, and constraints on the ads/cft duality. *arXiv preprint*, 2019. arXiv:1910.14646.

13 Z. Brakerski, P. Christiano, U. Mahadev, U. Vazirani, and T. Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. *IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 320–331, 2018. doi:10.1109/FOCS.2018.00038.

14 Zvika Brakerski and Omri Shmueli. (Pseudo) random quantum states with binary phase. In *Theory of Cryptography Conference*, pages 229–250. Springer, 2019.

15 Marcus Cramer, Martin B Plenio, Steven T Flammia, Rolando Somma, David Gross, Stephen D Bartlett, Olivier Landon-Cardinal, David Poulin, and Yi-Kai Liu. Efficient quantum state tomography. *Nature communications*, 1(1):149, 2010.

16 Toby S Cubitt, David Perez-Garcia, and Michael M Wolf. Undecidability of the spectral gap. *Nature*, 528(7581):207–211, 2015.

17 M. Fannes. A continuity property of the entropy density for spin lattice systems. *Communications in Mathematical Physics*, 31(4):291–294, December 1973. doi:10.1007/bf01646490.

18 Alexandru Gheorghiu and Matty J Hoban. Estimating the entropy of shallow circuit outputs is hard. *arXiv preprint*, 2020. arXiv:2002.12814.

19 Oded Goldreich and Salil Vadhan. Comparing entropies in statistical zero knowledge with applications to the structure of szk. In *Proceedings. Fourteenth Annual IEEE Conference on Computational Complexity (Formerly: Structure in Complexity Theory Conference)(Cat. No. 99CB36317)*, pages 54–73. IEEE, 1999.

20 Matthew B Hastings. An area law for one-dimensional quantum systems. *Journal of statistical mechanics: theory and experiment*, 2007(08):P08024, 2007.

21 Hsin-Yuan Huang, Richard Kueng, Giacomo Torlai, Victor V Albert, and John Preskill. Provably efficient machine learning for quantum many-body problems. *Science*, 377(6613):eabk3333, 2022.

22 Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In *Advances in Cryptology–CRYPTO 2018: 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19–23, 2018, Proceedings, Part III 38*, pages 126–152. Springer, 2018.

23 A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation.* American Mathematical Society, USA, 2002.

24 Zeph Landau, Umesh Vazirani, and Thomas Vidick. A polynomial time algorithm for the ground state of one-dimensional gapped local Hamiltonians. *Nature Physics*, 11(7):566–569, 2015.

25 Urmila Mahadev, Umesh Vazirani, and Thomas Vidick. Efficient certifiable randomness from a single quantum device, 2022. arXiv:2204.11353.

26 Chinmay Nirkhe, Umesh Vazirani, and Henry Yuen. Approximate low-weight check codes and circuit lower bounds for noisy ground states. *arXiv preprint*, 2018. arXiv:1802.07419.

27 Shinsei Ryu and Tadashi Takayanagi. Holographic derivation of entanglement entropy from the anti–de sitter space/conformal field theory correspondence. *Physical review letters*, 96(18):181602, 2006.

28 Amit Sahai and Salil Vadhan. A complete problem for statistical zero knowledge. *Journal of the ACM (JACM)*, 50(2):196–249, 2003.

29 Norbert Schuch, Ignacio Cirac, and Frank Verstraete. Computational difficulty of finding matrix product ground states. *Physical review letters*, 100(25):250501, 2008.

30 Salil Vadhan. Pseudorandomness. *Foundations and Trends® in Theoretical Computer Science*, 7(1–3):1–336, 2012.

31 Andreas Winter. Tight uniform continuity bounds for quantum entropies: Conditional entropy, relative entropy distance and energy constraints. *Communications in Mathematical Physics*, 347(1):291–313, March 2016. doi:10.1007/s00220-016-2609-8.

32 Mark Zhandry. Quantum minimalism, 2023. Talk at Simons Institute, https://www.youtube.com/live/7cqnrASfjco?si=1XlLZpqfaEsBEVp8.