


Accurate Self-Configuration of Rectangular Multiport Interferometers

Ryan Hamerly^{1,2,*}, Saumil Bandyopadhyay¹, and Dirk Englund¹

¹Research Laboratory of Electronics, MIT, 50 Vassar Street, Cambridge, Massachusetts 02139, USA

²NTT Research Inc., PHI Laboratories, 940 Stewart Drive, Sunnyvale, California 94085, USA

 (Received 21 January 2022; revised 27 June 2022; accepted 28 June 2022; published 5 August 2022)

Multiport interferometers based on integrated beamsplitter meshes are widely used in photonic technologies. While the rectangular mesh is favored for its compactness and uniformity, its geometry resists conventional self-configuration approaches, which are essential to programming large meshes in the presence of fabrication error. Here, we present a configuration algorithm, related to the 2×2 block decomposition of a unitary matrix, that overcomes this limitation. Our proposed algorithm is robust to errors, requires no prior knowledge of the process variations, and relies only on external sources and detectors. We show that self-configuration using this technique reduces the effect of fabrication errors by the same quadratic factor observed in triangular meshes. This relaxes a significant limit to the size of multiport interferometers, removing a major roadblock to the scaling of optical quantum and machine-learning hardware.

DOI: [10.1103/PhysRevApplied.18.024019](https://doi.org/10.1103/PhysRevApplied.18.024019)

I. INTRODUCTION

Large-scale programmable photonic circuits are the cornerstone of many emerging technologies, including quantum computing [1,2], machine learning acceleration [3–5], and microwave photonics [6,7]. One such circuit—the universal multiport interferometer, which functions as a linear optical input-output device with a programmable transfer matrix (Fig. 1)—is of special importance due to its generality and broad range of applications [8]. The most scalable designs involve meshes of Mach-Zehnder interferometers (MZIs): while the triangular RECK mesh [Fig. 1(a)] was initially employed and is straightforward to configure [9], more recently work has shifted to the CLEMENTS rectangle [Fig. 1(b)], which offers clear advantages of compactness, path-length uniformity, and reduced sensitivity to loss [10]. Much recent study has focused on scaling [11,12] and optimizing [13–16] MZI meshes based on the CLEMENTS design.

A major challenge to scaling MZI meshes is the presence of component errors due to fabrication imperfections. Errors cause each MZI transfer matrix to deviate from its programmed value [Fig. 1(c)]; since the overall circuit is a cascade of MZIs with $O(N)$ depth, the total error in the transfer matrix scales as $O(\sqrt{N})$, where N is the circuit size (assuming uncorrelated errors). At large mesh sizes, these errors (if left uncorrected) place unreasonable constraints on fabrication tolerances, ultimately limiting the

scaling of multiport interferometers. Error-correction techniques are therefore critical for large-scale programmable photonics. Global optimization [13–15] and *in situ* training [17] are promising in principle, but computationally inefficient and the optimization result is device specific. Local per-MZI correction is also effective, but requires prior characterization of the device errors [18,19]. A number of efficient “self-configuring” algorithms have been proposed [20–22], but these generally only work for triangular meshes and require large numbers of internal power monitors [23], a significant addition in hardware complexity. Recently, we proposed an efficient error-correction algorithm that does not require internal detectors or accurate precharacterization [24]. However, this algorithm only works for triangular (i.e., RECK) meshes, which excludes the more efficient CLEMENTS design.

In this article, we present a self-configuring strategy that is naturally adapted to the CLEMENTS design and requires no additional hardware complexity beyond external sources and (coherent) detectors. This algorithm proceeds by configuring the diagonals of the mesh, starting at the corners, in a manner that progressively zeroes out the elements of a target matrix through a sequence of Givens rotations. Numerical experiments on imperfect MZI meshes show that our algorithm is stable and robust, and reduces errors by at least a quadratic factor (while sufficiently small errors are corrected exactly), consistent with the behavior observed for triangular meshes [24]. This significantly relaxes the scaling constraints posed by imperfections in realistic MZI meshes. As an example, we consider the application to optical neural networks

*rhamerly@mit.edu

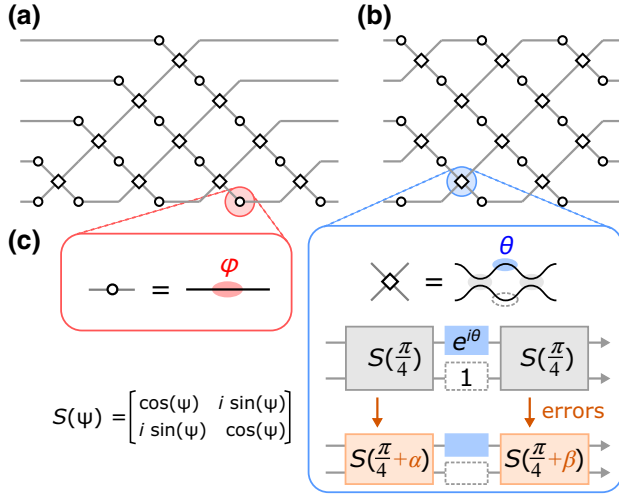


FIG. 1. (a) RECK and (b) CLEMENTS 5×5 MZI meshes. (c) Tunable building blocks: MZI and phase shifter. Effect of errors on the MZI transfer function.

and show how this correction scheme provides a path to overcome the no-go results of Ref. [25].

II. ALGORITHM

Our procedure is based on the diagonalization of a unitary using 2×2 Givens rotations [10], which we review here for clarity. Following Fig. 2(a), divide the mesh along the rising diagonal so that its transfer matrix becomes $U = VDW$, where D is a phase screen and V and W represent the upper and lower triangles, given by

$$V = (T_{21}T_{22})(T_{41} \cdots T_{44}) \cdots (T_{N-2,1} \cdots T_{N-2,N-2}),$$

$$W = (T_{N-1,N-1} \cdots T_{N-1,1}) \cdots (T_{33} \cdots T_{31})(T_{11}),$$

where T_{mn} is the 2×2 block unitary corresponding to the n th crossing (MZI and phase-shift pair) of the m th diagonal. Following the order in the figure, we “build up” the matrices (V , W) one block at a time while keeping track of $X = V^\dagger U W^\dagger$. Figure 2(b) shows the first few steps of this process. Updates to $W \rightarrow T_{mn} W$ right multiply $X \rightarrow X T_{mn}^\dagger$, and the phases of T_{mn} are chosen to zero an element in the lower-left corner of X . Likewise, updates to $V \rightarrow V T_{mn}^\dagger$ left multiply $X \rightarrow T_{mn}^\dagger X$. Following the order in Fig. 2(c), this procedure zeroes all elements in the lower-left triangle of X , enforcing diagonality. The remaining phases are read off from the diagonal elements.

While this procedure correctly sets the phases (θ , ϕ) for an ideal CLEMENTS mesh, it does not work in the presence of errors because the relationship between (θ , ϕ) and T_{mn} also depends on the splitting angle imperfections (α , β) [Fig. 1(c)], which are unknown. Here, we describe the procedure for programming CLEMENTS in the presence of errors.

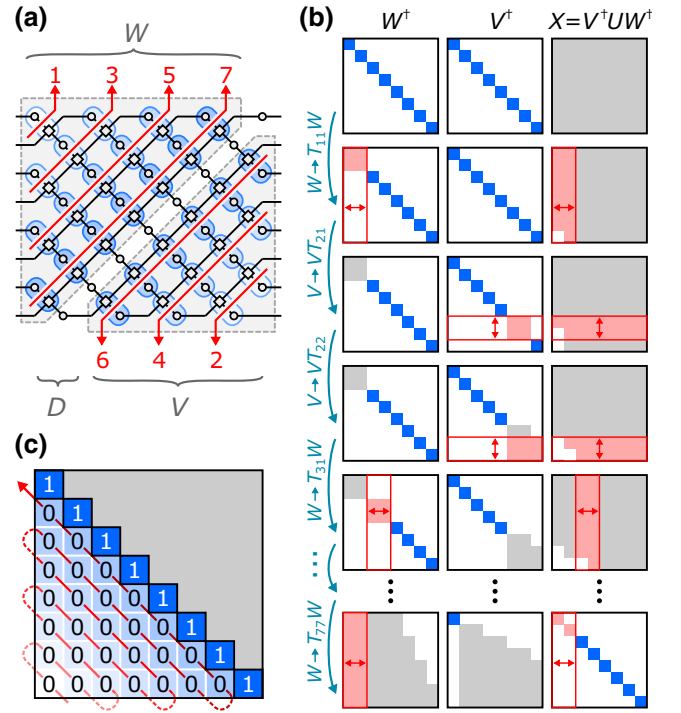


FIG. 2. (a) CLEMENTS mesh and order of the MZI configuration. (b) Building up the matrices V , W in a sequence of Givens rotations to diagonalize $X = V^\dagger U W^\dagger$. (c) Order of matrix elements zeroed by the procedure.

1. Initialize the MZIs to approximate the cross state ($\theta = 0$). An ideal cross state is not possible with errors, but an approximation will be fine.
2. Configure the crossings in the order given by Fig. 2(a). For each crossing in W (respectively V),
 - (a) perform the Givens rotation $X \rightarrow X T_{mn}^\dagger$ (respectively $X \rightarrow T_{mn}^\dagger X$) that zeroes the next element of X_{ij} in the sequence Fig. 2(c) [note that $(i, j) \neq (m, n)$];
 - (b) update $W \rightarrow T_{mn} W$ (respectively $V \rightarrow V T_{mn}^\dagger$);
 - (c) send input $\vec{a}_{\text{in}} = \vec{w}_j^*$ (the j th column of W^\dagger) into the device; see Fig. 3. The output $\vec{a}_{\text{out}}(\theta, \phi)$ depends on the phases being configured. Set (θ, ϕ) to zero the inner product $\langle \vec{v}_i | \vec{a}_{\text{out}}(\theta, \phi) \rangle$, where \vec{v}_i is the i th column of V .
3. Finally, set the diagonal: for each i , inject $\vec{a}_{\text{in}} = \vec{w}_i^*$ and adjust ϕ to satisfy $\arg(\langle \vec{v}_i | \vec{a}_{\text{out}}(\phi) \rangle) = \arg(X_{ii})$.

Of the main loop (step 2), parts (a)–(b) are just a restatement of the CLEMENTS factorization [10]. However, this is performed merely to keep track of matrices (V , W); we do not use the (θ , ϕ) provided by the algorithm. Instead, step 2(c) uses physical measurements to find the correct phases in the presence of hardware errors. This step, which amounts to zeroing matrix element $X_{ij} = v_i^* U w_j^*$ of the physical hardware, is visualized in Fig. 3.

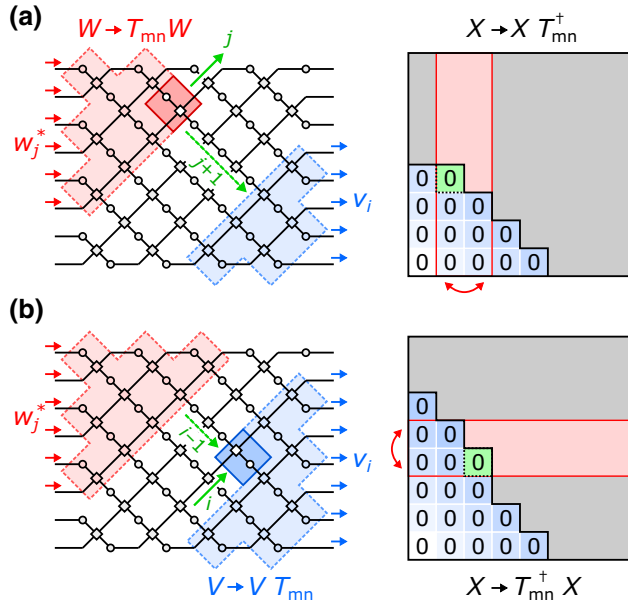


FIG. 3. Programming a physical CLEMENTS mesh in the presence of errors. (a) Procedure to set phases for a crossing in W . (b) Procedure for V .

1. When configuring a W [Fig. 3(a)], we input $\vec{a}_{\text{in}} = \vec{w}_j^*$ and program (θ, ϕ) to direct all the light to intermediate output j (solid arrow), zeroing the power that goes to $j + 1$ (dashed arrow). We do not have access to these intermediate outputs, but $j + 1$ connects to the input i of V ; therefore, $\langle \vec{v}_i | \vec{a}_{\text{out}} \rangle$ is a valid proxy for this field and zeroing it correctly configures the red block to match W .

2. When configuring V [Fig. 3(b)], we want \vec{v}_i to be the output from light at intermediate port i (solid arrow); therefore, the output from light at port $i - 1$ (dashed arrow) should be orthogonal to \vec{v}_i . We excite this field by inputting \vec{w}_j^* ; any errors in the unconfigured mesh will hit downstream inputs $k < i - 1$, but will not pollute input i . Therefore, zeroing $\langle \vec{v}_i | \vec{a}_{\text{out}} \rangle$ correctly configures the blue block of the mesh to match V , provided W is properly configured.

In both cases, the nulling signal $\langle \vec{v}_i | \vec{a}_{\text{out}} \rangle$ arises from light transmitted down the dashed diagonals in Fig. 3; a (near) cross-state initialization is needed to maintain a non-negligible transmitted power. As we discuss later, this initialization is limited by random phase shifts in the MZIs, which can be fixed with a one-time calibration. It is also relatively straightforward to extend our proposed scheme to triangular meshes, as described in Appendix A.

III. RESULTS

To test our algorithm, we perform numerical experiments on CLEMENTS meshes of size up to 128×128 . The algorithms are implemented in PYTHON and are available as part of the MESHES package [26]. We consider an error

model based on imperfections in splitter angles α, β [Fig. 1(c)], as these are the dominant effect of fabrication error in MZI-based silicon photonic circuits. For simplicity, consider the case of uncorrelated Gaussian errors so that the error magnitude can be characterized by a single variable $\sigma = \langle \alpha \rangle_{\text{rms}} = \langle \beta \rangle_{\text{rms}}$ (the case of correlated errors is treated in Appendix B; the qualitative results are the same because most correlations cancel out in ensemble averaging [24]). Target matrices are sampled uniformly over the Haar measure [27,28]. As a figure of merit, we consider the normalized matrix error $\mathcal{E} = \langle \|\Delta U\| \rangle_{\text{rms}} / \sqrt{N}$. For unitary matrices, $\mathcal{E} \in [0, 2]$ corresponds to the average relative error of a given matrix element U_{ij} .

In the uncorrected case, each MZI introduces a mean error $\langle \|\Delta U\| \rangle_{\text{rms}} = \sqrt{2}\sigma$. These errors add in quadrature, leading to an overall normalized error $\mathcal{E}_0 = \sqrt{2N}\sigma$, which grows with mesh size. This is understandable given that a circuit depth that grows as $O(N)$, with each layer contributing $O(\sigma)$ error and the layers adding in quadrature. Figure 4(a) plots the error as a function of σ for a 64×64 mesh, both without error correction (red) and with our algorithm (blue). The algorithm always improves the matrix fidelity, but there are two distinct regimes: for small σ , the error approaches machine precision, as errors can be corrected exactly. On the contrary, for large σ , the corrected error asymptotes to a finite value:

$$\mathcal{E}_c = \sqrt{2/3} N \sigma^2 = \frac{1}{\sqrt{6}} \mathcal{E}_0^2. \quad (2)$$

This form can be derived rigorously from the distribution of MZI splitting angles for Haar-sampled unitary matrices [24,29], where errors arise solely from MZIs whose target splitting ratios cannot be realized in the imperfect hardware. Since $\mathcal{E}_c \propto \mathcal{E}_0^2$, we can say that self-calibration leads to a *quadratic* suppression of errors: the smaller the initial error, the greater the benefit of error correction.

Figures 4(b) and 4(c) test the scalability of the algorithm by varying the mesh size N . Figure 4(b) shows that, as expected, as the mesh size grows, the corrected error asymptotes to the quadratic factor in Eq. (2). By plotting $\mathcal{E}_c / (\mathcal{E}_0^2 / \sqrt{6})$, Fig. 4(c) shows the boundary between the exactly correctable small-error regime ($\mathcal{E}_c = 0$) and the inexact large-error regime where Eq. (2) holds. The regimes meet where the coverage $\text{cov}(N)$ of $U(N)$ dips below unity [29]; we have previously shown that $\text{cov}(N) \sim e^{-N^3 \sigma^2 / 3}$ [24]; therefore, the boundary lies at roughly $N^3 \sigma^2 = 3$ [dashed curve in Fig. 4(c)].

To illustrate the benefits of this error reduction, we consider as an example deep neural network (DNN) inference on optical hardware. DNNs process data in a sequence of layers, each consisting of (linear) synaptic connections and (nonlinear) neuron activation [Fig. 5(a)]. One exciting possibility is to use photonics to accelerate this

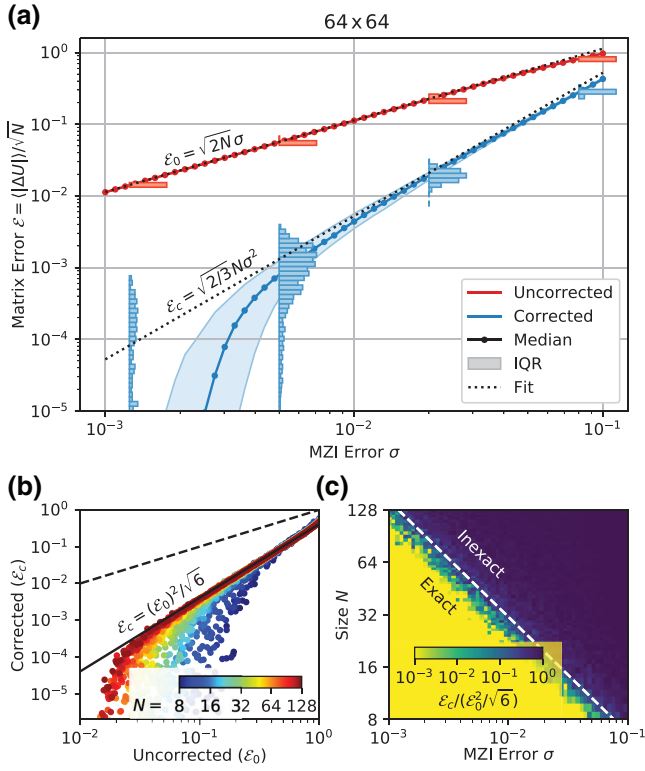


FIG. 4. (a) Corrected and uncorrected matrix errors as a function of σ for a 64×64 CLEMENTS mesh; dots show the median and shaded region is the interquartile range (IQR). (b) Scaling with mesh size, showing the quadratic suppression of errors due to correction. (c) Boundary between the regimes of exact and inexact error corrections.

process: encode the input in optical amplitudes, use a programmable MZI mesh to implement the synaptic weights, and perform the activations with an all-optical (or electro-optic) nonlinearity [3]. A major challenge is that useful learning tasks require large mesh sizes ($N > 100$), which are particularly susceptible to fabrication error; a recent study showed that accurate DNN inference might require unrealistic process tolerances in the hardware [25]. This challenge has spurred investigations into alternative proposals, which have their own limitations [4,5,30].

Figure 5(b) illustrates the advantage of error correction in a MZI-mesh DNN accelerator. Here, input images are preprocessed by a Fourier transform and fed into a two-layer DNN, with electro-optic neuron activations designed to approximate a complex modReLU [31,32]. Models with inner layer sizes $N = 64$ and $N = 256$ are trained on the MNIST digit dataset [33] using the NEUROPHOX package [34]. Details and code are provided in Appendix C and the Supplementary Material [35], respectively. These pre-trained models are then simulated numerically on imperfect CLEMENTS meshes to calculate the classification accuracy. Applying our correction algorithm increases the error

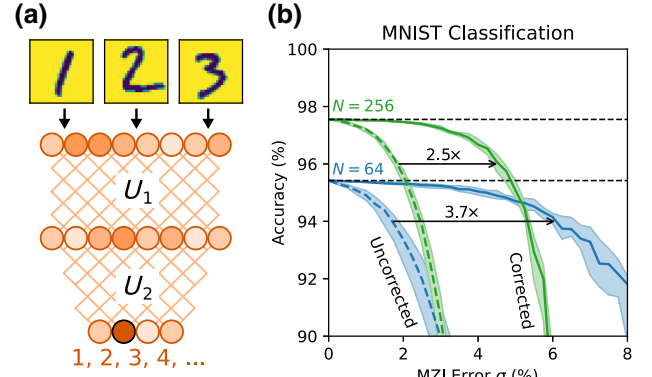


FIG. 5. (a) Two-layer DNN for MNIST classification where synaptic connections are represented by CLEMENTS meshes. (b) Simulated classification accuracy as a function of network size N and hardware MZI error σ .

tolerance of these DNNs by over 2 times. Note that directional couplers in silicon typically exhibit $\sigma \approx 2\%$ [36]; in this regime, the uncorrected DNNs show significant degradation, while error correction restores them to their canonical accuracy. Error correction may even allow the use of broad-band multimode interference couplers, which typically exhibit larger hardware errors.

IV. DISCUSSION

Beyond splitting-ratio errors, realistic meshes must contend with a range of additional imperfections, including undesired phase shifts, unbalanced losses, and limited programming precision. We analyze these error sources in Appendix D, with the main results shown in Fig. 6. Phase-shift and loss errors introduce a random transmission coefficient $e^{i\psi - \ell/2}$ to each waveguide segment, where $\langle \psi \rangle_{\text{rms}} = \sigma_{\text{ph}}$, $\langle \ell \rangle_{\text{rms}} = \sigma_{\ell}$. If known, phase errors can be corrected exactly with a one-time calibration. However, even without calibration, self-configuration still succeeds provided the errors are small enough, and we observe a sharp error-correction threshold of $\sqrt{N}\sigma_{\text{ph}} \lesssim O(1)$. This condition is related to the need to initialize the MZIs close enough to the near-cross state that a significant fraction of the light can propagate down the nulling diagonal (dashed green line in Fig. 3) to give a non-negligible signal $\langle \vec{v}_i | \vec{a}_{\text{out}} \rangle$.

Unlike phase errors, unbalanced losses and programming (DAC) errors are not correctable, as evidenced by the observed behavior $\mathcal{E}_c = \sqrt{N}\sigma_{\ell}$ and $\sqrt{N}\sigma_{\text{dac}}/2$, which follows the same \sqrt{N} scaling observed for uncorrected errors. Loss imbalance is limited by statistical fluctuations in surface roughness; an analysis based on Ref. [37] gives $\sigma_{\ell} \approx \pi^{-1}\alpha_{\text{wg}}\sqrt{L_c L \log(4L/L_c)}$, where α_{wg} is the waveguide loss, L is the waveguide length, and the scatterer size $L_c = \lambda/n_0$ depends on the wavelength and cladding refractive index. For $L = 200 \mu\text{m}$ and $\alpha_{\text{wg}} = 2 \text{ dB/cm}$,




FIG. 6. Comparison of hardware error types: \mathcal{E} plotted against mesh size. Here $\sigma_{bs} = 2\%$, $\sigma_{ph} = 0.2$, $\sigma_l = 5 \times 10^{-4}$, and $\sigma_{dac} = 2\pi/\sqrt{12} \times 2^{-B}$ with $B = 12$ bits of precision. See Appendix D for the derivation.

one calculates $\sigma_l = 5 \times 10^{-4}$, which is close to observed wafer-scale loss variations [38]. DAC error is set by the average truncation to B bits of a signal in the range $[0, 2\pi]$: $\sigma_{dac} = 2\pi/\sqrt{12} \times 2^{-B}$. While neither loss or DAC error are correctable, Fig. 6 clearly shows that these effects are about an order of magnitude smaller than the typical splitter errors, even after correction. These errors will only become relevant if the (correctable) unitary imperfections can be substantially reduced, e.g., with postfabrication trimming [39] or “perfect” MZI-doubled structures [40].

Our algorithm runs in $O(N^2)$ steps and requires only $O(N^3)$ floating point operations per second (FLOPs) of computation. This figure is at least an order of magnitude faster than *in situ* or direct numerical optimization schemes (Appendix E). While local error correction [18] and parallelized progressive methods [16] are faster still, these schemes respectively require calibration and internal detectors, a major challenge to their deployment on large-scale photonic meshes.

The technique presented in this paper requires only external sources and (coherent) detectors and does not rely on an accurate characterization of device errors. This method is based on the diagonalization of a unitary matrix by Givens rotations, with a specific set of measurements performed to ensure that the Givens rotations are properly implemented in the hardware. For sufficiently small hardware errors, our approach leads to perfect realization of the target matrix. For large errors, it achieves the same quadratic reduction $\mathcal{E} \rightarrow \mathcal{E}^2/\sqrt{6}$ observed for local correction algorithms [18] and self-configuration on triangular meshes [24]. For the benchmark optically accelerated DNNs studied here, the proposed technique increases the robustness to hardware error, particularly in the critical

region around $\sigma \approx 2\%$ characteristic of directional couplers in silicon.

One open question is increasing the robustness of error correction to nonunitary errors (unbalanced losses), as many emerging photonic devices often have undesired state-dependent loss [41–43]. In addition, extensions of the algorithm to more recently developed mesh geometries [44–48] may prove useful, as some geometries are less sensitive to hardware error and may be easier to scale up to large dimensions.

ACKNOWLEDGMENTS

S.B. is supported by a NSF Graduate Research Fellowship under Grant No. 1745302 and the Air Force Office of Scientific Research (AFOSR) under Grant No. FA9550-20-1-0113. D.E. acknowledges funding from AFOSR (Grants No. FA9550-20-1-0113 and No. FA9550-16-1-0391). R.H., S.B., and D.E. are inventors on patent application No. 96/196,301 assigned to MIT and NTT Research that covers techniques to suppress component errors in interferometer meshes.

APPENDIX A: TRIANGULAR MESHES

The key to self-configuring the CLEMENTS mesh [10] is realizing that it can be divided into two triangles with a phase screen in the center: $U = VDW$. By configuring one diagonal at a time, alternating between V and W , we can zero all the elements of the target matrix $X = V^\dagger UW^\dagger$, thus realizing the desired unitary. This procedure is simplified when configuring RECK [9], which can be expressed in terms of a single triangle:

$$U = D \left(\underbrace{\prod_{mn} T_{mn}}_W \right) \quad \text{or} \quad \left(\underbrace{\prod_{mn} T_{mn}}_V \right) D. \quad (\text{A1})$$

The first case $U = DW$ corresponds to a mesh with an output phase screen [Fig. 7(a)], while the mesh for $U = VD$ has an input phase screen [Fig. 7(b)]. The algorithm in Sec. II simplifies because we may substitute $V = I$ (respectively $W = I$). This leads to two cases.

1. For the first case [Fig. 7(a)], we work downstream from the leftmost MZI. For example, we can work down the falling diagonals, with the order given in the figure. Each step adds a MZI to $W \rightarrow T_{mn}W$, which updates the target matrix $X \rightarrow XT_{mn}^\dagger$ to zero an element in the upper triangle. With input field $a_{in} = w_j^*$ (the j th column of W^\dagger), the parameters (θ, ϕ) are set to zero the field at output i , where (i, j) is the index of the element X_{ij} being zeroed.

2. For the second case [Fig. 7(b)], we work upstream from the rightmost MZI, which performs the update $V \rightarrow T_{mn}V$, $X \rightarrow T_{mn}^\dagger X$. With light sent into the i th input, (θ, ϕ)

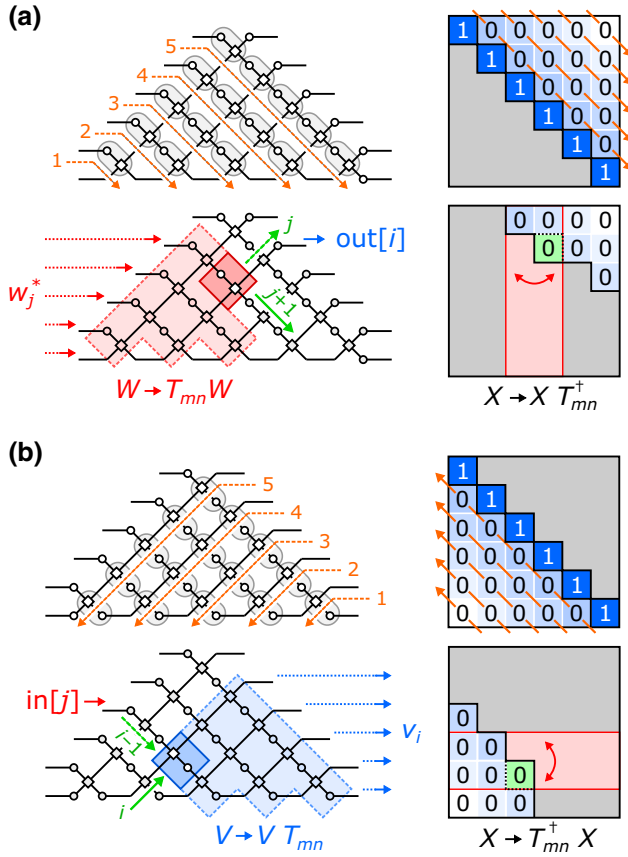


FIG. 7. Self-configuration procedure for the triangular RECK mesh. The procedure depends on whether the mesh layout has (a) an output phase screen or (b) an input phase screen.

are chosen to zero the dot product between the output field and v_i (the i th column of V).

After the crossings are configured, the phases of D can be obtained by inspection.

The MZI order in Fig. 7 is not unique. Any order that preserves causality (the set of configured and unconfigured MZIs must be causally separated) will produce a valid self-configuration.

Like the configuration of the CLEMENTS rectangle, the procedure above is designed, at each time step, to properly set the configured MZIs in the red (respectively blue) block to realize W (respectively V). The procedure in Fig. 1(a) is closely related to the reversed local light interference method (RELLIM) [22]. However, there are two important differences. First, RELLIM inputs columns of U^\dagger , while our scheme inputs columns of W^\dagger . As the MZIs are configured, W changes. Second, RELLIM assumes the existence of internal power detectors after each MZI (or the ability to precalibrate the downstream mesh so that internal fields can be read off from the outputs). By inputting a row of the target W^\dagger , light will exit the red block along only the j th and $(j + 1)$ th channels (provided all upstream MZIs are

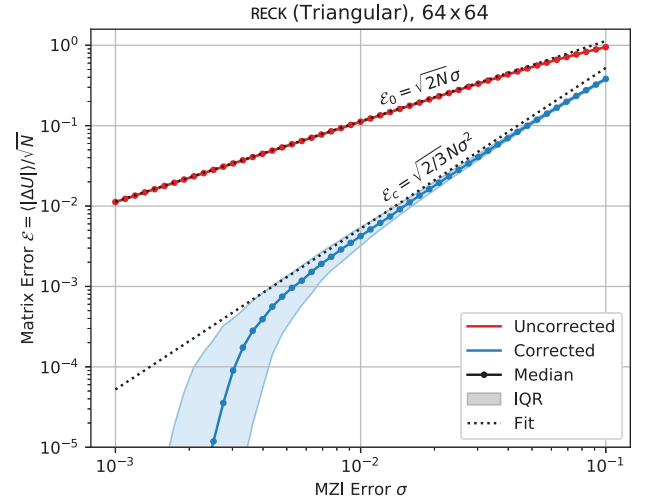


FIG. 8. Corrected and uncorrected matrix errors as a function of MZI error σ for a 64×64 RECK mesh.

correctly set). The former, which is set to zero by varying (θ, ϕ) , is read off directly from the i th output.

The procedures in Figs. 7(a) and 7(b) are reciprocal to each other and can also be related to the “ratio method” described in Ref. [24], since the process of zeroing matrix elements never depends on absolute amplitudes, only on their ratios. Moreover, the ratio method is also robust in the presence of large errors and achieves the same quadratic error reduction. However, when configuring each MZI, the ratio method required a sweep of the upstream phase shifter in order to subtract the common amplitude \bar{a} , a procedure that is not necessary here.

Figure 8 plots the matrix error $\mathcal{E} = \|\Delta U\|/\sqrt{N}$ as a function of MZI error σ for the RECK mesh. The result is identical to the case of the CLEMENTS mesh, Fig. 4, following the curves $\mathcal{E}_0 = \sqrt{2N}\sigma$, $\mathcal{E}_c = \sqrt{2/3N}\sigma^2$ predicted by theory [24]. Over the Haar measure, the distribution of MZI splitting angles is the same for both mesh geometries up to a reordering of the MZIs [29]; therefore, this correspondence is unsurprising.

The self-configuration procedures for RECK and CLEMENTS can be mapped to a general-purpose subroutine that self-configures any MZI mesh of the form $U = VDW$, provided that the geometry admits a matrix diagonalization by way of Givens rotations. This algorithm has been implemented in PYTHON with NUMBA extensions for numerical efficiency, and is available as part of the MESHES package [26]

APPENDIX B: CORRELATED ERRORS

In realistic MZI meshes, the splitter errors α_n, β_n will be strongly correlated, since the process variations that lead to errors (waveguide thickness and spacing, partial etch depth, slab height) all have correlation lengths much




FIG. 9. Corrected and uncorrected matrix errors as a function of correlated MZI error amplitude μ for a 128×128 CLEMENTS mesh.

longer than the size of a MZI. In general, the matrix error $\langle \|\Delta U\|^2 \rangle$ will depend both on the error amplitudes $(\langle \alpha_n^2 \rangle, \langle \beta_n^2 \rangle)$ as well as their correlations $(\langle \alpha_m \alpha_n \rangle, \text{etc.})$. For an individual matrix U , the dependence on correlations can be significant. However, we have previously shown that, in an *ensemble average* of matrices uniformly sampled over the Haar measure, this dependence becomes very small for most inter-MZI correlations because of the random phase shifts between pairs of MZIs [24, Appendix A]. Only intra-MZI correlations $\langle \alpha_n \beta_n \rangle$ have a significant effect on the ensemble-averaged matrix error.

Consider the extreme case of full correlation $\alpha_n = \beta_n = \mu$. This case is realized, for instance, when the dominant error source arises from operating the mesh away from the coupler design wavelength. In Ref. [24], the coverage, uncorrected error, and corrected error are calculated to be

$$\text{cov}(N) = e^{-(2/3)N^3\mu^2}, \quad (\text{B1})$$

$$\mathcal{E}_0 = 2\sqrt{N}\mu, \quad (\text{B2})$$

$$\mathcal{E}_c = \sqrt{8/9N}\mu^2. \quad (\text{B3})$$

Figure 9 plots \mathcal{E}_0 and \mathcal{E}_c against the error amplitude μ for a 128×128 CLEMENTS mesh. Small residual dependencies on correlation (proportional to $\theta_m \theta_n$) lead to a slight deviation in \mathcal{E}_0 , while the theoretical curve for \mathcal{E}_c matches very accurately. The behavior is qualitatively very close to that in the uncorrelated case.

APPENDIX C: NEURAL NETWORK MODEL

The optical neural network model is based on the architecture described in Ref. [16]. Images from the MNIST digit dataset are preprocessed with a Fourier transform,

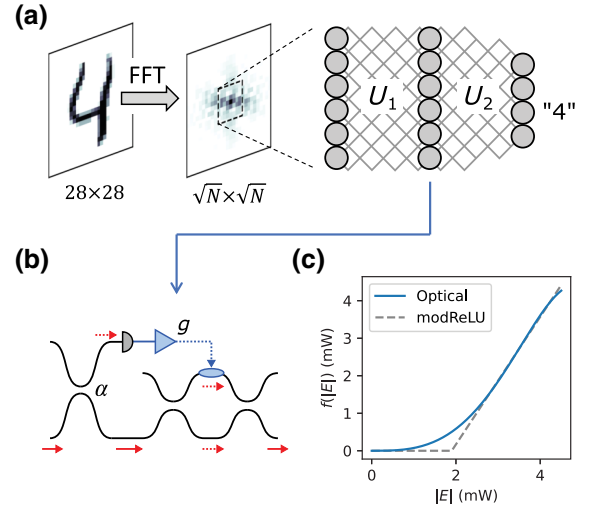


FIG. 10. Neural network model. (a) Dataflow: a 28×28 image is Fourier transformed and cropped to an $N \times N$ window; the complex amplitudes are then fed through a two-layer optical neural network. (b) Electro-optic nonlinearity implementing Eq. (C1). (c) Input-output relation for the nonlinearity, which approximates a complex modReLU.

which is cropped to a $\sqrt{N} \times \sqrt{N}$ window, where N is a model parameter that quantifies the size of the neural network. The light from this window (N input neurons) is fed into a two-layer optically accelerated DNN. This DNN consists of a single inner layer and two $N \times N$ unitary circuits, represented by CLEMENTS meshes [Fig. 10(a)].

The activation function at the inner layer is realized with an electro-optic nonlinearity: a fraction of each output field is fed into a detector that drives a Mach-Zehnder modulator, while the remaining light passes through the modulator [31]. This is shown in Fig. 10(b), and implements the activation function

$$f(E) = \sqrt{1 - \alpha} e^{-i(g|E|^2 + \phi - \pi)/2} \cos\left(\frac{1}{2}(g|E|^2 + \phi)\right), \quad (\text{C1})$$

where α is the power tap fraction, g is the modulator phase induced per unit optical power, and ϕ is the phase in the absence of power. Here, we choose $\alpha = 0.1$, $g = \pi/20$, and $\phi = \pi$, which causes $f(E)$ to approximate the form of a complex modReLU [49] in the right power regime [Fig. 10(c)].

Models of sizes $N = 64$ and $N = 256$ are trained using the NEUROPHOX package [34]. Code and model parameters are provided in the Supplementary Material [35].

APPENDIX D: ADDITIONAL HARDWARE AND CALIBRATION ERRORS

We find that beamsplitter imperfections are the dominant source of error in most situations. Additional hardware imperfections include: (1) random phase shifts ,

which are perfectly correctable provided the error is below a given threshold, and (2) unbalanced loss and finite programming precision, which are in general uncorrectable. The sections below analyze these effects in detail, concluding that, for realistic hardware, errors in the former category are well below the threshold for perfect correction, while errors in the latter category are small enough that they can be ignored.

1. Perfectly correctable errors: phase shifts

While multiport interferometers are length balanced to ensure that all paths accumulate the same phase shift (in the absence of heater power), in practice there will be phase errors due to localized imperfections. If these errors are known, they can be calibrated away and absorbed into the programmable phase shifts (θ, ϕ) of each MZI. Various calibration procedures have been reported in the literature [50,51]. However, even without precalibration, the method reported in this paper can exactly correct for phase-shifter errors provided that their magnitude falls below a given threshold.

To start, consider an error model with random phase shifts on both the internal and external arms, i.e., a given MZI has the 2×2 unitary

$$T(\theta, \phi) \rightarrow S\left(\frac{\pi}{4} + \beta\right) \begin{bmatrix} e^{i(\theta+\psi_1)} & 0 \\ 0 & e^{i\psi_2} \end{bmatrix} S\left(\frac{\pi}{4} + \alpha\right) \times \begin{bmatrix} e^{i(\phi+\psi_3)} & 0 \\ 0 & e^{i\psi_4} \end{bmatrix}, \quad (\text{D1})$$

where errors are uncorrelated and Gaussian: $(\alpha, \beta) \sim N(0, \sigma_{\text{bs}})$ and $\psi_i \sim N(0, \sigma_{\text{ph}})$. In Fig. 11, we numerically perform self-configuration on such faulty CLEMENTS meshes, varying both mesh size N and phase error σ_{ph} , for the cases of ideal ($\sigma_{\text{bs}} = 0$) and nonideal ($\sigma_{\text{bs}} = 0.03$) splitters. In the absence of error correction, the matrix error is calculated to be $\mathcal{E}_0 = \sqrt{2N(\sigma_{\text{ph}}^2 + \sigma_{\text{bs}}^2)}$, which is dominated by the phase shifts when $\sigma_{\text{ph}} > \sigma_{\text{bs}}$. With self-configuration, we see a sharp phase transition between a regime where error correction succeeds ($\mathcal{E}_c = \sqrt{2/3}N\sigma_{\text{bs}}^2$) and fails ($\mathcal{E}_c \approx 1$). This threshold scales as $\sigma_{\text{ph}} \sim 3/\sqrt{N}$, and varies only slightly with the splitter error σ_{bs} .

This threshold arises from imperfect initialization of the mesh, as phase errors cause the initialized MZIs to deviate from the cross state. While the algorithm does not rely on these MZIs being exact crossings, the measurement steps consist of nulling optical signals that depend on fields propagating down a long chain of crosslike MZIs (dashed lines in Fig. 3). As the (amplitude) transmission of an imperfect crossing goes as $t_{\text{MZI}} \sim e^{-(\psi_1 - \psi_2)^2/2}$ for small errors, the average transmission after N stages (there are approximately N crossings during the initial self-configuration steps) is $t = e^{-N\sigma_{\text{ph}}^2}$.

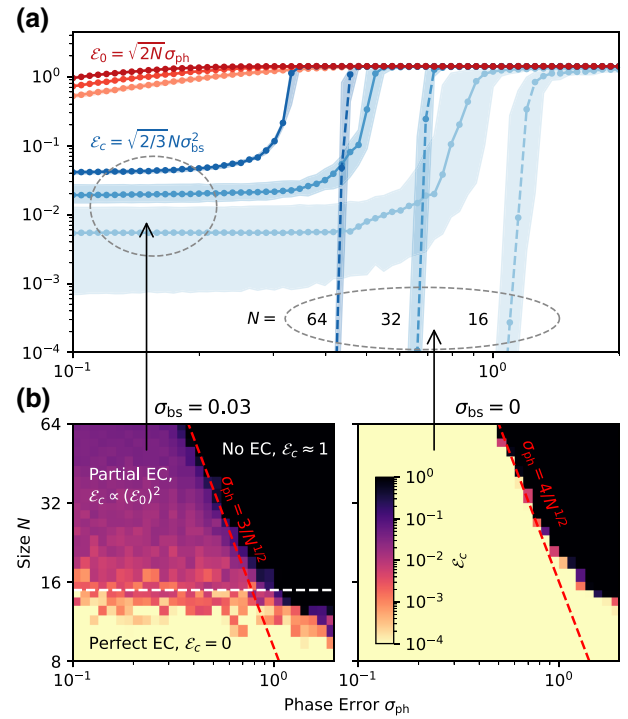


FIG. 11. Effect of phase errors on self-configuration. (a) Error as a function of σ_{ph} for $\sigma_{\text{bs}} \in \{0, 0.03\}$, both uncorrected and corrected. Note that uncorrected error saturates to unity for $\sigma_{\text{ph}} \gtrsim 0.2$. (b) Corrected error as a function of N and σ_{ph} , showing the sharp error correction threshold at $\sigma_{\text{ph}} \sim 3/N^{1/2}$.

The algorithm fails for $\sigma_{\text{ph}} > 4/\sqrt{N}$ even with perfect 50:50 splitters because $t < 10^{-7}$, which is below 32-bit machine precision. In practice, zeroing to such weak signals is impractical in the presence of noise, so realistically one requires $N\sigma_{\text{ph}}^2 \lesssim 2$ in order to maintain a reasonable amplitude for the nulling signal $t \gtrsim 0.1$. This is a factor of 2–3 below the thresholds in Fig. 11, so self-configuration should work to high accuracy.

In a recent 32×32 SOI switch network, Suzuki *et al.* [52] measured a trimming power variability ($\sigma = 1.52$ mW) an order of magnitude lower than the switching power ($P_{\pi} = 18.1$ mW). Putting these together, we calculate $\sigma_{\text{ph}} = 0.2$. This is small enough to comfortably admit self-configuration on meshes up to size $N \approx 64$ ($N\sigma^2 \approx 2.5$). To configure larger meshes, an initial calibration step [50,51] will be needed to comfortably reduce the phase errors so that $N\sigma_{\text{ph}}^2 \lesssim 2$. This calibration need only be performed once, as the phase shifts can be saved as the initial “cross” state for self-configuring to any matrix U .

2. Uncorrectable errors: loss and DAC precision

Loss errors. A uniform waveguide loss does not lead to errors in balanced MZI structures; however, in practice losses can be slightly nonuniform. We can model loss errors by introducing additional amplitude scalings to the

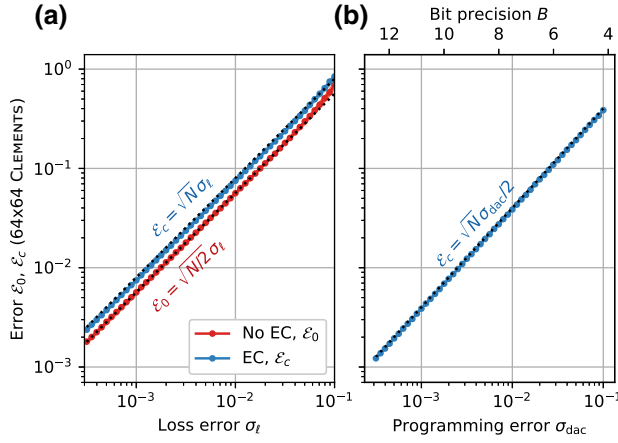


FIG. 12. Uncorrectable errors for a 64×64 CLEMENTS mesh. (a) Effect of loss error on a mesh with and without error correction (EC). (b) Programming error due to finite DAC bit precision or noise.

internal and external arms of the MZI as

$$T(\theta, \phi) \rightarrow S\left(\frac{\pi}{4}\right) \begin{bmatrix} e^{i\theta - \ell_1/2} & 0 \\ 0 & e^{-\ell_2/2} \end{bmatrix} S\left(\frac{\pi}{4}\right) \times \begin{bmatrix} e^{i\phi - \ell_3/2} & 0 \\ 0 & e^{-\ell_4/2} \end{bmatrix}, \quad (\text{D2})$$

where $\ell_i \sim N(0, \sigma_\ell)$ are the loss errors. The uncorrected error is $\mathcal{E}_0 = \sqrt{N/2}\sigma_\ell$. We can perform self-configuration on the mesh to attempt to correct the errors, but in practice the error increases by a factor of $\sqrt{2}$ to $\mathcal{E}_c = \sqrt{N}\sigma_\ell$ [Fig. 12(a)]. The matrix error increases because self-configuration is attempting to correct a nonunitary perturbation with a unitary one, in order to correctly set the splitting ratios of the MZIs. This procedure is generally successful at setting the MZIs to the target splitting ratios. However, a nonunitary matrix is not defined solely by its splitting ratio and output phases, so the logic behind error correction breaks down. In fact, the loss error and the corresponding unitary “correction” are orthogonal, so the overall matrix errors add up, leading to the additional factor of $\sqrt{2}$.

Loss errors likely rule out the use of doped-Si thermal phase shifters [43], as alignment and doping density variations lead to a significant wafer-scale variation 0.23 ± 0.13 dB ($\sigma_\ell = 0.03$). However, most SOI platforms use TiN phase shifters, where the heating element is placed sufficiently far above ($\Delta y \gtrsim 0.5 \mu\text{m}$) as to not interact with the waveguide mode. In such phase shifters, loss is determined entirely by the waveguide loss. Wafer-level statistics for typical SOI processes show that $\alpha = 2.1 \pm 0.25$ dB/cm for standard processing, while H_2 thermal annealing can reduce this to 0.1 ± 0.04 dB/cm. Using these figures, for a $200\text{-}\mu\text{m}$ thermal tuner, $\sigma_\ell = 1.1 \times 10^{-3}$ and 1.8×10^{-4} , respectively [38]. These values are based

on wafer-scale variations, so the actual die-scale values relevant to moderate-sized meshes may be much smaller.

Fundamentally, loss variations are limited by the statistics of sidewall scattering. Consider a waveguide segment of length L with scatterers of size $\approx \lambda/n_{\text{eff}}$; if these scatterers are randomly and independently placed, the number in a given segment will follow a Poisson distribution with mean $\langle n \rangle = L/\lambda$ and standard deviation $\langle \Delta n \rangle_{\text{rms}} = \sqrt{L/\lambda}$. Dividing these quantities, the ratio of loss variation to average loss will be $(\Delta\alpha/\alpha)_L \sim \sqrt{\lambda/L}$. A more rigorous calculation based on the scattering theory of Lacey and Payne [37] (see Appendix D3) gives

$$\frac{\langle \Delta\alpha \rangle_{\text{rms}}}{\langle \alpha \rangle} = \frac{1}{\pi} \sqrt{\frac{\log(4n_0L/\lambda)}{n_0L/\lambda}}, \quad (\text{D3})$$

where n_0 is the cladding index. For a $200\text{-}\mu\text{m}$ phase shifter at $1.55 \mu\text{m}$ with $\alpha = 2$ dB/cm, Eq. (D3) gives $\sigma_\ell = 5 \times 10^{-4}$, roughly consistent with the wafer-scale variations in Ref. [38]. These loss variations are low enough to allow scaling up to $N \approx 1000$ with matrix errors of at most a few percent.

Programming and DAC errors. The phase shifts (θ, ϕ) are always programmed to finite bit precision; moreover, finite signal-to-noise ratios in the nulling signals may limit the programming accuracy. To model these effects, in Fig. 12(b) we self-configure the mesh while introducing a random perturbation to every θ and ϕ after setting it. For example, if the phase is set by a DAC with bounds $[0, 2\pi]$ and bit precision B , truncation to B bits leads to a phase error $\sigma_{\text{dac}} = 2\pi/\sqrt{12} \times 2^{-B}$.

Naively, we would expect these errors to add randomly and in quadrature to give $\sqrt{N}\sigma_{\text{dac}}$. However, subsequent steps in self-configuration correct for exactly half of this error. To see how, note that, analogous to Ref. [18], for any physical MZI

$$T(\theta, \phi) = S\left(\frac{\pi}{4}\right) \begin{bmatrix} e^{i\theta} & 0 \\ 0 & 1 \end{bmatrix} S\left(\frac{\pi}{4}\right) \begin{bmatrix} e^{i\phi} & 0 \\ 0 & 1 \end{bmatrix}, \quad (\text{D4})$$

there exist (ψ_1, ψ_2) such that T is represented by a symmetric splitter with phase shifts on its outputs:

$$T = \begin{bmatrix} e^{i\psi_1} & 0 \\ 0 & e^{i\psi_2} \end{bmatrix} \underbrace{\begin{bmatrix} \cos(\theta/2) & i \sin(\theta/2) \\ i \sin(\theta/2) & \cos(\theta/2) \end{bmatrix}}_{T_{\text{sym}}(\theta)}. \quad (\text{D5})$$

These output phase shifts get corrected in subsequent self-configuration steps. Therefore, the residual error to matrix U comes entirely from the errors in the T_{sym} blocks

$$\|\Delta U\|^2 = \sum_{mn} \|T'_{\text{sym}}(\theta_{mn}) \Delta\theta_{mn}\|^2 = \frac{1}{2} \sum_{mn} \Delta\theta_{mn}^2, \quad (\text{D6})$$

which averages to $\mathcal{E}_c = \sqrt{N}\sigma_{\text{dac}}/2$. With $B \geq 10$ bits of precision, one can scale to $N \approx 1000$ with matrix errors of at most a few percent.

3. Statistical origin of scattering loss variations

Sidewall roughness scattering dominates loss in tightly confined silicon waveguides. This loss is calculated with the volume-current method, where boundary perturbations lead to current sources that scatter into the far field. Lacey and Payne [37, Eq. 16a] showed that, in two dimensions, the loss is given by

$$\alpha = C \int_0^\pi \tilde{R}(\beta - n_0 k_0 \cos \theta) d\theta, \quad (\text{D7})$$

where $k_0 = \omega/c$, $\beta = n_{\text{eff}}k_0$, n_0 is the cladding index, and C is a constant of proportionality. The loss depends on the roughness statistics through $\tilde{R}(k)$, which is the Fourier transform of the autocorrelation function $R(u) = \langle f(z)f(z+u) \rangle$ of the sidewall perturbation $f(z)$. The integral in Eq. (D7) corresponds to the sum over the amplitudes of all scatterers $R(\Delta k_z)$ that phase match the guided mode $k_z = \beta$ to free-space modes $k_z = n_0 k_0 \cos \theta$ that lie within the light cone $|k_z| < n_0 k_0$.

While Eq. (D7) is usually used to calculate the average waveguide loss, sidewall roughness is a statistical quantity and short waveguides will have large relative roughness variations. We can also use Eq. (D7) to calculate the statistical loss properties of finite-length waveguides by replacing the loss perturbation with a periodic function:

$$f^{(L)}(z) = \sum_{m>0} \sqrt{\pi \tilde{R}(k_m)} f_m e^{ik_m z} + \text{c.c.} \quad (\text{D8})$$

Here $k_m = 2\pi m/L$ is the m th Fourier series mode, while f_m is a complex-valued Gaussian with unit norm in both quadratures, i.e., $\langle \text{Re}[f_m]^2 \rangle = \langle \text{Im}[f_m]^2 \rangle = 1$, so that $\langle |f_m|^2 \rangle = 2$. We calculate the autocorrelation and its spectral function:

$$R^{(L)}(u) = \frac{2\pi}{L} \sum_m \tilde{R}(k_m) |f_m|^2 \cos(k_m u), \quad (\text{D9})$$

$$\tilde{R}^{(L)}(k) = \frac{\pi}{L} \sum_m \tilde{R}(k_m) |f_m|^2 \delta(k - k_m). \quad (\text{D10})$$

Note that in the limit $L \rightarrow \infty$, $\tilde{R}^{(L)}(k) \rightarrow \tilde{R}(k)$ in the sense of distributions, i.e., for any smooth $g(k)$, we have $\int \tilde{R}^{(L)}(k)g(k)dk \rightarrow \int \tilde{R}(k)g(k)dk$, as the former becomes a Riemann sum for the latter and the statistical fluctuations of $|f_m|^2$ average out when many modes are summed.

We can rewrite Eq. (D7) as

$$\alpha = C \int_{\beta-n_0k_0}^{\beta+n_0k_0} \tilde{R}(k) \frac{1}{\underbrace{\sqrt{(n_0k_0)^2 - (\beta - k)^2}}_{w(\beta-k)}} dk. \quad (\text{D11})$$

Following Eq. (D10), the scattering loss becomes a discrete sum:

$$\alpha = C \frac{\pi}{L} \sum_m \tilde{R}(k_m) |f_m|^2 w(\beta - k_m). \quad (\text{D12})$$

Note that α depends linearly on the $|f_m|^2$, which are independent and identically distributed random variables. Applying the moments $\langle |f_m|^2 \rangle = 2$, $\langle (\Delta |f_m|^2)^2 \rangle = 4$, we can calculate the mean and variance of α . These yield discrete sums

$$\langle \alpha \rangle = C \frac{\pi}{L} \sum_m \tilde{R}(k_m) \langle |f_m|^2 \rangle w(k_m), \quad (\text{D13})$$

$$\langle \Delta \alpha^2 \rangle = C^2 \frac{\pi^2}{L^2} \sum_m \tilde{R}(k_m)^2 \langle (\Delta |f_m|^2)^2 \rangle w(k_m)^2, \quad (\text{D14})$$

that may be replaced with integrals in the limit of moderately large L :

$$\begin{aligned} \langle \alpha \rangle &\rightarrow C \int \tilde{R}(k) w(k) dk, & \langle \Delta \alpha^2 \rangle \\ &\rightarrow \frac{2\pi C^2}{L} \int \tilde{R}(k)^2 w(k)^2 dk. \end{aligned} \quad (\text{D15})$$

To solve Eq. (D15), we need to specify the roughness model. For example, a Gaussian model $R(s) = e^{-s^2/2L_c^2}$ would yield $\tilde{R}(k) \sim e^{-k^2 L_c^2/2}$, while an exponential model $R(s) = e^{-|s|/L_c}$ would yield $\tilde{R}(k) \sim 1/(1+k^2 L_c^2)$. In both cases the model is characterized by a scatterer size L_c . Here we assume that the scatterers are much smaller than the optical wavelength $L_c \ll \lambda$, so that we can approximate the spectral density as a constant $\tilde{R}(k) \rightarrow R_0$, and the model details become irrelevant (in practice, $L_c \sim \lambda$ is possible, but the results will be qualitatively similar). Making this assumption, we find that $\langle \alpha \rangle = \pi R_0 C$, while the expression for $\langle \Delta \alpha^2 \rangle$ formally diverges:

$$\langle \Delta \alpha^2 \rangle = \frac{2\pi R_0^2}{L} \int_{-n_0k_0}^{n_0k_0} \frac{1}{(n_0k_0)^2 - k^2} dk \rightarrow \infty. \quad (\text{D16})$$

However, this divergence is merely logarithmic, and the integral is actually just the approximation to a discrete sum, Eq. (D14), over a mode set with spacing $2\pi/L$. Therefore, one should actually impose a cutoff and shrink the bounds

TABLE I. Mean and statistical variation of waveguide loss for a range of lengths, assuming that $\alpha = 2$ dB/cm loss. Ratio $\Delta\alpha/\alpha$ comes from Eq. (D18). Normalized nonunitary error magnitude σ_ℓ is given by $\sigma_\ell = \Delta\alpha L/(4.34 \text{ dB})$.

	L (μm)			
	50	100	200	400
$\Delta\alpha/\alpha$	0.106	0.080	0.060	0.044
αL (dB)	0.01	0.02	0.04	0.08
$\Delta\alpha L$ (dB)	0.0011	0.0016	0.0023	0.0036
σ_ℓ ($\times 10^{-4}$)	2.5	3.7	5.5	8.2

of the integral to $\pm(n_0 k_0 - \pi/L)$, in which case Eq. (D16) reduces to

$$\langle \Delta\alpha^2 \rangle = \frac{2\pi R_0^2 C^2}{n_0 k_0 L} \log(2n_0 k_0 L/\pi) = \frac{\lambda R_0^2 C^2}{n_0 L} \log(4n_0 L/\lambda), \quad (\text{D17})$$

where $\lambda = 2\pi/k_0$ is the optical wavelength. Although $\langle \alpha \rangle$ and $\langle \Delta\alpha^2 \rangle$ are functions of model-dependent constants R_0 and C , their ratio depends only on the refractive index, wavelength, and length of the waveguide segment:

$$\frac{\langle \Delta\alpha \rangle_{\text{rms}}}{\langle \alpha \rangle} = \frac{1}{\pi} \sqrt{\frac{\log(4n_0 L/\lambda)}{n_0 L/\lambda}}. \quad (\text{D18})$$

Table I lists expected losses and their variance for a range of phase-shifter lengths and a typical waveguide loss of 2 dB/cm. The ratio $\Delta\alpha/\alpha \lesssim 0.1$ is also consistent with wafer-scale measurements of waveguide loss variations [38]. Note that $\sigma_\ell < 10^{-3}$ for reasonable phase-shifter lengths, so the expected error $\mathcal{E}_c = \sqrt{N}\sigma_\ell$ is at most a few percent even for meshes as large as $N = 256$.

APPENDIX E: ALGORITHMIC EFFICIENCY

A large number of photonic mesh programming algorithms have been proposed in the literature; see Appendix B of Ref. [24] for a general overview. Table II summarizes some of the leading approaches in terms of both their computational cost, generality (applicable mesh geometries). This list is not exhaustive, but gives a representative sample that illustrates the trade-offs in the field. With the exception of the algorithms recently reported by us [24], to date, all optimization schemes require either (1) accurate precalibration of the hardware errors or (2) $O(N^2)$ internal photodetectors used to monitor power at intermediate points on the mesh. Our algorithm uniquely lacks both requirements, so it can be performed on uncalibrated “zero-change” photonic hardware, requiring only coherent control of the input fields and coherent detection at the outputs, and works for both RECK and CLEMENTS meshes.

Even in the absence of these considerations, our algorithm is competitive in terms of computational

TABLE II. Comparison of a representative sample of optimization algorithms. Scaling of computational cost plotted for *in situ* resources (number of MVM calls) and in-silico resources (CFLOPs), where N is the mesh size, T is the number of optimization steps, and B is the SGD batch size.

		Mesh ^a	MVMs	CFLOPs	Caveats ^b
Progressive	[21]	R	N^2	N^3	(D)
RELLIM	[16,22]	R C	$N-N^2$	N^3	(D)
<i>In situ</i>	[17]	R C	NT	$N^3 T$	(D)
SGD	[15]	R C	BT	$N^2 BT$	(C)
Numerical	[13,14]	R C	NT	$N^3 T$	(C)
Local	[18]	R C	–	N^2	(C)
Direct	[24]	R	N^2	N^3	(S)
Ratio	[24]	R	N^2	N^3	
This work		R C	N^2	N^3	

^aMesh types include (R)ECK and (C)LEMENTS.

^bCaveats include (D) requires internal detectors, (C) requires precalibration of errors, (S) stability issues in the presence of errors.

resources. The required resources of an algorithm depend on whether it is performed *in situ* (for correction of unknown errors) or in silico (for known, calibrated errors). For *in situ* algorithms, the number of matrix-vector multiplication (MVM) calls is most relevant. Our algorithm takes approximately $2N^2$ calls (four calls per MZI), which is comparable to progressive self-configuration and the initial REllim proposal [21,22] [although REllim with internal detectors can be parallelized to run in $O(N)$ time [16]]. Alternatively, numerical optimization with T time steps takes $3NT$ calls (three calls per step, first to estimate U , next to back propagate $U - \hat{U}$, and finally a forward-pass step to compute the gradient with respect to phase shifts [17]). Figure 13(a) shows the convergence of a 64×64 CLEMENTS mesh in terms of MVMs under the limited-memory, Broyden, Fletcher, Goldfarb and Shanno bound-constrained (LBFGSB) algorithm [53]. Convergence depends strongly on the mesh initialization [15]. Initializing to Haar-random unitaries is a significant improvement over random phase shifts, since a mesh with random phase shifts has a banded structure that leads to vanishing gradients with respect to matrix elements far from the diagonal. Even with Haar initialization, the algorithm takes thousands of steps and millions of MVMs to converge to the accuracy \mathcal{E}_c reached by our algorithm. Initializing to the phases of an ideal (error-free) CLEMENTS mesh helps considerably, but optimization still takes 10 times more calls. However, for CLEMENTS meshes, it appears that the global minimum can be as much as 5 times lower than the self-configuration result; as a result, numerical optimization may still be helpful as a refinement technique, where the mesh is initialized using self-configuration, followed by a numerical procedure to improve this solution.




FIG. 13. Convergence of LBFGSB for an imperfect CLEMENTS mesh (64×64 , $\sigma = 0.02$), minimizing the L_2 norm. The convergence curve depends on the initial solution, and is compared with the result from self-configuration. (a) Plot in terms of *in situ* MVMs to solution. (b) In-silico CFLOPs to solution.

Since a single optimization step requires N MVM calls, stochastic gradient descent (SGD) has been proposed as a solution to speed up matrix optimization, where a batch of $B < N$ random columns are used rather than the whole matrix. However, SGD tends to trade-off batch size for iteration count, so the overall resource requirement is higher [15].

Figure 13(b) shows the performance gap with respect to the in-silico case, where accurate calibration of the errors (α, β) allows the phases (θ, ϕ) to be computed numerically. Here, optimization protocols require $O(N^3)$ complex floating point operations per second (CFLOPs), equivalent to one matrix-matrix multiplication per time step, leading to a scaling of $N^3 T$ (SGD scales as $N^2 B T$). Self-configuration runs by performing approximately $N^2/2$ Givens rotations with $4N$ CFLOPs each, for a total of $2N^3$ CFLOPs, and thus runs about 10^2 times faster. As before, numerical optimization can still be helpful as a means for further refinement of the solution.

However, for in-silico optimization, the local correction method [18] is superior, as it takes only N^2 CFLOPs and is parallelizable to N time steps.

- Oguma, M. Itoh, Graham D. Marshall, Mark G. Thompson, Jonathan C. F. Matthews, Toshikazu Hashimoto, Jeremy L. O'Brien, and Anthony Laing, Universal linear optics, *Science* **349**, 711 (2015).
- [2] H.-S. Zhong *et al.*, Quantum computational advantage using photons, *Science* **370**, 1460 (2020).
- [3] Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund, and Marin Soljačić, Deep learning with coherent nanophotonic circuits, *Nat. Photonics* **11**, 441 (2017).
- [4] A. N. Tait, T. F. Lima, E. Zhou, A. X. Wu, M. A. Nahmias, B. J. Shastri, and P. R. Prucnal, Neuromorphic photonic networks using silicon photonic weight banks, *Sci. Rep.* **7**, 7430 (2017).
- [5] 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).
- [6] D. Marpaung, C. Roeloffzen, R. Heideman, A. Leinse, S. Sales, and J. Capmany, Integrated microwave photonics, *Laser Photon. Rev.* **7**, 506 (2013).
- [7] L. Zhuang, C. G. Roeloffzen, M. Hoekman, K.-J. Boller, and A. J. Lowery, Programmable photonic signal processor chip for radiofrequency applications, *Optica* **2**, 854 (2015).
- [8] N. C. Harris, J. Carolan, D. Bunandar, M. Prabhu, M. Hochberg, T. Baehr-Jones, M. L. Fanto, A. M. Smith, C. C. Tison, P. M. Alsing, and Dirk Englund, Linear programmable nanophotonic processors, *Optica* **5**, 1623 (2018).
- [9] M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani, Experimental Realization of Any Discrete Unitary Operator, *Phys. Rev. Lett.* **73**, 58 (1994).
- [10] 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).
- [11] N. C. Harris *et al.*, in *2020 Optical Fiber Communications Conference and Exhibition (OFC)* (IEEE, San Diego, CA, 2020), p. 1.
- [12] C. Ramey, in *2020 IEEE Hot Chips 32 Symposium (HCS)* (IEEE Computer Society, Palo Alto, CA, 2020), p. 1.
- [13] 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).
- [14] J. Mower, N. C. Harris, G. R. Steinbrecher, Y. Lahimi, and D. Englund, High-fidelity quantum state evolution in imperfect photonic integrated circuits, *Phys. Rev. A* **92**, 032322 (2015).
- [15] S. Pai, B. Bartlett, O. Solgaard, and D. A. B. Miller, Matrix Optimization on Universal Unitary Photonic Devices, *Phys. Rev. Appl.* **11**, 064044 (2019).
- [16] S. Pai, I. Williamson, T. W. Hughes, M. Minkov, O. Solgaard, S. Fan, and D. A. B. Miller, Parallel programming of an arbitrary feedforward photonic network, *IEEE J. Sel. Top. Quantum Electron.* (2020).
- [17] T. W. Hughes, M. Minkov, Y. Shi, and S. Fan, Training of photonic neural networks through in situ backpropagation and gradient measurement, *Optica* **5**, 864 (2018).
- [18] S. Bandyopadhyay, R. Hamerly, and D. Englund, Hardware error correction for programmable photonics, *Optica* **8**, 1247 (2021).

[1] J. Carolan, C. Harrold, C. Sparrow, E. Martín-López, N. J. Russell, J. W. Silverstone, P. J. Shadbolt, N. Matsuda, M.

- [19] 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 preprint [arXiv:2103.03183](https://arxiv.org/abs/2103.03183) (2021).
- [20] D. A. Miller, Self-aligning universal beam coupler, *Opt. Express* **21**, 6360 (2013).
- [21] D. A. Miller, Self-configuring universal linear optical component, *Photonics Res.* **1**, 1 (2013).
- [22] D. A. Miller, Setting up meshes of interferometers—reversed local light interference method, *Opt. Express* **25**, 29233 (2017).
- [23] S. Grillanda, M. Carminati, F. Morichetti, P. Ciccarella, A. Annoni, G. Ferrari, M. Strain, M. Sorel, M. Sampietro, and A. Melloni, Non-invasive monitoring and control in silicon photonics using CMOS integrated electronics, *Optica* **1**, 129 (2014).
- [24] R. Hamerly, S. Bandyopadhyay, and D. Englund, Stability of Self-Configuring Large Multiport Interferometers, *Phys. Rev. Appl.* **18**, 024018 (2022).
- [25] M. Y.-S. Fang, S. Manipatruni, C. Wierzynski, A. Khosrowshahi, and M. R. DeWeese, Design of optical neural networks with component imprecisions, *Opt. Express* **27**, 14009 (2019).
- [26] R. Hamerly, Meshes: tools for modeling photonic beam-splitter mesh networks, Online at: <https://github.com/QPG-MIT/meshes> (2021).
- [27] A. Haar, Der massbegriff in der theorie der kontinuierlichen gruppen, *Ann. Math.* **34**, 147 (1933).
- [28] W.-K. Tung, *Group Theory in Physics: An Introduction to Symmetry Principles, Group Representations, and Special Functions in Classical and Quantum Physics* (World Scientific Publishing Company, Singapore, 1985).
- [29] N. J. Russell, L. Chakhmakchyan, J. L. O'Brien, and A. Laing, Direct dialling of Haar random unitary matrices, *New J. Phys.* **19**, 033007 (2017).
- [30] L. Bernstein, A. Sludds, R. Hamerly, V. Sze, J. Emer, and D. Englund, Freely scalable and reconfigurable optical hardware for deep learning, *Sci. Rep.* **11**, 1 (2021).
- [31] I. A. Williamson, T. W. Hughes, M. Minkov, B. Bartlett, S. Pai, and S. Fan, Reprogrammable electro-optic nonlinear activation functions for optical neural networks, *IEEE J. Sel. Top. Quantum Electron.* **26**, 1 (2019).
- [32] H. Zhang, M. Gu, X. Jiang, J. Thompson, H. Cai, S. Paesani, R. Santagati, A. Laing, Y. Zhang, and M. Yung *et al.*, An optical neural chip for implementing complex-valued neural network, *Nat. Commun.* **12**, 1 (2021).
- [33] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* **86**, 2278 (1998).
- [34] S. Pai, Neurophox: A simulation framework for unitary neural networks and photonic devices, Online at: <https://github.com/solgaardlab/neurophox> (2020).
- [35] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevApplied.18.024019> for neural network models and source code.
- [36] J. C. Mikkelsen, W. D. Sacher, and J. K. Poon, Dimensional variation tolerant silicon-on-insulator directional couplers, *Opt. Express* **22**, 3145 (2014).
- [37] J. Lacey and F. Payne, Radiation loss from planar waveguides with random wall imperfections, *IEE Proc. J-Optoelectron.* **137**, 282 (1990).
- [38] Q. Wilmart, S. Brisson, J.-M. Hartmann, A. Myko, K. Ribaud, C. Petit-Etienne, L. Youssef, D. Fowler, B. Charbonnier, C. Sciancalepore, Erwine Pargon, Stephane Bernabe, and Bertrand Szlag, A complete Si photonics platform embedding ultra-low loss waveguides for O- and C-band, *J. Lightwave Technol.* **39**, 532 (2020).
- [39] X. Chen, M. M. Milosevic, D. J. Thomson, A. Z. Khokhar, Y. Franz, A. F. Runge, S. Mailis, A. C. Peacock, and G. T. Reed, Post-fabrication phase trimming of Mach-Zehnder interferometers by laser annealing of germanium implanted waveguides, *Photonics Res.* **5**, 578 (2017).
- [40] D. A. Miller, Perfect optics with imperfect components, *Optica* **2**, 747 (2015).
- [41] X. Li, N. Youngblood, C. Ríos, Z. Cheng, C. D. Wright, W. H. Pernice, and H. Bhaskaran, Fast and reliable storage using a 5 bit, nonvolatile photonic memory cell, *Optica* **6**, 1 (2019).
- [42] D. M. Gill, S. S. Patel, M. Rasras, K.-Y. Tu, A. E. White, Y.-K. Chen, A. Pomerene, D. Carothers, R. L. Kamocsai, and C. M. Hill *et al.*, CMOS-compatible Si-ring-assisted Mach-Zehnder interferometer with internal bandwidth equalization, *IEEE. J. Sel. Top. Quantum Electron.* **16**, 45 (2009).
- [43] N. C. Harris, Y. Ma, J. Mower, T. Baehr-Jones, D. Englund, M. Hochberg, and C. Galland, Efficient, compact and low loss thermo-optic phase shifter in silicon, *Opt. Express* **22**, 10487 (2014).
- [44] V. J. López-Pastor, J. S. Lundeen, and F. Marquardt, Arbitrary optical wave evolution with Fourier transforms and phase masks, arXiv preprint [arXiv:1912.04721](https://arxiv.org/abs/1912.04721) (2019).
- [45] S. Fldzhyan, M. Y. Saygin, and S. Kulik, Optimal design of error-tolerant reprogrammable multiport interferometers, *Opt. Lett.* **45**, 2632 (2020).
- [46] M. Y. Saygin, I. V. Kondratyev, I. V. Dyakonov, S. A. Mironov, S. S. Straupe, and S. P. Kulik, Robust Architecture for Programmable Universal Unitaries, *Phys. Rev. Lett.* **124**, 010501 (2020).
- [47] R. Tanomura, R. Tang, S. Ghosh, T. Tanemura, and Y. Nakano, Robust integrated optical unitary converter using multiport directional couplers, *J. Lightwave Technol.* **38**, 60 (2019).
- [48] J. Polcari, in *Advanced Research in Naval Engineering* (Springer, Cham, Switzerland, 2018), p. 35.
- [49] M. Arjovsky, A. Shah, and Y. Bengio, in *International conference on machine learning* (PMLR, New York, NY, 2016), p. 1120.
- [50] S. Suda, H. Matsuura, K. Tanizawa, K. Suzuki, K. Ikeda, H. Kawashima, and S. Namiki, in *Photonics in Switching* (Optica Publishing Group, New Orleans, LA, 2017), p. PTu3C-5.
- [51] C. Alexiev, J. C. Mak, W. D. Sacher, and J. K. Poon, Calibrating rectangular interferometer meshes with external photodetectors, *OSA Continuum* **4**, 2892 (2021).
- [52] 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 (2018).
- [53] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu, A limited memory algorithm for bound constrained optimization, *SIAM J. Sci. Comput.* **16**, 1190 (1995).