

Phase-Injected Topology Optimization for Scalable and Interferometrically Robust Photonic Integrated Circuits

Alec M. Hammond¹, Joel Slaby¹, Michael Probst¹, and Stephen E. Ralph^{1*}

¹School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30308, USA

*stephen.ralph@ece.gatech.edu

We present a novel device-design methodology that produces interferometrically robust photonic integrated circuits on commercial foundry platforms. This new approach, which we call phase-injected topology optimization, leverages reciprocity to both reduce the computational load and condition the corresponding optimization problem and is compatible with existing adjoint-based design frameworks. We experimentally validate our methodology by designing, fabricating, and testing a 90° optical hybrid that operates over C-band and occupies just $8\ \mu\text{m} \times 8\ \mu\text{m}$. We measured our device across three different wafers, and observe $< \pm 7^\circ$ of random variability, and a net -10° systematic phase deviation between simulation and experiment across all four output arms.

I. INTRODUCTION

Large-scale photonics integration provides a promising platform for high-density, high-bandwidth interferometric systems on a chip which are essential for next-generation telecommunications [1], neuromorphic computing [2–6], and quantum information processing [7]. Indeed, the apogee of photonic design lies in achieving integration densities that scale on par with integrated electronics, incorporating vast networks of thousands or even millions of photonic devices. Despite significant advances in commercial foundry processing [8], however, two fundamental challenges continue to impede this effort: (1) fabrication variability (as well as environmental changes) often renders large interferometric systems inoperable without the use of dynamic tuners dispersed throughout the chip [9], a costly compromise not amenable to scaling; (2) optimized integrated photonic devices are difficult to design due to the enormous number of degrees of freedom available to photonic designers, often requiring sophisticated “inverse-design” algorithms that, while effective, are typically limited by fundamental tradeoffs between design dimensionality, device footprint, functional complexity, computational cost, and realizable performance [10].

To overcome these challenges, we present and experimentally validate a novel *phased-injected* topology optimization (TO) paradigm uniquely capable of designing interferometrically stable integrated photonic devices robust to various forms of manufacturing variability and operating conditions. We show that for many interferometric devices, reciprocity enables alternative problem formulations that reduce the corresponding computational complexity required by the device optimization (Sec. II). As a result, we can efficiently design fundamental building blocks with reliable, arbitrary phase and amplitude transformations that readily scale to fully passive interferometric networks on chip. To validate our approach, we designed and fabricated a broadband and compact 90° hybrid on the GlobalFoundries 45CLO silicon photonics commercial foundry process. We mea-

sured the devices fabricated on three separate wafers and identified tight phase tolerances across all four output arms (Sec. III). Specifically, we measured $< \pm 7^\circ$ of random variability, and a net -10° systematic phase deviation between simulation and experiment across all four output arms.

While phase-injected TO relies on recent advances in photonic adjoint-variable algorithms and software tools [11–14], this work focuses on the *problem formulation* itself, using first principles (e.g. reciprocity and superposition) to identify ideal figures of merit (FOM) and optimization trajectories for interferometrically sensitive devices. The vast design space ensures that many possible designs (and consequently, problem formulations) exist, but not all formulations yield equal (or even comparable) device performance or computational complexity. Properly formulating a particular optimization problem (e.g. how the designer’s goals and constraints are mathematically composed) is often more computationally impactful than the underlying adjoint tooling or optimization algorithm. To this end, we anticipate that phase-injected TO will enable the design of large-scale photonic transformers, digital logic circuits, and holographic networks on commercial foundry platforms.

II. PHASE-INJECTED TOPOLOGY OPTIMIZATION

We now describe phase-injected TO within the context of a frequency-domain, density-based TO framework, although our work generalizes to any device-design approach (e.g. level-set optimization) that iteratively refines the device’s geometry using adjoint gradients. Density-based TO parameterizes the design at each simulation voxel, evolving the corresponding material until some user-specified FOM (e.g. transmitted and reflected power) is satisfied [15]. When paired with an adjoint-variable method (AVM) [16, 17], which efficiently computes gradients of the FOM w.r.t. all the above-mentioned degrees of freedom (DOF) using just two Maxwell solves, standard nonlinear gradient-based op-




Figure 1. Comparison of different optimization problem formulations, each of which is capable of designing a phase splitter. A straightforward approach involves minimizing the error-vector magnitude (EVM) for both output arms, where the corresponding error vector is described by the difference in simulated and ideal (complex) scattering parameters (a). This approach requires two distinct FOMs (whose evolution is depicted in the inset) and also constrains the optimizer, requiring it to descend toward a unique target value in the complex plane for each arm. To alleviate this restriction, we recognize that only the relative phase between each arm is relevant in this context, and that a ratio of the output (complex) scattering parameters captures both the desired magnitude and phase profile. To this end we define a new EVM, which compares the simulated splitting ratio to the ideal splitting ratio, and the accompanying insertion loss (IL), (b). Both of these techniques establish a traditional adjoint problem where the phase splitter operates in a “forward configuration” (d). To further simplify the optimization problem, we leverage reciprocity and formulate the problem in reverse such that the desired magnitude and phase profile are injected as *inputs* (e). The corresponding optimization problem results in a simple minimization of the insertion loss, a purely real quantity (c).

timization algorithms can readily traverse the enormous cost space. While the flexibility and freedom offered by density-based TO often produce compact devices that yield unprecedented performance [10], ensuring the final device is fully manufacturable [18–20] and robust [21–24] requires sophisticated optimization constraints.

Conventional AVM-design frameworks establish a “forward problem”, which evaluates the FOM of the device, and an “adjoint problem”, which is needed to compute the FOM gradients w.r.t the DOF. The forward problem setup consists of a “design region” where the geometry is allowed to evolve, current sources (e.g. planewave sources and eigenmode sources) which excite the structure, and one or more FOM regions (e.g. the mode overlap domains) that record the response of the sources’ interaction with physical geometry. To design a simple integrated photonic splitter, one might use an eigenmode source to excite the input waveguide of the splitter, and perform eigenmode overlap integrals with the steady-state response of each output. The corresponding adjoint problem places new “adjoint” current sources at the FOM regions, which are essentially the derivative of the FOM with respect to the electromagnetic fields from the forward run [10, 25].

While the general procedure is rather straightforward,

designing real devices and subsystems requires extensive experience and designer intuition, as several different FOMs can describe the same performance metrics for a given problem but exhibit significantly different cost landscapes due to the overall conditioning of the problem. As such, certain devices are better suited with particular problem formulations than others. The primary innovation behind phase-injected TO lies in strategically considering the role of the forward and adjoint current sources when formulating the optimization problem itself.

To better illustrate this subtle nuance, we consider an example problem where one seeks to design a simple splitter with an asymmetric (i.e. arbitrary) magnitude and phase profile between its two output arms. To achieve the proper device characteristics, we can formulate various FOMs that each favor different aspects of the device performance in exchange for varying levels of computational complexity or optimization “stiffness” (where the Hessian matrix of the corresponding cost function is ill-conditioned). For example, one could define two separate objective functions (one for each arm), seeking to minimize the error of the target complex-valued scattering parameters [1]. Alternatively, we can formulate the problem in *reverse*, by injecting into the output ports

the proper phase profile we desire and simply minimizing the insertion loss (IL) of the device itself. This approach, inspired by an experimental technique used to characterize and tune interferometric devices [26], minimizes a physical quantity (the insertion loss) rather than an abstract error-vector magnitude (EVM). While the corresponding FOM for this forward problem does not yield any explicit performance metrics regarding the device’s magnitude and phase response, the minimized IL implicitly ensures (due to reciprocity) that when operating in the conventional forward configuration, the splitter will exhibit the desired phase profile. Fig. 1 compares each of these optimization problem formulations and their respective tradeoffs. More detail regarding the explicit FOM derivations and corresponding optimization problems is provided in Supp. Mat. SI.

Whereas conventional design formulations rely on the adjoint sources to convey any knowledge of the phase-sensitive FOM, phase-injected TO encodes this information directly within the forward sources. As such, this approach is most advantageous when designing devices that exhibit an asymmetric transfer matrix (where the number of inputs does not match the number of outputs) and only require a relative phase relationship between the (conventional) output ports. Fig. 2 illustrates the evolution of a 50/50 phase splitter with a 180° phase offset between the output arms, designed using phase-injected TO. The final structure is fully binary, operates across 100 nm of optical bandwidth, and adheres to common silicon photonic fabrication constraints [18]. Importantly, the device was designed to be robust to over- and under-etching.

Ensuring that the final device is both broadband and robust to fabrication variability is essential for large-scale interferometric operation. To overcome this, designers can perform a “worst-case” optimization over multiple “design fields” which describe various random or systematic perturbations [24, 27]. For example, by incorporating and over- and under-etching tolerance into the design space, we demonstrated robust silicon photonic splitters exhibiting strongly asymmetric splitting ratios [21]. Accounting for *all* possible sources of variability (either due to manufacturing or operating conditions) leads to a combinatorial explosion which is infeasible to implement even on large compute clusters. Incorporating one particular form of device robustness into the design process, however, often reduces sensitivity to *other* forms of manufacturing or operating variability, such that only a few worst-case perturbations must be considered. For example, broadband optimization itself prevents any high-Q resonant behavior, which is inherently sensitive to even minor fabrication variability. Similarly, designing devices resilient to changes in the underlying material models (i.e. refractive index) also ensures that the final device will perform as expected in spite of material imperfections (due to the fabrication process) or temperature variations (due to operating conditions).

Phase-injected TO lends itself well toward incorporat-




Figure 2. Design evolution (a) and optimization evolution (b) of a 180-degree, 50/50 splitter designed using phase-injected TO. Light enters into a $3\ \mu\text{m} \times 3\ \mu\text{m}$ design region from two arms 180 degrees out of phase before interfering and recombining into the single output. The final design is shown with the steady-state field profiles operating in *forward* mode, demonstrating that the device works as expected regardless of the configuration. The binarization constant (β) is held constant for 70 iterations (1 epoch), after which the binarization factor increases and the optimization algorithm restarts, as indicated by significant drops in performance at iteration 70 and iteration 140. Design rule constraints are activated during the final optimization epoch (iteration 140). Once the fabrication constraints are fully activated, the optimization problem stiffens and is unable to achieve the same level of convergence seen within the previous epochs. Nonetheless, the final splitter exhibits just $< \pm 7^\circ$ of phase error across 100 nm of bandwidth, and is robust to over- and under-etching perturbations of ± 20 nm (see Supp. Mat. SI for additional optimization parameters and robust performance).

ing these design considerations into the problem formulation. For example, broadband optimization is easily achievable using a hybrid time-/frequency-domain AVM that timesteps Maxwell’s equations in the time domain, while computing the respective adjoint and gradient information in the frequency domain using a discrete-time Fourier transform [11]. This technique can compute gradients of the phase-injected FOM at an arbitrary number of wavelength points with no additional cost, provided the specified wavelength points are far enough apart to




Figure 3. Steady-state response (a) of the final 90° optical hybrid (b) at $\theta = 0^\circ, 180^\circ, 90^\circ$, and 270° . The hybrid mixes the two input signals with the four quadrature phase states in order to decode the relative phase offset between them.

satisfy the Fourier uncertainty principle on a reasonable time-scale (determined by the designer). Similarly, “in-the-loop” software packages exist, allowing the desire to run multiple forward and Maxwell solves for different design configurations in parallel [11]. In other words, the cost of additional design-field gradients, which are typically optimized in a minimax sense [28], is linear with each additional compute node the designer is able to allocate. Furthermore, the computational efficiency of the framework allows designers to design many such devices in tandem using the same optimization framework. Additional detail provided in Supp. Mat. SI describes the broadband, robust design of multiple phase splitters with different output phase profiles.

III. DEVICE DESIGN AND EXPERIMENTAL RESULTS

The 90° optical hybrid is often used to decode the relative phase difference between two input signals, and is a fundamental building block for various classical and quantum communication architectures [29]. Fundamen-

tally, the two inputs are interfered at the four quadrature phase states: $0^\circ, 180^\circ, 90^\circ$, and 270° . Balanced detectors measure the relative output power of each arm and recreate the in-phase and quadrature (IQ) components, thereby characterizing the inputs’ relative phase difference. Integrated 90° optical hybrids are either all-passive [29–31], or require dynamic tuning to ensure a stable and consistent phase response [32]. In both cases, the conventional design methodology is to first identify various constituent components, like multi-mode interferometers (MMI), directional couplers, and y-branch power splitters, and then cascade these individual components until the desired phase transformation is achieved. As such, a compact and comprehensive device capable of achieving the target transformation without any on-chip tuning or additional components is of great interest.

While there are many ways to formulate an optimization problem that designs an integrated 90° optical hybrid, the corresponding asymmetric transfer function lends itself well to phase-injected topology optimization. For example, one could naively formulate an independent FOM for each input-output relationship and ensure that both the final phase and magnitude response behave as intended. The resulting optimization problem, however, requires the simultaneous optimization of eight independent FOMs that are heavily constrained by the need to satisfy a magnitude *and* phase requirement. In contrast, we drive the hybrid’s output arms with a forward source described by the proper magnitude and phase profile. This alternate approach demonstrates two superior advantages to the naive formulation: (1) the total number of FOMs is reduced from eight to just two, one for each input arm; (2) each FOM must simply maximize the *power* (not the magnitude and phase) in each input arm, significantly simplifying the design problem.

Using the above formulation, we designed a 3D, $8\ \mu\text{m} \times 8\ \mu\text{m}$ 90° optical hybrid that operates over the optical C-band. We ensured the final device was compatible with thousands of foundry-specified design rule checks (DRC) by placing minimum linewidth, linespacing, curvature, area, and enclosed area constraints during the last epoch of the optimization [18]. Fig. 3 illustrates the final 3D design, along with the steady-state field performance for each of the four quadrature states.

The devices were fabricated across three different wafers using the GlobalFoundries 45CLO CMOS silicon photonics process [33, 34]. To characterize the final phase response of the device, we used an interferometric structure on chip. Fig. 4 describes the fabricated test structure and a comparison of the ideal, simulated and measured performance for each output arm. We measured $< \pm 7^\circ$ of random variability, and a net 10° systematic phase deviation between simulation and experiment across all four output arms. The experimental transfer function yields a smooth and continuous response from input to output for all measured wavelengths, even though the device was only optimized at five discrete wavelength points (Sec. V C).




Figure 4. Experimental results for the 90° optical hybrid. An interferometric on-chip test structure (a) was used to characterize the broadband performance of each hybrid output arm. The simulated (orange) and measured (green) transfer function were used to determine the phase error as a function of input angle (b). By using the 0° port as a reference, the broadband phase performance for the remaining three ports is compared to the simulation (c-e). Best- and worst-case measurements across all three wafers are depicting by the light green shaded region. A colored inset of the hybrid’s geometry describes which port each measurement corresponds to.

Two ports exhibit almost no systematic error (0° and 270°), whereas the other two ports exhibit a notable yet similar systematic shift of 15-20°. The root cause of the systematic error is likely due differences in simulated and realized layer thickness and refractive index models. In future work, we plan to design multilayer structures leveraging both the standard silicon-on-insulator, along with the poly-silicon layer, to provide additional wave action needed to minimize insertion loss and mitigate the broadband phase errors. Higher resolution forward and adjoint simulations, coupled with higher-order subpixel smoothing routines and more sophisticated material models [11] should result in a closer match between simulation and experiment.

IV. CONCLUSION

We present a new photonics inverse design methodology well-suited for compact, interferometric devices robust to fabrication variability. In particular, we highlight the importance of considering the forward and adjoint current sources when formulating the optimization problem, and show that for many devices, a proper choice yields an optimization problem requiring significantly fewer Maxwell solves yielding devices robust to variability. We validate our approach by designing and fabricating a compact 90° optical hybrid on a commercial

foundry platform and characterize its transfer function across three wafers. The device was explicitly designed for broadband operation, which additionally reduced its sensitivity to fabrication variability. Unlike other optical hybrids reported in the literature, our design is all-passive, comprised of a single component, and occupies a footprint orders of magnitude smaller than the state of the art. We foresee this methodology as a standard design tool for producing compact, phase transformation blocks needed to create very-large scale interferometric transformers.

V. METHODS

A. Device simulation

All devices were simulated using the free, open-source FDTD software package Meep [35]. Gradients for multiple wavelength points were computed in parallel using our hybrid time-/frequency-domain adjoint solver [11]. We used a free/open-source implementation [36] of the globally convergent method of moving asymptotes (MMA) [37] optimization algorithm to design all devices. We ran all simulations and optimizations on various Intel Xeon Gold 6226 2.7 GHz CPU nodes using resources provided by the Partnership for an Advanced

Computing Environment (PACE) at the Georgia Institute of Technology [38].

B. Asymmetric splitter design

We designed four different (2D) 50/50 splitters exhibiting unique phase offsets: $\Delta\phi = 0^\circ, 45^\circ, 90^\circ$, and 180° (Fig. S1). The simulations emulated a fictitious foundry with a minimum feature size of 80 nm and dispersionless Si and SiO₂ refractive index models for the solid and void regions respectively [18]. The splitters' design region was $3\ \mu\text{m} \times 3\ \mu\text{m}$, with a simulation resolution of 30 pixels/ μm and a design grid resolution of 60 pixels/ μm , such that 32,400 DOF were used throughout the design process (not including the auxiliary parameter, t). The input and output waveguides were $0.5\ \mu\text{m}$ wide, and the output waveguides were separated by $1\ \mu\text{m}$ edge-to-edge.

We ran each optimization for 70 iterations over 3 epochs (210 iterations total) using 12 cores of a compute node, such that each optimization terminated after about 2 hours. Each epoch was defined by a binarization constant such that $\beta = 8, 16, 32$ respectively, and was restarted using the final design of the previous epoch as an initial condition. The design was uniformly initialized ($\rho = 0.5$) for the first epoch. We used a conic filter radius of 80 nm throughout the optimization process. Minimum linewidth and linespacing constraints were activated during the last optimization epoch, using $\eta_e = 0.75$, $\eta_d = 0.25$, and $c = 8500$.

The optimization was performed over 10 discrete wavelength points from 1500 nm to 1600 nm. To further improve the robustness of each device, the optimization was additionally performed over three design field variants (i.e. eroded, ideal, and dilated). We used the harmonic filters to produce the eroded and dilated design fields with $\alpha = 10^{-4}$ and a filter support defined by a circle with a radius of 20 nm, which accounted for an over-/under-etch tolerance of 20 nm uniformly (Supp. Mat. SI, Fig. S2).

C. Optical hybrid design

We designed a 3D, 90° optical hybrid for the GlobalFoundries 45CLO commercial foundry silicon photonics process [33]. The design region was $8\ \mu\text{m} \times 8\ \mu\text{m}$, with a simulation resolution of 30 pixels/ μm and a design grid resolution of 60 pixels/ μm , such that 230,400 DOF were used throughout the design process (not including the auxiliary parameter, t). All waveguides are single-mode with a width of $0.5\ \mu\text{m}$. The two input waveguides feed into the hybrid from the same side. The in-phase arms exit from one side of the hybrid, and the quadrature arms exit from the other side such that no waveguide crossings are required. All input and output waveguide pairs are $0.5\ \mu\text{m}$ apart, and are routed away from the structure using standard S-bends. The

design was only parameterized along a 2D cross-section and extruded such that it spanned the thickness of the silicon-on-insulator layer. Using 8 nodes (each containing 24 cores), our hybrid time-/frequency-domain Maxwell adjoint solver converged for each forward and adjoint run after 15 minutes on average. Two FOMs describing the hybrid's transfer function were optimized in parallel (4 compute nodes/FOM). Unlike the phase splitters, only the nominal design field was simulated each design iteration (i.e. no over-/under-etching).

The optimization was performed over 5 discrete wavelength points from 1530 nm to 1565 nm (spanning C-band) and consisted of 70 iterations for 3 epochs, followed by an additional 100 iterations during the final epoch (310 iterations total) to enable better performance once the fabrication constraints were activated. Each epoch was defined by a binarization constant, such that $\beta = 8, 16, 32$ respectively. The design was uniformly initialized ($\rho = 0.5$). We used a conic filter radius of 90 nm throughout the optimization process. Minimum linewidth and linespacing constraints were activated during the last optimization epoch, using $\eta_e = 0.7$, $\eta_d = 0.25$, and $c = 8500$. The minimum area and minimum enclosed-area constraints were also activated during the last epoch.

D. Device fabrication

The optical hybrid was fabricated using the GlobalFoundries 45CLO CMOS platform [33, 34]. The process monolithically integrates electronics and photonics on 300 mm SOI wafers. The SOI layer (RX) is 160 nm thick and the buried-oxide layer (BOX) is $2\ \mu\text{m}$ thick. The optical hybrid and corresponding test structures were fabricated on three separate wafers.

E. Device measurement & characterization

To properly measure the phase relationship of the optical hybrid, we designed a mismatched Mach-Zehnder interferometer (MZI) using a standard directional coupler as the input and the inverse-designed optical hybrid as the output. The MZI's path length mismatch ($\Delta L = 115\ \mu\text{m}$) was chosen to ensure enough interference fringes spanned the band of interest without spanning below the resolution of the tunable laser source (Keysight 81642A). The four output arms of the optical hybrid fed into individual grating couplers and a single fiber array (OZ Optics) was used to measure all outputs simultaneously. The relative difference in peak transmission for each channel directly corresponds to the phase offset introduced by the optical hybrid (i.e. its transfer function). We used a Hilbert transform to determine the relative phase between output arms induced by the optical hybrid as a function of wavelength (Fig. S3).

DISCLOSURES

The authors declare no conflicts of interest.

FUNDING

AMH was supported by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate (NDSEG) Fellowship Program. JS, MP, and SER were supported by the Georgia Electronic Design Center of the Georgia Institute of Technology. This material is based upon work supported in part by the National Science Foundation (NSF) Center “EPICA” under Grant No. 2052808, <https://epica.research.gatech.edu/>.

ACKNOWLEDGMENTS

The authors would like to thank Steven G. Johnson, Ardavan Oskooi, Yusheng Bian, Michal Rakowski, and Roderick A. Augur for useful discussions. The authors would also like to thank GlobalFoundries for providing silicon fabrication through the 45CLO university program. This research was supported in part through research cyberinfrastructure resources and services provided by the Partnership for an Advanced Computing Environment (PACE) at the Georgia Institute of Technology, Atlanta, Georgia, USA.

-
- [1] L. Chrostowski and M. Hochberg, *Silicon photonics design: from devices to systems*. Cambridge University, 2015.
 - [2] D. Pérez, I. Gasulla, L. Crudgington, D. J. Thomson, A. Z. Khokhar, K. Li, W. Cao, G. Z. Mashanovich, and J. Capmany, “Multipurpose silicon photonics signal processor core,” *Nature communications*, vol. 8, no. 1, pp. 1–9, 2017.
 - [3] W. Bogaerts, D. Pérez, J. Capmany, D. A. Miller, J. Poon, D. Englund, F. Morichetti, and A. Melloni, “Programmable photonic circuits,” *Nature*, vol. 586, no. 7828, pp. 207–216, 2020.
 - [4] 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 *et al.*, “Linear programmable nanophotonic processors,” *Optica*, vol. 5, no. 12, pp. 1623–1631, 2018.
 - [5] 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,” *Nature Photonics*, vol. 11, no. 7, pp. 441–446, 2017.
 - [6] A. N. Tait, T. F. De Lima, E. Zhou, A. X. Wu, M. A. Nahmias, B. J. Shastri, and P. R. Prucnal, “Neuromorphic photonic networks using silicon photonic weight banks,” *Scientific reports*, vol. 7, no. 1, pp. 1–10, 2017.
 - [7] A. Deshpande, A. Mehta, T. Vincent, N. Quesada, M. Hinsche, M. Ioannou, L. Madsen, J. Lavoie, H. Qi, J. Eisert *et al.*, “Quantum computational advantage via high-dimensional gaussian boson sampling,” *Science advances*, vol. 8, no. 1, p. eabi7894, 2022.
 - [8] M. Liehr, M. Baier, G. Hoeffler, N. M. Fahrenkopf, J. Bowers, R. Gladhill, P. O’Brien, E. Timurdogan, Z. Su, and F. Kish, “Foundry capabilities for photonic integrated circuits,” in *Optical Fiber Telecommunications VII*. Elsevier, 2020, pp. 143–193.
 - [9] W. Bogaerts and L. Chrostowski, “Silicon photonics circuit design: methods, tools and challenges,” *Laser & Photonics Reviews*, vol. 12, no. 4, p. 1700237, 2018.
 - [10] S. Molesky, Z. Lin, A. Y. Piggott, W. Jin, J. Vucković, and A. W. Rodriguez, “Inverse design in nanophotonics,” *Nature Photonics*, vol. 12, no. 11, pp. 659–670, Oct. 2018.
 - [11] A. M. Hammond, A. Oskooi, M. Chen, Z. Lin, S. G. Johnson, and S. E. Ralph, “High-performance hybrid time/frequency-domain topology optimization for large-scale photonics inverse design,” *Optics Express*, vol. 30, no. 3, pp. 4467–4491, 2022.
 - [12] T. W. Hughes, I. A. Williamson, M. Minkov, and S. Fan, “Forward-mode differentiation of maxwell’s equations,” *ACS Photonics*, vol. 6, no. 11, pp. 3010–3016, 2019.
 - [13] L. Su, D. Vercruysse, J. Skarda, N. V. Sapra, J. A. Petykiewicz, and J. Vucković, “Nanophotonic inverse design with SPINS: Software architecture and practical considerations,” *Applied Physics Reviews*, vol. 7, no. 1, p. 011407, 2020.
 - [14] A. Michaels and E. Yablonovitch, “Leveraging continuous material averaging for inverse electromagnetic design,” *Optics express*, vol. 26, no. 24, pp. 31 717–31 737, 2018.
 - [15] O. Sigmund and K. Maute, “Topology optimization approaches,” *Structural and Multidisciplinary Optimization*, vol. 48, no. 6, pp. 1031–1055, 2013.
 - [16] G. Strang, *Computational Science and Engineering*. Wellesley, MA: Wellesley-Cambridge Press, 2007.
 - [17] S. G. Johnson, “Notes on adjoint methods for 18.335,” 2012.
 - [18] A. M. Hammond, A. Oskooi, S. G. Johnson, and S. E. Ralph, “Photonic topology optimization with semiconductor-foundry design-rule constraints,” *Optics Express*, vol. 29, no. 15, pp. 23 916–23 938, 2021.
 - [19] M. Zhou, B. S. Lazarov, F. Wang, and O. Sigmund, “Minimum length scale in topology optimization by geometric constraints,” *Computer Methods in Applied Mechanics and Engineering*, vol. 293, pp. 266–282, 2015.
 - [20] D. Vercruysse, N. V. Sapra, L. Su, R. Trivedi, and J. Vucković, “Analytical level set fabrication constraints for inverse design,” *Scientific reports*, vol. 9, no. 1, pp. 1–7, 2019.
 - [21] A. M. Hammond, J. Slaby, G. Saha, and S. E. Ralph, “Robust topology optimization for foundry-photonics inverse design: Examining compact and arbitrary power splitters,” in *2021 European Conference on Optical Communication (ECOC)*. IEEE, 2021, pp. 1–4.
 - [22] R. E. Christiansen, B. S. Lazarov, J. S. Jensen, and O. Sigmund, “Creating geometrically robust designs for

- highly sensitive problems using topology optimization,” *Structural and Multidisciplinary Optimization*, vol. 52, no. 4, pp. 737–754, 2015.
- [23] E. W. Wang, D. Sell, T. Phan, and J. A. Fan, “Robust design of topology-optimized metasurfaces,” *Optical Materials Express*, vol. 9, no. 2, pp. 469–482, 2019.
- [24] X. Qian and O. Sigmund, “Topological design of electromechanical actuators with robustness toward over- and under-etching,” *Computer Methods in Applied Mechanics and Engineering*, vol. 253, pp. 237–251, 2013.
- [25] O. D. Miller, “Photonic design: From fundamental solar cell physics to computational inverse design,” *arXiv preprint arXiv:1308.0212*, 2013.
- [26] A. M. Hammond, I. W. Frank, and R. M. Camacho, “Error correction in structured optical receivers,” *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 24, no. 6, pp. 1–8, 2018.
- [27] F. Wang, B. S. Lazarov, and O. Sigmund, “On projection methods, convergence and robust formulations in topology optimization,” *Structural and Multidisciplinary Optimization*, vol. 43, no. 6, pp. 767–784, 2011.
- [28] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University, 2004.
- [29] H. Guan, Y. Ma, R. Shi, X. Zhu, R. Younce, Y. Chen, J. Roman, N. Ophir, Y. Liu, R. Ding *et al.*, “Compact and low loss 90 optical hybrid on a silicon-on-insulator platform,” *Optics Express*, vol. 25, no. 23, pp. 28 957–28 968, 2017.
- [30] W. Yang, M. Yin, Y. Li, X. Wang, and Z. Wang, “Ultra-compact optical 90 hybrid based on a wedge-shaped 2×4 mm coupler and a 2×2 mm coupler in silicon-on-insulator,” *Optics express*, vol. 21, no. 23, pp. 28 423–28 431, 2013.
- [31] R. Halir, G. Roelkens, A. Ortega-Moñux, and J. Wangüemert-Pérez, “High-performance 90 hybrid based on a silicon-on-insulator multimode interference coupler,” *Optics letters*, vol. 36, no. 2, pp. 178–180, 2011.
- [32] K. N. Nguyen, P. J. Skahan, J. M. Garcia, E. Lively, H. N. Poulsen, D. M. Baney, and D. J. Blumenthal, “Monolithically integrated dual-quadrature receiver on inp with 30 nm tunable local oscillator,” *Optics Express*, vol. 19, no. 26, pp. B716–B721, 2011.
- [33] K. Giewont, K. Nummy, F. A. Anderson, J. Ayala, T. Barwicz, Y. Bian, K. K. Dezfulian, D. M. Gill, T. Houghton, S. Hu, B. Peng, M. Rakowski, S. Rauch, J. C. Rosenberg, A. Sahin, I. Stobert, and A. Stricker, “300-mm monolithic silicon photonics foundry technology,” *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 25, no. 5, pp. 1–11, 2019.
- [34] M. Rakowski, C. Meagher, K. Nummy, A. Aboketaf, J. Ayala, Y. Bian, B. Harris, K. Mclean, K. McStay, A. Sahin *et al.*, “45nm cmos-silicon photonics monolithic technology (45clo) for next-generation, low power and high speed optical interconnects,” in *Optical Fiber Communication Conference*. Optical Society of America, 2020, pp. T3H–3.
- [35] A. F. Oskooi, D. Roundy, M. Ibanescu, P. Bermel, J. D. Joannopoulos, and S. G. Johnson, “MEEP: A flexible free-software package for electromagnetic simulations by the FDTD method,” *Computer Physics Communications*, vol. 181, no. 3, pp. 687–702, 2010.
- [36] S. G. Johnson, “The NLOpt nonlinear-optimization package,” <http://github.com/stevengj/nlopt>.
- [37] K. Svanberg, “A class of globally convergent optimization methods based on conservative convex separable approximations,” *SIAM journal on optimization*, vol. 12, no. 2, pp. 555–573, 2002.
- [38] PACE, *Partnership for an Advanced Computing Environment (PACE)*, 2017. [Online]. Available: <http://www.pace.gatech.edu>
- [39] A. W. Snyder and J. Love, *Optical Waveguide Theory*. Springer Science & Business Media, 2012.
- [40] J. Schmidhuber, “Deep learning in neural networks: An overview,” *Neural Networks*, vol. 61, pp. 85–117, Jan. 2015.
- [41] K. Svanberg and H. Svärd, “Density filters for topology optimization based on the pythagorean means,” *Structural and Multidisciplinary Optimization*, vol. 48, no. 5, pp. 859–875, 2013.

**SUPPLEMENTARY INFORMATION TO
PHASE-INJECTED TOPOLOGY OPTIMIZATION FOR SCALABLE AND INTERFEROMETRICALLY
ROBUST PHOTONIC INTEGRATED CIRCUITS**

Alec M. Hammond¹, Joel Slaby¹, Michael Probst¹, and Stephen E. Ralph^{1*}

¹School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30308, USA

* stephen.ralph@ece.gatech.edu

SECTION I: COMPARISON OF OPTIMIZATION PROBLEM FORMULATIONS

Here, we mathematically establish the optimization framework used to implement phase-injected topology optimization and compare this framework to various other possible formulations. Throughout our analysis, which focuses on both the phase splitters and the 90° optical hybrid (although is completely generalizable), we seek to achieve three goals: (1) derive an optimization problem that is smoothly differentiable, allowing for off-the-shelf, gradient-based optimization algorithms; (2) ensure the resulting search space is smooth and well-conditioned, which minimizes the number of iterations needed for convergence; (3) minimize the computational cost of each optimization iteration.

Unlike many other photonics TO problems, all *phase-injected* TO problems are “minimax” (or alternatively, “maximin”) optimization problems [28] over the design variables $\boldsymbol{\rho}$ involving N distinct objective functions subject to Maxwell’s equations spanning M frequency points:

$$\begin{aligned} & \min_{\boldsymbol{\rho}} [\max_n \{f_n(\mathbf{E})\}] & n \in \{1, 2, \dots, N\} \\ \text{s.t. } & \nabla \times \frac{1}{\mu_0 \mu_r} \nabla \times \mathbf{E} - \omega_m^2 \epsilon_0 \epsilon_r(\boldsymbol{\rho}) \mathbf{E} = -i\omega_m \mathbf{J} & m \in \{1, 2, \dots, M\} \\ & 0 \leq \boldsymbol{\rho} \leq 1 \\ & g_k(\boldsymbol{\rho}) \leq 0 & k \in \{1, 2, \dots, K\} \end{aligned} \tag{S1}$$

where f_n is each objective function dependent on the spatial, steady-state field pattern $\mathbf{E}$, ϵ is the relative permittivity as a function of the density design variables at each point in space, $\mathbf{J}$ is the current density, and g_k is the k^{th} constraint function. In order to cast the above optimization problem into a differentiable form, we use an epigraph formulation [18, 37] by introducing the auxiliary parameter, $t \in \mathbb{R}$ with additional nonlinear constraints:

$$\begin{aligned} & \min_{\boldsymbol{\rho}, t} t \\ \text{s.t. } & \nabla \times \frac{1}{\mu_0 \mu_r} \nabla \times \mathbf{E} - \omega_m^2 \epsilon_0 \epsilon_r(\boldsymbol{\rho}) \mathbf{E} = -i\omega_m \mathbf{J} & m \in \{1, 2, \dots, M\} \\ & 0 \leq \boldsymbol{\rho} \leq 1 \\ & f_n(\mathbf{x}) - t \leq 0 & n \in \{1, 2, \dots, N\} \\ & g_k(\boldsymbol{\rho}) \leq 0 & k \in \{1, 2, \dots, K\} \end{aligned} \tag{S2}$$

When designing interferometrically robust devices, this approach offers an important advantage over alternative optimization formulations which aggregate multiple FOMs into a final objective function (e.g. using the L^1 or L^2 norm). Namely, a minimax framework applied to broadband optimization ensures each wavelength point of interest is optimized in tandem, forcing a “flat-band” response. The optimizer identifies *which* wavelength point exhibits the worst behavior and is able to target that wavelength’s performance during the next iteration. In contrast, a mean-square error (MSE) formulation (i.e. the L^2 norm) cannot identify or mitigate the effect of outliers, and is forced to optimize the device’s *aggregate* performance. Similarly, the minimax formulation identifies performance discrepancies across different design fields (e.g. when designing phase splitters robust to over- and under-etching, Sec. II) and/or different FOMs (e.g. when designing the 90° optical hybrids which require a different FOM for each input arm, Sec. III).

With the motivation behind a minimax framework established, the goal now is to identify the minimum number of FOMs, f_n , needed to characterize the device, as each FOM will require its own gradient computation. Furthermore, we want to ensure each FOM is properly conditioned, such that the optimizer can readily traverse the corresponding cost space. All objective functions relevant to this work involve manipulation of mode fields within the input and output waveguides (e.g. values within a scattering or transfer matrix). We define a complex-amplitude mode coefficient, $\alpha_{m,n}^{\pm}$, of the m^{th} mode for either the forward (+) or backward (−) directions in the n^{th} port of the device. The mode coefficient itself is determined by an overlap integral over a cross-section A [39],

$$\alpha_{m,n}^{\pm} = C \int_A \left[\widehat{\mathbf{E}}^*(r) \times \widehat{\mathbf{H}}_{m,n}^{\pm}(r) + \widehat{\mathbf{E}}_{m,n}^{\pm}(r) \times \widehat{\mathbf{H}}^*(r) \right] \cdot \widehat{\mathbf{n}} dA, \tag{S3}$$

where $\widehat{\mathbf{E}}(r)$ and $\widehat{\mathbf{H}}(r)$ are the Fourier-transformed monitor fields at a particular frequency, $\widehat{\mathbf{E}}_{m,n}^{\pm}(r)$ and $\widehat{\mathbf{H}}_{m,n}^{\pm}(r)$ are the mode profiles corresponding to the m^{th} mode at the n^{th} port at the same frequency for the forward- (+) and backward- (−) propagating modes computed based on the chosen waveguide geometry, and C is a normalization constant chosen such that

$$|\alpha_{m,n}^{\pm}|^2 = P_{m,n}^{\pm}, \quad (\text{S4})$$

where $P_{m,n}^{\pm}$ is the total power propagating in that particular mode.

We note that each FOM must be real-valued and must ensure that the magnitude and phase transformation of the device operates as desired. To better understand the nuances behind the choice of FOM, we first identify the desired transfer function of an arbitrary phase splitter, described by

$$T = \frac{\sqrt{2}}{2} \begin{bmatrix} a_1 e^{-j\theta_1} \\ a_2 e^{-j\theta_2} \end{bmatrix}, \quad (\text{S5})$$

where a_1 and a_2 describe the field amplitude of each output arm, and θ_1 and θ_2 describe the *absolute* phase delay. By defining the transfer function in this way, we assume the splitter operates from “left to right”, such that light is injected into the input arm, and then split accordingly to the two output arms, where the appropriate phase shift is applied.

Naively, one could define four, real-valued FOMs that capture the magnitude and phase of each element of the above transfer matrix,

$$\begin{aligned} f_0 &= \left(\left| \frac{\alpha_{1,2}^+}{\alpha_{1,1}^+} \right|^2 - 0.5 \right)^2, \\ f_1 &= \left(\left| \frac{\alpha_{1,3}^+}{\alpha_{1,1}^+} \right|^2 - 0.5 \right)^2, \\ f_2 &= \left| \angle \left(\frac{\alpha_{1,2}^+}{\alpha_{1,1}^+} \right) - \theta_1 \right|^2, \\ f_3 &= \left| \angle \left(\frac{\alpha_{1,3}^+}{\alpha_{1,1}^+} \right) - \theta_2 \right|^2, \end{aligned} \quad (\text{S6})$$

where port 1 corresponds to the input waveguide, ports 2 and 3 correspond to the output waveguides, and we evaluate the squared error to preserve a differentiable form. We note that the relative scales of f_0 and f_1 (which range between 0 and 1) are drastically different than that of f_2 and f_3 , which depend on the defined units of the complex angle, and must be appropriately normalized and scaled. Once this is properly implemented, we could potentially reduce the number of simultaneous Maxwell solves from four to one by casting an aggregated FOM, consisting of a simple linear combination of the above terms. However, this approach loses many of the advantages of a minimax optimization, as discussed earlier, since the optimization is likely to focus on simpler aspects of the FOM (e.g. power) to compensate for the lower performance of other aspects (e.g. phase).

A more elegant (and computationally efficient) formulation is to instead directly minimize the error-vector magnitude (EVM) between the simulated (complex) transmission coefficient and the desired transmission coefficient such that

$$\begin{aligned} f_0 &= \left| \frac{\alpha_{1,2}^+}{\alpha_{1,1}^+} - \frac{\sqrt{2}}{2} e^{-j\theta_1} \right|^2, \\ f_1 &= \left| \frac{\alpha_{1,3}^+}{\alpha_{1,1}^+} - \frac{\sqrt{2}}{2} e^{-j\theta_2} \right|^2. \end{aligned} \quad (\text{S7})$$

As before, these individual FOMs could be aggregated into a single FOM as a function of wavelength. Unlike our naive approach above, the EVM does not require any additional scaling or conditioning — an important improvement. However, the current formulation places an unnecessary constraint on the final device’s performance. Namely, it forces both the output arms to have a specific phase relative to the input (i.e. an absolute phase response). While this is

important when designing devices that must be cascaded together, many other devices, like the 90° optical hybrid, only require a relative phase profile between their output arms.

In these unique cases, we can further simplify the FOM by redefining the EVM as a measure of the error between the simulated complex ratio of both output arms and the target ratio, such that

$$f_0 = \left| \frac{\alpha_{1,2}^+}{\alpha_{1,3}^+} - R \right|^2, \quad (\text{S8})$$

where R is the target ratio. To ensure that the device's insertion loss (IL) is also minimized, we define an explicit constraint function (independent of the auxiliary parameter, t) of the form

$$g_0 = 1 - \left| \frac{\alpha_{1,2}^+}{\alpha_{1,1}^+} \right|^2 - \left| \frac{\alpha_{1,3}^+}{\alpha_{1,1}^+} \right|^2 \quad (\text{S9})$$

which will require its own forward and adjoint simulations to compute the corresponding gradient. The designer has some flexibility over how this constraint is implemented. For example, a particular threshold could be defined (e.g. a maximum insertion loss of 1 dB) which prevents any progress on the actual FOM until the insertion loss criteria is satisfied. In many ways, this additional threshold parameter is analogous to a scalar weighting parameter used in the augmented approaches described earlier, and subject to the same drawbacks. Alternatively, one could incorporate some dependence on the auxiliary parameter itself [18], to better condition the relationship between the IL constraint and the original FOM.

Finally, we can leverage a *phase-injected* formulation to maximally simplify the resulting optimization problem. To do so, we operate the device in reverse, which yields the transpose of the above transfer function

$$T^H = \frac{\sqrt{2}}{2} \begin{bmatrix} a_1 e^{j\theta_1} & a_2 e^{j\theta_2} \end{bmatrix}. \quad (\text{S10})$$

If we intentionally drive the (new) input arms with the desired magnitude and phase, an “optimal” device will maximize power through the (new) output arm. Thanks to reciprocity, if we minimize the device's insertion loss, the final device will operate as expected in *both* forward and backward configurations. The resulting FOM reduces to

$$f_0 = 1 - \left| \frac{\alpha_{1,2}^+}{\alpha_{1,1}^+} \right|^2 - \left| \frac{\alpha_{1,3}^+}{\alpha_{1,1}^+} \right|^2, \quad (\text{S11})$$

which is strictly positive. This last optimization formulation only requires a single FOM and a single gradient evaluation per wavelength and design field variation. Additionally, its relative definition for input phase is significantly easier to achieve than the alternative absolute definition.

We can replicate the same analysis as above using the 90° optical hybrid. By analyzing its transfer function,

$$T_{90^\circ} = \frac{1}{4} \begin{bmatrix} 1 & 1 \\ 1 & e^{-j\pi} \\ 1 & e^{-j\frac{\pi}{2}} \\ 1 & e^{j\frac{\pi}{2}} \end{bmatrix}, \quad (\text{S12})$$

and choosing to operate it in reverse, we see that the phase-injected formulation requires just two FOMs optimized in parallel (one for each column vector of T).

SECTION II: OPTIMIZATION PARAMETERS

In this section, we describe the optimization parameters used to design both the arbitrary phase splitters and the 90° optical hybrid. The parameterization approach follows that of traditional density-based TO methodologies [15], with some important modifications described in [18]. The “latent” DOF (ρ), which describe each voxel of the design

domain, are first convolved with a 2D conic filter kernel, $w(\mathbf{x})$, described by

$$w(\mathbf{x}) = \begin{cases} \frac{1}{a} \left(1 - \frac{|\mathbf{x} - \mathbf{x}_0|}{R}\right) & \mathbf{x} \in \mathcal{N} \\ 0 & \mathbf{x} \notin \mathcal{N} \end{cases}, \quad (\text{S13})$$

where $\mathcal{N}$ is a circle of radius R , $\mathbf{x}_0$ is the center of $\mathcal{N}$, and a is a normalization factor such that $\int w(\mathbf{x}) = 1$. Filtering the latent design field *regularizes* the final design, such that the resulting optimization problem is better conditioned and a smoother structure emerges [15]. In addition, the proper choice of filter kernel works in tandem with the geometric length constraints (described below) to ensure the final design is manufacturable.

The filtered DOF ($\tilde{\rho}$) are then projected to a quasi-binary value using

$$\bar{\rho} = \frac{\tanh(\beta\eta) + \tanh(\beta(\tilde{\rho} - \eta))}{\tanh(\beta\eta) + \tanh(\beta(1 - \eta))}, \quad (\text{S14})$$

where $\bar{\rho}$ is the projected design field, and β and η are the threshold parameters [27]. As β increases, the projected design becomes more binary. The norm of the resulting gradient, however, approaches 0 as $\beta \rightarrow \infty$ due to the resulting discontinuity in the projection function. As such, we use finite values of β throughout the optimization process. Furthermore, we use *multiple* values of β , restarting the optimization algorithm each time β is updated, as this corresponds to an entirely new optimization problem. This update scheme forces the resulting geometry to binarize while simultaneously pivoting the optimization algorithm out of any local minima it may be occupying. We call each of these optimization evolutions “epochs”, as they closely resemble the iteration scheme used while training artificial neural networks [40]. Since we are unable to *force* a binary design using β , we rely on the corresponding geometric constraints (described below) to ensure the final design is both binary and manufacturable.

The other threshold parameter, η , is used to erode ($\eta > 0.5$) or dilate ($\eta < 0.5$) the projected design, a useful feature when simultaneously optimizing over *multiple* design fields (e.g. for robust optimization). Unfortunately, the resulting morphological transforms heavily depend on the dynamic range of the *latent* design field, and ensuring a fixed, systematic erosion or dilation is nearly impossible.

To overcome this issue, we use the approach introduced in Ref. 18, which assumes $\eta = 0.5$ for all elements of ρ , but involves an additional nonlinear filter step *after* the projection. The harmonic erosion filter [41] is defined by

$$\mathcal{E}_{\mathcal{N}}(\bar{\rho}) = \left(\frac{1}{\bar{\rho} + \alpha} * w \right)^{-1} - \alpha, \quad (\text{S15})$$

where α is a nonlinear threshold parameter. The harmonic dilation is defined by

$$\mathcal{D}_{\mathcal{N}}(\bar{\rho}) = 1 - \left(\frac{1}{1 - \bar{\rho} + \alpha} * w \right)^{-1} - \alpha, \quad (\text{S16})$$

The filter w is a uniform kernel and its support determines the amount of erosion/dilation that occurs in a particular direction. α is a regularization parameter to ensure differentiability; the exact dilation/erosion operation is obtained in the $\alpha \rightarrow 0$ limit.

The final permittivity (ϵ) is then interpolated from the projected (and possible eroded/dilated) DOF using

$$\epsilon_r(\bar{\rho}) = \epsilon_{\min} + \bar{\rho}(\epsilon_{\max} - \epsilon_{\min}) \quad (\text{S17})$$

where $\epsilon_{\min}$ is the permittivity of the “void” region (SiO_2) and $\epsilon_{\max}$ is the permittivity of the “solid” region (Si).

To ensure manufacturability, we enforce *geometric* constraints (i.e. minimum linewidth and linespacing) and area constraints (i.e. minimum area and enclosed area). The minimum-linewidth constraint, $g_{LW} \leq 0$, is described by the function

$$g_{LW} = \frac{1}{n} \sum_{i \in \mathbb{N}} I_i^{LW}(\rho_i) \cdot [\min\{(\tilde{\rho}_i - \eta_e), 0\}]^2 \quad (\text{S18})$$

where $I_i^{LW}(\rho)$ is an indicator function that identifies the inflection region of the solid phase:

$$I_i^{LW}(\rho) = \bar{\rho} \cdot \exp\left(-c|\nabla \bar{\rho}|^2\right). \quad (\text{S19})$$

where c is a dampening term that dictates the “strength” of the indicator function. The minimum-linespacing

constraint, $g_{LS} \leq 0$, is described by a similar function

$$g_{LS} = \frac{1}{n} \sum_{i \in \mathbb{N}} I_i^{LS}(\rho_i) \cdot [\min\{(\eta_d - \tilde{\rho}_i), 0\}]^2 \quad (\text{S20})$$

where $I_i^{LS}(\rho)$ is an indicator function that identifies the inflection region of the void phase:

$$I_i^{LS}(\rho) = (1 - \bar{\rho}) \cdot \exp\left(-c|\nabla \tilde{\rho}|^2\right). \quad (\text{S21})$$

To determine η_e and η_d from the foundry minimum linewidth (l_w) and minimum spacing (l_s), the following relations hold [24]:

$$\eta_e = \begin{cases} \frac{1}{4} \left(\frac{l_w}{R}\right)^2 + \frac{1}{2}, & \frac{l_w}{R} \in [0, 1] \\ -\frac{1}{4} \left(\frac{l_w}{R}\right)^2 + \frac{l_w}{R}, & \frac{l_w}{R} \in [1, 2] \\ 1, & \frac{l_w}{R} \in [2, \infty) \end{cases} \quad (\text{S22})$$

$$\eta_d = \begin{cases} \frac{1}{2} - \frac{1}{4} \left(\frac{l_s}{R}\right)^2, & \frac{l_s}{R} \in [0, 1] \\ 1 + \frac{1}{4} \left(\frac{l_s}{R}\right)^2 - \frac{l_s}{R}, & \frac{l_s}{R} \in [1, 2] \\ 0, & \frac{l_s}{R} \in [2, \infty) \end{cases}, \quad (\text{S23})$$

where R is the user-specified radius of a conic filter. This allows one to arbitrarily choose the filter radius and derive the subsequent threshold parameters (η_e and η_d) from the above relations.

Both the linewidth and linespacing constraints also place implicit constraints on the corresponding minimum curvature for both solid and void regions. Assuming a circular filter is used (as is the case in this work) and that the constraint is satisfied, the resulting topology cannot contain circular elements whose diameters are smaller than the corresponding lengthscale (e.g. linewidth for solid regions and linespacing for void regions). Consequently, the corresponding minimum radius of curvature $\kappa_{w,s}$ is defined by

$$\kappa_{w,s} = \frac{l_{w,s}}{2}. \quad (\text{S24})$$

We incorporated minimum-area constraints [18] using the constraint function, g_A , such that

$$g_A = \int \bar{\rho} I_A(\bar{\rho}) d\bar{\rho}, \quad (\text{S25})$$

where $\bar{\rho}$ are the projected design parameters and $I_A(\bar{\rho})$ is an indicator function that is nonzero in regions enclosing all the *islands* that are too small to be reliably fabricated. Similarly, we incorporated minimum-enclosed-area constraints using the constraint function, g_{EA} , such that

$$g_{EA} = \int (1 - \bar{\rho}) I_{EA}(1 - \bar{\rho}) d\bar{\rho} \quad (\text{S26})$$

where I_{EA} is an indicator function that marks all the regions of the topology containing *holes* that are too small to be reliably fabricated.

SECTION III: SUPPLEMENTARY FIGURES




Figure S1. Steady-state performance (a), broadband phase response (b), and broadband insertion loss (c) for various phase splitters designed to operate at $\Delta\phi = 0^\circ$, 45° , 90° , and 180° . Each device was designed to be robust to ± 20 nm of over-/under-etch variability. Worst-case performance bounds are dictated by the light-blue region, and the mean performance is dictated by the solid-blue line.




Figure S2. Steady-state field performance for the nominal 180° phase splitter, compared to its over-etched (top row) and under-etched (bottom row) variants at $\lambda = 1.50 \mu\text{m}$, $\lambda = 1.55 \mu\text{m}$, and $\lambda = 1.60 \mu\text{m}$. The effects due to over- and under-etching were simulated using the morphological erosion and dilation filters, with a perturbation of $\pm 20 \text{ nm}$. Despite the fabrication variability, the device still performs as intended (i.e. 180° phase offset between the output arms) across the entire band of interest, as shown in Fig. S1.




Figure S3. Phase extraction procedure for the optical hybrid. Using a tunable laser source and a standard fiber array, we simultaneously measured the all four output channels over the optical C-band (a). We subsequently removed the amplitude envelope induced by the grating couplers using a bandpass filter (c). Finally, we extracted the relative phase as a function of wavelength using a Hilbert transform, such that Port 4 (0°) served as the frame of reference for the remaining ports. The resulting broadband transfer function can then be used to characterize the device's demodulation capabilities and overall performance.




Figure S4. Distributions of phase error across the measured spectra and all three wafers (denoted by W_1 , W_2 , and W_3) for three of the output arms (Port 4 is used as a reference) of the optical hybrid. Systematic variability are dictated by shifts in the median phase error (red line), whereas random variability is described by the spread of the distribution ($< \pm 7^\circ$ for each arm). As systematic errors tend to dominate the source of phase error, future design iterations should optimize over multiple design fields, such as over-/under-etching, changes in the refractive index, or the layer thickness itself.