Frontiers of Quantum Complexity *Instructor: Anand Natarajan, Anurag Anshu*

Lecture 18: Renormalization Group Flow

November 4, 2025 *Scribe: Vetri Vel*

1 Background

Renormalization group (RG) flow is an important idea in quantum field theory which studies how the parameters of a theory change as the energy or momentum scale is varied. However, we will focus on RG flow in the context of algorithms for finding low energy states of local Hamiltonians. RG flow algorithms have a recursive nature and often leverage physical insights. These two different notions of RG flow share the idea of coarse-graining a system and integrating out or discarding degrees of freedom.

Wilson famously used numerical RG flow to solve the Kondo problem [1] (see also [2]). We will cover an example of the related real-space RG flow. We will then cover the more widely used density matrix RG flow (DMRG) [3]. Finally, we will cover the rigorous RG flow algorithm for 1-dimensional ground states [4, 5]. This algorithm is important to Hamiltonian complexity because it rigorously proves that the ground state of a 1-dimensional gapped Hamiltonian can be computed in polynomial time. This stands in contrast to the general local Hamiltonian problem (LHP) which is QMA-hard. An RG flow-like construction was also used to prove a 2D area law for frustration free, locally gapped systems [6].

2 Real-space RG flow

We will give a simple example of a real-space RG flow algorithm. It can also be thought of as an ansatz known as the multi-scale entanglement renormalization ansatz (MERA) [7]. Consider a 1-dimensional local Hamiltonian

$$H = \sum_{i=1}^{n-1} h_{i,i+1} \tag{1}$$

where the goal is to find the ground state. Figure 1 shows the algorithm, or equivalently the structure of the ansatz. At a high level, the algorithm works by repeatedly diagonalizing interactions between neighboring qubits and keeping only the two low energy states, effectively defining a new coarse-grained qubit.

This algorithm works well on some systems, but fails on some simple systems. As Anand pointed out during class, it may fail on some classical 1-dimensional constraint satisfaction problems (CSPs) because it does not use a dynamic programming approach.

Anand's Challenge: Find a CSP Hamiltonian where this algorithm fails.

An illustrative example where this algorithm works well is the transverse-field Ising model (TFIM),

$$H_{\text{TFIM}} = -J \sum_{i=1}^{n-1} Z_i \otimes Z_{i+1} - h \sum_{i=1}^{n} X_i \tag{2}$$

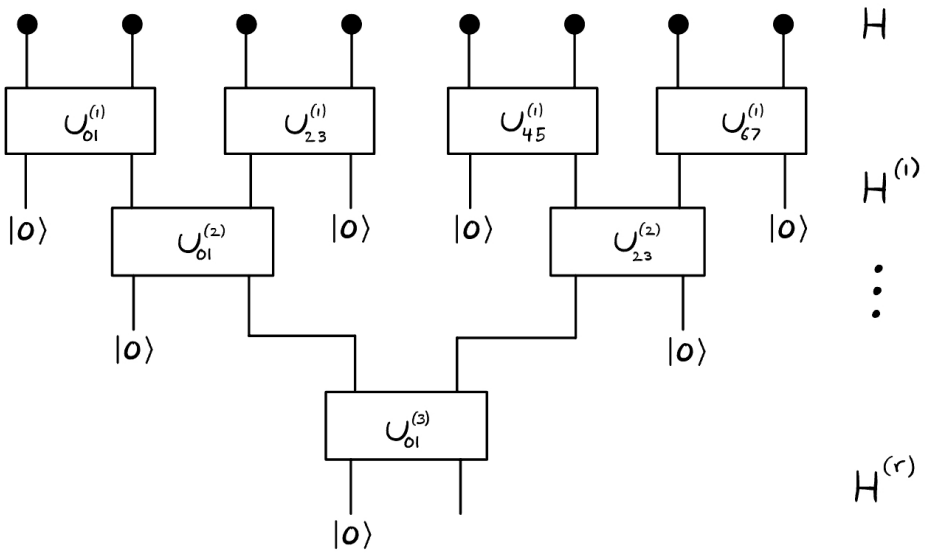

Figure 1: Schematic of real-space RG flow

which corresponds to an Ising model with nearest-neighbor interactions and magnetic field along the $\hat{x}$ direction. This Hamiltonian is exactly solvable and its behavior is described by the single parameter $g = J/h$. For $g < 1$ and $g > 1$, the Hamiltonian is gapped, but the gap vanishes at $g = 1$. Note that when $g \to 0$, the transverse field dominates and the ground state is $|+\rangle^{\otimes n}$ and when $g \to \infty$, the spin couplings dominate and the ground states are classical strings. There is a phase transition at $g = 1$ where the competing terms are equally balanced.

To see how the algorithm works on this model, consider qubits labeled 1 through 6. We can think of the Hamiltonian on qubits 2 and 3 as

$$h_{23} = JZ_2 \otimes Z_3 - h(X_2/2 + X_3/2) \tag{3}$$

which has eigenstates $|\psi_0\rangle$, $|\psi_1\rangle$, $|\psi_2\rangle$, $|\psi_3\rangle$ in ascending energy. We define U to be a unitary such that $U|\psi_0\rangle = |0\rangle_2 |0\rangle_3$ and $U|\psi_1\rangle = |0\rangle_2 |1\rangle_3$ and define the projector

$$\Pi_{23} = U|0\rangle\langle 0|_2 U^\dagger \tag{4}$$

which projects onto the two lowest eigenstates of h_{23}. In a sense, we are combining qubits 2 and 3 into a single qubit whose states are $|\psi_0\rangle$ and $|\psi_1\rangle$. We then define the coarse-grained Hamiltonian

$$H^{(1)} = \Pi_{23} h_{12} \Pi_{23} + (\Pi_{23}\Pi_{45}) h_{34} (\Pi_{23}\Pi_{45}) + \Pi_{45} h_{56} \Pi_{45}$$

where the Hamiltonian has been projected onto the low-energy spaces using Π_{23} and Π_{45}. The three terms correspond to the three green lines in Figure 2. For example, $(\Pi_{23}\Pi_{45})h_{34}(\Pi_{23}\Pi_{45})$ represents the interaction between the new coarse-grained qubits. Amazingly, it turns out that $H^{(1)}$ maintains the TFIM form,

$$H^{(1)} = -J^{(1)} \sum_{i=1}^{n/2-1} Z_i \otimes Z_{i+1} - h^{(1)} \sum_{i=1}^{n/2} X_i \tag{5}$$

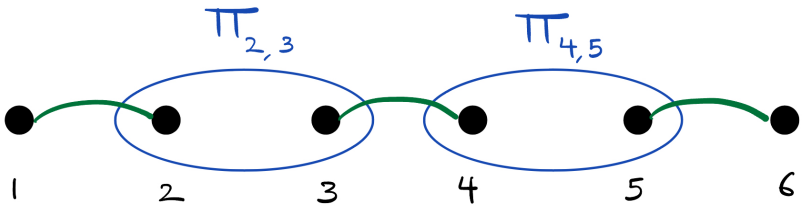

Figure 2: Illustration of real-space coarse-graining

where $g^{(1)} = J^{(1)}/h^{(1)}$ is a function of the original $g = J/h$, represented by $g^{(1)} = f(g)$. We see that repeatedly coarse-graining causes g to change according to

$$g \to f(g) \to f(f(g)) \to \dots$$

and it turns out that f is such that $f(1) = 1$ and for $g \neq 1$, it rapidly pushes g away from 1. This means that for $g \neq 1$, it takes only $r \sim \log(1/\epsilon)$ coarse-graining steps to get $g^{(r)}$ such that a product state is ϵ-close to the ground state, as mentioned above. This means that in the gapped case, approximate ground states can be prepared with constant depth, starting from a product state. However, the $g = 1$ gapless case requires depth $r = \log n$ in this construction. It is an open question whether we can generalize this finding.

Open question: Can all 1-dimensional gapped ground states be prepared in $O(1)$ depth? More precisely, can they be prepared in $f(\epsilon)$ depth for $1 - \epsilon$ fidelity? If not, what about logarithmic depth?

Note that this algorithm greedily keeps low energy subspaces without thinking about the overall ground state. This could lead it to accidentally project out the global ground state and it is straightforward to think of toy Hamiltonians where this occurs.

3 Density Matrix Renormalization Group (DMRG)

DMRG avoids some of the issues with the simple real-space RG flow algorithm. In this approach, we focus on the ground state instead of the Hamiltonian. The key assumption is that 1-dimensional gapped ground states have area law entanglement, which was only proved by Hastings many years after the advent of DMRG.

3.1 Infinite DMRG

This approach assumes the Hamiltonian is translationally invariant. Suppose that we want to find the ground state on n particles. We first define m (typically poly n) to be the maximum retained dimension of subspaces. We will assume the particles are qubits but this can be relaxed to qudits. The approach of infinite DMRG is as follows,

1. Start with the Hilbert space S_1 for one particle
2. For $k = 1, 2, 3 \dots$

- Take two sets of k particles (each restricted to m dimensional subspaces) and add two particles between them, as in figure 3
- Project the Hamiltonian according to S_k and solve for the ground state $|\Omega\rangle$ on the $2k+2$ particles in $4m^2$ dimensions (the dimension of the tensor product of each of the spaces)
- Find the Schmidt decomposition of $|\Omega\rangle$ along the center cut shown in figure 3, namely $|\Omega\rangle = \sum_{i=1}^{2m} \lambda_i |L_i\rangle \otimes |R_i\rangle$ where λ_i are in descending order
- Define S_{k+1} to be the subspace spanned by the first m eigenvectors $|L_i\rangle$

The process typically stops when the block bases S_k converge to within some tolerance ϵ.

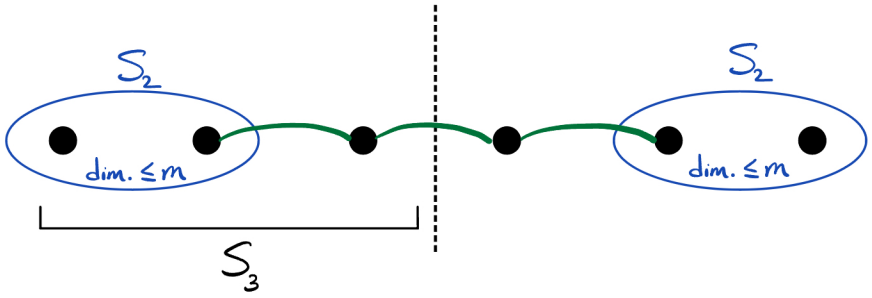

Figure 3: Going from a subspace on 2 particles to a subspace on 3 in infinite DMRG.

3.2 Finite DMRG

Finite DMRG starts with a constant matrix product state (MPS) ansatz of constant bond dimension, typically random or an educated guess. It then sweeps back and forth across individual tensors in the MPS, optimizing each one to minimize the energy of the state. Usually only a few sweeps are needed to converge to a low-energy state.

4 Rigorous RG flow

In 2015, Landau, Vidick, and Vaziriani gave the first rigorous polytime algorithm for finding gapped 1-dimensional ground states using the crucial notion of a viable set [4]. In 2017, they gave a polytime RG flow algorithm that has stronger properties than the previous algorithm [5]. It works for 1-dimensional local Hamiltonians and does not require translational invariance.

Viable sets: A subspace (Π_S) of dimension $\operatorname{poly} n$ for each continuous subset $S \subseteq [n]$ such that it can be represented efficiently using a basis of MPS states and $(\Pi_S \otimes I_{\bar{S}}) |\Omega\rangle \approx |\Omega\rangle$.

We expect viable sets to exist since we know that 1-dimensional gapped ground states can be represented by MPSs of bond dimension $\operatorname{poly} n$. Given this MPS and any continuous subset S of the qubits, an MPS basis for the viable set is given by fixing the bonds to the left and right of the MPS fragment in all $\operatorname{poly} n$ possible ways.

Now we cover how viable sets are merged. Suppose we have two adjacent viable sets for S_1 and S_2. We could try getting a viable set for $S = S_1 \cup S_2$ by simply taking the subspace given

by $\Pi_{S_1} \otimes \Pi_{S_2}$, which has an MPS basis. However, it would have dimension D^2, where D is the dimension of each subspace. This motivates a different approach. We take a constant number of vectors $|v_1\rangle, \ldots, |v_k\rangle$ from Π_{S_1} and $|w_1\rangle, \ldots, |w_k\rangle$ from Π_{S_2}. Note that if Π_{S_1} is viable with 0.99 accuracy, each $|v_i\rangle$ is viable with expected accuracy $0.99/D$. Therefore, we will lose accuracy unless we apply an AGSP to bring the few states closer to the ground state. To do this, we first find the Schmidt decomposition of the AGSP,

$$K = \sum_{i=1}^{D'} (K_S)_i \otimes (K_{\overline{S}})_i \tag{6}$$

and we define Π_S as projecting onto the following subspace,

$$\text{span}\{(K_S)_i |v_1\rangle |w_1\rangle, \ldots, (K_S)_i |v_k\rangle |w_k\rangle\}_{i=1}^{D'} \tag{7}$$

which has dimension $D = D'k$. We use a (D', Δ)-AGSP with $D'\Delta < 1/2$ so that it effectively projects states with sufficiently high overlap with $|\Omega\rangle$ onto the low-energy space while controlling the dimension of Π_S. This guarantees that our merging procedure on two viable sets produces another highly accurate viable set. We keep repeating this merging of viable sets until we get the whole space.

References

[1] Kenneth G. Wilson. The renormalization group: Critical phenomena and the kondo problem. *Rev. Mod. Phys.*, 47:773–840, 1975.

[2] Sergey Bravyi and David Gosset. Complexity of quantum impurity problems. *Comm. Math. Phys.*, 356:451–500, 2017.

[3] Steven R. White. Density matrix renormalization group algorithms with a single center site. *Phys. Rev. B*, 72:180403, 2005.

[4] Zeph Landau, Umesh Vazirani, and Thomas Vidick. A polynomial time algorithm for the ground state of one-dimensional gapped local hamiltonians. *Nature Physics*, 11:561–564, 2015.

[5] Itai Arad, Zeph Landau, Umesh Vazirani, and Thomas Vidick. Rigorous rg algorithms and area laws for low energy eigenstates in 1d. *Comm. Math. Phys.*, 356:65–105, 2017.

[6] 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 (STOC 2022)*. ACM, 2022.

[7] Guifre Vidal. A class of quantum many-body states that can be efficiently simulated. *Phys. Rev. Lett.*, 101:110501, 2008.