Video Article

Measurement of Quantum Interference in a Silicon Ring Resonator Photon Source

Jeffrey A. Steidle¹, Michael L. Fanto^{1,2}, Stefan F. Preble¹, Christopher C. Tison^{2,3,4}, Gregory A. Howland², Zihao Wang¹, Paul M. Alsing²

¹Microsystems Engineering, Rochester Institute of Technology

²Air Force Research Laboratory, Rome, NY

³Department of Physics, Florida Atlantic University

⁴Quanterion Solutions Incorporated

Correspondence to: Jeffrey A. Steidle at jas2518@rit.edu

URL: <https://www.jove.com/video/55257>

DOI: [doi:10.3791/55257](https://doi.org/10.3791/55257)

Keywords: Engineering, Issue 122, Silicon, Photonics, Quantum, Interference, Ring Resonator, Photon Source

Date Published: 4/4/2017

Citation: Steidle, J.A., Fanto, M.L., Preble, S.F., Tison, C.C., Howland, G.A., Wang, Z., Alsing, P.M. Measurement of Quantum Interference in a Silicon Ring Resonator Photon Source. *J. Vis. Exp.* (122), e55257, doi:10.3791/55257 (2017).

Abstract

Silicon photonic chips have the potential to realize complex integrated quantum information processing circuits, including photon sources, qubit manipulation, and integrated single-photon detectors. Here, we present the key aspects of preparing and testing a silicon photonic quantum chip with an integrated photon source and two-photon interferometer. The most important aspect of an integrated quantum circuit is minimizing loss so that all of the generated photons are detected with the highest possible fidelity. Here, we describe how to perform low-loss edge coupling by using an ultra-high numerical aperture fiber to closely match the mode of the silicon waveguides. By using an optimized fusion splicing recipe, the UHNA fiber is seamlessly interfaced with a standard single-mode fiber. This low-loss coupling allows the measurement of high-fidelity photon production in an integrated silicon ring resonator and the subsequent two-photon interference of the produced photons in a closely integrated Mach-Zehnder interferometer. This paper describes the essential procedures for the preparation and characterization of high-performance and scalable silicon quantum photonic circuits.

Video Link

The video component of this article can be found at <https://www.jove.com/video/55257/>

Introduction

Silicon is showing great promise as a photonics platform for quantum information processing^{1,2,3,4,5}. One of the vital components of quantum photonic circuits is the photon source. Photon-pair sources have been developed from silicon in the form of micro-ring resonators made via a third-order nonlinear process, spontaneous four-wave mixing (SFWM)^{6,7,8}. These sources are capable of producing pairs of indistinguishable photons, which are ideal for experiments involving photon entanglement⁹.

It is important to note that ring resonator sources can operate with both clockwise and counter-clockwise propagation, and the two different propagation directions are generally independent from one another. This allows a single ring to function as two sources. When optically pumped from both directions, these sources generate the following entangled state:

$$|\psi\rangle_{ring} = \frac{1}{2} \left[\hat{a}_{cw}^{\dagger 2} + \hat{a}_{ccw}^{\dagger 2} \right] |vac\rangle$$

where $\hat{a}_{cw}^{\dagger}$ and $\hat{a}_{ccw}^{\dagger}$ are the independent creation operators for clockwise- and counterclockwise-propagating bi-photons, respectively. This is a very desirable form of entangled state known as a NOON state (N=2)¹⁰.

Passing this state through an on-chip Mach-Zehnder interferometer (MZI) results in the state:

$$|\psi\rangle_{out} = \frac{1}{4} e^{-i\theta} \left[\hat{a}_{cw}^{\dagger 2} - \hat{a}_{ccw}^{\dagger 2} + 2i\hat{a}_{cw}^{\dagger}\hat{a}_{ccw}^{\dagger} + \left(-\hat{a}_{cw}^{\dagger 2} + \hat{a}_{ccw}^{\dagger 2} + 2i\hat{a}_{cw}^{\dagger}\hat{a}_{ccw}^{\dagger} \right) e^{i2\theta} \right] |vac\rangle$$

This state oscillates between maximum coincidence and zero coincidence at twice the frequency of classical interference in an MZI, effectively doubling the sensitivity of the interferometer¹⁰. Here, we present the procedure used to test such an integrated photon source and MZI device.

Protocol

NOTE: This protocol assumes that the photonic chip has already been fabricated. The chip described here (shown in **Figure 1A**) was fabricated at the Cornell University NanoScale Science & Technology Facility using standard processing techniques for silicon photonic devices¹¹. These include the use of silicon-on-insulator wafers (composed of a 220 nm-thick silicon layer, a 3- μm layer of silicon dioxide, and a 525 μm -thick silicon substrate), electron-beam lithography to define the strip waveguides (500 nm-wide), and the plasma-enhanced chemical vapor deposition of the silicon dioxide cladding ($\sim 3\text{ }\mu\text{m}$ -thick). The micro-ring resonators were designed with an inner radius of 18.5 μm and a waveguide-to-ring gap of 150 nm. Figures of merit for this device include loss, quality factor, free spectral range, and dispersion.

1. Photonic Chip Preparation

- Place a small amount of wax on a cross-sectioning polishing mount and heat it to $\sim 130\text{ }^{\circ}\text{C}$.
NOTE: The amount of wax to be used depends upon the size of the sample being mounted. There must be enough wax to keep the chip immobile, while too much will result in wax on the chip facets.
- Place the photonic chip on the portion of the polishing mount with the wax. Ensure that the wax is melted completely so that the chip is flat against the mount. Use plastic tweezers when handling the chip to avoid damaging the facets.
- Allow the mount to cool in ambient air so that the wax solidifies. Cooling faster than this may result in damage to the chip.
- Polishing the chip facets.
NOTE: It is important to choose the correct lapping pad as starting with a pad that is too aggressive can result in polishing away more of the chip than desired.
 - Attach the polishing mount to the polisher and polish for only a few seconds. A pad with a 3- μm roughness has shown to be a good starting point for silicon chips with facet lengths of $\sim 1\text{ cm}$.
 - Remove the polishing mount and inspect the chip facet to determine how level the chip is mounted.
NOTE: A microscope is useful for measuring the distance between the ends of the waveguides and the facet of the chip. These measurements allow for the angle between the facet and the waveguide to be determined.
 - Make the necessary adjustments to the micrometers on the polisher in order to improve the leveling of the chip.
 - Repeat steps 1.4.1-1.4.3 until the facet of the chip and the waveguides are within 0.15° of being orthogonal to each other.
 - Polish the chip in steps of $\sim 50\text{ }\mu\text{m}$, inspecting the chip between each step to monitor the distance remaining, until there is $\sim 100\text{ }\mu\text{m}$ left to polish. If at any point the cladding appears to be delaminating from the surface, ensure that the pad is rotating so as to polish from the top of the chip to the bottom.
NOTE: It may also help to use a polishing lubricant instead of water. This delamination is a result of stress in the cladding and is an indication that the fabrication process needs to be optimized.
 - Change to a 1- μm lapping pad and polish until there is $\sim 20\text{ }\mu\text{m}$ remaining.
 - Change to a 0.5- μm pad and continue polishing for another 15 μm .
 - Use a 0.1- μm pad for the final 5 μm to ensure a smooth facet. Microscope images of the facet of a silicon photonic chip before and after polishing are shown in **Figure 2**.
- Heat the mount with the attached chip to $\sim 130\text{ }^{\circ}\text{C}$ to allow the wax to melt.
- Once the wax is completely melted, remove the chip from the mount and allow it to cool slowly.
- Clean any of the remaining wax from the chip using acetone, isopropanol, and water.

2. Preparation of Fiber Pigtailed

- Strip any buffer or coating from the end of a single mode fiber (SMF) pigtail and from one end of an ultra-high numerical aperture (UHNA) fiber.
- Clean the bare ends of the fibers with a mixture of acetone and methanol.
- Cleave the bare ends of both fibers with a commercial fiber cleaver.
- Fusion splice the cleaved end of the fibers. A recipe for splicing SMF to UHNA fiber is shown in **Table 1**.
- Slide a protective sleeve over the splice and place it in the sleeve oven to permanently attach it to the fiber.
- Repeat steps 2.1-2.5 to prepare a total of three fibers.

3. Configuration of the Testing Setup

NOTE: A diagram of the testing setup is shown in **Figure 1B**. The mount for the chip is a copper pedestal that is in contact with a thermo-electric cooler (TEC). There is a microscope fitted with both visible and infrared (IR) cameras for viewing the photonic chip.

- Place a small amount of wax on the chip mount and apply voltage to the TEC to melt the wax.
- Place the chip on the melted wax, ensuring that it is sitting flat on the mount.
- Remove the voltage from the TEC and allow the mount and chip to cool slowly.
- Attach each of the spliced fibers to a fiber v-groove with polyimide tape and mount a single v-groove to each of the 3-axis stages using the manufacturer-supplied mounting hardware.
- Fiber edge coupling.
 - Connect the three fibers to their respective components: one to the optical output of the laser and the other two to optical power meters.
 - Adjust the microscope so that it is focused on the chip where the waveguides reach the edge.

3. Position the fibers close to the chip edge so that they are in the view of the visible camera and adjust their heights so that the core of each fiber is in focus.
4. Adjust the horizontal positioning of the fibers with the stage micrometers so that they are lined up with the waveguides.
5. Turn on the optical output of the laser and tune the horizontal and vertical micrometer positions of the input fiber until the light is coupling into the waveguide. This will show up on the IR camera as scattering along the input waveguide.
6. Tune the wavelength of the laser to a point where the micro-ring resonator is lit up on the camera. This indicates that the resonance condition is being satisfied and that the light is reaching the output waveguides.
7. Adjust the horizontal and vertical micrometer positions of the output fibers until there is a measurable amount of light extending from the waveguides to the power meters.
8. Maximize the power to both detectors by manipulating the horizontal and vertical micrometer positions of the three fibers.
9. Further maximize the power to the detectors by making fine adjustments to the horizontal and vertical fiber positions using the piezo controllers.
10. Use the piezo controllers to move the fibers slightly closer to the chip. Make sure not to use the micrometers to push the fibers onto the chip, as doing so will likely damage the cleaved ends of the fibers.
11. Repeat steps 3.5.9 and 3.5.10 until the fibers are firmly pressed against the sides of the chip.

NOTE: Excessive scattered light from the waveguides, coupled with poor waveguide transmission, can be an indication of waveguide defects. These can include, but are not limited to, material defect sites, stitching boundaries, and excessive waveguide roughness.

6. Place a fiber-based polarization controller between the laser and the chip. This allows for control of the polarization state that makes it to the chip. The waveguides, being wider than they are tall, aid in minimizing any on-chip polarization rotation.
7. Dispersion characterization.
 1. Tune the polarization of the laser output to maximize the coupling to the chip. This device was designed for transverse-electric (TE) polarization and, as such, the transverse-magnetic (TM) polarization has a much higher loss.
 2. Scan a tunable laser over the wavelength range of interest (1,510 nm to 1,600 nm in this case) and monitor the power meters. Any artifacts in the transmission spectra are likely due to a combination of the TM component of the polarization and etalon effects from the fiber-chip interface.
 3. Locate the resonant wavelengths in the spectrum and also extract the bandwidth of each resonance. This particular chip had bandwidths as small as 65 pm, which translates to quality factors (Q) of up to 23,000.
 4. Determine the free spectral range (FSR), the separation between resonances, for each adjacent pair of resonances. This particular device had an FSR of ~5 nm.
 5. Calculate the group index (n_g) of the guided mode for each value of the FSR using the following equation:

$$n_g \cong \frac{\lambda^2}{FSR * 2\pi * r}$$

where λ is the wavelength and r is the radius of the micro-ring resonator. The above equation is a first-order approximation of the group index.

6. Use the bandwidth of each resonance to determine a width (Δn_g) associated with each value of the group index.
7. Choose the wavelengths for the two pump lasers such that they coincide with resonances in the spectrum and have an odd number of resonances between them (**Figure 1C**).
8. Determine the wavelength of the degenerate bi-photons using the following equation:

$$\lambda_{Bi-Photons} = 2 \left(\frac{1}{\lambda_{Pump\ 1}} + \frac{1}{\lambda_{Pump\ 2}} \right)^{-1}$$

where $\lambda_{Pump\ 1}$ and $\lambda_{Pump\ 2}$ are the wavelengths of the pump photons.

9. Add a horizontal line on the plot of group index versus wavelength that extends between the two pump wavelengths (**Figure 3**). If it is possible for the line to sit within $n_g \pm \Delta n_g$ at all three wavelengths of interest simultaneously, the phase-matching condition is satisfied and photons can be generated via SFWM. If it is not possible, try choosing pump wavelengths that are closer together and check again.
8. Add a second tunable laser source and polarization controller to the setup and combine the optical outputs from both lasers with a 1 x 2 fiber combiner.
9. Add a series of fiber-based notch filters (enough of them achieve ~120 dB of attenuation) immediately before the chip.

NOTE: The filters allow both pump wavelengths to pass but reject the bi-photon wavelength. They help to remove excess noise (*i.e.*, broadband Raman scattering in the optical fiber) before coupling to the chip. The filter spectrum is shown in **Figure 1B**.
10. Add a series of fiber-based bandpass filters (enough of them to achieve ~150 dB of attenuation) immediately after the chip.

NOTE: The filters should be wide enough to allow the bi-photons to pass but narrow enough to reject the pump photons. Two sets of these are needed, one set for each output. The filter spectrum is shown in **Figure 1B**.
11. Send the rejected photons from each set of filters to separate power meters.

NOTE: These power meters are used to monitor the optical coupling to the chip and can also be used to determine if the pump lasers are remaining on-resonance.
12. Connect the individual optical output from each set of fiber-based filters to a single photon detector (SPD) and connect both electrical signal outputs from the SPDs to a coincidence correlator.
13. Cross a pair of tungsten probes and set the tips down on one of the spiral legs (~1 mm in length) of the MZI.
14. Connect a power supply to the two crossed probes such that they generate heat when voltage is applied. This will act as the phase-shifter for the MZI.

NOTE: See the discussion for a description of the more standardized method for the thermal tuning of photonic devices.

4. Measuring Two-photon Interference

1. Tune both of the pump lasers to the chosen wavelengths. Use the power meters that are monitoring the rejected pump photons to ensure that both lasers are tuned to the resonances. When the lasers are properly tuned to the desired resonances, the rejected signal from the filters will be maximized.
2. Set the optical power output from each laser to -3 dBm.
NOTE: This will result in <100 μ W at the chip. It is important to keep the pump power this low in order to minimize loss (from multi-photon absorption and free-carrier absorption) and maintain stability (by minimizing light-induced thermal shifts). PN junctions can be used to remove carriers from the waveguide to better accommodate higher pump powers.
3. Monitor the coincidence counts (synchronous singles across two ports) by integrating over ~220 ps about the peak of the data. A sufficient integration time has passed when a minimum of 100 coincidence counts have been collected.
Note: The integration window should be wide enough to account for the timing jitter of the SPDs.
4. Set the power supply for the phase-shifter to the initial voltage (e.g., 0 V).
5. Scan one of the tunable lasers over the entire wavelength range and use the power meters that are collecting the rejected pump photons to confirm the location of the resonances of interest. Set the pump lasers to the wavelengths corresponding to the desired resonances.
Note: It is important to complete this step every time the phase shifter voltage is changed as the thermal tuning can result in small shifts in the resonant wavelengths.
6. Collect the resulting data (single-photon counts as well as the coincidence counts) from the coincidence correlator for the previously chosen integration time. Here, an integration time of 90 s was chosen with a timing resolution of 32 ps.
7. Increase the voltage applied to the phase-shifter by 5 mV.
8. Repeat steps 4.4-4.6 until data has been collected for the desired range of voltages.
NOTE: The maximum voltage was limited to 2.4 V due to rapid degradation of the probes above this voltage.
9. Integrate the coincidence peaks for each power-supply voltage over ~220 ps to determine the total number of coincidences (**Figure 4**).
10. Integrate over 320 ns away from the coincidence peak to obtain the accidental coincidences. Use this result to calculate the number of accidentals in the coincidence peak.
11. Fit the singles counts from each detector with the following modified sine function:

$$f(V) = A * \sin[(B * e^{C*V} + D) * (V + E)] + F$$

where A, B, C, D, E, and F are the fit parameters. This fit is necessary due to the nonlinear relationship between the voltage and the induced thermal shift (relative phase).

12. Convert the independent variable to relative phase for all three sets of data (the singles counts from each detector and the coincidence counts) with the following equation:

$$\theta(V) = (B * e^{C*V} + D) * (V + E)$$

where B, C, D, and E are the fit parameters from step 4.11. This conversion is possible due to the well-known sinusoidal transfer function of an MZI¹².

13. Fit the coincidence data (with the relative phase as the independent variable) with the following sine function:

$$f(\theta) = A * \sin(2 * \theta) + B$$

where A and B are the fit parameters.

14. Calculate the visibility of each interference pattern with the following equation:

$$V = \frac{f(\theta)_{max} - f(\theta)_{min}}{f(\theta)_{max} + f(\theta)_{min}}$$

where $f(\theta)_{max}$ and $f(\theta)_{min}$ are the maximum and minimum values of $f(\theta)$, respectively. A visibility of 1 corresponds to a perfect interference pattern.

Representative Results

Individual photon counts from each detector, as well as the coincidence counts, were collected as the relative phase between the two paths was tuned. The individual counts (**Figure 5A**) show the classical interference pattern from an MZI with visibilities of $94.5 \pm 1.6\%$ and $94.9 \pm 0.9\%$. The coincidence measurements (**Figure 5B**) show the quantum interference of the entangled state, as is evident by the oscillation at twice the frequency of the classical interference pattern, with a visibility of $93.3 \pm 2.0\%$ ($96.0 \pm 2.1\%$ with the accidentals subtracted). To confirm that the photons are primarily being generated in the ring, the pumps were configured into two resonances that would require the bi-photons to be generated at a wavelength not supported by the ring. The orange line in **Figure 5B** confirms that, with such a configuration, there are no significant coincidences. **Figure 6** shows the coincidence counts for the available resonance pairs that are symmetric in frequency about the resonance corresponding to the desired bi-photons. In all cases, the 2θ dependence of the relative phase is evident.

Figure 1: Experimental Testbed for the Silicon Waveguide Circuit. (a) Image of the silicon photonic quantum chip indicating the propagation direction of the photons. The inset is an energy conservation diagram for the four-wave mixing process that occurs within the ring. (b) Experimental setup used to test the silicon photonic circuit. (c) Transmission spectrum of the micro-ring cavity, with arrows indicating the pumping configuration as well as the wavelength of the generated bi-photons. [Please click here to view a larger version of this figure.](#)

Figure 2: Facet Improvement from Polishing. Images of the facet of a silicon photonic chip (a) after fabrication but before polishing and (b) after polishing. [Please click here to view a larger version of this figure.](#)

Figure 3: Characterization of Waveguide Dispersion. Plot of the wavelength dependence of the group index. The red-shaded region is representative of the bandwidth of the resonances and allows for the easy evaluation of the phase-matching condition. The green dashed line is horizontal and lies completely within the shaded region, demonstrating that the phase-matching condition is satisfied. The fact that the data is flat across the entire range is confirmation of zero dispersion. [Please click here to view a larger version of this figure.](#)

Figure 4: Measurement of Coincident Photons. Plot of the coincidence peak measured with the time-correlator with an integration time of 90 s and a timing resolution of 32 ps. The red dashed lines indicate the edges of the coincidence window, in which there are a total of 459 coincidences. [Please click here to view a larger version of this figure.](#)

Figure 5: Classical and Quantum Interference Measurements. (a) Classical light demonstrating the typical interference pattern from an MZI as the relative phase between the two paths is varied. (b) Coincidence correlation measurements showing the 2θ dependence of the relative phase. [Please click here to view a larger version of this figure.](#)

Figure 6: Bi-Photon Visibility for Various Pump Wavelength Configurations. Plots of the coincidences correlation measurements and calculated visibilities for the pumping wavelengths of (a) 1,513.5 nm and 1,592.0 nm, (b) 1,518.2 nm and 1,586.9 nm, (c) 1,522.9 nm and 1,581.8 nm, (d) 1,527.7 nm and 1,576.7 nm, (e) 1,532.4 nm and 1,571.6 nm, and (f) 1,537.2 nm and 1,566.6 nm. In all cases, the 2θ dependence of the relative phase is evident. [Please click here to view a larger version of this figure.](#)

SMF to UHNA Fiber Splicing Parameters	
Align: Core	Focus: Auto
ECF: Off	Auto Power: Off
Cleave Limit: 1°	Core Angle Limit: 1°
Cleaning Arc: 150 ms	Gap: 15 µm
Gapset Position: Center	Prefuse Power: 20 bit
Prefuse Time: 180 ms	Overlap: 10 µm
Arc1 Power: 20 bit	Arc1 Time: 18,000 ms
Arc2: Off	Rearc Time: 800 ms
Taper Splice: Off	

Table 1: Settings for Fusion Splicing SMF to UHNA Fiber.

Discussion

There are multiple challenges for the field of integrated photonics to overcome in order for complex and scalable systems of photonic devices to be feasible. These include, but are not limited to: tight fabrication tolerances, isolation from environmental instabilities, and minimization of all forms of loss. There are critical steps in the above protocol that help to minimize the loss of photonic devices.

One of the most crucial requirements in minimizing loss is closely matching the optical modes of the fibers and the waveguides. Part of the difficulty stems from the large mode field diameter (MFD) of SMF (~10 µm). On the integrated device side, there is a 500 nm-wide silicon waveguide with a much smaller MFD (<1 µm). This mode transition between the fiber and the waveguide can be improved in two ways: the SMF to a length of UHNA fiber or placing an inverse taper to the edge of the photonic chip. The spliced region between the SMF and UHNA fiber acts as a mode converter, reducing the size of the mode to ~3 µm. The inverse taper is used to expand the mode on the chip by reducing the width of the waveguide as it approaches the facet. This chip uses a linear taper from the 500 nm waveguides to the 150 nm tips (at the chip facets), with a transition length of 300 µm. The tapering of the waveguide width to the edge of the chip results in a decrease in the effective index of the optical mode and, in turn, the mode expands.

The polishing of the chip facet is also very important in mitigating optical loss. Two concerns while polishing are stopping at the desired surface and delaminating the upper cladding material. Ideally, the final position of the facet would be precisely at the end of the taper. However, this is quite difficult to achieve, and for that reason, the tip of the taper is extended by 100 µm so that the polishing can be stopped a few microns before the taper starts. If too little material is removed, the mode will not be captured as efficiently by the taper. If too much material is removed, there will be a larger mode mismatch at the fiber/chip interface, and more of the light will be lost. The other main concern is the delamination of the upper cladding. If there are problems with fabrication (cleanliness or excessive stress in the cladding), the cladding may not adhere to the substrate at the edge of the chip. When delamination happens to occur at one of the waveguides, it will result in very poor coupling efficiencies. If it is noticed during polishing, a polishing lubricant other than water can often improve results.

There is room for improvement in the above protocol. The largest improvement would come from using a more standard method for thermally tuning the device. The method used here was a result of a simplified fabrication process that did not include any metal layers. Typically, a resistive metal layer is used for heater elements, and a highly-conductive metal layer is used for contact pads and wires from the pads to the heater elements. A stage can then be used to set down probes onto the pads, allowing a voltage to be applied to the heaters. This enables a greater level of control and stability. A silicon photonic chip similar to what was tested here but with metal heaters is shown in the accompanying video.

There are other methods of coupling light to the photonic chip. For this work, edge coupling was used. Other common methods include free space coupling and grating coupling. Free space coupling relies upon bulk optical elements to align and focus the beam into the waveguide at the edge of the chip. The drawback with coupling in this manner is that it can be very difficult to optimize the alignment of the beam, and there will always be a reflection at the interface due to the index difference. Grating couplers scatter the light from the waveguide vertically, so that the end of a fiber can be placed at the surface of the cladding to couple to the device. These also have some issues, including difficult alignment (the fiber is often in the line of sight of the microscope) and higher losses. Fiber edge coupling is not perfect either. Pressing the fibers against the chip can damage the ends of the fibers, and both the fibers and the chip edge need to be cleaned frequently. The benefit of fiber edge coupling is that the alignment is far easier than the other two methods and is capable of achieving lower losses.

As the complexity of optical systems increases, the only feasible way for them to scale into a stable platform is in an integrated system, much like the path of electronic technology. The challenge is merging the integrated photonics platform with the bulk and fiber-based optical systems that are already deployed. With the utilization of photon-based quantum information systems, where the information space scales exponentially (compared to the linear scaling of classical systems), phase stability and low-loss integrated photonic technologies are paramount for success. The protocol we have described serves as an initial path forward for advancing this emerging technology.

Disclosures

We have nothing to disclose.

Acknowledgements

This work was performed in part at the Cornell University NanoScale Science and Technology Facility, a member of the National Nanotechnology Infrastructure Network, which is supported by the National Science Foundation (Grant ECCS-1542081). We acknowledge support for this work from the Air Force Research Lab (AFRL). This material is based upon work partially supported by the National Science Foundation under Award No. ECCS14052481.

References

1. Silverstone, J. W., *et al.* On-chip quantum interference between silicon photon-pair sources. *Nat. Photon.* **8** (2), 104-108at (2014).
2. Harris, N. C., *et al.* Integrated Source of Spectrally Filtered Correlated Photons for Large-Scale Quantum Photonic Systems. *Phys. Rev. X.* **041047**, 1-10 (2014).
3. Grassani, D., *et al.* Micrometer-scale integrated silicon source of time-energy entangled photons. *Optica.* **2** (2), 88 (2015).
4. Najafi, F., *et al.* Scalable Integration of Single-Photon Detectors. *Nat. Commun.* **6**, 1-8 (2015).
5. Dutt, A., *et al.* On-Chip Optical Squeezing. *Phys. Rev. Appl.* **3** (4), 1-7 (2015).
6. Azzini, S., *et al.* Ultra-low power generation of twin photons in a compact silicon ring resonator. *Opt. Express.* **20** (21), 23100-23107 (2012).
7. Clemmen, S., *et al.* Continuous wave photon pair generation in silicon-on-insulator waveguides and ring resonators : erratum. *Opt. Express.* **17** (19), 16558 (2009).
8. Engin, E., *et al.* Photon pair generation in a silicon micro-ring resonator with reverse bias enhancement. *Opt. Express.* **21** (23), 27826-27834 (2013).
9. Steidle, J. a., *et al.* High spectral purity silicon ring resonator photon-pair source. *Proc. of SPIE.* **9500**, 950015 (2015).
10. Preble, S. F., *et al.* On-Chip Quantum Interference from a Single Silicon Ring-Resonator Source. *Phys. Rev. Appl.* **4**, 021001 (2015).
11. Cao, L., Aboketaf, A. A., & Preble, S. F. CMOS compatible micro-oven heater for efficient thermal control of silicon photonic devices. *Opt. Commun.* **305**, 66-70 (2013).
12. Chrostowski, L., & Hochberg, M. *Silicon Photonics Design.*, Section 4.3 Cambridge University Press (2013).