SYMMETRY PROTECTED SUBSPACES IN QUANTUM SIMULATIONS

by

Caleb G. Rotello

A thesis submitted to the Faculty and the Board of Trustees of the Colorado School of Mines in partial fulfillment of the requirements for the degree of Master of Science (Quantum Engineering).

Golden, Colorado

Date _____

Signed: _____

Caleb G. Rotello

Signed: _____

Dr. Eliot Kapit
Thesis Advisor

Golden, Colorado

Date _____

Signed: _____

Dr. Eliot Kapit
Professor and Director
Department of Quantum Engineering

ABSTRACT

This work demonstrates an efficient algorithm which leverages transitive closure on graphs to identify symmetry protected subspaces of quantum unitary operators in the σ^z basis. The algorithm's time complexity is linear to the size of the subspace. This subspace constrains states in quantum time evolution to a smaller subspace. If this subspace is small enough, a classical computer can complete the quantum simulation task; however, if this subspace has exponential scaling and is thus infeasible to calculate, a quantum computer is required to complete the same simulation task. In this regime, the symmetry protected subspace is leveraged to reduce quantum computation errors with post-selection. To do this, a probabilistic algorithm is presented which can determine if a measured state is within the subspace in polynomial time. Both algorithms are benchmarked on three quantum simulations: the Heisenberg XXX Hamiltonian, T2 Quantum Cellular Automata (QCA), and F4 QCA. QCA simulations are implemented as update rules which define quantum time evolution. Finally, the paper concludes by implementing post-selection on a quantum circuit emulator with noise with the three aforementioned simulations.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF SYMBOLS

Hamiltonian . . . . . H

Number of spins, total system . . . . . N

Number of spins, local operator L

Operator . . . . . $\hat{}$

Hadamard . . . . . $\hat{H}$

Pauli operator . . . . . $\sigma^{x/y/z}$

Projector . . . . . $\hat{P}$

Complete list of subspaces . . . . . G

Complete list of local subspaces on operator at spin i . . . . . G_i^L

Symmetry protected subspace . . . . . F

αth symmetry protected subspace on operator at spin i from local transition β . . . . . $F_{i,\beta,\alpha}^L$

jth symmetry protected subspace in a list . . . . . F_j

Set of states from $|a\rangle$ with non-zero probability after i time steps . . . . . $T_{|a\rangle}^i$

Graph, edges are explored by unitary time evolution . . . . . P

Path in P between a, b . . . . . $R(a, b)$

KL-Divergence . . . . . D_{KL}

Quantum evolution statistical divergence function . . . . . $\mathcal{D}$

Discrete time steps . . . . . t

LIST OF ABBREVIATIONS

Noisy intermediate scale quantum . . . . . NISQ

Quantum Cellular Automata . . . . . QCA

Uniform random noise . . . . . urn

Uniform random subspace . . . . . urs

ACKNOWLEDGMENTS

I would like to thank Dr. Eliot Kapit for his excellent advisement over the last year and giving me the opportunity to work on this project; his vision and trust has provided me with the most intellectually satisfying portion of my academic career. I would also like to thank him for his perseverance and dedication to his students and the Quantum Engineering program through much personal hardship; this dedication has proved essential as I wandered through the quantum dark while writing this thesis.

Thank you to the committee members Dr. Zhexuan Gong and Dr. Mohammed Hadi, along with Dr. Dinesh Mehta for serving as committee chair. A special thanks to the folks at the National Renewable Energy Laboratory; to Drs. Eric Jones and Peter Graf for hours listening to my (often nonsensical) presentations and providing more grounded goals and pathways. On a related note, thank you to the Kapit Research Group and everyone in our weekly journal club for useful conversations.

I would also like to thank Dr. Seth Tucker, Dr. Toni Lefton, the *High Grade* editing staff, and the Creative Writing Club for being the outlet to the other side of my brain. Thank you all being a community for the often overlooked but completely essential part of life.

Thank you to my friends within the quantum engineering community at Mines: Hakan Ayaz, Hunter Solomon, Matt Tennery, George Gratan, and Brandon Barton. Our bouldering sessions, board games, and sushi nights have made this program feel less like a degree and more like a home. I hope I can support each of you as you have supported me. Also thank you Jake Vossen, Ethan Holen, Jorgen Heilbron, and Kieran Lewis for being my friends who get me outside and moving, show me odd tv shows, and spend a detrimental amount of time discussing the collective unconscious and symbolic meta-narrative. Thank you Jeff and Nathan at Higher Grounds, for keeping the Mines graduate students going with good coffee.

Thank you to my habibis Tyler Rotello, Anthony Wanjiru, Micah Thomas, and Hakan Ayaz for being brothers to me. This work is built on our shared chai, late night video game sessions, and Turkish coffee.

Thank you to my family unit Tyler, Sophia, Tammy, and Steve Rotello. You are each my biggest supporter in your own way, and are all the apple of my eye, my grounding to the past, and path to the future.

For Grandpa Ray Palm and Papa Rocco Rotello, the head and the heart

CHAPTER 1

INTRODUCTION TO SYMMETRY FOR QUANTUM SYSTEMS

Symmetry is a concept fundamental to physics. Many theories like the standard model, time crystals [1], condensed matter physics [2], quantum field theory, and topological error correction [3] are built around conserving or breaking them. Additionally, symmetry is often useful to discover and leverage in quantum systems, and subfields have emerged which study quantum symmetry and leverage it to study novel physics [4] [5] [6] [7] [8].

This work addresses two ways in which symmetry has improved modern quantum simulations; this is by no means a comprehensive list. The algorithms presented later will seek to address both of these symmetry-provided advantages.

First, symmetry can make quantum simulation with a classical algorithm more feasible through the construction of a smaller computational basis. Classical algorithms to simulate quantum mechanics scale exponentially with the number of spins in the system; each additional spin adds another 2 dimensional degree of freedom, which doubles the computational overhead to store the wavefunction and perform operations on it. If it is known that the wavefunction only ever occupies a smaller set of states, which is the case with symmetry, the computational space can be restricted to this smaller set.

Second, symmetry also improves quantum simulation on a quantum computer with post-selection. This is a procedure where if a quantum system has a known conserved quantity that can be evaluated from a single measurement, measurements which violate this conserved quantity can be removed from the final calculations. In noisy intermediate scale quantum (NISQ) [9] devices, errors exponentially proliferate and destroy coherent quantum information. Post-selection offers a way of making this error less destructive by extracting a probability distribution closer to the ideal one, but is restricted because (1) it cannot detect errors which do not break the symmetry and (2) if a calculation requires a given number of measurements, and post-selection removes an exponential number of measurements because they break the symmetry, then exponentially more measurements are required to reach the same fidelity.

The advantages of symmetry also come with a large cost. Historically, symmetry must be intuitively identified or guessed for general unitary operators. If the entirety of the symmetry is lower order and obvious, such as total spin conservation, this is a trivial task. In general this is not the case, so this work will present an algorithm which identifies the symmetry protected subspaces observable in a given basis. It will also present an algorithm which can ascertain with small uncertainty if two states are within the same subspace.

The subspace finding algorithm begins with a starting state and efficiently finds the set of states with non-zero amplitude under one step in unitary time evolution. Then, repeat the process on each found state to find all states with non-zero amplitude under a second step in unitary time evolution. Repeat until there are no new states discovered; every state seen by this process has a non-zero probability of appearing as a basis state in the quantum wavefunction. The shared subspace algorithm works by taking two nodes, the start and measured states, in a graph of all basis states in a subspace and finding an efficient path between them to assert that they exist in the same subspace. If there is no path, they do not share the same symmetry protected subspace.

To tell the story, this work applies the algorithm to the simulated quantum dynamics of three quantum systems. The Heisenberg XXX model $\sum_i \sigma_i^x \sigma_{i+1}^x + \sigma_i^y \sigma_{i+1}^y + \sigma_i^z \sigma_{i+1}^z$, which has total spin conservation given by the symmetry operator $\sum_i \sigma_i^z$; the T2 Quantum Cellular Automata (QCA) [10] [11] $\sum_{i=1}^{N-2} \hat{H}_i \sum_{\alpha,\beta\in\{0,1\},\alpha+\beta=1} \hat{P}_{i-1}^{\alpha} \hat{P}_{i+1}^{\beta}$, which has the more complicated symmetry operator $\sum_i \sigma_i^z \sigma_{i+1}^z$; the F4 QCA [11] $\sum_{i=2}^{N-3} \hat{H}_i \sum_{\alpha,\beta,\gamma,\delta\in\{0,1\},\alpha+\beta+\delta+\gamma=2} \hat{P}_{i-2}^{\alpha} \hat{P}_{i-1}^{\beta} \hat{P}_{i+1}^{\gamma} \hat{P}_{i+2}^{\delta}$, for which there was no known symmetry operator.

The symmetry protected subspace finding algorithm in Chapter 2 shows that by restricting to the σ^z basis, an automated routine is developed to discover symmetries which are observable in that basis, and then examines the scaling of these symmetries to define when the algorithm is classically unmanageable. Post-selection can still be used when the restricted basis is computationally unfeasible, so Chapter 3 proposes a probabilistic post-selection with the greedy shared subspace algorithm, and provides numerical analysis of the algorithm's failure conditions. Finally, Chapter 4 applies this algorithm to the three relevant quantum models using a classically simulated quantum circuit with noise.

CHAPTER 2
AUTOMATED DETECTION OF SYMMETRY PROTECTED SUBSPACES

Quantum symmetries are observations of a conserved quantity, such that a symmetry operator $\hat{S}$ of a unitary $\hat{A}$ obeys the equation

$$\left[\hat{S}, \hat{A}\right] |\psi\rangle = 0 \tag{2.1}$$

where $|\psi\rangle$ is an eigenvector of $\hat{S}$. Equation 2.1 shows that if $|\psi'\rangle = \hat{A}\,|\psi\rangle$, where $|\psi'\rangle$ is also an eigenvector of $\hat{S}$, then $|\psi\rangle$ and $|\psi'\rangle$ will share an eigenvalue, which is the conserved quantity. No general form to calculate this $\hat{S}$ exists and the search space of all possible $\hat{S}$ operators is impossibly large, so it is usually discovered with clever intuition or guessing [12].

In lieu of the statement above, an alternative interpretation of symmetry must be used. For post-selection and reduced basis construction, it is useful to have all eigenvectors of $\hat{S}$ which share an eigenvalue, without knowing the underlying conserved quantity. These degenerate eigenvectors of $\hat{S}$ can be described by the defining the notion of a symmetry protected subspace.

Given a subspace F spanned by a subset of the basis vectors of an orthonormal basis U, define the projector onto the basis F

$$\hat{P}_F = \sum_{j \in F} |j\rangle \langle j| \tag{2.2}$$

F is a symmetry protected subspace for $\hat{A}$ iff

$$[\hat{P}_F, \hat{A}] = 0 \tag{2.3}$$

Which means that for $|\psi'\rangle = \hat{A}\,|\psi\rangle$, if $|\psi\rangle$ can be decomposed into a linear combination of basis vector in F, $|\psi'\rangle$ can as well. Every basis vector $|u\rangle$ of U will have a corresponding F, even if it is the only vector in that symmetry protected subspace; the total set of symmetry protected subspaces F spans U. This chapter describes how to find both the total set of F and individual F for a given time evolution unitary. Sections 2.1-2.3 detail how to find these subspaces by leveraging transitive closure and kronecker products if U is taken to be the σ^z basis, while Section 2.4 examines the post-selection and reduced simulation basis applications of the algorithm, and how they can perform in the use cases of the three considered models; Heisenberg XXX, T2 QCA, F4 QCA.

2.1 A Graph Approach to Unitary Operators

Choose an orthonormal basis U with basis vectors $\{|u\rangle\}$. If a unitary operator is applied once to a basis vector $|u\rangle$, giving $|\psi\rangle = \hat{A}\,|u\rangle$, then $|\psi\rangle$ has a basis vector decomposition

$$|\psi\rangle = \sum_{u' \in U} \alpha_{u'} \left|u'\right\rangle \tag{2.4}$$

With this single time step and decomposition for each $|u\rangle \in U$, an unweighted directed graph D of the Hamiltonian can be described. The nodes in the graph are given by basis states $\{|u\rangle\}$, and an edge $|u\rangle \to |u'\rangle$ exists if $\alpha_{u'} \neq 0$. For the current work, this graph does not need to formally implement any quantum dynamics, *just if a state-to-state transition exists*, and can thus have unweighted edges. An example of such a graph for a fractional iSWAP unitary operator

$$\text{iSWAP}(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & i\sin(\theta) & 0 \\ 0 & i\sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{2.5}$$

and implemented as iSWAP=$\sum_{i=0}^{N-2} \text{iSWAP}_{i,i+1}$ can be seen in Figure 2.1.a. D_{iSWAP} has directed edges between each basis states which move to each other under one step in unitary time evolution.

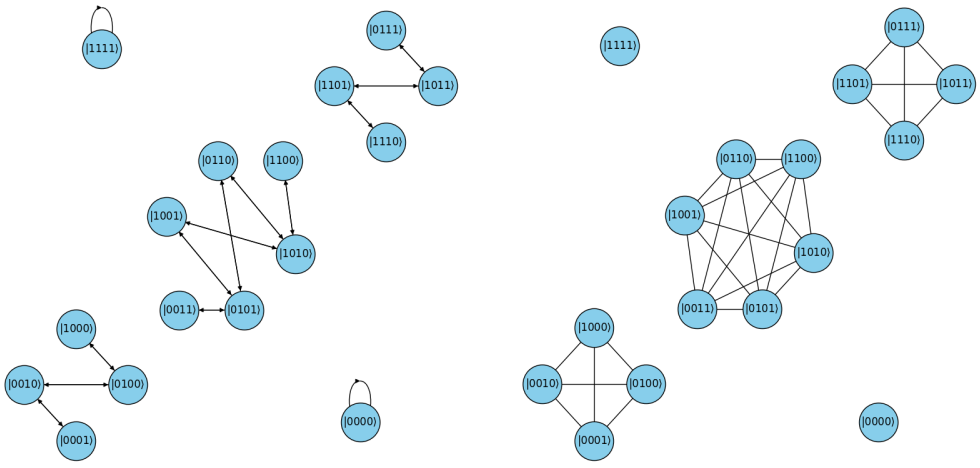

(a) (b)

Figure 2.1 (a) The unitary operator graph D_{iSWAP} of a fractional iSWAP network of 4 spins. Nodes are basis states of the σ^z basis, and edges represent a non-zero probability of transition in a single step of unitary time evolution. (b) The cluster graph G_{iSWAP} of the iSWAP network in (a) Each edge indicates that states can see each other at some point in unitary time evolution.

By adopting this graph structure, the matrix-vector multiplication operations can be left behind in favor of graph and set operations.

2.1.1 Defining Subspaces with Transitive Closure

Edges in D have the transitive property, which arises from a transitive property in unitary time evolution: if $\langle b| \hat{A} |a\rangle \neq 0$, and $\langle c| \hat{A} |b\rangle \neq 0$, then $\langle c| \hat{A}\hat{A} |a\rangle \neq 0$. This transitive property exists because the amplitudes of each edge were dropped, which means destructive interference on any basis state is no longer possible. This transitive property is established for all states in basis U for the unitary operator. The transitive closure of D is taken, which creates a new graph G, where an edge $|u\rangle \rightarrow |u'\rangle$ exists in G iff there is some path in D which connects the two nodes. In the language of unitary time evolution, an edge in G indicates

$$|u\rangle \rightarrow \left|u'\right\rangle \in G \Rightarrow \exists\, t \; s.t. \left\langle u'\right| [\hat{A}]^t \, |u\rangle \neq 0 \tag{2.6}$$

The transitive closure of the iSWAP graph in Figure 2.1.a can be seen in Figure 2.1.b on Page 4. Notice how each between state of a complete subgraph of G_{iSWAP}, total spin is conserved.

Unitary time evolution follows time-reversal symmetry, so each edge in D is bidirectional and D can always be practically considered undirected. The transitive closure of an undirected graph is a cluster graph, so G is a cluster graph. Cluster graphs are the disjoint union of complete graphs. Each node in a given complete graph can only transition to other nodes in the same complete graph, and each node is a basis vector, so this cluster graph which was built by $\hat{A}$ construction satisfies Equation 2.3 if each complete graph is a symmetry protected subspace F.

All-to-all connectivity is always implied in the transitive closure of an undirected graph; therefore, each complete sub-graph of G can be considered a set of basis vectors, so each complete sub-graph forms a symmetry protected subspace F, and G is the total set of subspaces.

$$G = \{F_1, F_2, \ldots F_M\} \tag{2.7}$$

These symmetry protected subspaces are $O(2^N)$ to articulate for a unitary of N spins in Hilbert space $\mathcal{H}_{2^N}$ with a breadth first search to find the transitive closure. While the exponential scaling of Hilbert space is always a limit for computation, this demonstrates that segregating the states which given unitary $\hat{A}$ acts on is linear to the number of states for a given graph D of a unitary operator.

Properly implemented, only one node from each F will start the breadth first search. Efficient breadth first search is linear in the number of nodes plus number of edges. The number of basis vector to basis vector transitions grows linearly with the number of basis vectors, because unitarys for quantum simulation are generally very sparse, so the number of edges grows linearly with the number of nodes. Therefore, if an efficient record of searched nodes is kept with hashing, G is $O(2^N)$ to compute for N spins.

2.2 Union of Subspaces of Local Operators to Create a Global Subspace

The method outlined in Section 2.1 gives an algorithm which can find every symmetry protected subspace for the graph $D_{\hat{A}}$ of some unitary operator $\hat{A}$ in time linear to the size of the Hilbert subspace the operator lives in. However, constructing the graph $D_{\hat{A}}$ is classically intractable if $\hat{A}$ is an interaction of many spins, because edges are discovered with matrix-vector multiplication. So an alternative approach must be taken to leverage transitive closure on basis vectors and form the total set of symmetry protected subspaces.

An N-spin unitary operator can be expressed in the form below, where each $\hat{A}_i^L$ is an operator on a few spins L starting with spin i,

$$\hat{A}^N = \sum_{i=0}^{N-1} c_i \hat{A}_i^L \tag{2.8}$$

where c_i is a complex coefficient and not all $\hat{A}_i$ are equivalent. Formally, each L is then dependent on i. Notationally this is too cluttered and the examined models are all spatially homogenous, so L is left uniform. Each coefficient c_i can be dropped, as the algorithm is only interested in possible states, not their actual probability. If $\hat{A}^N$ can be expressed such that each $\hat{A}_i^L$ is efficiently computable, the total set of symmetry protected subspaces G_i^L of each $\hat{A}_i^L$ can be computed through the same transitive closure process detailed in Section 2.1.1, yielding a list of cluster graphs which are not disjoint from each other $G' = \{G_0^L, \dots G_{N-1}^L\}$. By leveraging the transitive property of unitary operators again, the transitive closure of G' will yield the total set of symmetry protected subspaces G.

This process is notationally dense and verbose, so to keep the algorithm more tangible an example of this process for a fractional iSWAP operator is worked along with the explanation, by taking it from a single operator on 2 spins to a chain of 3 operators on 4 spins. This example

starts with the smallest computable G^2 for a single iSWAP operator:

$$\begin{aligned} F_\alpha =& \{|00\rangle\} \\ F_\beta =& \{|10\rangle, |01\rangle\} \\ F_\gamma =& \{|11\rangle\} \end{aligned}$$

Each $G_i^L \in G'$ has a corresponding set of local symmetry protected subspaces $G_i^L = \{F_{i,0}^L, F_{i,1}^L \dots F_{i,M}^L\}$. This is formed through a kronecker product scaling function, defined in Section 2.2.1, on the group G^L. In the iSWAP chain, G_0^L gives every possible interchange of basis states under action of $\text{iSWAP}_{0,1}$:

$$\begin{aligned} G_0^L = \{ & \\ & F_{0,\beta,0}^L = \{|1000\rangle, |0100\rangle\}, \\ & F_{0,\beta,1}^L = \{|1001\rangle, |0101\rangle\}, \\ & F_{0,\beta,2}^L = \{|1010\rangle, |0110\rangle\}, \\ & F_{0,\beta,3}^L = \{|1011\rangle, |0111\rangle\} \\ & \} \end{aligned}$$

Each $F_{i,j}^L$ is a complete graph and the jth local symmetry protected subspace of the local operator acting on spin i, and all G_i^L are the local total sets of subspaces; a subspace in this G_i^L is every interaction between basis states under evolution of only $\hat{A}_i^L$. Notationally, j is a tuple with the elements α, β for the $F\alpha \in G^L$ and the βth local subspace corresponding to it. If there are N local total sets of subspaces which each have 2^N vectors, some local symmetry protected subspace $F_{i,j}^L \in G_i^L$ will share each vector with another local subspace with every other local total set. To take a local symmetry $F_{i,j}^L$ and make it global, recursively merge it with all other local symmetries it shares a vector with, as demonstrated with the iSWAP example in Figure 2.2 on Page 8. Each local subspace only allows the interchange of a few basis states, but these basis states are shared with other local subspaces of other local operators. By following a path through local subspaces which share basis states, a full subspace is created through transitive closure.

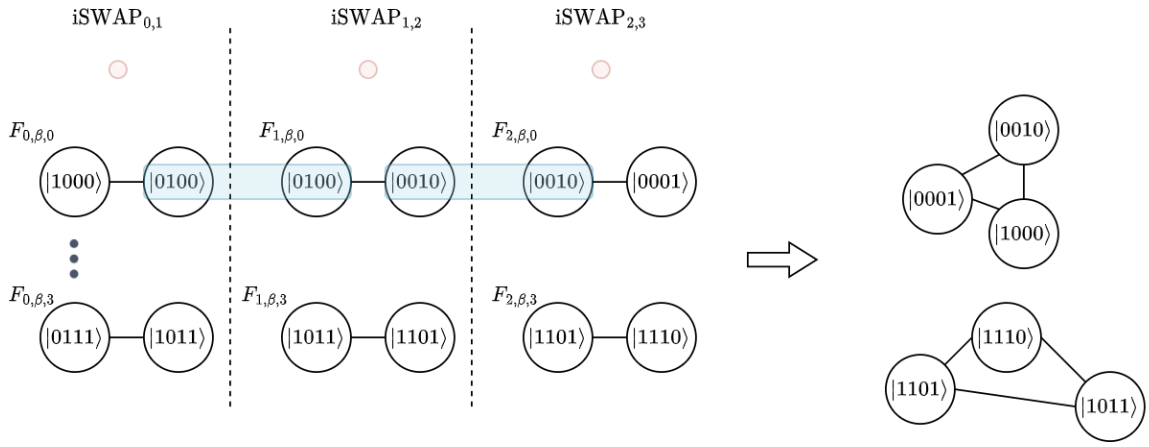

Figure 2.2 Each local subspace for each local operator in the iSWAP chain. Non-disjoint subspaces are merged to create global ones.

As stated above and shown in Figure 2.2 with the iSWAP example, if each G_i^L is a collection of local symmetry protected subspaces $\{F_{i,j}^L\}$ the global symmetry protected subspaces will be given by

$$F_k = \cup\{F_{i,j}^L \cup F_{i',j'}^L \mid \left|F_{i,j}^L \cap F_{i',j'}^L\right| \neq 0\} \tag{2.9}$$

which is the union of all $F_{i,j}^L$ which are not disjoint. If two sets $F_{i,j}^L$ and $F_{i',j'}^L$ are not disjoint, then there exists a path from each $|u\rangle \in F_{i,j}^L$ to each $|u'\rangle \in F_{i',j'}^L$, and its transitive closure is $F_{i,j}^L \cup F_{i',j'}^L$. Therefore, if a unitary operator can be decomposed into local operators $\hat{A}_i^L$, such that each $\hat{A}_i^L$ has a computable symmetry subspace structure, the total set of symmetry protected subspaces G is computable by following the transitive closure of symmetry protected subspaces of $\hat{A}_i^L$.

$$G = \{F_k\} \tag{2.10}$$

The finer details of this process are given in Sections 2.2.1 - 2.2.2.

2.2.1 The Kronecker Product Scaling Function

A crux of the algorithmic process of closing local subspaces is scaling a local operator to act on a single spin in a global spin system. Take an operator $\hat{A}^L$ acting on a local set of L spins such that $\hat{A}^L$ has dimension 2^L. This operator is acting on spins $[0, L-1]$. To get $\hat{A}^L$ to act on spins $[i, i+L-1]$ in a larger system of N spins, a scaling function f of the kronecker product with the identity is defined for matrices, where $f(\hat{A}^L, i) = \hat{A}_i^L$. This is given below by

$$f(\hat{A}^L, i) := I_{2^i} \otimes \hat{A}^L \otimes I_{2^{N-L-i}} \tag{2.11}$$

Scaling function f preserves all relationships between basis vectors given by Equation 2.4, but scales the basis vectors to a higher dimensional subspace of Hilbert space. This transitive closure of a graph approach is built around the interaction of subspaces, so it is imperative to understand the way f acts on vectors. Define $\hat{A}^L$ to act on basis vectors in $U^L \in \mathcal{H}_{2^L}$ such that $\sum_{|u'\rangle \in U} \alpha_{|u'\rangle} |u'\rangle = \hat{A}^L |u\rangle$. The operator $\hat{A}^L_i$ characterizes dynamics around spin i in $\mathcal{H}_{2^N}$, and is given by $f(\hat{A}^L, i) = \hat{A}^L_i$. If the σ^z basis is the basis of choice, the relationships defined in Equation 2.4 are re-written in the higher dimensional space as

$$\sum_{|u'_N\rangle \in U^N} \alpha_{|u'_N\rangle} |u'_N\rangle = \hat{A}^L_i |u_N\rangle , \tag{2.12}$$

$$|u_N\rangle = |a\rangle \otimes |u_L\rangle \otimes |b\rangle , \ |a\rangle \in \mathrm{Col}(I_{2^i}), |b\rangle \in \mathrm{Col}(I_{2^{N-L-i}}) \tag{2.13}$$

This is a notationally dense formalism, so back to the iSWAP example. iSWAP shows that $|01\rangle \leftrightarrow |10\rangle$, irrespective of the dimension of the system it sits in. To make iSWAP an operator that acts on spin i in a larger system of spins, it must also encode the identity on all other spins $\neq i, i+1$. This is what 2.13 does. Repeat for every spin in the system, and a set of local symmetry protected subspaces is defined. Figure 2.3 on Page 10 works through this for the iSWAP acting on spins 1 and 2. Notice how the scaling preserves the exchange of states $|01\rangle \leftrightarrow |10\rangle$ in the middle two spins while accounting for all permutations of how the identity acts on spins 0 and 3.

A new set of basis vectors can be established with

$$U^N = \{|u_N\rangle\} = \{|a\rangle \otimes |u\rangle \otimes |b\rangle : |u\rangle \in U^L, |a\rangle \in \mathrm{Col}(I_{2^i}), |b\rangle \in \mathrm{Col}(I_{2^{N-L-i}})\} \tag{2.14}$$

Because each permutation of $|a\rangle, |u\rangle, |b\rangle$ is unique, each vector $|u_N\rangle \in U^N$ is also unique, so $\mathrm{Dim}(U^N) = 2^N$ and U^N can span $\mathcal{H}_{2^N}$. Each $|u_N\rangle$ can be decomposed into a unique $|a\rangle, |u_L\rangle, |b\rangle$ where $f_v(|a\rangle, |u_L\rangle, |b\rangle) = |u_N\rangle$. From Equation 2.13 this is

$$f_v(|a\rangle, |u\rangle, |b\rangle) := |a\rangle \otimes |u_L\rangle \otimes |b\rangle \tag{2.15}$$

and f_v is a true function which has a one-to-one mapping between $\mathcal{H}_{2^L}$ and $\mathcal{H}_{2^N}$.

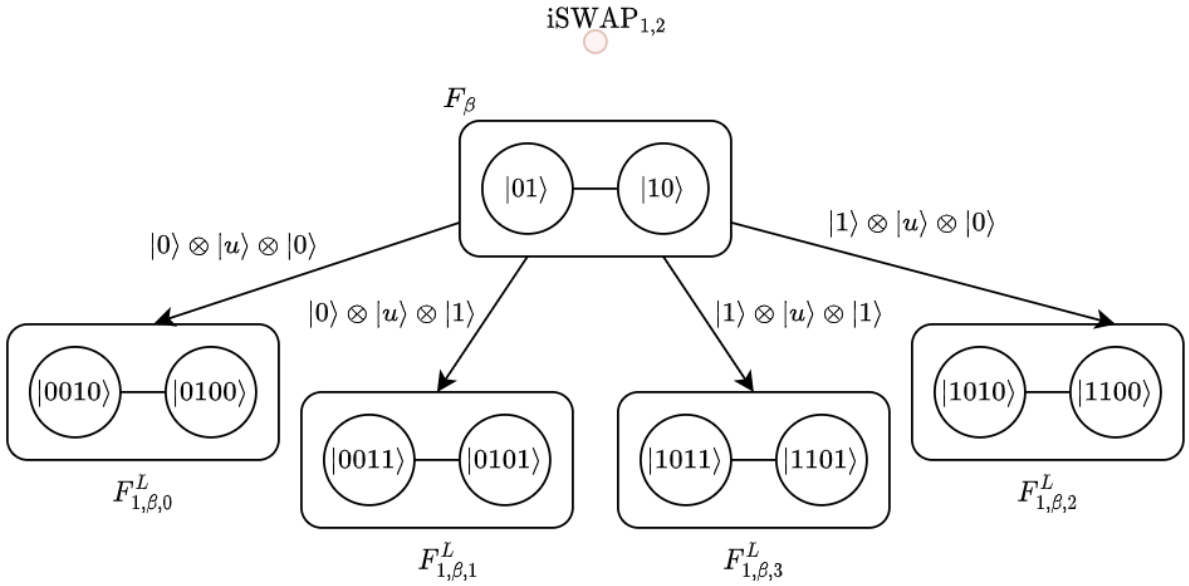

Figure 2.3 The vector scaling function in Equation 2.15 acting on the operator staring at spin $i = 1$.

2.2.2 Scaling Local Subspaces

The scaling function for vectors can be applied to sets of symmetry protected subspaces for a local operator A^L to define the set in a larger system of spins. Equation 2.13 shows that f preserves transitivity, and Equation 2.14 shows that f creates unique basis vectors in U^N. Therefore, to modify the G^L of A^L to describe G^L_i for $\hat{A}^L_i$, for each $F^L_j \in G^L$ apply the same f_v for vectors. This generates a new G^L_i, which has corresponding local symmetry protected subspaces $\{F^L_{i,j,ab}\}$. Each F^L_j is a collection of basis vectors, and so can be scaled with f_v. Each F^L_j is scaled explicitly with the formula

$$F^L_{i,j,ab} = \{\{f_v(|a\rangle, |u\rangle, |b\rangle) : |u\rangle \in F^L_j\} : |a\rangle \in \mathrm{Col}(I_{2^i}), |b\rangle \in \mathrm{Col}(I_{2^{N-L-i}})\} \tag{2.16}$$

which creates 2^{N-L} local subspaces for each F^L_j. These are the local total subspaces for each operator in the system, and by repeating the scaling function for every operator G' mentioned in Section 2.2 a collection of subspaces is form which can be transitively closed.

For the iSWAP example, this is shown in Figure 2.3. $F^L_{1,\beta,1}$ is the result of 2.16 with $|a\rangle = |0\rangle$ and $|b\rangle = |1\rangle$, $F^L_{1,\beta,2}$ is the result with $|a\rangle = |1\rangle$, $|b\rangle = |0\rangle$, and so on.

Once this total set of local symmetry protected subspaces is created, it can be transitively closed with the procedure mentioned in Section 2.2 and illustrated in Figure 2.2 on the previous page.

2.3 Subspace of a Single State

It is often useful to know the symmetry protected subspace $F_{|u\rangle}$ of a single state $|u\rangle$; this is the case for both post-selection and reduced basis construction. $F_{|u\rangle}$ is the transitive closure of state $|u\rangle$ under the unitary $\hat{A}$, but the unitary is sizeable and cannot be efficiently computed or evolved with. By borrowing methods developed in Sections 2.2 - 2.2.2, this subspace will be shown to be $O(N \times |F_{|u\rangle}|)$ to compute. The exact algorithm in pseudocode is given below, and the arguments for each step will follow.

Algorithm 1 Subspace of a single state

```
F ← ∅ // the actual subspace
F′ ← {|u⟩} // the new transitions
F″ ← ∅ // transitions from new transitions
while F′ ≠ ∅ do
    for |a⟩ ∈ F′ do
        F ← |a⟩
        T¹_|a⟩ // all states possible under a single time step of the unitary from |a⟩
        for |a′⟩ ∈ T¹_|a⟩ do
            if |a′⟩ ∉ F then
                F″ ← |a′⟩
            end if
        end for
    end for
    F′ ← ∅
    F′ ← F″
    F″ ← ∅
end while
```

The single state $|u\rangle$ has a set of states which are possible transitions under one discrete time step of the unitary; call this set $T^1_{|u\rangle}$. This set is difficult (but not impossible) to compute exactly because it requires a step of unitary time evolution. However, by taking $|u\rangle$ *backwards* through the kronecker scaling function for each operator $\hat{A}_i^L$ in $\hat{A} = \sum \hat{A}_i^L$, $|u\rangle$ can check its possible transitions under that just the operator at i, and scale the states of its transitions back to the larger system. This gives the set of states which $|u\rangle$ can transition to under one time step of the unitary $\hat{A}$. Define this set as $T^1_{|u\rangle}$. The formation of this set is shown by the iSWAP example in

Figure 2.4.

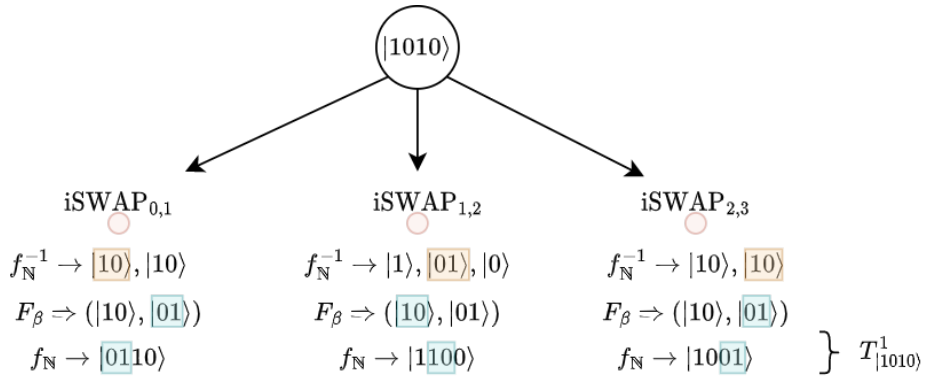

Figure 2.4 Given a state $|1010\rangle$, reverse the scaling for each local unitary, identify local transitions, and scale local transitions back to the larger system. Orange is the starting state, green is a state which a transition exists to, in the smaller basis.

Each state $|t\rangle \in T^1_{|u\rangle}$ has its own transitions which are one discrete time step away. Let the set $T^2_{|u\rangle}$ be each of these transitions two steps away, or $T^2_{|u\rangle} = T^1_{|u\rangle} \cup \{T^1_{|t\rangle} : |t\rangle \in T^1_{|u\rangle}\}$. This defines the recursion relation in Equation 2.19 and is illustrated in Figure 2.5 on the next page.

$$T^2_{|u\rangle} = T^1_{|u\rangle} \cup \{T^1_{|t\rangle} : |t\rangle \in T^1_{|u\rangle}\} \tag{2.17}$$

$$\vdots \tag{2.18}$$

$$T^i_{|u\rangle} = T^{i-1}_{|u\rangle} \cup \{T^1_{|t\rangle} : |t\rangle \in T^{i-1}_{|u\rangle}\} \tag{2.19}$$

with the stop condition that $T^{i+1}_{|u\rangle} = T^i_{|u\rangle} \Leftrightarrow T^i_{|u\rangle} = F_{|u\rangle}$; when no new transitions are discovered, then the full symmetry protected subspace is known. This has time complexity $O(N \times |F_{|u\rangle}|)$; each $T^1_{|t\rangle}$ is calculated by finding each local symmetry protected subspace of $|t\rangle$, which is $O(N)$, and $T^1_{|t\rangle}$ needs to be calculated at most 1 time for each $|t\rangle \in F_{|u\rangle}$, giving $O(N \times |F_{|u\rangle}|)$.

With the process for computing the full subspace in hand, the next few paragraphs are centered around the inverse scaling function and how it is used to find states a single step away. Recall that the basis U has been chosen as the σ^z basis. In this basis, each basis vector has an equivalent natural number representation which is the index of its single 1 value, also known as its binary index, which lets an inverse of f_v be easily computed. Under the hood, a kronecker product of a vector with the identity can be represented as changing the indices of the vector. This allows the vector scaling function f_v in Equation 2.15 to be written as Equation 2.20, where

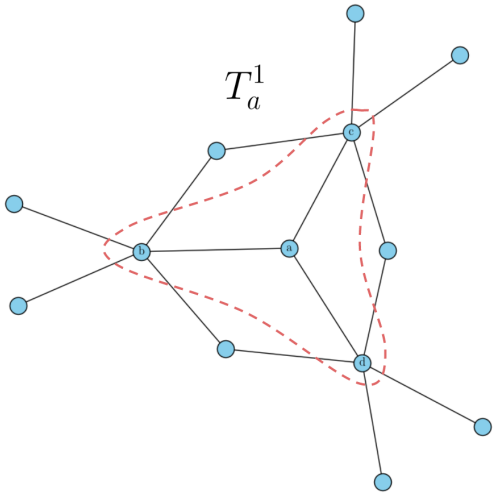

(a) The single transitions from a, T_a^1

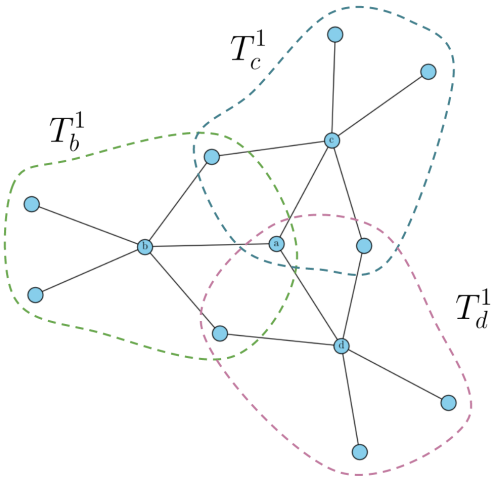

(b) The single transitions of each state in T_a^1

$T_a^2 = T_a^1 \cup T_b^1 \cup T_c^1 \cup T_d^1$

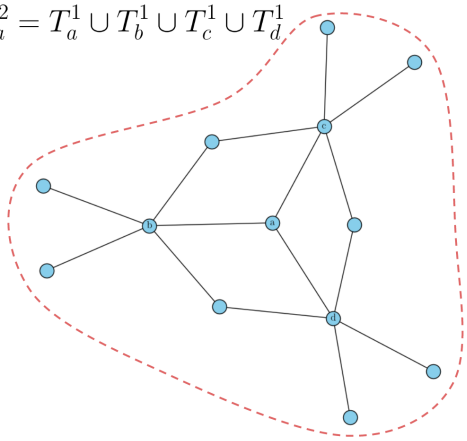

(c) The double transitions of a, T_a^2, which is the union of all states in T_b^1, T_c^1, T_d^1

Figure 2.5 The figure shows how given a which can transition to b, c, d in one time step, a recursive process can be followed to get states reachable in more and more time steps.

$|a\rangle, |b\rangle, |u_L\rangle$ have been exchanged for $a, b, u_L \in \mathbb{N}$.

$$f_{\mathbb{N}}(a, u_L, b) := 2^{N-L-i}u_L + 2^{N-i}a + b \tag{2.20}$$

Section 2.2.1 shows that each $f_{\mathbb{N}}(a, u_L, b)$ has a unique result. To get the single transitions T_u^1 from u, for each operator A_i^L which has a complete cluster graph G_i^L find the inverse of $f_{\mathbb{N}}$, which is given by

$$u_L = \lfloor 2^{L+i-N}u - 2^L \lfloor 2^{i-N}u \rfloor \rfloor \tag{2.21}$$
$$a = \lfloor 2^{i-N}u \rfloor \tag{2.22}$$
$$b = u - 2^{N-L-i}u_L - 2^{N-i}a \tag{2.23}$$

For each spin site in the unitary, there is a subspace $u_L \in F^L_{u_L}$, $F^L_{u_L} \in G^L$ given by the single operator A^L at i. This subspace $F^L_{u_L}$ can be scaled with the a, b found by Equation 2.21, giving the local subspace $F^L_{i,u}$ in the system of N spins. Equation 2.20 can be re-written as

$$f_{u'}(u, u_L, u'_L) := u + 2^{N-L-i}(u'_L - u_L) \tag{2.24}$$

and applied to each item in $F^L_{u_L}$ such that

$$F^L_{i,u} = \{f_{u'}(u, u_L, u'_L) : u'_L \in F^L_u\} \tag{2.25}$$

which is repeated to find $F^L_{i,u}$ for each spin site. Because u is in each local symmetry space of each $F^L_{i,u}$, the union of each of these operators gives T^1_u. Formally;

$$T^1_u = \cup\{F^L_{i,u} : i \in [0, N-1]\} \tag{2.26}$$

T^1_u is $O(N)$ to compute. With an efficient way to find single discrete time step transitions from u in hand, a symmetry protected subspace of u can be formed using the recursion relation in Equation 2.19. One iteration of Equation 2.19 can be seen in Figure 2.5 on Page 13.

Once the stop condition $T^{i+1}_u = T^i_u$ is reached, every $|u'\rangle \in T^i_u$ only has transitions to other vectors already present in T_u; therefore, T_u has been transitively closed and is the symmetry protected subspace of $|u\rangle$, $F_{|u\rangle}$.

2.4 Usage and Limits

While a pseudo-analytical approach for all symmetry which can be observed in a given basis is interesting in its own right, the aims of this algorithm are more practical in nature, and as such its limits are characterized. As mentioned in Chapter 1, this work uses the Heisenberg XXX, T2 Quantum Cellular Automata, and F4 Quantum Cellular Automata models to characterize how symmetry protected subspaces form and scale. Refer to Figure 2.6 on the next page to see group sizes at each number of spins. For the studied models, the maximum subspace size was asymptotically limited by 2^{N-L}, where L is the number of spins interacting in the local operators

in Equation 2.8.

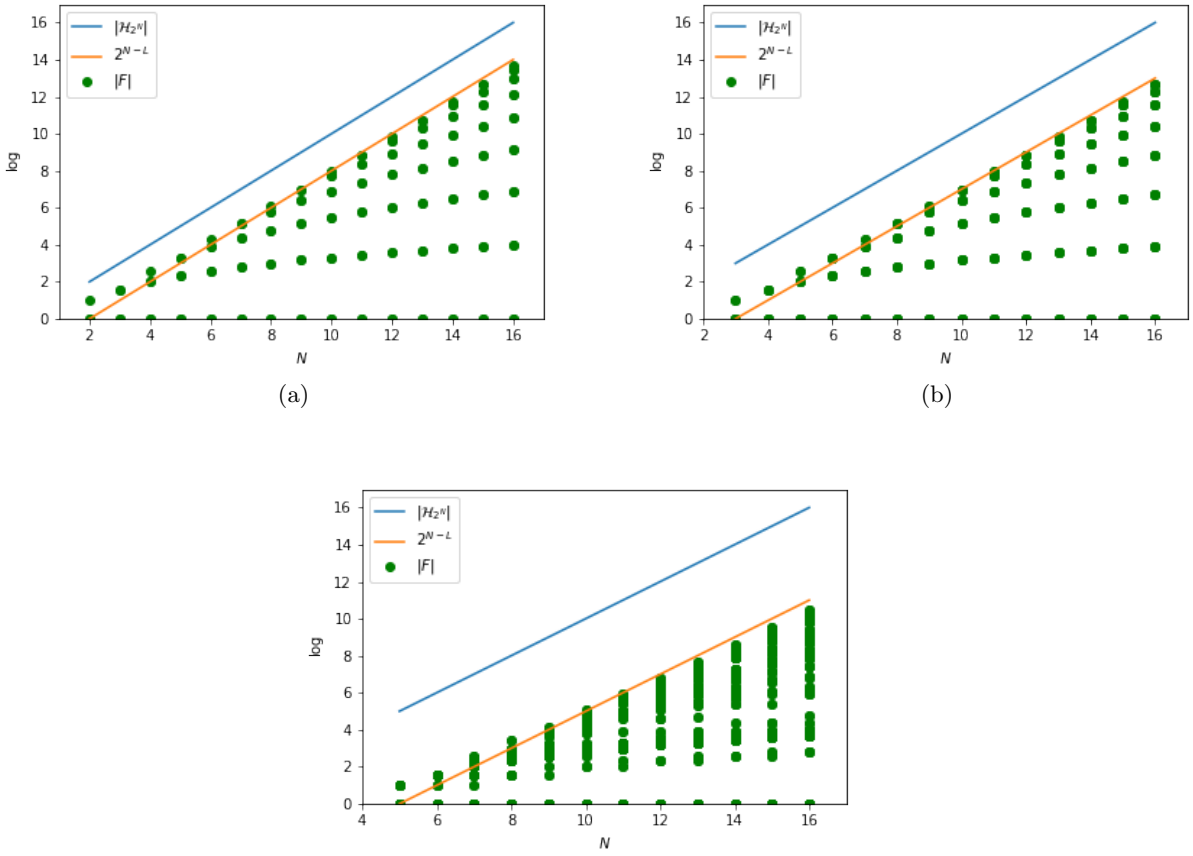

(a) (b)

(c)

Figure 2.6 Subspace sizes for (a) Heisenberg XXX (b) T2 QCA (c) F4 QCA.

To do post-selection of a circuit starting in state $|u\rangle$, $F_{|u\rangle}$ must be computed. For each measurement $|u'\rangle$, check if $|u'\rangle$ is in $F_{|u\rangle}$. This check is efficient, because $F_{|u\rangle}$ can be stored with hashing or binary trees. As seen in Figure 2.6, $|F_{|u\rangle}|$ grows with $O(2^{N-L})$ in the worst case. Therefore, for quantum algorithms simulating many spins, an alternative approximate post-selection is suggested in Chapter 3.

This procedure can be used for reduced basis creation if $|\mathcal{H}_{2^N}|$ is too large but $|F_{|u\rangle}|$ is tractable at the target number of spins. Figure 2.6 shows that if there is a reduced basis, it can extend classical simulation L spins beyond the current limit.

2.5 Future Work

The automated procedure used in this chapter leverages many assumptions, such as removing amplitudes from a basis vector decomposition and restricting to the σ^z basis. This algorithm would need to be modified to account for non-σ^z basis symmetries, which could be essential for some post-selection applications where an alternate basis is used for measurement or where a symmetry only exists in an alternate basis.

CHAPTER 3

POST-SELECTION IN LARGE SYMMETRY PROTECTED SUBSPACES

Post-selection is a method for error mitigation which can identify when a specific measurement of a quantum computer can only be the result of error [13] [10]. It relies on observing a conserved quantity across time evolution, which is implemented through a symmetry operator $\hat{S}$ which commutes with the unitary evolution operator $\hat{A}$ such that $[S, \hat{A}] = 0$. As described in Chapter 2, these symmetry operators are difficult to discover. Through the use of symmetry protected subspaces defined in the aforementioned chapter, a subspace of basis states $F \subseteq \mathcal{H}_{2^N}$ can be defined where the unitary time evolution of a state which starts in the subspace F will always live within F.

Forming a subspace is only the first problem; as discussed in Section 2.4, the worst case for subspace size $|F|$ is $O(2^{N-L})$, which is intractable to compute even with a near-linear algorithm for large numbers of spins. This post-selection process offers a bit of hope; the entire subspace is not necessarily needed, just a way of validating if a given measurement exists in the subspace as the starting state.

This chapter focus on addressing the question: given two states, do they exist in the same symmetry protected subspace? Section 3.1 discusses the pathfinding approach to verifying a measurement, and Section 3.2 discusses how greedy algorithms can give a probabilistic argument for solving the pathfinding problem, and when that argument fails. For verification, three unitary operators; Heisenberg XXX, T2 Quantum Cellular Automata, and F4 QCA, are used and their conditions for failure discussed in 3.2.2.

3.1 Post-selection as a Pathfinding Problem

Post-selection can be stated as a pathfinding problem; if two states are in the same subspace, then there has to be a path of basis vectors connecting the two states by the transitivity relation discussed in Section 2.1.1. Therefore, if some path $R(|a\rangle, |b\rangle)$ can be discovered, where $|a\rangle$ is the start state and $|b\rangle$ is the measurement, it is proven that $\langle b| [\hat{A}]^t |a\rangle \neq 0$ for some t. Given the full subspace F, which can be interpreted as a complete graph, this path has length 1 and is trivial to compute. In the case where F is too large, the pathfinding question becomes more nuanced.

Recall from Section 2.3 that the single time step transitions from a single state $|a\rangle$ are given by $T^1_{|a\rangle}$, which is $O(N)$ to compute. Define a graph P which has edges $|a\rangle \rightarrow |a'\rangle\,,\ |a'\rangle \in T^1_{|a\rangle}$. P will never be fully computed; instead of computing all of P, edges in P are found from the current node with this function. With a given state $|a\rangle$, P can be queried to find every edge from that state by computing $T^1_{|a\rangle}$. If $T^1_{|a\rangle}$ is used to find the transitive closure of P, $P \equiv F$.

There are $O(N)$ edges given by $|a\rangle$ in $T^1_{|a\rangle}$ from a node $|a\rangle$; if there are $O(2^{N-L})$ nodes in P, then P is densely connected and the path between two nodes is $O(N)$ edges [14]. This is supported by the idea that one state is $O(N)$ time steps from another in quantum time evolution. Therefore, each pair of nodes in P is connected by a path with $O(N)$ edges. If the path $R(|a\rangle\,, |b\rangle)$ is known, the time complexity of traversing it is $O(N^2)$.

In general, the path R is not known, and finding it for two arbitrary states would require a brute force search of the entire P; this is the same time cost as just computing the subspace, which is unfeasible. But if R starting at $|a\rangle$ can be predicted, it is $O(N^2)$ to verify that b lies at the end of it. This is done by evaluating the current set of next edges in the graph P and deciding which one to take for each node in the path. Section 3.2 examines how to walk across P, starting at $|a\rangle$, by only predicting the next step in the sequence.

3.2 Greedy Algorithm for Post-selection

To traverse from $|a\rangle$ to a measured state $|b\rangle$ by taking path R through graph P and keeping $|R| \ll |P|$, R must traverse a linear number of nodes in P. This can be done if the next step in R is always known, which is achieved by using a greedy algorithm with the story that follows.

Use a greedy algorithm that intelligently guesses which edge to take towards $|b\rangle$ from the current $|c\rangle$ on the path $R(|a\rangle\,, |b\rangle)$. If it is used to predict a path $R(|a\rangle\,, |b\rangle)$ it needs to have a stopping condition; otherwise, the entire graph P will be explored in the case where $R(|a\rangle\,, |b\rangle)$ does not exist. With a greedy algorithm, if $|c\rangle$ is the current node in the path $R(|a\rangle\,, |b\rangle)$ with known neighbors given by T^1_c, then the next node $|c'\rangle \in T^1_{|c\rangle}$ is taken iff $|c'\rangle$ is closer to $|b\rangle$ than $|c\rangle$ is, by some distance metric. If there are no moves from $|c\rangle$ which bring the path closer to $|b\rangle$, then it is assumed that no path exists and $|b\rangle$ is not in the same subspace as $|a\rangle$.

Now a distance metric must be established. If the distance metric can quantify how close two states are from each other in unitary time evolution, the path is certain to verify.

It proves useful to use the binary index, with $|a\rangle \to a$, $|b\rangle \to b$, of each state to define the distance metric with old-fashioned subtraction. First, note that a small natural number distance between a, b does not imply that $|a\rangle, |b\rangle$ are close to each other in unitary time evolution. This means that just moving from a to b with natural number distances may lead to false negatives in the case where the next greedy move is the lowest natural number distance but not closest in time evolution. Instead, choose a node which a and b should always find with reasonable probability in a greedy pathfinding algorithm; if there are two paths $R(a, c)$, and $R(b, c)$ in an undirected graph, then there must be a path $R(a, b)$.

It is numerically shown that the absolute minimum or maximum of the binary indexes of the nodes in P almost always lies at the end of the greedy minimum/maximum finding algorithm, and can thus verify the existence of a path. The following details are all given in terms of a minimum, but are the same for a maximum. The greedy algorithm to find the minimum node from a is defined by the recursion relation

$$a^1 = \min(T_a^1) \tag{3.1}$$
$$a^i = \min(T_{a^{i-1}}^1) \tag{3.2}$$

and is halted by the condition $a^{i+1} = a^i$. Define the minimum node found by this algorithm to be $a_{\min}$. After computing $R(a, a_{\min})$ and $R(b, b_{\min})$, if $a_{\min} \equiv b_{\min}$ then there exists a path from a to b, and thus $b \in F_a$ and will not be thrown out by post-selection. If there are no false minima in a graph P, this algorithm will always ascertain if a, b live within the same symmetry protected subspace in $O(N^2)$ and is no longer approximate. If there are false minima which are not escaped by the method described in Section 3.2.1, post-selection has a chance of throwing out measurements in which no error occurred. This is destructive for an ideal quantum computer, as measurements which had no noise could still be discarded, which would act as a source of systematic noise for an ideal computation. However, this approximate post-selection scheme may still be valuable in the NISQ era, which is a question that will be explored further in Section 4.3.1.

3.2.1 Attempt to Escape False Minima

There is one small addition to the approximate post-selection algorithm which can significantly increase the chances of finding the true minimum of on a greedy path in P; by

considering T_a^{q+1}, where $q \geq 1$ but is still small, the path $R(a, a_{\min})$ has a chance of finding a better minimum. This can be implemented in two ways. The first is to take the recursive process for building on T_a^1 defined in Equation 2.17, and repeat it q times with $a^i = \min(T_{a^{i-1}}^{q+1})$. The second is to use a larger cluster graph G^{L+q} by using the graph building process described in Section 2.2 for each spin in the simulation. This work uses the second approach, but either achieves the same result. q adds an exponential scaling to the next-edges part of the algorithm, so it must be kept small as to not negatively impact performance.

3.2.2 Frequency of Failure and Failure Conditions

A numerical analysis of the three models: Heisenberg XXX, T2 QCA, and F4 QCA is done to identify when the approximate post-selection procedure can fail. As discussed in Section 3.2, false minima can be potentially dangerous for the application of post-selection. Define a rough subspace as one that has at least one false minimum. Up to $N = 17$, the total symmetry structure of each model is computed. Define the quantity $\sum |F_r|$ to be the sum of the sizes of all rough subspaces; if approximate post-selection always works, $\sum |F_r| = 0$

For the Heisenberg XXX model no rough subspaces were found, so approximate post-selection is equivalent to post-selection.

For the T2 Quantum Cellular Automata, rough subspaces only exist at $q = 0$. Otherwise, no rough subspaces were found. This behavior is shown in Figure 3.1 on Page 21, where $q > 0$ is essential to avoiding false minima. This plot shows that for $q = 1$, approximate post-selection is equivalent to post-selection.

For the F4 QCA, things get more interesting. Like the T2 QCA, $q = 0$ causes the rough subspaces to dominate the graph structure. However, unlike T2, all computed q values just cause marginal improvements but never cause $\sum |F_r|$ to truly reach zero. See Figure 3.2 on Page 21. Notice that while $\sum |F_r|$ decreases as q increases, the quantity never truly reaches zero. This means that approximate post-selection will always have a chance of removing ideal states in this model; while some subspaces of this model have no false minima, it is impossible to know if a given state a is in such a subspace without computing a full F_a. Fortunately, false minima are apparent at small N, so trivial numerical work can show if false minima can be expected for a given quantum algorithm.

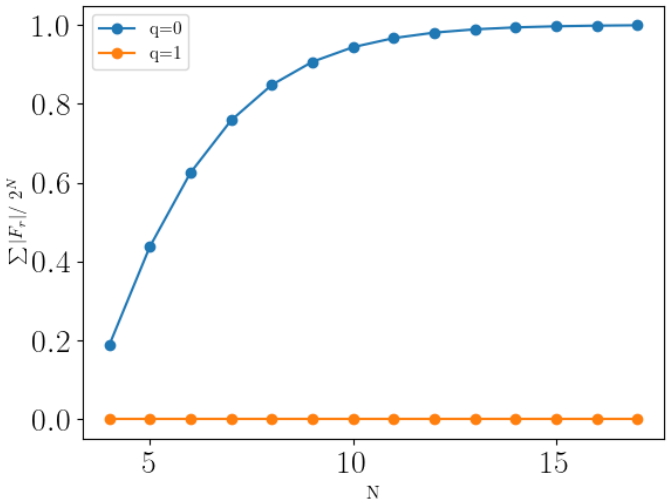

Figure 3.1 The portion of nodes in P which exist in a rough subspace and thus can take a greedy path to a false minimum, for the T2 QCA. Notice how for $q = 1$, this quantity drops to zero and remains there.

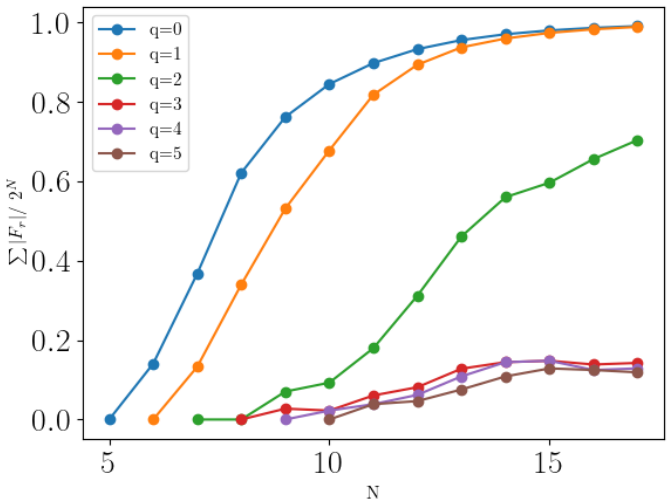

Figure 3.2 The portion of nodes in P which exist in a rough subspace and thus can take a greedy path to a false minimum, for the F4 QCA. For the limits this verification was carried to, $\sum |F_r|$ drops with increasing q but never truly reaches zero.

It is also possible that approximate post-selection creates a more intelligible result than a noisy circuit without it, despite potentially loosing results in which no errors occurred. Additionally, nodes which live in a rough subspace may still share a greedy minimum at smaller times, and thus coincidentally still have a path. However, this claim is still not analytically strong, and as such the only way to be really sure is to test, which is done for the same three models in Chapter 4.

CHAPTER 4

DEMONSTRATION OF POST-SELECTION

Simulations of the Heisenberg XXX Hamiltonian, T2 Quantum Cellular Automata [10], and F4 Quantum Cellular Automata [11] are performed in this chapter to benchmark how each algorithm can perform with post-selection implemented. These simulations are all a type of discrete time evolution on conditions seen in literature.

This chapter will use the Kullback-Leibler divergence function to quantify a given measurement distribution's divergence from the ideal quantum computation, followed by a description of how the ideal versus error simulations were performed. Each model will present its own simulation conditions. The data shows that post-selection raises the noise floor such that a computation with noise is guaranteed to remain closer to an ideal computation. There is also an examination of how approximate post-selection works in a subspace which is known to have false minima. Finally, a conclusion about post-selection, its use cases, and a discussion about its time complexity implications is given.

4.1 Methods for Demonstration

In quantum simulations, each time step has a corresponding sequence of measurements from which a probability distribution of states in the σ^z basis is formed. Once errors proliferate, this distribution becomes uniform random noise, given by Equation 4.1.

$$urn = \frac{1}{2^N} \sum_{i=0}^{2^N} |i\rangle \tag{4.1}$$

Kullback-Leibler divergence is a function given by Equation 4.2 that quantifies the difference between a probability distribution P and a reference probability distribution Q.

$$D_{KL}(P, Q) := \sum_{x \in X} P(x) \log\left(\frac{P(x)}{Qx}\right) \tag{4.2}$$

$|\psi_{\text{ideal}}(t)|^2$ is the probability distribution of the noise-free quantum simulation at time t, and $|\psi_{\text{noise}}(t)|^2$ is the probability distribution at the same time with noise. Define the function in Equation 4.3 which will smoothly decay from 1 to 0 as the simulation decays to uniform random noise [15].

$$\mathcal{D}(|\psi_{\text{noise}}|^2, |\psi_{\text{ideal}}|^2) := 1 - D_{KL}(|\psi_{\text{noise}}|^2, |\psi_{\text{ideal}}|^2)/D_{KL}(urn, |\psi_{\text{ideal}}|^2) \tag{4.3}$$

If post-selection is applied to a state which has decayed to uniform random noise, the relative probability distribution of states in the subspace F will remain uniformly random; the projector $\hat{P}_F$ applied to urn and renormalized gives the definition of a uniform random subspace.

$$urs = \frac{1}{|F|} \sum_{j \in F} |j\rangle \tag{4.4}$$

Each simulation verifies that the post-selected probability distribution converges to urs, because $\mathcal{D}$ of the post-selected distribution converges to $\mathcal{D}_{urs}$ of the uniform random subspace urs at sufficiently many measurements. Note that urs does not provide a solid noise floor, just a general trend that the simulation will converge to. Thus, small shifts in the $\mathcal{D}$ divergence from $\mathcal{D}_{urs}$, as seen in the following sections, are not surprising.

4.1.1 Methods

Data was collected with classical simulation of a quantum computer 10,000 measurements per time step. Error channels are defined by hand as the following; after every discrete step in the time evolution, apply the operator σ^x on each qubit with probability p followed by the operator σ^y with the same probability p. The total per qubit time step error, the probability of any σ^x, σ^y error on a single qubit after each Trotter step, is then given by the binomial distribution $\sum_{i=1}^{2} \binom{2}{i} p^i (1-p)^{2-i}$.

4.2 Known Symmetry Operator

This section computes $\mathcal{D}$ for each model, which had previously known symmetries.

4.2.1 Heisenberg XXX Hamiltonian

The Heisenberg XXX Hamiltonian is given by

$$H = \sum_{i=0}^{N-2} \sigma_i^x \sigma_{i+1}^x + \sigma_i^y \sigma_{i+1}^y + \sigma_i^z \sigma_{i+1}^z \tag{4.5}$$

and initialized in the state $|\psi(0)\rangle = |101010101010101\rangle$. This model has symmetry $\hat{S} = \sum_{i=0}^{N} \sigma_i^z$, which is well-studied and intuitive. The evolution unitary for it is defined as $\hat{A} = e^{iHdt}$. Results can be seen in Figure 4.1 on the next page; post-selection keeps the noisy data coherent for a

longer time. Figure 2.6 on Page 15 shows that the Heisenberg subspaces are fairly large; therefore, the *urs* of Heisenberg is very close to *urn*.

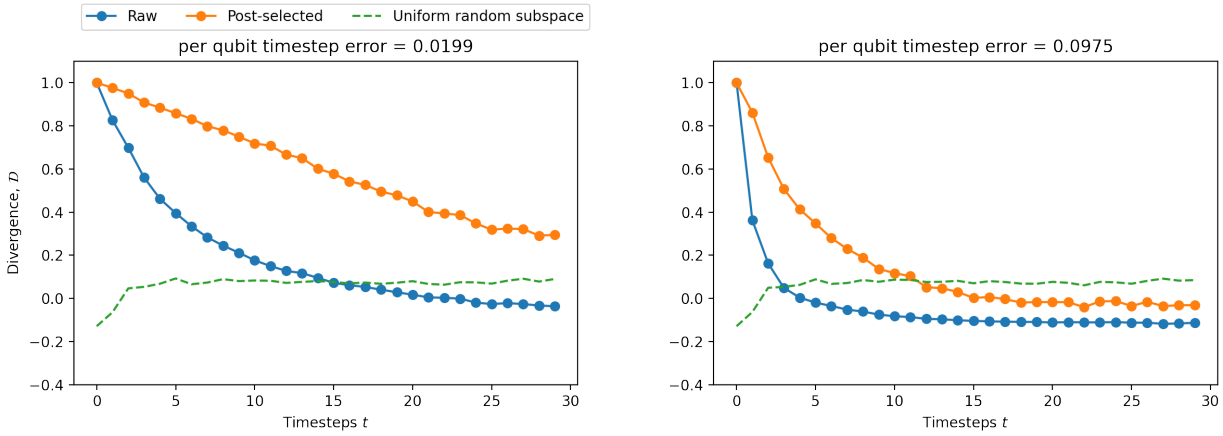

Figure 4.1 Divergence of Heisenberg XXX to uniform random noise.

4.2.2 T2 Quantum Cellular Automata

The T2 QCA applies a Hadamard gate $\hat{H}_i$ to a qubit if exactly one of its neighbors is in the $|1\rangle$ state. It is given by

$$\hat{A} = \sum_{i=1}^{N-2} \hat{H}_i \sum_{\alpha,\beta\in\{0,1\},\alpha+\beta=1} \hat{P}_{i-1}^{\alpha} \hat{P}_{i+1}^{\beta} \tag{4.6}$$

and initialized in $|\psi(0)\rangle = |0000000010000000\rangle$. This model has symmetry $\hat{S} = \sum_{i=0}^{N-1} \sigma_i^z \sigma_{i+1}^z$. Results can be seen in Figure 4.2 on the next page. Figure 2.6 on Page 15 shows that this model has a lower cap to its subspace size, so its uniform subspace noise is further from uniform random noise and closer to the ideal measurement distribution.

Additionally, time evolution of this model computes the $\langle n_i \rangle$, or the average number of times a $|1\rangle_i$ appears at a given qubit i. As this is a trivial expectation value to implement with σ^z measurement, it is included here in Figure 4.3 on Page 26 to characterize post-selection more. This example is useful, because it shows that post-selection can establish dynamics of a simulation which would have been otherwise hidden; the periodicity seen in the ideal example is invisible in the noisy data but appears in the post-selected data.

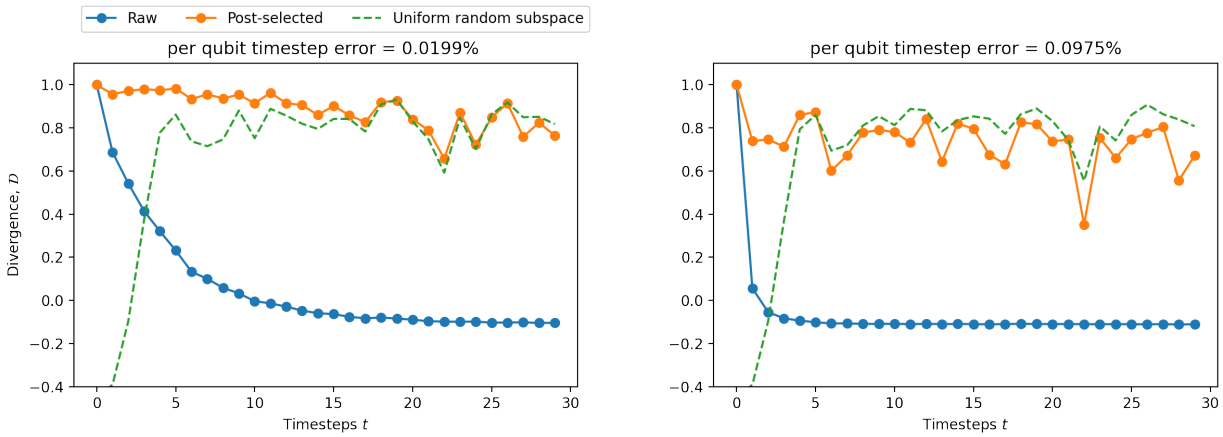

Figure 4.2 Divergence of T2 QCA. This model has a smaller $|F_{|\psi(0)\rangle}|$, which gives a higher uniform subspace noise.

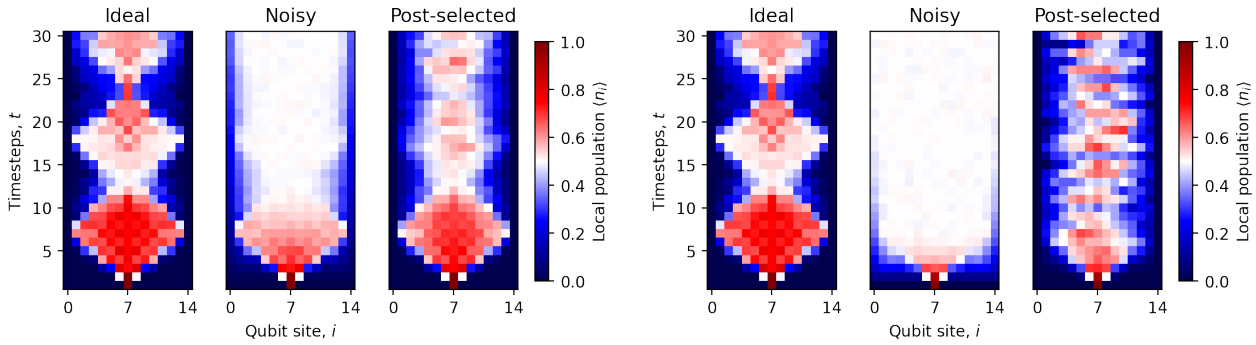

(a) (b)

Figure 4.3 Dynamics of T2 QCA. Per qubit time step error (a)= 0.0199 (b) = 0.0975.

4.3 Unknown Symmetry Operator - F4 Quantum Cellular Automata

This section computes $\mathcal{D}$ for an operator with no known $\hat{S}$, which has symmetry protected subspaces discovered by this algorithm.

The F4 QCA applies a Hadamard gate on a qubit if exactly two of its neighbors and next-nearest neighbors are in the $|1\rangle$ state. It is given by

$$\hat{A} = \sum_{i=2}^{N-3} \hat{H}_i \sum_{\alpha,\beta,\gamma,\delta\in\{0,1\},\alpha+\beta+\delta+\gamma=2} \hat{P}_{i-2}^{\alpha}\hat{P}_{i-1}^{\beta}\hat{P}_{i+1}^{\gamma}\hat{P}_{i+2}^{\delta} \tag{4.7}$$

and initialized in $|\psi(0)\rangle = |0000001010000000\rangle$. This start state lives in a subspace with no false minima, so post-selection is not approximate despite this system being covered mostly by rough

subspaces as seen in Figure 3.2 on Page 21. No $\hat{S}$ is previously known for this operator. However, this operator does have symmetry protected subspaces, discovered by the current work. Its post-selected $\mathcal{D}$ is given in Figure 4.4, with the standard convergence to *urs* observed.

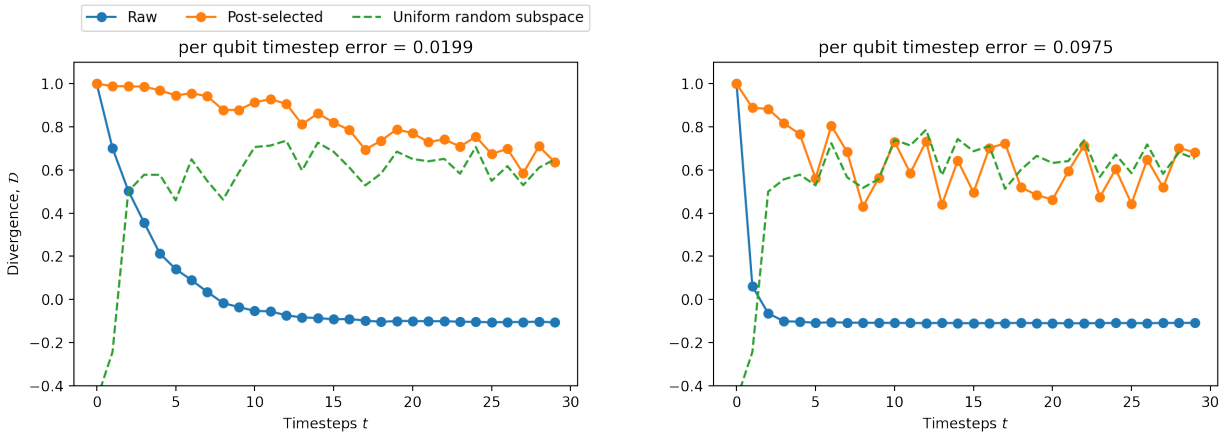

Figure 4.4 Divergence of F4 QCA. This unitary has no known analytical symmetries.

The $\langle n_i \rangle$ for each qubit, just as measured with the T2 QCA, is calculated and shown in Figure 4.5 on Page 28. Here, post-selection does not help observe nuanced properties of the unitary time evolution; this particular QCA is valuable because of the 'blinking' of the occupancy of qubit 7 in time. At one time step, it has $\langle n_7 \rangle \approx 0$, and the next it has $\langle n_7 \rangle \approx 1/2$. The subspace has no time dependence, so post-selection will accept states with either feature (provided they are in the subspace) regardless of time step. Because the circuit is overcome with noise and many errors will remain within the subspace, some errors will emphasize the feature opposite the current value of $\langle n_7 \rangle$ and the more nuanced observation is lost. Post-selection does maintain qualitative properties which are less nuanced; qubits 6 and 8 should remain in $\langle n_i \rangle \approx 1$ throughout the time evolution. This property is lost in the noisy time evolutions, but recovered at least in part with post-selection on the lower and higher error simulations.

Note that without the procedure developed in this paper, this data would remain in the noisy plots of time evolution.

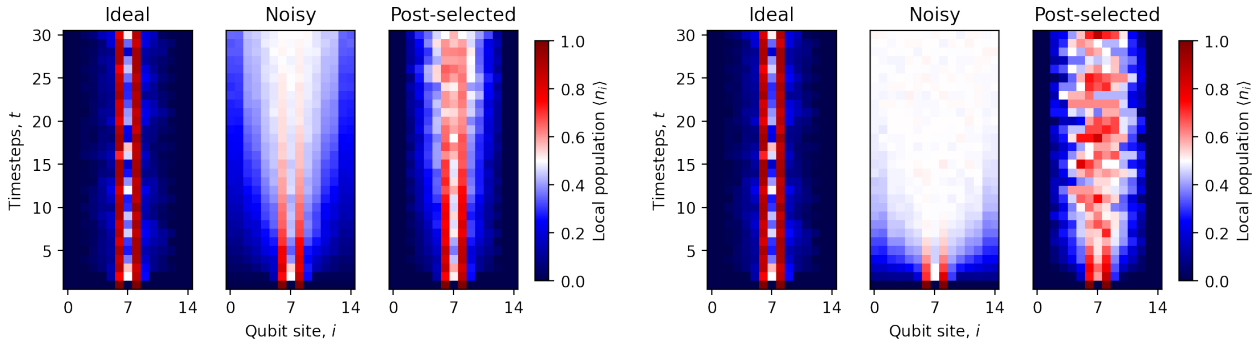

Figure 4.5 Dynamics of F4 QCA. Per qubit time step error (a) = 0.0199 (b) = 0.0975.

4.3.1 Subspace With Local Minima

Section 3.2.2 mentions that the F4 QCA has states which lie within the failure conditions for approximate post-selection. This section identifies a starting state whose subspace is a previously defined rocky subspace. This is the initial state $|\psi(0)\rangle = |101010111011111\rangle$, and post-selection is run with various q as defined in Section 3.2.1. In the studied system of 15 qubits, $L + q$ is close to N when $q = 7$, which implies that the q used for G^{L+q} is causing the precomputed total subspace to converge to the total subspace. In practice, the maximum computable q is likely not as close to the number of qubits.

This set of divergences is given for low per qubit time step error in Figure 4.6 and high in Figure 4.7 on Page 29. Finally, Figure 4.8 on Page 30 shows the $\mathcal{D}$ of using an approximate post-selection on a noise-free circuit. As can be seen in the figures, post-selection never gets worse than noisy data. On top of this, it is only destructive to an ideal wavefunction at larger t. Likely, this is because nodes which are close to the starting state in the graph P share an approximate minimum, so only a small set of transitions from the starting state are accepted.

As seen in the plots, only in the worst case scenario noise and with small q is approximate post-selection is worse than raw data. As such, the developed method for approximate post-selection is useful in all cases studied here.

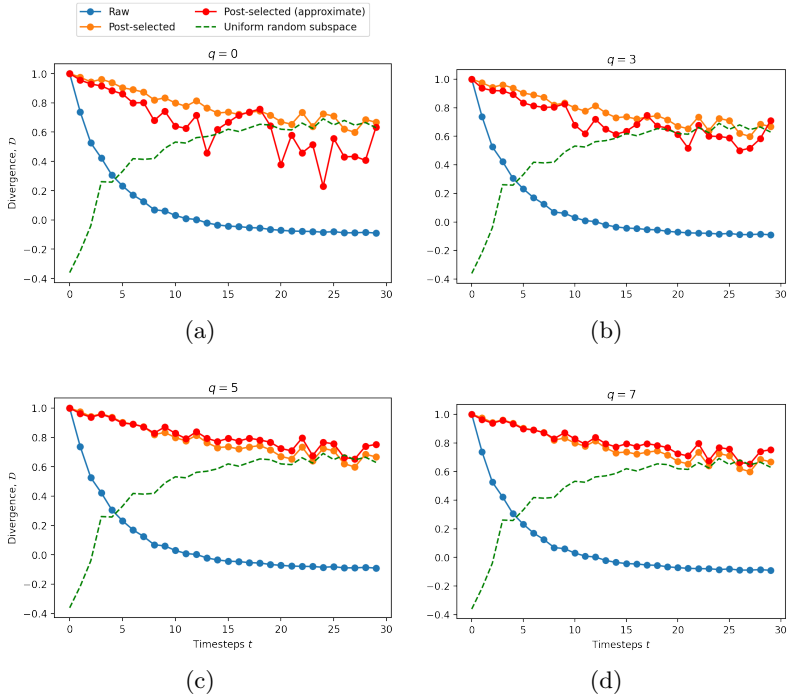

Figure 4.6 Approximate post-selection for per qubit time step error = 0.0199.

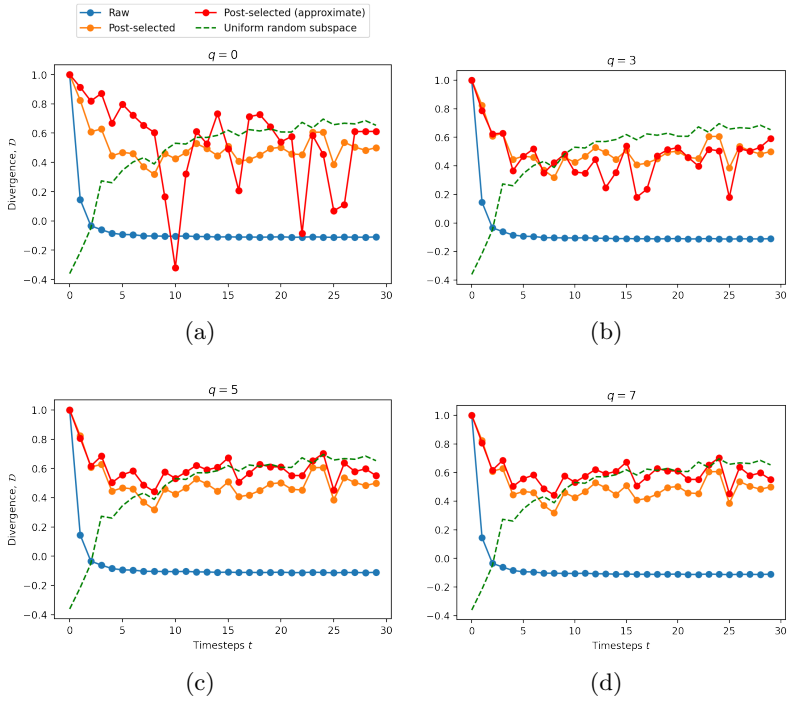

Figure 4.7 Approximate post-selection for per qubit time step error = 0.0975.

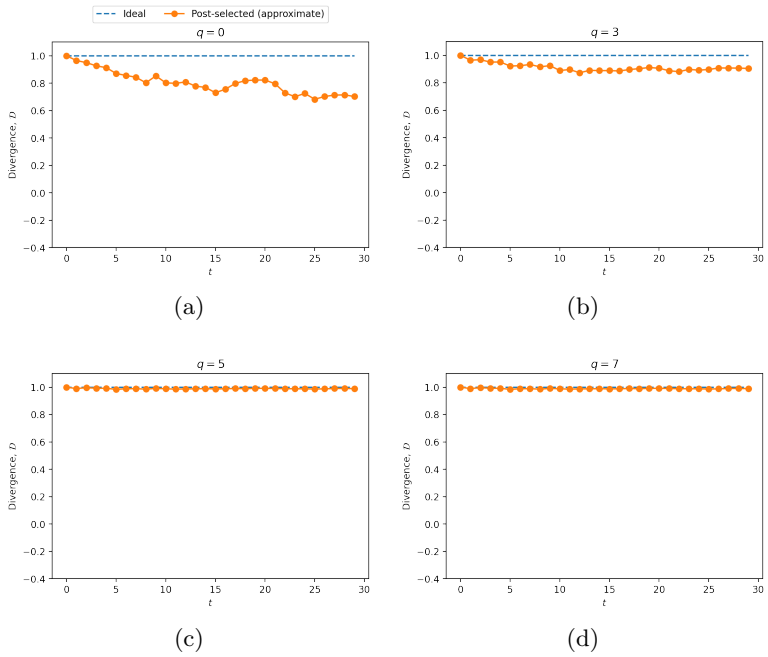

Figure 4.8 The divergence from using approximate post-selection on the ideal state.

4.4 Discussion and Conclusions

Post-selection has some nuances to its application which should be analyzed. Some quantum simulations involve analysis on the final probability distribution, such as to compute expectation values. As shown in this chapter, even without increasing the number of measurements post-selection achieves a probability distribution significantly closer to the ideal one, and is therefore extremely useful when a simulation can implement it. However, many quantum algorithms require a certain number of measurements to achieve a desired fidelity; post-selection then adds an exponential number of measurements to achieve the same fidelity. Future work could explore the true time complexity bound of this statement and analyze how much worse the fidelity is if not every error is replaced with a new measurement.

Symmetry protected subspaces have many applications to quantum information, which future work could explore. In quantum optimizers, initial states are chosen to either limit or maximize the simulations search space; if the symmetry structure of a system is known, these states can be chosen more intentionally, which could improve optimizer performance [16]. Because this algorithm can identify if there exists symmetry in a given basis, a basis can be chosen which has

symmetry apparent from this algorithm. This would allow either smaller, more efficient simulations or for post-selection. This algorithm can also discover symmetry which was not already known; if this is the case with any fundamental model, it could imply the existence of unknown properties in that model.

This work presented an algorithm to compute the symmetry protected subspace, which a starting state remains in under unitary time evolution. These subspaces likely have interesting implications for quantum systems where they were previously hidden. Using this algorithm, symmetry was discovered in the F4 QCA, a quantum system which previously had no known symmetry. Because the full subspace of a state is linear in the size of the subspace, which scales exponentially, an approximate algorithm for post-selection algorithm was presented and its conditions for failure shown. Finally, this approximate algorithm was used to post-select emulated quantum simulation data with error, which demonstrated that even approximate post-selection leads to a probability distribution closer to the ideal simulation.

REFERENCES

[1] Xiao Mi, Matteo Ippoliti, Chris Quintana, Ami Greene, Zijun Chen, Jonathan Gross, Frank Arute, Kunal Arya, Juan Atalaya, Ryan Babbush, Joseph C. Bardin, Joao Basso, Andreas Bengtsson, Alexander Bilmes, Alexandre Bourassa, Leon Brill, Michael Broughton, Bob B. Buckley, David A. Buell, Brian Burkett, Nicholas Bushnell, Benjamin Chiaro, Roberto Collins, William Courtney, Dripto Debroy, Sean Demura, Alan R. Derk, Andrew Dunsworth, Daniel Eppens, Catherine Erickson, Edward Farhi, Austin G. Fowler, Brooks Foxen, Craig Gidney, Marissa Giustina, Matthew P. Harrigan, Sean D. Harrington, Jeremy Hilton, Alan Ho, Sabrina Hong, Trent Huang, Ashley Huff, William J. Huggins, L. B. Ioffe, Sergei V. Isakov, Justin Iveland, Evan Jeffrey, Zhang Jiang, Cody Jones, Dvir Kafri, Tanuj Khattar, Seon Kim, Alexei Kitaev, Paul V. Klimov, Alexander N. Korotkov, Fedor Kostritsa, David Landhuis, Pavel Laptev, Joonho Lee, Kenny Lee, Aditya Locharla, Erik Lucero, Orion Martin, Jarrod R. McClean, Trevor McCourt, Matt McEwen, Kevin C. Miao, Masoud Mohseni, Shirin Montazeri, Wojciech Mruczkiewicz, Ofer Naaman, Matthew Neeley, Charles Neill, Michael Newman, Murphy Yuezhen Niu, Thomas E. O'Brien, Alex Opremcak, Eric Ostby, Balint Pato, Andre Petukhov, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vladimir Shvarts, Yuan Su, Doug Strain, Marco Szalay, Matthew D. Trevithick, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Juhwan Yoo, Adam Zalcman, Hartmut Neven, Sergio Boixo, Vadim Smelyanskiy, Anthony Megrant, Julian Kelly, Yu Chen, S. L. Sondhi, Roderich Moessner, Kostyantyn Kechedzhi, Vedika Khemani, and Pedram Roushan. Time-crystalline eigenstate order on a quantum processor. *Nature*, 601(7894):531–536, nov 2021. doi: 10.1038/s41586-021-04257-w. URL https://doi.org/10.1038%2Fs41586-021-04257-w.

[2] Aron Beekman, Louk Rademaker, and Jasper van Wezel. An introduction to spontaneous symmetry breaking. *SciPost Physics Lecture Notes*, dec 2019. doi: 10.21468/scipostphyslectnotes.11. URL https://doi.org/10.21468%2Fscipostphyslectnotes.11.

[3] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86:032324, Sep 2012. doi: 10.1103/PhysRevA.86.032324. URL https://link.aps.org/doi/10.1103/PhysRevA.86.032324.

[4] Minh C. Tran, Yuan Su, Daniel Carney, and Jacob M. Taylor. Faster digital quantum simulation by symmetry protection. *PRX Quantum*, 2(1), feb 2021. doi: 10.1103/prxquantum.2.010323. URL https://doi.org/10.1103%2Fprxquantum.2.010323.

[5] David T. Stephen, Henrik Dreyer, Mohsin Iqbal, and Norbert Schuch. Detecting subsystem symmetry protected topological order via entanglement entropy. *Physical Review B*, 100(11), sep 2019. doi: 10.1103/physrevb.100.115112. URL https://doi.org/10.1103%2Fphysrevb.100.115112.

[6] Tian-Shu Deng and Lei Pan. Fate of symmetry protected coherence in open quantum system. *Physical Review B*, 104(9), sep 2021. doi: 10.1103/physrevb.104.094306. URL https://doi.org/10.1103%2Fphysrevb.104.094306.

[7] Margarite L. LaBorde and Mark M. Wilde. Testing symmetry on quantum computers, 2021. URL https://arxiv.org/abs/2105.12758.

[8] Nanqing Dong, Michael Kampffmeyer, Irina Voiculescu, and Eric Xing. Negational symmetry of quantum neural networks for binary pattern classification. *Pattern Recognition*, 129: 108750, 2022. ISSN 0031-3203. doi: https://doi.org/10.1016/j.patcog.2022.108750. URL https://www.sciencedirect.com/science/article/pii/S003132032200231X.

[9] John Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, aug 2018. doi: 10.22331/q-2018-08-06-79. URL https://doi.org/10.22331%2Fq-2018-08-06-79.

[10] Eric B. Jones, Logan E. Hillberry, Matthew T. Jones, Mina Fasihi, Pedram Roushan, Zhang Jiang, Alan Ho, Charles Neill, Eric Ostby, Peter Graf, Eliot Kapit, and Lincoln D. Carr. Small-world complex network generation on a digital quantum processor, 2021. URL https://arxiv.org/abs/2111.00167.

[11] Logan E Hillberry, Matthew T Jones, David L Vargas, Patrick Rall, Nicole Yunger Halpern, Ning Bao, Simone Notarnicola, Simone Montangero, and Lincoln D Carr. Entangled quantum cellular automata, physical complexity, and goldilocks rules. *Quantum Science and Technology*, 6(4):045017, sep 2021. doi: 10.1088/2058-9565/ac1c41. URL https://doi.org/10.1088%2F2058-9565%2Fac1c41.

[12] Robert Fuller Fredrick Byron. *Mathematics of Classical and Quantum Physics*. Dover, 1992. ISBN 048667164X.

[13] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Andreas Bengtsson, Sergio Boixo, Michael Broughton, Bob B. Buckley, David A. Buell, Brian Burkett, Nicholas Bushnell, Yu Chen, Zijun Chen, Yu-An Chen, Ben Chiaro, Roberto Collins, Stephen J. Cotton, William Courtney, Sean Demura, Alan Derk, Andrew Dunsworth, Daniel Eppens, Thomas Eckl, Catherine Erickson, Edward Farhi, Austin Fowler, Brooks Foxen, Craig Gidney, Marissa Giustina, Rob Graff, Jonathan A. Gross, Steve Habegger, Matthew P. Harrigan, Alan Ho, Sabrina Hong, Trent Huang, William Huggins, Lev B. Ioffe, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Cody Jones, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Seon Kim, Paul V. Klimov, Alexander N. Korotkov, Fedor Kostritsa, David Landhuis, Pavel Laptev, Mike Lindmark, Erik Lucero, Michael Marthaler, Orion Martin, John M. Martinis, Anika Marusczyk, Sam McArdle, Jarrod R. McClean, Trevor McCourt, Matt McEwen, Anthony Megrant, Carlos Mejuto-Zaera, Xiao Mi, Masoud Mohseni, Wojciech Mruczkiewicz, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Hartmut Neven, Michael Newman, Murphy Yuezhen Niu, Thomas E. O'Brien, Eric Ostby, Bálint Pató, Andre Petukhov, Harald Putterman, Chris Quintana, Jan-Michael Reiner, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Doug Strain, Kevin J. Sung, Peter Schmitteckert, Marco Szalay, Norm M. Tubman, Amit Vainsencher, Theodore White, Nicolas Vogt, Z. Jamie Yao, Ping Yeh, Adam Zalcman, and Sebastian Zanker. Observation of separated dynamics of charge and spin in the fermi-hubbard model, 2020. URL https://arxiv.org/abs/2010.07965.

[14] Mark Newman. *Networks*, volume 2. Oxford University Press, 2018. ISBN 9780198805090.

[15] E. Kapit, P. Roushan, C. Neill, S. Boixo, and V. Smelyanskiy. Entanglement and complexity of interacting qubits subject to asymmetric noise. *Physical Review Research*, 2(4), oct 2020. doi: 10.1103/physrevresearch.2.043042. URL https://doi.org/10.1103%2Fphysrevresearch.2.043042.

[16] Sergey Bravyi, Alexander Kliesch, Robert Koenig, and Eugene Tang. Obstacles to variational quantum optimization from symmetry protection. *Physical review letters*, 125 (26):260505–260505, 2020. ISSN 0031-9007.