# PCOAST: A Pauli-based Quantum Circuit Optimization Framework

Jennifer Paykin\*<sup>‡</sup>, Albert T. Schmitz\*<sup>‡</sup>, Mohannad Ibrahim\*, Xin-Chuan Wu<sup>†</sup>, and A. Y. Matsuura\*

\* Intel Labs, Intel Corporation, Hillsboro, OR, USA

<sup>†</sup> Intel Labs, Intel Corporation, Santa Clara, CA, USA

<sup>‡</sup> These authors contributed equally. Email: {jennifer.paykin, albert.schmitz}@intel.com

Abstract—This paper presents the Pauli-based Circuit Optimization, Analysis, and Synthesis Toolchain (PCOAST), a framework for quantum circuit optimizations based on the commutative properties of Pauli strings. Prior work has demonstrated that commuting Clifford gates past Pauli rotations can expose opportunities for optimization in unitary circuits. PCOAST extends that approach by adapting the technique to mixed unitary and non-unitary circuits via generalized preparation and measurement nodes parameterized by Pauli strings. The result is the PCOAST graph, which enables novel optimizations based on whether a user needs to preserve the quantum state after executing the circuit, or whether they only need to preserve the measurement outcomes. Finally, the framework adapts a highly tunable greedy synthesis algorithm to implement the PCOAST graph with a given gate set.

PCOAST is implemented as a set of compiler passes in the Intel<sup>®</sup> Quantum SDK. In this paper, we evaluate its compilation performance against two leading quantum compilers, Qiskit and t|ket⟩. We find that PCOAST reduces total gate count by 32.53% and 43.33% on average, compared to the best performance achieved by Qiskit and t|ket⟩ respectively, two-qubit gates by 29.22% and 20.58%, and circuit depth by 42.02% and 51.27%.

Index Terms—Quantum Compiler, Quantum Circuit Optimization, Pauli Optimization, Classical-Quantum state

#### I. INTRODUCTION

Quantum circuit optimizations address an important challenge in the design of efficient quantum computing systems by reducing the number of operations required to execute quantum algorithms [1, 2, 3]. Optimizations fall in two main classes: local, peephole-style optimizations [4, 5, 6, 7, 8, 9], where local patterns of gates are replaced by other patterns; and global optimizations, where circuits are converted to an intermediate mathematical structure that highlights some semantic equivalence, simplified according to the rules of that structure, and synthesized back into a circuit that could be significantly different from the original. This paper presents a novel global optimization framework called PCOAST, a Pauli-based Circuit Optimization, Analysis, and Synthesis Toolchain, which successfully reduces total gate count, twoqubit gate count, and depth compared to the best performance of state-of-the-art optimizing compilers Qiskit [7] and t|ket) [10].

Recently, a class of global optimizations based on Pauli rotations have proved successful in reducing gate count for unitary circuits beyond the reach of more traditional local peephole optimizations [11, 12, 13]. These optimizations take advantage of the fact that unitary circuits can be decomposed into Clifford gates (generated by the Hadamard gate H, the phase gate S and the controlled-not gate CNOT); and non-Clifford gates represented by Pauli rotations  $Rot(P,\theta) = e^{-i\theta/2P}$ , where P is one the Pauli matrices X, Y, or Z. Because Cliffords always map Paulis to Paulis by conjugation, it is possible to push Clifford unitaries U past the non-Clifford Pauli rotations  $Rot(P,\theta)$  to produce a new rotation  $Rot(UPU^{\dagger},\theta)$ . Doing so can expose the fact that some rotations can be merged, reducing the number of gates required in the final circuit, even if the original gates did not appear directly next to each other in the original circuit [11]. The implementation of Pauli gadgets (the equivalent of Pauli rotations) in the t  $|ket\rangle$ compiler [10] has been shown to reduce both the number of 2-qubit entangling gates, relevant for Noisy Intermediate Scale Quantum (NISQ) workflows, and single-qubit non-Clifford gates, relevant for fault-tolerant workflows [12].

These optimizations are quite powerful, but they fall short in two main ways. First, they focus on unitary circuits, so equivalences enabled by non-unitary operations like preparations and measurements are not taken into account. For example, if a Z rotation occurs on a qubit that was prepared in a Zeigenstate, that rotation can be eliminated. Incorporating nonunitary gates enables several powerful optimizations, such as the ability to drop all unitary gates after measurement when coherence on the quantum device doesn't need to be preserved after measurement. Second, while optimizations reduce the number of rotations and put an upper bound on the size of the resulting circuit, the number of two-qubit gates required can be expensive if the rotations are synthesized in a suboptimal order. Schmitz et al. [13] and Li et al. [14] both explore synthesis in the context of Hamiltonian simulation, with the goal of automatically synthesizing a Hamiltonian, described as a product of Pauli rotations, into a circuit with a designated gate set. In particular, Schmitz et al. describe a greedy search algorithm that reduces this search problem to a variation of the traveling salesman problem [13].

In this paper, we show that by addressing these deficits, Pauli strings can be used not only as a standalone optimization pass, but as a cohesive optimization framework, PCOAST. Fig. 1 shows an example PCOAST workflow. First, the circuit in Fig. 1a is converted into the PCOAST graph in Fig. 1b. The graph highlights the commutativity of rotation nodes in relation to each other—if there is no dependency between two



(a) Example circuit. The c argument in MeasZ<sup>c</sup> indicates the classical variable that the measurement outcome is written to.



(b) The PCOAST graph generated by the example circuit. The groupings indicate nodes to be merged together, where  $\operatorname{Rot}(X_0, \theta_1)$  and  $\operatorname{Rot}(X_0, \theta_2)$  combine to  $\operatorname{Rot}(X_0, \theta_1 + \theta_2)$ , and  $\operatorname{Rot}(Z_1, \theta_3)$  is absorbed by  $\operatorname{Prep}(Z_1, X_1)$ . Note that the measurement to variable  $c_1$  has been transformed into a measurement of the first qubit,  $Z_0$ , due to the permutations of Clifford gates (F) past the measurement.



(c) The optimized PCOAST graph obtained when specifying a release outcome. The optimization has reduced the support of the measurement node  $\text{Meas}(Z_0X_1)$  to  $\text{Meas}(X_1)$  by recognizing that the  $Z_0X_1$  measurements can be reconstructed by measuring  $X_1$  and combining the outcomes with the measurement of  $Z_0$  classically. The measurement results of the optimized graph is guaranteed to produce the same probability distribution as Fig. 1b.



(d) Optimized released circuit synthesized from Fig. 1c, along with assignments to classical variables to account for the release outcome optimization.

Fig. 1: PCOAST optimization on an example circuit.

nodes, they commute with each other. In addition to the unitary Pauli rotations, our PCOAST graph contains non-unitary gates—preparation and measurement—that are parameterized by Pauli rotations and are subject to the same commutativity rules as rotations. To represent Cliffords, we use a compact representation as a Pauli frame F, otherwise known as a Pauli tableau [15], which emphasizes the behavior of the Clifford on Pauli arguments (Sec. IV-A). A summary of different types of nodes is shown in Fig. 2.

The addition of these non-unitary nodes enables a host of additional internal optimizations on PCOAST graphs. Users can choose between two optimization outcomes: either a *hold outcome*, where the optimizations preserve the semantics of the original circuit precisely; or a *release outcome*, where more aggressive optimizations can be applied as long as they produce the same measurement results. A release outcome will drop unitary gates that can be delayed until after measurement, with the guarantee that the measurement results will always be statistically equivalent. To achieve this, we introduce classical



Fig. 2: Types of PCOAST nodes in relation to each other. Note that Paulis themselves are not nodes, but are represented in PCOAST as Pauli frames.

remappings of measurement variables via what we call measurement space functions. For example, if a *release* outcome is specified for Fig. 1, it will be optimized to the PCOAST graph as shown in Fig. 1c using the measurement space function given in the bottom right.

Finally, we adapt Schmitz et al. [13]'s synthesis algorithm to determine both how to order commuting nodes, and how to decompose multi-qubit nodes using sequences of twoqubit entangling gates. The result is shown in Fig. 1d. We customize the synthesis algorithm with a number of heuristics to minimize cost according to a given cost model, map into a target gate set, and reduce the number of measurements required for a stabilizer search. Currently, synthesis primarily aims to minimize algorithm-level resource requirements like circuit depth, although the design allows for customization to prioritize other search criteria.<sup>1</sup>

This work makes the following contributions:

- We develop a semantics in which to describe the behavior of PCOAST nodes, terms, and graphs that incorporates both classical and quantum states.
- We introduce the key PCOAST data structures, including Pauli frames and the PCOAST graph.
- We present three major components of the PCOAST optimization: compiling a circuit to a PCOAST graph, optimizing the graph, and synthesizing a circuit back out.
- We implement PCOAST in C++ as a sequence of compiler passes in the Intel<sup>®</sup> Quantum Software Development Kit (SDK)<sup>2</sup> [16], and evaluate its compilation performance against two state-of-the-art optimizing quantum compilers, Qiskit [7] and t|ket> [10]. Our experimental results show that PCOAST reduces total gate count by 32.53% and 33.33% on average, compared to the best performance achieved by Qiskit and t|ket> respectively, two-qubit gates by 29.22% and 20.58%, and circuit depth by 42.02% and 51.27%.

An extended version of this paper gives full proofs for all lemmas and theorems [17].

<sup>&</sup>lt;sup>1</sup>The full implications of such customization, including hardware-aware layout, routing, and scheduling, are beyond the scope of this paper.

<sup>&</sup>lt;sup>2</sup>https://developer.intel.com/quantumsdk

# II. BACKGROUND

Quantum states are represented as density matrices  $\rho$ : positive semi-definite, Hermitian complex matrices with trace 1. Unitary transformations act on them via conjugation:  $U\rho U^{\dagger}$ . A density matrix is called a pure state if it can be written as the outer product of two state vectors  $|\phi\rangle\langle\phi|$ . If not, it is a mixed state and can be written as the weighted sum of pure states, representing a probability distribution over pure states. The behavior of a quantum circuit can be described as a function over density matrices known as a quantum channel [18].

# A. The Pauli group

A single-qubit Pauli p is one of X, Y, Z, or I, where

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \qquad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$
(1)

X, Y, and Z are Hermitian, meaning that pp = I, and satisfy XYZ = i. A consequence is that any two single-qubit Paulis either commute (written  $p_1 \perp p_2$ ), meaning  $p_1 \cdot p_2 = p_2 \cdot p_1$ ; or anticommute  $(p_1 \not\perp p_2)$ , meaning  $p_1 \cdot p_2 = -p_2 \cdot p_1$ . We write

$$\lambda(p_1, p_2) = \begin{cases} 0 & p_1 \perp p_2 \\ 1 & p_1 \not\perp p_2 \end{cases}$$
(2)

An *n*-qubit Pauli  $P = \alpha(p_0, \ldots, p_{n-1}) \in \mathcal{P}_n$  is the tensor product of *n* single-qubit Paulis scaled by  $\alpha \in \{1, -1, i, -i\}$ . Its support, supp(*P*), is the set of indices for which  $p_i \neq I$ . We write  $X_i$ ,  $Y_i$ , and  $Z_i$  for Paulis with support  $\{i\}$ , and thus the Pauli string  $X_0Z_2$  refers to (X, I, Z).

Multiplication can be lifted to *n*-qubit Paulis as follows:

$$P_1 \cdot P_2 = \alpha_1 \alpha_2 \sigma_0 \cdots \sigma_{n-1} (q_0, \dots, q_{n-1})$$
(3)

where  $P_i = \alpha_i(p_0^i, \dots, p_{n-1}^i)$  and  $p_i^1 \cdot p_i^2 = \sigma_i q_i$ . As a result, *n*-qubit Paulis form a group with identity  $I = (I, \dots, I)$ .

Commutativity can be lifted to *n*-qubit Paulis via a binary function 
$$\lambda(P, P') \in \{0, 1\}$$
 such that  $P \cdot P' = (-1)^{\lambda(P,P')} P' \cdot P$ :

$$\lambda(\alpha(p_0,\ldots,p_{n-1}),\alpha'(p'_0,\ldots,p'_{n-1})) = \sum_{i=0}^{n-1} \lambda(p_i,p'_i) \mod 2$$
(4)

We write  $P \perp P'$  if  $\lambda(P, P') = 0$  and  $P \not\perp P'$  if  $\lambda(P, P') = 1$ .

We are most interested in Hermitian Paulis where PP = I. Since single-qubit Paulis are all Hermitian, an *n*-qubit Pauli is Hermitian if and only if its coefficient is  $\pm 1$ . The *Hermitian product* of Hermitian Paulis is  $P_1 \odot P_2 = (-i)^{\lambda(P_1,P_2)}P_1 \cdot P_2$ : if  $P_1$  and  $P_2$  are both Hermitian, then so is  $P_1 \odot P_2$ .

#### **III. SEMANTICS**

As quantum channels, PCOAST nodes could be seen as transformations on quantum states. However, because PCOAST deals with mixed unitary and non-unitary circuits, it also must account for transformations on *classical states*, for example when writing measurement outcomes a classical registers. A classical state m is a finite sequence of assignments of classical variables c to boolean values  $b \in \{0, 1\}$ , written  $c_0 \leftarrow b_0; \cdots; c_{n-1} \leftarrow b_{n-1}$ . The boolean value associated with a variable is written m[c], and a finite set of classical states is referred to as a *measurement space*  $\mathcal{M}$ .

# A. Classical-quantum states

Instead of working with density matrices directly, we will operate over mixed *classical-quantum states* [19, 20]. A cqstate  $\gamma \in CQ^{\mathcal{M}} = \mathcal{M} \to \mathcal{C}^n \times \mathcal{C}^n$ , sometimes written  $m \mapsto \gamma_m$ , is a function from a classical state  $m \in \mathcal{M}$  to the quantum state of the system after the measurement outcome m is observed. The quantum state  $\gamma_m$  is represented as a *partial* density matrix  $\gamma_m$ , whose trace  $0 \leq \operatorname{tr}(\gamma_m) \leq 1$  corresponds to the probability of observing m. The sum of all the partial density matrices in the image of a cq-state is a full density matrix with trace 1.

As an example, the cq-state obtained from executing the circuit  $\operatorname{PrepZ}(0)$ ; H(0);  $\operatorname{MeasZ}^{c}(0)$  is  $m \mapsto \frac{1}{2} |m[c]\rangle \langle m[c]|$ . When it is clear from context, we write  $\rho$  for the constant

When it is clear from context, we write  $\rho$  for the constant cq-state  $\_ \mapsto \rho$ . Scaling and summation of cq-states over the same measurment space is defined pointwise.

PCOAST utilizes two different equivalence relations on cqstates. The *hold* relation completely preserves the quantum state corresponding to every classical state, while the *release* relation only requires that the probability of being in the same state, tr( $\gamma_m^i$ ), is the same for each classical state *m*.

$$\gamma^1 \equiv^{\text{hold}} \gamma^2 \iff \forall m \in \mathcal{M}, \ \gamma_m^1 = \gamma_m^2 \tag{5}$$

$$\gamma^1 \equiv^{\text{release}} \gamma^2 \iff \forall m \in \mathcal{M}, \ \text{tr}(\gamma_m^1) = \text{tr}(\gamma_m^2)$$
 (6)

In this context, the semantics of a quantum circuit C can be described as a classical-quantum channel—a linear function  $\llbracket C \rrbracket : CQ^{\mathcal{M}_1} \to CQ^{\mathcal{M}_2}$  between cq-states, where  $\mathcal{M}_1$  is the set of states exectuion may be in before C, and  $\mathcal{M}_2$  contains the states the program may be in after executing C.

# B. Sum-of-Pauli semantics

Classical-quantum channels will be used to describe the behavior of circuits and PCOAST nodes on classical variables. However, the majority of nodes do not affect classical states at all. In that case, their behavior can be naturally described as a *Pauli map*, a function from *n*-qubit Paulis to cq-states. Intuitively, the quantum component of the input cq-state will be decomposed into a sum of Pauli operators scaled by arbitrary complex values, which we call *Pauli vectors*.

**Lemma 1.** Every  $2^n \times 2^n$  complex matrix A can be decomposed into a Pauli vector  $A = \sum_i \alpha_i P_i$ .

Multiplication of Pauli vectors, written  $v_1 \cdot v_2$ , and conjugate transpose,  $v^{\dagger}$ , are defined in the expected way.

**Definition 2.** A *Pauli map* is a function  $f : \mathcal{P}_n \to CQ^{\mathcal{M}}$  from *n*-qubit Paulis to cq-states. It can be lifted to a cq-channel  $[f]: CQ^{\mathcal{M}_0} \to CQ^{\mathcal{M}_0 + \mathcal{M}}$  as<sup>3</sup>

$$[f](\gamma) = m_0; m \mapsto \sum_i \alpha_i \gamma_i(m) \tag{7}$$

where  $\gamma(m_0) = \sum_i \alpha_i P_i$  and  $\gamma_i = f(P_i) \in \mathbb{CQ}^M$ .

A Pauli vector v can be lifted to a Pauli map  $v^*(P) = vPv^{\dagger}$ , called the conjugation action of v. Scaling  $(\alpha m)$  and addition  $(v_1 + v_2)$  of Pauli maps are defined pointwise, and we say a

$${}^{3}\mathcal{M}_{1} + \mathcal{M}_{2} = \{m_{1}; m_{2} \mid m_{i} \in \mathcal{M}_{i}\}.$$

Pauli map *acts on* a Pauli vector by mapping it over every Pauli in the vector:  $m(\sum_i \alpha_i P_i) = \sum_i \alpha_i m(P_i)$ .

# IV. PCOAST DATA TYPES

PCOAST uses five types of nodes for different sorts of gates:

- Pauli rotations  $\operatorname{Rot}(P, \theta)$  for non-Clifford unitaries  $e^{-i\theta/2P}$
- Pauli preparations  $Prep(P_1, P_2)$
- Pauli measurements  $Meas^{c}(P)$
- Measurement space functions  $\boldsymbol{\mu}$
- Pauli frames F for Clifford unitaries

The semantics of PCOAST nodes are defined as Pauli maps  $[n] : \mathcal{P}_k \to CQ$ , except for measurement space functions, which are defined as cq-channels directly.

**Definition 3.** A Pauli rotation  $\operatorname{Rot}(P,\theta)$  consists of a Hermitian Pauli P on k qubits and a real number  $\theta \in \mathbb{R}$ . The semantics of a Pauli rotation  $[\operatorname{Rot}(P,\theta)]$  is given by the conjugation action of its corresponding unitary,  $[\operatorname{Rot}(P,\theta)]$ :<sup>4</sup>

$$[\operatorname{Rot}(P,\theta)](Q) = [\operatorname{Rot}(P,\theta)]Q[\operatorname{Rot}(P,-\theta)]$$
(8)

$$\left[\operatorname{Rot}(P,\theta)\right] = e^{-i\theta/2P} = \cos\left(\frac{\theta}{2}\right)I + i\sin\left(\frac{\theta}{2}\right)P \quad (9)$$

**Definition 4.** A Pauli preparation  $Prep(P_Z, P_X)$  consists of a pair of non-commutative Hermitian Paulis.

$$\begin{bmatrix} \operatorname{Prep}(P_Z, P_X) \end{bmatrix}(Q) \\ = \frac{1}{4} \left( (I + P_Z)Q(I + P_Z) + P_X(I - P_Z)Q(I - P_Z)P_X \right) \\ = \begin{cases} Q(I + P_Z) & \lambda(Q, P_Z) = \lambda(Q, P_X) = 0 \\ 0 & \text{otherwise} \end{cases}$$
(10)

Intuitively,  $Prep(P_Z, P_X)$  collapses the state in the eigensubspaces of  $P_Z$  and, if -1 is obtained, applies  $P_X$ . As an example,  $Prep(Z_i, X_i)$  prepares the *i*th qubit in the Z basis.

**Definition 5.** A Pauli measurement  $\text{Meas}^{c}(P)$  consists of a Hermitian Pauli acting as a measurement operator.

$$[\operatorname{Meas}^{c}(P)](Q) = m \mapsto \begin{cases} \frac{1}{4}(I+P)Q(I+P) & m[c] = 0\\ \frac{1}{4}(I-P)Q(I-P) & m[c] = 1 \end{cases}$$
$$= c \leftarrow b \mapsto \begin{cases} \frac{1}{2}Q(I+(-1)^{b}P) & \lambda(P,Q) = 0\\ 0 & \lambda(P,Q) = 1 \end{cases}$$
(11)

As an example,  $\text{Meas}^{c}(Z_i)$  measures qubit *i* in the *Z* basis. *c* is the classical register in which measurement is recorded.

**Definition 6.** A measurement space function  $\mu : \mathcal{M}_1 \to \mathcal{M}_2$  is a function between two measurement spaces. It is interpreted as a cq-channel  $\llbracket \mu \rrbracket : CQ^{\mathcal{M}_1} \to CQ^{\mathcal{M}_2}$  as follows:

$$\llbracket \mu \rrbracket(\gamma) = m_2 \mapsto \sum_{m_1 \in \mu^{-1}(m_2)} \gamma_{m_1}.$$
 (12)

As an example, in the circuit  $\operatorname{Prep}(Z_i, X_i)$ ;  $\operatorname{Meas}^c(Z_i)$ , the measurement can be optimized away, yielding  $\operatorname{Prep}(Z_i, X_i)$ ;  $\mu$  where  $\mu(m) = m$ ;  $c \leftarrow 0$ .

# A. Pauli Frames

Clifford unitaries satisfy the property that when acting on a Pauli P by conjugation, the result  $UPU^{\dagger}$  is still a Pauli. A Pauli frame<sup>5</sup> is a compact representation of a Clifford unitary defined by that conjugation—or more precisely, its inverse conjugation  $U^{\dagger}PU$ —on every base Pauli string  $Z_j$ ,  $X_j$ , and  $Y_j$ . For example, the inverse conjugation of  $U = \text{CNOT}_{0,1}$  is

$$\frac{j \quad U^{\dagger}Z_{j}U \quad U^{\dagger}X_{j}U \quad U^{\dagger}Y_{j}U}{0 \quad Z_{0} \quad X_{0}X_{1} \quad Y_{0}X_{1}} \\
1 \quad Z_{0}Z_{1} \quad X_{1} \quad Z_{0}Y_{1}$$
(13)

It suffices to store only the first two columns of this table: we can derive  $U^{\dagger}Y_{j}U = -i(U^{\dagger}Z_{j}U)(U^{\dagger}X_{j}U)$  since Y = -iZX. Thus, an *n*-qubit Clifford is represented by an  $n \times 2$  Pauli frame, where the first column stores  $U^{\dagger}Z_{j}U$  and the second column stores  $U^{\dagger}X_{j}U$ :

$$\begin{pmatrix} Z_0 & X_0 X_1 \\ Z_0 Z_1 & X_1 \end{pmatrix}$$
(14)

Note that whether we store  $U^{\dagger}PU$  or  $UPU^{\dagger}$  in the entries of the Pauli frame is a matter of style—the inverse of a Pauli frame can always be calculated to obtain one from the other [17]. However, the choice affects the efficiency of the lookup operation  $\vec{F}$  (Defn. 8 below). For efficiency in PCOAST, the lookup operation should implement the opposite of the semantic interpretation  $[\![F]\!](Q) = UQU^{\dagger}$  (Defn. 10).

**Definition 7.** A *Pauli frame* F on k qubits is a  $k \times 2$  array of Hermitian k-qubit Paulis:

$$F = \begin{pmatrix} \text{eff}Z_0 & \text{eff}X_0 \\ \vdots & \vdots \\ \text{eff}Z_{n-1} & \text{eff}X_{n-1} \end{pmatrix}$$
(15)

The arguments in the first column are called the *effective*  $Z_j$ Paulis, and the arguments in the second column the *effective*  $X_j$  Paulis, and they must respect all the same commutativity relations as the corresponding  $Z_j$  and  $X_j$  Paulis:

$$\lambda(\operatorname{eff} Z_i, \operatorname{eff} Z_j) = \lambda(\operatorname{eff} X_i, \operatorname{eff} X_j) = 0$$
(16)

$$\lambda(\operatorname{eff} Z_i, \operatorname{eff} X_j) = \lambda(\operatorname{eff} X_j, \operatorname{eff} Z_i) = \delta_{i,j}^{6}$$
(17)

**Definition 8.** The *lookup action* of F on  $P = \alpha(p_0, \ldots, p_{k-1})$ , written  $\vec{F}(P)$ , is the product of the effective entry of each  $p_j$ :

$$\vec{F}(\alpha(p_0,\ldots,p_{n-1})) = \alpha \prod_j \text{eff}p_j \tag{18}$$

where  $\operatorname{eff} I_j = I$  and  $\operatorname{eff} Y_j = \operatorname{eff} Z_j \odot \operatorname{eff} X_j$ .

**Lemma 9.** For every Pauli frame F there is a Clifford unitary  $U^F$ , unique up to overall phase, satisfying  $\overrightarrow{F}(P) = (U^F)^{\dagger} P U^F$  for any Pauli P.

<sup>5</sup>Pauli tableaus [15] were first introduced as a way to simulate stabilizer states generated entirely from Clifford gates and single-qubit measurements. Since then, Pauli tableaus have been used to represent Clifford circuits in general, not just for the purposes of stabilizer simulation. Following [13, 21], in this work we refer to Pauli tableaus as *Pauli frames* to emphasize their linear algebraic structure with regards to the Pauli group.

 ${}^{6}\delta_{i,j}$  is 0 if i = j and 1 otherwise.

<sup>&</sup>lt;sup>4</sup>Recall, if the classical state is not specified, it is assumed to be the constant cq-state  $\_ \mapsto \rho$ .

$$\frac{P \perp n}{\operatorname{Rot}(P,\theta) \perp n} \quad \frac{P_1 \perp n}{\operatorname{Prep}(P_1, P_2) \perp n} \quad \frac{P \perp n}{\operatorname{Meas}^c(P) \perp n}$$
$$\frac{F \circ F' = F' \circ F}{F \perp F'} \quad \frac{n \text{ not a Pauli frame} \quad n \perp F}{F \perp n}$$
$$\frac{\mu \perp \operatorname{Rot}(P,\theta)}{\mu \perp F} \quad \frac{\mu \perp \operatorname{Prep}(P_1, P_2)}{\mu \perp \operatorname{Prep}(P_1, P_2)}$$

Fig. 3: Definition of  $n_1 \perp n_2$  indicating when two PCOAST nodes commute. These definitions, given as inference rules, are understood as follows: the conclusion below the line holds if and only if all the hypotheses above the line hold.

**Definition 10.** The semantics of a Pauli frame is given by the Pauli map  $[\![F]\!](Q) = U^F Q (U^F)^{\dagger}$ .

The frame  $F^U$  associated with a Clifford unitary U is

$$\mathrm{eff} p_j^{F^U} = U^{\dagger} p_j U, \qquad (19)$$

where  $\operatorname{eff} p_j^F$  is entry of F corresponding to  $p_j$ . We refer to  $F^I$  as the *origin frame*.

A consequence of the commutativity rules for frames is that the lookup action preserves commutativity and composition.

$$\lambda(\vec{F}(P), \vec{F}(Q)) = \lambda(P, Q) \tag{20}$$

$$\vec{F}(PQ) = \vec{F}(P)\vec{F}(Q) \tag{21}$$

**Definition 11.** Composition of Pauli frames  $F_2 \circ F_1$  is defined as  $\operatorname{eff} p_i^{F_2 \circ F_1} = \overrightarrow{F_1}(\operatorname{eff} p_i^{F_2})$ . It satisfies  $U^{F_2 \circ F_1} = U^{F_2} U^{F_1}$ .

#### B. PCOAST Terms

A *PCOAST term* t is a sequence of PCOAST nodes  $n_0; \ldots; n_{k-1}$  with 1 indicating the empty sequence. The semantics of nodes can be lifted to terms  $[t_1]: CQ \rightarrow CQ$ .

We define an equivalence relation on PCOAST terms parameterized by a *hold* or *release* outcome *o*, corresponding to the two equivalence relations on cq-states.

$$t_1 \equiv^o t_2 \iff \forall \gamma, \ [t_1]](\gamma) \equiv^o [t_2]](\gamma) \tag{22}$$

If not specified, we assume *o* is the stronger *hold* outcome.

Intuitively, two PCOAST nodes commute exactly when their underlying Paulis commute. Formally, Fig. 3 defines a commutativity relation  $n \perp n'$  between nodes, aided by a helper relation  $Q \perp n$  (Fig. 4) that indicates when a Pauli commutes with a PCOAST node.

**Theorem 12.** If  $n_1 \perp n_2$  then  $n_1; n_2 \equiv n_2; n_1.^7$ 

$$\frac{Q \perp P}{Q \perp \operatorname{Rot}(P,\theta)} \qquad \frac{\overrightarrow{F}(Q) = Q}{Q \perp F}$$

$$\frac{Q \perp P_1}{Q \perp \operatorname{Prep}(P_1, P_2)} \qquad \frac{Q \perp P}{Q \perp \operatorname{Meas}^c(P)} \qquad \overline{Q \perp P}$$

Fig. 4: Definition of  $Q \perp n$ , when a Pauli Q commutes with a PCOAST node n. Recall that for two Paulis, we write  $P_1 \perp P_2$  if and only if  $\lambda(P_1, P_2) = 0$ .

## C. The PCOAST Graph

**Definition 13.** A PCOAST graph G = (V, E) is a directed acyclic graph whose vertices V are PCOAST nodes. For any nodes  $n_1$  and  $n_2$  that do not commute, there is either an edge from  $n_1$  to  $n_2$  or vice versa (but not both). There are no edges between commuting vertices.

As an example, in Fig. 1c there are edges from both measurement nodes to the the measurement space function node  $\mu$ , but there are no edges to the Pauli frame F' because both  $Z_0$  and  $X_1$  commute with F'.

Every topological ordering of a PCOAST graph G corresponds to a PCOAST term  $t^G$ . Even though topological orderings of a graph are not unique, they are equivalent to each other due to Thm. 12: if two vertices do not commute, there is an edge between them in one direction or the other, and that edge will be preserved by the topological ordering.

PCOAST graphs satisfy three key invariants: frameterminal, measurement-space terminal, and fully merged.

1) Frame-terminal graphs: A PCOAST graph is called frame-terminal when it contains a single Pauli frame node, and that frame has no outgoing edges. It is always possible to construct a frame-terminal graph because a Pauli frame F can be commuted past any other node n by transforming n into a new node  $\vec{F}(n)$  that satisfies  $F; n \equiv \vec{F}(n); F$ .

$$\vec{F}(n) = \begin{cases} F^{-1} \circ F' \circ F & n = F' \\ \operatorname{Rot}(\vec{F}(P), \theta) & n = \operatorname{Rot}(P, \theta) \\ \operatorname{Prep}(\vec{F}(P_Z), \vec{F}(P_X)) & n = \operatorname{Prep}(P_Z, P_X) \\ \operatorname{Meas}^c(\vec{F}(P)) & n = \operatorname{Meas}^c(P) \\ \mu & n = \mu \end{cases}$$
(23)

2) Measurement-space terminal graph: Similarly, a graph is called measurement-space terminal when it contains at most one measurement space function node, and that node has no outgoing edges. Since frames and measurement space functions always commute, this does not conflict with the graph being frame-terminal. To construct a measurementspace-terminal graph, it suffices to push all measurement space nodes past measurement nodes via the equivalence

$$\mu; \operatorname{Meas}^{c}(P) \equiv \operatorname{Meas}^{c'}(P); \mu'$$
(24)

where c' is fresh and  $\mu'(m) = m; c \leftarrow m(c')$ .

<sup>&</sup>lt;sup>7</sup>Note that this property does not hold in the other direction; the rules for  $n_1 \perp n_2$  are strictly more restrictive. In particular, the relation is not reflexive on  $\text{Prep}(P_Z, P_X)$  since  $P_Z \not\perp P_X$ .

$$\begin{split} & \operatorname{Rot}(P,\theta_1); \operatorname{Rot}(\pm P,\theta_2) \longrightarrow_M \begin{cases} F^{\operatorname{Rot}(P,\theta)} & \theta = \theta_1 \pm \theta_2 \text{ is a} \\ & \operatorname{multiple of } \frac{\pi}{2} \\ \operatorname{Rot}(P,\theta) & \operatorname{otherwise} \end{cases} \\ & \operatorname{Prep}(P_Z, P_X); \operatorname{Prep}(\pm P_Z, P_X') \longrightarrow_M \operatorname{Prep}(\pm P_Z, P_X') \\ & \operatorname{Prep}(P_Z, P_X); \operatorname{Rot}(\pm P_Z, \theta) \longrightarrow_M \operatorname{Prep}(P_Z, P_X) \\ & \operatorname{Meas}^{c_1}(P); \operatorname{Meas}^{c_2}((-1)^b P) \longrightarrow_M \operatorname{Meas}^{c_1}(P); c_2 \leftarrow c_1 + b \\ & \operatorname{Prep}(P_Z, P_X); \operatorname{Meas}^c((-1)^b P_Z) \longrightarrow_M \operatorname{Prep}(P, P_X); c \leftarrow b \\ & \operatorname{Rot}(\pm P, \theta); \operatorname{Meas}^c(P) \longrightarrow_M \operatorname{Meas}^c(P) \\ & F_1; F_2 \longrightarrow_M F_2 \circ F_1 \qquad \mu_1; \mu_2 \longrightarrow_M \mu_2 \circ \mu_1 \end{split}$$

Fig. 5: Rules for merging vertices. When  $\theta$  is a multiple of  $\frac{\pi}{2}$ , Rot $(P, \theta)$  is a Clifford unitary.



Fig. 6: Flow for constructing a PCOAST graph from a circuit, starting with the circuit on the left  $C^{\text{orig}}$  and ending with the optimized circuit on the right  $C^{\text{opt}}$ . The steps correspond to Sections V-A, V-B, and V-C).

3) Mergeable nodes: A pair of nodes  $(n_1, n_2)$  is called *mergeable* if there is no path between them in the graph and there is a step  $n_1; n_2 \longrightarrow_M t$  to a new Pauli term, as defined in Fig. 5. For such a merge to occur,  $n_1$  and  $n_2$  must have the same underlying Pauli arguments modulo  $\pm 1$ , in which case they can be combined into a single node, or in the case of measurement, two new but still simpler nodes.

# **Lemma 14.** If $t_1 \longrightarrow_M t_2$ then $t_1 \equiv t_2$ .

# V. THE PCOAST OPTIMIZATION

The PCOAST optimization, illustrated in Fig. 6, starts with a circuit, compiles it to a PCOAST graph, optimizes that graph, and then synthesizes a circuit again. To maintain the frame- and measurement-space-function-terminal invariants, we separate out the terminal frame F and measurement space function  $\mu$  from the rest of the graph G.

#### A. Compiling circuits to PCOAST graphs

The procedure CIRCTOGRAPH(C) (Fig. 7) produces a PCOAST graph G, a terminating Pauli frame F, and a terminating measurement space function  $\mu$  from a circuit C by moving gates into G or F respectively using a helper function ADDNODE(G, n) (Fig. 8). The algorithm maintains the loop invariant that  $t^C \equiv g; F; \mu; t^{C'}$ .

# B. Internal optimizations on PCOAST graphs

We give a brief overview of the internal optimizations on the PCOAST graphs here, where a detailed discussion is the contents of [21]. The optimizations are primarily concerned with the interfaces between unitary and non-unitary elements of the graph, and how these optimizations are leveraged

1: procedure CIRCTOGRAPH(C)  
2: 
$$G \leftarrow \emptyset; F \leftarrow F^{I}; \mu \leftarrow \mu^{0}$$
  
3: for each g in C do  
4: for each n in  $t^{g}$  do  
5: if  $n = \text{Meas}^{c}(P)$  then  
6:  $n \leftarrow \text{Meas}^{c'}(P); \mu \leftarrow \mu[c \mapsto c'] \triangleright c'$  fresh  
7: if  $n = \mu'$  then  $\mu \leftarrow \mu' \circ \mu$   
8: else if  $n = F'$  then  $F \leftarrow F' \circ F$   
9: else  
10:  $(G', F', \mu') \leftarrow \text{ADDNODE}(G, \overrightarrow{F}(n))$   
11:  $G \leftarrow G'; F \leftarrow F \circ F'; \mu \leftarrow \mu \circ \mu'$   
12: return  $(G, F, \mu)$ 

Fig. 7: The function CIRCTOGRAPH(C) takes as input a circuit C and returns a PCOAST graph G, a Pauli frame F, and a measurement space function  $\mu$  such that  $G; F; \mu \equiv t^C$ . We assume that every gate in C can be written in a straightforward way as a PCOAST term  $t^g$ .

- 1: procedure ADDNODE(G, n) 2: if n = F then return  $(G, F, \mu^0)$
- 2: else if  $n = \mu$  then return  $(G, F^I, \mu)$ 3: 4: else for each node  $n_0$  in G with outdegree 0 do 5: if  $n_0; n \longrightarrow_M n'_0$  OR  $n'_0; \mu_0$  then 6:  $G' \leftarrow \text{REMOVEVERTEX}(G, n_0)$ 7:  $(G'', F, \mu) \leftarrow \text{ADDNODE}(G', n'_0)$ 8: return  $(G'', F, \mu_0 \circ \mu)$ 9:  $G' \leftarrow \text{ADDVERTEX}(G, n')$ 10: for each node  $n_0$  in G' do 11: if  $n_0 \not \perp n'$  then  $G' \leftarrow \text{ADDEDGE}(G', n_0, n')$ return  $(G', F^I, \mu^0)$ 12:

Fig. 8: The procedure ADDNODE(G, n), where G is a PCOAST graph and n is a PCOAST node, returns an updated graph G', a Pauli frame F, and a measurement space function  $\mu$  such that  $G; n \equiv G'; F; \mu$ . The functions ADDVERTEX, RE-MOVEVERTEX, and ADDEDGE implement the corresponding simple graph operations. Note that if  $n_0; n \longrightarrow_M t$ , it is either the case that t is a single node  $n'_0$ , or is equal to  $n'_0; \mu_0$ .

depends on the desired outcome of the quantum program, release or hold. We give a brief outline of the optimizations here, and illustrate them in Fig. 9.

- 1) In both cases, reduce the size of the support for nodes which are either weakly or not at all dependent (as defined in [21]) on a preparation (Fig. 9a).
- 2) For a *hold* outcome, reduce the cost of the terminating frame in the presence of preparations with outdegree 0 in the graph (Fig. 9b).
- 3) For a *release* outcome (Fig. 9c):
  - a) Trivialize the terminating frame and remove nodes that can be commuted past all measurements.
  - b) If there are preparation nodes with outdegree 0,
    - i) Replace all measurement nodes with outdegree 0

with an equivalent and more efficient set of measurements and a terminal measurement map via a *stabilizer search* [21].

- ii) reduce the cost of the the frame generated by 3(b.i) in the presence of preparations with outdegree 0.
- iii) convert the graph back to frame-terminating.

**Theorem 15** ([21]). If G is optimized to G' with outcome o, then  $t^G \equiv^o t^{G'}$ .

# C. Synthesizing circuits from PCOAST graphs

After optimizing the PCOAST graph, the next step is to synthesize an equivalent circuit. In this section we summarize the method outlined in Ref. [13] along with the modifications needed to adapt it to PCOAST.

We will always start with a frame-terminal, fully merged PCOAST graph (G, F). The process starts by prepending the graph with an empty circuit and iteratively synthesizing gates from dependency-free nodes. If a node has a direct equivalent as a gate, such as  $Prep(Z_i, X_i)$ , that gate is added directly. If not, we select a Clifford gate g to add to the circuit, and transform the rest of the graph by  $F' = F^{g^{-1}}$  to obtain

$$C;g;F';G;F \equiv C;g;\overrightarrow{F'}(G);F \circ F'$$
(25)

where  $\vec{F'}(G)$  applies  $\vec{F'}(n)$  to each node in G.

Thus synthesis is primarily a method for selecting the sequence of Clifford gates, particularly two-qubit entangling (TQE) gates, that minimizes the gate cost, depth or other cost metric of the resulting circuit.<sup>8</sup> To do so, we adapt the search algorithm of [13] with what we call *search functions*.

- 1) NODECOST(n) returns the cost to implement n, which is zero if and only if a gate  $g^n$  implements it.
- 2) REDUCENODE(n) returns a set of TQE gates g that minimize the NODECOST of  $\overrightarrow{F^g}(n)$ . Repeated application will monotonically reduce NODECOST to 0.
- 3) GATECOST(C, G, F, g) returns the cost of the action of Eq. (25).
- 4) ADDGATE(C, G, F, g) performs the action of Eq. (25).

We start with synthesizing non-Clifford nodes (Fig. 10). The search always terminates due to the search functions' guarantees [13]. It should be clear that  $t^C; t^{G'}; t^{F'} \equiv t^G$ , as ADDGATE preforms a logical identity between the graph and the circuit.

1) Search function implementations: Search functions can be implemented differently based on complexity and cost. The basic implementation in the Intel Quantum SDK primarily considers the minimum number of TQE gates required to reduce the node cost to zero. This cost depends on whether the node is a *singlet* or a *factor* node.

Singlet nodes n(P) are those defined by a single Pauli operator such as  $Rot(P, \theta)$  or  $Meas^{c}(P)$ .

$$NODECOST(n(P)) = supp(P) - 1.$$
(26)

<sup>8</sup>For a given qubit pair there are  $3 \times 3 = 9$  TQE gates, generalizing CNOT and CZ, such that we have one Pauli operator basis for each qubit [13].

Factor nodes n(P,Q) are defined by two anti-commuting Pauli operators. We can understand factor nodes as *effective qubits* consisting of an effective Z and effective X, much like the rows of the Pauli frame. Thus far we have only introduced one type of factor node,  $Prep(P_Z, P_X)$ , but to synthesize Clifford circuits in Sec. V-C2, we break an *n*-qubit Pauli frame into *n* factor nodes, one for each row, called *local frames*.<sup>9</sup>

The node cost of a factor node is the number of TQE gates needed to reduce the effective qubit to an actual one. The *local* support matrix of an anti-commuting pair P, Q and qubit i is

$$\operatorname{supp}_{i}(P,Q) = \begin{pmatrix} \lambda(P,X_{i}) & \lambda(P,Z_{i}) \\ \lambda(Q,X_{i}) & \lambda(Q,Z_{i}) \end{pmatrix},$$
(27)

The determinant (mod 2) of the local support matrix is called the *local determinant*. We say a node has *weak support* on qubit *i* if its local support matrix is nonzero, but has zero determinant, and *strong support* if its local determinant is 1. Ref. [17] argues that for any factor node, the sum over all local determinants is 1 mod 2. In addition, there exist six TQE gates to reduce any two qubits with strong support to two with weak support, and exactly one to reduce one strong and one weak support to no support on the weak support qubit.<sup>10</sup> Reducing a Pauli factor node is then the process of reducing all strong support to a single qubit (the final qubit to which the node is reduced) and eliminating all weak support. Thus

NODECOST
$$(n(P,Q)) = \left(\sum_{i} \det(\operatorname{supp}_{i}(P,Q)) - 1\right)/2 + \left(\sum_{i} [\operatorname{supp}_{i}(P,Q) \neq 0] - 1\right)$$
 (28)

The implementation of REDUCENODE follows from the discussion in [17] around the selection of gates which reduce strong-strong and strong-weak qubit pairs.

GATECOST is the average change in NODECOST over all remaining nodes.<sup>11</sup> We use average change as opposed to average absolute cost as it is faster to compute. ADDGATE produces the equivalent change to the PCOAST graph.

The Intel Quantum SDK implements a variation of the basic cost functions, incorporating the following:

- A "schedule" search maintains an approximate ASAP scheduling of the circuit. GATECOST includes a parallelization credit with a tuneable parameter [13].<sup>12</sup>
- A "native gate" search targeting a particular gate set, which informs circuit cost. In the case of the the Intel Quantum SDK, the set is {CZ, RXY, MeasZ, PrepZ}.
- GATECOST includes a cost penalty for the action of the gate on the terminating frame.

<sup>9</sup>In the implementation, we also have 2-axis rotations Rot<sup>2</sup>( $P_1, P_2, \theta, \phi$ ), the generalization of the Intel Quantum SDK's RXY gate; see [17].

<sup>10</sup>There do exists cases where a TQE gate can reduce two qubits with weak support to one with weak support, but it is not guaranteed.

<sup>11</sup>One simple modification to GATECOST is to give greater weight to dependency-free nodes, as discussed in Section VI. There, we weight free nodes proportionally to the total number of nodes in the graph divided by the number of free nodes, meaning that weighting has diminishing impact as synthesis proceeds.

<sup>12</sup>The value of this is a fine-tuning modification studied in Sec. VI.



(a) Reducing node support in the presence of a preparation node (top) is the equivalent of removing a control-line when the control is initialized via a preparation (bottom).



(b) Reducing the cost of the Pauli Frame in the presence of a preparation (top) is the equivalent of removing a CZ when the control line is initialized via a preparation (bottom).



(c) Sequence of transformations applied in the case a release outcome.

Fig. 9: Examples of the PCOAST graph optimizations.

1:	<b>procedure</b> SEARCHNONCLIFFORD(G, F', hold/release)
2:	$C \leftarrow \varnothing, \ G' \leftarrow G, \ F' \leftarrow F$
3:	while $G'$ contains ( <i>release</i> ) non-measurement /
	(hold) non-Clifford nodes do
4:	for $n \in \text{BEGIN}(G')$ do
5:	if NODECOST $(n) = 0$ then
6:	if (hold) OR $n \notin \text{END}(G')$ OR
7:	n not a measurement node then
8:	$ADDGATE(C,G',F',g^n)$
9:	$G' \leftarrow \text{REMOVEVERTEX}(G', n)$
10:	$Min \leftarrow \arg \min_{n \in BEGIN(G')} NODECOST(n)$
11:	MinGate ← REDUCENODE(Min)
12:	$g_{min} \leftarrow \arg \min_{q \in MinGate} GATECOST(C, G', F', g)$
13:	$ADDGATE(C, \breve{G'}, F', g_{min})$
14:	return $(C, G', F')$

Fig. 10: Ultra Greedy Search Algorithm for Non-Clifford nodes of a PCOAST graph G. Returns a circuit C and terminating frame F' such that  $t^C; t^{G'}; t^{F'} \equiv t^G$ . BEGIN(G) is the set of dependency-free non-Clifford nodes in G and END(G) the set of nodes with outdegree 0.

• The search functions described here are also used to implement the stabilizer search algorithm of [21].

2) Adapting Ultra-Greedy Search to finalize the circuit: After synthesizing the non-Clifford nodes, Fig. 10 leaves us with (G', F'), where either (hold) G' is empty and F' is non-trivial, or (release) G' contains only mutually commuting measurements and F' can be discarded.

In the *hold* case, while methods exist for synthesizing a circuit for a Pauli frame/tableau [15, 22], we leverage the ultra-greedy algorithm to automatically take into account all search function considerations. As such, Pauli frame synthesis

is implemented by Fig. 10 with the rows of the Pauli frame reinterpreted as a set of mutually independent local-frame factor nodes, and a zero-cost node replacing the Clifford it represents. With a few additional promises made by the search functions (see [21] for details), it is guaranteed that when one local frame is reduced, it is independent of all other local frames. If the local frame is reduced to a different qubit than its row index, it is implemented by a swap. In some cases this swap can be implemented virtually and thus considered free, but future versions may include the swap cost when qubit swapping must be performed by gates.

In the *release* case, we can again adapt Fig. 10 for the measurements in G' by replacing the search functions with the stabilizer-search-templated versions and adding any measurement space function generated by the search. As measurement space functions are (binary) linear, they add no more than a polynomial-time cost (in number of measurements) to the overall quantum-classical computation [21].<sup>13</sup>

#### VI. EVALUATION

PCOAST is implemented in C++ as the core optimization of the Intel Quantum SDK, enabled by the (-O1) flag. In this section, we evaluate its performance against IBM's Qiskit [7] and Quantinuum's t $|\text{ket}\rangle$  [10] optimizing compilers.

#### A. Experimental Setup

**System:** Our experiments use an Intel Xeon<sup>®</sup> Platinum CPU (2.4GHz, 2TB RAM) and Python 3.10.

**Framework Setup:** We compare against the Qiskit transpiler with optimization levels 2 ( $qiskit_2$ ) and 3 ( $qiskit_3$ ). For t[ket), we compare with two predefined optimization

<sup>&</sup>lt;sup>13</sup>Measurement space functions are handled automatically in the Intel Quantum SDK: classical instructions are generated in the LLVM IR to appropriately map fresh measurement outcomes to classical variables.

TABLE I: Results for total gate count, two-qubit gates, and depth compared with Qiskit and t $|ket\rangle$ . The coloring indicates the range from best (dark green, bold) to worst (white).

			(		2Q Gates							Depth							
Benchmark	Ν	$PCOAST_1$	PCOAST <sub>FT</sub>	$qiskit_2$	qiskit <sub>3</sub>	$tket_1$	$tket_2$	$PCOAST_1$	PCOAST <sub>FT</sub>	$qiskit_2$	$qiskit_3$	$tket_1$	$tket_2$	PCOAST1	PCOAST <sub>FT</sub>	$qiskit_2$	qiskit <sub>3</sub>	$tket_1$	$tket_2$
H2_BK	4	47	47	117	117	79	76	13	13	38	33	18	17	31	31	95	86	66	60
H2_JW	4	58	36	195	165	91	81	19	12	56	40	20	16	25	21	126	103	65	54
H2_PM	4	53	53	119	106	72	71	15	15	38	30	16	15	25	25	99	80	63	58
LiH_BK	12	3143	2372	24985	21080	12017	10786	1308	964	8680	6834	3385	2858	1157	1027	18549	15111	9799	8665
LiH_JW	12	3173	2647	21239	20083	9471	8225	1295	1062	8064	6666	2616	2040	1194	1070	17027	14787	7654	6310
LiH_PM	12	3218	2402	22214	20897	11030	9795	1324	905	7640	6893	3092	2565	1182	999	16131	14674	9263	8021
BeH2_BK	14	7164	6016	54201	48857	26465	23834	3049	2407	18796	15936	7482	6240	2726	2430	39889	34861	21702	19029
BeH2_JW	14	6812	5967	54113	51123	23684	19974	2970	2363	21072	17528	6669	5044	2571	2372	44354	38564	19706	15619
BeH2_PM	14	7143	5806	59164	55629	29069	25923	3051	2260	20392	18418	8309	6913	2646	2340	42891	39019	24475	21316
qft_5	5	68	65	87	69	104	86	26	25	26	20	26	20	31	25	48	41	55	49
qft_10	10	262	210	334	289	416	371	112	89	105	90	105	90	81	70	103	96	120	114
qft_20	20	922	645	1269	1179	1631	1541	440	289	410	380	410	380	270	217	212	206	250	244
qft_30	30	2281	1280	2804	2642	3646	3466	1119	589	915	858	915	858	491	455	322	316	380	374
qft_50	50	6348	3150	7674	5862	9236	7666	3282	1489	2525	1898	2315	1898	1324	1165	542	536	640	634
grover_5	5	37	37	52	52	62	62	13	13	13	13	13	13	6	6	35	35	44	44
grover_10	10	162	145	187	187	230	230	56	48	49	49	49	49	11	11	94	94	124	124
grover_30	30	530	504	637	637	790	790	185	160	169	169	169	169	26	26	274	274	364	364
grover_80	80	1545	1442	1762	1762	2190	2190	559	467	469	469	469	469	63	63	721	721	960	960
grover_100	100	1936	1814	2212	2212	2750	2750	695	587	589	589	589	589	78	78	901	901	1200	1200
hea5_1_20	5	341	307	326	326	345	345	93	80	80	80	80	80	114	108	106	106	106	106
hea5_c_20	5	473	380	405	405	405	405	184	125	100	100	100	100	221	175	222	222	222	222
hea5_f_20	5	326	307	446	446	705	705	85	80	200	200	200	200	111	108	144	144	264	264
hea10_l_40	10	1591	1412	1451	1451	1490	1490	480	360	360	360	360	360	277	218	216	216	216	216
hea10_c_40	10	4348	2655	1610	1610	1610	1610	1749	977	400	400	400	400	924	659	842	842	842	842
hea10_f_40	10	1592	1412	2891	2891	5810	5810	476	360	1800	1800	1800	1800	275	218	489	489	1129	1129
hea20_1_50	20	4065	3772	3821	3821	3870	3870	1152	950	950	950	950	950	401	288	286	286	286	286
hea20_c_50	20	20533	4679	4020	4020	4020	4020	9048	1524	1000	1000	1000	1000	2207	1275	2052	2052	2052	2052
hea20_f_50	20	4063	3772	12371	12371	29520	29520	1147	950	9500	9500	9500	9500	401	288	1119	1119	2919	2919
qaoa_6_3	6	98	83	85	89	108	108	23	22	24	24	24	24	36	33	37	37	50	50
qaoa_6_6	6	97	86	93	93	122	122	25	24	28	28	28	28	37	34	51	51	76	76
qaoa_17_3	17	432	381	406	406	586	586	148	136	148	148	148	148	90	75	81	81	114	114
qaoa_17_6	17	731	659	761	761	1202	1202	253	220	324	324	324	324	128	124	161	161	247	247
qaoa_28_3	28	1007	887	952	952	1456	1456	356	328	384	384	384	384	151	116	156	156	225	225
qaoa_28_6	28	1746	1617	1870	1870	3052	3052	660	596	840	840	840	840	248	218	278	278	423	423
qaoa_40_3	40	2169	1882	2035	2035	3254	3254	865	725	884	884	884	884	251	222	279	279	418	418
qaoa_40_6	40	3456	3220	3876	3876	6460	6460	1361	1210	1800	1800	1800	1800	413	364	418	418	629	629

sequences:  $tket_1$ , comprised mainly of the *SynthesiseTket* pass; and  $tket_2$ , comprising of the *FullPeepholeOptimise* pass.<sup>14</sup>

For PCOAST we compare optimization level 1 (PCOAST<sub>1</sub>) and a version where we fine-tune some parameters of the search functions (PCOAST<sub>FT</sub>), as discussed in Sec. VI-C.

**Benchmarks:** We analyze the compilation performance for a total of 36 benchmarks of different configurations and sizes: the Unitary Coupled-Cluster Single and Double excitations (UCCSD) ansatz [23], Quantum Fourier Transform (QFT) [18], Grover's Diffusion operator [24], Hardware-Efficient Ansatz (HEA) [25], and the Quantum Approximate Optimization Algorithm (QAOA) ansatz [26].

For UCCSD, we construct ansatz for 3 molecules ( $H_2$ , LiH, and BeH<sub>2</sub>) obtained through 3 fermionic mapping techniques: Jordan-Wigner (JW) [27], Bravyi–Kitaev (BK) [28], and Parity Mapping (PM) [29]. We test HEA for different circuit sizes and entanglement arrangements: linear, circular, and full. Finally, we construct QAOA ansatz for graph MaxCut on Erdős-Rényi random graphs [30] of different sizes with edge probabilities of 0.3 and 0.6.

**Gate Set Conversions:** To guarantee fairness, we map both Qiskit and t|ket)'s compilation workflows to Intel's native gate set {RXY, CZ, RZ, I}. This is done in Qiskit by adding equivalence rules to the SessionEquivalenceLibrary class, such as connecting RXY (*RGate* in Qiskit) to other gates and decomposing CX into RXY; CZ; RXY. For t|ket) this is done via a *CustomRebase* pass that maps the gate set to t|ket)'s TK1 and CX gates, and applied at the end of *tket*<sub>1</sub> and *tket*<sub>2</sub>.

# B. Gate Counts & Circuit Depth

Table I shows the total gate count, two-qubit gates, and depth for PCOAST, Qiskit, and  $t|ket\rangle$  across all benchmarks. Results assume an all-to-all connected machine as a backend and are obtained from multiple runs to guarantee consistency.

The untuned implementation, PCOAST<sub>1</sub>, reduces the total gate count by 22.51% (respectively 33.70%), singlequbit gates by 22.06% (43.58%), two-qubit gates by 13.46%(2.62%), and depth by 36.02% (45.52%) compared to the best Qiskit (t|ket)) performance across all benchmarks.

PCOAST<sub>1</sub> performs exceptionally well with UCCSD, with average reductions in total gate count by 76.42% (55.51%), two-qubit gates by 72.52% (32.45%), and depth by 84.82% (74.30%) compared to the best Qiskit (t|ket)) results. This can be attributed to the fact that all fermionic mapping methods are equivalent up to conjugation by a Clifford, and thus are naturally captured by Pauli frames.

In some cases, PCOAST<sub>1</sub> is outperformed by Qiskit and t|ket) in two-qubit gate counts for QFT (34.71%) and HEA (34.14%). This is a result of PCOAST<sub>1</sub>'s circuit synthesis cost function optimizing mainly for circuit depth.

#### C. Fine-tuning PCOAST's Cost Function

To investigate the cases where PCOAST<sub>1</sub> performs poorly ( $qft_50$ ,  $hea10_c$  and  $hea20_c$ ), we modify the parallelization credit, as described in Footnote 12, away from its default value of 1.0, and we add the weighting modification discussed in Footnote 11. Additionally, we utilize a release outcome whenever a workload represents a full computation (UCCSD, HEA, and QAOA), not a subroutine. Table I's PCOAST<sub>FT</sub> shows the results of these modifications. We see that by fine-tuning the cost function, we are able to achieve better

<sup>&</sup>lt;sup>14</sup>In both cases we add other passes like *PauliSimp* and *OptimisePhaseGadgets* when they improve t|ket⟩'s performance.



Fig. 11: Normalized gate counts vs. compile time for each framework to evaluate the scalability across different benchmarks.

performance in all cases, where PCOAST<sub>FT</sub> reduces the gate count by 50.33%, 38.24%, and 76.82% for the respective workloads, as compared to PCOAST<sub>1</sub>.

Our empirical analysis reveals that setting the parallelization credit less than 0.1 tends to achieve the best results. As we increase its value beyond 1, performance plateaus due to the dominance of the parallelization part of the cost function.

Overall, PCOAST<sub>FT</sub> reduces total gate count, two-qubit gates, and depth by 16.79%, 20.5%, and 11.28% respectively compared to PCOAST<sub>1</sub>, and by 32.53%(43.33%), 29.22%(20.58%), and 42.05%(51.27%) compared to the best Qiskit (t|ket)) performance across all benchmarks.

# D. Scalability

Ref [17] argues that the complexity of PCOAST's search algorithm is  $\mathcal{O}(N^3|G|^2)$ , where N is the number of qubits and |G| the number of nodes. Table I shows that PCOAST can be applied to circuits with up to 100 qubits. To gain a comprehensive understanding of scalability and quality, we compared compilation time and performance (gate count) across various toolchains (Fig. 11), revealing that PCOAST's scalability is superior to other frameworks. Its data points consistently reside in the lower left region, indicating better results, with a few exceptions in HEA.

#### VII. RELATED WORK

Global quantum circuit optimizations, such as phase polynomials [31, 32, 6, 33], the ZX-calculus [34, 35, 36, 12], Pauli strings [14], and Pauli rotations [11, 13], leverage mathematical structures to reduce gate count. Most focus on unitary optimizations, with the exception of some ZX variants [37, 36]. Most similar to PCOAST graphs is [11]'s DAGs of Pauli rotations and [12]'s ZX-based Pauli gadgets. Both overlap with Sec. V-A when restricted to unitary gates, but neither use Pauli frames to represent Cliffords, nor address efficient Pauli gadget synthesis. With ZX-based approaches in particular, optimizations must maintain a "circuit-like" form, as not all ZX diagrams can be directly synthesized into gates [34, 36, 38]. In contrast, all PCOAST nodes are synthesizable, as synthesis is built into the framework itself.

Bottom-up synthesis methods construct parameterized circuits by iteratively adding gates and using numerical optimization algorithms for parameter determination [39, 40, 41, 42]. Compilation algorithms like QGo [43] and QUEST [44] leverage bottom-up synthesis for larger circuits, though scalability remains a concern because their search space increases exponentially with circuit size. In contrast, PCOAST scales well to large circuits, as demonstrated in Sec. VI-D.

Schmitz et al. [13] and Li et al. [14] both address circuit synthesis from Pauli strings in the context of Hamiltonian simulation. [13] is the basis of the PCOAST synthesis algorithm, extended to support non-unitary gates and custom cost functions that allow the ultra-greedy search to be fine-tuned. Li et al. incorporate hardware-aware optimization and scheduling passes into Hamiltonian synthesis and, though out of scope of this work, we will extend PCOAST search functions with such hardware-aware considerations in the near future.

## VIII. CONCLUSION

PCOAST is a novel optimization framework for mixed unitary and non-unitary quantum circuits that adapts the commutativity properties of Cliffords and Pauli strings to preparation and measurement gates in the PCOAST graph. Internal optimizations simplify the graph depending on whether the quantum state needs to be preserved (*hold*) or can be released (*release*) after circuit execution. Finally, a customizable greedy search algorithm finds an efficient gate implementation for the optimized PCOAST graph.

Implemented in the Intel Quantum SDK, PCOAST significantly reduces gate count, two-qubit gates, and depth in key benchmarks. With minor tuning, it reduces total gate count by between 32% (resp. 43%) two-qubit gates by 29% (21%), and depth by 42% (51%) compared to the best performance of Qiskit (resp. t|ket)). On applications for quantum chemistry, it reduces gate count by 79% (62%), two-qubit gates by 77% (54%), and depth by 85% (76%).

The framework leaves many avenues for future work. PCOAST can be used as an IR beyond circuit conversion for Hamiltonian simulation [13] and higher-order circuit transformations [45]. Future internal optimizations could include more advanced unitary optimizations such as singlet node to factor node merging and incorporating other representations like phase polynomials into PCOAST. For synthesis, immediate next steps will adapt state-of-the-art methods for limited connectivity in NISQ architectures by incorporating connectivity and noise into the search functions.

#### REFERENCES

[1] F. T. Chong, D. Franklin, and M. Martonosi, "Programming languages and compiler design for realistic quantum hardware," *Nature*, vol. 549, no. 7671, pp. 180–187, 2017.

- [2] A. K. Prasad, V. V. Shende, I. L. Markov, J. P. Hayes, and K. N. Patel, "Data structures and algorithms for simplifying reversible circuits," *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, vol. 2, no. 4, pp. 277–293, 2006.
- [3] B. Valiron, N. J. Ross, P. Selinger, D. S. Alexander, and J. M. Smith, "Programming the quantum future," *Communications of the ACM*, vol. 58, no. 8, pp. 52–61, 2015.
- [4] V. Kliuchnikov and D. Maslov, "Optimization of Clifford circuits," *Physical Review A*, vol. 88, no. 5, p. 052307, 2013.
- [5] N. Abdessaied, M. Soeken, and R. Drechsler, "Quantum circuit optimization by Hadamard gate reduction," in *Reversible Computation: 6th International Conference, RC 2014.* Springer, 2014, pp. 149–162.
- [6] Y. Nam, N. J. Ross, Y. Su, A. M. Childs, and D. Maslov, "Automated optimization of large quantum circuits with continuous parameters," *npj Quantum Information*, vol. 4, no. 1, p. 23, 2018.
- [7] Qiskit contributors, "Qiskit: An open-source framework for quantum computing," 2023. [Online]. Available: https://doi.org/10.5281/zenodo.2573505
- [8] J. Pointing, O. Padon, Z. Jia, H. Ma, A. Hirth, J. Palsberg, and A. Aiken, "Quanto: Optimizing quantum circuits with automatic generation of circuit identities," 2021, arXiv:2111.11387.
- [9] M. Xu, Z. Li, O. Padon, S. Lin, J. Pointing, A. Hirth, H. Ma, J. Palsberg, A. Aiken, U. A. Acar *et al.*, "Quartz: superoptimization of quantum circuits," in *Proceedings* of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation, 2022, pp. 625–640.
- [10] S. Sivarajah, S. Dilkes, A. Cowtan, W. Simmons, A. Edgington, and R. Duncan, "t|ket> : A retargetable compiler for NISQ devices," *Quantum Science and Technology*, vol. 6, no. 1, p. 014003, nov 2020. [Online]. Available: https://doi.org/10.1088/2058-9565/ab8e92
- [11] F. Zhang and J. Chen, "Optimizing T gates in Clifford+T circuit as  $\pi/4$  rotations around Paulis," 2019, arxiv:1903.12456.
- [12] A. Cowtan, S. Dilkes, R. Duncan, W. Simmons, and S. Sivarajah, "Phase gadget synthesis for shallow circuits," in 16th International Conference on Quantum Physics and Logic 2019. Open Publishing Association, 2019, pp. 213–228.
- [13] A. T. Schmitz, N. P. D. Sawaya, S. Johri, and A. Y. Matsuura, "Graph optimization perspective for low-depth Trotter-Suzuki decomposition," 2021, arxiv:2103.08602.
- [14] G. Li, A. Wu, Y. Shi, A. Javadi-Abhari, Y. Ding, and Y. Xie, "Paulihedral: A generalized block-wise compiler optimization framework for quantum simulation kernels," in *Proceedings of the 27th ACM International Conference on Architectural Support for Programming*

Languages and Operating Systems, ser. ASPLOS '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 554–569. [Online]. Available: https://doi.org/10.1145/3503222.3507715

- [15] S. Aaronson and D. Gottesman, "Improved simulation of stabilizer circuits," *Physical Review A*, vol. 70, no. 5, nov 2004. [Online]. Available: https://doi.org/10.1103/ PhysRevA.70.052328
- [16] P. Khalate, X.-C. Wu, S. Premaratne, J. Hogaboam, A. Holmes, A. Schmitz, G. G. Guerreschi, X. Zou, and A. Y. Matsuura, "An LLVM-based C++ compiler toolchain for variational hybrid quantumclassical algorithms and quantum accelerators," 2022, arXiv:2202.11142.
- [17] J. Paykin, A. T. Schmitz, M. Ibrahim, X.-C. Wu, and A. Y. Matsuura, "PCOAST: A Pauli-based quantum circuit optimization framework (extended version)," 2023, arXiv:2305.10966v2.
- [18] M. A. Nielsen and I. Chuang, "Quantum computation and quantum information," 2002.
- [19] Y. Feng and M. Ying, "Quantum Hoare logic with classical variables," ACM Transactions on Quantum Computing, vol. 2, no. 4, pp. 1–43, 2021.
- [20] P. Selinger, "Towards a quantum programming language," *Mathematical Structures in Computer Science*, vol. 14, no. 4, pp. 527–586, 2004.
- [21] A. T. Schmitz, M. Ibrahim, N. P. D. Sawaya, G. G. Guerreschi, J. Paykin, X.-C. Wu, and A. Y. Matsuura, "Optimization at the interface of unitary and non-unitary quantum operations in PCOAST," 2023, arXiv:2305.09843.
- [22] D. Maslov and M. Roetteler, "Shorter stabilizer circuits via bruhat decomposition and quantum circuit transformations," *IEEE Transactions on Information Theory*, vol. 64, no. 7, pp. 4729–4738, 2018.
- [23] P. K. Barkoutsos, J. F. Gonthier, I. Sokolov, N. Moll, G. Salis, A. Fuhrer, M. Ganzhorn, D. J. Egger, M. Troyer, A. Mezzacapo, S. Filipp, and I. Tavernelli, "Quantum algorithms for electronic structure calculations: Particle-hole Hamiltonian and optimized wave-function expansions," *Physical Review* A, vol. 98, no. 2, aug 2018. [Online]. Available: https://doi.org/10.1103/physreva.98.022322
- [24] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, ser. STOC '96. New York, NY, USA: Association for Computing Machinery, 1996, p. 212–219. [Online]. Available: https://doi.org/10.1145/237814.237866
- [25] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, "Hardwareefficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, no. 7671, pp. 242–246, Sep. 2017. [Online]. Available: https://doi.org/10.1038/nature23879
- [26] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," 2014,

arxiv:1411.4028.

- [27] P. Jordan and E. Wigner, "Über das paulische Äquivalenzverbot," Zeitschrift für Physik, vol. 47, no. 9-10, pp. 631–651, Sep. 1928. [Online]. Available: https://doi.org/10.1007/bf01331938
- [28] S. B. Bravyi and A. Y. Kitaev, "Fermionic quantum computation," *Annals of Physics*, vol. 298, no. 1, pp. 210–226, may 2002. [Online]. Available: https: //doi.org/10.1006/aphy.2002.6254
- [29] J. T. Seeley, M. J. Richard, and P. J. Love, "The Bravyi-Kitaev transformation for quantum computation of electronic structure," *The Journal of Chemical Physics*, vol. 137, no. 22, p. 224109, Dec. 2012. [Online]. Available: https://doi.org/10.1063/1.4768229
- [30] P. Erdös and A. Rényi, "On random graphs I," *Publicationes Mathematicae Debrecen*, vol. 6, pp. 290–297, 1959.
- [31] M. Amy, D. Maslov, and M. Mosca, "Polynomial-time T-depth optimization of Clifford+ T circuits via matroid partitioning," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 33, no. 10, pp. 1476–1489, 2014.
- [32] M. Amy, "Towards large-scale functional verification of universal quantum circuits," in *Proceedings of Quantum Physics and Logic (QPL)*, 2018.
- [33] A. Meijer-van de Griend and R. Duncan, "Architectureaware synthesis of phase polynomials for NISQ devices," 2020, arxiv:2004.06052.
- [34] A. Kissinger and J. van de Wetering, "PyZX: Large scale automated diagrammatic reasoning," *Electronic Proceedings in Theoretical Computer Science*, vol. 318, pp. 229– 241, 2020.
- [35] —, "Reducing the number of non-Clifford gates in quantum circuits," *Phys. Rev. A*, vol. 102, p. 022406, Aug 2020. [Online]. Available: https://doi.org/10.1103/ PhysRevA.102.022406
- [36] N. de Beaudrap, R. Duncan, D. Horsman, and S. Perdrix, "Pauli fusion: a computational model to realise quantum transformations from ZX terms," *Electronic Proceedings in Theoretical Computer Science*, vol. 318, pp. 85–105, May 2020. [Online]. Available: https://doi.org/10.4204/eptcs.318.6
- [37] A. Borgna, S. Perdrix, and B. Valiron, "Hybrid quantumclassical circuit simplification with the ZX-calculus," in *Programming Languages and Systems: 19th Asian Symposium, APLAS 2021, Chicago, IL, USA, October* 17–18, 2021, Proceedings 19. Springer, 2021, pp. 121– 139.
- [38] N. de Beaudrap, A. Kissinger, and J. van de Wetering, "Circuit extraction for ZX-diagrams can be #P-hard," 2022, arxiv:2202.09194.
- [39] M. G. Davis, E. Smith, A. Tudor, K. Sen, I. Siddiqi, and C. Iancu, "Towards optimal topology aware quantum circuit synthesis," in 2020 IEEE International Conference on Quantum Computing and Engineering (QCE). IEEE, 2020, pp. 223–234.

- [40] E. Younis, K. Sen, K. Yelick, and C. Iancu, "Qfast: Conflating search and numerical optimization for scalable quantum circuit synthesis," in 2021 IEEE International Conference on Quantum Computing and Engineering (QCE). IEEE, 2021, pp. 232–243.
- [41] P. Rakyta and Z. Zimborás, "Efficient quantum gate decomposition via adaptive circuit compression," 2022, arxiv:2203.04426.
- [42] L. Madden and A. Simonetto, "Best approximate quantum compiling problems," ACM Transactions on Quantum Computing, vol. 3, no. 2, pp. 1–29, 2022.
- [43] X.-C. Wu, M. G. Davis, F. T. Chong, and C. Iancu, "Reoptimization of quantum circuits via hierarchical synthesis," in 2021 International Conference on Rebooting Computing (ICRC). IEEE, 2021, pp. 35–46.
- [44] T. Patel, E. Younis, C. Iancu, W. de Jong, and D. Tiwari, "Quest: systematically approximating quantum circuits for higher output fidelity," in *Proceedings of the 27th* ACM International Conference on Architectural Support for Programming Languages and Operating Systems, 2022, pp. 514–528.
- [45] A. Schmitz, J. Paykin, and A. Matsuura, "A functional approach to the modular construction of quantum logic," *Bulletin of the American Physical Society*, 2023.