nature photonics

ARTICLES

https://doi.org/10.1038/s41566-022-01026-7

Check for updates

Mirror symmetric on-chip frequency circulation of light

Jason F. Herrmann[1]✉, Vahid Ansari[1], Jiahui Wang[1], Jeremy D. Witmer[1], Shanhui Fan[2] and Amir H. Safavi-Naeini[1]✉

Integrated circulators and isolators are important for developing on-chip optical technologies such as laser cavities, communication systems and quantum information processors. These devices seem to inherently require mirror symmetry breaking to separate backwards from forwards propagation, and thus existing implementations rely on magnetic materials or interactions driven by propagating waves. By contrast to past works, we exhibit a mirror-symmetric non-reciprocal device that comprises three coupled photonic resonators implemented in thin-film lithium niobate. Applying radiofrequency modulation, we drive conversion between the frequency eigenmodes of this system. We measure nearly 40 dB of isolation for approximately 75 mW of radiofrequency power near 1,550 nm. We simultaneously generate non-reciprocal conversion between all of the eigenmodes to demonstrate circulation. Mirror-symmetric circulation simplifies the fabrication and operation of non-reciprocal integrated devices. Finally, we consider applications of such on-chip isolators and circulators, such as full-duplex isolation within a single waveguide.

Circulators are non-reciprocal devices that allow propagating signals to cycle between three channels, for example, $in_1 \rightarrow out_3$, $in_3 \rightarrow out_2$, $in_2 \rightarrow out_1$, as shown in Fig. 1a. A circulator can isolate components from each other, that is, it can prevent reflections from returning to the source of a signal, simply by terminating one of its three channels through absorption. For example, if we terminate port 2 in Fig. 1a, signals will propagate from $in_1 \rightarrow out_3$, but $in_3 \rightarrow out_2$ would be terminated, thereby preventing reflections on the third port from propagating back to the first port. There is flexibility in selecting the input and output channels. In many realizations, the channels in_k/out_k are the incoming and outgoing waves in the waveguides connected to port k; however, other realizations are also possible. In this work, our input channels correspond with the incoming waves at different frequencies in one waveguide (on the left in Fig. 1b), whereas the output channels correspond with the outgoing waves at different frequencies at the other side of the waveguide (on the right in Fig. 1b). Unlike more typical three-waveguide configurations, this type of circulator is compatible with mirror symmetry[1]: by imposing the same scattering relations between the frequency channels for light travelling from right to left in the device, we obtain a second frequency circulation, which is related to the first by reflection about the centre axis (dashed line in Fig. 1b).

The mirror symmetry of our circulator means that its physical implementation can be simpler than more traditional circulators and isolators, which require mirror symmetry breaking. Many of these demonstrations use magnetic materials[2–7], travelling waves[8–18] or multiple modulators emitting with different phases[19–23] to induce a sense of direction and non-reciprocal propagation. As a consequence of mirror symmetry in our device, a single radiofrequency on-chip electro-optic modulator (EOM) is sufficient to generate circulation between three optical frequency channels.

Results

We demonstrate a device on a thin-film lithium niobate platform that consists of three coupled photonic racetrack resonators. The first resonator is coupled to a bus waveguide, as shown in Fig. 1c. The modes of racetrack resonators 1 (2) and 2 (3) are coupled at a rate μ_{12} (μ_{23}), whereas 1 and 3 are not directly coupled ($\mu_{13}=0$). Our device uses the TE_{00} guided mode, which on X-cut thin-film lithium niobate has large electro-optic coupling to fields parallel to both the chip surface and the crystal z-axis. Electrodes fabricated across each racetrack enable independent voltage bias tuning of the uncoupled or bare mode frequencies via the linear electro-optic effect[24]. This is necessary to counteract drift and fabrication disorder.

We set the desired operating point by using the constant voltage bias to tune all of the bare mode frequencies to be equal, which yields three evenly spaced resonances in the coupled or dressed basis, with frequencies ω_i. In an ideal system the spacing is $\Omega = \sqrt{2}\mu$, with $\mu \equiv |\mu_{12}| = |\mu_{23}|$. These three resonances couple to the bus waveguide through their spatial overlaps with the first racetrack, which leads to dips in the transmission spectrum, as shown in Fig. 1d. This overlap also means that radiofrequency modulation of the racetrack adjacent to the bus waveguide couples all three resonances together, as depicted by the schematic three-level system in the inset of Fig. 1d.

We apply radiofrequency modulation of the form:

$$V(t) = A_1 \cos(\Omega t + \phi_1) + A_2 \cos(2\Omega t + \phi_2).$$

The frequency Ω corresponds to the microwave modulation frequency, A_i are the amplitudes of the two modulation signals, ϕ_i are the modulation phases, and t is time. In an ideal system with equal intermodal couplings, μ, we would set Ω to $\sqrt{2}\mu$. Fabrication disorder and tuning imprecision means that the μ_{ij} values are slightly different. For example, $\mu_{12}/2\pi = 1.864$ GHz and $\mu_{23}/2\pi = 1.861$ GHz, as inferred for data similar to those in Fig. 1d; we therefore drive at $\Omega/2\pi = 2.63$ GHz, which is the average of the difference between the measured coupled resonance frequencies $\omega_1 - \omega_2$ and $\omega_2 - \omega_3$. Driving at Ω scatters light between dressed modes ω_1 and ω_2, whereas the 2Ω drive scatters light between modes ω_1 and ω_3.

[1]Ginzton Laboratory and Department of Applied Physics, Stanford University, Stanford, CA, USA. [2]Ginzton Laboratory and Department of Electrical Engineering, Stanford University, Stanford, CA, USA. ✉e-mail: jfherrm@stanford.edu; safavi@stanford.edu

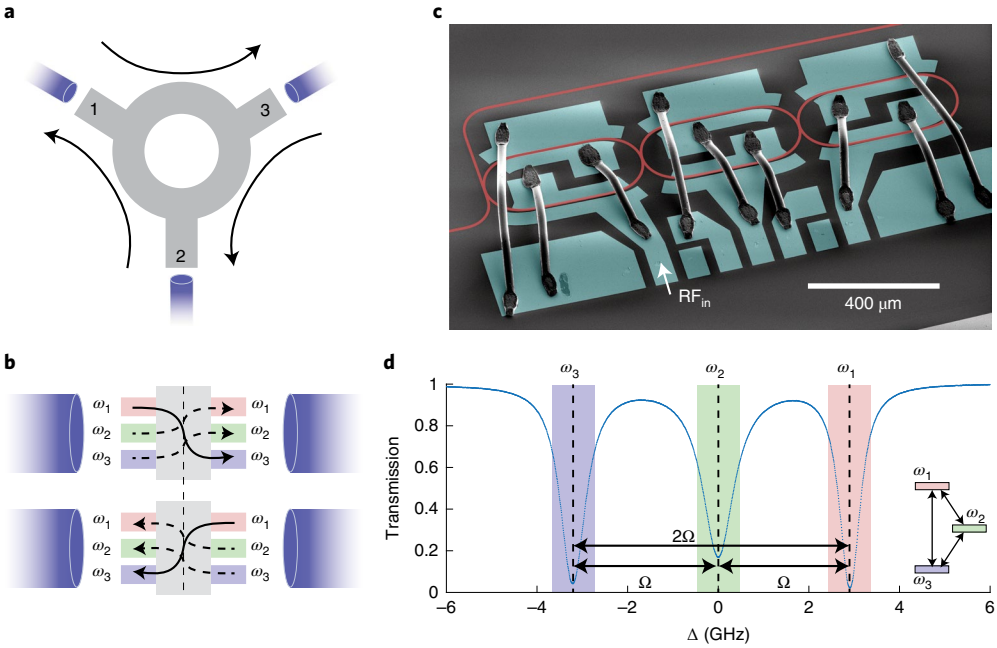

Fig. 1 | Device structure and resonant system. a, Schematic of a circulator consisting of three physical waveguide inputs/outputs. **b**, Schematic of a circulator in frequency space operating within a single waveguide. The circulator can operate equivalently by inputting light from the left- (top) or right-side (bottom) waveguide. The solid arrows emphasize the conversion from ω_1 to ω_3, the dashed arrows indicate alternative frequency conversions and the vertical dashed line depicts the structural mirror plane of the device. **c**, Scanning electron microscope image (false-coloured) of a device identical to that used in this experiment. Modulation is applied across both straight lengths of the first (left-most) racetrack. Constant voltage bias is applied to the electrodes across the second (middle) racetrack. Red, lithium niobate; blue, gold electrodes. The red lines overlapping with the waveguides are emboldened for visualization purposes. RF_{in}, radiofrequency input. **d**, Optical spectrum of the hybridized super-modes of the device operating at 1,543 nm. The inset is a schematic of the three-level system formed by the hybridized modes. Frequency conversion between any two modes may occur by two possible pathways; one path is a direct transition, whereas the other occurs in two steps, mediated by the third level.

For any pair of dressed modes, two possible transition pathways exist, as depicted by the schematic inset in Fig. 1d. One pathway is a direct transition, whereas the other is a two-step transition through the third mode. By varying the amplitudes and relative phases of the two radiofrequency tones, we enhance or suppress different pathways by generating interference. In the ideal disorder-free model, forward isolation (mode $\omega_1 \rightarrow \omega_3$) is maximized for the phase condition $2\phi_1 - \phi_2 = \pi/2$.

We characterize the performance of our circulator by measuring the scattering parameters S_{ij}, which quantify how the amplitude a_j of each incoming wave at frequency ω_j is converted to the amplitude of an outgoing wave b_i at frequency ω_i, propagating in the same direction in the opposite waveguide (depicted in Fig. 2a). We define an isolation parameter as

$$I_{ij} \equiv |S_{ij}|^2/|S_{ji}|^2.$$

This parameter characterizes the asymmetry between forwards and backwards frequency conversion for a pair of modes. As shown in Fig. 2b for channels 1 and 3, it is clear that forward scattering from $1 \rightarrow 3$ occurs efficiently, whereas backward scattering is strongly suppressed. This leads to I_{31} approaching 40 dB when the driving phase condition is $\pi/2$ and A_i is correctly tuned, as explained below. Similar isolation parameters are observed for other pairs of channels.

We obtain the scattering and isolation parameters described above by first characterizing the device's linear spectrum, with the on-chip radiofrequency modulation turned off. Fig. 3a depicts a schematic of the experimental set-up. We measure the optical transmission parameter $t(\omega)|_{RF,off}$ as a function of frequency, which we then use to infer and bias-voltage-tune the device's parameters (Fig. 3b,c). This transmission amplitude evaluated at frequencies ω_j corresponds to $S_{jj}|_{RF,off}$ as it describes transmission through the device without a change in frequency. We place a laser tone at frequency ω_0 (blue-detuned from all of the resonances) and feed it through an off-chip EOM before sending it to the device. Driving the off-chip EOM with a vector network analyser (VNA), we generate two sidebands at $\omega_0 \pm \Delta$. By sweeping the VNA modulation frequency, we move one sideband across the cavity response, which then beats against the optical feed-through, ω_0, on a subsequent radiofrequency photodiode connected to the VNA. Sweeping the VNA frequency allows us to see (almost in real-time) the hybridized mode structure of the device (Fig. 3b). Our device exhibits voltage bias drift as a result of photorefraction—a common challenge in lithium niobate devices[25–27]. This measurement technique enables us to observe and compensate for these changes, as the optical input power and constant voltage bias both affect the spectrum. We fit a complete model of the measurement—including the off-chip EOM transmission, cavity response and phase response measured by the VNA—to determine all of the device's optical parameters (loss rates, resonant frequencies and so on). An example of such a fit is depicted in Fig. 3c. We can also extract the sideband transmission $t(\omega)|_{RF,off}$ from these spectra.

We then measure the active device by turning on the on-chip radiofrequency drive and characterizing the scattering between frequency channels. As VNA measurements only find linear scattering parameters, characterizing the scattering between frequencies requires a different measurement scheme than that described above. Keeping the laser at ω_0, we modulate the off-chip EOM at frequency Δ to generate and input light to the chip at frequency

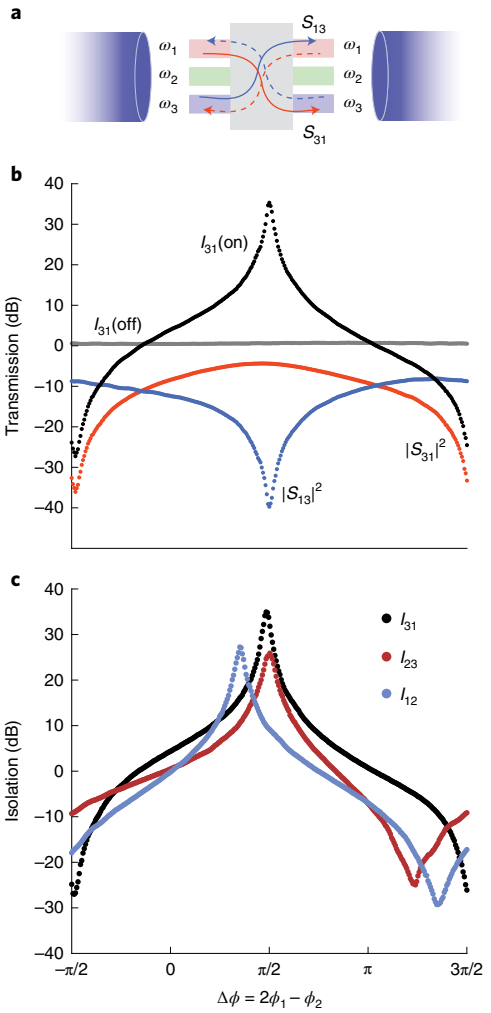

Fig. 2 | Isolation versus microwave phase condition $\Delta\phi = 2\phi_1 - \phi_2$. **a**, Schematic of frequency-domain circulator. Solid arrows indicate light input from the waveguide on the left (as measured), whereas dashed arrows indicate the same conversions when light is input from the waveguide on the right. **b**, Isolation for $\omega_1 \rightarrow \omega_3$ conversion. The line labelled 'off' corresponds to a separate measurement in which we decrease the Ω frequency drive (P_1) while keeping the 2Ω drive (P_2) high, indicating the importance of coherent interference of the two processes for non-reciprocity. The isolation measurement is taken for on-chip radiofrequency powers of $P_1 = 73.2$ mW and $P_2 = 59.0$ mW; $P_2 = 74.3$ mW for the separate 'off' measurement. **c**, Mutual isolation for all conversions observed for the 'on' power condition in **b**. Shifts between the peaks and the ideal phase condition stem from disorder in the mode hybridization.

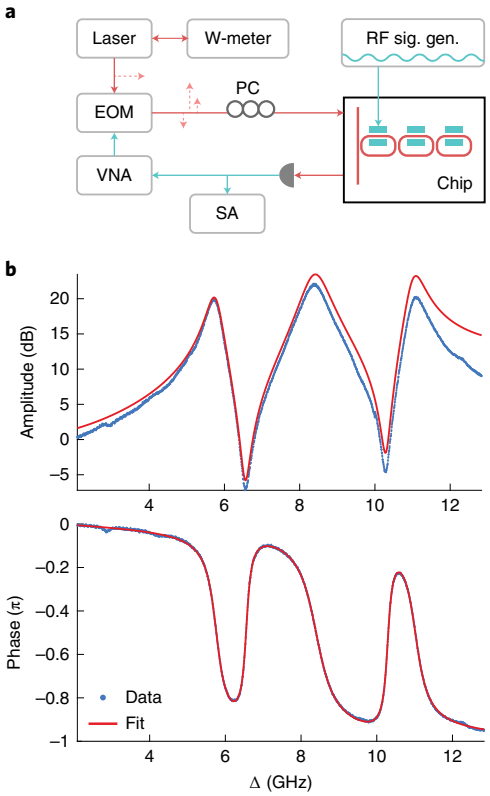

Fig. 3 | Experiment characterization scheme. a, Simplified schematic of the characterization set-up. The laser wavelength is locked to a Bristol Wavelength Meter (W-meter). An off-chip EOM is driven by a VNA to generate optical input tones to the chip. Separate radiofrequency signal generators (RF sig. gen.) are used to produce the on-chip radiofrequency signal for frequency conversion. The output from the chip is incident on a photodiode and recorded on both a VNA and a spectrum analyser (SA). **b**, Background-normalized amplitude and phase of the VNA spectrum produced by the off-chip EOM response filtered through the optical cavity. This is linear transmission through the device with on-chip radiofrequency turned off. The off-chip EOM is set to a bias point phase shift of approximately π. The fit is made to the phase response and used to predict the amplitude response.

$\omega_j = \omega_0 - \Delta$. We infer all of the optical powers in the different frequency channels by taking radiofrequency measurements of their beating against the feed-through light at ω_0. This presents a minor complication for inferring S_{jj}, as the beat notes generated by both off-chip EOM sidebands $\omega_0 \pm \Delta$ interfere at radiofrequency Δ. This ambiguity can be resolved by comparing measurements of S_{jj} with the independently measured scattering parameters (see the above paragraph) when radiofrequency modulation to the chip is turned off. Note that the other scattering parameters $S_{ij}, i \neq j$ can be inferred without this complication, and thus the isolation parameters I_{ij} are unaffected. For all scattering parameters, we record the radiofrequency power in the beat note that corresponds with optical frequency ω_i while varying the phase ϕ_2 of the radiofrequency drive sent to the chip; this power is proportional to $|b_i|^2$. We take the ratio of the measured power in ω_i to that of the unmodulated transmission at the signal frequency ω_j. By factoring out the contribution from the non-resonant input tone at $\omega_0 + \Delta$, we obtain the scattering parameters $|S_{ij}|^2 = |b_i/a_j|^2$ (see Supplementary Section 2).

The scattering parameter measurements shown in Fig. 2 are taken for optimized values of the on-chip radiofrequency drive amplitude A_i. When the Ω radiofrequency drive is turned down (that is, $A_1 \approx 0$) but the 2Ω drive is kept on, we do not observe isolation (dashed line in Fig. 2b). For tuned values of A_1 and A_2, we observe simultaneous isolation between the three pairs of dressed frequencies, demonstrating circulator-like behaviour, as depicted in Fig. 2c.

We next characterize the radiofrequency power-dependent operation of our device by varying the modulation amplitudes A_1 and A_2. For each power combination, we sweep the phase condition of the radiofrequency sources through multiple periods and extract the peak isolation observed for each transition. We generate maps of the isolation versus on-chip radiofrequency power (Fig. 4a–c). Each pixel is normalized by linear scattering parameters, analogously to the data in Fig. 2. The optical device characteristics used for this

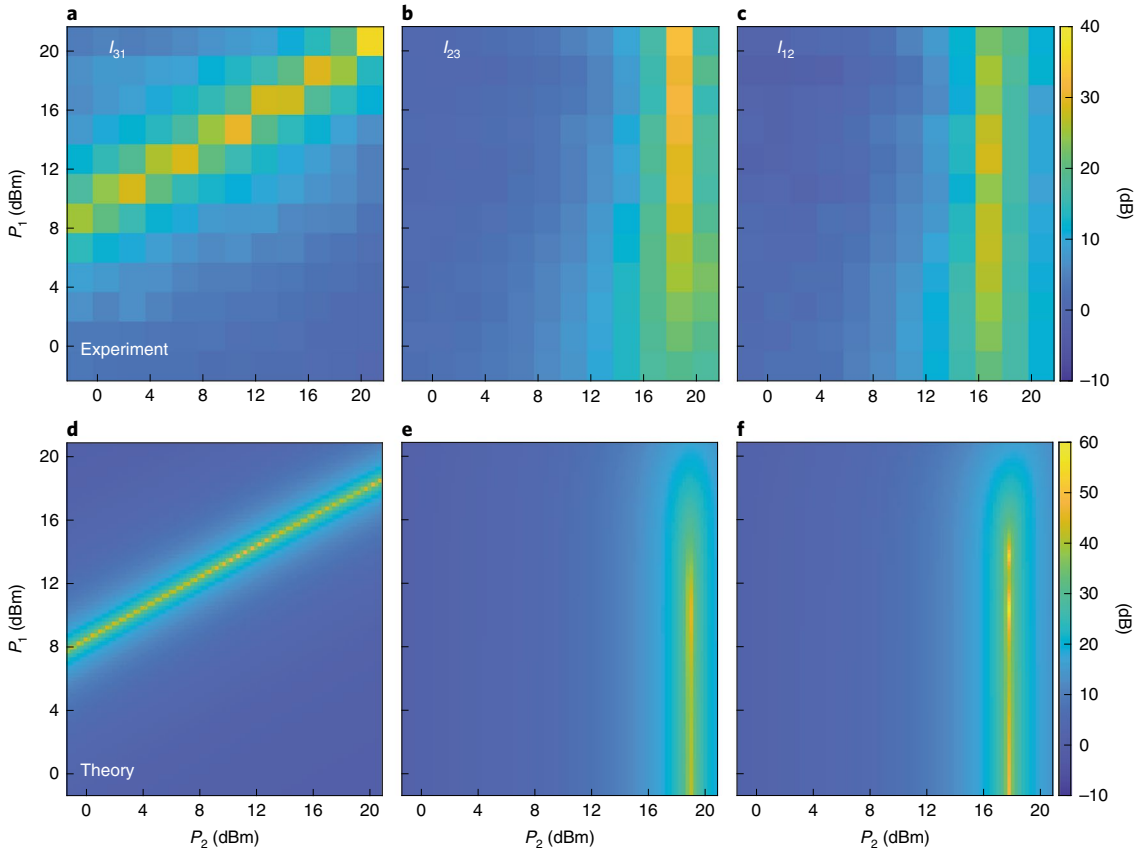

Fig. 4 | Isolation versus microwave power. a–f, We sweep the two radiofrequency drive powers, indicated here as P_1 (Ω drive) and P_2 (2Ω drive). For each power combination, the phase of the 2Ω drive (ϕ_2) is swept across multiple periods. Each pixel corresponds to the maximum observed isolation for each radiofrequency power combination for a given isolation parameter. The phase condition varies between pixels of different maps. The top row corresponds to measurements (**a–c**), whereas the bottom row shows the theoretical predictions of isolation (**d–f**). We attribute the differences between experiment and theory to system parameters shifting over time and as the radiofrequency power increases (see Supplementary Section 4). Predictions are based on the fit system parameters, which also have some degree of uncertainty. Colorbars indicate the isolation parameter, I_{ij}.

normalization are extracted from the VNA trace in Fig. 3, taken before varying the radiofrequency power.

Figure 4d–f depicts theoretical plots of device performance obtained from coupled mode theory and show good agreement with the measured trends. We attribute the differences between theory and experiment (for example, the slopes in the I_{31} maps and the locations of peak isolation in I_{12} and I_{23}) to mode drifts over the full course of measurement. This drift could emerge from either voltage bias drift, which introduces disorder in the mode hybridization, or from high-power radiofrequency-induced drift (see Supplementary Section 4). We assume this drift does not substantially affect the normalization factor we use for data processing, which is obtained from the initial VNA trace. We infer a high-frequency bare mode modulation rate of $g_{EO}/2\pi \approx 330$ MHz/V from the theoretical plots, where g_{EO} is the electro-optic tuning rate of the bare modes. This rate closely matches the directly measured low-frequency modulation rate $g_{EO,DC}/2\pi = 328.5$ MHz/V.

We can understand the trends in the isolation factor radiofrequency power dependence by studying the coupled mode theory description of the dynamics[1]. The radiofrequency modulation scatters photons between the dressed modes at a rate proportional to $\tilde{A}_i = g_{EO}A_i$. We observe strong enhancement in I_{31} in Fig. 4a,d along a locus corresponding to $\tilde{A}_1^2 \propto \tilde{A}_2$. This condition is found by noting that, for cancellation to occur, the rate at which direct $1 \rightarrow 3$ scattering occurs ($\tilde{A}_2$) needs to match the rate of indirect scattering $1 \rightarrow 2 \rightarrow 3$, which is approximately $\tilde{A}_1^2/\gamma_2$ (by second-order perturbation theory). Here, γ_2 is the linewidth (full-width at half-maximum) of the second dressed mode. A more detailed analysis[1] indicates that maximum I_{13} isolation is given by:

$$\tilde{A}_2 = \frac{\tilde{A}_1^2}{2\gamma_2},$$

where $\gamma_2 = (\kappa_1 + \kappa_3)/2$.

A similar line of reasoning explains the large isolation regions in Fig. 4b,c,e,f. For example, note that a peak I_{12} requires interference between two scattering processes—a direct process $1 \rightarrow 2$, occurring with a rate $\tilde{A}_1$, and an indirect process $1 \rightarrow 3 \rightarrow 2$, occurring at a rate $\tilde{A}_1\tilde{A}_2/\gamma_3$. This interference is maximized when the rates are nearly equal. A more detailed coupled mode analysis confirms that the ideal condition $\tilde{A}_2\tilde{A}_1 = 4\gamma_{1,3}\tilde{A}_1$, where $\gamma_{1,3} = (\kappa_1 + 2\kappa_2 + \kappa_3)/4$, maximizes the isolation parameter. The isolation parameters I_{12} and I_{23} are therefore maximized when:

$$\tilde{A}_2 = 4\gamma_{1,3},$$

independently of the direct scattering rate $\tilde{A}_1$. Inhomogeneity in the mode hybridization causes deviations from this exact condition for I_{23} versus I_{12}, that is, they appear at slightly different power conditions in Fig. 4b,c.

Discussion

Our platform's most substantial limit is due to optical loss in the cavities. In the ideal realization the dressed mode linewidths are entirely due to coupling out into the feed waveguide. When the on-chip radiofrequency drives are turned off, the absence of intrinsic loss means that the dressed modes are over-coupled, and thus $|S_{kk}|=1$. As we increase the radiofrequency driving power, the scattering between dressed modes appears as additional loss at the signal frequency. This moves the mode from being over-coupled into a critically coupled condition, in which the scattering-induced loss at mode k equals the mode's coupling rate to the waveguide, leading to $|S_{kk}|=0$. However, in the experimentally realized system, there is additional intrinsic loss in all three cavities, and the dressed modes are already close to being critically coupled when the radiofrequency drives are off. The radiofrequency-induced scattering then acts to increase transmission $|S_{kk}|$ as it moves the mode farther away from being critically coupled. This means that in contrast to the ideal circulator, the diagonal elements of the scattering matrix are non-zero in our realization. We expect that this problem can be largely eliminated with improvements to the optical Q of on-chip lithium niobate devices[28]. One possible way to circumvent this challenge in the shorter term is to include an additional resonator after the device at frequency ω_k to filter out any feed-through and make $S_{kk}=0$ for a single channel (see Supplementary Section 2). Another approach that does not require major improvements in Q is to increase the cavity–waveguide coupling and start within a more over-coupled regime. However, this approach would require larger radiofrequency powers and we are already power-limited. The off-chip microwave amplifier saturates at approximately 35 dBm. Moreover, as power to the device increases above roughly 120 mW, we observe dressed mode drifts on the order of megahertz, which disrupts the resonance condition required for high isolation (see Supplementary Section 4).

Another limit of our approach is due to its resonant nature. The isolation bandwidth of our device depends on the dressed mode linewidths, and we measured sustained isolation over more than a few hundred megahertz. This too can be increased, but at the cost of greater radiofrequency power as $\tilde{A}_k$ also scales with bandwidth. One strategy to make the modulation more efficient would be to include modulation into the third racetrack. Alternatively, as we only need narrowband radiofrequency modulation, we could use the extremely efficient optomechanical modulation schemes recently demonstrated on lithium niobate[29–31] to considerably reduce the needed radiofrequency power and realize larger bandwidth. Finally, we note that the bandwidth of the device is potentially larger than the linewidth of a single mode as there are families of modes repeating with the cavity free spectral range; under low disorder, these would also behave as circulators for signals at different sets of frequencies.

A novel feature of our device is its operation as a bidirectional isolator—a property that emerges from mirror symmetry and frequency-domain operation. For example, two transmitters/receivers can operate simultaneously along a single channel. Node 1 can transmit a signal down a waveguide at frequency ω_1 to node 2, which recieves at ω_3. This communication is isolated. Meanwhile, node 2 can then send a signal backwards along the waveguide at ω_1 to node 1, which also receives at ω_3; this communication is also isolated. With the standard optical isolator, this behaviour would require two optical channels, one for each direction of transmission.

Overall, we have demonstrated an integrated frequency isolator/circulator on thin-film lithium niobate, an emerging platform for classical and quantum photonics. We measured peak isolation of nearly 40 dB with 4.4 dB insertion loss for dual radiofrequency drive powers of $P_1=73.2$ mW and $P_2=59.0$ mW. Our device is reconfigurable, enabling isolation over a wide range of powers for different operating frequencies. For example, we also measure $\omega_1\rightarrow\omega_3$ isolation of more than 25 dB for dual radiofrequency powers of $P_1=7.32$ mW and $P_2=0.74$ mW, but with a commensurate increase of insertion loss (~16 dB). Furthermore, we also measured insertion loss as low as 4 dB for different power and isolation conditions. Ultimately, the mirror symmetry and frequency-domain operation of our device provide for novel applications as an isolator/circulator and frequency router in photonic circuits.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41566-022-01026-7.

Received: 14 September 2021; Accepted: 20 May 2022;
Published online: 07 July 2022

References

1. Wang, J., Herrmann, J. F., Witmer, J. D., Safavi-Naeini, A. H. & Fan, S. Photonic modal circulator using temporal refractive-index modulation with spatial inversion symmetry. *Phys. Rev. Lett.* **126**, 193901 (2021).
2. Bi, L. et al. On-chip optical isolation in monolithically integrated non-reciprocal optical resonators. *Nat. Photon.* **5**, 758–762 (2011).
3. Tzuang, L. D., Fang, K., Nussenzveig, P., Fan, S. & Lipson, M. Non-reciprocal phase shift induced by an effective magnetic flux for light. *Nat. Photon.* **8**, 701–705 (2014).
4. Srinivasan, K. & Stadler, B. J. Magneto-optical materials and designs for integrated TE-and TM-mode planar waveguide isolators: a review. *Opt. Mater. Express* **8**, 3307–3318 (2018).
5. Huang, D. et al. Electrically driven and thermally tunable integrated optical isolators for silicon photonics. *IEEE J. Sel. Top. Quantum Electron.* **22**, 271–278 (2016).
6. Yan, W. et al. Waveguide-integrated high-performance magneto-optical isolators and circulators on silicon nitride platforms. *Optica* **7**, 1555–1562 (2020).
7. Sobu, Y., Shoji, Y., Sakurai, K. & Mizumoto, T. GaInAsP/InP MZI waveguide optical isolator integrated with spot size converter. *Opt. Express* **21**, 15373–15381 (2013).
8. Kittlaus, E. A., Otterstrom, N. T., Kharel, P., Gertler, S. & Rakich, P. T. Non-reciprocal interband Brillouin modulation. *Nat. Photon.* **12**, 613–619 (2018).
9. Kittlaus, E. A. et al. Electrically driven acousto-optics and broadband non-reciprocity in silicon photonics. *Nat. Photon.* **15**, 43–52 (2021).
10. Kim, S., Sohn, D. B., Peterson, C. W. & Bahl, G. On-chip optical non-reciprocity through a synthetic hall effect for photons. *APL Photon.* **6**, 011301 (2021).
11. Shen, Z. et al. Experimental realization of optomechanically induced non-reciprocity. *Nat. Photon.* **10**, 657–661 (2016).
12. Ruesink, F., Miri, M.-A., Alu, A. & Verhagen, E. Nonreciprocity and magnetic-free isolation based on optomechanical interactions. *Nat. Commun.* **7**, 1–8 (2016).
13. Kang, M. S., Butsch, A. & Russell, P. S. J. Reconfigurable light-driven opto-acoustic isolators in photonic crystal fibre. *Nat. Photon.* **5**, 549–553 (2011).
14. Kim, J., Kuzyk, M. C., Han, K., Wang, H. & Bahl, G. Non-reciprocal Brillouin scattering induced transparency. *Nat. Phys.* **11**, 275–280 (2015).
15. Dong, C.-H. et al. Brillouin-scattering-induced transparency and non-reciprocal light storage. *Nat. Commun.* **6**, 1–6 (2015).
16. Hafezi, M. & Rabl, P. Optomechanically induced non-reciprocity in microring resonators. *Opt. Express* **20**, 7672–7684 (2012).
17. Tian, H. et al. Magnetic-free silicon nitride integrated optical isolator. *Nat. Photon.* **15**, 828–836 (2021).
18. Sohn, D. B., Örsel, O. E. & Bahl, G. Electrically driven optical isolation through phonon-mediated photonic Autler–Townes splitting. *Nat. Photon.* **15**, 822–827 (2021).
19. Fang, K., Yu, Z. & Fan, S. Photonic Aharonov–Bohm effect based on dynamic modulation. *Phys. Rev. Lett.* **108**, 153901 (2012).
20. Lira, H., Yu, Z., Fan, S. & Lipson, M. Electrically driven nonreciprocity induced by interband photonic transition on a silicon chip. *Phys. Rev. Lett.* **109**, 033901 (2012).
21. Doerr, C. R., Dupuis, N. & Zhang, L. Optical isolator using two tandem phase modulators. *Opt. Lett.* **36**, 4293–4295 (2011).
22. Yu, Z. & Fan, S. Complete optical isolation created by indirect interband photonic transitions. *Nat. Photon.* **3**, 91–94 (2009).
23. Dostart, N., Gevorgyan, H., Onural, D. & Popović, M. A. Optical isolation using microring modulators. *Opt. Lett.* **46**, 460–463 (2021).

24. Weis, R. S. & Gaylord, T. K. Lithium niobate: summary of physical properties and crystal structure. *Appl. Phys. A* **37**, 191–203 (1985).
25. McKenna, T. P. et al. Cryogenic microwave-to-optical conversion using a triply resonant lithium-niobate-on-sapphire transducer. *Optica* **7**, 1737–1745 (2020).
26. Jiang, H. et al. Fast response of photorefraction in lithium niobate microresonators. *Opt. Lett.* **42**, 3267–3270 (2017).
27. Xu, Y. et al. Mitigating photorefractive effect in thin-film lithium niobate microring resonators. *Opt. Express* **29**, 5497–5504 (2021).
28. Shams-Ansari, A. et al. Probing the limits of optical loss in ion-sliced thin-film lithium niobate. In *Conference on Lasers and Electro-Optics* STh4J.4 (Optical Society of America, 2021); http://www.osapublishing.org/abstract.cfm?URI=CLEO_SI-2021-STh4J.4
29. Jiang, W. et al. Efficient bidirectional piezo-optomechanical transduction between microwave and optical frequency. *Nat. Commun.* **11**, 1166 (2020).
30. Shao, L. et al. Microwave-to-optical conversion using lithium niobate thin-film acoustic resonators. *Optica* **6**, 1498–1505 (2019).
31. Sarabalis, C. J., McKenna, T. P., Patel, R. N., Van Laer, R. & Safavi-Naeini, A. H. Acousto-optic modulation in lithium niobate on sapphire. *APL Photon.* **5**, 086104 (2020).

Methods

Fabrication. The device is fabricated in a two-mask process from 500-nm-thick film of lithium niobate atop a sapphire handle. The components are air-clad and the electrodes are fabricated as a Ti:Au metal bilayer. The first mask, for photonics fabrication, is defined from hydrogen silsequioxane, a negative-tone electron beam resist, patterned with 100 kV electron beam lithography (JEOL JBX-6300FS). We transfer the patterns into the lithium niobate with argon ion mill etching (IntlVac ion mill), followed by an acid cleaning procedure. We confirm the device's optical performance before fabricating the electrodes. The electrodes are defined with a standard photoresist lift-off bilayer. The patterns are written using direct-write lithography (Heidelberg MLA150) and metal is evaporated (Kurt J. Lesker LAB18) before solvent-based lift-off. Finally, we wirebond on-chip to connect electrodes across the optics, thereby defining proper modulation polarity on-chip (West Bond 7476E).

Characterization. The device is optically pumped from a telecommunications wavelength diode laser (Santec TSL-550). We lock the optical pump blue-detuned from the modes at 1,543.6 nm (via a Bristol Wavelength Meter). The light is passed through a polarization control wheel and in-line fibre polarizer to maximize transmission of transverse electric (TE) polarized light. This light is passed into a commercial EOM. We drive the EOM from a VNA (R&S ZNB20) to generate sidebands. The pump and sidebands are then passed through another polarization controller, into a variable optical attenuator and then a power meter before being passed onto the chip. By sweeping the frequency of the VNA output (or by driving a particular frequency), we can sweep the sideband across (or directly drive) the optical dressed modes of the system. The light undergoes modulation on-chip. Output light from the optical waveguide passes through an erbium-doped fibre amplifier (EDFA) and an X-switch, which switches our detector between a photodiode for linear optical characterization, and a fast photodiode (Optilab PD-40M), which records beat tones between the optical pump, the EOM sideband and converted sidebands when applying on-chip modulation. The output of the fast photodiode passes through a bias-tee. One arm passes back to the VNA to record the broad-band response of the cavity, and the other arm passes to a spectrum analyser (R&S FSW26) to record converted sideband powers with high precision. This set-up enables us to simultaneously observe the full cavity response on the VNA, almost in real-time, while adjusting constant voltage bias across the device and manipulating the optical pump, while also observing sideband powers at particular frequencies when driving on-chip modulation.

To characterize our device mirror symmetry, we separately insert a second X-switch and polarization control just before the device. We record transmission through the device from both waveguide propagation directions, demonstrating equivalence of scattering matrix elements (see Supplementary Section 2).

On-chip modulation is driven from two pulsed signal generators (PSG, Keysight E8257D). One PSG drives the Ω tone while the other drives the 2Ω tone. We vary the relative phase between these sources. The clocks of the PSGs are locked together, and these are in turn locked to the clock of the FSW26. The two PSG outputs are combined at a power splitter and passed through a high-power microwave amplifier. The output of the amplifier is then passed to a probe, which is contacted to the on-chip electrical pads. Importantly, the rate of direct phase modulation on the PSG is much faster than the global phase drift of the system, enabling us to visualize power modulations at sideband frequencies on the FSW26.

Data availability

The data comprising Fig. 4 and Supplementary Fig. 8 are available on Zenodo at https://doi.org/10.5281/zenodo.6537345. Additional data generated and analysed in this study are available from the corresponding author on reasonable request.

Acknowledgements

J.F.H. acknowledges support from the National Science Foundation Graduate Research Fellowship Program (grant no. DGE-1656518). V.A. acknowledges support by the Stanford Q-FARM Bloch Fellowship Program. We acknowledge the support of an AFOSR MURI project (grant no. FA9550-18-1-0379) and the National Science Foundation under award no. ECCS-1820938. Part of this work was performed at the Stanford Nano Shared Facilities (SNSF), supported by the National Science Foundation under award no. ECCS-2026822. Work was performed in part at the nano@stanford laboratories, which are supported by the National Science Foundation as part of the National Nanotechnology Coordinated Infrastructure under award no. ECCS-1542152. We would like to thank W. Jiang and C. J. Sarabalis for insightful and helpful discussions.

Author contributions

J.F.H. fabricated the device. J.F.H. and V.A. led the experimental effort. J.W. developed the device operating theory and characterized theoretical device performance. J.F.H, J.D.W. and J.W. determined physical device designs. J.D.W. assisted in early experimentation. A.H.S.-N and S.F. provided experimental and theoretical guidance and support for this experiment.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41566-022-01026-7.

Correspondence and requests for materials should be addressed to Jason F. Herrmann or Amir H. Safavi-Naeini.

Peer review information *Nature Photonics* thanks the anonymous reviewers for their contribution to the peer review of this work.