High-resolution aliasing-free optical beam steering: supplementary material

DAVID N. HUTCHISON, JIE SUN, JONATHAN K. DOYLEND, RANJEET KUMAR, JOHN HECK, WOOSUNG KIM, CHRISTOPHER T. PHARE, AVI FESHALI, AND HAISHENG RONG*

Intel Labs, Intel Corporation, 2200 Mission College Blvd., Santa Clara, California 95054, USA

*Corresponding author: haisheng.rong@intel.com

Published 5 August 2016

This document provides supplementary information to “High-resolution aliasing-free optical beam steering,” <http://dx.doi.org/10.1364/optica.3.000887>. We discuss the emitter placement, how we condense the beam and calibrate the measured image, the measured beam divergence, the measured and theoretical wavelength tuning efficiency, and the channel power distribution used in our simulations. © 2016 Optical Society of America

<http://dx.doi.org/10.1364/optica.3.000887.s001>

1. Emitters

The 1/e decay length for the emitters is both measured (by observing a decaying power profile along the emitter in a top-view infrared microscope image) and is calculated to be about 90 μm . The emitters are longer than 1 mm which is more than enough to ensure practically all the light couples out along its length.

The uniformly-spaced OPA was designed with a spacing of $d_0 = 7.246 \mu\text{m}$, which has a lobe spacing at wavelength λ of $\sin \theta = \lambda/d_0$, resulting in about $\theta = 10^\circ$ near normal, near $\lambda = 1.3 \mu\text{m}$. For the non-uniform OPA we chose pseudo-random emitter positions with an average spacing of the same d_0 . Its emitter positions were uniformly distributed random numbers – through an iterative procedure, different emitter positions were simulated with various standard deviations until a figure of merit (< -10 dB side-lobe power at $\pm 45^\circ$ beam deflection) was satisfied. As a result, a particular set of emitter positions with standard deviation 1.1 μm was chosen. The standard deviation cannot be too large, because the closest two waveguides in the final set need to have a spacing of a few wavelengths in order to mitigate on-chip cross-talk between emitters. The pitch for the closest two in our chosen set was 5.4 μm .

2. Condensing the beam

The emitter phases are initially arbitrary after passing through the star coupler, S-bends, and other components, so we observe an arbitrary pattern of constructive and destructive interference in

the far field along the elliptical line defining the ψ axis (Fig. 3(b)). To align the phases and therefore coalesce the beam we use a gradient-search algorithm. First, each channel is given a random starting phase and a reference camera image is taken. Next, one-by-one, each channel takes a step in phase, a camera image is captured, and then that channel is returned to its original phase. Each of these camera images yields a figure of merit which is defined as the power in the desired angle (in the vicinity of the desired pixel position) divided by the power at other angles. After sampling the gradient in all 128 channels in this manner, each figure of merit is compared to the figure of merit from the original reference image to determine the overall direction of the gradient in 128-dimensional space, a step is taken in the direction of the gradient, and then the cycle repeats from the beginning. The optimization is therefore limited by the camera's capture speed of 10 Hz.

After several such iterations (usually fewer than 10) the beam has condensed to the desired angle and the final phases are stored in a lookup table. By recalling the phases from the lookup table in sequence, the beam can be steered from one position to another.

In our design, since the light in each channel travels a different distance between the phase tuners and the emitters, if the beam is condensed and then the wavelength changed by more than a few nanometers, the phases will no longer be aligned and the beam no longer well condensed. That means the phases must be aligned at each desired wavelength. However this limitation only affects how long it takes to set up the lookup table. Once the phases and wavelength are stored in the lookup table this does not impact how quickly the condensed beam can be recalled.

3. Far-field image calibration

As described in the Letter, after adjusting the lenses and cameras to capture the desired field-of-view, we first calibrate the far-field image precisely by measuring an optical phased array (OPA) with uniformly-spaced emitters and comparing the measured pixel positions to the known lobe angles from simulation. Taking the image in Fig. 3(d) as an example, we plot its calibration curve in Fig. S1. The vertical axis is the sine of the simulated lobe angles rather than just the angle, since we are projecting the far-field on to a plane. A very linear relationship is observed (R^2 of the straight-line fit is 0.99997). After the camera image has been calibrated in this way we can measure other devices, such as the non-uniformly spaced OPA.

Fig. S1. An example camera image calibration, taken from the image in Fig. 3(d).

4. Divergence in ψ and θ

The extracted beam divergence in both axes for the 75 measurements in Fig. 3(e) is shown in Fig. S2.

Fig. S2. Measurements of beam divergence in both axes. (a) Cross-sections in ψ of the condensed beam, showing 75 separate measurements of the beam steered to different angles across about 10° . (b)-(c) The beam widths from those measurements in the θ and ψ axes respectively.

As discussed in the letter, the divergence in θ is $0.142 \pm 0.005^\circ$ (the uncertainty value is the standard deviation). For comparison, this is about 2.5 milliradians, approaching the typical beam divergence of a laser pointer. This divergence is, to our knowledge, the smallest ever reported for a silicon emitter.

The divergence in ψ is $0.14 \pm 0.02^\circ$. The smallest divergence in ψ was 0.11° , matching the expected beam divergence from simulation. Next we expand upon the reasons for the slightly wider-than-optimal average beam width in ψ . This is probably due to a number of factors, for example:

(A) About 10% of the channels were not connected in that measurement due to initial difficulties with landing such a large number of electrical probes. This also leads to some side-lobe (background) power as can be seen in Fig. 3(e). In later measurements such as those shown Fig. 3(b)-(d) and Fig. 5(a), we improved our testing and achieve 100% connection.

(B) The algorithm is searching for a global maximum in 128-dimensional space – since this is a large number of dimensions it is overwhelmingly more probable that it will find a local maximum rather than the global maximum, resulting in imperfect phase alignment.

(C) Thermal crosstalk may make perfect phase alignment difficult or impossible in several ways. Firstly, a step in one dimension (the phase in one channel) is not necessarily independent of the other dimensions (the phase in adjacent channels). Secondly, it reduces the availability of possible phase combinations – consider for example the case where the i -th channel has a large phase offset (near 2π) and the $(i+1)$ -th channel is set to a small phase offset (near zero). In that case thermal cross-talk will set a limit on the minimum phase in the $(i+1)$ -th channel. Adding 2π to the $(i+1)$ -th channel is possible, but difficult because that introduces its own cross-talk (for example, consider the effect on the $(i+2)$ -th channel).

We have seen that the beam width in ψ is not always at its theoretical minimum due to a number of factors, so the average divergence of a large number of experiments may be the most realistic number to use when assessing future high-resolution beam-steering systems. The beam width also depends on the steering angle, but that is a small effect for the small range of angles we consider in Fig. S2.

5. Side-lobe detail

When the main lobe is well-condensed and not saturated on the camera, there is typically a significant part of the side-lobe power outside of the camera's dynamic range. This usually isn't important – it is more important to know the ratio of the non-saturated main lobe to the *tallest* side-lobe (sometimes known as side-lobe level, SLL) rather than to see many side-lobes. The tallest side-lobe is always within the camera's dynamic range. However, to see more of the side-lobe profile, we may simply increase the laser power. In Fig. S3 we first condense the beam as in Fig. 5, then take curved cross-sections of the image at successively higher laser power.

Fig. S3. Image cross-sections as a function of laser power. (a) Typical curved cross-section showing main lobe and side-lobes for

a non-uniform OPA. (b) One of the largest side-lobes. (c) The main lobe. The legend applies to (a)-(c). (d) The main lobe (red, saturated for most of the measurements) and eight tallest side-lobes (other colors) as a function of laser power.

Fig. S3(a) shows the beam from our non-uniform OPA, condensed at $\psi = 0^\circ$, and (b), (c) show a side-lobe and the main lobe respectively. Part (a) just shows three laser power levels for clarity while (b)-(c) show all measured power levels, but they show the same dataset. In (a)-(c), pixels reporting values below the pixel sensitivity are reported as zero by our camera – thus those lines are drawn to negative infinity on this log scale.

In Fig. S3 (d) we extract the height of the main lobe and eight tallest side-lobes versus laser power. It can be seen that the SLL is about 8.6 dB and the worst from the five datasets in Fig. 5 is 8.9 dB. The exact SLL varies each time the beam is condensed and these are only typical numbers. These SLL numbers are similar to other groups' devices, and we do not claim the best SLL ever achieved. Instead, our goal was extremely high steering range and small beam size, rather than optimizing our design for the lowest SLL.

As expected, the side-lobes in (d) scale linearly with power, but the lines are not perfectly straight. We attribute this to a number of factors. As can be seen from (b)-(c), the lobes are typically only a couple of pixels wide for our widest FOV. There is some 'dead' space between each pixel and it is possible that the far-field features move slightly between measurements. In general Fig. S3 highlights the dual difficulties of simultaneously measuring a very small beam and a large FOV, and of simultaneously measuring a tall main lobe and small side-lobe features.

6. Tuning efficiency in θ

From the grating equation [1]

$$\sin \theta = n_{eff} - \frac{N\lambda}{\Lambda} \quad (S1)$$

where n_{eff} is the effective index simulated by finite-element method (Fig. S4(a)), Λ is the grating pitch, and N is a non-negative integer (unity, in our case), we get the theoretical emission angle θ (Fig. S4(b)). Since the emitter consists of repeating notches in the top surface of a waveguide, the cross-section varies along its length and so too does its effective index. However the average of the effective index with and without the notch is a reasonable substitute for the emitter's effective index.

Fig. S4. Simulated wavelength tuning. (a) The fundamental TE mode of the waveguide. The scale bar is 500 nm. (b) The waveguide effective index, simulated using Finite Element Method. (c) The expected outcoupling angle. (d) The tuning efficiency.

By taking the derivative $d\theta/d\lambda$ we calculate the tuning efficiency to be $0.159^\circ/\text{nm}$ at a wavelength of $1.3\ \mu\text{m}$ (Fig. S4(c)). This is in fairly good agreement with the measured value of $0.17^\circ/\text{nm}$ from sets of images like Fig. 4(d). The discrepancy probably comes from a slightly different actual waveguide cross-section and therefore effective index compared to the idealized simulated geometry (for example slightly sloping sidewalls or slightly different thickness, resulting in a different effective index).

7. Star coupler and channel power distribution

The star coupler (Fig. 1(b) and Fig. S5) begins as a $1\ \mu\text{m}$ wide rib waveguide that is terminated to launch the light into the $0.4\ \mu\text{m}$ -thick slab. Radially $1.24\ \text{mm}$ away from the tip of the launch waveguide, we begin the 128 output waveguides. Each output waveguide consists of a simple linear taper that funnels down, ending in our standard single-mode waveguide (dimensions are given in the main text).

Fig. S5. Scanning electron microscope (SEM) images of the star coupler. (a) Zoomed out. The circle is the SEM aperture, the launch waveguide is on the left and the 128 output waveguides are on the right. (b)-(d) Successively higher magnification SEM images.

The light from the launch waveguide can be expected to spread out in the slab as approximately as a Gaussian, thus the optical power in the output channels is quasi-Gaussian. Relative to the star coupler axis, the outermost outputs are at $\pm 15.3^\circ$ in the slab, so designed in order to get approximately 1% of the power in the outermost outputs relative to the central waveguide, thus capturing most of the launched power.

For our simulations to be accurate we needed to know what power distribution we had across the channels, but this is difficult to measure directly on a full device since the light is confined in the waveguides and there are no power monitors on chip. To measure this power distribution we tested 12 standalone star couplers across a $300\ \text{mm}$ wafer and recorded the transmission through the odd-numbered channels (Fig. S6(a)). The star coupler has a mean insertion loss of about $0.2\ \text{dB}$ (calculated by summing the average measured output powers, multiplying by 2 because only the odd-numbered channels were measured, and accounting for the grating coupler and waveguide losses). The phase tuners are not

very efficient ($80 \text{ mW}/\pi$) and have not yet been optimized; the edge of the doped heating strip is $2 \mu\text{m}$ from the waveguide and thermal isolation could be improved. Thermo-optic and other phase tuners are well-studied in the literature and a more efficient phase-tuner design would improve cross-talk.

Since we observe a low insertion loss but a bumpy profile across the channels, we attribute this irregularity to multimode interference rather than poor coupling to the output waveguides. Multimode interference may manifest itself if multiple slab modes are excited at the launch point and interfere at the outputs. We believe this can be improved by controlling the field we launch into the slab, nevertheless simulations show this irregularity doesn't greatly affect the expected far-field profile. As demonstrated in Fig. S6(b), if instead of the measured mean power profile we use a smooth Gaussian power profile with approximately the same width as the bumpy measured profile, the far-field simulation result is not significantly modified. This makes sense because the far-field is the Fourier transform of the near-field, so the lobe divergence should depend more on the width of the collection of emitters rather than fluctuations in the individual emitters' amplitude.

Fig. S6. Star coupler channel power distribution characterization. (a) The measured insertion loss of each channel (IL; black) relative to a single-input single-output reference device. The box plots indicate the statistics for 12 measured star couplers across a 300 mm wafer. Only the odd-numbered channels were measured, and the even-numbered channels were assumed by symmetry. We connect the means (green) and also show a Gaussian with approximately the same width (blue). (b) The far-field simulated in both cases. This demonstrates that the irregularity in the star coupler channel power does not significantly affect the resultant far field, which is mostly dependent on the general width of the entire emitting aperture rather than on the details of the power distribution in each channel.

References

1. R. Waldhäusl, B. Schnabel, P. Dannberg, E. B. Kley, A. Bräuer, W. Karthe, "Efficient coupling into polymer waveguides by gratings," *Appl. Opt.* **36**, 9383-9390 (1997).