

Hardware-Efficient Universal Linear Transformations for Optical Modes in the Synthetic Time Dimension

Jasvith Raj Basani^{1,2,*}, Chaohan Cui^{1,3,*†}, Jack Postlewaite¹, Edo Waks^{1,2} and Saikat Guha^{1,3}

¹*Department of Electrical and Computer Engineering, University of Maryland, College Park, Maryland 20742, USA*

²*Institute for Research in Electronics and Applied Physics, and Joint Quantum Institute, University of Maryland, College Park, Maryland 20742, USA*

³*James C. Wyant College of Optical Sciences, University of Arizona, Tucson, Arizona 85721, USA*



(Received 26 August 2025; revised 4 January 2026; accepted 4 March 2026; published 8 April 2026)

Recent progress in photonic information processing has spurred strong demand in scalable and reconfigurable photonic circuitry. Conventional spatially meshed multiport interferometers require a number of components growing quadratically with the system size, posing a fundamental scaling challenge ahead. Here, we introduce a hardware-efficient synthetic time-domain photonic processor that achieves at least an exponential reduction in hardware component count for implementing arbitrary linear transformations. The processor's dynamic connectivity allows systematic pruning, minimizing optical loss while preserving all-to-all connectivity. We benchmark our architecture on the task of boosted Bell-state measurements—a protocol essential for linear optical quantum computation, and show that it exceeds thresholds for universal cluster-state quantum computation under realistic hardware constraints. We link the device performance to the geometry of multiphoton transport, showing that localization effects from redundant, imperfect hardware may enhance robustness to coherent errors. Our design establishes a practical pathway toward near-term, scalable, and reconfigurable photonic processors in the synthetic time dimension.

DOI: [10.1103/ycj-d7pb](https://doi.org/10.1103/ycj-d7pb)

I. INTRODUCTION

Linear programmable photonic circuits with multiple inputs and outputs are fundamental building blocks throughout the development of advanced photonic processors, playing a crucial role in both classical and quantum information processing. In the classical regime, such circuits have enabled acceleration of linear algebraic computations, offering significant advantages in machine learning and AI workloads [1–4]. In the quantum domain, their capacity to implement unitary transformations among multiple modes enables high-dimensional quantum logic [5–10], complex control [11–13], and precise long-range interaction of quantized photonic fields [14–16]. This capability underpins quantum-advantageous protocols in

photonic quantum computing, communication, and sensing, including boson sampling [17,18], quantum sensing with dynamic learning [19,20], and superadditive laser communication systems [21–24].

Traditionally, the most common approach to constructing programmable photonic circuits involves the use of a multiport interferometer. This approach utilizes a mesh of beamsplitters and phase shifters to implement arbitrary unitary transformations [25,26]. Although widely demonstrated, such architectures require $\mathcal{O}(N^2)$ programmable elements to implement arbitrary unitary operations on N modes, resulting in substantial hardware overhead even on integrated platforms [14–16,27], rendering large-scale implementations cumbersome.

In addition to scalability challenges, these individual components are highly susceptible to fabrication imperfections [28–31]. Their individual errors accumulate as the system scales, thus introducing a challenging trade-off between scalability and achieved fidelity [32–34]. The complexity is further exacerbated when extending such architectures to processors utilizing other optical degrees of freedom, such as time bins. In such cases, the implementation requires an intricate mode sorter with $\mathcal{O}(N)$

*These authors contributed equally.

†Contact author: chaohan@umd.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

phase-stabilized optical paths, which presents formidable practical challenges in large-scale systems.

To address these challenges, alternative architectures that exploit the large-scale multiplexing capabilities of photonics have been proposed [13,35–37]. In particular, processing information encoded within photonic time bins can significantly reduce the hardware overhead by time-multiplexing the optical components while maintaining full programmability. One popular approach [38–43] utilizes a pair of nested short and long optical delay lines to construct arbitrary linear transformations via the Reck-Zeilinger decomposition [25]. The time complexity required to compile arbitrary unitary matrices using the nested loop architecture scales as $\mathcal{O}(N^2)$. More recently, an alternative approach [44] used optically induced nonlinearities and birefringent materials to perform arbitrary unitary transformations using a synchronized pulsed pump and $\mathcal{O}(N)$ optical components in the cascaded layout.

In this manuscript, we introduce the *generalized Green Machine* (gGM), a flexible time-domain photonic processor architecture for implementing programmable linear transformations on optical modes encoded in the synthetic time dimension. This recursive architecture trades complexity in the spatial domain for complexity in the temporal domain while offering enhanced flexibility. It relies on only a single Mach-Zehnder interferometer (MZI) pairing with switchable delay lines to perform interference among the time-binned modes. Unlike spatial-mode multiport interferometers, where $\mathcal{O}(N^2)$ active MZIs create a scaling bottleneck in practice, the generalized Green Machine requires only $\mathcal{O}(\log_2 N)$ to $\mathcal{O}(1)$ delay lines, depending on the adopted connectivity. Therefore, the generalized Green Machine enjoys at least an exponential reduction in the hardware component count compared to traditional spatial mode interferometer meshes.

We provide detailed prescriptions for programming the generalized Green Machine to achieve desired unitary matrices and numerically show its robustness to coherent beamsplitter errors in the MZI with one-photon and two-photon scattering. After the end-to-end characterization, we benchmark its performance by the boosted Bell-state measurement (BSM) [45,46], a fundamental protocol in linear optical quantum computing and networking. We show that under practical hardware conditions, our proposed architecture is capable of surpassing a specific percolation threshold needed for the fusion-based generation of large-scale quantum cluster states using current photonic technology [47]. Owing to its flexible connectivity, the number of recursive rounds required to implement boosted BSM is significantly lower than that of the nested loop architecture restricted to the Clements decomposition [38]. We have further explored that these results are linked to the geometrical nature of our design in multiphoton transport and interference.

II. PARALLEL INTERFERENCE AND RECURSIVE ARCHITECTURE

The layout of the generalized Green Machine is depicted in Fig. 1(a). The input photonic state is encoded over complex-valued amplitudes across serialized time bins, with a time interval of τ . The MZI is parameterized by two reconfigurable phases (θ, ϕ) , allowing individual interference of each pair of time bins. An output port of the second switch is connected to one of the input ports, allowing photonic fields to be recursively propagated within this system or to drop out of the system. This structure enables dynamic programming of connectivity among the time-bin modes and overall circuit depth, allowing it to be configured into well-established interferometer architectures [25,26,48].

To demonstrate the underlying working principle of the generalized Green Machine, we begin with an illustration of the interference performed in a single round trip. To achieve this function, the first 2×2 optical switch divides the input time bins into two sequences, each directed to a distinct path mode, with one sequence leading the other. The leading sequence is then delayed by an appropriate delay line to align with the lagging sequence. When the corresponding pair of time-bin modes (i, j) arrives, the programmable MZI parameters, θ_i and ϕ_i , are set to the values determined by the decomposition of the target unitary operation. After interference, two output time-bin sequences exit the MZI simultaneously. The delay line in the second arm then delays the output sequence at the bottom port, postponing it to avoid overlapping with the other sequence in time. At the end of this round, the second switch concatenates the two sequences into one. Figure 1(b) illustrates the schematics of stepwise operations applied to 8 time-bin modes to couple fourth-nearest and nearest-neighbor modes.

Therefore, the unitary transformation for the n th round can be represented as

$$U^{(n)}(\vec{\theta}, \vec{\phi}) = \prod T_{ij}(\theta_i, \phi_i), \quad (1)$$

where T_{ij} is the 2×2 unitary matrix among selected pairs of time bins i and j , parametrized as

$$T_{ij}(\theta, \phi) = ie^{i\theta/2} \begin{bmatrix} e^{i\phi} \sin(\theta/2) & \cos(\theta/2) \\ e^{i\phi} \cos(\theta/2) & -\sin(\theta/2) \end{bmatrix}. \quad (2)$$

The output state, after undergoing the unitary transformation $U^{(n)}$ in the n th round, is fed back into the apparatus for the next round. This process is repeated until the desired global unitary transformation is performed, and the second switch can be programmed to send the combined output sequence to the drop-off port.

By compiling interference patterns in time (via multiple programmed recursive rounds), the generalized




FIG. 1. Schematic and operation of the generalized Green Machine. (a) Illustration of the hardware components to construct the generalized Green Machine consisting of switches, delay lines, and programmable beamsplitters. Information is encoded serially over time bins of width τ . (b) Stepwise operations indicating the processes applied to the eight time bins over two spatial modes to interfere with the fourth nearest-neighbor modes. The second delay line delays the bottom arm, equivalent to shift the time frame of the top arm to the front. (c) Fully programmable unitary transformations implemented on time-bin modes over multiple round trips in the sine-cosine fractal configuration (top) and the Clements configuration (bottom).

Green Machine can be configured to perform arbitrary unitary transformations. The sine-cosine fractal (SCF) architecture [48], shown in Fig. 1(c), is composed of stages with 2^k nearest-neighbor connectivity, where $k \in [0, 1, 2, \dots, \log_2 N - 1]$. To realize the SCF mesh architecture, $\log_2 N$ delay lines of length $2^k \tau$ are sufficient. Alternatively, the device can also be programmed into the Clements architecture [26] by programming alternate stages to couple the nearest-neighbor modes. In this case, a *single* delay line of length τ is sufficient. The stepwise operations, indicating the processes applied to the time-bin modes to configure them into Clements and SCF configurations, are detailed in Appendix A.

III. PERFORMANCE UNDER IMPERFECTIONS

Errors in the generalized Green Machine stem from two main sources. The first arises from component imperfections, which cause the beamsplitters in the MZI to deviate from their ideal 50:50 splitting ratio [28–31, 49]. The second is the imbalanced insertion loss introduced by the switchable delay modules and MZI. We discuss the impact of coherent errors in this section, and the impact of loss in Appendix C.

Imperfect splitting in the beamsplitter ratios, caused by variations in the fabrication process, introduces errors in

the programmed unitary matrix. Unlike spatial-mode interferometers, the generalized Green Machine utilizes only a single MZI, resulting in correlated errors among all the time-bin modes. Under the influence of these correlated errors, the transfer matrix implemented by the generalized Green Machine U_{GM} deviates from the ideal targeted unitary matrix U_{ideal} . To quantify the impact of these coherent errors, we evaluate the average state infidelity of a group of few-photon states $|\psi\rangle$ that evolve through U_{GM} as

$$\bar{\mathcal{E}} = 1 - \bar{\mathcal{F}} = 1 - \langle |\langle \psi | U_{\text{ideal}}^\dagger U_{\text{GM}} | \psi \rangle|^2 \rangle. \quad (3)$$

The impact of component imperfections on the Green Machine's ability to implement large-scale unitaries is benchmarked by simulating the scattering of a single photon across $N \in [16, 64, 256]$ time bins. For a given N , we sample 10^4 unitaries from the Haar measure, each with a random input state. The beamsplitter errors for each circuit, (α, β) , are drawn independently from the Gaussian distribution $\mathcal{N}(0, \sigma)$. Figure 2(a) illustrates a schematic of an MZI with the two beamsplitters perturbed by component errors (α, β) . In Fig. 2(b), we plot the state infidelity $\mathcal{E}$ as a function of beamsplitter error σ . The median and interquartile ranges (colored bands) are plotted alongside analytical predictions for spatial-mode interferometers




FIG. 2. Performances of the generalized Green Machine under beamsplitter errors. (a) Model of MZI with coherent errors (α, β) causing deviations from the ideal 50:50 splitting. (b) State infidelity as a function of beamsplitter errors σ for a single-photon state scattered among $N = 256$ time bins. The simulated performance outperforms the dashed lines that mark the scaling law of traditional spatial-mode interferometers. (c) Histograms of the infidelities of the two-photon boson sampling among up to $N = 32$ time bins. Infidelities due to the correlated model of errors from the generalized Green Machine are shown in the darker color (right half), fitted with the dotted line. The brighter (left half) histograms illustrate the distribution of infidelities under uncorrelated errors (as seen in traditional multiport interferometer architectures) in sine-cosine fractal mesh, and are fitted with the dashed line.

with uncorrelated errors, which scale as $\mathcal{E} \approx \frac{1}{2}N\sigma^2$ (black lines). See Appendix B for the analytical derivation).

Then, to explore its potential in multiphoton transport and boson sampling, we benchmark the generalized Green Machine on sampling an $N_{\text{ph}} = 2$ photon state from random unitaries of up to size $N = 64$. Figure 2(c) compares distributions of sampling infidelities for the Green Machine (right half, darker) and a conventional spatially meshed multiport interferometer (left half, brighter) with uncorrelated errors under two noise levels: $\sigma = 0.001$ plotted in green and $\sigma = 0.01$ plotted in blue. The dotted

and dashed lines are fitted to the median values of these distributions, and follow the scaling law:

$$\mathcal{E} \propto N_{\text{ph}}N\sigma^2. \quad (4)$$

As shown in Figs. 2(b) and 2(c), for both cases, the generalized Green Machine reduces the mean infidelity by a factor of $\sim\sqrt{2}$ across all system sizes by virtue of its correlated error model (discussed in Appendix B). Despite a marginally larger variance, the resulting broader distribution increases the likelihood of postselecting high-performance hardware. Notably, the analytical scaling of spatial-mode interferometer meshes under uncorrelated coherent errors remains identical, regardless of whether the decomposition follows the Clements, Reck-Zeilinger, or SCF architecture [29,30]. In practice, physical implementations of the generalized Green Machine may experience a combination of both correlated and uncorrelated errors. Uncorrelated noise may arise from fluctuations in control signals or the thermal drift of the MZIs. If this noise is uncorrectable, the fidelity of the generalized Green Machine will typically fall between the optimal correlated-noise limit (with a $\sqrt{2}$ advantage) and the standard baseline performance of a spatial mesh with completely uncorrelated noise, depending on the relative dominance of these error sources.

While Fig. 2(c) compares the infidelity of spatial-mode interferometer meshes (uncorrelated errors) with that of the generalized Green Machine (correlated errors), an additional advantage of the Green Machine emerges when hardware error-correction techniques are applied [28,48]. These techniques further improve the scaling of corrected matrix error from $\mathcal{O}(N\sigma^2)$ to $\mathcal{O}(\sqrt{N \log_2 N}\sigma^2)$ providing a substantial reduction in coherent-error accumulation. Figure 11 provides a direct comparison of matrix errors across all four regimes: spatial-mode interferometers with uncorrelated errors, the generalized Green Machine with correlated errors, and both schematics after hardware error correction.

IV. TOWARD ROBUST BOOSTED BELL-STATE MEASUREMENT

A near-term application that immediately benefits from our architecture is the boosted BSM [45]. The BSM aims to project a dual-rail-encoded photon pair onto four Bell states: $|\Psi_{\pm}\rangle = (|01\rangle \pm |10\rangle)/\sqrt{2}$ and $|\Phi_{\pm}\rangle = (|00\rangle \pm |11\rangle)/\sqrt{2}$, which is a fundamental operation for the fusion-based generation of a large-scale cluster state [47], quantum teleportation, and entanglement swapping [50]. The standard BSM circuit with a simple beamsplitter operation can only deterministically distinguish between $|\Psi_{\pm}\rangle$ from measurement results, which places an upper bound on the success rate at 50% [51]. This success rate can be boosted by incorporating ancillary single-photon




FIG. 3. Green Machine implementation of boosted Bell-state measurement. (a) Circuit of the boosted Bell-state measurement implemented by a minimum of three stages of the Green Machine on two dual rail qubits and four ancillary modes to achieve a successful prediction probability of 75%. The blue shaded regions indicate sections of the Green Machine circuit set to the identity operation. (b) Probability of successful prediction as a function of MZI coherent-error characteristics σ comparing the performance of the 3-stage Green Machine with the 7-stage Clements mesh. (c) Error given successful prediction as a function of MZI coherent-error characteristics σ .

modes [45] and a larger multiport interferometer. An example of such a circuit and its equivalent implementation using three stages of the generalized Green Machine is shown in Fig. 3(a). The positions of the dual rail qubits are denoted by $|\cdot\rangle_{1/2}$ while the ancillary modes are denoted by rails populated with a single photon, indicated by the green pulse. This circuit increases the BSM success rate to 75% in ideal cases—surpassing the percolation threshold of 67.2% necessary for universal photonic quantum computing with cluster states [47], and increases the efficiency of entanglement generation for quantum networks. An added advantage of using the time-bin architecture is that the ancillary single photons can be sequentially generated from a selected single-photon emitter, ensuring consistent

high indistinguishability necessary for high-fidelity quantum interference.

The performance of boosted BSM is characterized by two figures of merit: (1) the *success heralding rate*, defined as the probability of heralding the projection into one Bell state according to its detection signature, and (2) the *error rate given heralded event*, which quantifies the misidentification rate due to measurement crosstalk after heralding. We employ a Bayesian inference model that assigns posterior probabilities to each detection signature, distinguishing between two outcomes:

- (1) *Decode*, where the event projects the received pattern onto one of the four Bell states if the posterior probability is greater than the decision threshold. This contributes to the success heralding rate. A nonzero probability of detecting the same click patterns by the other Bell states contributes to the increased error rate given this successful prediction.
- (2) *Discard*, where no assignment is made if the posteriors are inconclusive. This results from an inability to decide among all four Bell states, and reduces the heralding success rate.

The success rate predicted by the Bayesian model is subject to variation due to several factors. These include sources of crosstalk in the time-domain boosted BSM circuit, such as coherent beamsplitter errors or tunable parameters like circuit depth and decision threshold. Our results are plotted over 1000 samples taken from circuits with emulated noises, as quantile box plots with the mean value indicated by the red diamond.

First, we evaluate the impact of coherent errors by sampling the errors (α, β) from a normal distribution $\mathcal{N}(0, \sigma)$. Results plotted in Figs. 3(b) and 3(c) compare the successful prediction rate (and a corresponding increase in the error rate) for the boosted BSM circuit implemented using the Green Machine and the Clements mesh. For the Green Machine (indicated in red), at a splitter error of $\sigma = 3\%$, the worst-case success rate drops below the percolation threshold. The average success rate drops to this threshold at $\sigma \sim 10\%$, where more than a quarter of the sampled circuits have their heralding success rates below the threshold. In contrast, an eight-mode spatial Clements interferometer mesh requires the full circuit depth of all seven stages to implement the equivalent matrix transformation. More stages accumulate more uncorrelated MZI error, further degrading the fidelity of the boosted BSM circuit. As shown in Figs. 3(b) and 3(c), the average success rate of the spatial-domain Clements implementation falls to the percolation threshold at an error rate of $\sigma \sim 4\%$, whereas the average success rate of the Green Machine hits the threshold at an error rate of $\sigma \sim 10\%$. For the boosted BSM error rates given heralded events, the Clements implementation is also more sensitive to the MZI error than the Green Machine.




FIG. 4. Successful heralding rate and error rate given heralded event for boosted Bell-state measurement. Probability of successful heralding and error given heralded event as a function of: (a) and (b) circuit depth, varying from maximally pruned to fully expressive; (c) and (d) decision threshold with 2% MZI error.

Subsequently, we vary the circuit depth by adding redundant stages in the configuration of the SCF mesh. Contrary to expectations that deeper, noisier circuits degrade performance, we observe in Figs. 4(a) and 4(b) a nonmonotonic improvement in both success and error rates. This counterintuitive behavior is linked to multiphoton transport phenomena, which we analyze in the following section. In practice, the three-stage Green Machine may still be the best option, considering the additional loss brought by redundant stages.

Finally, we sweep the decision threshold, which is used to determine the decoding strategy employed in our inference model. By reducing this threshold, the successful prediction rate can increase at the expense of a higher error rate, as shown in Figs. 4(c) and 4(d). With $\sigma = 2\%$ MZI error, the transition edge between inference strategies becomes less sharp, which can serve as a tuning knob for globally optimizing the overall performance of fusion-based photonic quantum computing in practice.

V. SELF-SIMILAR QUANTUM TRANSPORT

The performance of the generalized Green Machine is strongly linked to the geometrical nature of quantum transport through modes [52]. An example of this behavior is distinctly visible in our simulations of the boosted BSM in Fig. 4(a), where the success rate drops at a circuit depth of 5, and increases significantly at a depth of 7. To explain this phenomenon, we study photon transport through the generalized Green Machine with it configured to the SCF mesh.

We model this transport by tracking the photon statistics of single and two-photon states hopping from one input time-bin mode to other modes after each round averaged over 100 noisy circuits sampled with $\sigma = 2\%$. An initial state of one or two photons (either a two-photon Fock state or two unentangled indistinguishable photons) is used as the input in our simulations. This optical state interferes across the time-binned modes as it evolves through the stages, with an MZI that is always set to 50:50 splitting ratio. These amplitudes are plotted in Fig. 5 for 8 (similar to the boosted BSM circuit) and 32 time-bin modes. The versatility of the generalized Green Machine enables us to contrast the transport dynamics under different connectivity configurations. This fundamental difference in transport dynamics observed between the SCF and Clements configurations is an expected consequence of their internal connectivity and unit-cell structures. The SCF configuration possesses a fractal, self-similar structure that leads to complex and nonmonotonic transport patterns, arising from re-interference. On the other hand, the Clements architecture interferes only nearest-neighbor modes which results in systematic diffused mode-mixing.

These simulation results reveal distinct, visually structured regimes of transport, ranging from fully diffused to strongly localized in the SCF configuration. For the case of single-photon transport in an 8-mode system, illustrated in Fig. 5(a), we observe a fully diffused regime at stage 5, where the probability of detecting the photon is uniformly distributed across all modes. In contrast, at stage 7, the optical state becomes highly localized, with significant amplitude confined to just two modes. This alternating behavior between delocalized and localized regimes persists in larger systems, such as the 32-mode configuration shown in Fig. 5(b), where the transport patterns reflect the self-similar structure of the SCF mesh.

The introduction of disorder—particularly static perturbations arising from beamsplitter imperfections—modifies these transport characteristics. This disorder has been shown to cause coherent particles to become exponentially localized—a phenomenon commonly referred to as Anderson localization [52,53]. In our system, this would correspond to the exponential localization of the optical mode to its respective time bin. We attribute the nonmonotonic behavior of the boosted BSM success rate in Figs. 4(a) and 4(b) to such localization effects. Specifically, the robustness




FIG. 5. Photon transport through the generalized Green Machine. (a) Single photon transport through the 8-mode circuit with all MZIs setting to 50:50, showing a localized state at $L = 7$, and a diffused state at $L = 5$. (b) Single photon transport with varying degrees of localization in a 32-mode circuit. The heatmap indicates the probability of measuring a photon at each mode after each stage. (c) and (d) Two-photon transport depicting varying degrees of localization in the presence of Hong-Ou-Mandel interference, indicating that the localization and diffusion still happen with quantum interferences. The heatmap indicates the probability of two-photon meeting at the same mode after each stage. (e) Single photon transport in a 32-mode Clements configured circuit, showing no distinct regimes of transport. (f) and (g) Two-photon transport statistics (probability of two-photon meeting at the same mode) through the Clements configuration.

observed at stage 7 may stem from localization-induced protection, in contrast to the fragility of the diffused state at stage 5, as shown in Fig. 5(a). The impact of these localization effects and their ability to protect quantum information from coherent errors remains the subject of our future work.

The effects resulting from this self-similarity also produce distinct regimes in the case of multiphoton transport. As an example, we simulate two-photon transport in a 32-mode system, shown in Figs. 5(c) and 4(d). The resulting interference pattern exhibits a similar fractal pattern, with populations alternating between the single and two-photon excitation manifolds, when the input state is a two-photon Fock state [see Fig. 5(c)]. In another example, when modes 15 and 16 are excited using identical single photons, only at 21st stage do we see two photons meeting at the same mode, where the probability is distributed uniformly over all modes, as shown in Fig. 5(d). These are in sharp contrast to both single- and two-photon transport across the Clements configuration with identical input states, as

shown in Figs. 5(e), 5(f), and 5(g), where the photon evolution trajectories diffuses from the input modes.

VI. DISCUSSION AND OUTLOOK

We have presented an architecture for the *generalized Green Machine*, a hardware-efficient, universal time-domain linear optical processor that utilizes dual switches. This architecture is naturally compatible with both Clements [26], SCF mesh [48], or a hybrid of both, to express arbitrary linear unitary matrices in distinct symmetries with a minimum number of round trips. We have numerically simulated its performance under practical imperfections in the task of boosted BSM, single-photon, and two-photon transport, in which our new architecture gains a constant scaling advantage in average fidelity compared to spatial-mode multipoint interferometers. Other than discrete-variable quantum photonics, the proposed architecture also works for continuous-variable quantum photonics as it operates without unintended vacuum input

TABLE I. Comparison of performance metrics for approaches to unitary transformations on optical modes. Throughput density is defined in units of multiply–accumulate (MAC) operations per hardware per pass. ‡ We emphasize that the throughput of Motes’ architecture is half of this work. Here τ denotes the temporal spacing of neighbor time bins (including bin size and guard band if exists), η_{bs} denotes the loss in dB per Mach-Zehnder interferometer, while η_i and η_o denote the loss rate in dB/m for the inner and outer delay lines respectively, and c stands for the speed of light in the delay lines. We note that the architecture of Bouchard *et al.* [44] can be naturally generalized into the broader operational framework of the Green Machine architecture.

Architecture	Hardware Complexity	Throughput density	Compilation time	Loss scaling
This work (Clements)	$\mathcal{O}(1)$	$\mathcal{O}(N)$	$\mathcal{O}(N^2\tau)$	$\mathcal{O}(N(\eta_{bs} + \eta_i + c\tau\eta_oN))$
This work (SCF)	$\mathcal{O}(\log_2 N)$	$\mathcal{O}(N/\log_2 N)$	$\mathcal{O}(N^2\tau)$	$\mathcal{O}(N(\eta_{bs} + c\tau\eta_i\log_2 N + c\tau\eta_oN))$
Clements (spatial) [26]	$\mathcal{O}(N^2)$	$\mathcal{O}(1)$	$\mathcal{O}(N)$	$\mathcal{O}(N\eta_{bs})$
Motes <i>et al.</i> [38]	$\mathcal{O}(1)$	$\mathcal{O}(N/2)^\ddagger$	$\mathcal{O}(N^2\tau)$	$\mathcal{O}(N(\eta_{bs} + \eta_i + c\tau\eta_oN))$
Bouchard <i>et al.</i> [44]	$\mathcal{O}(N)$	$\mathcal{O}(N)$	$\mathcal{O}(N\tau)$	$\mathcal{O}(N(\eta_{bs} + \eta_i))$

modes, which leads to new opportunities for all-photonic generation of cat and GKP states [54], quantum computing [55,56], and new practical quantum photonic applications [57] focusing on time-domain information processing.

The two leading figures of merit that can be used to benchmark across linear optical processors are (1) the hardware complexity, defined as the number of hardware components required to implement an arbitrary $N \times N$ unitary matrix, and (2) the throughput density, defined as the number of multiply accumulate (MAC) operations per round-trip per amount of hardware required. In terms of hardware complexity, we require at best $\mathcal{O}(1)$ components in the Clements configuration or $\mathcal{O}(\log_2 N)$ components in the SCF configuration, providing at least an exponential reduction in the amount of hardware. Since our device operates on all N modes in a single round trip, we achieve a linear scaling in throughput density, which is typically realized only with hyper-multiplexing. The comparison with other competing architectures is exhibited in Table I. An additional advantage of the generalized Green Machine is that, since it can be programmed into well-understood interferometric configurations, it is amenable to self-configuration techniques [29,30,48], making it more robust to hardware imperfections.

The generalized Green Machine also features flexibility as it allows for reducing the depth of the linear optical circuit by pruning [48,58]—simply by reducing the number of recursive rounds. On the contrary, spatial-mode linear processors cannot be dynamically pruned—this would involve either bypassing the redundant hardware or physically detaching it. By programming the pruning scheme, the all-to-all connectivity can be maintained while minimizing loss and latency. This is particularly beneficial for implementing classes of transformations with specific symmetries. For instance, achieving an eight-mode boosted BSM circuit requires only three recursive rounds, while achieving an $N \times N$ Hadamard transform requires only $\log_2 N$ recursive rounds [21]. The proof-of-concept demonstration of a 16-mode Hadamard transform in a fiber-optical setup has been shown in Ref. [22].

This capability to dynamically program the connectivity, depth, and drop-out ports of the circuit enables the generalized Green Machine to be programmed into more general (i.e., nonunitary) beamsplitter meshes [59] such as the diamond [60,61] or path-independent loss (PILOSS) architectures [62,63]. Large-scale and high-dimensional transformations can be performed by spatially multiplexing our device. We discuss this hypermultiplexed architecture in Appendix D and evaluate its computational speed for accelerating matrix products in classical machine learning tasks.

The loss of the programmable linear optical circuit constitutes another critical figure of merit. Quantum applications relying on multiphoton transport usually experience a severe decline in performance due to such loss. Here, we analyze the loss tolerance for boosted BSM against a baseline success heralding rate exceeding the percolation threshold of 67.2%. In the worst-case scenario—postselecting for the detection of all six photons, the success heralding rate scales as $P = 0.75\eta^6$ with η the total transmission of the circuit. This imposes a minimum transmission requirement of $\eta \geq 0.98$ (loss ≤ 0.08 dB), which is achievable with near-term integrated photonic technologies [64–66], nonlinear optics [44], or using free-space optics [17,67]. Notably, estimates from Refs. [68,69] suggest that this bound may be further relaxed where losing photons does not destroy all the information. In contrast, a spatial-mode Clements interferometer requires a depth of seven layers to realize the same transformation, which exacerbates the drop in success rate and the rise in error rates given heralded event, rendering practical implementation significantly more challenging.

Similarly, optical loss compromises the quantum advantage of boson sampling by rendering the sampling problem classically tractable [70]. Our proposal’s dynamic connectivity, however, allows us to realize Haar-random matrices with the minimum necessary optical path length through well-established optimization routines. This ability to dynamically prune circuits directly reduces total optical depth and accumulated loss while perfectly preserving the

all-to-all connectivity required for universality—a capability typically unavailable to spatial-mode interferometer meshes.

In practice, the optical loss of the generalized Green Machine stems primarily from fast optical switches and optical delay lines. As detailed in Table I, the recursive architecture introduces additional loss through $\mathcal{O}(N)$ passes of the outer-loop delay line, scaling as $c\tau\eta_o\mathcal{O}(N^2)$. However, for a system with $N = 100$ optical modes and a time-bin width of $\tau = 100$ ps, this accumulated loss is approximately 0.04 dB using single-mode fiber at 1550 nm (≈ 200 m with a loss rate of 0.2 dB/km). This is negligible compared to the measured 0.1 dB per MZI facet made with thin-film Barium Titanate (BTO) [66]. Therefore, using present-day hardware, our architecture offers loss figures comparable to—or, with layer pruning, even lower than those of state-of-the-art fast-programmable spatial interferometer meshes.

The generalized Green Machine achieves a significantly reduced hardware footprint at the cost of temporal overhead, with its compilation time scaling as $\mathcal{O}(N^2\tau)$. This contrasts with spatial-mode interferometer meshes, where compilation is time-of-flight limited and scales as $\mathcal{O}(N)$. However, the $\mathcal{O}(N^2\tau)$ scaling for the generalized Green Machine represents a worst-case upper bound. While in practice, circuits can be dynamically pruned to minimize the round trips required for a specific target unitary.

To put this overhead into perspective, we consider a system of size $N = 100$ utilizing state-of-the-art optical switches ($\tau \approx 4.3$ ps) [44]. The latency to compile an arbitrary unitary is approximately 43 ns, which is comparable to the detection and data acquisition times of a spatial-domain processor. Nevertheless, for extremely large-scale systems where quadratic scaling overtakes the switching speed (e.g., $N^2 > 1 \text{ sec}/\tau$ at $N \sim 10^5$), this compilation time could become the dominant constraint on throughput.

ACKNOWLEDGMENTS

All authors acknowledge the DARPA PHENOM program for the support. C.C. and S.G. also acknowledge the DARPA QuANET program for partial support and thank the support from the Engineering Research Center for Quantum Networks (CQN) under NSF Grant No. EEC-1941583 for synergistic research support.

DATA AVAILABILITY

The data that support the findings of this article are openly available [71].

APPENDIX A: COMPILING TEMPORAL MODE TRANSFORMATIONS

The gGM architecture we introduce in the main text compiles transformations stage-wise on time-bin modes




FIG. 6. Stepwise operations to couple n th nearest time-bin modes. Illustration of the switching scheme and the impact of interference on eight time-bin modes. These steps can be used to compile global unitary operations via the SCF configuration using (a), (b), and (c) or the Clements configuration via (c) and (d).

to perform a desired unitary operation. Depending on the preprogrammed configuration being used (Clements, SCF, or otherwise), each stage couples the n th nearest-neighbor modes. In Fig. 6, we illustrate the stepwise operations, indicating the processes being applied to eight time-bin modes. Operations demonstrated in Figs. 6(a), 6(b), and 6(c) are sufficient to realize arbitrary 8-mode unitaries in the SCF configuration by coupling fourth-nearest, second-nearest, and nearest-neighbor modes. Similarly, operations in Figs. 6(c) and 6(d) are sufficient to realize the Clements configuration by coupling odd and even nearest-neighbor modes.

These stage-wise interference patterns can be compiled to realize a fully expressive unitary, or pruned meshes that cannot express the entire $SU(N)$ unitary group. In our analysis of the boosted Bell-state measurement, we interpolate between the circuit shown in Fig. 3(a) of the main text and the fully expressive SCF mesh following the scheme proposed in Ref. [48]. Figure 7 illustrates the order in which interference patterns are compiled.

APPENDIX B: DERIVATION OF ERROR SCALING

We are interested in the infidelity of the output state introduced by coherent errors arising from beamsplitter




FIG. 7. Connectivity as a function of number of stages. (a)–(e) Schematics of the connectivity among eight time-binned modes implementing different circuit depths from three to seven stages as performed in our analysis of the boosted Bell-state measurement.

imperfections in unitaries implemented by the generalized Green Machine. Consider a single-photon state $|\psi\rangle = \sum c_i \hat{a}_i^\dagger |0\rangle$ that evolves under a unitary implemented by the Green Machine. The notation $\hat{a}_i$ ($\hat{a}_i^\dagger$) is the annihilation (creation) operator for a bosonic excitation in the i th time bin, and c_i is its corresponding probability amplitude such that $\sum |c_i|^2 = 1$.

We quantify the infidelity of the output state introduced by small coherent errors arising from beamsplitter imperfections. We consider a quantum state $|\psi\rangle$ evolving under a $N \times N$ unitary transformation U . Due to fabrication imperfections, the implemented unitary deviates from the target ideal unitary U_{ideal} as $U = U_{\text{ideal}} + \Delta U$. By averaging the state infidelity over a group of random input states, the infidelity $\bar{\mathcal{E}} = 1 - \bar{\mathcal{F}}$ is dominated to the second order by the variance of the error operator,

$$\bar{\mathcal{E}} \approx \frac{1}{4N} \|\Delta U\|^2 = \frac{1}{4N} \text{Tr}(\Delta U^\dagger \Delta U), \quad (\text{B1})$$

where $\|\cdot\|$ denotes the Frobenius norm. The coefficient $1/4N$ normalizes the infidelity to $[0, 1]$, matching the definition of the average state infidelity. This total error ΔU is accumulated by all the imperfections of every MZI. To analyze the scaling of ΔU , we first determine the form of the perturbation ΔU_ℓ at a single MZI.

A general physical MZI is parametrized by two angles/phases ϕ and θ . θ determines the splitting ratio which is susceptible to beamsplitter coherent errors α and β (coherent errors that affect ϕ can be compensated by calibration). We model this noisy l th MZI (2-by-2) unitary $U_l(\theta_l, \alpha_l, \beta_l) = e^{-i(\beta_l + \pi/4)\sigma_x} e^{-i\frac{\theta_l}{2}\sigma_z} e^{-i(\alpha_l + \pi/4)\sigma_x}$ with the target splitting angle θ_l , errors α_l, β_l sampled from $\mathcal{N}(0, \sigma)$, and $\sigma_{x,y,z}$ representing 2-by-2 Pauli matrices. Expanding around the target ideal unitary U_0 , the first-order error perturbation (for small α_l and β_l) is

$$\Delta U_\ell \approx - \left[(\alpha_\ell + \beta_\ell) \cos \frac{\theta_\ell}{2} + i(\alpha_\ell - \beta_\ell) \sin \frac{\theta_\ell}{2} \sigma_y \right] + \mathcal{O}(\alpha_\ell^2, \beta_\ell^2, \alpha_\ell \beta_\ell). \quad (\text{B2})$$

Now, combining the independent errors into symmetric ($S_\ell = \alpha_\ell + \beta_\ell$) and antisymmetric ($A_\ell = \alpha_\ell - \beta_\ell$) error components, it yields

$$\Delta U_\ell \approx - \left[S_\ell \cos \frac{\theta_\ell}{2} + iA_\ell \sin \frac{\theta_\ell}{2} \sigma_y \right]. \quad (\text{B3})$$

Now, we average over all possible $\theta_\ell \in [0, \pi]$ for sampled Haar-random unitaries under the same set of the error S_ℓ and A_ℓ . For Clements architecture implementing Haar-random unitaries [72], the distribution of splitting angles clusters tightly near $\theta \approx 0$ (the bar state). In this limit, contribution from the antisymmetric $\sin(\theta/2)A_\ell$ is negligible. The error is thus dominated by the symmetric component $\|\Delta U_\ell\|^2 \approx 2S_\ell^2$. Since α_ℓ and β_ℓ are statistically independent variables with zero mean and RMS amplitude σ ($\text{Var}(\alpha_\ell) = \text{Var}(\beta_\ell) = \sigma^2$), their variances add: $\text{Var}(S_\ell) = \text{Var}(A_\ell) = \text{Var}(\alpha_\ell) + \text{Var}(\beta_\ell) = 2\sigma^2$. Therefore, averaging over all sampled errors, we have $\mathbb{E}[\|\Delta U_\ell\|^2] = 2 \text{Var}(S_\ell) = 4\sigma^2$.

When $N(N-1)/2$ individual MZIs compose a standard Clements mesh, all the physical errors α_ℓ, β_ℓ can be assumed to be uncorrelated and independently sampled from $\mathcal{N}(0, \sigma)$. The expected error in the first order, $\Delta U_{\text{Clements}}$ scales the Euclidean distance of all independent MZI errors $\{\Delta U_\ell\}$. Thus, the expected squared norm of the error is

$$\begin{aligned} \bar{\mathcal{E}}_{\text{Clements}} &\approx \frac{1}{4N} \mathbb{E}[\|\Delta U_{\text{Clements}}\|^2] \\ &\approx \frac{1}{4N} \sum_{\ell=1}^{N(N-1)/2} \mathbb{E}[\|\Delta U_\ell\|^2] \\ &= \frac{1}{2} (N-1) \sigma^2 \approx \frac{N}{2} \sigma^2. \end{aligned} \quad (\text{B4})$$

In the recursive time-bin gGM-SCF architecture, the error is defined by a single static vector $\vec{\epsilon} = (S, A)$ fixed for the entire mesh. The total error operator is the coherent




FIG. 8. Scaling of state infidelity as a function of MZI error. Numerical simulations and analytically derived scaling for a single photon state scattered among upto $N = 256$ time bins, with uncorrelated errors (Clements mesh) and correlated errors (gGM). The numerical simulations agree with the analytical solution for large unitaries.

sum of propagated local errors,

$$\Delta U_{\text{gGM}} \approx -S \sum_{\ell=1}^{N(N-1)/2} \underbrace{U_{\text{post}}^{(\ell)} \left(\cos \frac{\theta_{\ell}}{2} \mathbb{I}^{(\ell)} \right) U_{\text{pre}}^{(\ell)}}_{K_{\ell}^{(S)}} - iA \sum_{\ell=1}^{N(N-1)/2} \underbrace{U_{\text{post}}^{(\ell)} \left(\sin \frac{\theta_{\ell}}{2} \sigma_y^{(\ell)} \right) U_{\text{pre}}^{(\ell)}}_{K_{\ell}^{(A)}}. \quad (\text{B5})$$

Unlike the uncorrelated case, the deviation caused by the two error constants S and A accumulates across $N - 1$ layers of MZIs. Within each layer, the MZI error applies in parallel. Therefore, the cross-correlation across layers contributes to the overall error.

From the simulation results shown in Fig. 8, we empirically find that for gGM-SCF, the expected total error scales with the Euclidean norm of the component error terms (quadrature sum) rather than their algebraic sum, which is distinct from the uncorrelated error model used for Clements decomposition,

$$\begin{aligned} \bar{\mathcal{E}}_{\text{gGM}} &\approx \frac{1}{4N} \mathbb{E} [|\Delta U_{\text{gGM}}|^2] \\ &\approx \frac{(N-1)}{8} \sqrt{\mathcal{E}_S^2 + \mathcal{E}_A^2} \\ &\approx \frac{N}{8} \sqrt{\text{Var}(S)^2 + \text{Var}(A)^2}. \end{aligned} \quad (\text{B6})$$

For the case where $\text{Var}(S_{\ell}) = \text{Var}(A_{\ell}) = 2\sigma^2$, the expected infidelity (error variance) for the generalized Green Machine scales as

$$\bar{\mathcal{E}}_{\text{gGM}} \approx \frac{N}{2\sqrt{2}} \sigma^2, \quad (\text{B7})$$

which gives us the $\sqrt{2}$ improvement in the mean fidelity seen in Fig. 2(b) of the main text. Furthermore, a broader distribution of infidelity (larger variances) is also seen in the histograms from Figs. 2(b) and 8.

Finally, we extend this single-particle analysis to the case of an input state $|\psi\rangle$ containing N_{ph} photons. In a linear lossless optical network, the transformation is particle-conserving; a coherent phase error ϕ in the mesh creates a phase rotation $e^{i\phi}$ on each traversing particle. For the many-body state, this accumulates into a global phase shift $U(\phi) \sim e^{iN_{\text{ph}}\phi}$. The total error operator ΔU is determined by the derivative of this transformation with respect to the phase error, $\partial U / \partial \phi$. In the second-quantized formalism, this promotes the error generator to the total photon number operator $\hat{n} = \sum \hat{a}_i^{\dagger} \hat{a}_i$. Since the infidelity is dominated by the variance of this error generator, and errors accumulate additively across the independent excitations in the limit of small perturbations, the total infidelity scales linearly with the photon number for both correlated and uncorrelated error models. Combining this particle scaling with the mesh scaling derived above yields the total infidelity,

$$\mathcal{E} \propto N_{\text{ph}}(N-1)\sigma^2 \approx N_{\text{ph}}N\sigma^2. \quad (\text{B8})$$

APPENDIX C: INACCESSIBLE STATES FROM INTRA-MZI IMBALANCE OF LOSS

In practice, the unbalanced transmission coefficients on each arm of the MZI, denoted by (Γ_1, Γ_2) in Fig. 9(a), resulting in a reduced fraction of the unitary group that can be realized. This drop in expressivity results in groups of states that cannot be prepared by the linear optical circuit. These inaccessible output states that arise as a result of coherent errors and unbalanced loss can be determined by evaluating the range of admissible splitting ratios implemented by the end-to-end transfer matrix of our single

MZI beamsplitter [31]. For a given 2×2 transfer matrix U implemented on a set of time-bin modes, the complex-valued splitting ratio is defined as $s = U_{11}/U_{12}$. The range

of admissible splitting ratios in the presence of splitter error (α, β) and unbalanced transmission coefficients (Γ_1, Γ_2) on each arm reduces to

$$\begin{aligned} & \frac{\Gamma_1 [\cos|\alpha - \beta| - \sin|\alpha + \beta|] - \Gamma_2 [\cos|\alpha - \beta| + \sin|\alpha + \beta|]}{\Gamma_1 [\cos|\alpha + \beta| - \sin|\alpha - \beta|] + \Gamma_2 [\cos|\alpha + \beta| + \sin|\alpha - \beta|]} \leq |s| \\ & \leq \frac{\Gamma_1 [\cos|\alpha - \beta| - \sin|\alpha + \beta|] + \Gamma_2 [\cos|\alpha - \beta| + \sin|\alpha + \beta|]}{\Gamma_1 [\cos|\alpha + \beta| - \sin|\alpha - \beta|] - \Gamma_2 [\cos|\alpha + \beta| + \sin|\alpha - \beta|]}, \end{aligned} \quad (\text{C1})$$




FIG. 9. An illustration of accessible regions of tunable beam-splitter's splitting ratios restricted by loss and coherent errors. (a) Clement mesh circuit implemented by the recursive structure of the generalized Green Machine unrolled in time, identifying the two imbalanced loss factors Γ_1 and Γ_2 taken into account. (b) Riemann sphere representation of accessible splitting ratios, with the forbidden regions shown in red. Purely unbalanced loss [transmission coefficients chosen to be $(\Gamma_1 = 0.5, \Gamma_2 = 0.9)$ for illustrative purposes] introduces forbidden regions around the poles. (c) Coherent errors [chosen to be $(\alpha = 0.15, \beta = 0.0)$] perturb the forbidden regions around the poles asymmetrically, which could happen in spatially meshed linear optical circuits. (d) Lower degree of imbalance in the loss [transmission coefficients chosen to be $(\Gamma_1 = 0.7, \Gamma_2 = 0.9)$] on each arm shrinks the forbidden regions.

which reduces to $\tan|\alpha + \beta| \leq |s| \leq \cot|\alpha - \beta|$ in the lossless setting [31] and $(\Gamma_1 - \Gamma_2)/(\Gamma_1 + \Gamma_2) \leq |s| \leq (\Gamma_1 + \Gamma_2)/(\Gamma_1 - \Gamma_2)$ in the absence of coherent errors.

Accessible splitting ratios over a single stage in the presence of these errors can be visualized on a Riemann Sphere, illustrated in Fig. 9. Under purely unbalanced loss where $(\Gamma_1 = 0.5, \Gamma_2 = 0.9)$ chosen for illustrative purposes, forbidden regions near the poles emerge, indicated in red in Fig. 9(b), showing regions where perfect nulling of the power is impossible. In the presence of both unbalanced loss and coherent errors, the forbidden regions that emerge are asymmetric, given by Eq. (C1). Figure 9(c) illustrates this case with $(\alpha = 0.15, \beta = 0.0)$, where the inaccessible region at the $s = 0$ pole expands, while the region at the $s = \infty$ pole contracts. This suggests that the forbidden region at one of the poles can be eliminated entirely, allowing the realization of the perfect cross state for improved matrix fidelity [31,73,74].

When the degree of imbalance in loss on each arm is reduced (i.e., as $\Gamma_1/\Gamma_2 \rightarrow 1$), the inaccessible regions shrink, shown in Fig. 9(d). Perfectly balanced loss eliminates the forbidden regions completely, allowing access to the full range of splitting ratios while resulting in only a scaling down of the output power. Imbalanced loss arising from interfacing each time bin with delay lines asymmetrically over multiple stages [such as the Clements configuration shown in Fig. 9(a)] cascades into a complex chain of errors, further reducing circuit expressivity. This expressivity can be parameterized by two coherent-error parameters α and β , two intra-MZI loss parameters Γ_1 and Γ_2 , and two delay line loss rates, which is the subject of our future work.

APPENDIX D: LARGE-SCALE MULTIPLEXING FOR TENSOR OPERATIONS

While the results in this manuscript have mainly focused on applications involving quantum computing, this architecture is also well suited to implement the large-scale and high-dimensional unitaries necessary for machine learning accelerators. A metric commonly used to evaluate the




FIG. 10. Multiplexed architecture and computational speed for high-dimensional tensor operations. (a) Multiplexed architecture by cascading a single-stage generalized Green Machine for high-dimensional unitary transformations for data encoded in spatial and temporal degrees of freedom. (b) Number of multiply accumulate operations per second in matrix-vector multiplication as a function of the bin width, as decided by the switching speed. The experimental implementation of the Green Machine is marked in Ref. [22] at a bin width of 10 ns. State-of-the-art switching speeds from Ref. [44] is marked, at bin widths of ≈ 4.3 ps.

performance of these accelerators is the number of multiply accumulate operations per second (MACs). Here, we evaluate the computational speed of matrix-vector multiplications implemented by the generalized Green Machine as a function of the width of the time bins used.

Figure 10(a) illustrates a schematic for a spatially and temporally multiplexed Green Machine architecture, consisting of multiple stages of the generalized Green Machine introduced in the main text of this manuscript. By stacking M individual stages together, where each stage operated on N time-bin modes, this processor can be used to perform $M \times M \times N$ tensor operations. Figure 10(b) plots the computational speed of a single stage and the multiplexed stages as a function of the width of the time bins. We mark the speeds of each architecture using the dotted line at a previous experimental implementation [22] from the authors and state-of-the-art switching speeds [44].




FIG. 11. Scaling of the matrix errors for the Clements architecture (with uncorrelated errors $\sigma = 0.02$), the generalized Green Machine (with correlated errors $\sigma = 0.02$) in the absence and presence of hardware error correction. The improvement from correlated errors is shown in the blue histograms, where the Green Machine provides a constant improvement over the Clements mesh. Additional use of hardware error-correction techniques provides a further significant scaling improvement, as indicated in purple.

APPENDIX E: PERFORMANCE UNDER HARDWARE ERROR CORRECTION

The main text of this manuscript demonstrates an improvement in the robustness to coherent errors of the generalized Green Machine compared to the spatial mode Clements architecture. This improvement arises from the correlated nature of errors that is a result of the repeated use of a single MZI. We show in Fig. 2 of the main text that under correlated errors, the scaling law for the state fidelity/matrix error remains identical, with an improvement in the constant prefactor. In Fig. 11 we plot the scaling laws to compare the error tolerance of the Clements architecture to the generalized Green Machine under hardware error correction [28]. In the figure, the darker distributions on the right plot the matrix error for the Green Machine, while the lighter distributions to the left plot the matrix error for the Clements mesh. The blue histograms illustrate the improvement arising from the correlated nature of errors. Additional use of error correction techniques provides a significant scaling improvement, as indicated by the purple histograms.

- [1] Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund *et al.*, Deep learning with coherent nanophotonic circuits, *Nat. Photonics* **11**, 441 (2017).
- [2] S. Bandyopadhyay, A. Sludds, S. Krastanov, R. Hamerly, N. Harris, D. Bunandar, M. Streshinsky, M. Hochberg, and D. Englund, Single-chip photonic deep neural network with forward-only training, *Nat. Photonics* **18**, 1335 (2024).
- [3] S. R. Ahmed, R. Baghdadi, M. Bernadskiy, N. Bowman, R. Braid, J. Carr, C. Chen, P. Ciccarella, M. Cole, J. Cooke

- et al.*, Universal photonic artificial intelligence acceleration, *Nature* **640**, 368 (2025).
- [4] S. Hua, E. Divita, S. Yu, B. Peng, C. Roques-Carmes, Z. Su, Z. Chen, Y. Bai, J. Zou, Y. Zhu *et al.*, An integrated large-scale photonic accelerator with ultralow latency, *Nature* **640**, 361 (2025).
 - [5] Y.-H. Luo, H.-S. Zhong, M. Erhard, X.-L. Wang, L.-C. Peng, M. Krenn, X. Jiang, L. Li, N.-L. Liu, C.-Y. Lu *et al.*, Quantum teleportation in high dimensions, *Phys. Rev. Lett.* **123**, 070505 (2019).
 - [6] C. Zhang, J. F. Chen, C. Cui, J. P. Dowling, Z. Y. Ou, and T. Byrnes, Quantum teleportation of photonic qudits using linear optics, *Phys. Rev. A* **100**, 032330 (2019).
 - [7] C. Cui, K. P. Seshadreesan, S. Guha, and L. Fan, High-dimensional frequency-encoded quantum information processing with passive photonics and time-resolving detection, *Phys. Rev. Lett.* **124**, 190502 (2020).
 - [8] Y. Chi, J. Huang, Z. Zhang, J. Mao, Z. Zhou, X. Chen, C. Zhai, J. Bao, T. Dai, H. Yuan *et al.*, A programmable qudit-based quantum processor, *Nat. Commun.* **13**, 1166 (2022).
 - [9] J. R. Basani, M. Y. Niu, and E. Waks, Universal logical quantum photonic neural network processor via cavity-assisted interactions, *npj Quantum Info.* **11**, 142 (2025).
 - [10] J. Carolan, C. Harrold, C. Sparrow, E. Martín-López, N. J. Russell, J. W. Silverstone, P. J. Shadbolt, N. Matsuda, M. Oguma, M. Itoh *et al.*, Universal linear optics, *Science* **349**, 711 (2015).
 - [11] J. Romero, J. P. Olson, and A. Aspuru-Guzik, Quantum autoencoders for efficient compression of quantum data, *Quantum Sci. Technol.* **2**, 045001 (2017).
 - [12] B. Fischer, M. Chemnitz, B. MacLellan, P. Roztock, R. Helsten, B. Wetzel, B. E. Little, S. T. Chu, D. J. Moss, J. Azaña *et al.*, Autonomous on-chip interferometry for reconfigurable optical waveform generation, *Optica* **8**, 1268 (2021).
 - [13] S. Ou, K. Xue, L. Zhou, C. Lee, A. Sludds, R. Hamerly, K. Zhang, H. Feng, R. Koppurapu, E. Zhong *et al.*, Hypermultiplexed integrated photonics-based optical tensor processor, *Sci. Adv.* **11**, eadu0228 (2025).
 - [14] C. Sparrow, E. Martín-López, N. Maraviglia, A. Neville, C. Harrold, J. Carolan, Y. N. Joglekar, T. Hashimoto, N. Matsuda, J. L. O'Brien *et al.*, Simulating the vibrational quantum dynamics of molecules using photonics, *Nature* **557**, 660 (2018).
 - [15] X. Qiang, X. Zhou, J. Wang, C. M. Wilkes, T. Loke, S. O'Gara, L. Kling, G. D. Marshall, R. Santagati, T. C. Ralph *et al.*, Large-scale silicon quantum photonics implementing arbitrary two-qubit processing, *Nat. Photonics* **12**, 534 (2018).
 - [16] J. Bao, Z. Fu, T. Pramanik, J. Mao, Y. Chi, Y. Cao, C. Zhai, Y. Mao, T. Dai, X. Chen *et al.*, Very-large-scale integrated quantum graph photonics, *Nat. Photonics* **17**, 573 (2023).
 - [17] L. S. Madsen, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, J. F. Bulmer, F. M. Miatto, L. Neuhaus, L. G. Helt, M. J. Collins *et al.*, Quantum computational advantage with a programmable photonic processor, *Nature* **606**, 75 (2022).
 - [18] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu *et al.*, Quantum computational advantage using photons, *Science* **370**, 1460 (2020).
 - [19] Y. Xia, W. Li, Q. Zhuang, and Z. Zhang, Quantum-enhanced data classification with a variational entangled sensor network, *Phys. Rev. X* **11**, 021047 (2021).
 - [20] Z.-H. Liu, R. Brunel, E. E. Østergaard, O. Cordero, S. Chen, Y. Wong, J. A. Nielsen, A. B. Bregnsbo, S. Zhou, H.-Y. Huang *et al.*, Quantum learning advantage on a scalable photonic platform, *Science* **389**, 1332 (2025).
 - [21] S. Guha, Structured optical receivers to attain superadditive capacity and the Holevo limit, *Phys. Rev. Lett.* **106**, 240502 (2011).
 - [22] C. Cui, J. Postlewaite, B. N. Saif, L. Fan, and S. Guha, Superadditive communication with the green machine as a practical demonstration of nonlocality without entanglement, *Nat. Commun.* **16**, 3760 (2025).
 - [23] M. Rosati and A. Solana, Joint-detection learning for optical communication at the quantum limit, *Opt. Quantum* **2**, 390 (2024).
 - [24] M. Rosati and G. Cincotti, A Fourier machine for quantum optical communications, *J. Lightwave Technol.* **43**, 1 (2025).
 - [25] M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani, Experimental realization of any discrete unitary operator, *Phys. Rev. Lett.* **73**, 58 (1994).
 - [26] W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, and I. A. Walmsley, Optimal design for universal multiport interferometers, *Optica* **3**, 1460 (2016).
 - [27] J. Carolan, M. Mohseni, J. P. Olson, M. Prabhu, C. Chen, D. Bunandar, M. Y. Niu, N. C. Harris, F. N. Wong, M. Hochberg *et al.*, Variational quantum unsampling on a quantum photonic processor, *Nat. Phys.* **16**, 322 (2020).
 - [28] S. Bandyopadhyay, R. Hamerly, and D. Englund, Hardware error correction for programmable photonics, *Optica* **8**, 1247 (2021).
 - [29] R. Hamerly, S. Bandyopadhyay, and D. Englund, Accurate self-configuration of rectangular multiport interferometers, *Phys. Rev. Appl.* **18**, 024019 (2022).
 - [30] R. Hamerly, S. Bandyopadhyay, and D. Englund, Stability of self-configuring large multiport interferometers, *Phys. Rev. Appl.* **18**, 024018 (2022).
 - [31] R. Hamerly, S. Bandyopadhyay, and D. Englund, Asymptotically fault-tolerant programmable photonics, *Nat. Commun.* **13**, 6831 (2022).
 - [32] R. Burgwal, W. R. Clements, D. H. Smith, J. C. Gates, W. S. Kolthammer, J. J. Renema, and I. A. Walmsley, Using an imperfect photonic network to implement random unitaries, *Opt. Express* **25**, 28236 (2017).
 - [33] S. P. Kumar, L. Neuhaus, L. G. Helt, H. Qi, B. Morrison, D. H. Mahler, and I. Dhand, Mitigating linear optics imperfections via port allocation and compilation, *arXiv:2103.03183*.
 - [34] J. Ewaniuk, J. Carolan, B. J. Shastri, and N. Rotenberg, Imperfect quantum photonic neural networks, *Adv. Quantum Technol.* **6**, 2200125 (2023).
 - [35] R. Hamerly, S. Bandyopadhyay, A. Sludds, Z. Chen, L. Bernstein, S. K. Vadlamani, J. Basani, R. Davis, and D. Englund, in *AI and Optical Data Sciences IV* (SPIE, San Francisco, California, United States, 2023), Vol. 12438, p. 1243802.
 - [36] J. R. Basani, M. Heuck, D. R. Englund, and S. Krashtanov, All-photonic artificial-neural-network processor via nonlinear optics, *Phys. Rev. Appl.* **22**, 014009 (2024).

- [37] R. Hamerly, L. Bernstein, A. Sludds, M. Soljačić, and D. Englund, Large-scale optical neural networks based on photoelectric multiplication, *Phys. Rev. X* **9**, 021032 (2019).
- [38] K. R. Motes, A. Gilchrist, J. P. Dowling, and P. P. Rohde, Scalable boson sampling with time-bin encoding using a loop-based architecture, *Phys. Rev. Lett.* **113**, 120501 (2014).
- [39] Y. He, X. Ding, Z.-E. Su, H.-L. Huang, J. Qin, C. Wang, S. Unsleber, C. Chen, H. Wang, Y.-M. He *et al.*, Time-bin-encoded boson sampling with a single-photon device, *Phys. Rev. Lett.* **118**, 190501 (2017).
- [40] S. Sempere-Llagostera, R. B. Patel, I. A. Walmsley, and W. S. Kolthammer, Experimentally finding dense subgraphs using a time-bin encoded Gaussian boson sampling device, *Phys. Rev. X* **12**, 031045 (2022).
- [41] M. Monika, F. Nosrati, A. George, S. Sciara, R. Fazili, A. L. Marques Muniz, A. Bisianov, R. Lo Franco, W. J. Munro, M. Chemnitz *et al.*, Quantum state processing through controllable synthetic temporal photonic lattices, *Nat. Photonics* **19**, 95 (2025).
- [42] S. Yu, Z.-P. Zhong, Y. Fang, R. B. Patel, Q.-P. Li, W. Liu, Z. Li, L. Xu, S. Sagona-Stophel, E. Mer *et al.*, A universal programmable Gaussian boson sampler for drug discovery, *Nat. Comput. Sci.* **3**, 839 (2023).
- [43] F. Pegoraro, P. Held, J. Lammers, B. Brecht, and C. Silberhorn, Demonstration of a photonic time-multiplexed C-NOT gate, [arXiv:2412.02478](https://arxiv.org/abs/2412.02478).
- [44] F. Bouchard, K. Fenwick, K. Bonsma-Fisher, D. England, P. J. Bustard, K. Heshami, and B. Sussman, Programmable photonic quantum circuits with ultrafast time-bin encoding, *Phys. Rev. Lett.* **133**, 090601 (2024).
- [45] F. Ewert and P. van Loock, 3/4-efficient Bell measurement with passive linear optics and unentangled ancillae, *Phys. Rev. Lett.* **113**, 140403 (2014).
- [46] N. Hauser, M. J. Bayerbach, S. E. D'Aurelio, R. Weber, M. Santandrea, S. P. Kumar, I. Dhand, and S. Barz, Boosted Bell-state measurements for photonic quantum computation, *npj Quantum Inf.* **11**, 41 (2025).
- [47] M. Pant, D. Towsley, D. Englund, and S. Guha, Percolation thresholds for photonic quantum computing, *Nat. Commun.* **10**, 1070 (2019).
- [48] J. R. Basani, S. K. Vadlamani, S. Bandyopadhyay, D. R. Englund, and R. Hamerly, A self-similar sine-cosine fractal architecture for multiport interferometers, *Nanophotonics* **12**, 975 (2023).
- [49] S. K. Vadlamani, D. Englund, and R. Hamerly, Transferable learning on analog hardware, *Sci. Adv.* **9**, eadh3436 (2023).
- [50] P. Dhara, D. Englund, and S. Guha, Entangling quantum memories via heralded photonic Bell measurement, *Phys. Rev. Res.* **5**, 033149 (2023).
- [51] J. Calsamiglia and N. Lütkenhaus, Maximum efficiency of a linear-optical Bell-state analyzer, *Appl. Phys. B* **72**, 67 (2001).
- [52] N. C. Harris, G. R. Steinbrecher, M. Prabhu, Y. Lahini, J. Mower, D. Bunandar, C. Chen, F. N. Wong, T. Baehr-Jones, M. Hochberg *et al.*, Quantum transport simulations in a programmable nanophotonic processor, *Nat. Photonics* **11**, 447 (2017).
- [53] P. W. Anderson, Absence of diffusion in certain random lattices, *Phys. Rev.* **109**, 1492 (1958).
- [54] D. Su, C. R. Myers, and K. K. Sabapathy, Conversion of Gaussian states to non-Gaussian states using photon-number-resolving detectors, *Phys. Rev. A* **100**, 052301 (2019).
- [55] S. Konno, W. Asavanant, F. Hanamura, H. Nagayoshi, K. Fukui, A. Sakaguchi, R. Ide, F. China, M. Yabuno, S. Miki *et al.*, Logical states for fault-tolerant quantum computation with propagating light, *Science* **383**, 289 (2024).
- [56] H. Aghaee Rad, T. Ainsworth, R. Alexander, B. Altieri, M. Askarani, R. Baby, L. Banchi, B. Baragiola, J. Bourassa, R. Chadwick *et al.*, Scaling and networking a modular photonic quantum computer, *Nature* **638**, 912 (2025).
- [57] S. Guha, T. C. John, Z. Gong, and P. Basu, Quantum-enhanced quickest change detection of transmission loss, *Phys. Rev. Lett.* **135**, 210801 (2025).
- [58] S. Yu and N. Park, Heavy tails and pruning in programmable photonic circuits for universal unitaries, *Nat. Commun.* **14**, 1853 (2023).
- [59] R. Hamerly, J. R. Basani, A. Sludds, S. K. Vadlamani, and D. Englund, Toward the information-theoretic limit of programmable photonics, *APL Photonics* **10** (2025).
- [60] S. Mosca, F. Bilotti, A. Toscano, and L. Vegni, A novel design method for Blass matrix beam-forming networks, *IEEE Trans. Antennas Propag.* **50**, 225 (2002).
- [61] C. Taballione, T. A. Wolterink, J. Lugani, A. Eckstein, B. A. Bell, R. Grootjans, I. Visscher, D. Geskus, C. G. Roeloffzen, J. J. Renema *et al.*, 8×8 reconfigurable quantum photonic processor based on silicon nitride waveguides, *Opt. Express* **27**, 26842 (2019).
- [62] K. Suzuki, K. Tanizawa, T. Matsukawa, G. Cong, S.-H. Kim, S. Suda, M. Ohno, T. Chiba, H. Tadokoro, M. Yanagihara *et al.*, Ultra-compact 8×8 strictly-non-blocking Si-wire PILOSS switch, *Opt. Express* **22**, 3887 (2014).
- [63] K. Suzuki, R. Konoike, J. Hasegawa, S. Suda, H. Matsuura, K. Ikeda, S. Namiki, and H. Kawashima, Low-insertion-loss and power-efficient 32×32 silicon photonics switch with extremely high- δ silica PLC connector, *J. Lightwave Technol.* **37**, 116 (2019).
- [64] L. He, M. Zhang, A. Shams-Ansari, R. Zhu, C. Wang, and L. Marko, Low-loss fiber-to-chip interface for lithium niobate photonic integrated circuits, *Opt. Lett.* **44**, 2314 (2019).
- [65] L. Chang, M. H. Pfeiffer, N. Volet, M. Zervas, J. D. Peters, C. L. Manganelli, E. J. Stanton, Y. Li, T. J. Kippenberg, and J. E. Bowers, Heterogeneous integration of lithium niobate and silicon nitride waveguides for wafer-scale photonic integrated circuits on silicon, *Opt. Lett.* **42**, 803 (2017).
- [66] PsiQuantum team, A manufacturable platform for photonic quantum computing, *Nature* **641**, 876 (2025).
- [67] N. T. Arnold, M. Victora, M. E. Goggin, and P. G. Kwiat, in *Quantum Computing, Communication, and Simulation III* (SPIE, San Francisco, California, United States, 2023), Vol. 12446, pp. 25–30.
- [68] A. Melkozerov, A. Avanesov, I. Dyakonov, and S. Straupe, Analysis of optical loss thresholds in the fusion-based quantum computing architecture, *APL Quantum* **1**, 036119 (2024).

- [69] N. Maring, A. Fyrrillas, M. Pont, E. Ivanov, P. Stepanov, N. Margaria, W. Hease, A. Pishchagin, A. Lemaître, I. Sagnes *et al.*, A versatile single-photon-based quantum computing platform, *Nat. Photonics* **18**, 603 (2024).
- [70] C. Oh, M. Liu, Y. Alexeev, B. Fefferman, and L. Jiang, Classical algorithm for simulating experimental Gaussian boson sampling, *Nat. Phys.* **20**, 1461 (2024).
- [71] J. Basani, Cascaded optical systems approach to neural networks (CasOptAx), 2024, <https://github.com/JasvithBasani/CasOptAx>.
- [72] N. J. Russell, L. Chakhmakhchyan, J. L. O’Brien, and A. Laing, Direct dialling of Haar random unitary matrices, *New J. Phys.* **19**, 033007 (2017).
- [73] D. A. B. Miller, Perfect optics with imperfect components, *Optica* **2**, 747 (2015).
- [74] K. Suzuki, G. Cong, K. Tanizawa, S.-H. Kim, K. Ikeda, S. Namiki, and H. Kawashima, Ultra-high-extinction-ratio 2×2 silicon optical switch with variable splitter, *Opt. Express* **23**, 9086 (2015).