

Programmable Basis-Selective Measurement for Photonic Qudits Using a Single-Qudit Gate

Mitchell A. Thornton, Hiva Shahoei, Cooper Shapard, R. Mason Tuller and Duncan L. MacFarlane

Darwin Deason Institute for Cyber Security

Department of Electrical and Computer Engineering

Southern Methodist University

Dallas, Texas, USA

{mitch,hshahoei,cshapard,mtuller,dmacfarlane}@smu.edu

Abstract—A programmable measurement primitive for frequency-bin photonic qudits is described. The method places a programmable single-qudit gate immediately ahead of a fixed photo-detection stage and programs the gate to realize the inverse transform associated with a desired measurement basis, converting a detector that is native to the computational basis into a reconfigurable basis-selective measurement device. We use a reconfigurable single-qudit gate as a starting point and derive measurement operators, quantify expected measurement fidelity and robustness, and demonstrate through simulation that the architecture can realize measurement in a complete set of mutually unbiased bases (MUBs) for dimension $d = 4$. The result is a hardware-oriented path toward programmable qudit readout for integrated photonic quantum information processing applications such as QKD and high-dimensional cluster state processing.

Index Terms—frequency-bin qudits, quantum photonics, programmable measurement, mutually unbiased bases, cluster states, measurement-based quantum computing, electro-optic modulation

I. INTRODUCTION

Frequency-bin quantum photonics offers an appealing route to high-dimensional quantum information processing because a single spatial mode can support a plurality of computational basis states distinguished only by optical frequency. This permits dense encoding, compatibility with fiber transport, and direct alignment with integrated frequency-comb sources [1]–[4]. At the same time, working with multicolor photon states imposes a special burden on the measurement hardware. The basis states are colors rather than paths or polarizations, so arbitrary projective measurements require coherent movement of probability amplitudes among spectral lines before the photo-detector performs its native basis discrimination.

This requirement becomes particularly important in measurement-based quantum computing (MBQC). In the one-way or cluster-state model, computational progress is driven by a sequence of projective measurements whose bases are selected according to the intended algorithm and, in general, conditioned on earlier outcomes [5], [6]. Thus, basis selection is not incidental to the computation; it is the mechanism by which the computation proceeds. Recent work has strengthened this motivation in the photonic setting. Optical frequency combs have long been viewed as attractive cluster-state resources [7]–[9], and very recent experimental results

have shown resource-efficient photonic quantum computation using high-dimensional cluster states, indicating that high-dimensional photonic MBQC is now an active experimental direction rather than only a theoretical possibility [10].

The purpose of this paper is to show that a programmable single-qudit gate can serve directly as a basis-selective measurement front end when placed immediately ahead of photo-detection. The gate architecture uses a demultiplexer (DeMUX), a parallel bank of nested Mach-Zehnder modulators (MZMs), and a re-multiplexer (ReMUX) to realize a programmable single-qudit operation with favorable architectural scaling to higher dimension and with a favorable loss structure relative to serial frequency processors. Here, we treat the gate not simply as a state-evolution element but as a measurement primitive that converts a fixed detector into a reconfigurable analyzer capable of measuring with respect to a selected qudit basis. To substantiate this claim beyond the Chrestenson gate reported in [1], we simulate the architecture programmed to realize each of the four nontrivial mutually unbiased bases (MUBs) for dimension $d = 4$, thereby demonstrating that the measurement front end supports a complete MUB set sufficient for quantum state tomography and higher-dimensional quantum key distribution (QKD).

We note that the observation that a pre-measurement unitary can rotate an arbitrary basis into the detector-native basis is itself well known [6], [11], and that programmable high-dimensional measurement has been demonstrated in other photonic degrees of freedom, notably the spatial domain [10]. Indeed, because linear optics is reciprocal, any programmable state-preparation circuit is in principle a programmable analyzer when traversed toward the detector. The contribution of this paper is therefore not the reversal itself, but the device-level case that a specific parallel frequency-comb qudit gate survives the reversal in practice: a complete, informationally significant family of measurement bases is reachable at high fidelity under realistic device parameterization, with quantified sensitivity to the dominant error channels, and at an optical depth that is fixed as the dimension grows. Because the gate accommodates increasing dimension d by adding parallel branches at fixed optical depth and with minimal added loss, it is architecturally well-suited for applications that process high-dimensional single qudits or cluster states that require

arbitrary measurement bases; the engineering constraints that accompany this scaling are examined in Section IX.

II. BACKGROUND AND RELATED WORK

The recent literature on frequency-bin quantum information establishes both the promise of the encoding and the difficulty of general control. Kues *et al.* demonstrated on-chip generation of high-dimensional entangled quantum states together with coherent control [3]. In other work, Lu *et al.* reviewed the use of frequency-comb qudits for quantum information processing and emphasized the role of electro-optic modulation in transferring amplitudes among bins [2]. A comprehensive treatment of the frequency-bin paradigm, covering source generation, gate architectures, and applications, is provided in [12]. Experimental realizations of frequency-domain beam splitters and tritters, along with broader quantum frequency processor architectures, have shown that high-fidelity transformations are possible but generally require carefully orchestrated combinations of modulation and filtering [13]–[15]. Lu *et al.* demonstrated fully arbitrary control of frequency-bin qubits with fidelities exceeding 0.98 across the full Bloch sphere [16], and subsequently realized high-dimensional discrete Fourier transform gates for $d = 4$ in the serial QFP architecture [17]. More recently, programmable nano-engineered devices have shown that frequency-bin state generation and state control can be pushed onto a chip [18], while programmable qudit photonic processors have highlighted the value of monolithic integration of preparation, transformation, and projection [19].

A parallel thread in the literature concerns programmable projective measurement. Kumar *et al.* analyzed optimal quantum-programmable projective measurements with coherent states [20], Bouchard *et al.* proposed efficient arbitrary-basis measurement for ultrafast time-bin qudits [11], and Feng *et al.* recently demonstrated a programmable higher-dimensional measurement on an integrated photonic platform [21]. Programmable mode-selective projection has also been realized through engineered nonlinear frequency conversion: the quantum pulse gate developed in the Silberhorn group projects onto programmable temporal-mode superpositions [22], [23] and has been applied to high-dimensional measurement tasks including temporal-mode QKD [24]. The present work pursues the same operational goal for frequency-bin encodings but through a linear, parallel electro-optic architecture, trading the mode selectivity of nonlinear conversion for direct RF programmability and fixed optical depth. These results collectively make clear that measurement programmability is itself a meaningful hardware objective; however, they do not address the specific issue of a frequency-bin qudit front end built from an architecture that already supports programmable single-qudit gates.

The MBQC literature further clarifies why such a front end is useful. Raussendorf and Briegel introduced the one-way quantum computer, and Raussendorf, Browne, and Briegel provided the detailed cluster-state model in which adaptive local measurements realize computation [5], [6]. Menicucci *et al.* subsequently showed that optical frequency combs,

a coherent comb of discrete, high quality factor frequency components separated by a constant frequency usually termed the Free Spectral Range (FSR), could support one-way quantum computing resources [7]. Chen, Menicucci, and Pfister reported multipartite entanglement of 60 modes in a quantum optical frequency comb [8], while Roslund *et al.* demonstrated wavelength-multiplexed quantum networks with ultrafast frequency combs [9]. Most recently, Lib and Bromberg demonstrated resource-efficient photonic quantum computation with high-dimensional cluster states and showed that high-dimensional encoding can reduce resource overhead and support instantaneous intra-photon feed-forward [10]. For the present work, these results collectively imply that a compact programmable measurement layer for high-dimensional photonic states is timely and potentially enabling.

III. MEASUREMENT PRINCIPLE

Let the detector-native basis of a d -dimensional frequency-bin receiver be

$$\mathcal{B}_0 = \{|0\rangle, |1\rangle, \dots, |d-1\rangle\}, \quad (1)$$

where each basis vector corresponds to a resolved frequency bin incident on a detector channel. Suppose the desired measurement basis is

$$\mathcal{B}_\phi = \{|\phi_k\rangle : |\phi_k\rangle = \mathbf{U}|k\rangle, k = 0, \dots, d-1\}, \quad (2)$$

for some unitary transformation $\mathbf{U}$. Then measurement with respect to $\mathcal{B}_\phi$ is described by the projectors

$$\Pi_k = |\phi_k\rangle\langle\phi_k| = \mathbf{U}|k\rangle\langle k|\mathbf{U}^\dagger. \quad (3)$$

If an input state ρ is first evolved by $\mathbf{U}^\dagger$ and then detected in the native basis, the probability of outcome k is

$$p(k) = \text{Tr}[|k\rangle\langle k|\mathbf{U}^\dagger\rho\mathbf{U}] \quad (4)$$

$$= \text{Tr}[\mathbf{U}|k\rangle\langle k|\mathbf{U}^\dagger\rho] \quad (5)$$

$$= \text{Tr}[\Pi_k\rho]. \quad (6)$$

The cyclic property of the trace is used in the step from (4) to (5). Thus, the pre-measurement unitary $\mathbf{U}^\dagger$ converts a fixed computational-basis detector into a basis-selective detector.

This result acquires practical significance when an efficient hardware platform exists for realizing the required $\mathbf{U}^\dagger$ over a chosen family of qudit bases. In the present case, the programmable single-qudit gate supplies that functionality. The detector itself remains fixed. The measurement basis is specified instead through the RF programming of the electro-optic gate immediately ahead of the detector. Because the architecture is electrically programmable, one may update the basis from shot to shot or between successive operations subject to the speed of the control electronics.

IV. GATE ARCHITECTURE

The parallel frequency-comb qudit processor described in [1] provides the hardware platform for the measurement front end proposed here. Because that work is a conference paper that may not yet be accessible to all readers, we summarize the architecture in sufficient detail to make the present contribution self-contained.




Fig. 1: Parallel qudit quantum frequency processor architecture from [1]. A multicolor photon qudit input state is demultiplexed into its N constituent frequency components. Each component is processed by a parallel nested MZM branch that mixes energy among frequency bins. The processed components are remultiplexed to produce the output state $|\Psi_1\rangle = \mathbf{U}|\Psi_0\rangle$.

A. Structure

The gate comprises three stages as illustrated in Fig. 1. First, a DeMUX separates the constituent frequency components of the incoming multicolor photon. For a d -dimensional qudit whose basis states correspond to d frequency bins spaced at multiples of a free spectral range (FSR) $\Delta\omega$, the DeMUX routes each frequency component to a distinct parallel branch. Second, each branch contains a nested Mach-Zehnder modulator (MZM) that applies a programmable electro-optic mixing operation. As detailed in Fig. 2, a nested MZM consists of two MZIs arranged inside a larger interferometric structure, with each arm containing a phase modulator (PM) driven by an RF signal $p_{RF}(t)$ followed by a phase shifter (PS) biased by a DC voltage $p_{DC}(t)$. This nested configuration provides more precise control of phase shifts and interference patterns than a single-arm phase modulator or a basic two-arm MZM. Third, a ReMUX recombines the processed frequency components into a single output spatial mode.

Fig. 2 illustrates the internal structure of the electro-optic mixer/modulator elements that comprise each parallel branch. The top portion shows the abstract “EO Mixer/Modulator” block as it appears in Fig. 1. Below it, three levels of implementation detail are shown. A single-arm electro-optic phase modulator (PM) followed by a DC-biased phase shifter (PS) constitutes the basic building block, where $p_{RF}(t)$ is the RF drive signal and $p_{DC}(t)$ is the bias voltage. Two such arms arranged as a Mach-Zehnder interferometer (left) form a basic MZM that provides interference-based amplitude and phase control. The nested MZM structure used in the present work (right) comprises four arms arranged as two MZI pairs within a larger interferometric structure. This nested configuration provides 26 control parameters per branch: twelve modulation depths, twelve drive phases, and two bias phases. Both MZM and nested MZMs are widely used in optical communications, and a plurality of these may be integrated together on a single




Fig. 2: Internal structure of the electro-optic mixer/modulator branches. Top: abstract “EO Mixer/Modulator” block from Fig. 1. Left column: a single-arm phase modulator–phase shifter (PM–PS) chain, and a basic two-arm MZM. Right: the nested MZM used in this work, comprising four PM–PS arms arranged as two MZI pairs. Each arm is driven by an RF signal $p_{RF}(t)$ and a DC bias $p_{DC}(t)$, yielding 26 programmable parameters per branch. Importantly, the phase shift of a branch due to the optical transit time is inconsequential in the processing of the signal unless there are different transit times from branch to branch caused, perhaps, by the finite resolution of the fabrication tolerancing. Such fabrication tolerances may be compensated, after calibration, by an additional voltage on the two bias phases.

substrate.

B. Parameterization and Control

Each nested MZM branch is driven by RF signals that can be represented as a truncated Fourier series,

$$m(t) = \sum_{n=1}^{d-1} a_n \cos(\omega_{m,n}t + \phi_n), \quad (7)$$

where $\omega_{m,n} = n\delta$ with δ equal to the FSR of the frequency comb. The modulation depths a_n , drive phases ϕ_n , and bias phases of the nested MZM collectively serve as the control variables. For a four-dimensional ($d = 4$) gate, each branch has 26 parameters: twelve modulation depths (three per arm across four arms), twelve drive phases, and two bias phases. The full four-branch gate thus requires 104 parameters. These parameters are numerically optimized to realize a target unitary transfer matrix.

C. Optimization

The control parameter values are obtained through particle swarm optimization (PSO) that minimizes a cost function comprising two terms. The in-band cost c_{in} is the Frobenius norm between the target transfer matrix and the circuit transfer matrix as given in (8),

$$c_{in} = \|\mathbf{T}_{\text{target}} - \mathbf{T}_{\text{circ}}\|^2, \quad (8)$$

while the out-of-band cost c_{out} penalizes spectral leakage into undesired sidebands. The total PSO cost is $c_{\text{pso}} = c_{\text{in}} + c_{\text{out}}$. As reported in [1], the optimization achieved gate fidelities exceeding 0.995 for the four-dimensional Chrestenson operator C_4 , where fidelity is defined using the standard expression in terms of density matrices ρ and σ ,

$$F(\rho, \sigma) = \left(\text{Tr} \sqrt{\sqrt{\rho} \sigma \sqrt{\rho}} \right)^2. \quad (9)$$

D. Favorable Properties

Three properties of this architecture are relevant to its use as a measurement front end. First, the parallel arrangement reduces the serial optical depth relative to the cascaded EOM-plus-filter stages of the serial quantum frequency processor (QFP) approach [13], [25]. In a serial architecture, parasitic coupling losses accumulate multiplicatively with each stage, whereas the parallel arrangement distributes a single coupling loss across each branch and recombines the branches only once at the output. Second, modulation and filtering are combined within the MZM structures rather than assigned to separate serial stages, reducing component count. Third, scaling to higher qudit dimension d is accomplished by adding parallel branches rather than fundamentally redesigning the gate cell.

V. GATE-BASED MEASUREMENT FRONT END

Viewed as a measurement front end, the parallel gate architecture described in Section IV performs the inverse basis transformation required by Section III. If a measurement in basis $\mathcal{B}_\phi$ is desired, the gate is programmed to implement $\mathbf{U}^\dagger$ such that $\mathbf{U}|k\rangle = |\phi_k\rangle$. The detector bank then reports the classical outcome corresponding to the native basis state $|k\rangle$, which is identical to the outcome of a projective measurement in $\mathcal{B}_\phi$.

This interpretation suggests three practical arrangements.

- 1) A *detector-bank arrangement* in which the ReMUX output is spectrally resolved into d detector channels. This directly realizes the full projective measurement $\{\Pi_k\}$.
- 2) A *single-port heralding arrangement* in which the gate is programmed so that only one selected projector is mapped onto a monitored output, useful for binary accept/reject tests or adaptive hypothesis discrimination.
- 3) A *time-sequenced adaptive arrangement* in which the gate settings are updated between successive detections according to classical feed-forward logic, appropriate to MBQC and adaptive tomography.

The first arrangement is the most natural for qudit MBQC. The second may be useful when only one basis projector is required at a given stage. The third highlights the role of the gate as a measurement compiler that translates a logical basis choice into physical detector-basis coordinates.

We show the architecture is flexible with respect to the family of accessible measurement bases. The C_4 Chrestenson basis is one concrete target due to previous work that reports a high-fidelity realization. More generally, any basis connected

to the computational basis through a realizable single-qudit unitary may be addressed. This includes mutually unbiased bases (MUBs), task-specific eigenbases, and bases selected adaptively during tomography, discrimination, or MBQC. For $d = 4$, there exist five MUBs [26]; the extent to which all five are reachable with the nested-MZM architecture at high fidelity is addressed in Sections VI and VIII. The principal limitations on the set of accessible bases are practical: achievable modulation bandwidth, phase stability, insertion loss, and calibration complexity.

VI. MUTUALLY UNBIASED BASES FOR $d = 4$

The family of measurement bases that the front end must support determines the range of quantum information tasks it can address. Among the most important families for quantum informatics are the mutually unbiased bases (MUBs). In this section we define MUBs, summarize why they are significant, and construct the complete set for $d = 4$ that serves as the simulation target in Section VIII.

A. Definition

Two orthonormal bases $\mathcal{B}_a = \{|a_j\rangle\}$ and $\mathcal{B}_b = \{|b_k\rangle\}$ of a d -dimensional Hilbert space are *mutually unbiased* if and only if

$$|\langle a_j | b_k \rangle|^2 = \frac{1}{d} \quad \forall j, k. \quad (10)$$

In other words, if a system is prepared in any basis state of $\mathcal{B}_a$ and then measured in $\mathcal{B}_b$, every outcome is equally likely. The two bases extract maximally complementary information about the quantum state.

B. Significance

MUBs arise in several core quantum information tasks. For quantum state tomography, a complete set of $d + 1$ MUBs provides exactly $d^2 - 1$ independent measurement outcomes, which is the minimum required to reconstruct an arbitrary d -dimensional density matrix [27], [28]. This makes a complete MUB set a minimal informationally complete measurement and thus the optimal strategy for tomography in terms of minimizing statistical error for a given number of measurements. In quantum key distribution, MUBs underpin the security of the *BB84* protocol and its higher-dimensional generalizations: an eavesdropper who measures in the wrong basis gains no information due to the uniform overlap condition of (10), and higher-dimensional MUB-based QKD protocols achieve improved key rates and noise tolerance [29]. MUBs also saturate entropic uncertainty relations that are foundational to security proofs, and they appear naturally in MBQC as the measurement families that implement different logical operations on encoded information [6], [30].

C. Maximum Number

For a Hilbert space of dimension d , the maximum number of pairwise mutually unbiased bases is $d + 1$. This bound is achieved whenever d is a prime or a prime power, and explicit constructions are known for those cases [27], [31]. For $d = 4 = 2^2$, a prime power, there exist exactly five MUBs.

D. Construction for $d = 4$

We construct the complete set of five MUBs for $d = 4$ using the stabilizer formalism for two-qubit systems. Identifying the four-dimensional basis states with the two-qubit computational basis via $|0\rangle \leftrightarrow |00\rangle$, $|1\rangle \leftrightarrow |01\rangle$, $|2\rangle \leftrightarrow |10\rangle$, $|3\rangle \leftrightarrow |11\rangle$, the 15 nontrivial two-qubit Pauli operators can be partitioned into five groups of three mutually commuting operators [31], [32]. The simultaneous eigenbasis of each group defines one MUB. The five groups and their associated unitary matrices $\mathbf{U}_m$ (such that $\mathcal{B}_m = \{\mathbf{U}_m|k\rangle : k = 0, \dots, 3\}$) are as follows.

Group 0 ($\sigma_z \otimes \mathbf{I}$, $\mathbf{I} \otimes \sigma_z$, $\sigma_z \otimes \sigma_z$): The simultaneous eigenbasis is the computational basis itself, so $\mathbf{U}_0 = \mathbf{I}_4$.

Group 1 ($\sigma_x \otimes \mathbf{I}$, $\mathbf{I} \otimes \sigma_x$, $\sigma_x \otimes \sigma_x$): The eigenbasis is the tensor product of single-qubit Hadamard eigenstates. The corresponding unitary is

$$\mathbf{U}_1 = \mathbf{H} \otimes \mathbf{H} = \frac{1}{2} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{bmatrix}. \quad (11)$$

Group 2 ($\sigma_x \otimes \sigma_z$, $\sigma_z \otimes \sigma_y$, $\sigma_y \otimes \sigma_x$): This group produces an entangled basis. Simultaneous diagonalization yields

$$\mathbf{U}_2 = \frac{1}{2} \begin{bmatrix} 1 & 1 & 1 & 1 \\ i & -i & i & -i \\ 1 & 1 & -1 & -1 \\ -i & i & i & -i \end{bmatrix}. \quad (12)$$

Group 3 ($\sigma_z \otimes \sigma_x$, $\sigma_y \otimes \sigma_z$, $\sigma_x \otimes \sigma_y$): This group also produces an entangled basis,

$$\mathbf{U}_3 = \frac{1}{2} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ i & -i & i & -i \\ -i & i & i & -i \end{bmatrix}. \quad (13)$$

Group 4 ($\sigma_y \otimes \mathbf{I}$, $\mathbf{I} \otimes \sigma_y$, $\sigma_y \otimes \sigma_y$): The eigenbasis is the tensor product of single-qubit $\mathbf{Y}$ -eigenstates. Defining the single-qubit rotation $\mathbf{R} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ i & -i \end{bmatrix}$, we obtain

$$\mathbf{U}_4 = \mathbf{R} \otimes \mathbf{R} = \frac{1}{2} \begin{bmatrix} 1 & 1 & 1 & 1 \\ i & -i & i & -i \\ i & i & -i & -i \\ -1 & 1 & 1 & -1 \end{bmatrix}. \quad (14)$$

All ten pairwise overlaps $|\langle \psi_j^{(m)} | \psi_k^{(n)} \rangle|^2 = \frac{1}{4}$ for $m \neq n$ have been verified numerically. We note that the four-dimensional Chrestenson operator $\mathbf{C}_4$ reported in [1] is mutually unbiased with respect to the computational basis $\mathcal{B}_0$ and with $\mathcal{B}_2$, but it does not satisfy the mutual unbiasedness condition with $\mathcal{B}_1$, $\mathcal{B}_3$, or $\mathcal{B}_4$, and is therefore not a member of this particular complete MUB set. The $\mathbf{C}_4$ results thus represent a sixth, independent unitary target that complements the MUB set.

VII. MEASUREMENT ROLE IN CLUSTER-STATE PROCESSING

In cluster-state or one-way quantum computing, computation proceeds by preparing a highly entangled resource state and then performing a sequence of single-particle measurements. For qubits, the measurement basis may be parameterized by a rotation angle on the Bloch sphere. For qudits, the corresponding measurement family occupies a larger manifold, but the operational principle is the same: the chosen basis implements the desired logical step, up to known byproduct operators that are handled classically or through subsequent adaptive choices [6], [30].

A frequency-bin cluster-state processor therefore requires a measurement unit that can repeatedly perform the following tasks: (i) select a prescribed basis, (ii) measure the incident qudit with respect to that basis, and (iii) update the next basis in light of previous outcomes. The proposed gate-plus-detector structure supports each of these tasks. The selected basis is encoded into the control settings of the single-qudit gate. The photodetection stage provides the projective outcome. The control system may then update the gate parameters before the next qudit arrives. In this sense, the present architecture supplies the qudit analog of a reconfigurable readout layer.

The relevance of this point has increased due to recent experiments on high-dimensional photonic cluster states. Lib and Bromberg showed that multiple logical qubits may be encoded on a single photon through high-dimensional spatial structure, reducing resource requirements and allowing instantaneous feed-forward within a photon [10]. Their implementation did not use frequency bins, but it demonstrates that high-dimensional cluster-state photonics is a practical and current research topic. A programmable frequency-bin measurement front end directly addresses the analogous basis-selection problem for frequency-comb qudits.

VIII. RESULTS AND PERFORMANCE INTERPRETATION

A. Simulation Methodology

To evaluate the measurement front end across the complete MUB set, we program the parallel nested-MZM architecture to realize each of the four nontrivial MUB unitaries $\mathbf{U}_1$ through $\mathbf{U}_4$ defined in Section VI, using the same PSO-based optimization methodology described in Section IV-C and reported in [1]. For each target unitary $\mathbf{U}_m$, the gate is programmed to implement $\mathbf{U}_m^\dagger$, which is the premeasurement transform required to convert a native computational-basis detector into a $\mathcal{B}_m$ -selective measurement device. The PSO minimizes the cost function $c_{\text{pso}} = c_{\text{in}} + c_{\text{out}}$ over the 26 parameters of each nested MZM branch (104 parameters total for the four-branch gate). Gate fidelity is computed for each target using the column-averaged metric

$$\bar{F} = \frac{1}{d} \sum_{k=0}^{d-1} |\langle \psi_{k,\text{ideal}} | \psi_{k,\text{circ}} \rangle|^2, \quad (15)$$

where $|\psi_{k,\text{ideal}}\rangle$ and $|\psi_{k,\text{circ}}\rangle$ are the k -th column vectors of the target and realized transfer matrices, respectively. A

TABLE I: Gate fidelity and Monte Carlo sensitivity for each MUB target unitary and the Chrestenson operator. Fidelities are lower bounds over the PSO optimization outcomes. Sensitivities are mean fidelity changes over 100,000 trials (Gaussian parameter noise, $\sigma = 2\%$ of nominal), reported for perturbation of all parameters and of phase parameters only.

Target	Type	Fidelity $\bar{F}$	Sensitivity $\Delta\bar{F} (\times 10^{-3})$	
			All params	Phase only
C_4	Fourier (DFT)	> 0.995 [1]	-22.7	-15.5
U_1	Product ($\mathbf{H} \otimes \mathbf{H}$)	> 0.976	-62.6	-51.6
U_2	Entangled	> 0.988	-33.9	-28.3
U_3	Entangled	> 0.980	-25.6	-15.3
U_4	Product ($\mathbf{R} \otimes \mathbf{R}$)	> 0.991	-54.3	-25.8

Monte Carlo sensitivity analysis is also performed for each MUB target, following the same protocol as in [1]: each of the 104 optimized parameters is independently perturbed with additive Gaussian noise (zero mean, standard deviation equal to 2% of the nominal value) over 100,000 trials. The resulting fidelity distributions characterize the robustness of each MUB measurement to parameter variations.

Including the C_4 Chrestenson gate already reported in [1], the simulation campaign covers five nontrivial unitary targets: C_4 plus the four MUB unitaries U_1 through U_4 . Together with the trivial computational-basis measurement (requiring no gate), these targets demonstrate the architecture's capacity for programmable measurement across a complete set of informationally significant bases.

B. MUB Measurement Fidelities

The gate fidelities obtained for each MUB target are reported in Table I.

C. MUB Unitaries Sensitivity Analysis

A Monte Carlo sensitivity analysis was performed on the optimized EOM parameters for each of the four MUB unitaries to quantify robustness to fabrication and control variations by the same protocol in [1]. For each MUB, 100,000 independent perturbations were applied to (i) all parameters, (ii) phase parameters only, and (iii) modulation-strength parameters only. The resulting distributions of gate fidelity change are shown in the sensitivity plots for MUB1–MUB4 in Figure 3 (labeled according to the corresponding unitary U_k).

Substantial variation in sensitivity is observed across the four bases. The implementation of MUB1 ($U_1 = \mathbf{H} \otimes \mathbf{H}$), whose matrix consists entirely of real-valued $\pm 1/2$ entries, exhibits the highest sensitivity, with an average fidelity change of -62.6×10^{-3} (standard deviation 32.6×10^{-3}) when all parameters are perturbed. This degradation is overwhelmingly driven by the phase parameters (-51.6×10^{-3} , std. dev. 31.8×10^{-3}), while modulation-strength parameters contribute only -4.07×10^{-3} . MUB4 follows with an intermediate sensitivity of -54.3×10^{-3} (all parameters), again dominated by phases. In contrast, MUB2 and MUB3 are progressively more robust, with all-parameter averages of -33.9×10^{-3} and $-25.6 \times$

10^{-3} , respectively (phase-only: -28.3×10^{-3} and -15.3×10^{-3}).

Across all four MUBs, phase perturbations dominate fidelity loss by roughly an order of magnitude over modulation-strength variations, whose impact remains small and comparable (≈ -1.5 to -4×10^{-3}). This variation arises directly from the structural differences among the target unitaries. U_1 relies exclusively on relative phases of 0 or π (sign flips) to produce its characteristic pattern of constructive and destructive interference. Small phase deviations therefore disrupt the precise cancellations more severely. The remaining bases (U_2 , U_3 , and U_4) incorporate explicit factors of i and $-i$ (quadrature $\pi/2$ phase shifts) that appear to align with EOM operating points where the fidelity landscape is flatter, rendering the implementations inherently more tolerant to small phase errors. The histograms confirm that fidelity reductions remain modest in all cases, with distributions peaked near zero and only modest negative tails, underscoring the practical viability of the architecture as a pre-measurement basis selector while highlighting the need for tighter phase stabilization when deploying MUB1.

D. C_4 Sensitivity Baseline

The sensitivity analysis reported in [1] for the C_4 gate provides a baseline for interpreting the MUB results. Fig. 4 shows the distribution of gate fidelity changes over 100,000 trials when all parameters are perturbed, with an embedded table containing the average fidelity change and standard deviation for all parameters and for two parameter subsets: phase parameters only and modulation-strength parameters only. Perturbing all parameters produces a mean fidelity change of -22.7×10^{-3} ; phase errors alone produce approximately -15.5×10^{-3} with a minimum fidelity near 0.899, whereas modulation-strength errors produce only a mean change of -0.8×10^{-3} with minimum fidelity remaining near 0.995. This clear asymmetry between phase and amplitude sensitivity has a direct implication for measurement hardware: basis-selective readout should prioritize RF phase stability, calibration, and thermal control before investing disproportionate effort in amplitude trimming.

E. Measurement Fidelity Interpretation

To interpret gate fidelity as a measurement figure of merit, consider an intended basis $\mathcal{B}_\phi$ with ideal pre-measurement transform $U^\dagger$ and an implemented transform $\tilde{U}^\dagger$. The effective measurement operators become

$$\tilde{\Pi}_k = \tilde{U}|k\rangle\langle k|\tilde{U}^\dagger. \quad (16)$$

The probability of obtaining outcome k when the true projector is Π_k and the state is ρ is $\tilde{p}(k) = \text{Tr}[\tilde{\Pi}_k \rho]$. The statistical distance between the ideal and realized outcome distributions is bounded by the operator distance between the corresponding POVM elements. Specifically, for each outcome k ,

$$|p(k) - \tilde{p}(k)| = |\text{Tr}[(\Pi_k - \tilde{\Pi}_k)\rho]| \leq \|\Pi_k - \tilde{\Pi}_k\|_1, \quad (17)$$





(a) U_1 (b) U_2 (c) U_3 (d) U_4

Fig. 3: Monte Carlo sensitivity analysis of the four MUB implementations. Each subplot shows the distribution of gate fidelity change for 100,000 independent parameter perturbations (Gaussian noise, $\sigma = 2\%$ of nominal). The embedded tables report the average fidelity change and standard deviation when perturbing (top row) all parameters, (middle row) phase parameters only, and (bottom row) modulation-strength parameters only.

where $\|\cdot\|_1$ denotes the trace norm. Because the projectors Π_k and $\tilde{\Pi}_k$ are both rank-one, their trace-norm distance can be expressed as

$$\|\Pi_k - \tilde{\Pi}_k\|_1 = 2\sqrt{1 - |\langle\phi_k|\tilde{\phi}_k\rangle|^2}, \quad (18)$$

where $|\tilde{\phi}_k\rangle = \tilde{\mathbf{U}}|k\rangle$. The column-averaged fidelity $\bar{F}_{\text{meas}} = \frac{1}{d} \sum_k |\langle\phi_k|\tilde{\phi}_k\rangle|^2$ reported in Table I thus provides a direct bound on measurement quality. For $\bar{F}_{\text{meas}} > 0.995$ (as achieved for C_4), the average trace-norm error per projector is bounded by $2\sqrt{1 - 0.995} \approx 0.141$, indicating that individual measurement outcome probabilities deviate from ideal by at most approximately 14% in the worst case. In practice the deviation is substantially smaller since the bound in (17) is over all states ρ .

F. Comparison with Serial Architecture

Table II compares the parallel architecture with the serial QFP approach that has been the predominant architecture for frequency-bin unitary operations [13], [25], [33].

The serial QFP achieves universality through alternating layers of electro-optic modulation and Fourier-transform pulse shaping, with each additional layer adding optical depth and coupling loss. The parallel architecture trades this serial depth for spatial parallelism: each frequency component is processed independently in its own branch. For the measurement application, this distinction is especially important because projective measurement is typically the final stage in a protocol and any avoidable front-end loss directly reduces the detection probability.

The serial QFP has been experimentally validated for fully




Fig. 4: Histogram of C_4 gate fidelity changes over 100,000 Monte Carlo trials with all parameters perturbed (Gaussian noise, $\sigma = 2\%$ of nominal). The embedded table reports the average fidelity change and standard deviation when perturbing key parameter subsets.

TABLE II: Comparison of parallel and serial architectures for frequency-bin measurement.

Property	Serial QFP	Parallel (this work)
Optical depth	Grows with number of EOM/filter stages	Fixed at one DeMUX + one MZM + one ReMUX per branch
Loss scaling	Multiplicative across cascaded stages	Single coupling loss per branch; combined once at output
Mixing and filtering	Separate serial stages	Combined within each MZM
Scaling to d	Additional serial stages increase depth	Additional parallel branches; unchanged depth
Programming	Interleaved EOM drives and filter settings	Independent per-branch RF parameter sets

arbitrary qubit control with fidelity exceeding 0.98 [16] and for high-dimensional DFT gates at $d = 4$ [17]. The present work achieves comparable fidelities (> 0.976 for all MUBs) in simulation for the parallel architecture. While the serial architecture has the advantage of experimental maturity, the parallel architecture offers the structural advantages summarized in Table II, which become increasingly significant as d grows and as measurement-stage loss becomes a critical constraint in protocols requiring high detection efficiency.

IX. PRACTICAL CONSIDERATIONS AND SCALING

A. Frequency Plan and Efficiency

Two practical constraints named by the frequency-bin literature deserve explicit treatment: bin spacing and efficiency. The RF drive of (7) places tones at integer multiples $n\delta$ of the comb FSR up to $(d-1)\delta$, so the FSR is bounded above by the usable electro-optic and RF-chain bandwidth divided by $d-1$. The FSR is bounded below by the spectral resolvability

of the DeMUX/ReMUX and of the detection stage: adjacent bins must be separated cleanly enough that inter-bin crosstalk and unfiltered sideband leakage remain small relative to the target fidelities. The out-of-band cost term c_{out} in the PSO objective already penalizes leakage into undesired sidebands at design time; residual crosstalk in fabricated hardware adds to the phase-dominated error budget quantified in Section VIII.

Regarding efficiency, the parallel arrangement fixes the insertion budget seen by a photon at three elements, one DeMUX, one nested MZM branch, and one ReMUX, independent of dimension, whereas serial quantum frequency processors accumulate the insertion loss of $\mathcal{O}(d)$ modulator and pulse-shaper stages multiplicatively [13], [25]. End-to-end measurement efficiency additionally includes the detector: for QKD key rates and for heralded MBQC operation, the detection efficiency multiplies the front-end transmission, so the fixed optical depth of the parallel gate translates directly into a higher efficiency ceiling at a given component quality. Calibration is the remaining practical burden: each branch carries 26 programmable parameters, and branch-to-branch transit-time mismatch arising from fabrication tolerances is compensated after calibration through the two bias phases, as noted in Fig. 2.

B. Scaling with Dimension

Although the numerical optimization reported here is carried out at $d = 4$, the architecture is in principle dimension independent: increasing the dimension is accomplished by replicating parallel processing branches rather than by lengthening a cascade of optical components. The architectural cost is therefore explicit. Each added dimension adds one parallel EOM branch, and each of the DeMUX and ReMUX grows by one input/output port per added branch; no redesign of the branch cell is required, and the optical depth experienced by a photon is unchanged because every frequency component still traverses exactly one DeMUX, one branch, and one ReMUX. The control burden grows quadratically with dimension: the drive series of (7) allocates $d-1$ tones per arm, so each branch carries $8(d-1) + 2$ parameters (26 at $d = 4$) and the full gate requires $d[8(d-1) + 2]$ parameters in total. The classical optimization cost of the PSO grows with this parameter count, and drive-tone frequencies extend to $(d-1)\delta$, tying the scaling to the frequency-plan constraint above.

Constant optical depth is not, however, the same as cost-free scaling. Several engineering burdens grow with d even though the number of components in any single optical path does not. The RF distribution network must deliver $d-1$ tones to each of the four arms of every branch, so RF fan-out, routing complexity, and channel count grow with the branch count, and the aggregate RF power and instantaneous bandwidth demanded of the drive electronics grow with the number and spectral spread of those tones. Because the sensitivity results of Section VIII identify relative phase as the dominant error channel, the requirement that these $d-1$ tones remain mutually synchronized and phase stable across all d branches becomes the principal control challenge at higher dimension

rather than a secondary one. On the optical side, the DeMUX and ReMUX must resolve a larger number of bins within the same usable bandwidth, which tightens the insertion-loss and inter-bin crosstalk specifications of the wavelength-division multiplexing components in exactly the manner described above.

For $d = 8$, the RF distribution network delivers seven tones to each of the eight branches. We used particle swarm optimization with a swarm size of 100 and a maximum of 100 iterations with a maximum stall of 30 for each of the eight nontrivial MUB matrices, $\mathbf{U}_1$ through $\mathbf{U}_8$. The mean fidelity achieved was 0.9651 under this reduced optimization budget, providing a favorable indication of the ability of the circuit to scale to higher dimensions. Further scaling investigations are ongoing.

X. DISCUSSION

There are at least three reasons why this measurement interpretation is useful. First, it packages a programmable single-qudit gate as a reusable subsystem with an immediately understandable systems role. A fixed detector plus a programmable front end is a cleaner abstraction than a different custom analyzer for each basis family. Second, it aligns naturally with cluster-state processing in which measurement basis choice is the instruction set of the computation. Third, it suggests a path toward integrated photonic measurement fabrics in which source, gate, and detector are co-designed rather than treated as unrelated laboratory modules.

A practical consideration concerns the reconfiguration speed of the measurement front end. The electro-optic modulators in the nested MZM branches can be driven at RF frequencies, so in principle the measurement basis can be updated on timescales set by the RF electronics rather than by slower thermal or mechanical tuning. For MBQC protocols requiring adaptive basis selection conditioned on prior measurement outcomes, this electronic reconfigurability is a significant advantage over measurement schemes that require physical reconfiguration of optical components.

The MUB simulation results have a direct implication for quantum state tomography. A measurement device that can implement all $d + 1$ MUBs for $d = 4$, including the trivial computational basis, provides a minimal informationally complete measurement. Each basis yields $d - 1 = 3$ independent probability estimates, and five bases yield $5 \times 3 = 15 = d^2 - 1$ independent parameters, exactly the number required to reconstruct an arbitrary four-dimensional density matrix. The same capability underpins higher-dimensional QKD protocols in which security relies on the uniform overlap property of (10). Experimental validation using the TFLN-on-SiN platform is in preparation, targeting the $\mathbf{H} \otimes \mathbf{H}$ basis as the initial hardware demonstration due to its real-valued transfer matrix and compatibility with standard RF drive configurations.

A further observation concerns the structural differences among the MUB target unitaries. Two of the four nontrivial MUBs ($\mathbf{U}_1 = \mathbf{H} \otimes \mathbf{H}$ and $\mathbf{U}_4 = \mathbf{R} \otimes \mathbf{R}$) are tensor products of single-qubit operations, while $\mathbf{U}_2$ and $\mathbf{U}_3$ produce entangled

bases. The simulation results show that the entangled-basis unitaries $\mathbf{U}_2$ and $\mathbf{U}_3$ achieve comparable or superior fidelity to the product-basis unitaries $\mathbf{U}_1$ and $\mathbf{U}_4$, and exhibit greater robustness to parameter perturbations. This suggests that the complex phase structure of entangled bases aligns favorably with the EOM operating points, and that the architecture does not preferentially favor product-form targets.

XI. CONCLUSION

A programmable basis-selective measurement architecture for frequency-bin photonic qudits has been described. The architecture places a programmable single-qudit gate directly in front of a fixed photodetection stage and programs the gate to implement the inverse transform associated with the desired basis. In this manner, the detector performs projective measurement with respect to an arbitrary selected basis supported by the gate. The contribution is hardware-oriented and is based on the reinterpretation of an accepted parallel frequency-comb qudit processor architecture as a measurement front end.

To substantiate the programmability claim beyond a single gate realization, we have simulated the architecture programmed to realize each of the four nontrivial MUBs for $d = 4$, constructed via the two-qubit Pauli stabilizer formalism. Together with the previously reported Chrestenson gate and the trivial computational basis, these five nontrivial targets demonstrate that the measurement front end supports a complete set of informationally significant bases. The MUB measurement capability has direct implications for quantum state tomography and higher-dimensional QKD, where a complete MUB set provides minimal informational completeness and underpins protocol security.

Previously reported device modeling indicates high-fidelity operation and identifies RF phase stability as the dominant control concern. The parallel optical arrangement offers a favorable loss structure relative to serial frequency processors, and the architecture scales to higher qudit dimension by adding parallel branches without increasing optical depth. A preliminary scaling investigation yielded an average fidelity among the eight MUBs of 0.9651 for the $d = 8$ case and these simulations are ongoing. Taken together, these observations support the view that programmable single-qudit gates may serve not only as state-evolution operators, but also as practical qudit readout layers for integrated photonic quantum information systems.

ACKNOWLEDGMENT

The authors gratefully acknowledge the device modeling work of Caden Crowson and Alex Tiu.

REFERENCES

- [1] D. L. MacFarlane, H. Shahoei, R. M. Tuller, C. Shapard, C. Crowson, and M. A. Thornton, "Parallel architecture of a frequency comb qudit quantum processor," in *Proc. 56th IEEE Int. Symp. on Multiple-Valued Logic (ISMVL)*, 2026, to appear.
- [2] H.-H. Lu, A. M. Weiner, P. Lougovski, and J. M. Lukens, "Quantum information processing with frequency-comb qudits," *IEEE Photonics Technology Letters*, vol. 31, no. 23, pp. 1858–1861, 2019.

- [3] M. Kues *et al.*, “On-chip generation of high-dimensional entangled quantum states and their coherent control,” *Nature*, vol. 546, pp. 622–626, 2017.
- [4] Y. Zhang *et al.*, “A silicon photonic processor for frequency-bin qubits and qudits,” *Nature Communications*, vol. 12, p. 2233, 2021.
- [5] R. Raussendorf and H. J. Briegel, “A one-way quantum computer,” *Physical Review Letters*, vol. 86, no. 22, pp. 5188–5191, 2001.
- [6] R. Raussendorf, D. E. Browne, and H. J. Briegel, “Measurement-based quantum computation on cluster states,” *Physical Review A*, vol. 68, p. 022312, 2003.
- [7] N. C. Menicucci, S. T. Flammia, and O. Pfister, “One-way quantum computing in the optical frequency comb,” *Physical Review Letters*, vol. 101, p. 130501, 2008.
- [8] M. Chen, N. C. Menicucci, and O. Pfister, “Experimental realization of multipartite entanglement of 60 modes of a quantum optical frequency comb,” *Physical Review Letters*, vol. 112, p. 120505, 2014.
- [9] J. Roslund, R. M. de Araújo, S. Jiang, C. Fabre, and N. Treps, “Wavelength-multiplexed quantum networks with ultrafast frequency combs,” *Nature Photonics*, vol. 8, pp. 109–112, 2014.
- [10] O. Lib and Y. Bromberg, “Resource-efficient photonic quantum computation with high-dimensional cluster states,” *Nature Photonics*, vol. 18, pp. 1218–1224, 2024.
- [11] F. Bouchard *et al.*, “Measuring ultrafast time-bin qudits,” *Physical Review A*, vol. 107, p. 022618, 2023.
- [12] H.-H. Lu, M. Liscidini, A. L. Gaeta, A. M. Weiner, and J. M. Lukens, “Frequency-bin photonic quantum information,” *Optica*, vol. 10, no. 12, pp. 1655–1671, 2023.
- [13] H.-H. Lu *et al.*, “Electro-optic frequency beam splitters and tritters for high-fidelity photonic quantum information processing,” *Physical Review Letters*, vol. 120, p. 030502, 2018.
- [14] P. Imany *et al.*, “High-dimensional optical quantum logic in large operational spaces,” *npj Quantum Information*, vol. 5, p. 59, 2019.
- [15] N. B. Lingaraju *et al.*, “Adaptive bandwidth management for quantum frequency processors,” *Optica*, vol. 8, no. 3, pp. 329–337, 2021.
- [16] H.-H. Lu, E. M. Simmerman, P. Lougovski, A. M. Weiner, and J. M. Lukens, “Fully arbitrary control of frequency-bin qubits,” *Physical Review Letters*, vol. 125, no. 12, p. 120503, 2020.
- [17] H.-H. Lu, N. B. Lingaraju, D. E. Leaird, A. M. Weiner, and J. M. Lukens, “High-dimensional discrete Fourier transform gates with a quantum frequency processor,” *Optics Express*, vol. 30, no. 6, pp. 10 126–10 134, 2022.
- [18] M. Clementi *et al.*, “Programmable frequency-bin quantum states in a nano-engineered silicon device,” *Nature Communications*, vol. 14, p. 176, 2023.
- [19] Y. Chi *et al.*, “A programmable qudit photonic processor,” *Nature Communications*, vol. 13, p. 1166, 2022.
- [20] C. Kumar, F. Lacerda, and S. Pirandola, “Optimal quantum-programmable projective measurement with linear optics,” *Physical Review A*, vol. 104, p. 052427, 2021.
- [21] L.-T. Feng *et al.*, “Programmable higher-dimensional measurement on an integrated photonic platform,” *Physical Review Letters*, vol. 134, p. 010601, 2025.
- [22] A. Eckstein, B. Brecht, and C. Silberhorn, “A quantum pulse gate based on spectrally engineered sum frequency generation,” *Optics express*, vol. 19, no. 15, pp. 13 770–13 778, 2011.
- [23] B. Brecht, D. V. Reddy, C. Silberhorn, and M. G. Raymer, “Photon temporal modes: a complete framework for quantum information science,” *Physical Review X*, vol. 5, no. 4, p. 041017, 2015.
- [24] L. Serino, J. Gil-Lopez, M. Stefszky, R. Ricken, C. Eigner, B. Brecht, and C. Silberhorn, “Realization of a multi-output quantum pulse gate for decoding high-dimensional temporal modes of single-photon states,” *PRX quantum*, vol. 4, no. 2, p. 020306, 2023.
- [25] J. M. Lukens and P. Lougovski, “Frequency-encoded photonic qubits for scalable quantum information processing,” *Optica*, vol. 4, no. 1, pp. 8–16, 2017.
- [26] I. Bengtsson *et al.*, “Mutually unbiased bases and Hadamard matrices of order six,” *Journal of Mathematical Physics*, vol. 48, no. 5, p. 052108, 2007.
- [27] W. K. Wootters and B. D. Fields, “Optimal state-determination by mutually unbiased measurements,” *Annals of Physics*, vol. 191, no. 2, pp. 363–381, 1989.
- [28] T. Durt, B.-G. Englert, I. Bengtsson, and K. Życzkowski, “On mutually unbiased bases,” *International Journal of Quantum Information*, vol. 8, no. 4, pp. 535–640, 2010.
- [29] N. J. Cerf, M. Bourennane, A. Karlsson, and N. Gisin, “Security of quantum key distribution using d -level systems,” *Physical Review Letters*, vol. 88, p. 127902, 2002.
- [30] D. L. Zhou, B. Zeng, Z. Xu, and C. P. Sun, “Quantum computation based on d -level cluster state,” *Physical Review A*, vol. 68, p. 062303, 2003.
- [31] S. Bandyopadhyay, P. O. Boykin, V. Roychowdhury, and F. Vatan, “A new proof for the existence of mutually unbiased bases,” *Algorithmica*, vol. 34, no. 4, pp. 512–528, 2002.
- [32] J. Lawrence, Č. Brukner, and A. Zeilinger, “Mutually unbiased binary observable sets on n qubits,” *Physical Review A*, vol. 65, p. 032320, 2002.
- [33] M. Kues *et al.*, “Quantum optical microcombs,” *Nature Photonics*, vol. 13, pp. 170–179, 2019.