Catalytic Embeddings of Quantum Circuits

Matthew Amy

School of Computing Science, Simon Fraser University, Burnaby, BC V5A 1S6, Canada

Matthew Crawford

Department of Mathematics, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, USA

Andrew N. Glaudell

Photonic Inc., Coquitlam, BC V3K 6T1, Canada
Department of Mathematical Sciences, George Mason University, Fairfax, VA 22030, USA
Quantum Science and Engineering Center, George Mason University, Fairfax, VA 22030, USA and
Booz Allen Hamilton, Washington, DC 20005, USA

Melissa L. Macasieb

Booz Allen Hamilton, Washington, DC 20005, USA

Samuel S. Mendelson

Strategic & Computing Systems Department, Naval Surface Warfare Center, Dahlgren Division, Dahlgren, VA 22448, USA

Neil J. Ross

Department of Mathematics and Statistics, Dalhousie University, Halifax, NS B3H 4R2, Canada

If a set $\mathbb{G}$ of quantum gates is countable, then the operators that can be exactly represented by a circuit over $\mathbb{G}$ form a strict subset of the collection of all unitary operators. When $\mathbb{G}$ is universal, one circumvents this limitation by resorting to repeated gate approximations: every occurrence of a gate which cannot be exactly represented over $\mathbb{G}$ is replaced by an approximating circuit. Here, we introduce catalytic embeddings, which provide an alternative to repeated gate approximations. With catalytic embeddings, approximations are relegated to the preparation of a fixed number of reusable resource states called catalysts. Because the catalysts only need to be prepared once, when catalytic embeddings exist they always produce shorter circuits, in the limit of increasing gate count and target precision. In the present paper, we lay the foundations of the theory of catalytic embeddings and we establish several of their structural properties. In addition, we provide methods to design catalytic embeddings, showing that their construction can be reduced to that of a single fixed matrix when the gates involved have entries in well-behaved rings of algebraic numbers. Finally, we showcase some concrete examples and applications. Notably, we show that catalytic embeddings generalize a technique previously used to implement the Quantum Fourier Transform over the Clifford+T gate set with $O(n)$ gate approximations.
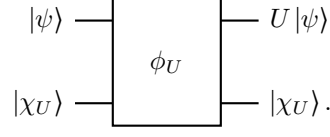
I. INTRODUCTION

Certain gate sets in quantum computing have become particularly distinguished due to their rich mathematical structure, prominence in proposals for scalable quantum computing, and practical convenience. To perform tasks reliably on a quantum computer, it is often important to understand the *expressivity* of a given gate set $\mathbb{G}$. This is done by characterizing the unitary operators which can be exactly implemented as a circuit over $\mathbb{G}$. In the context of fault-tolerant quantum computation, such characterizations have been possible through constructive number-theoretic methods [1–4]. These characterizations of expressivity enabled major progress in the theory of quantum circuits, including exact synthesis algorithms for circuits on one or more qubits [1–6], powerful circuit optimization strategies [7], and optimally efficient approximation methods [8, 9]. Evidenced by this growing body of work, number-theoretic characterizations of expressivity have come to play a central role in a large number of practical frameworks and strategies for quantum compilation.

Number-theoretic characterizations are particularly crucial in the development of efficient approximation methods. Gate sets that are well-suited for fault-tolerant quantum computation are finite and, as a result, can only express a countable set of distinct operators. This is in contrast to arbitrary unitary evolution, where operators may come from an uncountable continuum. In order to perform universal quantum computation, one therefore needs to rely on approximations to extend the expressivity of a gate set: if an operator U cannot be implemented exactly over the gate set $\mathbb{G}$, then one looks for a circuit V over $\mathbb{G}$ which approximates U for a desired norm and precision. A number-theoretic characterization of the gate set simplifies this process by allowing the operator U to be approximated by a matrix with entries in the number ring $\mathcal{R}$ associated with $\mathbb{G}$. In practice, this approximation process is then repeated
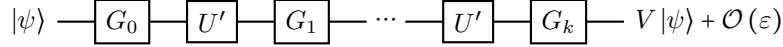
as often as the operator U is needed.

Here, we revisit the problem of extending the expressivity of a gate set. We introduce the method of *catalytic embeddings* which provides an alternative to repeated gate approximations in a number of important contexts. For a unitary operation U and a quantum gate set $\mathbb{G}$, a catalytic embedding of U over $\mathbb{G}$ is given by a circuit ϕ_U over $\mathbb{G}$ and a quantum state $|\chi_U\rangle$ called a *catalyst* such that, for all quantum states $|\psi\rangle$ of appropriate dimension, we have

The power of such a construction becomes evident in the case where U does not normally have an exact circuit representation over $\mathbb{G}$. Suppose we wish to implement the composite unitary operation $V = G_k U G_{k-1} U \cdots G_1 U G_0$ to a target precision ε, where each G_j is a unitary with an exact implementation over the gate set $\mathbb{G}$. If V acts on more than a few qubits and there are no additional obvious circuit optimizations, we are left with few options besides replacing each U independently.
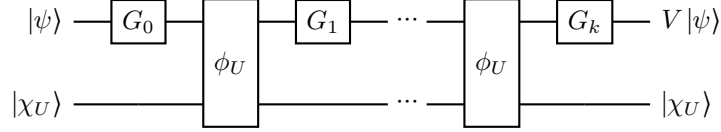
The standard approach to this problem is to use repeated approximations. The operator U is approximated by some circuit U' over $\mathbb{G}$ up to ε/k, which then guarantees that the circuit below implements V to the desired precision.
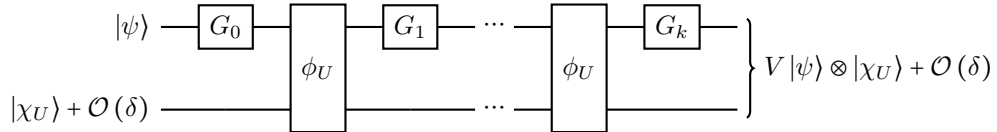

Using the Solovay-Kitaev algorithm, or an improved method if it is available for $\mathbb{G}$, we can find such a U' with $\mathcal{O}\left(\log^c(k/\varepsilon)\right)$ gates in $\mathbb{G}$, where c is a constant that depends on $\mathbb{G}$ and the chosen approximation technique. Writing T for the gate count of the $k+1$ circuits G_j after exact synthesis, the total gate count of the entire circuit is given by

$$T + k \cdot \mathcal{O}\left(\log^c(k/\varepsilon)\right) = T + \mathcal{O}\left(k\log^c(k/\varepsilon)\right).$$

Now suppose that ϕ_U and $|\chi_U\rangle$ define a catalytic embedding for U over $\mathbb{G}$. Then, for any $|\psi\rangle$, we have:

That is, in the presence of the catalyst $|\chi_U\rangle$, the above circuit implements V on its top register. Importantly, no elements in the circuit are dependent on ε. In practice, we cannot assume to have direct access to the catalyst $|\chi_U\rangle$, but we can instead consider the action of this circuit on an approximate catalyst $|\chi\rangle + \mathcal{O}(\delta)$. By linearity, we then have:

Thus, we get the desired action if we take $\delta = \varepsilon$, and so we must approximate $|\chi_U\rangle$ to precision ε. Writing N for the gate count of the fixed circuit ϕ_U and, as before, T for the gate count of the $k+1$ circuits G_j, and applying the Solovay-Kitaev algorithm (or an alternative) to approximately prepare $|\chi_U\rangle$, we find the total gate count to be

$$T + Nk + \mathcal{O}\left(\log^c(1/\varepsilon)\right) = T + \mathcal{O}\left(k + \log^c(1/\varepsilon)\right).$$

Thus, given a circuit where k approximations need to be made and an overall precision of ε is desired, catalytic embeddings will incur the asymptotic cost savings

$$\mathcal{O}\left(k\log^c(k/\varepsilon)\right) \to \mathcal{O}\left(k + \log^c(1/\varepsilon)\right).$$

In other words, when catalytic embeddings exist, they will always outperform repeated gate approximations, as k and $1/\varepsilon$ increase.

In the present paper, we lay the foundations of the theory of catalytic embeddings. We define a very general notion of catalytic embedding and study its abstract properties. We then progressively endow catalytic embeddings with

additional structure, culminating in the notion of *linear catalytic embedding*. This provides an axiomatization of a very general type of catalytic embedding that is well-suited for gate sets admitting a number-theoretic characterization, such as the Clifford+T, Clifford+CS, or Toffoli-Hadamard gate sets. We then introduce a method to design linear catalytic embeddings and we show that their construction can be reduced to that of a single fixed matrix when the gates involved have entries in well-behaved rings of algebraic numbers. Throughout, we showcase concrete examples and applications.

Because catalytic embeddings have clear connections to algebraic number theory, they provide a convenient framework through which to study and compare gate sets that admit a number-theoretic characterization. In the present paper, we define the framework of catalytic embeddings and study its applications to the synthesis of quantum circuits. We believe, however, that catalytic embeddings may find applications in other areas of the theory of quantum circuits, such as the study of the Clifford hierarchy [10] or the classification of certain subgroups of the unitary group [11, 12].

The remainder of the paper is organized as follows. In Section II, we define notation and provide the reader with a very brief introduction to topics from algebraic number theory and the theory of quantum computation. In Section III, we formally define catalytic embeddings. Afterwards, we introduce linear catalytic embeddings in Section IV and study their properties in Section V. We introduce a particularly convenient method for constructing linear catalytic embeddings of unitary matrices in Section VI. Finally, we provide some concluding remarks and outline avenues for future work in Section VII.

II. PRELIMINARIES

We begin by introducing some notation and terminology. We assume that the reader has some familiarity with algebraic number theory and the theory of quantum computation. Further details on these topics can be found in [13, 14] and [15], respectively.

A. Kroneckerian Number Rings

Recall that a *number field* is a finite degree field extension of the field $\mathbb{Q}$ of rational numbers. An injective field homomorphism from a number field $\mathcal{K}$ into the field $\mathbb{C}$ of complex numbers is an *embedding* of $\mathcal{K}$ into $\mathbb{C}$.

We will be interested in the number fields which can be endowed with an unambiguous notion of complex conjugation. These are precisely the *Kroneckerian* number fields.

Definition 1 (Kroneckerian Number Field). A number field $\mathcal{K}$ is *Kroneckerian* if there exists an involution $c : \mathcal{K} \to \mathcal{K}$ such that, for any embedding $\sigma : \mathcal{K} \to \mathbb{C}$, the following diagram commutes:

$$
\begin{array}{ccc}
\mathcal{K} & \xrightarrow{\ c\ } & \mathcal{K} \\
\sigma \downarrow & & \downarrow \sigma \\
\mathbb{C} & \xrightarrow{\ (.)^{\dagger}\ } & \mathbb{C}
\end{array}
$$

The involution c in Definition 1 is uniquely determined by the number field $\mathcal{K}$ and can be interpreted as the analogue of complex conjugation in $\mathcal{K}$. For this reason, and by a slight abuse of notation, we use $(.)^{\dagger}$ to denote c in what follows.

Remark 1. Let $\mathcal{K}$ be a number field. Then $\mathcal{K}$ is said to be *totally real* if $\sigma[\mathcal{K}] \subseteq \mathbb{R}$ for every embedding σ of $\mathcal{K}$, and *totally imaginary* if $\sigma[\mathcal{K}] \nsubseteq \mathbb{R}$ for every embedding σ of $\mathcal{K}$. Moreover, $\mathcal{K}$ is said to be *CM* if it is a degree two totally imaginary extension of a totally real field. Kroneckerian number fields can also be defined as number fields that are either totally real or CM. The latter definition of Kroneckerian number fields is the standard one in the literature (see, e.g., [14]) and is equivalent to Definition 1. We choose to define Kroneckerian number fields as in Definition 1 because this definition emphasizes the fact that Kroneckerian number fields are those in which complex conjugation can be unambiguously defined.

We will not only be interested in number fields but also in *number rings*, which we take to be subrings of number fields (viewed as rings). Recall that the field of fractions of a number ring is the smallest field (up to isomorphism) containing that ring. We introduce a notion of *Kroneckerian number ring* to deal with complex conjugation in number rings.

Definition 2 (Kroneckerian Number Ring)**.** A number ring $\mathcal{R}$ is *Kroneckerian* if its field of fractions $\text{Frac}\,(\mathcal{R})$ is a Kroneckerian field and if $r^\dagger \in \mathcal{R}$ for every $r \in \mathcal{R}$.

An injective ring homomorphism from a number ring $\mathcal{R}$ into the ring $\mathbb{C}$ of complex numbers is an *embedding* of $\mathcal{R}$ into $\mathbb{C}$.

We assume in the remainder of this paper that all number rings and number fields exist in some ambient field $\mathcal{L}$, which is not required to be Kroneckerian. For concreteness, one can consider all of the rings and fields used in this paper as existing in the field of algebraic numbers.

B. Matrices and Unitary Groups

The addition, multiplication, direct sum, and tensor (or Kronecker) product of matrices are defined as usual. If $\mathcal{R}$ is a ring, we write $\mathcal{M}_n\,(\mathcal{R})$ for the associative $\mathcal{R}$-algebra of $n \times n$ matrices with entries in $\mathcal{R}$ and $\mathcal{M}\,(\mathcal{R})$ for the collection of all square matrices with entries in $\mathcal{R}$. That is,

$$\mathcal{M}\,(\mathcal{R}) = \bigcup_{n>0} \mathcal{M}_n\,(\mathcal{R}).$$

If $\mathcal{R}$ is a Kroneckerian number ring, we can extend conjugation componentwise from $\mathcal{R}$ to $\mathcal{M}\,(\mathcal{R})$. We then write $U^\dagger$ for the conjugate transpose (or adjoint) of $U \in \mathcal{M}\,(\mathcal{R})$ and we say that U is *Hermitian* if $U = U^\dagger$, *unitary* if $U^{-1} = U^\dagger$, and *normal* if $UU^\dagger = U^\dagger U$.

For a fixed positive integer n, the collection of $n \times n$ unitary matrices over a Kroneckerian number ring $\mathcal{R}$ forms a group known as the *unitary group of degree n over $\mathcal{R}$*.

Definition 3 (Unitary Group)**.** Let $\mathcal{R}$ be a Kroneckerian number ring. The *unitary group of degree n over $\mathcal{R}$*, denoted $\mathcal{U}_n\,(\mathcal{R})$, is the group whose elements are n-dimensional unitary matrices over $\mathcal{R}$. That is,

$$\mathcal{U}_n\,(\mathcal{R}) = \left\{\, U \in \mathcal{M}_n\,(\mathcal{R}) \mid U^\dagger U = UU^\dagger = I \,\right\}.$$

We write $\mathcal{U}\,(\mathcal{R})$ for the collection of all unitary matrices over $\mathcal{R}$. That is,

$$\mathcal{U}\,(\mathcal{R}) = \bigcup_{n>0} \mathcal{U}_n\,(\mathcal{R}) = \left\{\, U \in \mathcal{M}\,(\mathcal{R}) \mid U^\dagger U = UU^\dagger = I \,\right\}.$$

Definition 3 is the reason for our interest in Kroneckerian number rings. Indeed, because conjugation is well-defined in Kroneckerian number rings, one can meaningfully talk about unitary matrices and unitary groups over these rings.

Note that any embedding $\rho : \mathcal{R} \to \mathbb{C}$ extends componentwise to a group homomorphism $\rho : \mathcal{U}_n\,(\mathcal{R}) \to \mathcal{U}_n\,(\mathbb{C})$ satisfying $\rho(\mathcal{U}_n\,(\mathcal{R})) = \mathcal{U}_n\,(\rho(\mathcal{R}))$. Note moreover that $\mathcal{U}\,(\mathcal{R})$ is closed under direct sum and tensor product. In other words, if $U, V \in \mathcal{U}\,(\mathcal{R})$, then $U \oplus V, U \otimes V \in \mathcal{U}\,(\mathcal{R})$.

Remark 2. The collection $\mathcal{U}\,(\mathcal{R})$ is equipped with a partially defined operation of composition (the multiplication of matrices), a tensor product, and a *dagger* (the conjugate transpose). These operations can be shown to endow $\mathcal{U}\,(\mathcal{R})$ with the structure of a *strict dagger symmetric monoidal groupoid* [16–18].

C. Quantum States and Quantum Evolution

Quantum states are described by vectors in a Hilbert space. Throughout the paper, we denote the Hilbert space of dimension m by $\mathcal{H}_m$, assuming that the underlying field is $\mathbb{C}$ unless otherwise stated.

We assume the existence of a preferred orthonormal basis called the *computational basis* and we represent the computational basis for $\mathcal{H}_m$ by the collection of quantum states $\mathcal{B}_m = \{\, |a\rangle \mid a \in \{\, 0, 1, \ldots, m-1 \,\} \,\}$. We further assume that the computational basis for the Hilbert space $\mathcal{H}_m \otimes \mathcal{H}_n \cong \mathcal{H}_{mn}$ is given by

$$\{\, |a\rangle \otimes |b\rangle = |an + b\rangle \mid |a\rangle \in \mathcal{B}_m, |b\rangle \in \mathcal{B}_n \,\}.$$

The evolution of a quantum state is described by a norm-preserving linear transformation on a Hilbert space. Thus, in any orthonormal basis, and in particular in the computational basis, this evolution can be represented by a unitary matrix. A simple but ubiquitous example of this kind of evolution is the transformation $\text{swap}(m, n) : \mathcal{H}_m \otimes \mathcal{H}_n \to \mathcal{H}_n \otimes \mathcal{H}_m$ whose action on computational basis elements $|a\rangle \in \mathcal{B}_m$ and $|b\rangle \in \mathcal{B}_n$ is given by:

$$\text{swap}(m, n) : |a\rangle \otimes |b\rangle \mapsto |b\rangle \otimes |a\rangle.$$

D. Circuits

We now introduce *circuits*, which are a widespread notation for building complex linear operators in quantum information (see, e.g., [15]). Our presentation differs slightly from the standard one and, in particular, emphasizes the distinction between a circuit and the matrix it represents.

Definition 4 (Gates). A *gate* is a symbol. Every gate G is associated with a positive integer n, called the *dimension* of G, and an n-dimensional unitary matrix eval(G), called the *evaluation* of G. A *gate set* is a set of gates.

Definition 5 (Circuits). Let $\mathbb{G}$ be a gate set. *Circuits over* $\mathbb{G}$ and their *dimension* are defined inductively as follows.

- For every $n \in \mathbb{N}^*$, I_n is a circuit over $\mathbb{G}$ of dimension n.

- For every $m, n \in \mathbb{N}^*$, swap(m, n) is a circuit over $\mathbb{G}$ of dimension mn.

- If G is a gate in $\mathbb{G}$ of dimension n, then G is a circuit over $\mathbb{G}$ of dimension n.

- If C and D are circuits over $\mathbb{G}$ of dimension n, then $(C \circ D)$ is a circuit over $\mathbb{G}$ of dimension n.

- If C and D are circuits over $\mathbb{G}$ of dimensions m and n, respectively, then $(C \otimes D)$ is a circuit over $\mathbb{G}$ of dimension mn.

We write $\mathcal{C}(\mathbb{G})$ for the collection of circuits over $\mathbb{G}$ and $\mathcal{C}_n(\mathbb{G})$ for the circuits over $\mathbb{G}$ whose dimension is n.

In Definition 5, I_n, swap(m, n), $\circ$, and $\otimes$ are treated as uninterpreted symbols. From this perspective, a circuit over a gate set $\mathbb{G}$ is nothing but a word over the alphabet

$$\mathbb{G} \cup \{\, I_n \mid n \in \mathbb{N}^* \,\} \cup \{\, \mathrm{swap}(m, n) \mid m, n \in \mathbb{N}^* \,\} \cup \{\, \circ, \otimes, (,) \,\}$$

satisfying the constraints stated in Definition 5. For example, if $\mathbb{G} = \{\, A, B, C \,\}$ and the gates A, B, and C, have dimension 1, 2, and 4, respectively, then $((I_2 \otimes ((A \otimes I_2) \circ B)) \circ C)$ is a circuit over $\mathbb{G}$ of dimension 4. In what follows, we always omit the outermost bracket around a circuit. That is, we write $(I_2 \otimes ((A \otimes I_2) \circ B)) \circ C$ instead of $((I_2 \otimes ((A \otimes I_2) \circ B)) \circ C)$.

By Definition 4, every gate in a gate set comes with an evaluation. We now extend this evaluation from gates to circuits containing these gates.

Definition 6 (Evaluation). Let $\mathbb{G}$ be a gate set. The *evaluation* function $e : \mathcal{C}(\mathbb{G}) \to \mathcal{U}(\mathbb{C})$ is defined inductively as follows.

- For every $n \in \mathbb{N}$, $e(I_n) = I_n$.

- For every $m, n \in \mathbb{N}$, $e(\mathrm{swap}(m, n)) = \mathrm{swap}(m, n)$.

- If G is a gate in $\mathbb{G}$, then $e(G) = \mathrm{eval}(G)$.

- If C and D are circuits over $\mathbb{G}$, then $e(C \circ D) = e(C) \circ e(D)$.

- If C and D are circuits over $\mathbb{G}$, then $e(C \otimes D) = e(C) \otimes e(D)$.

Definitions 4, 5 and 6 are meant to carefully distinguish between a circuit (which is a word over some alphabet) and the unitary matrix it denotes (its evaluation). Note that while Definition 5 treated I_n, swap(m, n), $\circ$, and $\otimes$ as uninterpreted symbols, Definition 6 fixes the evaluation of these symbols to be natural one: I_n is evaluated as the identity matrix of dimension n, etc.

A common alternative to Definition 5 is to denote circuits using diagrams, rather than words. In diagrammatic notation, the circuit I_n is represented by a wire labelled by n on both sides, and the circuit swap(m, n) is represented by two crossing wires labelled m and n:

$$n \;\text{———}\; n \qquad \text{and} \qquad \begin{matrix} m \;\text{——} \\ n \;\text{——} \end{matrix}\!\!\bowtie\!\!\begin{matrix} \text{——}\; n \\ \text{——}\; m. \end{matrix}$$
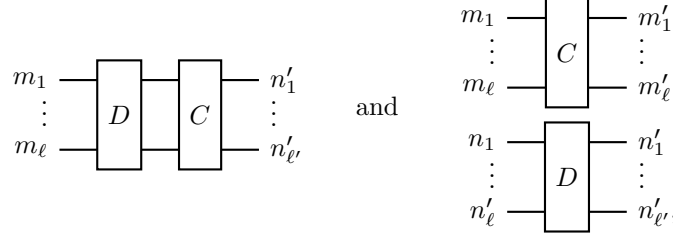
For every $G \in \mathbb{G}$ of dimension n, the circuit G is represented by a box:

$$n \;\text{—}\; \boxed{G} \;\text{—}\; n.$$

Finally, if C and D are circuits respectively represented as

$$
\begin{array}{c}
m_1 \\ \vdots \\ m_\ell
\end{array}
\boxed{C}
\begin{array}{c}
m'_1 \\ \vdots \\ m'_\ell
\end{array}
\qquad \text{and} \qquad
\begin{array}{c}
n_1 \\ \vdots \\ n_{\ell'}
\end{array}
\boxed{D}
\begin{array}{c}
n'_1 \\ \vdots \\ n'_{\ell'},
\end{array}
$$

then the circuits $C \circ D$ and $C \otimes D$ are respectively represented as

$$
\begin{array}{c}
m_1 \\ \vdots \\ m_\ell
\end{array}
\boxed{D}\boxed{C}
\begin{array}{c}
n'_1 \\ \vdots \\ n'_{\ell'}
\end{array}
\qquad \text{and} \qquad
\begin{array}{c}
m_1 \\ \vdots \\ m_\ell \\[4pt] n_1 \\ \vdots \\ n_\ell
\end{array}
\begin{array}{c}
\boxed{C} \\[6pt] \boxed{D}
\end{array}
\begin{array}{c}
m'_1 \\ \vdots \\ m'_\ell \\[4pt] n'_1 \\ \vdots \\ n'_{\ell'},
\end{array}
$$

where $C \circ D$ is well-defined if and only if $\ell = \ell'$ and $m'_j = n_j$ for every $1 \le j \le \ell$. We omit the wire labels if they are unimportant or can be inferred from context.
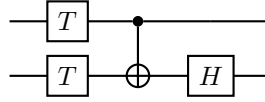
Example 1. The $\{X, CX, CCX\}$ gate set consists of the gates X, CX, and CCX, where

$$
e(X) = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad e(CX) = \mathrm{diag}(I_2, e(X)) \quad \text{and} \quad e(CCX) = \mathrm{diag}(I_4, e(CX)).
$$

Example 2. The *Clifford+T* gate set consists of the gates H, T, and CX, where $X := H \circ (T \circ (T \circ (T \circ (T \circ H))))$ and

$$
e(H) = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \qquad e(T) = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}, \qquad \text{and} \qquad e(CX) = \mathrm{diag}(I_2, e(X))
$$

The Clifford+T circuit $C = ((I_2 \otimes H) \circ CX) \circ (T \otimes T)$ can be diagrammatically represented as follows (using the standard convention for the CX gate)

and the circuit C evaluates to

$$
e(C) = \frac{1}{\sqrt{2}} \begin{bmatrix}
1 & e^{i\pi/4} & 0 & 0 \\
1 & -e^{i\pi/4} & 0 & 0 \\
0 & 0 & e^{i\pi/4} & e^{i\pi/2} \\
0 & 0 & -e^{i\pi/4} & e^{i\pi/2}
\end{bmatrix}.
$$

Remark 3. The notion of circuit introduced in this section differs from the usual one (as defined, say, in [15]). Firstly, the circuits introduced here can act (when evaluated) on Hilbert spaces of arbitrary dimensions, whereas circuits in the quantum computing literature are often restricted to spaces of dimension 2^n for some nonnegative integer n. The reader familiar with quantum circuits can think of the circuits described above as acting not only on qubits but, more generally, on a mixture of qudits of varying (finite) dimensions. Secondly, circuits are traditionally considered up to certain transformations. For example, according to Definitions 5 and 6, if A and B are two gates of dimension 2, then $A \otimes B$, $(A \otimes I_2) \circ (I_2 \otimes B)$, and $(I_2 \otimes B) \circ (A \otimes I_2)$ are three distinct circuits representing the same matrix. In contrast, circuits in the literature are often defined in a way that equates these three circuits.

Definition 7. Let $\mathscr{C}$ be a collection of circuits. We define $\mathcal{U}(\mathscr{C})$ as the image of $\mathscr{C}$ under the evaluation function e. That is, $\mathcal{U}(\mathscr{C}) = e[\mathscr{C}]$. We further define $\mathcal{U}_n(\mathscr{C})$ as $\mathcal{U}_n(\mathscr{C}) = \mathcal{U}(\mathscr{C}) \cap \mathcal{U}_n(\mathbb{C})$.

The set $\mathcal{U}(\mathcal{C}(\mathbb{G}))$ is the collection of all unitary matrices that can be represented by a circuit over $\mathbb{G}$. If $\mathscr{U}$ is a collection of unitary matrices, then we sometimes interpret $\mathscr{U}$ as a gate set by introducing, for every $U \in \mathscr{U}$ of dimension n, a gate G_U of dimension n with $e(G_U) = U$. In such a case, we then have $\mathscr{U} \subseteq \mathcal{U}(\mathcal{C}(\mathscr{U}))$. If $\mathscr{U}$ is closed under composition and tensor products, and contains all identities and swaps, then we in fact have $\mathscr{U} = \mathcal{U}(\mathcal{C}(\mathscr{U}))$.

By definition, the evaluation function $e : \mathscr{C} \to \mathcal{U}(\mathscr{C})$ is surjective. It is not injective, however, since many different circuits evaluate to the same unitary matrix. An *exact synthesis function for* $\mathscr{C}$ is a function which assigns, to each unitary matrix in $\mathcal{U}(\mathscr{C})$, a unique circuit representing that matrix.

Definition 8 (Exact Synthesis). An *exact synthesis function* for a collection of circuits $\mathscr{C}$ is a function $s : \mathcal{U}(\mathscr{C}) \to \mathscr{C}$ such that $e \circ s = I_{\mathcal{U}(\mathscr{C})}$. An algorithm computing an exact synthesis function is an *exact synthesis algorithm*.

The notion of exact synthesis introduced in Definition 8 is sometimes known as *ancilla-free* exact synthesis.

Example 3. Several exact synthesis algorithms have been introduced in the literature. In particular, exact synthesis algorithms exist for the gate sets $\{X, CX, CCX\}$ and $\{H, T, CX\}$ discussed in Examples 1 and 2 (see [2, 19]). In both cases, the exact synthesis algorithm relies on (and in fact establishes) a characterization of the unitary matrices that can be represented over the gate set. In particular, for $n \geq 4$,

- the elements of $\mathcal{U}_{2^n}(\mathcal{C}(\{X, CX, CCX\}))$ are exactly the permutations matrices of dimension 2^n that have determinant 1 (i.e., $\mathcal{U}_{2^n}(\mathcal{C}(\{X, CX, CCX\})) = A_{2^n}$), and

- the elements of $\mathcal{U}_{2^n}(\mathcal{C}(\{H, T, CX\}))$ are exactly the elements of $\mathcal{U}_{2^n}(\mathbb{Z}[1/2, e^{2\pi i/8}])$ that have determinant 1.

The exact synthesis functions of [2, 19], and the accompanying exact synthesis algorithms, can be adapted to $n < 4$ by varying the condition on the determinant. In fact, as we will discuss below, the conditions on the determinant can be lifted altogether through the use of ancillas.

III. CATALYTIC EMBEDDINGS

Recall from Section I that we are interested in the problem of constructing a circuit that, in the presence of a well-chosen resource state, implements a desired operator. To make this intuition precise, we now introduce *catalytic embeddings*.

Definition 9. Let C be an n-dimensional circuit and let $\mathscr{C}$ be a collection of circuits. A k-*dimensional catalytic embedding of C in $\mathscr{C}$* is a pair (ϕ, Π) consisting of an (nk)-dimensional circuit $\phi \in \mathscr{C}$ and a nonzero orthogonal projector Π on $\mathcal{H}_k$ satisfying the following *catalytic condition*:

$$e(\phi)(I \otimes \Pi) = e(C) \otimes \Pi.$$

If (ϕ, Π) is a k-dimensional catalytic embedding of a circuit C, and $|\chi\rangle$ is such that $\Pi|\chi\rangle = |\chi\rangle$, then the catalytic condition ensures that for any $|\psi\rangle$ we have

$$e(\phi)|\psi\rangle|\chi\rangle = e(\phi)(I \otimes \Pi)|\psi\rangle|\chi\rangle = (e(C) \otimes \Pi)|\psi\rangle|\chi\rangle = (e(C)|\psi\rangle)(\Pi|\chi\rangle) = (e(C)|\psi\rangle)|\chi\rangle. \tag{1}$$

The circuit ϕ therefore acts as C in the presence of $|\chi\rangle$. Moreover, $|\chi\rangle$ remains unchanged by the application of ϕ. For these reasons, and in accordance with related work [20], we refer to any quantum state $|\chi\rangle \in \mathcal{H}_k$ with $\Pi|\chi\rangle = |\chi\rangle$ as a *catalyst for C over $\mathscr{C}$*. Moreover, if (ϕ, Π) is an embedding of a circuit C, we sometimes refer to ϕ as the *embedding* and to Π as the *catalytic projector*. Note that a catalytic embedding can be constructed from a catalyst. Indeed, if ϕ and $|\chi\rangle$ jointly satisfy Equation (1) and $|\chi\rangle$ is a unit vector, then $(\phi, |\chi\rangle\langle\chi|)$ is a catalytic embedding of C.

As the propositions below show, the (evaluation of the) embedding of a circuit C can always be written as a block-diagonal matrix, up to a change of basis.

Proposition 1. *Let C be a circuit and $\mathscr{C}$ be a collection of circuits. If (ϕ, Π) is a catalytic embedding of C in $\mathscr{C}$, then $(I \otimes \Pi)e(\phi) = e(C) \otimes \Pi$.*

Proof. Note that $(e(C) \otimes I)(I \otimes \Pi) = e(C) \otimes \Pi = (I \otimes \Pi)(e(C) \otimes I)$. The catalytic condition therefore yields

$$e(\phi)(I \otimes \Pi)(e(C)^\dagger \otimes I) = I \otimes \Pi.$$

Since Π is orthogonal, applying $(.)^\dagger$ on both sides of the above equation gives $(e(C) \otimes I)(I \otimes \Pi)e(\phi)^\dagger = I \otimes \Pi$, which then implies the desired equation by left-multiplication with $e(\phi)$. $\square$

Proposition 2. *Let C be a circuit and $\mathscr{C}$ be a collection of circuits. If (ϕ, Π) is a catalytic embedding of C in $\mathscr{C}$, then there exists a unitary U such that*

$$e(\phi) \sim \underbrace{e(C) \oplus e(C) \oplus \cdots \oplus e(C)}_{\mathrm{rank}(\Pi)} \oplus U,$$

where $\sim$ denotes equality up to conjugation by a unitary.

Proof. Suppose that C is an n-dimensional circuit and that (ϕ, Π) is a k-dimensional catalytic embedding, so that $e(\phi)$ acts on $\mathcal{H} = \mathcal{H}_{nk}$. Consider the projectors $P = I \otimes \Pi$ and $P^\perp = I - P$. Note that $PP^\perp = P^\perp P = 0$. Note moreover that, by the catalytic condition and Proposition 1, $Pe(\phi) = e(\phi)P$. We can therefore decompose the action of $e(\phi)$ on $\mathcal{H}$ into its action on the subspaces associated with P and $P^\perp$ as follows:

$$
\begin{aligned}
e(\phi) &= (P + P^\perp)e(\phi)(P + P^\perp) \\
&= Pe(\phi)P + Pe(\phi)P^\perp + P^\perp e(\phi)P + P^\perp e(\phi)P^\perp \\
&= Pe(\phi)P + e(\phi)PP^\perp + P^\perp Pe(\phi) + P^\perp e(\phi)P^\perp \\
&= Pe(\phi)P + P^\perp e(\phi)P^\perp,
\end{aligned}
$$

Now let r be the rank of Π and let $Q = \operatorname{diag}(I_r, 0_{k-r})$. By the spectral theorem, there exists a unitary matrix V such that $V \Pi V^\dagger = Q$. Let $W = (V \otimes I_n) \circ \operatorname{swap}(n, k)$. Then W is unitary and

$$
We(\phi)W^\dagger = WPe(\phi)PW^\dagger + WP^\perp e(\phi)P^\perp W^\dagger = Q \otimes e(C) + [(I - Q) \otimes I]We(\phi)W^\dagger[(I - Q) \otimes I],
$$

since $Pe(\phi)P = e(\phi)P = e(C) \otimes \Pi$. Thus $We(\phi)W^\dagger = \operatorname{diag}(I_r \otimes e(C), U)$ for some matrix U. Since $e(\phi)$, $I_r \otimes e(C)$, and W are unitary, U must also be unitary. $\qquad\square$

Rather than focus on the embedding of a single circuit, we will focus on the embedding of collections of circuits.

Definition 10. Let $\mathscr{C}$ and $\mathscr{C}'$ be two collections of circuits. A *catalytic embedding of $\mathscr{C}$ in $\mathscr{C}'$* is a collection $\{\,(\phi_C, \Pi_C) \mid C \in \mathscr{C}\,\}$, where, for each $C \in \mathscr{C}$, the pair (ϕ_C, Π_C) is a catalytic embedding of C in $\mathscr{C}'$.

A catalytic embedding $\{\,(\phi_C, \Pi_C) \mid C \in \mathscr{C}\,\}$ as in Definition 10 implicitly specifies two functions. The function $\phi : C \mapsto \phi_C$, which assigns an embedding to every circuit in $\mathscr{C}$, and the function $\Pi : C \to \Pi_C$, which assigns a catalytic projector to every circuit in $\mathscr{C}$. For brevity, we sometimes write $(\phi, \Pi) : \mathscr{C} \to \mathscr{C}'$ to refer to the catalytic embedding $\{\,(\phi_C, \Pi_C) \mid C \in \mathscr{C}\,\}$ of $\mathscr{C}$ in $\mathscr{C}'$.

Definition 11. We say that a catalytic embedding $(\phi, \Pi) : \mathscr{C} \to \mathscr{C}'$ is *homogeneous* when $\Pi_C = \Pi_D$ for every $C, D \in \mathscr{C}$.

If $(\phi, \Pi) : \mathscr{C} \to \mathscr{C}'$ is a homogeneous catalytic embedding, then there exists an integer k such that, for every $C \in \mathscr{C}$, (ϕ_C, Π_C) is a k-dimensional catalytic embedding. We then call k the *dimension* of (ϕ, Π). A catalytic embedding $(\phi, \Pi) : \mathscr{C} \to \mathscr{C}'$ can always be homogenized when $\mathscr{C}$ is finite. This can be done, for example, by replacing Π by $\bigotimes \Pi_C$ and amending ϕ appropriately.

Example 4. An exact synthesis function in the sense of Definition 8 is a 1-dimensional catalytic embedding of $\mathcal{U}(\mathscr{C})$ (viewed as a gate set) in $\mathscr{C}$. More generally, exact synthesis results that rely on ancillas, such as the most general algorithms introduced in [2, 19], can also be viewed as catalytic embeddings: [19] introduces a 2-dimensional catalytic embedding of S_{2^n} in $\mathcal{C}_{2^{n+1}}(\{\,X, CX, CCX\,\})$, and [2] introduces a 2-dimensional catalytic embedding of $\mathcal{U}_{2^n}(\mathbb{Z}[1/2, e^{2\pi i/8}])$ in $\mathcal{C}_{2^{n+1}}(\{\,H, T, CX\,\})$. In both cases, the embeddings are homogeneous. The embedding of [2] uses the projector $|0\rangle\langle 0|$, whereas that of [19] utilizes the projector I_2.
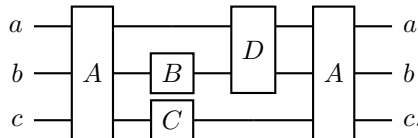
Definition 12. The *concatenation* of two catalytic embeddings $(\phi, \Pi) : \mathscr{C} \to \mathscr{C}'$ and $(\phi', \Pi') : \mathscr{C}' \to \mathscr{C}''$ is the catalytic embedding $(\phi', \Pi') \circ (\phi, \Pi) : \mathscr{C} \to \mathscr{C}''$ defined by

$$
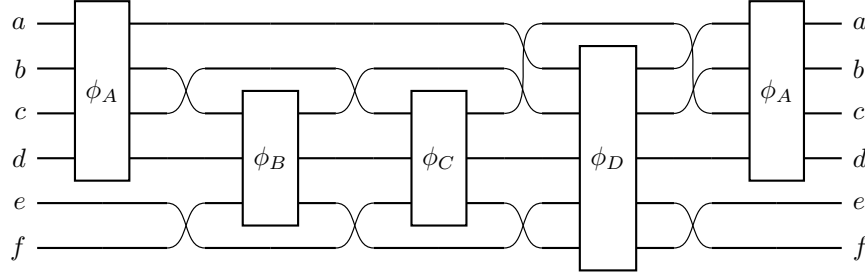(\phi', \Pi') \circ (\phi, \Pi) = (\phi' \circ \phi, \Pi \otimes \Pi').
$$

It is straightforward to verify that the concatenation of two catalytic embeddings is indeed a catalytic embedding. If (ϕ, Π) and (ϕ', Π') are two catalytic embeddings of dimension k and k', respectively, then $(\phi', \Pi') \circ (\phi, \Pi)$ is a catalytic embedding of dimension kk'. Moreover, the concatenation of embeddings is associative and preserves homogeneity.

We will be especially interested in catalytic embeddings when $\mathscr{C} = \mathcal{C}(\mathbb{G})$ and $\mathscr{C}' = \mathcal{C}(\mathbb{H})$, for some gate sets $\mathbb{G}$ and $\mathbb{H}$. One often thinks of the gates in a gate set $\mathbb{G}$ as generators for the circuits in $\mathcal{C}(\mathbb{G})$. It is therefore natural to expect that a catalytic embedding of $\mathbb{G}$ in some collection of circuits $\mathscr{C}'$ might extend to a catalytic embedding of $\mathcal{C}(\mathbb{G})$ in $\mathscr{C}'$. This is indeed the case, as the following example illustrates.

Example 5. Consider the following circuit F over some gate set $\mathbb{G} = \{\,A, B, C, D\,\}$:

Suppose that we have a catalytic embedding of $\mathbb{G}$ in $\mathscr{C}$ such that (ϕ_A, Π_1), $(\phi_B, \Pi_1 \otimes \Pi_3)$, $(\phi_C, \Pi_1 \otimes \Pi_2)$, and $(\phi_D, \Pi_1 \otimes \Pi_3 \otimes \Pi_2)$ are catalytic embeddings of A, B, C, and D in $\mathscr{C}$ with $\Pi_1 : \mathcal{H}_d \to \mathcal{H}_d$, $\Pi_2 : \mathcal{H}_e \to \mathcal{H}_e$, and $\Pi_3 : \mathcal{H}_f \to \mathcal{H}_f$. We can then embed F in $\mathscr{C}$ as (ϕ, Π) where $\Pi = \Pi_1 \otimes \Pi_2 \otimes \Pi_3$ and ϕ is the circuit below.

A catalytic embedding $(\phi, \Pi) : \mathbb{G} \to \mathscr{C}$ can always be extended to a catalytic embedding $\mathcal{C}(\mathbb{G}) \to \mathscr{C}$, e.g., as in Example 5. Such an extension takes a particularly nice form when the catalytic embedding (ϕ, Π) is homogeneous. In this case, the same catalyst can be used for any circuit over $\mathbb{G}$.

Definition 13. Let $(\phi, \Pi) : \mathbb{G} \to \mathscr{C}$ be a homogeneous catalytic embedding of dimension k of a gate set $\mathbb{G}$ in a collection of circuits $\mathscr{C}$ and let C be a circuit over $\mathbb{G}$. Then the *catalytic embedding of C in $\mathscr{C}$ induced by (ϕ, Π)* is the pair $(\overline{\phi}_C, \overline{\Pi}_C)$ where $\overline{\Pi}_C = \Pi$ and $\overline{\phi}$ is defined by induction on C as follows.

- If $C = I_n$ for some $n \in \mathbb{N}$, then $\overline{\phi}_C = I_n \otimes I_k$.

- If $C = \text{swap}(n, m)$ for some $n, m \in \mathbb{N}$, then $\overline{\phi}_C = \text{swap}(n, m) \otimes I_k$.

- If $C = G$ for some $G \in \mathbb{G}$, then $\overline{\phi}_C = \phi_G$.

- If $C = (C_1 \circ C_2)$, then $\overline{\phi}_C = (\overline{\phi}_{C_1} \circ \overline{\phi}_{C_2})$.

- If $C = (C_1 \otimes C_2)$, with C_1 of dimension m and C_2 of dimension n, then

$$\overline{\phi}_C = (((I_m \otimes \phi_{C_2}) \circ (I_m \otimes \text{swap}(k, n))) \circ (\phi_{C_1} \otimes I_n)) \circ (I_m \circ \text{swap}(n, k)).$$

We write $\overline{(\phi, \Pi)}$ for the catalytic embedding $\left\{ (\overline{\phi}_C, \overline{\Pi}_C) \mid C \in \mathcal{C}(\mathbb{G}) \right\} : \mathcal{C}(\mathbb{G}) \to \mathscr{C}$.

It can be verified that the induced catalytic embedding introduced in Definition 13 is a well-defined homogeneous catalytic embedding.

Example 6. Consider the circuit F from Example 5 and assume that F was given as

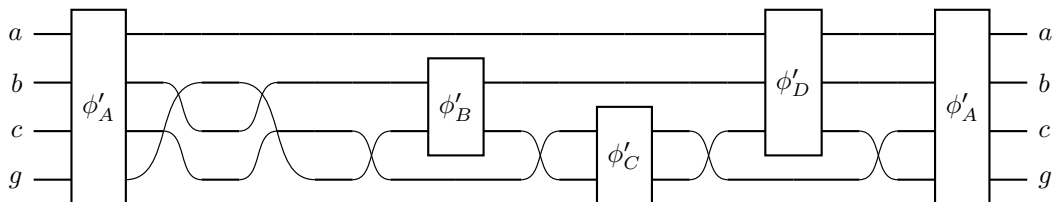$$F = A \circ ((D \otimes I_c) \circ ((I_a \otimes (B \otimes C)) \circ A)).$$

We can define a homogeneous catalytic embedding of $\mathbb{G}$ in $\mathscr{C}$ of dimension $g = d \cdot e \cdot f$ via

$$(\phi', \Pi) = \{(\phi'_A, \Pi), (\phi'_B, \Pi), (\phi'_C, \Pi), (\phi'_D, \Pi)\}$$

where $\Pi = \Pi_1 \otimes \Pi_2 \otimes \Pi_3$ and

$$\begin{aligned}
\phi'_A &= \phi_A \otimes I_{ef} \\
\phi'_B &= (I_{bd} \otimes \text{swap}(f, e)) \circ ((\phi_B \otimes I_e) \circ (I_{bd} \otimes \text{swap}(e, f))) \\
\phi'_C &= \phi_C \otimes I_f \\
\phi'_D &= (I_{abd} \otimes \text{swap}(f, e)) \circ (\phi_D \circ (I_{abd} \otimes \text{swap}(e, f))).
\end{aligned}$$

The corresponding circuit diagram for the image of F under the action of the catalytic embedding induced by (ϕ', Π) is the circuit depicted below.

It is a straightforward exercise to show that this circuit is equivalent to the naïve construction of Example 5 by using well-known properties of swap operations.

Induced catalytic embeddings can be used to concatenate embedings of gate sets. Suppose, for example, that $\mathbb{F}$, $\mathbb{G}$, and $\mathbb{H}$ are three gate sets and that $(\phi, \Pi): \mathbb{F} \to \mathcal{C}(\mathbb{G})$ and $(\phi', \Pi'): \mathbb{G} \to \mathcal{C}(\mathbb{H})$ are catalytic embeddings. We can then define a catalytic embedding $(\phi'', \Pi''): \mathbb{F} \to \mathcal{C}(\mathbb{H})$ by setting

$$(\phi'', \Pi'') = \overline{(\phi', \Pi')} \circ (\phi, \Pi).$$

Of course, the catalytic embedding (ϕ'', Π'') itself induces a catalytic embedding $\overline{(\phi'', \Pi'')}: \mathcal{C}(\mathbb{F}) \to \mathcal{C}(\mathbb{H})$.

IV. LINEAR CATALYTIC EMBEDDINGS

Catalytic embeddings, as introduced in the previous section, are very general. In principle, a catalytic embedding can be defined for arbitrary collections of circuits $\mathscr{C}$ and $\mathscr{C}'$ and is not required to preserve any structure, beyond that imposed by the catalytic condition. This generality can make it rather daunting to construct catalytic embeddings that might prove useful in any way. We thus turn our attention to catalytic embeddings that preserve the structure that may be present in underlying the collections of matrices $\mathcal{U}(\mathscr{C})$ and $\mathcal{U}(\mathscr{C}')$.

Definition 14. We say that a catalytic embedding $(\phi, \Pi): \mathscr{C} \to \mathscr{C}'$ is *strong* when,

$$e \circ \phi(C) = e \circ \phi(D) \iff e(C) = e(D)$$

for all $C, D \in \mathscr{C}$.

Proposition 3. *Let $(\phi, \Pi): \mathscr{C} \to \mathscr{C}'$ be a catalytic embedding. Then (ϕ, Π) is strong if, and only if, there exists an injective function $\mu: \mathcal{U}(\mathscr{C}) \to \mathcal{U}(\mathscr{C}')$ such that the following diagram commutes:*

$$
\begin{array}{ccc}
\mathscr{C} & \xrightarrow{\ \ \phi\ \ } & \mathscr{C}' \\
\downarrow{\scriptstyle e} & & \downarrow{\scriptstyle e} \\
\mathcal{U}(\mathscr{C}) & \xrightarrow[\ \ \mu\ \ ]{} & \mathcal{U}(\mathscr{C}').
\end{array}
$$

Moreover, the function $\mu: \mathcal{U}(\mathscr{C}) \to \mathcal{U}(\mathscr{C}')$ is uniquely determined by ϕ.

Proof. For the left-to-right direction, let $s: \mathcal{U}(\mathscr{C}) \to \mathscr{C}$ be any exact synthesis function and define $\mu = e \circ \phi \circ s$. Then, for any $C \in \mathscr{C}$, we have $\mu \circ e(C) = e \circ \phi \circ s \circ e(C)$. The function s is an exact synthesis function, so that $e \circ s$ is trivial and therefore that $e \circ (s \circ e)(C) = (e \circ s) \circ e(C) = e(C)$. Because (ϕ, Π) is strong, this implies that $e \circ \phi \circ s \circ e(C) = e \circ \phi(C)$ and therefore that $\mu \circ e(C) = e \circ \phi(C)$, so that the diagram in Proposition 3 commutes. Reasoning similarly, we get

$$\mu(U) = \mu(V) \implies e \circ \phi \circ s(U) = e \circ \phi \circ s(V) \implies e \circ s(U) = e \circ s(V) \implies U = V,$$

so that μ is injective.

For the right-to-left direction, assume that there exists an injective function $\mu: \mathcal{U}(\mathscr{C}) \to \mathcal{U}(\mathscr{C}')$ making the diagram in Proposition 3 commute. The equality $e(U) = e(V)$ implies $\mu \circ e(U) = \mu \circ e(V)$, which implies $e \circ \phi(U) = e \circ \phi(V)$ by commutativity of the diagram. Conversely, the equality $e \circ \phi(U) = e \circ \phi(V)$ implies $\mu \circ e(U) = \mu \circ e(V)$ by commutativity of the diagram, which implies $e(U) = e(V)$ by injectivity of μ. Hence, (ϕ, Π) is strong.

To see that μ is uniquely defined, note that if $\mu': \mathcal{U}(\mathscr{C}) \to \mathcal{U}(\mathscr{C}')$ makes the diagram in Proposition 3 commute then, for any $U \in \mathcal{U}(\mathscr{C}')$, we have $\mu'(U) = \mu' \circ (e \circ s)(U) = \mu' \circ e(s(U)) = e \circ \phi(s(U)) = \mu(U)$. $\qquad \square$

Note that in defining the function μ in Proposition 3, one can choose any exact synthesis function. Indeed, if s and s' are two such functions, then the uniqueness of μ implies that $e \circ \phi \circ s = e \circ \phi \circ s'$.

In practical contexts, one typically studies collections of circuits of the form $\mathcal{C}(\mathbb{G})$ and $\mathcal{C}(\mathbb{H})$, for some gate sets $\mathbb{G}$ and $\mathbb{H}$ of interest. In these cases the collections $\mathcal{C}(\mathbb{G})$ and $\mathcal{C}(\mathbb{H})$ are closed under several circuit-building operations. Consequently, the collections $\mathcal{U}(\mathcal{C}(\mathbb{G}))$ and $\mathcal{U}(\mathcal{C}(\mathbb{H}))$ are endowed with some structure. The next proposition shows that, for collections of this form, the map μ associated with a strong induced catalytic embedding preserves much of this structure.

Proposition 4. *Let $\mathbb{G}$ and $\mathbb{H}$ be two gate sets, let $(\phi, \Pi) : \mathcal{C}(\mathbb{G}) \to \mathcal{C}(\mathbb{H})$ be a strong induced k-dimensional catalytic embedding, and let $\mu : \mathcal{U}(\mathcal{C}(\mathbb{G})) \to \mathcal{U}(\mathcal{C}(\mathbb{H}))$ be as in Proposition 3. Then, for any $U, V \in \mathcal{U}(\mathcal{C}(\mathbb{G}))$ with $\mathrm{rank}(U) = \ell$ and $\mathrm{rank}(V) = m$,*

1. *$\mu(U)(I_\ell \otimes \Pi) = U \otimes \Pi$,*

2. *$\mu(I_n) = I_{kn}$ for any positive integer n,*

3. *$\mu(VU) = \mu(V)\mu(U)$ if $\ell = m$, and*

4. *$\mu(V \otimes U) = (\mathrm{swap}(\ell, m) \otimes I_k)(I_\ell \otimes \mu(V))(\mathrm{swap}(m, \ell) \otimes I_k)(I_m \otimes \mu(U))$.*

Proof. Recall from Proposition 3 that for any exact synthesis function $s : \mathcal{U}(\mathcal{C}(\mathbb{H})) \to \mathcal{C}(\mathbb{H})$, we have $\mu = e \circ \phi \circ s$. We take advantage of this here by choosing a convenient exact synthesis function for each one of the assertions to be established. For the first assertion, let s be arbitrary. By the catalytic condition and the fact that $e \circ s$ is trivial, we have

$$\mu(U)(I_\ell \otimes \Pi) = e \circ \phi \circ s(U)(I_\ell \otimes \Pi) = e \circ s(U) \otimes \Pi = U \otimes \Pi$$

so that the assertion follows. For the second assertion, take an exact synthesis map s such that $s(I_n)$ is the circuit I_n. Then, $\mu(I_n) = e \circ \phi \circ s(I_n) = e \circ \phi(I_n) = e(I_n \otimes I_k) = I_{nk}$ by Definition 13, since (ϕ, Π) is an induced catalytic embedding. For the third assertion, let s be a synthesis map such that $s(VU) = s(V) \circ s(U)$. Then, writing μ as $\mu = e \circ \phi \circ s$, we get

$$\mu(VU) = e \circ \phi \circ s(VU) = e \circ \phi(s(V) \circ s(U)) = e((\phi \circ s(V)) \circ (\phi \circ s(U))) = (e \circ \phi \circ s(V)) \circ (e \circ \phi \circ s(U)) = \mu(V)\mu(U),$$

using Definitions 6 and 13. For the final assertion, note that $V \otimes U = (V \otimes I_\ell)(I_m \otimes U) = (I_m \otimes U)(V \otimes I_\ell)$. Thus, the fourth assertion follows from the third, as long as we can show that $\mu(I_m \otimes U) = I_m \otimes \mu(U)$ and that

$$\mu(V \otimes I_\ell) = (\mathrm{swap}(\ell, m) \otimes I_k)(I_\ell \otimes \mu(V))(\mathrm{swap}(m, \ell) \otimes I_k).$$

To show that $\mu(I_m \otimes U) = I_m \otimes \mu(U)$, let s be such that $s(I_m \otimes U) = I_m \otimes s(U)$. We then have, using Definitions 6 and 13, and properties of the unitary $\mathrm{swap}(n, k)$,

$$
\begin{aligned}
\mu(I_m \otimes U) &= e \circ \phi \circ s(I_m \otimes U) \\
&= e \circ \phi(I_m \otimes s(U)) \\
&= e(((I_m \otimes \phi(s(U))) \circ (I_m \otimes \mathrm{swap}(k, n))) \circ (I_{mk} \otimes I_n)) \circ (I_m \circ \mathrm{swap}(n, k)) \\
&= e(I_m \otimes \phi(s(U))) \\
&= I_m \otimes e(\phi(s(U))) \\
&= I_m \otimes \mu(U).
\end{aligned}
$$

To show that $\mu(V \otimes I_\ell) = (\mathrm{swap}(\ell, m) \otimes I_k)(I_\ell \otimes \mu(V))(\mathrm{swap}(m, \ell) \otimes I_k)$, one can reason analogously, choosing an exact synthesis function s satisfying $s(V \otimes I_\ell) = s(V) \otimes I_\ell$. $\qquad\square$

Remark 4. Both $\mathcal{U}(\mathcal{C}(\mathbb{G}))$ and $\mathcal{U}(\mathcal{C}(\mathbb{H}))$ have the structure of a groupoid. Proposition 4 shows that the map μ is a faithful groupoid functor from $\mathcal{U}(\mathcal{C}(\mathbb{G}))$ to $\mathcal{U}(\mathcal{C}(\mathbb{H}))$.

It is natural to seek a converse statement to Proposition 4. Indeed, this would provide conditions under which a function $\mathcal{U}(\mathcal{C}(\mathbb{G})) \to \mathcal{U}(\mathcal{C}(\mathbb{H}))$ can be used to define a catalytic embedding $\mathcal{C}(\mathbb{G}) \to \mathcal{C}(\mathbb{H})$. This is the goal of the following proposition.

Proposition 5. *Let $\mathbb{G}$ and $\mathbb{H}$ be two gate sets and let $s : \mathcal{U}(\mathcal{C}(\mathbb{H})) \to \mathcal{C}(\mathbb{H})$ be an exact synthesis function. If $\Pi : \mathcal{H}_k \to \mathcal{H}_k$ is an orthogonal projector of nonzero rank and $\mu : \mathcal{U}(\mathcal{C}(\mathbb{G})) \to \mathcal{U}(\mathcal{C}(\mathbb{H}))$ is a function such that, for any $U, V \in \mathcal{U}(\mathcal{C}(\mathbb{G}))$ with $\mathrm{rank}(U) = \ell$ and $\mathrm{rank}(V) = m$,*

1. *$\mu(U)(I_\ell \otimes \Pi) = U \otimes \Pi$,*

2. *$\mu(I_n) = I_{kn}$ for any positive integer n,*

3. *$\mu(VU) = \mu(V)\mu(U)$ if $\ell = m$, and*

4. *$\mu(V \otimes U) = (\mathrm{swap}(\ell, m) \otimes I_k)(I_\ell \otimes \mu(V))(\mathrm{swap}(m, \ell) \otimes I_k)(I_m \otimes \mu(U))$,*

then $\overline{(\phi, \Pi)}$, *where* $\phi = s \circ \mu \circ e$, *is a strong and homogeneous k-dimensional catalytic embedding of* $\mathcal{C}(\mathbb{G})$ *in* $\mathcal{C}(\mathbb{H})$.

Proof. Write $\phi = s \circ \mu \circ e$. For $G \in \mathbb{G}$, we have

$$e(\phi(G))(I \otimes \Pi) = e \circ s \circ \mu \circ e(G)(I \otimes \Pi) = \mu(e(G))(I \otimes \Pi) = e(G) \otimes \Pi$$

so that $(\phi, \Pi) : \mathbb{G} \to \mathcal{C}(\mathbb{H})$ is a homogeneous catalytic embedding of dimension k. Write $\overline{(\phi, \Pi)}$ for the induced catalytic embedding. To show that $\overline{(\phi, \Pi)}$ is strong, we use Proposition 3. By the condition 1, μ is injective so that we only need to show that the diagram of Proposition 3 commutes. Let $U \in \mathcal{C}(\mathbb{G})$. By Definition 5, U is a well-formed word over $\mathbb{G} \cup \{ I_n \mid n \in \mathbb{N}^* \} \cup \{ \operatorname{swap}(m,n) \mid m, n \in \mathbb{N}^* \} \cup \{ \circ, \otimes, (,) \}$ and $\overline{\phi}(U)$ is the word obtained from U by Definition 13. Note that, for $G \in \mathbb{G}$, we have

$$e \circ \overline{\phi}(G) = e \circ \phi(G) = e \circ s \circ \mu \circ e(G) = \mu \circ e(G),$$

so that $e \circ \overline{\phi}$ and $\mu \circ e$ agree on $\mathbb{G}$. Since e is multiplicative and behaves trivially with respect to identities and swaps, conditions 2–4 ensure that

$$e \circ \overline{\phi}(U) = \mu \circ e(U).$$

Hence, the diagram commutes as desired. $\qquad\square$

Definition 15 (Lifting). Let μ and ϕ be as in Proposition 5. The catalytic embedding ϕ is called the *lifting* of μ.

Proposition 5 suggests that to find a catalytic embedding between $\mathcal{C}(\mathbb{G})$ and $\mathcal{C}(\mathbb{H})$, we can turn to the corresponding collection of matrices. Note that for any gate set $\mathbb{G}$, we can always find a number ring $\mathcal{S}$ such that $\mathcal{U}(\mathbb{G}) \subseteq \mathcal{U}(\mathcal{S})$. In fact, as discussed earlier, some important gate sets in quantum computing can be identified as exactly the set of unitary matrices of appropriate size which have entries in a particular number ring. These gate sets, such as the Clifford+T gate set, are often used for the fault-tolerant implementation of unitary operations. These properties naturally motivate the study of maps between unitary matrices over number rings, paying particular attention to maps into number rings associated with fault-tolerant gate sets. We thus consider what additional structure can be added to facilitate studying such maps. Unitary groups over rings are only closed under multiplication. However, when representing the elements of these groups in the computational basis as matrices over number rings, it is often convenient to use the additive structure of the underlying algebra of matrices. We thus focus on maps which respect this additive structure on matrices whilst also preserving unitarity.

Proposition 6. *Let $\mathcal{R}$ and $\mathcal{S}$ be Kroneckerian number rings with* $\operatorname{Frac}(\mathcal{R}), \operatorname{Frac}(\mathcal{S}) \subseteq \mathcal{K}$ *for a field $\mathcal{K}$, $\mathcal{T} = \mathcal{R} \cap \mathcal{S}$, $A, B \in \mathcal{M}(\mathcal{S})$, and $C \in \mathcal{M}(\mathcal{T})$. Suppose that $\Phi : \mathcal{M}(\mathcal{S}) \to \mathcal{M}(\mathcal{R})$ satisfies the following conditions:*

1. *$\Phi(AB) = \Phi(A)\Phi(B)$ when AB is defined*

2. *$\Phi(A + B) = \Phi(A) + \Phi(B)$ when $A + B$ is defined*

3. *$\Phi(I_n) = I_{kn}$ for some fixed k*

4. *$\Phi(A^\dagger) = \Phi(A)^\dagger$*

5. *$\Phi(C \otimes A) = C \otimes \Phi(A)$*

6. *There exists an orthogonal projector $\Pi : \mathcal{H}_k \to \mathcal{H}_k$ with $\operatorname{rank}(\Pi) > 0$ such that $\Phi(M)(I \otimes \Pi) = M \otimes \Pi$ where $\dim(M) = \dim(I)$ for all $M \in \mathcal{M}(\mathcal{S})$*

Then $\Phi(\mathcal{U}(\mathcal{S})) \subset \mathcal{U}(\mathcal{R})$ and the map $\mu : \mathcal{U}(\mathcal{S}) \to \mathcal{U}(\mathcal{R})$ with $\mu(U) := \Phi(U)$ can be lifted to a homogeneous, strong catalytic embedding of $\mathcal{C}(\mathcal{S})$ in $\mathcal{C}(\mathcal{R})$ of dimension k.

Proof. Let $U \in \mathcal{U}(\mathcal{S})$. By conditions 3 and 4 we have

$$\Phi(U)\Phi(U)^\dagger = \Phi(U)\Phi(U^\dagger) = \Phi(UU^\dagger) = \Phi(I) = I_k,$$

and so Φ maps unitary matrices to unitary matrices and μ is well-defined. Next, we would like to show that μ satisfies the conditions presented in Proposition 5 so that we can lift it to a strong and homogeneous catalytic embedding. By

inspection, we see that the only condition not immediately satisfied by our assumptions is the action of μ on $V \otimes U$ for $U, V \in \mathcal{U}(\mathcal{S})$. For $\operatorname{rank}(U) = \ell$ and $\operatorname{rank}(V) = m$, we calculate

$$\begin{aligned}
\mu(V \otimes U) &= \mu(\operatorname{swap}(l, m) \circ (I_\ell \otimes V) \circ \operatorname{swap}(m, l) \circ (I_m \otimes U)) \\
&= \mu(\operatorname{swap}(l, m)) \mu(I_\ell \otimes V) \mu(\operatorname{swap}(m, l)) \mu(I_m \otimes U) \\
&= \mu(\operatorname{swap}(l, m)) \otimes I_1) \mu(I_\ell \otimes V) \mu(\operatorname{swap}(m, l) \otimes I_1) \mu(I_m \otimes U) \\
&= (\operatorname{swap}(l, m)) \otimes \mu(I_1))(I_\ell \otimes \mu(V))(\operatorname{swap}(m, l) \otimes \mu(I_1))(I_m \otimes \mu(U)) \\
&= (\operatorname{swap}(l, m)) \otimes I_k)(I_\ell \otimes \mu(V))(\operatorname{swap}(m, l) \otimes I_k)(I_m \otimes \mu(U))
\end{aligned}$$

and thus all conditions of Proposition 5 hold. Therefore, we can lift μ to a homogeneous and strong catalytic embedding of $\mathcal{C}(\mathcal{S})$ over $\mathcal{C}(\mathcal{R})$ of dimension k. $\qquad \square$

As it might not be immediately obvious from the proof of Proposition 6, it is worth considering one consequence of condition 5. Let $U, V \in \mathcal{C}(\mathcal{S})$ be such that $e(V) \in \mathcal{U}(\mathcal{R})$ with $\dim(V) = \ell$ and $\dim(U) = m$. Lifting μ as in Proposition 5 to ϕ, we get

$$\begin{aligned}
e \circ \phi(V \otimes U) &= (\operatorname{swap}(\ell, m) \otimes I_k)(I_\ell \otimes e \circ \phi(V))(\operatorname{swap}(m, \ell) \otimes I_k)(I_m \otimes e \circ \phi(U)) \\
&= (\operatorname{swap}(\ell, m) \otimes I_k)(I_\ell \otimes \mu \circ e(V))(\operatorname{swap}(m, \ell) \otimes I_k)(I_m \otimes \mu \circ e(U)) \\
&= \mu(e(V) \otimes e(U)) \\
&= \Phi(e(V) \otimes e(U)) \\
&= e(V) \otimes \Phi(e(U)) \\
&= e(V) \otimes \mu \circ e(U) \\
&= e(V) \otimes e \circ \phi(U).
\end{aligned}$$

The above derivation shows that any lifting of μ acts trivially on the elements of $\mathcal{C}(\mathcal{R})$ which are equivalent to an element of $\mathcal{C}(\mathcal{S})$. This is consistent with the type of action we might desire in practice: we have no obvious need to embed gates or circuits which we already have direct access to.

We focus on studying maps that satisfy the conditions in Proposition 6 for the remainder of this document and so we present the following definitions.

Definition 16 (Linear catalytic embedding, Pre-embedding)**.** Let Φ and μ be defined as in Proposition 6 and let ϕ be the catalytic embedding that results from lifting μ. We say that ϕ is a *linear* catalytic embedding and that Φ is the *pre-embedding* of ϕ.

While we have introduced embeddings with varying amounts of structure, we might wonder whether such constructions exist in practice. In fact, there are distinct instances of each type of embedding, as the following example highlights.

Example 7. Let α be a primitive third root of unity. Consider the gate set

$$\mathbb{G} = \{R, X\} \quad \text{with} \quad e(R) = \begin{bmatrix} 1 & 0 \\ 0 & \alpha \end{bmatrix}, \ e(X) = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

Below are different examples of catalytic embeddings of R and X in $\mathcal{C}(\mathbb{Q})$.

1. $e(\phi_R) = \begin{bmatrix} I_3 & 0 \\ 0 & \Lambda \end{bmatrix}$, $e(\phi_X) = \begin{bmatrix} 0 & I_3 \\ I_3 & 0 \end{bmatrix}$, $\Pi = \frac{1}{3} \begin{bmatrix} 1 & \alpha & \alpha^2 \\ \alpha^2 & 1 & \alpha \\ \alpha & \alpha^2 & 1 \end{bmatrix}$, for $\Lambda = \frac{1}{3} \begin{bmatrix} -2 & -2 & 1 \\ 1 & -2 & -2 \\ -2 & 1 & -2 \end{bmatrix}$

2. $e(\phi_R) = \begin{bmatrix} I_3 & 0 \\ 0 & \Lambda \end{bmatrix}$, $e(\phi_X) = \begin{bmatrix} 0 & I_3 \\ I_3 & 0 \end{bmatrix}$, $\Pi = \frac{1}{3} \begin{bmatrix} 1 & \alpha & \alpha^2 \\ \alpha^2 & 1 & \alpha \\ \alpha & \alpha^2 & 1 \end{bmatrix}$, for $\Lambda = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$

3. $e(\phi_R) = \begin{bmatrix} I_4 & 0 \\ 0 & \Lambda \end{bmatrix}$, $e(\phi_X) = \begin{bmatrix} 0 & I_4 \\ I_4 & 0 \end{bmatrix}$, $\Pi = \frac{1}{6} \begin{bmatrix} 3 & 1+2\alpha & 1+2\alpha & 1+2\alpha \\ 1+2\alpha^2 & 3 & 1+2\alpha & 1+2\alpha^2 \\ 1+2\alpha^2 & 1+2\alpha^2 & 3 & 1+2\alpha \\ 1+2\alpha^2 & 1+2\alpha & 1+2\alpha^2 & 3 \end{bmatrix}$, for $\Lambda = \frac{1}{2} \begin{bmatrix} -1 & -1 & -1 & -1 \\ 1 & -1 & -1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \end{bmatrix}$

Each of these embeddings is homogeneous, and induces a catalytic embedding $\overline{(\phi, \Pi)}$ from $\mathcal{C}(\mathbb{G})$ to $\mathcal{C}(\mathbb{Q})$. These induced embeddings highlight the different properties a catayltic embedding may have:

- Embedding 1 is not strong: $e(R \circ R \circ R) = e(X \circ X)$ while $e(\phi_R \circ \phi_R \circ \phi_R) \neq e(\phi_X \circ \phi_X)$.

- Embedding 2 is strong, but not linear: $e(R \circ X \circ R) + e(R \circ R \circ X \circ R \circ R) = -e(X)$ while $e(\phi_R \circ \phi_X \circ \phi_R) + e(\phi_R \circ \phi_R \circ \phi_X \circ \phi_R \circ \phi_R) \neq -e(\phi_X)$

- Embedding 3 is linear.

These statements can be verified by direct computation.

We conclude this section with a number of simple properties of pre-embeddings and linear catalytic embeddings. First, we start with two lemmas that aid in characterizing possible constructions for linear catalytic embeddings.

Lemma 1. *Let $\mathcal{R}$ and $\mathcal{S} \subset \mathcal{T}$ be Kroneckerian subrings of a number field. If $\Phi : \mathcal{M}(\mathcal{T}) \to \mathcal{M}(\mathcal{R})$ is the pre-embedding for a linear catalytic embedding, then $\Phi|_{\mathcal{S}} : \mathcal{M}(\mathcal{S}) \to \mathcal{M}(\mathcal{R})$ is the pre-embedding for a linear catalytic embedding.*

While Lemma 1 may appear rather trivial, it is nonetheless quite useful for proving the (non-)existence of certain linear catalytic embeddings. For example, the contrapositive of Lemma 1 implies that if no such $\Phi|_{\mathcal{S}}$ can be constructed, then Φ cannot exist. This is especially powerful in instances where the elements of $\mathcal{T}$ are simpler to describe than those of $\mathcal{S}$ (or vice-versa).

Lemma 2. *Let $\mathcal{R}$ and $\mathcal{S}$ be Kroneckerian subrings of a number field. If a pre-embedding $\Phi : \mathcal{M}(\mathcal{S}) \to \mathcal{M}(\mathcal{R})$ for a linear catalytic embedding exists, then $\mathcal{S} \cap \mathcal{R} = \mathcal{S} \cap \mathrm{Frac}(\mathcal{R})$.*

Proof. We have $\mathcal{S} \cap \mathcal{R} \subseteq \mathcal{S} \cap \mathrm{Frac}(\mathcal{R})$, and so we just need to show that there is no $s \in \mathcal{S} \cap \mathrm{Frac}(\mathcal{R})$ with $s \notin \mathcal{S} \cap \mathcal{R}$. Suppose that such an s existed. There always exists some nonzero $t \in \mathcal{S} \cap \mathcal{R}$ such that $ts \in \mathcal{S} \cap \mathcal{R}$, which simultaneously implies

$$\Phi(ts) = t\Phi(s) \quad \text{and} \quad \Phi(ts) = tsI.$$

As $\mathcal{R}$ is an integral domain, by assumption we would conclude that $\Phi(s) = sI$, which is absurd as $s \in \mathcal{S} \cap \mathrm{Frac}(\mathcal{R})$ but $s \notin \mathcal{S} \cap \mathcal{R}$ implies $s \notin \mathcal{R}$. Thus, there can be no such s. $\qquad\square$

Lemma 2 precludes "reducing" the set of denominators that appear in number rings via embedding. This has implications for embedding to restricted gate sets such as Clifford+T, whose corresponding unitaries only permit particular denominators. We end this section with two important results for the implementation of linear catalytic embeddings. First, we show that linear catalytic embeddings respect the direct sum operation.

Proposition 7. *Let $\mathcal{R}$ and $\mathcal{S}$ be Kroneckarian number rings and $\Phi : \mathcal{M}(\mathcal{S}) \to \mathcal{M}(\mathcal{R})$ be the pre-embeddding of a linear catalytic embedding. Then, for $A, B \in \mathcal{M}(\mathcal{S})$,*

$$\Phi(A \oplus B) = \Phi(A) \oplus \Phi(B).$$

Proof. Let E_{ij} the standard basis matrices for $\mathcal{M}(\mathcal{S})$. We can write

$$A = \sum_{i,j=1}^{k} A_{ij} E_{ij} \text{ and } B = \sum_{i,j=1}^{m} B_{ij} E_{ij}$$

for k and m the dimension of A and B respectively. Then

$$A \oplus B = \sum_{i,j=1}^{k} A_{ij} E_{ij} + \sum_{i,j=k+1}^{m+k} B_{ij} E_{ij}$$

and

$$\Phi(A \oplus B) = \sum_{i,j=1}^{k} E_{ij} \otimes \Phi(A_{ij}) + \sum_{i,j=k+1}^{m+k} E_{ij} \otimes \Phi(B_{ij}).$$

We also have

$$\Phi(A) = \sum_{i,j=1}^{k} E_{ij} \otimes \Phi(A_{ij}) \text{ and } \Phi(B) = \sum_{i,j=1}^{m} E_{ij} \otimes \Phi(B_{ij}).$$

and so

$$\Phi(A) \oplus \Phi(B) = \sum_{i,j=1}^{k} E_{ij} \otimes \Phi(A_{ij}) + \sum_{i,j=k+1}^{m+k} E_{ij} \otimes \Phi(B_{ij}) = \Phi(A \oplus B). \qquad\square$$

Proposition 7 is useful when considering the action of linear catalytic embeddings with respect to controlled operations. When lifting a pre-embedding to a linear catalytic embedding ϕ, we can always construct a lift that preserves the controlled unitary structure between the control and target registers. In other words, for a gate G which applies U to a target register conditioned on a control register being in state $|a\rangle$, we can lift to the gate $\phi(G)$ which applies $\phi(U)$ to the target and catalyst registers conditioned on the control register being in state $|a\rangle$.

Finally, we show that the concatenation of two linear catalytic embeddings is a linear catalytic embedding, under a minor assumption.

Proposition 8. *Let $\mathcal{R}$, $\mathcal{S}$, and $\mathcal{T}$ be Kroneckarian subrings of a number field with $\mathcal{T} \cap \mathcal{R} \subseteq \mathcal{S}$ and $\Phi_1 : \mathcal{M}(\mathcal{T}) \to \mathcal{M}(\mathcal{S})$ and $\Phi_2 : \mathcal{M}(\mathcal{S}) \to \mathcal{M}(\mathcal{R})$ be pre-embeddings for linear catalytic embeddings. Then the concatenation $\Phi_2 \circ \Phi_1$ is a pre-embedding for a linear catalytic embedding.*

Proof. We show that the composition $\Phi_2 \circ \Phi_1$ satisfies the conditions of a pre-embedding.

1., 2. Properties 1 and 2 follow from the composition of ring homomorphism being a homomorphism.

3. Let k_1 and k_2 be such that $\Phi_1(I_n) = I_{k_1 n}$ and $\Phi_2(I_n) = I_{k_2 n}$ for all n. Then $\Phi_2 \circ \Phi_1(I_n) = \Phi_2(I_{k_1 n}) = I_{k_1 k_2 n}$ for all n.

4. Let $A \in \mathcal{M}(\mathcal{S})$. Then $\Phi_2 \circ \Phi_1(A^\dagger) = \Phi_2(\Phi_1(A)^\dagger) = (\Phi_2 \circ \Phi_1(A))^\dagger$.

5. Let $C \in \mathcal{M}(\mathcal{T} \cap \mathcal{R})$. Then $C \in \mathcal{M}(\mathcal{T} \cap \mathcal{S})$ and $C \in \mathcal{M}(\mathcal{S} \cap \mathcal{R})$ as $\mathcal{T} \cap \mathcal{R} \subset \mathcal{S}$ implies $C \in \mathcal{M}(\mathcal{R}), \mathcal{M}(\mathcal{S}), \mathcal{M}(\mathcal{T})$. So

$$\Phi_2 \circ \Phi_1(C \otimes A) = \Phi_2(C \otimes \Phi_1(A)) = C \otimes \Phi_2 \circ \Phi_1(A).$$

6. Let Π_1 and Π_2 be the projectors that satisfy the catalytic condition for Φ_1 and Φ_2 with dimension k_1 and k_2 respectively. Then $\Pi_1 \otimes \Pi_2 : \mathcal{H}_{k_1 k_2} \to \mathcal{H}_{k_1 k_2}$ with $\mathrm{rank}(\Pi_1 \otimes \Pi_2) = \mathrm{rank}(\Pi_1)\mathrm{rank}(\Pi_2) > 0$, and

$$(\Pi_1 \otimes \Pi_2)^\dagger(\Pi_1 \otimes \Pi_2) = (\Pi_1^\dagger \otimes \Pi_2^\dagger)(\Pi_1 \otimes \Pi_2) = (\Pi_1^\dagger \Pi_1) \otimes (\Pi_2^\dagger \Pi_2) = \Pi_1 \otimes \Pi_2$$

so that $\Pi_1 \otimes \Pi_2$ is a nonzero orthogonal projector on $\mathcal{H}_{k_1 k_2}$. Checking that it satisfies the catalytic condition for $A \in \mathcal{M}(\mathcal{T})$, we have

$$\begin{aligned}
\Phi_2 \circ \Phi_1(A)(I \otimes (\Pi_1 \otimes \Pi_2)) &= \Phi_2(\Phi_1(A))((I \otimes I_{k_1}) \otimes \Pi_2)((I \otimes \Pi_1) \otimes I_{k_2}) \\
&= (\Phi_1(A) \otimes \Pi_2)((I \otimes \Pi_1) \otimes I_{k_2}) \\
&= (A \otimes \Pi_1) \otimes \Pi_2 \\
&= A \otimes (\Pi_1 \otimes \Pi_2). \qquad\qquad \square
\end{aligned}$$

Proposition 8 allows us to concatenate sequences of linear catalytic embeddings together to simplify circuits. We will give an example of this process in the following section.

V. PROPERTIES OF LINEAR CATALYTIC EMBEDDINGS

In this section, we establish several important properties of linear catalytic embeddings. We show that linear catalytic embeddings can often be thought of as linear representations. We also prove that a linear catalytic embedding can be associated with a family of catalysts which can be used to perform non-trivial transformations of circuits. Finally, we provide a few simple examples to highlight the construction of linear catalytic embeddings.

Recall that if $\mathcal{T}$ is a subring of $\mathcal{S}$, then $\mathcal{S}$ is a $\mathcal{T}$-module, and that a generating set for $\mathcal{S}$ over $\mathcal{T}$ is a set $\Gamma \subseteq \mathcal{S}$ such that every element of $\mathcal{S}$ can be written as a finite $\mathcal{T}$-linear combination of elements of Γ. We begin this section by showing that a linear catalytic embedding is determined by its action on a generating set.

Proposition 9. *Let $\mathcal{R}$ and $\mathcal{S}$ be Kroneckarian number rings with $\mathcal{T} = \mathcal{R} \cap \mathcal{S}$ and let $\Phi : \mathcal{M}(\mathcal{S}) \to \mathcal{M}(\mathcal{R})$ be a pre-embedding. If Γ is a generating set for $\mathcal{S}$ over $\mathcal{T}$, then the action of Φ on $\mathcal{M}(\mathcal{S})$ is completely determined by its action on Γ.*

Proof. Let Γ be a generating set for $\mathcal{S}$ over $\mathcal{T}$. Since Γ generates $\mathcal{S}$ as a $\mathcal{T}$-module, for $s \in \mathcal{S}$ we can write

$$s = \sum_{g \in \Gamma} t^{(g)} g,$$

for $t^{(g)} \in \mathcal{T}$. For each i, j, let $E_{ij} \in \mathcal{M}(\mathcal{S})$ be the standard basis matrix, i.e., the matrix whose (i, j) entry is 1 and whose other entries are all 0. Suppose that $M \in \mathcal{M}(\mathcal{S})$. Then, for each i, j, we can write the i, j-th entry of M as

$$M_{i,j} = \sum_{g \in \Gamma} t_{i,j}^{(g)} g.$$

By applying the properties of a pre-embedding and the fact that scalar multiplication is equivalent to taking the tensor product with a 1×1 matrix, we then have

$$\Phi(M) = \Phi\left(\sum_{i,j} M_{i,j} E_{i,j}\right) = \Phi\left(\sum_{i,j} \sum_{g \in \Gamma} t_{i,j}^{(g)} g \cdot E_{i,j}\right) = \sum_{i,j} \sum_{g \in \Gamma} \Phi\left(t_{i,j}^{(g)} E_{i,j} \otimes g\right) = \sum_{i,j} \sum_{g \in \Gamma} t_{i,j}^{(g)} E_{i,j} \otimes \Phi(g).$$

Hence, for all $M \in \mathcal{M}(\mathcal{S})$, $\Phi(M)$ is completely determined by the action of Φ on Γ. $\qquad\square$

In much the same way that linear operators are determined by their action on a basis, linear catalytic embeddings are determined by their action on a small set that generates $\mathcal{M}(\mathcal{S})$. However, unlike linear operators, linear catalytic embeddings are not free to act in any way on a generating set. The set of matrices $\mathcal{M}(\mathcal{S})$ has more structure than than that of a module and linear catalytic embeddings have to preserve this additional structure. Our next theorem addresses this additional structure and shows that, in many important cases, a linear catalytic embedding is completely determined by a single element. In proving this result we use the *the Galois closure* of a field extension $\mathcal{F} \backslash \mathcal{E}$, which is the smallest field in which irreducible polynomials in $\mathcal{E}$ with a linear factor in $\mathcal{F}$ can be factored completely. For further details, we direct the interested reader to [21].

Theorem 1. *Let $\mathcal{R}$ and $\mathcal{S}$ be Kroneckarian number rings with $\mathcal{T} = \mathcal{R} \cap \mathcal{S}$. Let $\alpha \in \mathcal{S}$ such that $\mathrm{Frac}(\mathcal{S}) = \mathrm{Frac}(\mathcal{T})[\alpha]$ with minimal polynomial $p \in \mathrm{Frac}(\mathcal{T})[x]$ over $\mathrm{Frac}(\mathcal{T})$. We can construct the pre-embedding of a linear catalytic embedding of $\mathcal{C}(\mathcal{S})$ in $\mathcal{C}(\mathcal{R})$ if and only if we can construct $\Lambda \in \mathcal{M}(\mathcal{R})$ satisfying the following properties:*

1. *Λ is normal*

2. *$p(\Lambda) = 0$*

3. *Λ has α as one of its eigenvalues*

4. *there exists a generating set Γ for $\mathcal{S}$ over $\mathcal{T}$ such that for every $g \in \Gamma$, written uniquely as the sum over powers of α as*

$$g = \sum_{j=0}^{d-1} c_j \alpha^j$$

for some $c_j \in \mathrm{Frac}(\mathcal{T})$, the matrix

$$\sum_{j=0}^{d-1} c_j \Lambda^j$$

is a matrix over $\mathcal{R}$.

Proof. ($\Rightarrow$) Let $\Phi : \mathcal{M}(\mathcal{S}) \to \mathcal{M}(\mathcal{R})$ be a pre-embedding of a linear catalytic embedding. Let $\Lambda = \Phi(\alpha)$.

1. By the properties of linear catalytic embeddings and commutativity of $\mathcal{S}$,

$$\Phi(\alpha)\Phi(\alpha)^\dagger = \Phi(\alpha)\Phi(\alpha^\dagger) = \Phi(\alpha\alpha^\dagger) = \Phi(\alpha^\dagger\alpha) = \Phi(\alpha^\dagger)\Phi(\alpha) = \Phi(\alpha)^\dagger\Phi(\alpha)$$

and so Λ is a normal matrix.

2. Let $p \in \mathrm{Frac}(\mathcal{T})[x]$ be the minimal (monic) polynomial of α over $\mathrm{Frac}(\mathcal{T})$. There exists nonzero $t \in \mathcal{T}$ such that $t \cdot p = q \in \mathcal{T}[x]$ by clearing denominators. By the properties of linear catalytic embeddings

$$0 = \Phi(0) = \Phi(q(\alpha)) = q(\Phi(\alpha)) = t \cdot p(\Phi(\alpha)).$$

As $\mathrm{Frac}(\mathcal{T})$ is an integral domain, $p(\Lambda) = 0$.

3. Because Φ is the pre-embedding of a linear catalytic embedding, there exists a nonzero projector Π such that $\Phi(\alpha)(I \otimes \Pi) = \alpha \otimes \Pi$. Let $v \neq 0$ be in the image of Π. We have

$$\Phi(\alpha)(I \otimes \Pi)v = (\alpha \otimes \Pi)v$$
$$\Phi(\alpha)v = \alpha v$$
$$\Lambda v = \alpha v.$$

Thus, α is an eigenvalue of Λ.

4. Let Γ be a generating set for $\mathcal{S}$ over $\mathcal{T}$ and $g \in \Gamma$. There exist $c_j \in \mathrm{Frac}\,(\mathcal{T})$ such that

$$g = \sum_{j=0}^{d-1} c_j \alpha^j.$$

There exists nonzero $t \in \mathcal{T}$ such that $t \cdot c_j \in \mathcal{T}$ for all c_j by clearing denominators. By the properties of linear catalytic embeddings and by $\mathrm{Frac}\,(\mathcal{T})$ an integral domain, we have

$$tg = \sum_{j=0}^{d-1} tc_j \alpha^j$$

$$\Phi(tg) = \Phi\left(\sum_{j=0}^{d-1} tc_j \alpha^j\right)$$

$$t\Phi(g) = \sum_{j=0}^{d-1} tc_j \Phi(\alpha)^j$$

$$\Phi(g) = \sum_{j=0}^{d-1} c_j \Phi(\alpha)^j$$

$$= \sum_{j=0}^{d-1} c_j \Lambda^j \in \mathcal{M}(\mathcal{R})$$

because $\Phi(g) \in \mathcal{M}(\mathcal{R})$.

($\Leftarrow$) Suppose there exists Λ that satisfies the properties listed above and let Γ be a generating set for $\mathcal{S}$ over $\mathcal{T}$ for which property 4 is satisfied. As Λ is such that $p(\Lambda) = 0$, all eigenvalues of Λ are roots of p. We define the function $\Phi : \mathcal{M}(\mathcal{S}) \to \mathcal{M}(\mathcal{R})$ as follows:

$$\Phi(g) = \sum_{j=0}^{d-1} c_j^{(g)} \Lambda^j \text{ where } g = \sum_{j=0}^{d-1} c_j^{(g)} \alpha^j \text{ for all } g \in \Gamma \text{ and}$$

$$\Phi(M) = \sum_{i,j} \sum_{g \in \Gamma} t_{i,j}^{(g)} E_{i,j} \otimes \Phi(g) \text{ where } M = \sum_{i,j} \sum_{g \in \Gamma} t_{i,j}^{(g)} E_{i,j} \otimes g.$$

The function Φ is well-defined as each element of $\mathcal{S}$ can be written uniquely as a linear combination of powers of α over $\mathrm{Frac}\,(\mathcal{T})$. By extension, we can also uniquely write each element of $\mathcal{M}_n(\mathcal{S})$ as a linear combination of powers of α over $\mathcal{M}_n(\mathrm{Frac}\,(\mathcal{T}))$. Thus for arbitrary $M \in \mathcal{M}_n(\mathcal{S})$, we can always write

$$\Phi : M = \sum_{j=0}^{d-1} M_j \alpha^j \mapsto \sum_{j=0}^{d-1} M_j \otimes \Lambda^j$$

for some unique $M_j \in \mathcal{M}_n(\mathrm{Frac}\,(\mathcal{T}))$. We show that Φ satisfies the properties of the pre-embedding of a linear catalytic embedding. Let $A, B \in \mathcal{M}(\mathcal{S})$, and $C \in \mathcal{M}(\mathcal{T})$.

1. The space $\mathcal{H}_k$ is spanned by any linearly independent choice of eigenvectors of Λ. Since $p(\Lambda) = 0$, each such eigenvector has a corresponding eigenvalue $\sigma(\alpha)$ for σ a $\mathrm{Frac}\,(\mathcal{T})$-fixing automorphism of the Galois closure of $\mathrm{Frac}\,(\mathcal{S})$ as an extension of $\mathrm{Frac}\,(\mathcal{T})$. For arbitrary $|v\rangle \in \mathcal{H}_n$ and eigenvector $|\sigma(\alpha)\rangle$ of Λ, we have

$$\Phi(M)(|v\rangle \otimes |\sigma(\alpha)\rangle) = \sum_{j=0}^{d-1} M_j |v\rangle \otimes \Lambda^j |\sigma(\alpha)\rangle = \sum_{j=0}^{d-1} M_j |v\rangle \otimes \sigma(\alpha)^j |\sigma(\alpha)\rangle = \sigma(M) |v\rangle \otimes |\sigma(\alpha)\rangle.$$

Because σ is a homomorphism and vectors of the form $|v\rangle \otimes |\sigma(\alpha)\rangle$ span $\mathcal{H}_{kn}$, we conclude $\Phi(AB) = \Phi(A)\Phi(B)$ when AB is defined.

2. By definition of Φ and linearity of the tensor product, $\Phi(A + B) = \Phi(A) + \Phi(B)$ when $A + B$ is defined.

3. I_n has the unique decomposition as a linear combination of powers of α over $\mathrm{Frac}\,(\mathcal{T})$ of $I_n \alpha^0$. Thus

$$\Phi(I_n) = I_n \otimes \Lambda^0 = I_n \otimes I_k = I_{kn}.$$

4. For any $M \in \mathcal{M}_n(\mathcal{S})$ we write $M = \sum_{j=0}^{d-1} M_j \alpha^j$ for unique $M_j \in \mathcal{M}_n(\mathrm{Frac}\,(\mathcal{T}))$. This implies

$$M^\dagger = \sum_{j=0}^{d-1} M_j^\dagger (\alpha^\dagger)^j = \sum_{j=0}^{d-1} N_j \alpha^j$$

again for some unique $N_j \in \mathcal{M}_n(\mathrm{Frac}\,(\mathcal{T}))$ since $M^\dagger \in \mathcal{M}_n(\mathcal{S})$ by $\mathcal{S}$ Kroneckerian. Consider eigenvector $|\sigma(\alpha)\rangle$ of Λ for σ a $\mathrm{Frac}\,(\mathcal{T})$-fixing automorphism of the Galois closure of $\mathrm{Frac}\,(\mathcal{S})$ as an extension of $\mathrm{Frac}\,(\mathcal{T})$. As $\mathcal{S}$ and $\mathcal{T}$ are Kroneckerian, the Galois closure of $\mathrm{Frac}\,(\mathcal{S})$ as an extension of $\mathrm{Frac}\,(\mathcal{T})$ is necessarily Kroneckerian. As Λ is normal, $\Lambda^\dagger$ has the same eigenvectors as Λ with $\Lambda^\dagger |\sigma(\alpha)\rangle = \sigma(\alpha)^\dagger |\sigma(\alpha)\rangle = \sigma(\alpha^\dagger) |\sigma(\alpha)\rangle$ where we have used the defining property of Kroneckerian fields. For arbitrary $|v\rangle \in \mathcal{H}_n$ we have

$$\Phi(M)^\dagger (|v\rangle \otimes |\sigma(\alpha)\rangle) = \sum_{j=0}^{d-1} M_j^\dagger |v\rangle \otimes (\Lambda^\dagger)^j |\sigma(\alpha)\rangle = \sum_{j=0}^{d-1} M_j^\dagger |v\rangle \otimes \sigma(\alpha^\dagger)^j |\sigma(\alpha)\rangle = \sigma(M^\dagger) |v\rangle \otimes |\sigma(\alpha)\rangle.$$

On the other hand,

$$\Phi(M^\dagger)(|v\rangle \otimes |\sigma(\alpha)\rangle) = \sum_{j=0}^{d-1} N_j |v\rangle \otimes \Lambda^j |\sigma(\alpha)\rangle = \sum_{j=0}^{d-1} N_j |v\rangle \otimes \sigma(\alpha)^j |\sigma(\alpha)\rangle = \sigma(M^\dagger) |v\rangle \otimes |\sigma(\alpha)\rangle.$$

We conclude $\Phi(M)^\dagger = \Phi(M^\dagger)$ since vectors of the form $|v\rangle \otimes |\sigma(\alpha)\rangle$ span $\mathcal{H}_{kn}$.

5. By definition of Φ, $\Phi(C \otimes A) = C \otimes \Phi(A)$.

6. We have α is an eigenvalue of Λ. Let Π be the projector onto the eigenspace of Λ corresponding to α. Then

$$\Phi(M)(I \otimes \Pi) = \sum_{j=0}^{d-1} M_j \otimes \Lambda^j \Pi = \sum_{j=0}^{d-1} M_j \otimes \alpha^j \Pi = M \otimes \Pi. \qquad \square$$

Theorem 1 provides necessary and sufficient conditions for producing a linear catalytic embedding. More pertinently, it gives us a road map for constructing linear catalytic embeddings. If we are able to find a matrix Λ with the properties listed above, we can extend the map $\alpha \mapsto \Lambda$ to a linear catalytic embedding on all of $\mathcal{M}(\mathcal{S})$. Alternatively, if we can show that no such Λ exists, then this likewise precludes the existence of a linear catalytic embeddings. While Theorem 1 shows how a linear catalytic embedding is determined by its action on a single element, the following theorem gives a complete description of the image of a catalytic embedding.

Theorem 2. *Let $\mathcal{R}$ and $\mathcal{S}$ be Kroneckarian number rings, $\mathcal{K}$ be the Galois closure of $\mathrm{Frac}\,(\mathcal{S})/\mathrm{Frac}\,(\mathcal{S} \cap \mathcal{R})$, $\Phi : \mathcal{M}(\mathcal{S}) \to \mathcal{M}(\mathcal{R})$ be the pre-embedding of a k-dimensional linear catalytic embedding, and $\{\tau_m\}_{m=1}^n$ be the set of automorphisms on $\mathcal{K}\,\mathrm{Frac}\,(\mathcal{R})$ fixing $\mathrm{Frac}\,(\mathcal{R})$. There exists a set of automorphisms $\{\sigma_\ell\}_{\ell=1}^j$ on $\mathcal{K}$ fixing $\mathrm{Frac}\,(\mathcal{S} \cap \mathcal{R})$ and corresponding projectors $\{\Pi_\ell\}_{\ell=1}^j \subset \mathcal{M}_k(\mathcal{K}\,\mathrm{Frac}\,(\mathcal{R}))$ with $1 \le j \le [\mathcal{K} \cap \mathrm{Frac}\,(\mathcal{R}) : \mathrm{Frac}\,(\mathcal{S} \cap \mathcal{R})]$ such that*

1. *$\{\tau_m(\Pi_\ell) \mid 1 \le \ell \le j, 1 \le m \le n\}$ is a complete set of mutually-orthogonal projectors,*

2. *$\Phi(M)(I \otimes \tau_m(\Pi_\ell)) = (\tau_m \circ \sigma_\ell(M)) \otimes \tau_m(\Pi_\ell)$, and*

3. *$\Phi(M) = \sum_{\ell=1}^j \mathrm{tr}_{\mathrm{Frac}\,(\sigma_\ell(\mathcal{S})\,\mathcal{R})/\mathrm{Frac}\,(\mathcal{R})}(\sigma_\ell(M) \otimes \Pi_\ell)$.*

Moreover, the action of the circuit $\Phi(M)$ is completely determined by its action on the set $\{\Pi_\ell\}_{\ell=1}^j$.

Proof. Let $\alpha \in \mathcal{S}$ such that $\mathrm{Frac}\,(\mathcal{S}) = \mathrm{Frac}\,(\mathcal{S} \cap \mathcal{R})[\alpha]$ with minimal polynomial $p \in \mathrm{Frac}\,(\mathcal{S} \cap \mathcal{R})[x]$ with degree d. By the definition of Galois closure, $\mathcal{K}$ is the splitting field of p. Suppose p splits into irreducible factors $p = p_1 p_2 \ldots p_r$ over $\mathrm{Frac}\,(\mathcal{R})$. We shall bound this r value later. As p cannot have repeated roots, no p_ℓ has repeated roots nor can two factors share any common roots. Let λ_ℓ be any single root of p_ℓ for $1 \le \ell \le r$. Without loss of generality, we can assume $\lambda_1 = \alpha$. For each λ_ℓ, there exists an automorphism $\sigma_\ell : \mathcal{K} \to \mathcal{K}$ that fixes $\mathrm{Frac}\,(\mathcal{S}) \cap \mathrm{Frac}\,(\mathcal{R})$ such that $\sigma_\ell(\alpha) = \lambda_\ell$. Without loss of generality, we take σ_1 to be identity.

By Theorem 1, there exists a normal matrix $\Lambda \in \mathcal{M}(\mathcal{R})$ such that

$$p(\Lambda) = 0, \Phi(\alpha) = \Lambda, \text{ and } \Phi(M) = \sum_{i=0}^{d-1} M_i \otimes \Lambda^i \text{ where } M = \sum_{i=0}^{d-1} M_i \otimes \alpha^i \text{ and } M_i \in \mathcal{M}(\operatorname{Frac}(\mathcal{R})).$$

Let $q \in \mathcal{R}[x]$ be the characteristic polynomial of Λ with degree k. Since $p(\Lambda) = 0$, without loss of generality we can take $q = \pm p_1^{d_1} p_2^{d_2} \dots p_j^{d_j}$ with $d_\ell > 0$ and $j \le r$. This is consistent with our choice to make α a root of p_1, as $\Phi(\alpha) = \Lambda$ and thus α is an eigenvalue of Λ. For $1 \le \ell \le j$, observe that the eigenvector equation

$$(\Lambda - \lambda_\ell I)v = 0$$

has coefficients in the field $\operatorname{Frac}(\mathcal{R})[\lambda_\ell] = \operatorname{Frac}(\sigma_\ell(\mathcal{S})\mathcal{R}) \subseteq \mathcal{K}\operatorname{Frac}(\mathcal{R})$. As $\operatorname{Frac}(\sigma_\ell(\mathcal{S})\mathcal{R})$ is Kroneckerian, we can find a spanning set for all such v over $\operatorname{Frac}(\sigma_\ell(\mathcal{S})\mathcal{R})^k$ and apply the Gram-Schmidt procedure to obtain (unnormalized) eigenvectors $\{v_i^{(\ell)}\}_{i=1}^{d_\ell}$ with $v_i^{(\ell)} \in \operatorname{Frac}(\sigma_\ell(\mathcal{S})\mathcal{R})^k$. This basis can be used to define the projector onto the eigenspace of λ_ℓ,

$$\Pi_\ell = \sum_{i=1}^{d_\ell} \frac{v_i^{(\ell)} \left(v_i^{(\ell)}\right)^\dagger}{\langle v_i^{(\ell)}, v_i^{(\ell)} \rangle} \in \mathcal{M}_k(\operatorname{Frac}(\sigma_\ell(\mathcal{S})\mathcal{R})) \subseteq \mathcal{M}_k(\mathcal{K} \cap \operatorname{Frac}(\mathcal{R})),$$

which is necessarily orthogonal. This immediately implies that application of any $\operatorname{Frac}(\mathcal{R})$-fixing $\mathcal{K}\operatorname{Frac}(\mathcal{R})$-automorphism τ to any Π_ℓ must also be an orthogonal projector, as

$$\tau_m(\Pi_\ell)^\dagger \tau_m(\Pi_\ell) = \tau_m(\Pi_\ell^\dagger)\tau_m(\Pi_\ell) = \tau_m(\Pi_\ell^\dagger \Pi_\ell) = \tau_m(\Pi_\ell).$$

As $\mathcal{K}\operatorname{Frac}(\mathcal{R})/\operatorname{Frac}(\mathcal{R})$ is a Galois extension with intermediate field $\operatorname{Frac}(\sigma_\ell(\mathcal{S})\mathcal{R})$, $\operatorname{Gal}(\mathcal{K}\operatorname{Frac}(\mathcal{R})/\operatorname{Frac}(\mathcal{R}))$ has a subgroup G_ℓ which fixes $\operatorname{Frac}(\sigma_\ell(\mathcal{S})\mathcal{R})$. By definition, G_ℓ necessarily acts as identity on λ_ℓ. Let $\tau_a, \tau_b \in \operatorname{Gal}(\mathcal{K}\operatorname{Frac}(\mathcal{R})/\operatorname{Frac}(\mathcal{R}))$. We necessarily have

$$\tau_a G_\ell = \tau_b G_\ell \iff \tau_b^{-1}\tau_a \in G_\ell \iff (\tau_b^{-1}\tau_a)(\lambda_\ell) = \lambda_\ell \iff \tau_a(\lambda_\ell) = \tau_b(\lambda_\ell).$$

By the fundamental theorem of Galois theory, we have

$$|\operatorname{Gal}(\mathcal{K}\operatorname{Frac}(\mathcal{R})/\operatorname{Frac}(\mathcal{R})) : G_\ell| = [\operatorname{Frac}(\sigma_\ell(\mathcal{S})\mathcal{R}) : \operatorname{Frac}(\mathcal{R})]$$

distinct cosets of G_ℓ in $\operatorname{Gal}(\mathcal{K}\operatorname{Frac}(\mathcal{R})/\operatorname{Frac}(\mathcal{R}))$, implying there are precisely $[\operatorname{Frac}(\sigma_\ell(\mathcal{S})\mathcal{R}) : \operatorname{Frac}(\mathcal{R})]$ elements of $\{\tau_m(\lambda_\ell)\}_{m=1}^n$, each corresponding to exactly one root of p_ℓ. By application of the automorphism τ_m to the eigenprojector equation for Λ, we find

$$\Lambda \tau_m(\Pi_\ell) = \tau_m(\Lambda \Pi_\ell) = \tau_m(\lambda_\ell \Pi_\ell) = \tau_m(\lambda_\ell)\tau_m(\Pi_\ell)$$

so that not only is $\tau_m(\Pi_\ell)$ an orthogonal projector, it is an eigenprojector of Λ with eigenvalue $\tau_m(\lambda_\ell)$. As τ_m is invertible, $\tau_m(\Pi_\ell)$ must project onto the full eigenspace of $\tau_m(\sigma_\ell)$. Note that again there are only $[\operatorname{Frac}(\sigma_\ell(\mathcal{S})\mathcal{R}) : \operatorname{Frac}(\mathcal{R})]$ distinct elements for the set $\{\tau_m(\Pi_\ell) \mid 1 \le m \le n\}$, one for each root of p_ℓ.

As Λ is normal, its eigenspaces are mutually orthogonal, and hence we conclude that $\{\tau_m(\Pi_\ell) \mid 1 \le \ell \le j, 1 \le m \le n\}$ is a set of pairwise orthogonal projectors. Moreover, as there is a maximal projector for each root of q and hence eigenvalue of Λ, this set must also be complete, which proves the first part of the theorem. For the second statement, direct computation yields

$$\Phi(M)\left(I \otimes \tau_m(\Pi_\ell)\right) = \left(\sum_{i=0}^{d-1} M_i \otimes \Lambda^i\right)\left(I \otimes \tau_m(\Pi_\ell)\right) = \sum_{i=0}^{d-1} M_i \otimes \left(\Lambda^i \tau_m(\Pi_\ell)\right) = \sum_{i=0}^{d-1} M_i \otimes \tau_m(\sigma_\ell(\alpha))^i \tau_m(\Pi_\ell)$$

$$= \sum_{i=0}^{d-1} M_i \tau_m(\sigma_\ell(\alpha))^i \otimes \tau_m(\Pi_\ell) = \sum_{i=0}^{d-1} \tau_m(\sigma_\ell(M_i \alpha^i)) \otimes \tau_m(\Pi_\ell) = \tau_m \circ \sigma_\ell\left(\sum_{i=0}^{d-1} M_i \alpha^i\right) \otimes \tau_m(\Pi_\ell)$$

$$= \tau_m \circ \sigma_\ell(M) \otimes \tau_m(\Pi_\ell).$$

Choosing one coset representative $\overline{\tau}_m^{(\ell)}$ for each coset of $\operatorname{Gal}(\mathcal{K}\operatorname{Frac}(\mathcal{R})/\operatorname{Frac}(\mathcal{R}))/G_\ell$, completeness of our projectors implies

$$I = \sum_{\ell=1}^{j} \sum_{m=1}^{\deg(p_\ell)} \overline{\tau}_m^{(\ell)}(\Pi_\ell).$$

Using this relation alongside the second statement, we conclude

$$\Phi(M) = \Phi(M)(I \otimes I) = \Phi(M)\left(I \otimes \sum_{\ell=1}^{j} \sum_{m=1}^{\deg(p_\ell)} \overline{\tau}_m^{(\ell)}(\Pi_\ell)\right)$$

$$= \Phi(M)\left(\sum_{\ell=1}^{j} \sum_{m=1}^{\deg(p_\ell)} I \otimes \overline{\tau}_m^{(\ell)}(\Pi_\ell)\right) = \sum_{\ell=1}^{j} \sum_{m=1}^{\deg(p_\ell)} \Phi(M)\left(I \otimes \overline{\tau}_m^{(\ell)}(\Pi_\ell)\right)$$

$$= \sum_{\ell=1}^{j} \sum_{m=1}^{\deg(p_\ell)} \overline{\tau}_m^{(\ell)}(\sigma_\ell(M)) \otimes \overline{\tau}_m^{(\ell)}(\Pi_\ell) = \sum_{\ell=1}^{j} \mathrm{tr}_{\mathrm{Frac}(\sigma_\ell(\mathcal{S})\mathcal{R})/\mathrm{Frac}(\mathcal{R})}(\sigma_\ell(M) \otimes \Pi_\ell)$$

by definition of the field trace, and hence we have proved the third major statement. Note that $\Phi(M)$ is completely decomposed into its action on orthogonal subspaces of the catalyst space by this statement, each determined by an automorphism of $\mathcal{K}\,\mathrm{Frac}(\mathcal{R})$ and one projector from $\{\Pi_\ell\}_{\ell=1}^{j}$. Again, completeness allows us to determine the action of $\Phi(M)$ for any input state for the catalyst space, proving the fourth statement.

Finally, we bound $j \le r$, the number of irreducible factors of p over $\mathrm{Frac}(\mathcal{R})$. As $\mathcal{K}$ is Galois over $\mathrm{Frac}(\mathcal{S} \cap \mathcal{R})$, there is a subgroup $H \le \mathrm{Gal}(\mathcal{K}/\mathrm{Frac}(\mathcal{S} \cap \mathcal{R}))$ that fixes $\mathcal{K} \cap \mathrm{Frac}(\mathcal{R})$. Because H fixes $\mathcal{K} \cap \mathrm{Frac}(\mathcal{R})$, elements of H can *only* permute the roots of irreducible polynomials over $\mathcal{K} \cap \mathrm{Frac}(\mathcal{R})$. Therefore, because $\sigma_{\ell_1}(\alpha)$ and $\sigma_{\ell_2}(\alpha)$ are roots of different polynomials when $\ell_1 \ne \ell_2$ by assumption, $\sigma_{\ell_1} H \ne \sigma_{\ell_2} H$ and each $\sigma_\ell H$ defines a different coset of H. By the fundamental theorem of Galois theory, there are precisely $[\mathcal{K} \cap \mathrm{Frac}(\mathcal{R}) : \mathrm{Frac}(\mathcal{S} \cap \mathcal{R})]$ cosets, and hence $j \le r = [\mathcal{K} \cap \mathrm{Frac}(\mathcal{R}) : \mathrm{Frac}(\mathcal{S} \cap \mathcal{R})]$. $\qquad\square$

Breaking down Theorem 2, we see a complete characterization of the structure of a linear catalytic embedding. Given a linear catalytic embedding, statement one says that the catalyst subsystem naturally decomposes into an orthonomal basis of "catalyst-like" states. The second statement shows that embedded circuits fix this basis of the catalyst subsystem. In other words, the embedded circuit is a controlled operator in the "catalyst basis." But how does the embedded circuit behave in the presence of the other basis states? The second, third, and fourth statements tell us that an embedded circuit acts like the original circuit composed with an automorphism of a Galois field. These automorphisms map $\mathcal{S}$ to an isomorphic copy of $\mathcal{S}$, thereby changing the original unitary to some isomorphic copy. These automorphisms also relate many of the projectors, so that at most $[\mathcal{K} \cap \mathrm{Frac}(\mathcal{R}) : \mathrm{Frac}(\mathcal{S} \cap \mathcal{R})]$ need to be computed separately. This bound is a measure of how polynomials factor over $\mathrm{Frac}(\mathcal{R})$ as opposed to $\mathrm{Frac}(\mathcal{S} \cap \mathcal{R})$. If certain irreducible polynomials over $\mathrm{Frac}(\mathcal{S} \cap \mathcal{R})$ factor over $\mathrm{Frac}(\mathcal{R})$, more states will need to be computed. We can characterize precisely when the construction simplifies, as the following corollary shows:

Corollary 1. *Let $\mathcal{R}$ and $\mathcal{S}$ be Kroneckarian number rings, $\mathcal{K}$ be the Galois closure of $\mathrm{Frac}(\mathcal{S})/\mathrm{Frac}(\mathcal{S} \cap \mathcal{R})$, $\Phi : \mathcal{M}(\mathcal{S}) \to \mathcal{M}(\mathcal{R})$ be the pre-embedding of a linear catalytic embedding, and $\{\tau_m\}_{m=1}^{n}$ be the set of automorphisms on $\mathcal{K}\,\mathrm{Frac}(\mathcal{R})$ fixing $\mathrm{Frac}(\mathcal{R})$. If $\mathrm{Frac}(\mathcal{S} \cap \mathcal{R}) = \mathcal{K} \cap \mathrm{Frac}(\mathcal{R})$, then there exists a single orthogonal projector $\Pi \in \mathcal{M}(\mathrm{Frac}(\mathcal{S}\mathcal{R}))$ such that*

1. *$\{\tau_m(\Pi) \mid 1 \le m \le n\}$ is a complete set of mutually-orthogonal projectors with $[\mathrm{Frac}(\mathcal{S}) : \mathrm{Frac}(\mathcal{S} \cap \mathcal{R})]$ distinct elements,*

2. *$\Phi(M)(I \otimes \tau_m(\Pi)) = \tau_m(M) \otimes \tau_m(\Pi)$ and,*

3. *$\Phi(M) = \mathrm{tr}_{\mathrm{Frac}(\mathcal{S}\mathcal{R})/\mathrm{Frac}(\mathcal{R})}(M \otimes \Pi)$.*

Moreover, the action of the circuit $\Phi(M)$ is completely determined by its action on Π.

Proof. As $\mathcal{K}$ is a Galois extension of $\mathrm{Frac}(\mathcal{S} \cap \mathcal{R}) = \mathcal{K} \cap \mathrm{Frac}(\mathcal{R})$, $\mathcal{K}$ and $\mathrm{Frac}(\mathcal{R})$ are linearly disjoint as extensions of their intersection $\mathcal{K} \cap \mathrm{Frac}(\mathcal{R})$. Any subfield of $\mathcal{K}$ is necessarily also linearly disjoint with $\mathrm{Frac}(\mathcal{R})$ as an extension of $\mathcal{K} \cap \mathrm{Frac}(\mathcal{R})$, and thus $[\mathrm{Frac}(\mathcal{S}\mathcal{R}) : \mathrm{Frac}(\mathcal{R})] = [\mathrm{Frac}(\mathcal{S}) : \mathrm{Frac}(\mathcal{S} \cap \mathcal{R})]$. Applying Theorem 2 with $\mathrm{Frac}(\mathcal{S} \cap \mathcal{R}) = \mathcal{K} \cap \mathrm{Frac}(\mathcal{R})$, we have $1 \le j \le [\mathcal{K} \cap \mathrm{Frac}(\mathcal{R}) : \mathrm{Frac}(\mathcal{S} \cap \mathcal{R})] = 1$, and so $j = 1$ such that the only automorphism of $\mathcal{K} \cap \mathrm{Frac}(\mathcal{R})$ fixing $\mathrm{Frac}(\mathcal{S} \cap \mathcal{R})$ is identity. We can thus write $\sigma_1(M) = M$ and $\Pi_1 = \Pi$, and the corollary follows by noting the number of distinct projectors is precisely $[\mathrm{Frac}(\mathcal{S}\mathcal{R}) : \mathrm{Frac}(\mathcal{R})]$ given the proof of Theorem 2. $\qquad\square$

The hypothesis $\mathrm{Frac}(\mathcal{S} \cap \mathcal{R}) = \mathcal{K} \cap \mathrm{Frac}(\mathcal{R})$ is satisfied in most cases of practical interest. For example, this condition holds trivially when $\mathcal{R} \subset \mathcal{S}$, which coincides with the notion that our goal with linear catalytic embeddings is to "simplify" the ring over which we apply circuits. When $\mathrm{Frac}(\mathcal{S} \cap \mathcal{R}) \ne \mathcal{K} \cap \mathrm{Frac}(\mathcal{R})$, we are not so much simplifying circuits as we are working (at least partially) over an isomorphic copy of $\mathcal{S}$. For the remainder of this section, we thus focus on illuminating the consequences of Theorem 1 and Corollary 1 with some examples to show linear catalytic embeddings at work.

Example 8. Let $\mathcal{S} = \mathbb{Q}[\sqrt{5}]$ and $\mathcal{R} = \mathbb{Q}$ so that $\mathcal{R}, \mathcal{S}$ are both their own field of fractions. The element $\sqrt{5}$ plays the role of α in Theorem 1 with minimal polynomial $p = x^2 - 5$. We need to find a normal matrix Λ such that $\Lambda^2 - 5I = 0$. The matrix

$$\Lambda = \begin{bmatrix} 1 & 2 \\ 2 & -1 \end{bmatrix}$$

has characteristic polynomial p, and so $\Lambda^2 - 5I = 0$. Additionally, Λ has $\sqrt{5}$ as an eigenvalue and the set $\{1, \sqrt{5}\}$ is a generating set satisfying condition 4 of Theorem 1. The pre-embedding $\Phi : \mathcal{M}(\mathbb{Q}[\sqrt{5}]) \to \mathcal{M}(\mathbb{Q})$ is constructed by extending the map $\sqrt{5} \mapsto \Lambda$ as:

$$\Phi(M) = M_0 \otimes \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + M_1 \otimes \begin{bmatrix} 1 & 2 \\ 2 & -1 \end{bmatrix},$$

where $M = M_0 + M_1 \sqrt{5}$ with $M_0, M_1 \in \mathcal{M}(\mathbb{Q})$. As every quadratic number field extension is Galois, we have $\mathcal{K} = \mathcal{S} = \mathbb{Q}[\sqrt{5}]$, whose automorphisms as an extension of $\mathbb{Q}$ are given by $\{\text{id}, \tau\}$ for $\tau : \sqrt{5} \mapsto -\sqrt{5}$ since $-\sqrt{5}$ is the only remaining root of p. Additionally, $\mathcal{R} \subset \mathcal{S}$ implies we can apply Corollary 1 to determine the projectors and action of this embedding. We can easily compute the projector onto the eigenspace of Λ corresponding to $\sqrt{5}$:

$$\Pi = \begin{bmatrix} \frac{1}{2}\left(1 + \frac{\sqrt{5}}{5}\right) & \frac{\sqrt{5}}{5} \\ \frac{\sqrt{5}}{5} & \frac{1}{2}\left(1 - \frac{\sqrt{5}}{5}\right) \end{bmatrix}.$$

Checking the action of our embedding, we thus have

$$\Phi(M)(I \otimes \Pi) = M_0 \otimes (I\Pi) + M_1 \otimes (\Lambda\Pi) = M_0 \otimes \Pi + M_1 \otimes (\sqrt{5}\Pi) = M \otimes \Pi$$

as expected. Per Corollary 1, we should have that

$$\Pi = \begin{bmatrix} \frac{1}{2}\left(1 + \frac{\sqrt{5}}{5}\right) & \frac{\sqrt{5}}{5} \\ \frac{\sqrt{5}}{5} & \frac{1}{2}\left(1 - \frac{\sqrt{5}}{5}\right) \end{bmatrix} \xrightarrow{\tau} \begin{bmatrix} \frac{1}{2}\left(1 - \frac{\sqrt{5}}{5}\right) & -\frac{\sqrt{5}}{5} \\ -\frac{\sqrt{5}}{5} & \frac{1}{2}\left(1 + \frac{\sqrt{5}}{5}\right) \end{bmatrix} = \tau(\Pi)$$

is itself a projector so that $\{\Pi, \tau(\Pi)\}$ is a mutually orthogonal and complete set of projectors with cardinality $[\mathbb{Q}[\sqrt{5}] : \mathbb{Q}] = 2$. It is indeed clear that $\tau(\Pi) \neq \Pi$, and furthermore it is straightforward to verify that $\Pi^\dagger \Pi = \Pi$, $\tau(\Pi)^\dagger \tau(\Pi) = \tau(\Pi)$, $\tau(\Pi)^\dagger \Pi = 0$, and $\Pi + \tau(\Pi) = I$ as expected.

We can also check that our embedding suffices for extraction of the matrix $\tau(M) = M_0 - \sqrt{5}M_1$ via $\tau(\Pi)$. Indeed, we have that $\Lambda\tau(\Pi) = -\sqrt{5}\tau(\Pi)$ so that

$$\Phi(M)(I \otimes \tau(\Pi)) = M_0 \otimes (I\tau(\Pi)) + M_1 \otimes (\Lambda\tau(\Pi)) = M_0 \otimes \tau(\Pi) + M_1 \otimes (-\sqrt{5}\tau(\Pi)) = \tau(M) \otimes \tau(\Pi).$$

Finally, by completeness of our projectors Π and $\tau(\Pi)$, we have

$$\Phi(M) = \Phi(M)(I \otimes (\Pi + \tau(\Pi))) = M \otimes \Pi + \tau(M) \otimes \tau(\Pi) = \text{tr}_{\mathbb{Q}[\sqrt{5}]/\mathbb{Q}}(M \otimes \Pi)$$

which confirms that the image of M under Φ is given by the expected field trace.

Example 9. Let $\omega = e^{2\pi i/8}$ and $\mathcal{R} = \mathbb{D} = \mathbb{Z}[1/2]$, $\mathcal{S} = \mathbb{D}[i = \omega^2]$, and $\mathcal{T} = \mathbb{D}[\omega]$. We will construct the pre-embeddings of linear catalytic embeddings $\Phi_1 : \mathcal{M}(\mathcal{T}) \to \mathcal{M}(\mathcal{S})$ and $\Phi_2 : \mathcal{M}(\mathcal{S}) \to \mathcal{M}(\mathcal{R})$ to construct the pre-embedding $\Phi_2 \circ \Phi_1 : \mathcal{M}(\mathcal{T}) \to \mathcal{M}(\mathcal{R})$. In light of number-theoretic characterizations of certain fault-tolerant gate sets [2, 4], this construction yields a structure-preserving map from the Clifford-T gate set ($\mathbb{D}[\omega]$) to subsets $\mathbb{D}[i]$ and $\mathbb{D}$ of the Clifford+CS and Hadamard+Toffoli gate sets, respectively.

We begin by finding the pre-embedding $\Phi_1 : \mathcal{M}(\mathcal{T}) \to \mathcal{M}(\mathcal{S})$. In this case, Frac $(\mathbb{D}[\omega]) = \mathbb{Q}[\omega]$ and Frac $(\mathbb{D}[i]) = \mathbb{Q}[i]$. The element ω plays the role of α from Theorem 1 with minimal polynomial $p_1 = x^2 + i$, and the set $\{1, \omega\}$ plays the role of Γ. The normal matrix

$$\Lambda_1 = \begin{bmatrix} 0 & 1 \\ i & 0 \end{bmatrix}$$

has p_1 as its characteristic polynomial, and so satisfies p_1 as required. In addition ω is an eigenvalue of Λ_1 and so Λ_1 satisfies all the properties of Theorem 1. The linear catalytic embedding $\Phi_1 : \mathcal{M}(\mathcal{T}) \to \mathcal{M}(\mathcal{S})$ is constructed by extending the map $\omega \mapsto \Lambda_1$ as follows:

$$\Phi_1(M) = M_0 \otimes I + M_1 \otimes \Lambda_1 \text{ where } M = M_0 + M_1\omega \text{ with } M_0, M_1 \in \mathcal{M}(\mathcal{S}).$$

We have that the extension $\mathbb{Q}[\omega]/\mathbb{Q}[i]$ is Galois so that there are two distinct automorphisms of $\mathbb{Q}[\omega]$ which fix $\mathbb{Q}[i]$. These are given by

$$\{\text{id} : \omega \mapsto \omega, \tau_1 : \omega \mapsto -\omega, \}.$$

As we have $\mathcal{S} \subset \mathcal{T}$, we can apply Corollary 1. The projector Π_1 is the projector onto the subspace of Λ_1 corresponding to ω given by

$$\Pi_1 = \frac{1}{2} \begin{bmatrix} 1 & \omega^7 \\ \omega & 1 \end{bmatrix}.$$

It is a straightforward exercise to check that $\tau_1(\Pi_1)$ is the projector onto the eigenspace of Λ_1 corresponding to $\tau_1(\omega) = -\omega$, and furthermore that these projectors are mutually orthogonal, complete, and satisfy their respective catalytic conditions.

Constructing the pre-embedding $\Phi_2 : \mathcal{M}(\mathbb{D}[i]) \to \mathcal{M}(\mathbb{D})$ is remarkably similar. As $\text{Frac}\,(\mathbb{D}) = \mathbb{Q}$, the element i plays the role of α from Theorem 1 with minimal polynomial $p_2 = x^2 + 1$, and the set $\{1, i\}$ plays the role of Γ. The normal matrix

$$\Lambda_2 = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

satisfies the conditions of Theorem 1, and so we can define $\Phi_2 : \mathcal{M}(\mathcal{S}) \to \mathcal{M}(\mathcal{R})$ by

$$\Phi_2(M) = M_0 \otimes I + M_1 \otimes \Lambda_2 \text{ where } M = M_0 + M_1 i \text{ with } M_0, M_1 \in \mathcal{M}(\mathcal{R}).$$

Again, the extension $\mathbb{Q}[i]/\mathbb{Q}$ is Galois so that there are two distinct automorphisms of $\mathbb{Q}[i]$ which fix $\mathbb{Q}$. These are

$$\{\text{id} : i \mapsto i, \tau_2 : i \mapsto -i, \}.$$

Since $\mathcal{R} \subset \mathcal{S}$, we can again apply Corollary 1. The projector Π_2 projects onto the subspace of Λ_2 corresponding to i given by

$$\Pi_2 = \frac{1}{2} \begin{bmatrix} 1 & i^3 \\ i & 1 \end{bmatrix}.$$

As before, $\tau_2(\Pi_2)$ projects onto the eigenspace of Λ_2 corresponding to $\tau_2(i) = -i$, and the projectors are mutually orthogonal, complete, and satisfy their respective catalytic conditions.

We now check that the concatenation $\Phi_2 \circ \Phi_2$ behaves as expected and yields a pre-embedding $\Phi : \mathcal{M}(\mathcal{T}) \to \mathcal{M}(\mathcal{R})$ for a linear catalytic embedding. Given

$$M = M_0 + \omega M_1 + \omega^2 M_2 + \omega^3 M_3$$

for the basis $\{1, \omega, \omega^2, \omega^3\}$ of $\mathcal{T}$ as an $\mathcal{R}$ module we have

$$\Phi(M) = \Phi_2 \circ \Phi_1(M)$$
$$= \Phi_2\big((M_0 + \omega^2 M_2) \otimes I_2 + (M_1 + \omega^2 M_3) \otimes \Lambda_1\big)$$
$$= M_0 \otimes I_4 + M_1 \otimes \Phi_2(\Lambda_1) + M_2 \otimes I_2 \otimes \Lambda_2 + M_3 \otimes \Phi_2(i\Lambda_1).$$

Computing $\Phi_2(\Lambda_1)$, we have

$$\Phi_2(\Lambda_1) = \Lambda = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ -1 & 0 & 0 & 0 \end{bmatrix}.$$

In fact, we have $\Lambda^2 = I_2 \otimes \Lambda_2$ and $\Lambda^3 = \Phi_2(i\Lambda_1)$ so that really

$$\Phi(M) = M_0 \otimes \Lambda + M_1 \otimes \Lambda + M_2 \otimes \Lambda^2 + M_3 \otimes \Lambda^3.$$

We thus identify Λ as the normal matrix $\Phi(\omega)$, and indeed the characteristic polynomial of Λ is $x^4 + 1$, which is precisely the minimal polynomial of ω over $\mathbb{Q}$. Furthermore, the projector

$$\Pi_1 \otimes \Pi_2 = \Pi = \frac{1}{4} \begin{bmatrix} 1 & \omega^6 & \omega^7 & \omega^5 \\ \omega^2 & 1 & \omega & \omega^7 \\ \omega & \omega^7 & 1 & \omega^6 \\ \omega^3 & \omega & \omega^2 & 1 \end{bmatrix}$$

projects onto the ω eigenspace of Λ as expected. The automorphisms of $\mathbb{Q}[\omega]$ which fix $\mathbb{Q}$ are

$$\left\{ \mathrm{id} : \omega \mapsto \omega, \sigma_1 : \omega \mapsto \omega^3, \sigma_2 : \omega \mapsto \omega^5, \sigma_3 : \omega \mapsto \omega^7 \right\}.$$

and their action on Π yield orthogonal projectors which are mutually orthogonal, complete, and satisfy the appropriate catalytic conditions. Thus we have constructed the desired pre-embedding via concatenation.

VI.   STANDARD CATALYTIC EMBEDDINGS

Ring extensions are often built by taking quotients. This is especially true in the context of the number rings encountered in fault-tolerant quantum computing [4]. In these cases, there is a convenient method to build linear catalytic embeddings, which can be seen as a simplification of Theorem 1. We call the linear catalytic embeddings defined in this way *standard catalytic embeddings*.

A.   Definitions and Properties

We start by introducing a generalization of the notion of companion matrix which will play, in the context of standard catalytic embeddings, the role of matrix Λ in Theorem 1, as seen in Examples 8 and 9.

Definition 17 (Pseudo-Companion Matrix)**.** Let $\mathcal{R}$ be a number ring and let $p \in \mathcal{R}[X]$ be a monic polynomial over $\mathcal{R}$. A matrix $\Lambda \in \mathcal{M}(\mathcal{R})$ is a *pseudo-companion matrix for p* if the characteristic polynomial of Λ is $\pm p^c$, for some positive integer c.

Proposition 10. *Let $\mathcal{R} \subset \mathcal{R}[\alpha]$ be a Kroneckerian, integral extension of number rings such that α has minimal polynomial $p \in \mathrm{Frac}(\mathcal{R})[x]$ with coefficients in $\mathcal{R}$ and degree d. Suppose $\Lambda \in \mathcal{M}(\mathcal{R})$ is a normal pseudo-companion matrix for p. If $A \in \mathcal{M}(\mathcal{R}[\alpha])$, then:*

1. *there exist unique $A_i \in \mathcal{M}(\mathcal{R})$ such that*

$$A = \sum_{i=0}^{d-1} A_i \alpha^i,$$

2. *the map $\Phi : \mathcal{M}(\mathcal{R}[\alpha]) \to \mathcal{M}(\mathcal{R})$ given by*

$$A \mapsto \sum_{i=0}^{k} A_i \otimes \Lambda^i$$

   is the pre-embedding for a linear catalytic embedding as defined in Proposition 6.

This result follows from Theorem 1, but we outline the proof here.

Proof. Let $A, B \in \mathcal{M}(\mathcal{R}[\alpha])$ and $C \in \mathcal{M}(\mathcal{R})$. Because p is irreducible over $\mathrm{Frac}(\mathcal{R})$ and has coefficients in $\mathcal{R}$, $\mathcal{R}[\alpha]$ is a free module over $\mathcal{R}$ with basis $\{\alpha^i\}_{i=0}^{d-1}$. Therefore, for each $A \in \mathcal{M}(\mathcal{R}[\alpha])$, there exist unique $A_i \in \mathcal{M}(\mathcal{R})$ such that

$$A = \sum_{i=0}^{d-1} A_i \alpha^i.$$

We now show that Φ satisfies the conditions of a pre-embedding of a linear catalytic embedding.

1. Because Λ is normal and has characteristic polynomial $\pm p^m$ for some positive integer m, the eigenvalues of Λ are all the roots of p, each occurring with multiplicity m. In addition, the eigenvectors of Λ span $\mathcal{H}_k$. Let $|x\rangle$ be an eigenvector of Λ with eigenvalue λ and $|v\rangle \in \mathcal{H}_n$ be an arbitrary state. Because p is irreducible, there always exists a ring isomorphism $\sigma : \mathcal{R}[\alpha] \to \mathcal{R}[\lambda]$ that fixes $\mathcal{R}$ and maps $\sigma(\alpha) = \lambda$. Then

$$\Phi(M)(|v\rangle \otimes |x\rangle) = \sum_{i=0}^{d-1} M_i |v\rangle \otimes \Lambda^i |x\rangle = \sum_{i=0}^{d-1} M_i |v\rangle \otimes \lambda^i |x\rangle = \sum_{i=0}^{d-1} M_i |v\rangle \otimes \sigma(\alpha)^i |x\rangle = (\sigma(M)|v\rangle) \otimes |x\rangle.$$

   Because σ is a ring isomorphism and vectors of the form $|v\rangle \otimes |x\rangle$ span $\mathcal{H}_{kn}$, we conclude that $\Phi(AB) = \Phi(A)\Phi(B)$ when AB is defined.

2. $\Phi(A) + \Phi(B) = \sum_{i=0}^{d-1} A_i \otimes \Lambda^i + \sum_{i=0}^{d-1} B_i \otimes \Lambda^i = \sum_{i=0}^{d-1} (A_i + B_i) \otimes \Lambda^i = \Phi(A + B)$.

3. $\Phi(I_n) = \Phi(I_n \alpha^0) = I_n \otimes \Lambda^0 = I_{nk}$.

4. Because Λ is normal, if λ is an eigenvalue of Λ with eigenvector $|x\rangle$, then $\lambda^\dagger$ is an eigenvalue of $\Lambda^\dagger$ with eigenvector $|x\rangle$. We have

$$M^\dagger = \sum_{i=0}^{d-1} M_i^\dagger (\alpha^i)^\dagger \quad \text{and} \quad M^\dagger = \sum_{i=0}^{d-1} N_i \alpha^i$$

for some unique $N_i \in \mathcal{R}$. If $|v\rangle \in \mathcal{H}_n$ is some arbitrary state and σ is defined as above,

$$\Phi(M)^\dagger(|v\rangle \otimes |x\rangle) = \sum_{i=0}^{d-1} M_i^\dagger |v\rangle \otimes (\Lambda^i)^\dagger |x\rangle = \sum_{i=0}^{d-1} M_i^\dagger |v\rangle \otimes (\lambda^i)^\dagger |x\rangle = \sum_{i=0}^{d-1} M_i^\dagger |v\rangle \otimes \sigma(\alpha^i)^\dagger |x\rangle = (\sigma(M^\dagger) |v\rangle) \otimes |x\rangle.$$

On the other hand,

$$\Phi(M^\dagger)(|v\rangle \otimes |x\rangle) = \sum_{i=0}^{d-1} N_i |v\rangle \otimes \Lambda^i |x\rangle = \sum_{i=0}^{d-1} N_i |v\rangle \otimes \lambda^i |x\rangle = \sum_{i=0}^{d-1} N_i |v\rangle \otimes \sigma(\alpha)^i |x\rangle = \sigma(M^\dagger) |v\rangle \otimes |x\rangle.$$

We conclude that $\Phi(M)^\dagger = \Phi(M^\dagger)$, since vectors of the form $|v\rangle \otimes |x\rangle$ span $\mathcal{H}_{kn}$.

5. $\Phi(C \otimes A) = \Phi\left(\sum_{i=0}^{d-1} C \otimes A_i \alpha^i\right) = \sum_{i=0}^{d-1} C \otimes A_i \otimes \Lambda^i = C \otimes \Phi(A)$.

6. Because Λ is a pseudo-companion matrix for p, it has α as an eigenvalue. Let Π be the projector of onto the eigenspace of Λ corresponding α. Then

$$\Phi(A)(I \otimes \Pi) = \left(\sum_{i=0}^{d-1} A_i \otimes \Lambda^i\right)(I \otimes \Pi) = \sum_{i=0}^{d-1} A_i \otimes (\alpha^i \Pi) = \sum_{i=0}^{d-1} A_i \alpha^i \otimes \Pi = \Phi(A) \otimes \Pi. \qquad \square$$

Corollary 2. *Let $\mathcal{R}$, α, Λ, and Φ be defined as in Proposition 10. If $U \in \mathcal{U}(\mathcal{R}[\alpha])$, then $\Phi(U) \in \mathcal{U}(\mathcal{R})$.*

Proof. By Proposition 10, $\Phi(U^\dagger) = \Phi(U)^\dagger$, $\Phi(UV) = \Phi(U)\Phi(V)$, and $\Phi(I) = I_n$ for some n. Suppose $U \in \mathcal{U}(\mathcal{R}[\alpha])$, then

$$\Phi(U)\Phi(U)^\dagger = \Phi(U)\Phi(U^\dagger) = \Phi(UU^\dagger) = \Phi(I) = I_n,$$

and $\Phi(U)$ is unitary. $\qquad \square$

Definition 18 (Standard catalytic embedding)**.** Let $\mathcal{R}$, α, Λ, and Φ be defined as in Proposition 10, and let ϕ be the catalytic embedding obtained from the lifting of $\Phi|_{\mathcal{U}(\mathcal{R}[\alpha])}$. We call this embedding a *standard* catalytic embedding.

The difficulty in constructing a standard catalytic embedding hinges on finding a normal pseudo-companion matrix over a given number ring. On the surface, this might not seem immediately difficult. Each polynomial over a number ring automatically admits a companion matrix, and hence a pseudo-companion matrix. The normality condition ensures unitarity of the resulting standard catalytic embedding, and is where any difficulty in constructing such an embedding lies. In fact, generalizing to pseudo-companion matrices (rather than simply companion matrices) is *necessary* to accommodate the normality condition, as the following remark shows.

Remark 5. An extension of $\mathbb{Q}$ by a totally real algebraic number θ yields a totally real number field, and this extension $\mathbb{Q}[\theta]$ necessarily satisfies the assumptions of Theorem 3. However, there exist θ whose minimal polynomials do not permit a normal companion matrix with entries in $\mathbb{Q}$ [22]. An example of such a polynomial is $X^2 - 3$, because 3 cannot be written as a sum of two squares from $\mathbb{Q}$. In fact, this is a relatively universal phenomenon, as randomly selecting an irreducible quadratic polynomial over $\mathbb{Q}$ will yield a polynomial p that does not permit a normal companion matrix with high probability. Using pseudo-companion matrices, rather than the usual companion matrices, skirts this issue.

We also might wonder under what conditions a generic number ring extension $\mathcal{R} \subset \mathcal{S}$ permits an $\alpha \in \mathcal{S}$ as in Proposition 10 so that $\mathcal{S} = \mathcal{R}[\alpha]$. As it turns out, these conditions are rather strong. Nonetheless, constructing ring extensions in this way coincides with many ring extensions of practical use. For example, quadratic extensions and extensions by roots of unity generically have the required structure.

Example 10. We show how to implement a fifth root of unity $\alpha = e^{2\pi i/5}$ over the Clifford$+T$ gate set using a standard catalytic embedding. The Clifford$+T$ gate set has a characterization given by $\mathcal{U}(\mathcal{R})$ where $\mathcal{R} = \mathbb{Z}[1/2, \sqrt{2}, i]$ [2]. Therefore, we need to construct a matrix whose characteristic polynomial is a power of the minimal polynomial of α over $\mathcal{R}$. The minimal polynomial of α over $\mathcal{R}$ is $p(x) = x^4 + x^3 + x^2 + x + 1$. The matrix

$$\Lambda = \frac{1}{2} \begin{bmatrix} -1+i & 1 & 0 & i \\ 1 & i & i & i \\ 0 & i & -1-i & 1 \\ i & i & 1 & -i \end{bmatrix}$$

is normal, is a pseudo-companion matrix for p, and has entries in $\mathcal{R}$. The mapping

$$U = \sum_{k=0}^{4} A_k e^{k\frac{2\pi i}{5}} \mapsto \sum_{k=0}^{4} A_k \otimes \Lambda^k$$

can then be lifted to a standard catalytic embedding of $\mathcal{C}(\mathcal{R}[\alpha])$ into $\mathcal{C}(\mathcal{R})$. The catalyst for this embedding is the eigenvector of Λ corresponding to the eigenvalue α.

Proposition 11. *Let $\mathcal{R} \subset \mathcal{R}[\alpha] \subset \mathcal{R}[\beta]$ be integral Kroneckerian extensions. If $\Phi_1 : \mathcal{M}(\mathcal{R}[\beta]) \to \mathcal{M}(\mathcal{R}[\alpha])$ and $\Phi_2 : \mathcal{M}(\mathcal{R}[\alpha]) \to \mathcal{M}(\mathcal{R})$ are pre-embeddings for standard catalytic embeddings, then $\Phi = \Phi_2 \circ \Phi_1$ is a pre-embedding for a standard catalytic embedding.*

Proof. Let p be the minimal polynomial for β over $\mathcal{R}$. Let $\Lambda = \Phi_2 \circ \Phi_1(\beta) \in \mathcal{M}(\mathcal{R})$. Then $0 = \Phi_2 \circ \Phi_1(p(\beta)) = p(\Phi_2 \circ \Phi_1(\beta)) = p(\Lambda)$. Because p is irreducible over $\mathcal{R}$, the characteristic polynomial of Λ must be a power of p, so Λ is a pseudo-companion matrix for p. Because Φ_2 and Φ_1 are pre-embeddings,

$$\Lambda\Lambda^{\dagger} = \Phi_2 \circ \Phi_1(\beta)\Phi_2 \circ \Phi_1(\beta^{\dagger}) = \Phi_2 \circ \Phi_1(\beta\beta^{\dagger}) = \Phi_2 \circ \Phi_1(\beta^{\dagger}\beta) = \Phi_2 \circ \Phi_1(\beta^{\dagger})\Phi_2 \circ \Phi_1(\beta) = \Lambda^{\dagger}\Lambda,$$

thus Λ is normal. Let $M \in \mathcal{M}(\mathcal{R}[\beta])$. There exist M_i such that $M = \sum_{i=1}^{d-1} M_i\beta^i$ where d is the degree of p. By linearity of Φ_1 and Φ_2,

$$\Phi_2 \circ \Phi_1(M) = \sum_{i=1}^{d-1} M_i \otimes \Phi_2 \circ \Phi_1(\beta)^i = \sum_{i=1}^{d-1} M_i \otimes \Lambda^i.$$

Thus $\Phi = \Phi_2 \circ \Phi_1$ is a pre-embedding of a standard catalytic embedding. $\qquad\square$

Example 11. In Example 10, while Λ is, in fact, a companion matrix for p, it is not immediately obvious how to arrive at Λ given p. Here we show how to apply Proposition 11 to construct Λ using an intermediate ring extension.

The ring $\mathcal{R}[\cos(2\pi/5)]$ sits between $\mathcal{R}[\alpha]$ and $\mathcal{R}$. First, we find a matrix Λ_1 for the embedding from $\mathcal{U}(\mathcal{R}[\alpha])$ to $\mathcal{U}(\mathcal{R}[\cos(2\pi/5)])$. The minimal polynomial of α over $\mathcal{R}[\cos(2\pi/5)]$ is $q(x) = x^2 + 2\cos(2\pi/5)x + 1$. Our matrix must therefore have α and $\overline{\alpha}$ as its only eigenvalues. The matrix

$$\frac{1}{2} \begin{bmatrix} 1 & 1+2\cos(2\pi/5) \\ 1+2\cos(2\pi/5) & -1 \end{bmatrix}$$

has eigenvalues $\pm\sin(2\pi/5)$ (since $\sin^2(2\pi/5) = (1/2)^2 + (1/2+\cos(2\pi/5))^2$). Then, by multiplying by i and adding $\cos(2\pi/5) \cdot I$ to this matrix, we obtain the matrix

$$\Lambda_1 = \frac{i}{2} \begin{bmatrix} 1 & 1+2\cos(2\pi/5) \\ 1+2\cos(2\pi/5) & -1 \end{bmatrix} + \begin{bmatrix} \cos(2\pi/5) & 0 \\ 0 & \cos(2\pi/5) \end{bmatrix} = \frac{1}{2} \begin{bmatrix} i+2\cos(2\pi/5) & i+2i\cos(2\pi/5) \\ i+2i\cos(2\pi/5) & -i+2\cos(2\pi/5) \end{bmatrix},$$

whose eigenvalues are $\cos(2\pi/5) \pm i\sin(2\pi/5) = \alpha, \overline{\alpha}$. Consequently, Λ_1 has the characteristic polynomial, q.

Next, we find a matrix Λ_2 for the embedding from $\mathcal{U}(\mathcal{R}[\cos(2\pi/5)])$ to $\mathcal{U}(\mathcal{R})$. The minimal polynomial of $\cos(2\pi/5)$ over $\mathcal{R}$ is $r(x) = x^2 + \frac{1}{2}x + \frac{1}{4}$. The matrix

$$\Lambda_2 = \frac{1}{2} \begin{bmatrix} -1 & 1 \\ 1 & 0 \end{bmatrix}$$

has r as its characteristic polynomial.

Finally, we embed Λ_1 using Λ_2 in order to obtain

$$\Lambda_1 = \frac{1}{2}\begin{bmatrix} i & i \\ i & -i \end{bmatrix} + \begin{bmatrix} 1 & i \\ i & 1 \end{bmatrix} \cdot \cos(2\pi/5) \longmapsto \frac{1}{2}\begin{bmatrix} i & i \\ i & -i \end{bmatrix} \otimes I_2 + \begin{bmatrix} 1 & i \\ i & 1 \end{bmatrix} \otimes \Lambda_2 = \frac{1}{2}\begin{bmatrix} -1+i & 1 & 0 & i \\ 1 & i & i & i \\ 0 & i & -1-i & 1 \\ i & i & 1 & -i \end{bmatrix} = \Lambda.$$

If $|\psi_1\rangle$ is the eigenvector of Λ_1 corresponding to α and $|psi_2\rangle$ is the eigenvector of Λ_2 corresponding to $\cos(2\pi/5)$, then the catalyst for Λ is given by $|\psi_1\rangle \otimes |\psi_2\rangle$.

Theorem 3. *Let $\mathcal{R} \subset \mathcal{R}[\alpha]$ be a Kroneckerian, integral extension of number rings such that α has a minimal polynomial over $\mathrm{Frac}(\mathcal{R})$ with coefficients in $\mathcal{R}$. If $\phi : \mathcal{C}(\mathcal{R}[\alpha]) \to \mathcal{C}(\mathcal{R})$ is a linear catalytic embedding, then ϕ is a standard catalytic embedding.*

Proof. Let Φ be the pre-embedding of ϕ. The set $\{\alpha^j\}_{j=0}^k$ is a generating set for $\mathcal{R}[\alpha]$ over $\mathcal{R}$ for some k. Let $\Phi(\alpha) = \Lambda$ and $p \in \mathcal{R}$ be the minimal polynomial of α. By Theorem 1, Λ is normal and $p(\Lambda) = 0$. Because the characteristic polynomial of Λ has coefficients in $\mathrm{Frac}(\mathcal{R})$ and is divisible by p with p irreducible over $\mathrm{Frac}(\mathcal{R})$, the characteristic polynomial of Λ is a plus or minus a power of p. Therefore, Λ is a normal pseudo-companion matrix for p. Thus, by Proposition 9 Φ is completely determined by its action on $\{\alpha^j\}_{j=0}^k$, and so Φ is the pre-embedding of a standard catalytic embedding. $\square$

Standard catalytic embeddings are useful because they reduce the problem of constructing linear catalytic embeddings to the problem of finding normal pseudo-companion matrices for irreducible polynomials. The latter problem is, in general, simpler. When the rings in question are actually fields, there is an explicit method for constructing such matrices, and thus an explicit methods for building standard catalytic embeddings [23].

B. Circuits Including Order-3 Z-Rotations

One can obtain explicit reductions in the resources needed for fault-tolerant quantum computing using standard catalytic embeddings. We consider computations over the Clifford+T gates set, a gate set commonly used in fault-tolerant quantum computing. Let ω_3 be a third root of unity and E be the rotation

$$E := \begin{bmatrix} 1 & 0 \\ 0 & \omega_3 \end{bmatrix}.$$

It was shown in [2] that the Clifford+T gate set corresponds to the set of unitary operations $\mathcal{U}(\mathbb{D}[\omega_8])$ where ω_8 is an eighth root of unity and $\mathbb{D} = \mathbb{Z}[1/2]$. Since $\omega_3 \notin \mathbb{D}[\omega_8]$, it follows that E cannot be implemented directly by a circuit over Clifford+T.

While quantum computation over Clifford+T+E may not be typical, in some circumstances it may be useful to extend the Clifford+T gate set with a phase gate of order 3. In the standard approach, one would implement the E gate by approximation over Clifford+T. With embeddings, we can instead directly implement E using Clifford+T gates by embedding $\mathcal{U}(\mathbb{D}[\omega_8, \omega_3])$ in $\mathcal{U}(\mathbb{D}[\omega_8])$.

Because $\omega_3 \notin \mathbb{D}[\omega_8]$ and $\omega_3^2 + \omega_3 + 1 = 0$, we need to find a normal, pseudo-companion matrix for the polynomial $x^2 + x + 1$ with entries in the ring $\mathbb{D}[\omega_8]$. Explicitly, we can observe that

$$\Lambda = \frac{1}{2}\begin{bmatrix} -1-i & 1-i \\ -1-i & -1+i \end{bmatrix}$$

is a normal pseudo-companion matrix for ω_3 over $\mathbb{D}[\omega_8]$ with $i = \omega_8^2$, giving a standard catalytic embedding. Note also that

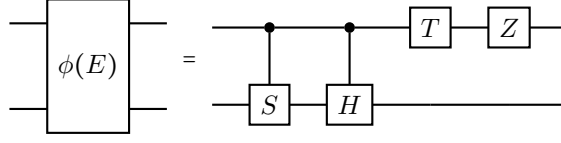$$|v\rangle = \frac{1}{\sqrt{3+\sqrt{3}}}\begin{bmatrix} -\omega_3 - i\omega_3^2 \\ 1 \end{bmatrix}$$

is in the ω_3 eigenspace of Λ and the corresponding projector is given by $|v\rangle\langle v|$. Finally, we note that

$$\Lambda = \omega_8^5 HS.$$

Since $E = |0\rangle\langle 0| + \omega_3 |1\rangle\langle 1|$, we have

$$\phi(E) = |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes \Lambda = |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes (\omega_8^5 HS).$$

In circuit form, the right hand side amounts to a controlled $\omega_8^5 HS$ gate, as synthesized below (a controlled ω_8^5 gate can be implemented as a $T^5 = ZT$ gate):

Using known constructions for the controlled-H and -S gates [24] gives a Clifford+T implementation of $\phi(E)$ using 6 T gates, which can be further reduced to 4 T gates using standard techniques (e.g. [25]).

For circuits with a large number of E gates, this gives a significant reduction in T-count compared with the usual method of repeated approximations. In particular, for a circuit with m E gates, assuming a practical overall precision of $\varepsilon = 10^{-15}$ [26], and using asymptotically optimal Clifford+T approximations [8] of E, this gives a T-count of

$$m \cdot 3 \log_2(m/\varepsilon) \approx 3m \cdot (50 + \log_2 m).$$

By comparison, using approximations to prepare the single-qubit catalyst $|v\rangle$, we get a T-count of

$$6 \log_2(1/\varepsilon) + 4m \approx 300 + 4m$$

with the given embedding. In the limit of large m, the ratio of these two costs has an asymptotic scaling of

$$\frac{4}{3 \cdot (50 + \log_2 m)}.$$

Ignoring the $\log_2 m$ term, this reduces the T-count by 97% compared to the standard approach. Including this term for a reasonable number of E gates (say $m = 2^{20} \approx 10^6$), this reduces the T-count by 98% over the standard method. Were we to consider the arbitrary precision limit, this value becomes arbitrarily close to 100%. This highlights the power of catalytic embeddings to reduce gate counts in practice.

C. The Quantum Fourier Transform

The quantum Fourier transform (QFT) on n qubits is the unitary operation given by the matrix $\frac{1}{2^{n/2}}[\omega^{jk}]_{j,k=0}^{n-1}$ where $\omega = e^{2\pi i/2^n}$. It is well-known that the QFT can be realized as the circuit below, where R_k is the 2×2 diagonal matrix $R_k = \mathrm{diag}(1, e^{2\pi i/2^k})$ [15].

We construct a standard catalytic embedding to implement this circuit using the gate set $\langle H, X, CX, CCX \rangle$. We then show how to use n additional X gates and the same embedded circuit to implement the inverse quantum Fourier transform.

Our goal is to reduce the cost of the expensive Z-rotations. To do so, we will find an embedding $\phi : \mathcal{U}(\mathbb{Z}[e^{2\pi i/2^n}]) \to \mathcal{U}(\mathbb{N})$. We first construct a sequence of standard catalytic embeddings

$$\mathcal{U}(\mathcal{R}_n) \to \mathcal{U}(\mathcal{R}_{n-1}) \to \cdots \to \mathcal{U}(\mathcal{R}_k) \to \mathcal{U}(\mathcal{R}_{k-1}) \to \cdots \to \mathcal{U}(\mathcal{R}_2) \to \mathcal{U}(\mathcal{R}_1)$$
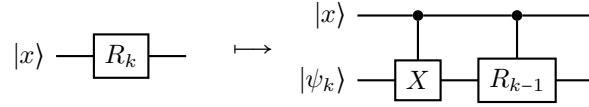
where $\mathcal{R}_k = \mathbb{Z}[e^{2\pi i/2^k}]$. Let us begin with $\phi_k : \mathcal{U}(\mathcal{R}_k) \to \mathcal{U}(\mathcal{R}_{k-1})$. The minimal polynomial for $e^{2\pi i/2^k}$ over $\mathcal{R}_{k-1}$ is $p_k(x) = x^2 - e^{2\pi i/2^{k-1}}$. Define

$$\Lambda_k = \begin{bmatrix} 0 & 1 \\ e^{2\pi i/2^{k-1}} & 0 \end{bmatrix} \qquad \text{and} \qquad |\psi_k\rangle = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ e^{2\pi i/2^k} \end{bmatrix}.$$
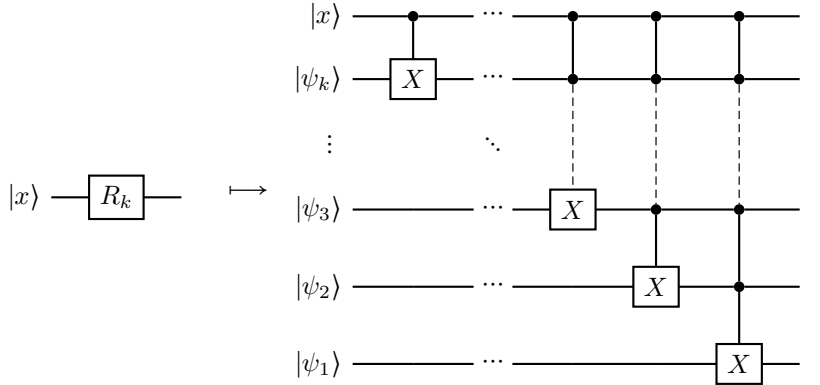
The matrix Λ_k is a normal companion matrix for p_k and has $|\psi_k\rangle$ as its corresponding catalyst. We have the following embedding of R_k:

$$\begin{bmatrix} 1 & 0 \\ 0 & e^{2\pi i/2^k} \end{bmatrix} \longmapsto \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & e^{2\pi i/2^{k-1}} & 0 \end{bmatrix}$$
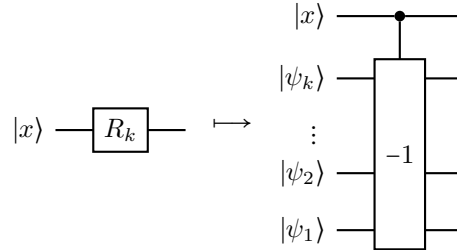
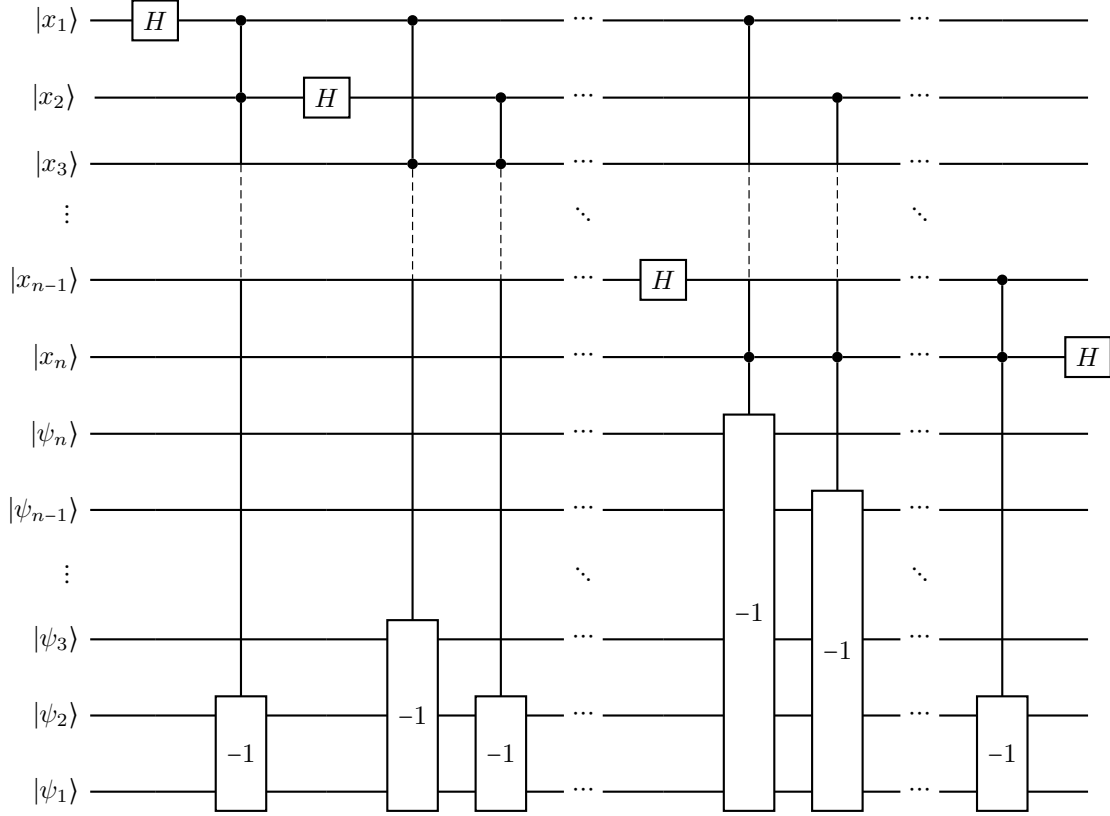Using controlled gates, we can represent this embedding diagrammatically as below.

Following this sequence of embeddings to its end, we obtain the following following embedding for R_k:

In our presentation above, we have gone one step further and embedded R_1 (the Z gate) as CX, which we note is not a linear catalytic embedding (because $\mathbb{N}$ is not a number ring) but nonetheless constitutes a catalytic embedding. On computational basis states, the action of the circuit above can be described as a controlled decrementer, where $|x\rangle$ is the control and the register $|\psi_1\rangle \otimes \cdots \otimes |\psi_k\rangle$ holds the target integer (represented as a bitstring for computational basis states). We therefore simplify our notation and write the embedding as below.

By Proposition 7, linear catalytic embeddings are well-behaved with respect to direct sums, and so the embedding of a controlled operation is a controlled operation. Similarly, despite not being a linear catalytic embedding the final catalytic embedding also respects direct sum structure. Putting all this together, we get the following embedding of the quantum Fourier transform:

Each of these controlled-decrement circuits can be implemented with X, CX, and CCX gates [19], and so we have implemented the quantum Fourier transform using the gate set $\langle H, X, CX, CCX \rangle$ in the presence of catalysts.

Compiled as a circuit over Clifford+T instead, the above implementation of the quantum Fourier transform is equivalent to the T-count efficient circuit given in [27, 28]. Each individual decrement can be seen as an inverse adder controlled by the bottom control qubit and subtracting the top control qubit from the ancilla qubits. Each sequence of decrements can be seen as an inverse adder taking as input the binary number represented by the top control bits and subtracting it from the ancilla bits. As an n-bit adder can be implemented with linear T complexity, and noting that $|\psi_k\rangle = R_k H |0\rangle$, the entire circuit may be implemented over Clifford+T with T count

$$O(n^2 + n \log_2(1/\epsilon))$$

compared to the standard approach of approximation, which would require T-count

$$O(n^2 \log_2(1/\epsilon)).$$

While this implementation of the QFT has been previously derived using phase gradients [27, 28], we have shown that the more general framework of catalytic embeddings suffices to reproduce it.

While it might appear that catalytic embeddings have merely reproduced the best known constructions of the QFT, we can in fact glean additional insight into the structure of those constructions. In light of Corollary 1, we know that the catalyst $|\psi_1\rangle \otimes \cdots \otimes |\psi_n\rangle$ produced in this construction should have orthogonal counterparts corresponding to alternative embeddings of $e^{2\pi i/2^n}$ in $\mathbb{C}$. For example, by applying an X gate to each $|\psi_k\rangle$ we see that

$$X \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ e^{2\pi/2^k} \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} e^{2\pi/2^k} \\ 1 \end{bmatrix} \sim \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ e^{-2\pi/2^k} \end{bmatrix}.$$

The resulting state is orthogonal to the original state since $\langle \psi_2 | X | \psi_2 \rangle = 0$, and by inspection such a tensor product of X gates induces the complex conjugation automorphism on ω. Therefore, $X |\psi_1\rangle \otimes \cdots \otimes X |\psi_n\rangle$ is one such alternative catalyst, and it is precisely such that it maps each ω in the QFT unitary to $\omega^\dagger$, inducing the complex conjugation automorphism on the circuit which happens to be equivalent to the inverse QFT. Thus, using the *same* embedded circuit along with n X gates, we can implement the inverse QFT.

VII. PERSPECTIVES

In this paper, we laid the foundations for the theory of catalytic embeddings. We believe that catalytic embeddings may find a variety of applications in the study of quantum circuits and, more generally, in the theory and practice of fault-tolerant quantum computation. As discussed in Section VI, there are cases where catalytic embeddings reproduce or beat existing quantum circuit constructions for specific operations. We are eager to see what other algorithmic primitives can be improved with catalytic embeddings. Approximate and exact synthesis methods also seem like prime candidates to bolster with the power of catalytic embeddings. To make the most of this framework, it is important to provide constructive methods for producing catalytic embeddings. In follow-up work [23], we provide such constructive methods in many cases of interest.

The structure-preserving nature of catalytic embeddings may provide insights into a number of open questions. Firstly, one may be able to use catalytic embeddings to better understand gate sets. By embedding a poorly understood gate set into a well-understood one (such as the Toffoli-Hadamard gate set), one could in principle transform results about the latter into results about the former. This approach may help in characterizing gate sets, finding relations for circuits, and deriving asymptotic lower bounds for resources. Further afield, catalytic embeddings seem to be a natural tool with which to tackle long-standing open questions about the Clifford hierarchy. Indeed, catalytic embeddings were (in part) born out of generalizing gate teleportation protocols. Even farther afield, there seems to be a growing body of evidence that various approaches to achieving fault-tolerant quantum computation share important properties. Catalytic embeddings may help in understanding what unifies these different approaches.

VIII. ACKNOWLEDGEMENTS

[1] V. Kliuchnikov, D. Maslov, and M. Mosca, Quantum Information & Computation **13**, 607 (2013), arXiv:1206.5236.
[2] B. Giles and P. Selinger, Physical Review A **87**, 032332 (2013), arXiv:1212.0506.
[3] S. Forest, D. Gosset, V. Kliuchnikov, and D. McKinnon, Journal of Mathematical Physics **56**, 082201 (2015), arXiv:1501.04944.
[4] M. Amy, A. N. Glaudell, and N. J. Ross, Quantum **4**, 252 (2020), arXiv:1908.06076.
[5] B. Giles and P. Selinger, arXiv preprint arXiv:1312.6584 (2013).
[6] A. N. Glaudell, N. J. Ross, and J. M. Taylor, npj Quantum Information **7**, 1 (2021), arXiv:2001.05997.
[7] M. Mosca and P. Mukhopadhyay, Quantum Science and Technology **7**, 10.1088/2058-9565/ac2d3a (2021), arXiv:2006.12440.
[8] N. J. Ross and P. Selinger, Quantum Information & Computation **16**, 901 (2016), arXiv:1403.2975.
[9] V. Kliuchnikov, A. Bocharov, M. Roetteler, and J. Yard, arXiv preprint arXiv:1510.03888 (2015).
[10] D. Gottesman and I. L. Chuang, arXiv preprint arXiv:quant-ph/9908010 (1999).
[11] S. Aaronson, D. Grier, and L. Schaeffer, arXiv preprint arXiv:1504.05155 (2015).
[12] D. Grier and L. Schaeffer, Quantum **6**, 734 (2022), arXiv:1603.03999.
[13] M. Artin, Algebra (Pearson Education, 2011).
[14] A. Schinzel, Polynomials with Special Regard to Reducibility, Encyclopedia of Mathematics and its Applications (Cambridge University Press, 2000).
[15] M. A. Nielsen and I. L. Chuang, Quantum Computation and Quantum Information, Cambridge Series on Information and the Natural Sciences (Cambridge University Press, 2000).
[16] S. Lane, Categories for the Working Mathematician, Graduate Texts in Mathematics (Springer, New York, NY, USA, 1998).
[17] P. Selinger, A survey of graphical languages for monoidal categories, in New Structures for Physics, edited by B. Coecke (Springer Berlin Heidelberg, Berlin, Heidelberg, 2011) pp. 289–355, arXiv:0908.3347.
[18] C. Heunen and J. Vicary, Categories for Quantum Theory: An Introduction (Oxford University Press, Oxford, UK, 2019).
[19] V. Shende, A. Prasad, I. Markov, and J. Hayes, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems **22**, 710 (2003), arXiv:quant-ph/0207001.
[20] M. Beverland, E. Campbell, M. Howard, and V. Kliuchnikov, Quantum Science and Technology **5**, 035009 (2020), arXiv:1904.01124.
[21] D. S. Dummit and R. M. Foote, Abstract algebra, Vol. 3 (Wiley Hoboken, 2004).
[22] G. Schmeisser, Linear Algebra and its Applications **193**, 11 (1993).
[23] M. Amy, M. Crawford, A. N. Glaudell, M. L. Macasieb, S. S. Mendelson, and N. J. Ross, Forthcoming (2023).
[24] M. Amy, D. Maslov, M. Mosca, and M. Roetteler, Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on **32**, 818 (2013), arXiv:1206.0758.

[25] M. Amy, D. Maslov, and M. Mosca, IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems **33**, 1476 (2014), arXiv:1303.2042.

[26] V. Kliuchnikov, D. Maslov, and M. Mosca, Phys. Rev. Lett. **110**, 190502 (2013), arXiv:1212.0822.

[27] C. Gidney, Turning gradients into additions into qfts (2016).

[28] Y. Nam, Y. Su, and D. Maslov, npj Quantum Information **6**, 1 (2020), arXiv:1803.04933.

[29] A. Kay, arXiv preprint arXiv:1809.03842 (2018).