

PHOTONIC PROCESSING SYSTEMS AND METHODS**CROSS-REFERENCE TO RELATED APPLICATIONS**

This application claims priority under 35 U.S.C. § 119(e) to U.S. Provisional Patent Application Serial No. 62/671,793, entitled "ALGORITHMS FOR TRAINING NEURAL NETWORKS WITH PHOTONIC HARDWARE ACCELERATORS," filed on May 15, 2018 under Attorney Docket No. L0858.70001US00, which is hereby incorporated herein by reference in its entirety.

This application also claims priority under 35 U.S.C. § 119(e) to U.S. Provisional Patent Application Serial No. 62/680,557, entitled "PHOTONICS PROCESSING SYSTEMS AND METHODS," filed on June 4, 2018 under Attorney Docket No. L0858.70000US00, which is hereby incorporated herein by reference in its entirety.

This application also claims priority under 35 U.S.C. § 119(e) to U.S. Provisional Patent Application Serial No. 62/689,022, entitled "CONVOLUTIONAL LAYERS FOR NEURAL NETWORKS USING PROGRAMMABLE NANOPHOTONICS," filed on June 22, 2018 under Attorney Docket No. L0858.70003US00, which is hereby incorporated herein by reference in its entirety.

This application also claims priority under 35 U.S.C. § 119(e) to U.S. Provisional Patent Application Serial No. 62/755,402, entitled "REAL-NUMBER PHOTONIC ENCODING," filed on November 11, 2018 under Attorney Docket No. L0858.70008US00, which is hereby incorporated herein by reference in its entirety.

This application also claims priority under 35 U.S.C. § 119(e) to U.S. Provisional Patent Application Serial No. 62/792,720, entitled "HIGH-EFFICIENCY DOUBLE-SLOT WAVEGUIDE NANO-OPTOELECTROMECHANICAL PHASE MODULATOR ," filed on January 15, 2019 under Attorney Docket No. L0858.70006US00, which is hereby incorporated herein by reference in its entirety.

This application also claims priority under 35 U.S.C. § 119(e) to U.S. Provisional Patent Application Serial No. 62/793,327, entitled "DIFFERENTIAL, LOW-NOISE HOMODYNE RECEIVER," filed on January 16, 2019 under Attorney Docket No. L0858.70004US00, which is hereby incorporated herein by reference in its entirety.

This application also claims priority under 35 U.S.C. § 119(e) to U.S. Provisional Patent Application Serial No. 62/834,743 entitled "STABILIZING LOCAL OSCILLATOR

PHASES IN A PHOTOCORE," filed on April 16, 2019 under Attorney Docket No. L0858.70014US00, which is hereby incorporated herein by reference in its entirety.

BACKGROUND

Conventional computation use processors that include circuits of millions of transistors to implement logical gates on bits of information represented by electrical signals. The architectures of conventional central processing units (CPUs) are designed for general purpose computing, but are not optimized for particular types of algorithms. Graphics processing, artificial intelligence, neural networks, and deep learning are a few examples of the types of algorithms that are computationally intensive and are not efficiently performed using a CPU.

Consequently, specialized processors have been developed with architectures better-suited for particular algorithms. Graphical processing units (GPUs), for example, have a highly parallel architecture that makes them more efficient than CPUs for performing image processing and graphical manipulations. After their development for graphics processing, GPUs were also found to be more efficient than CPUs for other memory-intensive algorithms, such as neural networks and deep learning. This realization, and the increasing popularity of artificial intelligence and deep learning, lead to further research into new electrical circuit architectures that could further enhance the speed of these algorithms.

SUMMARY

In some embodiments, a photonic processor is provided. The photonic processor may include a first array of interconnected variable beam splitters (VBSs) comprising a first plurality of optical inputs and a first plurality of optical outputs; a second array of interconnected VBSs comprising a second plurality of optical inputs and a second plurality of optical outputs; and a plurality of controllable optical elements, each of the plurality of these controllable optical elements coupling a single one of the first plurality of optical outputs of the first array to a respective single one of the second plurality of optical inputs of the second array.

In some embodiments, a photonic processing system is provided. The photonic processing system may include an optical encoder configured to encode an input vector into a first plurality of optical signals. The photonic processing system may also include a photonic processor configured to: receive the first plurality of optical signals, each of the first plurality

of signals received by a respective input spatial mode of a plurality of input spatial modes of the photonic processor; perform a plurality of operations on the first plurality of optical signals, the plurality of operations implementing a matrix multiplication of the input vector by a matrix; and output a second plurality of optical signals representing an output vector, each of the second plurality of signals transmitted by a respective output spatial mode of a plurality of output spatial modes of the photonic processor. The photonic processing system may also include an optical receiver configured to detect the second plurality of optical signals and output an electrical digital representation of the output vector.

In some embodiments, a method of optically performing matrix-vector multiplication is provided. The method may include: receiving a digital representation of an input vector; encoding, using an optical encoder, the input vector into a first plurality of optical signals; performing, using a processor, a singular value decomposition (SVD) of a matrix to determine a first, second, and third SVD matrix; controlling photonic processor comprising a plurality of variable beam splitters (VBS) to optically implement the first, second, and third SVD matrix; propagating the first plurality of optical signals through the photonic processor; detecting a second plurality of optical signals received from the photonic processor; and determining an output vector based on the detected second plurality of optical signals, wherein the output vector represents a result of the matrix-vector multiplication.

The foregoing apparatus and method embodiments may be implemented with any suitable combination of aspects, features, and acts described above or in further detail below. These and other aspects, embodiments, and features of the present teachings can be more fully understood from the following description in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Various aspects and embodiments of the application will be described with reference to the following figures. It should be appreciated that the figures are not necessarily drawn to scale. Items appearing in multiple figures are indicated by the same reference number in all the figures in which they appear.

FIG. 1-1 is a schematic diagram of a photonic processing system, in accordance with some non-limiting embodiments.

FIG. 1-2 is a schematic diagram of an optical encoder, in accordance with some non-limiting embodiments.

FIG. 1-3 is a schematic diagram of a photonic processor, in accordance with some non-limiting embodiments.

FIG. 1-4 is a schematic diagram of an interconnected variable beam splitter array, in accordance with some non-limiting embodiments.

FIG. 1-5 is a schematic diagram of a variable beam splitter, in accordance with some non-limiting embodiments.

FIG. 1-6 is a schematic diagram of a diagonal attenuation and phase shifting implementation, in accordance with some non-limiting embodiments.

FIG. 1-7 is a schematic diagram of an attenuator, in accordance with some non-limiting embodiments.

FIG. 1-8 is a schematic diagram of a power tree, in accordance with some non-limiting embodiments.

FIG. 1-9 is a schematic diagram of an optical receiver, in accordance with some non-limiting embodiments.

FIG. 1-10 is a schematic diagram of a homodyne detector, in accordance with some non-limiting embodiments.

FIG. 1-11 is a schematic diagram of a folded photonic processing system, in accordance with some non-limiting embodiments.

FIG. 1-12A is a schematic diagram of a wavelength-division-multiplexed (WDM) photonic processing system, in accordance with some non-limiting embodiments.

FIG. 1-12B is a schematic diagram of the frontend of the wavelength-division-multiplexed (WDM) photonic processing system of FIG. 1-12A, in accordance with some non-limiting embodiments.

FIG. 1-12C is a schematic diagram of the backend of the wavelength-division-multiplexed (WDM) photonic processing system of FIG. 1-12A, in accordance with some non-limiting embodiments.

FIG. 1-13 is a schematic diagram of a circuit for performing analog summation of optical signals, in accordance with some non-limiting embodiments.

FIG. 1-14 is a schematic diagram of a photonic processing system with column-global phases shown, in accordance with some non-limiting embodiments.

FIG. 1-15 is a plot showing the effects of uncorrected global phase shifts on homodyne detection, in accordance with some non-limiting embodiments.

FIG. 1-16 is a plot showing the quadrature uncertainties of coherent states of light, in accordance with some non-limiting embodiments.

FIG 1-17 is an illustration of matrix multiplication, in accordance with some non-limiting embodiments.

FIG. 1-18 is an illustration of performing matrix multiplication by subdividing matrices into sub-matrices, in accordance with some non-limiting embodiments.

FIG. 1-19 is a flowchart of a method of manufacturing a photonic processing system, in accordance with some non-limiting embodiments.

FIG. 1-20 is a flowchart of a method of manufacturing a photonic processor, in accordance with some non-limiting embodiments.

FIG. 1-21 is a flowchart of a method of performing an optical computation, in accordance with some non-limiting embodiments.

FIG. 2-1 is a flow chart of a process for training a latent variable model, in accordance with some non-limiting embodiments.

FIG. 2-2 is a flow chart of a process for configuring a photonics processing system to implement unitary transfer matrices, in accordance with some non-limiting embodiments.

FIG. 2-3 is a flow chart of a process for computing an error vector using a photonics processing system, in accordance with some non-limiting embodiments.

FIG. 2-4 is a flow chart of a process for determining updated parameters for unitary transfer matrices, in accordance with some non-limiting embodiments.

FIG. 2-5 is a flow chart of a process for updating parameters for unitary transfer matrices, in accordance with some non-limiting embodiments.

FIG. 3-1 is a flowchart of a method for computing a forward pass through a convolutional layer, in accordance with some non-limiting embodiments.

FIG. 3-2 is a flowchart of a method for computing a forward pass through a convolutional layer, in accordance with some non-limiting embodiments.

FIG. 3-3A is a flowchart of a method suitable for computing two-dimensional convolutions, in accordance with some non-limiting embodiments.

FIG. 3-3B is a flowchart of a method suitable for building a circulant matrix, in accordance with some non-limiting embodiments.

FIG. 3-4A illustrates a pre-processing step of building a filter matrix from input filter matrices including a plurality of output channels, in accordance with some non-limiting embodiments.

FIG. 3-4B illustrates building a circulant matrix from input matrices including a plurality of input channels, in accordance with some non-limiting embodiments.

FIG. 3-4C illustrates a two-dimensional matrix multiplication operation, in accordance with some non-limiting embodiments.

FIG. 3-4D illustrates a post-processing step of rotating vector rows, in accordance with some non-limiting embodiments.

FIG. 3-4E illustrates a post-processing step of vector row addition, in accordance with some non-limiting embodiments.

FIG. 3-4F illustrates reshaping an output matrix into multiple output channels, in accordance with some non-limiting embodiments.

FIG. 3-5 is a flowchart of a method for performing a one-dimensional Fourier transform, in accordance with some non-limiting embodiments.

FIG. 3-6 is a flowchart of a method for performing a two-dimensional Fourier transform, in accordance with some non-limiting embodiments.

FIG. 3-7 is a flowchart of a method for performing a two-dimensional Fourier transform, in accordance with some non-limiting embodiments.

FIG. 3-8 is a flowchart of a method for performing convolutions using Fourier transforms, in accordance with some non-limiting embodiments.

FIG. 4-1 is a block diagram illustrating an optical modulator, in accordance with some non-limiting embodiments.

FIG. 4-2A is a block diagram illustrating an example of a photonic system, in accordance with some non-limiting embodiments.

FIG. 4-2B is a flowchart illustrating an example of a method for processing signed, real numbers in the optical domain, in accordance with some non-limiting embodiments.

FIG. 4-3A is schematic diagram illustrating an example of a modulator that may be used in connection with the photonic system of FIG. 4-2A, in accordance with some non-limiting embodiments.

FIG. 4-3B is a plot illustrating the intensity and phase spectral responses of the modulator of FIG. 4-3A when no voltage is applied, in accordance with some non-limiting embodiments.

FIG. 4-3C is a plot illustrating the intensity and phase spectral responses of the modulator of FIG. 4-3A when driven at a certain voltage, in accordance with some non-limiting embodiments.

FIG. 4-3D is an example of an encoding table associated with the modulator of FIG. 4-3A, in accordance with some non-limiting embodiments.

FIG. 4-3E is a visual representation, in the complex plane, of the spectral response of the modulator of FIG. 4-3A, in accordance with some non-limiting embodiments.

FIGs. 4-4A through 4-4C are visual representations, in the complex plane, of different spectral responses at the output of an optical transformation unit, in accordance with some non-limiting embodiments.

FIG. 4-5 is a visual representation, in the complex plane, of how detection of an optical signal may be performed, in accordance with some non-limiting embodiments.

FIG. 4-6 is an example of a decoding table, in accordance with some non-limiting embodiments.

FIG. 4-7 is a block diagram of an optical communication system, in accordance with some non-limiting embodiments.

FIG. 4-8 is a plot illustrating an example of a power spectral density output by the modulator of FIG. 4-7, in accordance with some non-limiting embodiments.

FIG. 4-9 is a flowchart illustrating an example of a method for fabricating a photonic system, in accordance with some non-limiting embodiments.

FIG. 5-1 is a circuit diagram illustrating an example of a differential optical receiver, in accordance with some non-limiting embodiments.

FIG. 5-2 is a schematic diagram illustrating a photonic circuit that may be coupled with the differential optical receiver of FIG. 5-1, in accordance with some non-limiting embodiments.

FIG. 5-3A is a schematic diagram illustrating a substrate including a photonic circuit, photodetectors and a differential operational amplifier, in accordance with some non-limiting embodiments.

FIG. 5-3B is a schematic diagram illustrating a first substrate including a photonic circuit and photodetectors, and a second substrate including a differential operational amplifier, where the first and second substrates are flip-chip bonded to each other, in accordance with some non-limiting embodiments.

FIG. 5-3C is a schematic diagram illustrating a first substrate including a photonic circuit and photodetectors, and a second substrate including a differential operational amplifier, where the first and second substrates are wire bonded to each other, in accordance with some non-limiting embodiments.

FIG. 5-4 is a flowchart illustrating an example of a method for fabricating an optical receiver, in accordance with some non-limiting embodiments.

FIGs. 5-4A through 5-4F illustrate an example of a fabrication sequence for an optical receiver, in accordance with some non-limiting embodiments.

FIG. 5-5 is a flowchart illustrating an example of a method for receiving an optical signal, in accordance with some non-limiting embodiments.

FIG. 6-1A is a top view illustrating schematically a Nano-Opto-Electromechanical Systems (NOEMS) phase modulator, in accordance with some non-limiting embodiments.

FIG. 6-1B is a top view illustrating schematically a suspended multi-slot optical structure of the NOEMS phase modulator of FIG. 6-1A, in accordance with some non-limiting embodiments.

FIG. 6-1C is a plot illustrating an example of an optical mode arising in the suspended multi-slot optical structure of FIG. 6-1B, in accordance with some non-limiting embodiments.

FIG. 6-1D is a top view illustrating schematically a mechanical structure of the NOEMS phase modulator of FIG. 6-1A, in accordance with some non-limiting embodiments.

FIG. 6-1E is a top view illustrating schematically a transition region of the NOEMS phase modulator of FIG. 6-1A, in accordance with some non-limiting embodiments.

FIG. 6-2 is a cross-sectional view of the NOEMS phase modulator of FIG. 6-1A, taken in a yz-plane, and illustrating a suspended waveguide, in accordance with some non-limiting embodiments.

FIG. 6-3 is a cross-sectional view of the NOEMS phase modulator of FIG. 6-1A, taken in a xy-plane, and illustrating a portion of a suspended multi-slot optical structure, in accordance with some non-limiting embodiments.

FIGs. 6-4A through 6-4C are cross-sectional views illustrating how a suspended multi-slot optical structure can be mechanically driven to vary the widths of the slots between the waveguides, in accordance with some non-limiting embodiments.

FIG. 6-5 is a plot illustrating how the effective index of a suspended multi-slot optical structure may vary as a function of the width of a slot, in accordance with some non-limiting embodiments.

FIG. 6-6 is a flowchart illustrating an example of a method for fabricating a NOEMS phase modulator, in accordance with some non-limiting embodiments.

DETAILED DESCRIPTION

I. Photo-Core

A. Overview of Photonics-based Processing

The inventors have recognized and appreciated that there are limitations to the speed and efficiency of conventional processors based on electrical circuits. Every wire and transistor in the circuits of an electrical processor has a resistance, an inductance, and a capacitance that cause propagation delay and power dissipation in any electrical signal. For example, connecting multiple processor cores and/or connecting a processor core to a memory uses a conductive trace with a non-zero impedance. Large values of impedance limit the maximum rate at which data can be transferred through the trace with a negligible bit error rate. In applications where time delay is crucial, such as high frequency stock trading, even a delay of a few hundredths of a second can make an algorithm unfeasible for use. For processing that requires billions of operations by billions of transistors, these delays add up to a significant loss of time. In addition to electrical circuits' inefficiencies in speed, the heat generated by the dissipation of energy caused by the impedance of the circuits is also a barrier in developing electrical processors. The inventors further recognized and appreciated that using light signals, instead of electrical signals, overcomes many of the aforementioned problems with electrical computing. Light signals travel at the speed of light in the medium in which the light is traveling; thus the latency of photonic signals is far less of a limitation than electrical propagation delay. Additionally, no power is dissipated by increasing the distance traveled by the light signals, opening up new topologies and processor layouts that would not be feasible using electrical signals. Thus, light-based processors, such as a photonics-based processor may have better speed and efficiency performance than conventional electrical processors.

Additionally, the inventors have recognized and appreciated that a light-based processor, such as a photonics-based processor, may be well-suited for particular types of algorithms. For example, many machine learning algorithms, e.g. support vector machines, artificial neural networks, probabilistic graphical model learning, rely heavily on linear transformations on multi-dimensional arrays/tensors. The simplest example is multiplying vectors by matrices, which using conventional algorithms has a complexity on the order of $O(n^2)$, where n is the dimensionality of the square matrices being multiplied. The inventors

have recognized and appreciated that a photonics-based processor, which in some embodiment may be a highly parallel linear processor, can perform linear transformations, such as matrix multiplication, in a highly parallel manner by propagating a particular set of input light signals through a configurable array of beam splitters. Using such implementations, matrix multiplication of matrices with dimension $n=512$ can be completed in hundreds of picoseconds, as opposed to the tens to hundreds of nanoseconds using conventional processing. Using some embodiments, matrix multiplication is estimated to speed up by two orders of magnitude relative to conventional techniques. For example, a multiplication that may be performed by a state-of-the-art graphics processing unit (GPU) can be performed in about 10 ns can be performed by a photonic processing system according to some embodiments in about 200 ps.

To implement a photonics-based processor, the inventors have recognized and appreciated that the multiplication of an input vector by a matrix can be accomplished by propagating coherent light signals, e.g., laser pulses, through a first array of interconnected variable beam splitters (VBSs), a second array of interconnected variable beam splitters, and multiple controllable optical elements (e.g., electro-optical or optomechanical elements) between the two arrays that connect a single output of the first array to a single input of the second array. Details of a photonic processing system that includes a photonic processor are described below.

B. Photonic Processing System Overview

Referring to FIG. 1-1, a photonic processing system 1-100 includes an optical encoder 1-101, a photonic processor 1-103, an optical receiver 1-105, and a controller 1-107, according to some embodiments. The photonic processing system 1-100 receives, as an input from an external processor (e.g., a CPU), an input vector represented by a group of input bit strings and produces an output vector represented by a group of output bit strings. For example, if the input vector is an n -dimensional vector, the input vector may be represented by n separate bit strings, each bit string representing a respective component of the vector. The input bit string may be received as an electrical or optical signal from the external processor and the output bit string may be transmitted as an electrical or optical signal to the external processor. In some embodiments, the controller 1-107 does not necessarily output an output bit string after every process iteration. Instead, the controller 1-107 may use one or more output bit strings to determine a new input bit stream to feed through the components of

the photonic processing system 1-100. In some embodiments, the output bit string itself may be used as the input bit string for a subsequent iteration of the process implemented by the photonic processing system 1-100. In other embodiments, multiple output bit streams are combined in various ways to determine a subsequent input bit string. For example, one or more output bit strings may be summed together as part of the determination of the subsequent input bit string.

The optical encoder 1-101 is configured to convert the input bit strings into optically encoded information to be processed by the photonic processor 1-103. In some embodiments, each input bit string is transmitted to the optical encoder 1-101 by the controller 1-107 in the form of electrical signals. The optical encoder 1-101 converts each component of the input vector from its digital bit string into an optical signal. In some embodiments, the optical signal represents the value and sign of the associated bit string as an amplitude and a phase of an optical pulse. In some embodiments, the phase may be limited to a binary choice of either a zero phase shift or a π phase shift, representing a positive and negative value, respectively. Embodiments are not limited to real input vector values. Complex vector components may be represented by, for example, using more than two phase values when encoding the optical signal. In some embodiments, the bit string is received by the optical encoder 1-101 as an optical signal (e.g., a digital optical signal) from the controller 1-107. In these embodiments, the optical encoder 1-101 converts the digital optical signal into an analog optical signal of the type described above.

The optical encoder 1-101 outputs n separate optical pulses that are transmitted to the photonic processor 1-103. Each output of the optical encoder 1-101 is coupled one-to-one to a single input of the photonic processor 1-103. In some embodiments, the optical encoder 1-101 may be disposed on the same substrate as the photonic processor 1-103 (e.g., the optical encoder 1-101 and the photonic processor 1-103 are on the same chip). In such embodiments, the optical signals may be transmitted from the optical encoder 1-101 to the photonic processor 1-103 in waveguides, such as silicon photonic waveguides. In other embodiments, the optical encoder 1-101 may be disposed on a separate substrate from the photonic processor 1-103. In such embodiments, the optical signals may be transmitted from the optical encoder 1-101 to the photonic processor 103 in optical fiber.

The photonic processor 1-103 performs the multiplication of the input vector by a matrix M . As described in detail below, the matrix M is decomposed into three matrices using a combination of a singular value decomposition (SVD) and a unitary matrix decomposition. In some embodiments, the unitary matrix decomposition is performed with operations similar

to Givens rotations in QR decomposition. For example, an SVD in combination with a Householder decomposition may be used. The decomposition of the matrix M into three constituent parts may be performed by the controller 1-107 and each of the constituent parts may be implemented by a portion of the photonic processor 1-103. In some embodiments, the photonic processor 1-103 includes three parts: a first array of variable beam splitters (VBSs) configured to implement a transformation on the array of input optical pulses that is equivalent to a first matrix multiplication (see, e.g., the first matrix implementation 1-301 of FIG. 1-3); a group of controllable optical elements configured to adjust the intensity and/or phase of each of the optical pulses received from the first array, the adjustment being equivalent to a second matrix multiplication by a diagonal matrix (see, e.g., the second matrix implementation 1-303 of FIG. 1-3); and a second array of VBSs configured to implement a transformation on the optical pulses received from the group of controllable electro-optical element, the transformation being equivalent to a third matrix multiplication (see, e.g., the third matrix implementation 1-305 of FIG. 3).

The photonic processor 1-103 outputs n separate optical pulses that are transmitted to the optical receiver 1-105. Each output of the photonic processor 1-103 is coupled one-to-one to a single input of the optical receiver 1-105. In some embodiments, the photonic processor 1-103 may be disposed on the same substrate as the optical receiver 1-105 (e.g., the photonic processor 1-103 and the optical receiver 1-105 are on the same chip). In such embodiments, the optical signals may be transmitted from the photonic processor 1-103 to the optical receiver 1-105 in silicon photonic waveguides. In other embodiments, the photonic processor 1-103 may be disposed on a separate substrate from the optical receiver 1-105. In such embodiments, the optical signals may be transmitted from the photonic processor 103 to the optical receiver 1-105 in optical fibers.

The optical receiver 1-105 receives the n optical pulses from the photonic processor 1-103. Each of the optical pulses is then converted to electrical signals. In some embodiments, the intensity and phase of each of the optical pulses is measured by optical detectors within the optical receiver. The electrical signals representing those measured values are then output to the controller 1-107.

The controller 1-107 includes a memory 1-109 and a processor 1-111 for controlling the optical encoder 1-101, the photonic processor 1-103 and the optical receiver 1-105. The memory 1-109 may be used to store input and output bit strings and measurement results from the optical receiver 1-105. The memory 1-109 also stores executable instructions that, when executed by the processor 1-111, control the optical encoder 1-101, perform the matrix

decomposition algorithm, control the VBSs of the photonic processor 103, and control the optical receivers 1-105. The memory 1-109 may also include executable instructions that cause the processor 1-111 to determine a new input vector to send to the optical encoder based on a collection of one or more output vectors determined by the measurement performed by the optical receiver 1-105. In this way, the controller 1-107 can control an iterative process by which an input vector is multiplied by multiple matrices by adjusting the settings of the photonic processor 1-103 and feeding detection information from the optical receiver 1-105 back to the optical encoder 1-101. Thus, the output vector transmitted by the photonic processing system 1-100 to the external processor may be the result of multiple matrix multiplications, not simply a single matrix multiplication.

In some embodiments, a matrix may be too large to be encoded in the photonic processor using a single pass. In such situations, one portion of the large matrix may be encoded in the photonic processor and the multiplication process may be performed for that single portion of the large matrix. The results of that first operation may be stored in memory 1-109. Subsequently, a second portion of the large matrix may be encoded in the photonic processor and a second multiplication process may be performed. This “chunking” of the large matrix may continue until the multiplication process has been performed on all portions of the large matrix. The results of the multiple multiplication processes, which may be stored in memory 1-109, may then be combined to form the final result of the multiplication of the input vector by the large matrix.

In other embodiments, only collective behavior of the output vectors is used by the external processor. In such embodiments, only the collective result, such as the average or the maximum/minimum of multiple output vectors, is transmitted to the external processor.

C. Optical Encoder

Referring to FIG. 1-2, the optical encoder includes at least one light source 1-201, a power tree 1-203, an amplitude modulator 1-205, a phase modulator 1-207, a digital to analog converter (DAC) 1-209 associated with the amplitude modulator 1-205, and a 1-DAC 211 associated with the phase modulator 1-207, according to some embodiments. While the amplitude modulator 1-205 and phase modulator 1-207 are illustrated in FIG. 1-2 as single blocks with n inputs and n outputs (each of the inputs and outputs being, for example, a waveguide), in some embodiments each waveguide may include a respective amplitude modulator and a respective phase modulator such that the optical encoder includes n

amplitude modulators and n phase modulators. Moreover, there may be an individual DAC for each amplitude and phase modulator. In some embodiments, rather than having an amplitude modulator and a separate phase modulator associated with each waveguide, a single modulator may be used to encode both amplitude and phase information. While using a single modulator to perform such an encoding limits the ability to precisely tune both the amplitude and phase of each optical pulse, there are some encoding schemes that do not require precise tuning of both the amplitude and phase of the optical pulses. Such a scheme is described later herein.

The light source 1-201 may be any suitable source of coherent light. In some embodiments, the light source 1-201 may be a diode laser or a vertical-cavity surface emitting lasers (VCSEL). In some embodiments, the light source 1-201 is configured to have an output power greater than 10 mW, greater than 25 mW, greater than 50 mW, or greater than 75mW. In some embodiments, the light source 1-201 is configured to have an output power less than 100 mW. The light source 1-201 may be configured to emit a continuous wave of light or pulses of light (“optical pulses”) at one or more wavelengths (e.g., the C-band or O-band). The temporal duration of the optical pulses may be, for example, about 100 ps.

While light source 1-201 is illustrated in FIG. 1-2 as being on the same semiconductor substrate as the other components of the optical encoder, embodiments are not so limited. For example, the light source 1-201 may be a separate laser packaging that is edge-bonded or surface-bonded to the optical encoder chip. Alternatively, the light source 1-201 may be completely off-chip and the optical pulses may be coupled to a waveguide 1-202 of the optical encoder 1-101 via an optical fiber and/or a grating coupler.

The light source 1-201 is illustrated as two light sources 1-201a and 1-201b, but embodiments are not so limited. Some embodiments may include a single light source. Including multiple light sources 201a-b, which may include more than two light sources, can provide redundancy in case one of the light sources fails. Including multiple light sources may extend the useful lifetime of the photonic processing system 1-100. The multiple light sources 1-201a-b may each be coupled to a waveguide of the optical encoder 1-101 and then combined at a waveguide combiner that is configured to direct optical pulses from each light source to the power tree 1-203. In such embodiments, only one light source is used at any given time.

Some embodiments may use two or more phase-locked light sources of the same wavelength at the same time to increase the optical power entering the optical encoder

system. A small portion of light from each of the two or more light sources (e.g., acquired via a waveguide tap) may be directed to a homodyne detector, where a beat error signal may be measured. The beat error signal may be used to determine possible phase drifts between the two light sources. The beat error signal may, for example, be fed into a feedback circuit that controls a phase modulator that phase locks the output of one light source to the phase of the other light source. The phase-locking can be generalized in a master-slave scheme, where $N \geq 1$ slave light sources are phase-locked to a single master light source. The result is a total of $N + 1$ phase-locked light sources available to the optical encoder system.

In other embodiments, each separate light source may be associated with light of different wavelengths. Using multiple wavelengths of light allows some embodiments to be multiplexed such that multiple calculations may be performed simultaneously using the same optical hardware.

The power tree 1-203 is configured to divide a single optical pulse from the light source 1-201 into an array of spatially separated optical pulses. Thus, the power tree 1-203 has one optical input and n optical outputs. In some embodiments, the optical power from the light source 1-201 is split evenly across n optical modes associated with n waveguides. In some embodiments, the power tree 1-203 is an array of 50:50 beam splitters 1-801, as illustrated in FIG. 1-8. The number “depth” of the power tree 1-203 depends on the number of waveguides at the output. For a power tree with n output modes, the depth of the power tree 1-203 is $\text{ceil}(\log_2(n))$. The power tree 1-203 of FIG. 1-8 only illustrates a tree depth of three (each layer of the tree is labeled across the bottom of the power tree 1-203). Each layer includes 2^{m-1} beam splitters, where m is the layer number. Consequently, the first layer has a single beam splitter 1-801a, the second layer has two beam splitters 1-801b – 1-801c, and the third layer has four beam splitters 1-801d – 1-801g.

While the power tree 1-203 is illustrated as an array of cascading beam splitters, which may be implemented as evanescent waveguide couplers, embodiments are not so limited as any optical device that converts one optical pulse into a plurality of spatially separated optical pulses may be used. For example, the power tree 1-203 may be implemented using one or more multimode interferometers (MMI), in which case the equations governing layer width and depth would be modified appropriately.

No matter what type of power tree 1-203 is used, it is likely that manufacturing a power tree 1-203 such that the splitting ratios are precisely even between the n output modes will be difficult, if not impossible. Accordingly, adjustments can be made to the setting of the amplitude modulators to correct for the unequal intensities of the n optical pulses output by

the power tree. For example, the waveguide with the lowest optical power can be set as the maximum power for any given pulse transmitted to the photonic processor 1-103. Thus, any optical pulse with a power higher than the maximum power may be modulated to have a lower power by the amplitude modulator 1-205, in addition to the modulation to the amplitude being made to encode information into the optical pulse. A phase modulator may also be placed at each of the n output modes, which may be used to adjust the phase of each output mode of the power tree 1-203 such that all of the output signals have the same phase.

Alternatively or additionally, the power tree 1-203 may be implemented using one or more Mach-Zehnder Interferometers (MZI) that may be tuned such that the splitting ratios of each beam splitter in the power tree results in substantially equal intensity pulses at the output of the power tree 1-203.

The amplitude modulator 1-205 is configured to modify, based on a respective input bit string, the amplitude of each optical pulse received from the power tree 1-203. The amplitude modulator 1-205 may be a variable attenuator or any other suitable amplitude modulator controlled by the DAC 1-209, which may further be controlled by the controller 1-107. Some amplitude modulators are known for telecommunication applications and may be used in some embodiments. In some embodiments, a variable beam splitter may be used as an amplitude modulator 1-205, where only one output of the variable beam splitter is kept and the other output is discarded or ignored. Other examples of amplitude modulators that may be used in some embodiments include traveling wave modulators, cavity-based modulators, Franz-Keldysh modulators, plasmon-based modulators, 2-D material-based modulators and nano-opto-electro-mechanical switches (NOEMS).

The phase modulator 1-207 is configured to modify, based on the respective input bit string, the phase of each optical pulse received from the power tree 1-203. The phase modulator may be a thermo-optic phase shifter or any other suitable phase shifter that may be electrically controlled by the 1-211, which may further be controlled by the controller 1-107.

While FIG. 1-2 illustrates the amplitude modulator 1-205 and phase modulator 1-207 as two separate components, they may be combined into a single element that controls both the amplitudes and phases of the optical pulses. However, there are advantages to separately controlling the amplitude and phase of the optical pulse. Namely, due to the connection between amplitude shifts and phase shifts via the Kramers-Kronenig relations, there is a phase shift associated with any amplitude shift. To precisely control the phase of an optical pulse, the phase shift created by the amplitude modulator 1-205 should be compensated for using the phase modulator 1-207. By way of example, the total amplitude of an optical pulse

exiting the optical encoder 1-101 is $A = a_0 a_1 a_2$ and the total phase of the optical pulse exiting the optical encoder is $\theta = \Delta\theta + \Delta\varphi + \varphi$, where a_0 is the input intensity of the input optical pulse (with an assumption of zero phase at the input of the modulators), a_1 is the amplitude attenuation of the amplitude modulator 1-205, $\Delta\theta$ is the phase shift imparted by the amplitude modulator 1-205 while modulating the amplitude, $\Delta\varphi$ is the phase shift imparted by the phase modulator 1-207, a_2 is the attenuation associated with the optical pulse passing through the phase modulator 1-209, and φ is the phase imparted on the optical signal due to propagation of the light signal. Thus, setting the amplitude and the phase of an optical pulse is not two independent determinations. Rather, to accurately encode a particular amplitude and phase into an optical pulse output from the optical encoder 1-101, the settings of both the amplitude modulator 1-205 and the phase modulator 1-207 should be taken into account for both settings.

In some embodiments, the amplitude of an optical pulse is directly related to the bit string value. For example, a high amplitude pulse corresponds to a high bit string value and a low amplitude pulse corresponds to a low bit string value. The phase of an optical pulse encodes whether the bit string value is positive or negative. In some embodiments, the phase of an optical pulse output by the optical encoder 1-101 may be selected from two phases that are 180 degrees (π radians) apart. For example, positive bit string values may be encoded with a zero degree phase shift and negative bit string values may be encoded with a 180 degree (π radians) phase shift. In some embodiments, the vector is intended to be complex-valued and thus the phase of the optical pulse is chosen from more than just two values between 0 and 2π .

In some embodiments, the controller 1-107 determines the amplitude and phase to be applied by both the amplitude modulator 1-205 and the phase modulator 1-207 based on the input bit string and the equations above linking the output amplitude and output phase to the amplitudes and phases imparted by the amplitude modulator 1-204 and the phase modulator 1-207. In some embodiments, the controller 1-107 may store in memory 1-109 a table of digital values for driving the amplitude modulator 1-205 and the phase modulator 1-207. In some embodiments, the memory may be placed in close proximity to the modulators to reduce the communication temporal latency and power consumption.

The digital to analog converter (DAC) 1-209, associated with and communicatively coupled to the amplitude modulator 1-205, receives the digital driving value from the controller 1-107 and converts the digital driving value to an analog voltage that drives the

amplitude modulator 1-205. Similarly, the DAC 1-211, associated with and communicatively coupled to the phase modulator 1-207, receives the digital driving value from the controller 1-107 and converts the digital driving value to an analog voltage that drives the phase modulator 1-207. In some embodiments, the DAC may include an amplifier that amplifies the analog voltages to sufficiently high levels to achieve the desired extinction ratio within the amplitude modulators (e.g., the highest extinction ratio physically possible to implement using the particular phase modulator) and the desired phase shift range within the phase modulators (e.g., a phase shift range that covers the full range between 0 and 2π). While the DAC 1-209 and the DAC 1-211 are illustrated in FIG. 1-2 as being located in and/or on the chip of the optical encoder 1-101, in some embodiments, the DACs 1-209 and 1-211 may be located off-chip while still being communicatively coupled to the amplitude modulator 1-205 and the phase modulator 1-207, respectively, with electrically conductive traces and/or wires.

After modulation by the amplitude modulator 1-205 and the phase modulator 1-207, the n optical pulses are transmitted from the optical encoder 1-101 to the photonic processor 1-103.

D. Photonic Processor

Referring to FIG. 1-3, the photonic processor 1-103 implements matrix multiplication on an input vector represented by the n input optical pulse and includes three main components: a first matrix implementation 1-301, a second matrix implementation 1-303, and a third matrix implementation 1-305. In some embodiments, as discussed in more detail below, the first matrix implementation 1-301 and the third matrix implementation 1-305 include an interconnected array of programmable, reconfigurable, variable beam splitters (VBSs) configured to transform the n input optical pulses from an input vector to an output vector, the components of the vectors being represented by the amplitude and phase of each of the optical pulses. In some embodiments, the second matrix implementation 1-303 includes a group of electro-optic elements.

The matrix by which the input vector is multiplied, by passing the input optical pulses through the photonic processor 1-103, is referred to as M . The matrix M is a general $m \times n$ known to the controller 1-107 as the matrix that should be implemented by the photonic processor 1-103. As such, the controller 1-107 decomposes the matrix M using a singular value decomposition (SVD) such that the matrix M is represented by three constituent matrices: $M = V^T \Sigma U$, where U and V are real orthogonal $n \times n$ and $m \times m$ matrices,

respectively ($U^T U = U U^T = I$ and $V^T V = V V^T = I$), and Σ is an $m \times n$ diagonal matrix with real entries. The superscript “T” in all equations represents the transpose of the associated matrix. Determining the SVD of a matrix is known and the controller 1-107 may use any suitable technique to determine the SVD of the matrix M . In some embodiments, the matrix M is a complex matrix, in which case the matrix M can be decomposed into $M = V^\dagger \Sigma U$, where V and U are complex unitary $n \times n$ and $m \times m$ matrices, respectively $U^\dagger U = U U^\dagger = I$ and $V^\dagger V = V V^\dagger = I$), and Σ is an $m \times n$ diagonal matrix with real or complex entries. The values of the diagonal singular values may also be further normalized such that the maximum absolute value of the singular values is 1.

Once the controller 1-107 has determined the matrices U , Σ and V for the matrix M , in the case where the matrices U and V are orthogonal real matrices, the control may further decompose the two orthogonal matrices U and V into a series of real-valued Givens rotation matrices. A Givens rotation matrix $G(i, j, \theta)$ is defined component-wise by the following equations:

$$\begin{aligned} g_{kk} &= 1 \text{ for } k \neq i, j \\ g_{kk} &= \cos(\theta) \text{ for } k = i, j \\ g_{ij} &= -g_{ji} = -\sin(\theta), \\ g_{kl} &= 0 \text{ otherwise,} \end{aligned}$$

where g_{ij} represents the element in the i -th row and j -th column of the matrix G and θ is the angle of rotation associated with the matrix. Generally, the matrix G is an arbitrary 2×2 unitary matrix with determinant 1 (SU(2) group) and it is parameterized by two parameters. In some embodiments, those two parameters are the rotation angle θ and another phase value ϕ . Nevertheless, the matrix G can be parameterized by other values other than angles or phases, e.g. by reflectivities/transmissivities or by separation distances (in the case of NOEMS).

Algorithms for expressing an arbitrary real orthogonal matrix in terms of a product of sets of Givens rotations in the complex space are provided in M. Reck, *et al.*, “Experimental realization of any discrete unitary operator,” *Physical Review Letters* **73**, 58 (1994) (“Reck”), and W. R. Clements, *et al.*, “Optimal design for universal multiport interferometers,” *Optica* **3**, 12 (2016) (“Clements”), both of which are incorporated herein by reference in their entirety and at least for their discussions of techniques for decomposing a real orthogonal matrix in terms of Givens rotations. (In the case that any terminology used herein conflicts with the usage of that terminology in Reck and/or Clements, the terminology should be

afforded a meaning most consistent with how a person of ordinary skill would understand its usage herein.). The resulting decomposition is given by the following equation:

$$U = D \prod_{k=1}^n \prod_{(i,j) \in S_k} G(i,j,\theta_{ij}^{(k)})$$

where U is an $n \times n$ orthogonal matrix, S_k is the set of indices relevant to the k -th set of Givens rotations applied (as defined by the decomposition algorithm), $\theta_{ij}^{(k)}$ represents the angle applied for the Givens rotation between components i and j in the k -th set of Givens rotations, and D is a diagonal matrix of either +1 or -1 entries representing global signs on each component. The set of indices S_k is dependent on whether n is even or odd. For example, when n is even:

- $S_k = \{(1,2), (3,4), \dots, (n-1, n)\}$ for odd k
- $S_k = \{(2,3), (4,5), \dots, (n-2, n-1)\}$ for even k

When n is odd:

- $S_k = \{(1,2), (3,4), \dots, (n-2, n-1)\}$ for odd k
- $S_k = \{(2,3), (4,5), \dots, (n-1, n)\}$ for even k

By way of example and not limitation, the decomposition of a 4×4 orthogonal matrix can be represented as:

$$U = DG(1,2,\theta_{12}^{(1)})G(3,4,\theta_{34}^{(1)})G(2,3,\theta_{23}^{(2)})G(1,2,\theta_{12}^{(2)})G(3,4,\theta_{34}^{(2)})G(2,3,\theta_{23}^{(3)})$$

A brief overview of one embodiment of an algorithm for decomposing an $n \times n$ matrix U in terms of n sets of real-valued Givens rotations, which may be implemented using the controller 1-107, is as follows:

```

U' ← U
For i from 1 to n-1:
  If i is odd:
    For j=0 to i-1:
      Find  $G_{j,j+1}^T(\theta)$  that nullifies element  $U'_{n-j,j}$  i.e.  $\theta = \tan^{-1}(U'_{n-j,j}/U'_{n-j,j-1})$ 
       $U' \leftarrow U' G_{j,j+1}^T(\theta)$ 
  Else if i is even:
    For j=1 to i:
      Find  $G_{n-j-1,n-j}(\theta)$  that nullifies element  $U'_{n-j-1,j}$  i.e.  $\theta = \tan^{-1}(-U'_{n-j-1,j}/U'_{n-j-1,j-1})$ 
       $U' \leftarrow G_{n-j-1,n-j}(\theta) U'$ 

```

The resultant matrix U' of the above algorithm is lower triangular and is related to the original matrix U by the equation:

$$U = \prod_{(j,k) \in S_L} G_{jk}^T U' \prod_{(j,k) \in S_R} G_{jk} = D_U \prod_{(j,k) \in S} G_{jk},$$

where the label S_L labels the set of two modes connected by the VBS to the left of U' and the label S_R labels the set of two modes connected by the VBS to the right of U' . Because U is an orthogonal matrix, U' is a diagonal matrix with $\{-1,1\}$ entries along the diagonal. This matrix, $U' = D_U$, is referred to as a “phase screen.”

The next step of the algorithm, is to repeatedly find which $G_{jk}^T(\theta_1)D_U = D_U G_{jk}(\theta_2)$ is accomplished using the following algorithm, which may be implemented using the controller 1-107:

```

For every (j,k) in  $S_L$ :
    If  $U'_{jj}$  and  $U'_{kk}$  have different signs:
         $\theta_2 = -\theta_1$ 
    Else:
         $\theta_2 = \theta_1$ 

```

The above algorithm may also be used to decompose V and/or V^T to determine the m layers of VBS values and the associated phase screen.

The above concept of decomposing an orthogonal matrix into real-valued Givens rotation matrices can be expanded to complex matrices, e.g., unitary matrices rather than orthogonal matrices. In some embodiments, this may be accomplished by including an additional phase in the parameterization of the Givens rotation matrices. Thus, a general form of the Givens matrices with the addition of the additional phase term is $T(i,j,\theta,\phi)$, where

$$\begin{aligned}
 t_{kk} &= 1 \text{ for } k \neq i, j, \\
 t_{ii} &= e^{i\phi} \cos(\theta), \\
 t_{jj} &= \cos(\theta), \\
 t_{ij} &= -\sin(\theta), \\
 t_{ji} &= e^{i\phi} \sin(\theta), \\
 t_{kl} &= 0 \text{ otherwise,}
 \end{aligned}$$

where t_{ij} represents the i -th row and j -th column of the matrix T , θ is the angle of rotation associated with the matrix, and ϕ is the additional phase. Any unitary matrix can be decomposed into matrices of the type $T(i,j,\theta,\phi)$. By making the choice to set the phase $\phi = 0$, the conventional real-valued Givens rotation matrices described above are obtained. If, instead, the phase $\phi = \pi$, then a set of matrices known as Householder matrices are obtained. A Householder matrix, H , has the form $H = I - (v \otimes v)$, where I is the $n \times n$ identity

matrix, v is a unit vector, and $\otimes$ is the outer product. Householder matrices represent reflections about a hyperplane orthogonal to the unit vector v . In this parameterization the hyperplane is a two-dimensional subspace, rather than an $n - 1$ dimensional subspace as is common in defining Householder matrices for the QR decomposition. Thus, a decomposition of a matrix into Givens rotations is equivalent to a decomposition of the matrix into Householder matrices.

Based on the aforementioned decomposition of an arbitrary unitary matrix into a restricted set of Givens rotations, any unitary matrix can be implemented by a particular sequence of rotations and phase shifts. And in photonics, rotations may be represented by variable beam splitters (VBS) and phase shifts are readily implemented using phase modulators. Accordingly, for the n optical inputs of the photonic processor 1-103, the first matrix implementation 1-301 and the third matrix implementation 1-305, representing the unitary matrices of the SVD of the matrix M may be implemented by an interconnected array of VBSs and phase shifters. Because of the parallel nature of passing optical pulses through a VBS array, matrix multiplication can be performed in $O(1)$ time. The second matrix implementation 1-303 is a diagonal matrix of the SVD of the matrix M combined with the diagonal matrices D associated with each of the orthogonal matrices of the SVD. As mentioned above, each matrix D is referred to as a “phase screen” and can be labeled with a subscript to denote whether it is the phase screen associated with the matrix U or the matrix V . Thus, the second matrix implementation 303 is the matrix $\Sigma' = D_V \Sigma D_U$.

In some embodiments, the VBS unit cell of the photonic processor 1-103 associated with the first matrix implementation 1-301 and the third matrix implementation 1-305 may be a Mach-Zehnder interferometer (MZI) with an internal phase shifter. In other embodiments, the VBS unit cell may be a microelectromechanical systems (MEMS) actuator. An external phase shifter may be used in some embodiments to implement the additional phase needed for the Givens rotations.

The second matrix implementation 1-303, representing the diagonal matrix $D_V \Sigma D_U$ may be implemented using an amplitude modulator and a phase shifter. In some embodiments, a VBS may be used to split off a portion of light that can be dumped to variably attenuate an optical pulse. Additionally or alternatively, a controllable gain medium may be used to amplify an optical signal. For example, GaAs, InGaAs, GaN, or InP may be used as an active gain medium for amplifying an optical signal. Other active gain processes such as the second harmonic generation in materials with crystal inversion symmetric, e.g.

KTP and lithium niobate, and the four-wave mixing processes in materials that lack inversion symmetry, e.g. silicon, can also be used. A phase shifter in each optical mode may be used to apply either a zero or a π phase shift, depending on the phase screen being implemented. In some embodiments, only a single phase shifter for each optical mode is used rather than one phase shifter for each phase screen. This is possible because each of the matrices D_V , Σ , and D_U are diagonal and therefore commute. Thus, the value of each phase shifter of the second matrix implementation 1-303 of the photonic processor 1-103 is the result of the product of the two phase screens: $D_V D_U$.

Referring to FIG. 1-4, the first and third matrix implementation 1-301 and 1-305 are implemented as an array of VBSs 1-401, according to some embodiments. For the sake of simplicity, only $n=6$ input optical pulses (the number of rows) are illustrated resulting in a “circuit depth” (e.g., the number of columns) equal to the number of input optical pulses (e.g., six). For the sake of clarity, only a single VBS 1-401 is labeled with a reference numeral. The VBS are labeled, however, with subscripts that identify which optical modes are being mixed by a particular VBS and a super script labeling the associated column. Each VBS 1-401 implements a complex Givens rotation, $T(i,j,\theta,\phi)$, as discussed above, where i and j are equivalent to the subscript labels of the VBSs 1-401, θ is the rotation angle of the Givens rotation, and ϕ is the additional phase associated with the generalized rotation.

Referring to FIG. 1-5, each VBS 1-401 may be implemented using a MZI 1-510 and at least one external phase shifter 1-507. In some embodiments, a second external phase shifter 1-509 may also be included. The MZI 1-510 includes a first evanescent coupler 1-501 and a second evanescent coupler 1-503 for mixing the two input modes of the MZI 1-510. An internal phase shifter 1-505 modulates the phase θ in one arm of the MZI 1-510 to create a phase difference between the two arms. Adjusting the phase θ causes the intensity of light output by the VBS 1-401 to vary from one output mode of the MZI 1-510 to the other thereby creating a beam splitter that is controllable and variable. In some embodiments, a second internal phase shifter can be applied on the second arm. In this case, it is the difference between the two internal phase shifters that cause the output light intensity to vary. The average between the two internal phases will impart a global phase to the light that enter mode i and mode j . Thus the two parameters θ and ϕ may each be controlled by a phase shifter. In some embodiments, the second external phase shifter 1-509 may be used to correct for an unwanted differential phase across the output modes of the VBS due to static phase disorder.

In some embodiments, the phase shifters 1-505, 1-507 and 1-509 may include a thermo-optic, electro-optic, or optomechanic phase modulator. In other embodiments, rather than including an internal phase modulator 505 within an MZI 510, a NOEMS modulator may be used.

In some embodiments, the number of VBSs grows with the size of the matrix. The inventors have recognized and appreciated that controlling a large number of VBSs can be challenging and there is a benefit to sharing a single control circuit among multiple VBSs. An example of a parallel control circuit that may be used to control multiple VBSs is a digital-to-analog converter receives as an input a digital string that encodes the analog signal to be imparted on a specific VBS. In some embodiments, the circuit also receives a second input the address of the VBS that is to be controlled. The circuit may then impart analog signals on the addressed VBS. In other embodiments, the control circuit may automatically scan through a number of VBSs and impart analog signals on the multiple VBSs without being actively given an address. In this case, the addressing sequence is predefined such that it traverses the VBS array in known order.

Referring to FIG. 1-6, the second matrix implementation 1-303 implements multiplication by the diagonal matrix $\Sigma' = D_V \Sigma D_U$. This may be accomplished using two phase shifters 1-601 and 1-605 to implement the two phase screens and an amplitude modulator 1-603 to adjust the intensity of an associate optical pulse by an amount η . As mentioned above, in some embodiments only a single phase modulator 1-601 may be used, as the two phase screens can be combined together since the three constituent matrices that form Σ' are diagonal and therefore commute.

In some embodiments, the amplitude modulators 1-603 may be implemented using an attenuator and/or an amplifier. If the value of the amplitude modulation η is greater than one, the optical pulse is amplified. If the value of the amplitude modulation η is less than one, the optical pulse is attenuated. In some embodiments, only attenuation is used. In some embodiments, the attenuation may be implemented by a column of integrated attenuators. In other embodiments, as illustrated in FIG. 1-7, the attenuation 1-603 may be implemented using a MZI that includes two evanescent couplers 1-701 and 1-703 and a controllable internal phase shifter 1-705 to adjust how much of the input light is transmitted from the input of the MZI to a first output port 1-709 of the MZI. A second output port 1-707 of the MZI may be ignored, blocked or dumped.

In some embodiments, the controller 1-107 controls the value of each phase shifter in the photonic processor 1-103. Each phase shifter discussed above may include a DAC similar

to the DACs discussed in connection with the phase modulator 1-207 of the optical encoder 1-101.

The photonic processor 1-103 can include any number of input nodes, but the size and complexity of the interconnected VBS arrays 1-301 and 1-305 will increase as the number of input modes increases. For example, if there are n input optical modes, then the photonic processor 1-103 will have a circuit depth of $2n+1$, where the first matrix implementation 1-301 and the second matrix implementation 1-305 each has a circuit depth n and the second matrix implementation 1-303 has a circuit depth of one. Importantly, the complexity in time of performing a single matrix multiplication is not even linear with the number of input optical pulses – it is always $O(1)$. In some embodiments, this low order complexity afforded by the parallelization results in energy and time efficiencies that cannot be obtained using conventional electrical processors.

It is noted that, while embodiments described herein illustrate the photonic processor 1-103 as having n inputs and n outputs, in some embodiments, the matrix M implemented by the photonic processor 1-103 may not be a square matrix. In such embodiments, the photonic processor 1-103 may have a different number of outputs and inputs.

It is also noted that, due to the topology of the interconnections of the VBSs within the first and second matrix implementations 1-301 and 1-305, it is possible to subdivide the photonic processor 1-103 into non-interacting subsets of rows such that more than one matrix multiplication can be performed at the same time. For example, in the VBS array illustrated in FIG. 1-4, if each VBS 1-401 that couples optical modes 3 and 4 is set such that optical modes 3 and 4 do not couple at all (e.g., as if the VBSs 1-401 with subscript “34” were absent from FIG. 1-4) then the top three optical modes would operate completely independently from the bottom three optical modes. Such a subdivision may be done at a much larger scale with a photonic processor with a larger number of input optical modes. For example, an $n=64$ photonic processor may multiply eight eight-component input vectors by a respective 8×8 matrix simultaneously (each of the 8×8 matrices being separately programmable and controllable). Moreover, the photonic processor 1-103 need not be subdivided evenly. For example, an $n=64$ photonic processor may subdivide into seven different input vectors with 20, 13, 11, 8, 6, 4, and 2 components, respectively, each multiplied by a respective matrix simultaneously. It should be understood that the above numerical examples are for illustration purposes only and any number of subdivisions is possible.

Additionally, while the photonic processor 1-103 performs vector-matrix multiplication, where a vector is multiplied by a matrix by passing the optical signals through

the array of VBSs, the photonic processor 1-103 may also be used to perform matrix-matrix multiplication. For example, multiple input vectors may be passed through the photonic processor 1-103, one after the other, one input vector at a time, where each input vector represents a column of an input matrix. After optically computing each of the individual vector-matrix multiplications (each multiplication resulting in an output vector that corresponds to a column of an output column of the resulting matrix), the results may be combined digitally to form the output matrix resulting from the matrix-matrix multiplication.

E. Optical Receiver

The photonic processor 1-103 outputs n optical pulses that are transmitted to the optical receiver 1-105. The optical receiver 1-105 receives the optical pulses and generates an electrical signal based on the received optical signals. In some embodiments, the amplitude and phase of each optical pulse is determined. In some embodiments, this is achieved using homodyne or heterodyne detection schemes. In other embodiments, simple phase-insensitive photodetection may be performed using conventional photodiodes.

Referring to FIG. 1-9, the optical receiver 1-105 includes a homodyne detector 1-901, a transimpedance amplifier 1-903 and an analog-to-digital converter (ADC) 1-905, according to some embodiments. While the components are shown as one element for all optical modes in FIG. 1-9, this is for the sake of simplicity. Each optical mode may have a dedicated homodyne detector 1-901, a dedicated transimpedance amplifier 1-903 and a dedicated ADC 1-905. In some embodiments, a transimpedance amplifier 1-903 may not be used. Instead, any other suitable electronic circuit that converts a current to a voltage may be used.

Referring to FIG. 1-10, the homodyne detector 1-903 includes a local oscillator (LO) 1-1001, a quadrature controller 1-1003, a beam splitter 1-1005 and two detectors 1-1007 and 1-1009, according to some embodiments. The homodyne detector 1-903 outputs an electrical current that is based on the difference between the current output by the first detector 1-1007 and the second detector 1-1009.

The local oscillator 1-1001 is combined with the input optical pulse at the beam splitter 1-1005. In some embodiments, a portion of the light source 1-201 is transmitted via an optical waveguide and/or an optical fiber to the homodyne detector 1-901. The light from the light source 1-201 may itself be used as the local oscillator 1-1001 or, in other embodiments, the local oscillator 1-1001 may be a separate light source that uses the light from the light source 1-201 to generate a phase matched optical pulse. In some embodiments,

an MZI may replace the beam splitter 1-1005 such that adjustments can be made between the signal and the local oscillator.

The quadrature controller 1-1003 controls the cross-section angle in phase space in which the measurement is made. In some embodiments, the quadrature controller 1-1003 may be a phase shifter that controls the relative phase between the input optical pulse and the local oscillator. The quadrature controller 1-1003 is shown as a phase shifter in the input optical mode. But in some embodiments, the quadrature controller 1-1003 may be in the local oscillator mode.

The first detector 1-1007 detects light output by a first output of the beam splitter 1-1005 and the second detector 1-1009 detects light output by a second output of the beam splitter 1-1005. The detectors 1-1007 and 1-1009 may be photodiodes operated with zero bias.

A subtraction circuit 1-1011 subtracts the electrical current from the first detector 1-1007 from the electrical current from the second detector 1-1009. The resulting current therefore has an amplitude and a sign (plus or minus). The transimpedance amplifier 1-903 converts this difference in current into a voltage, which may be positive or negative. Finally, an ADC 1-905 converts the analog signal to a digital bit string. This output bit string represents the output vector result of the matrix multiplication and is an electrical, digital version of the optical output representation of the output vector that is output by the photonic processor 1-103. In some embodiments, the output bit string may be sent to the controller 1-107 for additional processing, which may include determining a next input bit string based on one or more output bit strings and/or transmitting the output bit string to an external processor, as described above.

The inventors have further recognized and appreciated that the components of the above-described photonic processing system 1-100 need not be chained together back-to-back such that there is a first matrix implementation 1-301 connected to a second matrix implementation 1-303 connected to a third matrix implementation 1-305. In some embodiments, the photonic processing system 1-103 may include only a single unitary circuit for performing one or more multiplications. The output of the single unitary circuit may be connected directly to the optical receiver 1-105, where the results of the multiplication are determined by detecting the output optical signals. In such embodiments, the single unitary circuit may, for example, implement the first matrix implementation 1-301. The results detected by the optical receiver 1-105 may then be transmitted digitally to a conventional

processor (e.g., processor 1-111) where the diagonal second matrix implementation 1-303 is performed in the digital domain using a conventional processor (e.g., 1-111). The controller 1-107 may then reprogram the single unitary circuit to perform the third matrix implementation 1-305, determine an input bit string based on the result of the digital implementation of the second matrix implementation, and control the optical encoder to transmit optical signals, encoded based on the new input bit string, through the single unitary circuit with the reprogrammed settings. The resulting output optical signals, which are detected by the optical receiver 105, are then used to determine the results of the matrix multiplication.

The inventors have also recognized and appreciated that there can be advantages to chaining multiple photonic processors 1-103 back-to-back, in series. For example, to implement a matrix multiplication $M = M_1 M_2$, where M_1 and M_2 are arbitrary matrices but M_2 changes more frequently than M_1 based on a changing input workload, the first photonic processor can be controlled to implement M_2 and the second photonic processor coupled optically to the first photonic processor can implement M_1 which is kept static. In this way, only the first photonic processing system needs to be frequently updated based on the changing input workload. Not only does such an arrangement speed up the computation, but it also reduces the number of data bits that travel between the controller 1-107 and the photonic processors.

F. Folded Photonic Processing System

In FIG. 1-1, in such an arrangement, the optical encoder 1-101 and the optical receiver 1-105 are positioned on opposite sides of the photonic processing system 1-100. In applications where feedback from the optical receiver 1-105 is used to determine the input for the optical encoder 1-101 for a future iteration of the process, the data is transferred electronically from the optical receiver 1-105 to the controller 1-107 and then to the optical encoder 1-101. The inventors have recognized and appreciated that reducing the distance that these electrical signals need to travel (e.g., by reducing the length of electrical traces and/or wires) results in power savings and lower latency. Moreover, there is no need for the optical encoder 1-101 and optical receiver 1-105 to be placed on opposite ends of the photonic processing system.

Accordingly, in some embodiments, the optical encoder 1-101 and the optical receiver 1-105 are positioned near one another (e.g., on the same side of the photonics processor 1-103) such that the distance electrical signals have to travel between the optical encoder 1-101 and the optical receiver 1-105 is less than the width of the photonics processor 1-103. This may be accomplished by physically interleaving components of the first matrix implementation 1-301 and the third matrix implementation 1-305 such that they are physically in the same portion of the chip. This arrangement is referred to as a “folded” photonic processing system because the light first propagates in a first direction through the first matrix implementation 1-301 until it reaches a physical portion of the chip that is far from the optical encoder 1-101 and the optical receiver 1-105, then folds over such that the waveguides turn the light to be propagating in a direction opposite to the first direction when implementing the third matrix implementation 1-305. In some embodiments, the second matrix implementation 1-303 is physically located adjacent to the fold in the waveguides. Such an arrangement reduces the complexity of the electrical traces connecting the optical encoder 1-101, the optical receiver 1-105, and the controller 1-107 and reduces the total chip area used to implement the photonic processing system 1-100. For example, some embodiments using the folded arrangement only use 65% of the total chip area that would be needed if the back-to-back photonic arrangement of FIG. 1-1 was used. This may reduce the cost and complexity of the photonic processing system.

The inventors have recognized and appreciated that there are not only electrical advantages to a folded arrangement, but also optical advantages. For example, by reducing the distance that the light signal has to travel from the light source to be used as a local oscillator for the homodyne detection, the time-dependent phase fluctuations of the optical signal may be reduced, resulting in higher quality detection results. In particular, by locating the light source and the homodyne on the same side of the photonics processor, the distance traveled by the light signal used for the local oscillator is no longer dependent on the size of the matrix. For example, in the back-to-back arrangement of FIG. 1-1, the distance traveled by the light signal for the local oscillator scales linearly with the size of the matrix, whereas the distance traveled in the folded arrangement is constant, irrespective of the matrix size.

FIG. 1-11 is a schematic drawing of a folded photonic processing system 1-1100, according to some embodiments. The folded photonic processing system 1-1100 includes a

power tree 1-1101, a plurality of optical encoders 1-1103a – 1-1103d, a plurality of homodyne detectors 1-1105a – 1-1105d, a plurality of selector switches 1-1107a – 1-1107d, a plurality of U-matrix components 1-1109a – 1-1109j, a plurality of diagonal-matrix components 1-1111a – 1-1111d, and a plurality of V-matrix components 1-1113a – 1-1113j. For the sake of clarity, not all components of the folded photonic processing system are shown in the figure. It should be understood that the folded photonic processing system 1-1100 may include similar components as the back-to-back photonic processing system 1-100.

The power tree 1-1101 is similar to the power tree 1-203 of FIG. 2 and is configured to deliver light from a light source (not shown) to the optical encoders 1-1103. However, a difference in the power tree 1-1101 and the power tree 1-203 is that the power tree delivers optical signals to the homodyne detectors 1-1105a directly. In FIG. 2, the light source 201 delivers a local oscillator signal to the homodyne detectors on the other side of the photonic processor by tapping off a portion of the optical signal from the light source and guiding the optical signal using a waveguide. In FIG. 1-11, the power tree 1-1101 includes a number of outputs that is equal to twice the number of spatial modes. For example, FIG. 1-11 illustrates only four spatial modes of a photonic processor, which results in eight output modes from the power tree 1-1101 – one output directing light to each optical encoder 1-1103 and one output directing light to each homodyne detector 1-1105. The power tree may be implemented, for example, using cascading beam splitters or a multimode interferometer (MMI).

The optical encoders 1-1103 are similar to the power tree optical encoder 1-101 of FIG. 1 and are configured to encode information into the amplitude and/or phase of the optical signals received from the power tree 1-1101. This may be achieved, for example as described in connection with the optical encoder 1-101 of FIG. 2.

The homodyne detectors 1-1105 are located between the power tree 1-1101 and the U-matrix components 1-1109. In some embodiments, the homodyne detectors 1-1105 are physically positioned in a column with the optical encoder 1-1103. In some embodiments, the optical encoders 1-1103 and the homodyne detectors 1-1105 may be interleaved in a single column. In this way, the optical encoders 1-1103 and the homodyne detectors 1-1105 are in close proximity to one another, reducing the distance of electrical traces (not shown) used to connect the optical encoders 1-1103 and the homodyne detectors 1-1105 and a controller (not

shown) which may be physically located adjacent to the column of the optical encoders 1-1103 and the homodyne detectors 1-1105.

Each of the optical encoders 1-1103 is associated with a respective homodyne detector 1-1105. Both the optical encoders 1-1103 and the homodyne detectors 1-1105 receive optical signals from the power tree 1-1101. The optical encoders 1-1103 use the optical signals to encode an input vector, as described above. The homodyne detectors 1-1105 use the received optical signals received from the power tree as the local oscillator, as described above.

Each pair of the optical encoders 1-1103 and the homodyne detectors 1-1105 is associated with and connected to a selector switch 1-1107 by a waveguide. The selector switches 1-1107a – 1-1107d may be implemented using, for example, a conventional 2x2 optical switch. In some embodiments, the 2×2 optical switch is a MZI with an internal phase shifter to control the MZI's behavior from a crossing to a bar. The switch 1-1107 is connected to a controller (not shown) to control whether an optical signal received from the optical encoder 1-1103 will be guided towards the U-matrix components 1-1109 or the V-matrix components 1-1113. The optical switch is also controlled to guide light received from the U-matrix components 1-1109 and/or the V-matrix components 1-1113 toward the homodyne detectors 1-1105 for detection.

The techniques for implementing matrix multiplication is similar in the photonic folded photonic processing system 1-1100 as was described above in connection with the back-to-back system, described in FIG. 1-3. A difference between the two systems is in the physical placement of the matrix components and the implementation of a fold 1-1120, where the optical signals change from propagating approximately left to right in FIG. 1-11 to propagating approximately right to left. In FIG. 1-11, the connections between components may represent waveguides. The solid-lined connections represent portions of waveguide where the optical signals are propagating from left to right, in some embodiments, and the dashed-lined connections represent portions of waveguide where the optical signals are propagating from right to left, in some embodiments. In particular, given this nomenclature, the embodiment illustrated in FIG. 1-11 is an embodiment where the selector switches 1-1107 guide the optical signals to the U-matrix components 1-1109 first. In other embodiments, the selector switches 1-1107 may guide the optical signals to the V-matrix components 1-1113

first, in which case the dashed lines would represent portions of waveguide where the optical signals are propagating from left to right, and the solid-lined connections would represent portions of waveguide where the optical signals are propagating from right to left.

The U-matrix of the SVD of a matrix M is implemented in photonic processing system 1-1100 using U-matrix components 1-1109 that are interleaved with the V-matrix components 1-1113. Thus, unlike the embodiment of the back-to-back arrangement illustrated in FIG. 1-3, all of the U-matrix components 1-1109 and the V-matrix components 1-1113 are not physically located in a respective self-contained array within a single physical area. Accordingly, in some embodiments, the photonic processing system 1-1100 includes a plurality of columns of matrix components and at least one of the columns contains both U-matrix components 1-1109 and V-matrix components 1-1113. In some embodiments, the first column may only have U-matrix components 1-1109, as illustrated in FIG. 1-11. U-matrix components 1-1109 are implemented similarly to the first matrix implementation 1-301 of FIG. 3.

Due to the interleaving structure of the U-matrix components 1-1109 and the V-matrix components 1-1113, the folded photonic processing system 1-1100 includes waveguide crossovers 1-1110 at various locations between the columns of matrix elements. In some embodiments, the waveguide crossovers can be constructed using adiabatic evanescent elevators between two or more layers in an integrated photonics chip. In other embodiments, the U-matrix and the V-matrix may be positioned on different layers of the same chip and the waveguide crossovers are not used.

After optical signals propagate through all of the U-matrix components 1-1109, the optical signals propagate to the diagonal-matrix components 1-1111, which are implemented similarly to the second matrix implementation 1-303 of FIG. 1-3.

After optical signals propagate through all of the diagonal-matrix components 1-1111, the optical signals propagate to the V-matrix components 1-1113, which are implemented similarly to the third matrix implementation 1-305 of FIG. 1-3. The V-matrix of the SVD of a matrix M is implemented in photonic processing system 1-1100 using V-matrix components 1-1113 that are interleaved with the U-matrix components 1-1109. Thus, all of the V-matrix components 1-1113 are not physically located in a single self-contained array.

After the optical signals propagate through all of the V-matrix components 1-1113, the optical signals return to the selector switch 1-1107, which guides the optical signals to the homodyne detectors 1-1105 for detection.

The inventors have further recognized and appreciated that by including selector switches after the optical encoders and before the matrix components, the folded photonic processing system 1-1100 allows efficient bi-directionality of the circuit. Thus, in some embodiments, a controller, such as the controller 1-107 described in connection with FIG. 1-1, may control whether the optical signals are multiplied by the U matrix first or the V^T matrix first. For an array of VBSs set to implement a unitary matrix U when propagating the optical signals from left to right, propagating the optical signals from right to left implements a multiplication by a unitary matrix U^T . Thus, the same settings for an array of VBSs can implement both U and U^T depending which way the optical signals are propagated through the array, which may be controlled using the selector switch in 1-1107. In some applications, such as back-propagation used to train a machine learning algorithm, it may be desirable to run optical signals through one or more matrices backwards. In other applications, the bi-directionality can be used to compute the operation of an inverted matrix on an input vector. For example, for an invertible $n \times n$ matrix M , an SVD results in $M = V^T \Sigma U$. The inverse of this matrix is $M^{-1} = U^T \Sigma^{-1} V$, where Σ^{-1} is the inverse of a diagonal matrix which can be computed efficiently by inverting each diagonal element. To multiply a vector by the matrix M , the switches are configured to direct the optical signals through the matrix U , then Σ , then V^T in a first direction. To multiply a vector by the inverse M^{-1} , the singular values are first set to program the implementation of the Σ^{-1} matrix. This constitutes changing the settings of only one column of VBSs instead of all $2n + 1$ columns of the photonic processor, which is the case for a single-directional photonic processing system such as the one illustrated in FIG. 1-3. The optical signals representing the input vector are then propagated through the matrix V^T , then Σ^{-1} , and then U in a second direction that is opposite the first direction. Using the selector switches 1-1107, the folded photonic processing system 1-1100 may be easily changed from implementing the U matrix (or its transpose) first and implementing the V^T matrix (or its transpose) first.

G. Wavelength Division Multiplexing

The inventors have further recognized and appreciated that there are applications where different vectors may be multiplied by the same matrix. For example, when training or using machine learning algorithms sets of data may be processed with the same matrix multiplications. The inventors have recognized and appreciated that this may be accomplished with a single photonic processor if the components before and after the photonic processor are wavelength-division-multiplexed (WDM). Accordingly, some embodiments include multiple frontends and backends, each associated with a different wavelength, while only using a single photonic processor to implement the matrix multiplication.

FIG. 1-12A illustrates a WDM photonic processing system 1-1200, according to some embodiments. The WDM photonic processing system 1-1200 includes N frontends 1-1203, a single photonic processor 1-1201 with N spatial modes, and N backends 1-1205.

The photonic processor 1-1201 may be similar to the photonic processor 1-103, with N input modes and N output modes. Each of the N frontends 1-1203 is connected to a respective input mode of photonic processor 1-1201. Similarly, each of the N backends 1-1205 is connected to a respective output mode of photonic processor 1-1201.

FIG. 1-12B illustrates details of at least one of the frontends 1-1203. As with the photonic processing system of other embodiments, the photonic processing system 1-1200 includes optical encoders 1-1211. But in this embodiment, there are M different optical encoders, where M is the number of wavelengths being multiplexed by the WDM photonic processing system 1-1200. Each of the M optical encoders 1-1211 receives light from a light source (not shown) that generates the M optical signals, each of a different wavelength. The light source may be, for example, an array of lasers, a frequency comb generator or any other light source that generates coherent light at different wavelengths. Each of the M optical encoders 1-1211 is controlled by a controller (not shown) to implement an appropriate amplitude and phase modulation to encode data into the optical signal. Then, the M encoded optical signals are combined into a single waveguide using an M:1 WDM 1-1213. The single waveguide then connects to the one of the N waveguides of the photonic processor 1-1201.

FIG. 1-12C illustrates details of at least one of the backends 1-1205. As with the photonic processing system of other embodiments, the photonic processing system 1-1200

includes detectors 1-1223, which may be phase-sensitive or phase-insensitive detectors. But in this embodiment, there are M different detectors 1-1223, where M is the number of wavelengths being multiplexed by the WDM photonic processing system 1-1200. Each of the M detectors 1-1223 receives light from a 1: M WDM 1-1221, which splits the single output waveguide from photonic processor 1-1201 into M different waveguides, each carrying an optical signal of a respective wavelength. Each of the M detectors 1-1223 may be controlled by a controller (not shown) to record measurement results. For example, each of the M detectors 1223 may be a homodyne detector or a phase insensitive photodetector.

In some embodiments, the VBSs in the photonic processor 1-1201 may be chosen to be non-dispersive within the M wavelengths of interest. As such, all the input vectors are multiplied by the same matrix. For example, an MMI can be used instead of a directional coupler. In other embodiments, the VBSs may be chosen to be dispersive within the M wavelengths of interest. In some applications related to stochastic optimization of the parameters of a neural network model, this is equivalent to adding noise when computing the gradient of the parameters; increased gradient noise may be beneficial for faster optimization convergence and may improve the robustness of a neural network.

While FIG. 1-12A illustrates a back-to-back photonic processing system, similar WDM techniques may be used to form a WDM folded photonic processor using the techniques described in relation to folded photonic processor 1-1100.

H. Analog Summation of Outputs

The inventors have recognized and appreciated that there are applications where it is useful to calculate the sum or the average of the outputs from the photonic processor 1-103 over time. For example, when the photonic processing system 1-100 is used to compute a more exact matrix-vector multiplication for a single data point, one may want to run a single data point through the photonic processor multiple times to improve the statistical results of the calculation. Additionally or alternatively, when computing the gradient in a backpropagation machine learning algorithm, one may not want a single data point determining the gradient, so multiple training data points may be run through photonic processing system 1-100 and the average result may be used to calculate the gradient. When using a photonic processing system to perform a batched gradient based optimization

algorithm, this averaging can increase the quality of the gradient estimate and thereby reduce the number of optimization steps required to achieve a high quality solution.

The inventors have further recognized and appreciated that the output signals may be summed in the analog domain, before converting the outputs to digital electrical signals. Thus, in some embodiments, a low pass filter is used to sum the outputs from the homodyne detectors. By performing the summation in the analog domain, the homodyne electronics may use a slow ADC rather than a costlier fast ADC (e.g., an ADC with high power consumption requirements) that would be required to perform a summation in the digital domain.

FIG. 1-13 illustrates a portion of an optical receiver 1-1300 and how a low pass filter 1-1305 may be used with a homodyne detector 1-1301, according to some embodiments. The homodyne detector 1-1301 performs a measurement of the field and phase of an incoming optical pulse. If k is the label for the different input pulses over time and there is a total of K inputs, the sum over k can be automatically performed in the analog domain using low-pass filter 1-1305. The main difference between this optical receiver 1-1300 and the optical receiver 1-105 illustrated in FIG. 1-9 is that the low-pass filter is after the transimpedance amplifier 1-1303 after the output of the homodyne detector. If a total of K signals (with components $y_i^{(k)}$) arrives at the homodyne detector within a single slow sampling period $T_s^{(slow)}$, the low-pass filter will have accumulated/removed the charges in the capacitor C according to the sign and value of $y_i^{(k)}$. The final output of the low-pass filter is proportional to $Y_i = \sum_{k=1}^K y_i^{(k)}$, which can be read *once* with a slower ADC (not shown) with a sampling frequency of $f_s^{(slow)} = 1/T_s^{(slow)} = f_s/K$, where f_s is the originally required sampling frequency. For an ideal system, the low-pass filter should have a 3-dB bandwidth: $f_{3\text{ dB}} = f_s^{(slow)}/2$. For a low-pass filter using an RC circuit as shown in the embodiment of FIG. 1-13, $f_{3\text{ dB}} = 1/(2\pi RC)$, and the values of R and C can be chosen to obtain the desired sampling frequency: $f_s^{(slow)}$.

In some embodiments both a fast ADC and a slow ADC may be present. In this context, a fast ADC is an ADC that is configured to receive and convert each individual analog signal into a digital signal (e.g., an ADC with a sampling frequency equal to or greater than the frequency at which the analog signals arrive at the ADC), and a slow ADC is an ADC that is configured to receive multiple analog signals and convert the sum or average of multiple received analog signals into a single digital signal (e.g., an ADC with a sampling

frequency less than the frequency at which the analog signals arrive at the ADC). An electrical switch may be used to switch the electrical signal from the homodyne detector and possibly transimpedance amplifier to the low-pass filter with a slow ADC or to the fast ADC. In this way, the photonic processing system of some embodiments may switch between performing analog summation using the slow ADC and measuring every optical signal using the fast ADC.

I. Stabilizing Phases

The inventors have recognized and appreciated that it is desirable to stabilize the phase of the local oscillator used for performing phase-sensitive measurements (e.g., homodyne detection) to ensure accurate results. The photonic processors of the embodiments described herein perform matrix operations by interfering light between N distinct spatial modes. The results are measured, in some embodiments, with phase sensitive detectors, such as homodyne or heterodyne detectors. Thus, to ensure the matrix operations are accurately performed, the phase imparted at various portions of the photonic processor should be as accurate as possible and the phase of the local oscillator used to perform phase-sensitive detection should be precisely known.

The inventors have recognized and appreciated that parallel interference operations, such as those performed within a single column of VBSs of the photonic processor, must not only impart the correct phases using the phase modulators controlling the relative phase within the MZI of the VBS and the phase and the relative phase of the output of the MZI, but each VBS in a column should impart the same global phase shift across all the spatial modes of photonic processor. In this application, the global phase shift for a column of VBSs in the photonic processor is referred to as the “column-global phase.” The column-global phase is the phase imparted due to effects not related to the programmed phases associated with the VBS, such as phases imparted due to propagation through the waveguide or phases due to temperature shifts. These phases need not be imparted exactly simultaneously within a column on VBSs, but only need be imparted as a result of traversing the column in question. Ensuring the column-global phase is uniform between the different spatial modes of the column is important because the output optical signals from one column will likely be interfered at one or more VBSs at a subsequent column. The subsequent interference – and

therefore the accuracy of the calculation itself – would be incorrect if the column-global phase at the previous columns is not uniform.

FIG. 1-14 illustrates the column-global phases and total global phase for a photonic processing system 1-1400. Similar to the above-described embodiments of photonic processing systems, the photonic processing system 1-1400 includes a U-matrix implementation 1-1401, a diagonal matrix implementation 1-1403, a V-matrix implementation 1-1405, and a plurality of detectors 1-1407a – 1-1407d. These implementations are similar to the first, second and third matrix implementations described above. For the sake of simplicity, only four modes of the photonic processing system 1-1400 are shown, though it should be understood that any larger number of modes may be used. Also for simplicity, only the VBSs associated with the U-matrix implementation 1-1401 are illustrated. The arrangement of components of the diagonal matrix implementation 1-1403 and a V-matrix implementation 1-1405 are similar to the third and fourth matrix implementations described above.

The U-matrix implementation 1-1401 includes a plurality of VBSs 1-1402, though only a single VBS 1-1402 is labeled for the sake of clarity. The VBSs are labeled, however, with subscripts that identify which optical modes are being mixed by a particular VBS and a superscript labeling the associated column.

As illustrated in FIG. 1-14, each column is associated with a column-global phase that is ideally uniform for every element of the column. For example, column 1 of the U-matrix implementation 1-1401 is associated with a column-global phase ϕ_{U_1} , column 2 of the U-matrix implementation 1-1401 is associated with a column-global phase ϕ_{U_2} , column 3 of the U-matrix implementation 1-1401 is associated with a column-global phase ϕ_{U_3} , and column 4 of the U-matrix implementation 1-1401 is associated with a column-global phase ϕ_4 .

In some embodiments, the column-global phases can be made uniform at least in part by implementing each VBS 1-1402 as a MZI in a push-pull configuration. Alternatively or additionally, external phase shifter can be added to the output of each MZI to correct for any phase error imparted from the internal phase elements of the MZIs (e.g., the phase shifters).

The inventors have further recognized and appreciated that even if the conditions are such that each column of the photonic processing system 1-1400 provides a uniform column-global phase, phases can be accrued as the signal propagates from the first column to the last. There is a global U-matrix phase, Φ_U , associated with the entire U-matrix implementation 1-

1401 and is equal to the sum of the individual column-global phase. Similarly, the diagonal-matrix implementation 1-1403 is associated with a global diagonal-matrix phase, Φ_S , and the V-matrix implementation 1-1405 is associated with a global diagonal-matrix phase, $\Phi_{V^\dagger}$. A total global phase Φ_G for the entire photonic processing system 1-1400 is then given by the sum of the three individual global matrix phases. This total global phase may be set to be uniform between all the output modes, but the local oscillator that is used for phase-sensitive detection did not propagate through the photonic processor and did not experience this total global phase. The total global phase Φ_G , if not accounted for, can lead to an error in the values read out by the homodyne detectors 1-1407a – 1-1407d.

The inventors have further recognized that errors in the multiplication operation may result from changes in temperature, which change a waveguide's effective refractive index n_{eff} . Accordingly, in some embodiments, either the temperature of each column is set to be uniform or stabilization circuits can be placed at each column such that the phases imparted to all the modes of a single column are actively tuned to be uniform. Additionally, as the light signal for the local oscillator propagates through a different part of the system, the temperature difference between different parts of the system can cause errors in the phase-sensitive measurements. The amount of phase difference between the signal and the local oscillator is $\Phi_T = \frac{2\pi}{\lambda} (n_{eff}(T_S)L_S - n_{eff}(T_{LO})L_{LO})$, where T_S and T_{LO} are the temperatures of the signal waveguide in the photonic processor and the local oscillator waveguide, respectively, $n_{eff}(T)$ is the effective index of refraction as a function of temperature, λ is the average wavelength of the light, and L_S and L_{LO} are the propagation lengths through the signal waveguide in the photonic processor and the local oscillator waveguide, respectively. Assuming that the difference in temperature $\Delta T = T_{LO} - T_S$ is small, then the effective index can be rewritten as: $n_{eff}(T_{LO}) \approx n_{eff}(T_S) + \frac{dn_{eff}}{dT} \big|_{T=T_S} \Delta T$. Therefore, the phase difference between the signal and the LO can be well approximated by $\Phi_T = \frac{2\pi}{\lambda} \frac{dn_{eff}}{dT} \big|_{T=T_S} L \cdot \Delta T$, which increases linearly with longer propagation length L . Therefore, for a sufficiently long propagation distance, a small change in temperature can result in a large phase shift (on the order of one radian). Importantly, the values of L_S does not need to be the same as the value of L_{LO} , and the maximum difference between the two is determined by the coherence length of the light source L_{coh} . For a light source with a bandwidth of $\Delta\nu$, the coherence length can

be well approximated by $L_{coh} \approx c_{eff} \Delta \nu$, where c_{eff} is the speed of light in the transmission medium. As long as the length difference between L_S and L_{LO} is much shorter than L_{coh} , interference between the signal and the local oscillator will be possible for the correct operation of the photonic processing system.

Based on the foregoing, the inventors have identified at least two sources of possible phase errors between the output signals of the photonic processor and the local oscillator used for homodyne detection in some embodiments. Thus, where an ideal homodyne detector would measure the magnitude and phase of the signal output by subtracting the outputs of the two photodetectors, resulting in a phase sensitive intensity output measurement of $I_{out} \propto |E_s| |E_{LO}| \cos(\theta_s - \theta_{LO} + \Phi_G + \Phi_T)$, where E_s is the electric field magnitude of the optical signal from the output of the photonic processor, E_{LO} is the electric field magnitude of the local oscillator, θ_s is the phase shift imparted by the photonic processor that is desired to be measured, Φ_G is the total global phase, and Φ_T is the phase shift caused by temperature differences between the local oscillator and the optical signal. Consequently, if the total global phase and the phase shift due to temperature differences are not accounted for, the result of the homodyne detection can be erroneous. Therefore, in some embodiments the total systematic phase error, $\Delta\Phi = \Phi_G + \Phi_T$, is measured and the system is calibrated based on that measurement. In some embodiments, the total systematic phase error includes contributions from other sources of error that are not necessarily known or identified.

According to some embodiments, the homodyne detectors may be calibrated by sending pre-computed test signals to the detectors and using the difference between the pre-computed test signals and the measured test signals to correct for the total systematic phase error in the system.

In some embodiments, rather than considering the total global phase, Φ_G , and the phase shift caused by temperature differences, Φ_T , as being related to the optical signals propagating through the photonic processor, they can be described as the signal not accruing any phase shift at all but the LO having a total systematic phase error $-\Delta\Phi$. FIG. 1-15 illustrates the effect on the results of the homodyne measurements in such a situation. The original (correct) vector of quadrature values of the signal $[x, p]^T$ is rotated by a rotation matrix parameterized by $\Delta\Phi$ producing an incorrect quadrature values of $[x', p']^T$.

Based on the rotation in quadrature due to the total systematic error, in some embodiments, the value of $\Delta\Phi$ is obtained as follows. First, a vector $\vec{v}_{in}$ is selected (e.g., a random vector), using, e.g., the controller 1-107. The vector is of a type that can be prepared by the optical encoders of the photonic processing system. Second, the output value of $\vec{v}_{out} = M \vec{v}_{in}$, where M is the matrix implemented by the photonic processor in the ideal case assuming that there is no unaccounted phase accrued of $\Delta\Phi$, is calculated using, for example, the controller 1-107 or some other computing device. As a result, each element of $\vec{v}_{out}$ corresponds to $x_k + ip_k$, where k labels each of the output modes of the photonic processor.

In some embodiments, loss in propagating the random vector through the photonic processor may be considered when calculating the theoretical prediction $x_k + ip_k$. For example, for a photonic processor with transmission efficiency η , the field signal of $x_k + ip_k$ will become $\sqrt{\eta}(x_k + ip_k)$.

Next, the random vector $\vec{v}_{in}$ is prepared by the optical encoder of the actual system, propagated through the photonic processor, and each element of the output vector is measured in both quadratures to obtain $x_k' + ip_k'$. The phase difference $\Delta\Phi_k$ between the local oscillator and the signal of output mode k is given by $\Delta\Phi_k = \tan^{-1} \left(\frac{p_k' x_k - x_k' p_k}{x_k x_k' + p_k p_k'} \right)$. (Generally, the phase difference $\Delta\Phi_k \neq \Delta\Phi_l$ for $k \neq l$ as the path length of the LO to the detector for mode k can be different to that for mode l).

Finally, the local oscillator phase shifter used to select the measurement quadrature of the homodyne detector is controlled to impart $\theta_{LO,k} = \Delta\Phi_k$. As a result, the axes (x, p) will align with the axes (x', p') , as illustrated in FIG. 1-15. The calibration may be checked at this stage to ensure it is accurate by propagating the vector $\vec{v}_{in}$ once again to see that the obtained measurement results are equal to the predicted $\vec{v}_{out}$ when both quadratures are measured.

Generally, the value of $\Delta\Phi_k$ can be determined more precisely if the field amplitude $|E_{S,k}| = \sqrt{x_k^2 + p_k^2} = \sqrt{x_k'^2 + p_k'^2}$ is as large as possible. For example, if the field $E_{S,k}$ is considered to be a coherent signal, e.g. from a laser source, then the optical signal may be theoretically modeled as a coherent state. The intuitive picture is given in FIG. 1-16, where the signal is the amplitude $|E_{S,k}|$ and the noise is given by the standard deviation of the Gaussian coherent state. The coherent state $|\alpha_k\rangle$ in mode k is the eigenstate of the annihilation operator a_k , i.e. $a_k|\alpha_k\rangle = \alpha_k|\alpha_k\rangle$. The electric field of mode k with a single

frequency ω is described by $E_{S,k}^{(+)}(t) = a_k e^{-i\omega t}$, which is also an eigenstate of the coherent state: $E_{S,k}^{(+)}(t)|\alpha_k\rangle = \alpha_k e^{-i\omega t}|\alpha_k\rangle$. A homodyne detector with a local oscillator of the same frequency ω performs the quadrature measurements $x_k = (a_k + a_k^\dagger)/2$ when $\theta_{LO} = 0$ and $p_k = (a_k - a_k^\dagger)/2i$ when $\theta_{LO} = \pi/2$. An ideal homodyne detector will find that these measurements have an intrinsic quantum noise of $\sqrt{\langle \Delta x_k^2 \rangle} = 1/2$ and $\sqrt{\langle \Delta p_k^2 \rangle} = 1/2$. This noise is related to the quantum uncertainties, and it can be reduced by squeezing on the quadratures. The precision at which the angle $\Delta\Phi_k$ can be determined is directly related to the signal-to-noise ratio (SNR) of these measurements. For a coherent-state signal $E_{S,k}$ with a total of N_{ph} photons (i.e. $E_{S,k} = |\sqrt{N_{ph}}e^{i\theta_S}\rangle$, the SNR of both x_k and p_k is upper bounded by: $SNR_x = \frac{\langle x_k \rangle^2}{\langle \Delta x_k^2 \rangle} \leq 4N_{ph}$ and $SNR_p = \frac{\langle p_k \rangle^2}{\langle \Delta p_k^2 \rangle} \leq 4N_{ph}$. (The bound of SNR_x is saturated when $\theta_S = 0$ or π , and the bound on SNR_p is saturated when $\theta_S = \pi/2$ or $3\pi/2$). Therefore, to increase the SNR and to determine the values of $\Delta\Phi_k$ more accurately, some embodiments may propagate a few different choices of vector $\vec{v}_{in}$ (e.g., multiple different random vectors). In some embodiments, the choices of $\vec{v}_{in}$ are chosen to maximize the amplitude $|E_{S,k}| = N_{ph}$ for one value of k at a time.

There may be phase drift during the operation of the photonic processing system, e.g., due to temperature fluctuations over time. Thus, in some embodiments, the aforementioned calibration procedure may be performed repeatedly during the operation of the system. For example, in some embodiments, the calibration procedure is performed regularly at a time scale that is shorter than the natural timescale of the phase drift.

The inventors have further recognized and appreciated that it is possible to perform signed matrix operations without the need of phase-sensitive measurements at all. Therefore, in applications, each homodyne detector at each output mode may be replaced by a direct photodetector which measures the intensity of the light at that output mode. As there is no local oscillator in such a system, the systematic phase error $\Delta\Phi$ is non-existent and meaningless. Thus, according to some embodiments, phase-sensitive measurements, such as homodyne detection, may be avoided such that the systematic phase error is insignificant. For example, when computing matrix operations of signed matrices and vectors, complex matrices and vectors, and hypercomplex (quaternion, octonion, and other isomorphisms (e.g.,

elements of unital algebra)) matrices and vectors using unsigned matrices do not require phase-sensitive measurements.

To illustrate how phase-sensitive measurements are not necessary, consider the case of performing matrix multiplication between a signed matrix M and a signed vector $\vec{v}_{in}$. To compute the value of signed output $\vec{v}_{out} = M \vec{v}_{in}$, the following procedure may be performed by, for example, the controller 1-107. First, the matrix M is split into M_+ and M_- , where M_+ (M_-) is a matrix that contains all the positive (negative) entries of M . In this case, $M = M_+ - M_-$. Second, the vector is split in a similar manner such that the vector $\vec{v}_{in} = \vec{v}_{in,+} - \vec{v}_{in,-}$, where $\vec{v}_{in,+}$ ($\vec{v}_{in,-}$) is a vector that contains all the positive (negative) entries of $\vec{v}_{in}$. As a result of the splittings, $\vec{v}_{out} = M \vec{v}_{in} = (M_+ - M_-)(\vec{v}_{in,+} - \vec{v}_{in,-}) = (M_+ \vec{v}_{in,+} + M_- \vec{v}_{in,-}) - (M_+ \vec{v}_{in,-} + M_- \vec{v}_{in,+})$. Each term of this final equation corresponds to a separate operation ($M_+ \vec{v}_{in,+}$, $M_- \vec{v}_{in,-}$, $M_+ \vec{v}_{in,-}$, and $M_- \vec{v}_{in,+}$) that may be performed individually by the photonic processing system. The output of each operation is a vector of a single (positive) sign, and therefore can be measured using a direct detection scheme without the need for homodyne detection. The photodetector scheme will measure the intensity, but the square root of the intensity may be determined, resulting in the electric field amplitude. In some embodiments, each operation is performed separately and the results are stored in a memory (e.g., memory 1-109 of controller 1-107) until all of the separate operations are performed and the results may be digitally combined to obtain the final result of the multiplication, $\vec{v}_{out}$.

The above scheme works because M_+ and M_- are both matrices of all positive entries. Similarly, $\vec{v}_{in,+}$ and $\vec{v}_{in,-}$ are both vectors of all positive entries. Therefore, the results of their multiplications will be vectors of all positive entries—regardless of the combination.

The inventors have further recognized and appreciated that the above splitting technique may be extended to complex-valued vectors/matrices, quaternion-valued vectors/matrices, octonion-valued vectors/matrices, and other hypercomplex representations. Complex numbers employ two different fundamental units $\{1, i\}$, Quaternions employ four different fundamental units $\{1, i, j, k\}$, and octonions employ eight fundamental units $\{e_0 \equiv 1, e_1, e_2, \dots, e_7\}$.

In some embodiments, a complex vector may be multiplied by a complex matrix without the need for phase-sensitive detection by splitting the multiplication into separate operations similar to the procedure described above for signed matrices and vectors. In the

case of complex numbers, the multiplication splits into 16 separate multiplications of all-positive matrices and all-positive vectors. The results of the 16 separate multiplications may then be digitally combined to determine the output vector result.

In some embodiments, a quaternion-valued vector may be multiplied by a quaternion-valued matrix without the need for phase-sensitive detection by splitting the multiplication into separate operations similar to the procedure described above for signed matrices and vectors. In the case of quaternion-valued numbers, the multiplication splits into 64 separate multiplications of all-positive matrices and all-positive vectors. The results of the 64 separate multiplications may then be digitally combined to determine the output vector result.

In some embodiments, a octonion-valued vector may be multiplied by a octonion -valued matrix without the need for phase-sensitive detection by splitting the multiplication into separate operations similar to the procedure described above for signed matrices and vectors. In the case of octonion -valued numbers, the multiplication splits into 256 separate multiplications of all-positive matrices and all-positive vectors. The results of the 256 separate multiplications may then be digitally combined to determine the output vector result.

The inventors have further recognized and appreciated that temperature-dependent phase Φ_T can be corrected by placing a temperature sensor next to each MZI of the photonic processor. The results of the temperature measurement may then be used as an input to a feedback circuitry that controls the external phases of each MZI. The external phases of the MZI are set to cancel the temperature-dependent phase accrued at every MZI. A similar temperature feedback loop can be used on the local oscillator propagation path. In this case, the temperature measurement results are used to inform the settings of the homodyne detector quadrature-selecting phase shifter to cancel the phase accrued by the local oscillator due to detected temperature effects.

In some embodiments, the temperature sensors can be those conventionally used in semiconductor devices, e.g. p-n junction or bipolar junction transistor, or they can be photonic temperature sensors, e.g. using resonators whose resonance changes with temperatures. External temperature sensors such as thermocouples or thermistors may also be used in some embodiments.

In some embodiments, the phases accrued may be directly measured by, for example, tapping some light at every column and performing homodyne detection with the same global

local oscillator. This phase measurement can directly inform the values of external phases used at each MZI to correct for any phase error. In the case of directly measured phase errors, the errors do not need to be column-global to be corrected.

J. Intermediary Computation for Large Data

The inventors have recognized and appreciated that the matrix vector product performed by the photonic processor 1-103, and/or any other photonic processor according to other embodiments described in the present disclosure, can be generalized into tensor (multidimensional array) operations. For example, the core operation of $M\vec{x}$ where M is a matrix and $\vec{x}$ is a vector can be generalized into a matrix-matrix product: MX where both M and X are matrices. In this particular example, consider the n -by- m matrix X to be a collection of m column vectors each consisting of n elements, i.e. $X = [\vec{x}_1, \vec{x}_2, \dots, \vec{x}_m]$. A photonic processor can complete the matrix-matrix product MX one column vector at a time with a total of m matrix-vector products. The computation can be distributed among multiple photonic processors as the computation is a linear operation, which is perfectly parallelizable, e.g., any one matrix-vector product output does not depend on the results of the other matrix-vector products. Alternatively, the computation can be performed by a single photonic processor serially over time, e.g., by performing each matrix-vector product one at a time and combining the results digitally after performing all of the individual matrix-vector multiplications to determine the result of the matrix-matrix product (e.g., by storing the results in an appropriate memory configuration).

The concept above can be generalized into computing a product (e.g., a dot product) between two multidimensional tensors. The general algorithm is as follows and may be performed, at least in part, by a processor such as the processor 1-111: (1) Take a matrix slice of the first tensor; (2) Take a vector slice of the second tensor; (3) Perform a matrix-vector product, using the photonic processor, between the matrix slice in step 1 and the vector slice in step 2, resulting in an output vector; (4) Iterate over the tensor indices from which the matrix slice (from step 1) was obtained and the tensor indices from which the vector slice (from step 2) was obtained. It should be noted that when taking the matrix slice and the vector slice (steps 1 and 2), multiple indices can be combined into one. For example, a matrix can be vectorized by stacking all the columns into a single column vector, and in general a

tensor can be matricized by stacking all the matrices into a single matrix. Since all the operations are fully linear, they are again can be highly parallelized where each of a plurality of photonic processor does not need to know whether the other photonic processors have completed their jobs.

By way of a non-limiting example, consider the multiplication between two three-dimensional tensors $C_{ijlm} = \sum_k A_{ijk} B_{klm}$. The pseudocode based on the prescription above is as follows:

- (1) Take a matrix slice: $A_i \leftarrow A[i, :, :]$;
- (2) Take a vector slice: $\overrightarrow{b_{lm}} \leftarrow B[:, l, m]$;
- (3) Compute $\overrightarrow{c_{ilm}} = A_i \overrightarrow{b_{lm}}$, where $C[i, :, l, m] \leftarrow \overrightarrow{c_{ilm}}$; and
- (4) Iterate over the indices i, l , and m to reconstruct the four-dimensional tensor C_{ijlm} ,

where the values of all elements indexed by j is fully determined with a single matrix-vector multiplication.

The inventors have further recognized and appreciated that the size of the matrices/vectors to be multiplied can be larger than the number of modes supported by the photonic processor. For example, a convolution operation in a convolutional neural network architecture may use only a few parameters to define a filter, but may consist of a number of matrix-matrix multiplications between the filter and different patches of the data. Combining the different matrix-matrix multiplications result in two input matrices that are larger than the size of the original filter matrix or data matrix.

The inventors have devised a method of performing matrix operations using the photonic processor when the matrices to be multiplied are larger than the size/the number of modes possessed by the photonic processor being used to perform the calculation. In some embodiments, the method involves using memory to store intermediate information during the calculation. The final calculation result is computed by processing the intermediate information. For example, as illustrated in FIG. 1-17, consider the multiplication 1-1700 between an $I \times J$ matrix A and a $J \times K$ matrix B to give a new matrix $C = AB$, which has $I \times K$ elements, using an $n \times n$ photonic processing system with $n \leq I, J, K$. In FIG. 1-17, the shaded elements simply illustrate that the element 1-1701 of matrix C is calculated using the elements of row 1-1703 of matrix A and column 1-1705 of matrix B . The method illustrated by FIG. 1-17 and 1-18 is as follows:

Construct $n \times n$ submatrix blocks of within matrices A and B . Label the blocks by the parenthesis superscript $A^{(ij)}$ and $B^{(jk)}$, where $i \in \{1, \dots, \text{ceil}(I/n)\}$, $j \in \{1, \dots, \text{ceil}(J/n)\}$, and $k \in \{1, \dots, \text{ceil}(K/n)\}$. When the values of I, J , or K are not divisible by n , the matrices may be padded with zeros such that the new matrix has dimensions that are divisible by n —hence the ceil function in the indexing of i, j , and k . In the example multiplication 1-1800 illustrated in FIG. 1-18, matrix A is split into six $n \times n$ submatrix blocks 1-1803 and matrix B is split into three $n \times n$ submatrix blocks 1-1805, resulting in a resulting matrix C that is comprised of two $n \times n$ submatrix blocks 1-1801.

To compute the $n \times n$ submatrix block $C^{(ik)}$ within matrix C , perform the multiplications $C^{(ik)} = \sum_{j=1}^{\text{ceil}(J/n)} A^{(ij)} B^{(jk)}$ in the photonic processor by, for example:

- (1) Controlling the photonic processor to implement the submatrix $A^{(ij)}$ (e.g., one of the submatrices 1-1803);
- (2) Encoding optical signals with the column vectors of one of the submatrices $B^{(jk)}$ (e.g., one of the submatrices 1-1805) and propagating the signals through the photonic processor;
- (3) Storing the intermediate results of each matrix-vector multiplication in memory;
- (4) Iterating over the values of j , repeating steps (a)–(c); and
- (5) Computing the final submatrix $C^{(ik)}$ (e.g., one of the submatrices 1-1801) by combining the intermediate results with digital electronics, e.g., a processor.

As described above and shown in FIG. 1-17 and FIG. 1-18, the method may include expressing the matrix multiplication using parentheses index notation and performing the operation of the matrix-matrix multiplication using the parentheses superscript indices instead of the subscript indices, which are used to describe the matrix elements in this disclosure. These parentheses superscript indices correspond to the $n \times n$ block of submatrix. In some embodiments, the method can be generalized to tensor-tensor multiplications by breaking up the multidimensional arrays into $n \times n$ submatrix block slices, for example, by combining this method with the tensor-tensor multiplication described above.

In some embodiments, an advantage of processing blocks of submatrices using a photonic processor with fewer number of modes is that it provides versatility with regards to the shape of the matrices being multiplied. For example, in a case where $I \gg J$, performing

singular value decompositions will produce a first unitary matrix of size I^2 , a second unitary matrix of size J^2 , and a diagonal matrix with J parameters. The hardware requirements of storing or processing I^2 matrix elements, which are much larger than the number of elements of the original matrix, can be too large for the number of optical modes included in some embodiments of the photonic processor. By processing submatrices rather than the entire matrix all at once, any size matrices may be multiplied without imposing limitations based on the number of modes of the photonic processor.

[0056] In some embodiments, the submatrices of B are further vectorized. For example, the matrix A may be first padded to a $[(n \cdot \lceil I/n \rceil) \times (n \cdot \lceil J/n \rceil)]$ matrix and then partitioned into a $[\lceil I/n \rceil \times \lceil J/n \rceil]$ grid of submatrices (each of size $[n \times n]$) and $A^{(ij)}$ is the $[n \times n]$ submatrix in the i^{th} row and j^{th} column of this grid, B has been first padded to a $[(n \cdot \lceil J/n \rceil) \times K]$ matrix and then partitioned into a $[\lceil J/n \rceil \times 1]$ grid of submatrices (each of size $[n \times K]$) and $B^{(j)}$ is the $[n \times K]$ submatrix in the j^{th} row of this grid, and C has been first padded to a $[(n \cdot \lceil J/n \rceil) \times K]$ matrix and then partitioned into a $[\lceil I/n \rceil \times 1]$ grid of submatrices (each of size $[n \times K]$) and $C^{(i)}$ is the $[n \times K]$ submatrix in the i^{th} row of this grid. In this vectorized form, the computation is denoted by: $C^{(i)} = \sum_{j=1}^{\lceil J/n \rceil} A^{(ij)} B^{(j)}$.

Using the above vectorization process, a photonic processor can compute any GEMM by loading $(\lceil I/n \rceil \cdot \lceil J/n \rceil)$ different matrices into the photonic array and, for each loaded matrix, propagating K different vectors through the photonic array. This yields $\lceil I/n \rceil \cdot \lceil J/n \rceil \cdot K$ output vectors (each comprised of n elements), a subset of which may be added together to yield the desired $[I \times K]$ output matrix, as defined by the equation above.

K. Precision of the Computation

The inventors have recognized and appreciated that the photonic processor 1-103, and/or any other photonic processor according to other embodiments described in the present disclosure, is an instance of analog computer and, as most data in this information age are stored in a digital representation, the digital precision of the computation performed by the photonic processor is important to quantify. In some embodiments, the photonic processor according to some embodiments performs a matrix-vector product: $\vec{y} = M\vec{x}$, where $\vec{x}$ is the input vector, M is an $n \times n$ matrix, and $\vec{y}$ is the output vector. In index notation, this multiplication is written as $y_i = \sum_{j=1}^n M_{ij}x_j$ which is the multiplication between n elements of

M_{ij} (iterate over j) and n elements of x_j (iterate over j) and then summing the results altogether. As the photonic processor is a physical analog system, in some embodiments the elements M_{ij} and x_j are represented with a fixed point number representation. Within this representation, if $M_{ij} \in \{0,1\}^{m_1}$ is an m_1 -bit number and $x_j \in \{0,1\}^{m_2}$ is an m_2 -bit number, then a total of $m_1 + m_2 + \log_2(n)$ bits are used to fully represent the resulting vector element y_i . In general, the number of bits used to represent the result of a matrix-vector product is larger than the number of bits required to represent the inputs of the operation. If the analog-to-digital converter (ADC) used is unable to read out the output vector at full precision, then the output vector elements may be rounded to the precision of the ADC.

The inventors have recognized and appreciated that constructing an ADC with a high bit-precision at bandwidths that correspond to the rate at which input vectors in the form of optical signals are sent through the photonic processing system can be difficult to achieve. Therefore, in some embodiments, the bit precision of the ADC may limit the bit precision at which the matrix elements M_{ij} and the vector element x_j are represented (if a fully precise computation is desired). Accordingly, the inventors have devised a method of obtaining an output vector at its full precision, which can be arbitrarily high, by computing partial products and sums. For the sake of clarity, it will be assumed that the number of bits needed to represent either M_{ij} or x_j is the same, i.e. $m_1 = m_2 = m$. However, this assumption can be obviated in general and does not limit the scope of embodiments of the present disclosure.

The method, according to some embodiments, as a first act, includes dividing the bit-string representation of the matrix element M_{ij} and the vector element x_j into d divisions with each division containing $k = m/d$ bits. (If k is not an integer, zeros may be appended until m is divisible by d .) As a result, the matrix element $M_{ij} = M_{ij}^{[0]}2^{k(d-1)} + M_{ij}^{[1]}2^{k(d-2)} + \dots + M_{ij}^{[d-1]}2^0$, where $M_{ij}^{[a]}$ is the k -bit value of the a -th most significant k -bit string of M_{ij} . In terms of bit string, one writes $M_{ij} = M_{ij}^{[0]}M_{ij}^{[1]} \dots M_{ij}^{[d-1]}$. Similarly, one can also obtain $x_j = x_j^{[0]}2^{k(d-1)} + x_j^{[1]}2^{k(d-2)} + \dots + x_j^{[d-1]}2^0$, where the vector element $x_j = x_j^{[0]}x_j^{[1]} \dots x_j^{[d-1]}$ in terms of its bit string. The multiplication $y_i = \sum_j M_{ij}x_j$ can be broken

down in terms of these divisions as: $y_i = \sum_{p=0}^{2^{(d-1)}} ((\sum_{a,b \in S_p} \sum_j M_{ij}^{[a]} x_j^{[b]}) 2^{2k(d-1)-pk})$, where the set S_p is the set of all integer values of a and b , where $a + b = p$.

The method, as a second act, includes controlling the photonic processor to implement the matrix $M_{ij}^{[a]}$ and propagating the input vector $x_j^{[b]}$, each of which is only k -bit precise, through the photonic processor in the form of encoded optical signals. This matrix-vector product operation performs $y_i^{[a,b]} = \sum_j M_{ij}^{[a]} x_j^{[b]}$. The method includes, storing the output vector $y_i^{[a,b]}$ which is precise up to $2k + \log_2(n)$ bits.

The method further includes iterating over the different values of a, b within the set S_p and repeating the second act for each of the different values of a, b and storing the intermediate results $y_i^{[a,b]}$.

As a third act, the method includes computing the final result $\sum_{a,b \in S_p} \sum_j M_{ij}^{[a]} x_j^{[b]} = \sum_{a,b \in S_p} y_i^{[a,b]}$ by summing over the different iterations of a and b with digital electronics, such as a processor.

The precision of the ADC used to capture a fully precise computation according to some embodiments of this method is only $2k + \log_2(n)$ bits, which is fewer than the $2m + \log_2(n)$ bits of precision needed if the computation is done using only a single pass.

The inventors have further recognized and appreciated that embodiments of the foregoing method can be generalized to operate on tensors. As previously described, the photonic processing system can perform tensor-tensor multiplications by using matrix slices and vector slices of the two tensors. The method described above can be applied to the matrix slices and vector slices to obtain the output vector slice of the output tensor at full precision.

Some embodiments of the above method use the linearity of the elementary representation of the matrix. In the description above, the matrix is represented in terms of its Euclidean matrix space and the matrix-vector multiplication is linear in this Euclidean space. In some embodiments, the matrix is represented in terms of the phases of the VBSs and therefore the divisions may be performed on the bit strings representing the phases, instead of the matrix elements directly. In some embodiments, when the map between the phases to the matrix elements is a linear map, then the relationship between the input parameters—the phases of the VBSs and the input vector elements in this case—and the output vector is linear. When this relationship is linear, the method described above is still applicable.

However, in general, a nonlinear map from the elementary representation of the matrix to the photonic representation may be considered, according to some embodiments. For example, the bit-string division of the Euclidean space matrix elements from their most-significant k -bit string to the least-significant k -bit string may be used to produce a series of different matrices that are decomposed to a phase representation and implementing using a photonic processor.

The divisions need not be performed on both the matrix elements and the input vector elements simultaneously. In some embodiments, the photonic processor may propagate many input vectors for the same matrices. It may be efficient to only perform the divisions on the input vectors and keep the VBS controls at a set precision (e.g., full precision) because the digital-to-analog converters (DACs) for the vector preparations may operate at a high bandwidth while the DACs for the VBSs may be quasi-static for multiple vectors. In general, including a DAC with a high bit precision at higher bandwidth is more difficult than designing one at a lower bandwidth. Thus, in some embodiments, the output vector elements may be more precise than what is allowed by the ADC, but the ADC will automatically perform some rounding to the output vector value up to the bit precision allowed by the ADC.

L. Method of Manufacture

Embodiments of the photonic processing system may be manufactured using conventional semiconductor manufacturing techniques. For example, waveguides and phase shifters may be formed in a substrate using conventional deposition, masking, etching, and doping techniques.

FIG. 1-19 illustrates an example method 1-1900 of manufacturing a photonic processing system, according to some embodiments. At act 1-1901, the method 1-1900 includes forming an optical encoder using, e.g., conventional techniques. For example, a plurality of waveguides and modulators may be formed in a semiconductor substrate. The optical encoder may include one or more phase and/or amplitude modulators as described elsewhere in this application.

At act 1-1903, the method 1-1900 includes forming a photonic processor and optically connecting the photonic processor to the optical encoder. In some embodiments, the photonic

processor is formed in the same substrate as the optical encoder and the optical connections are made using waveguides formed in the substrate. In other embodiments, the photonic processor is formed in a separate substrate from the substrate of the optical encoder and the optical connection is made using optical fiber.

At act, 1-1905, the method 1-1900 include forming an optical receiver and optically connecting the optical receiver to the photonic processor. In some embodiments, the optical receiver is formed in the same substrate as the photonic processor and the optical connections are made using waveguides formed in the substrate. In other embodiments, the optical receiver is formed in a separate substrate from the substrate of the photonic processor and the optical connection is made using optical fiber.

FIG. 1-20 illustrates an example method 1-2000 of forming a photonic processor, as shown in act 1-1903 of FIG. 1-19. At act 1-2001, the method 1-2000 includes forming a first optical matrix implementation, e.g., in a semiconductor substrate. The first optical matrix implementation may include an array of interconnected VBSs, as described in the various embodiments above.

At act 1-2003, the method 1-2000 include forming a second optical matrix implementation and connecting the second optical matrix implementation to the first optical matrix implementation. The second optical matrix implementation may include one or more optical components that are capable of controlling the intensity and phase of each optical signal received from the first optical matrix implementation, as described in the various embodiments above. The connections between the first and second optical matrix implementation may include waveguides formed in the substrate.

At act 1-2005, the method 1-2000 includes forming a third optical matrix implementation and connecting the third optical matrix implementation to the second optical matrix implementation. The third optical matrix implementation may include an array of interconnected VBSs, as described in the various embodiments above. The connections between the second and third optical matrix implementation may include waveguides formed in the substrate.

In any of the above acts, the components of the photonic processor may be formed in a same layer of the semiconductor substrate or in different layers of the semiconductor substrate.

M. Method of Use

FIG. 1-21 illustrates a method 1-2100 of performing optical processing, according to some embodiments. At act 1-2101, the method 1-2100 includes encoding a bit string into optical signals. In some embodiments, this may be performed using a controller and optical encoder, as described in connection with various embodiments of this application. For example, a complex number may be encoded into the intensity and phase of an optical signal.

At act 1-2103, the method 1-2100 includes controlling a photonic processor to implement a first matrix. As described above, this may be accomplished by having a controller perform an SVD on the matrix and break the matrix into three separate matrix components that are implemented using separate portions of a photonic processor. The photonic processor may include a plurality of interconnected VBSs that control how the various modes of the photonic processor are mixed together to coherently interfere the optical signals when they are propagated through the photonic processor.

At act 1-2105, the method 1-2100 includes propagating the optical signals through the optical processor such that the optical signals coherently interfere with one another in a way that implements the desired matrix, as described above.

At act, 1-2107, the method 1-2100 includes detecting output optical signals from the photonic processor using an optical receiver. As discussed above, the detection may use phase-sensitive or phase-insensitive detectors. In some embodiments, the detection results are used to determine a new input bit string to be encoded and propagated through the system. In this way, multiple calculations may be performed in serial where at least one calculation is based on the results of a previous calculation result.

II. Training Algorithm

The inventors have recognized and appreciated that for many matrix-based differentiable program (e.g., neural network or latent-variable graphical model) techniques, the bulk of the computational complexity lies in matrix-matrix products that are computed as layers of the model are traversed. The complexity of a matrix-matrix product is $O(IJK)$,

where the two matrices have dimension I -by- J and J -by- K . Moreover, these matrix-matrix products are performed in both the training stage and the evaluation stage of the model.

A deep neural network (i.e., a neural network with more than one hidden layer) is an example of a type of matrix-based differentiable program that may employ some of the techniques described herein. However, it should be appreciated that the techniques described herein for performing parallel processing may be used with other types of matrix-based differentiable programs including, but not limited to, Bayesian networks, Trellis decoders, topic models, and Hidden Markov Models (HMMs).

The success of deep learning is in large part due to the development of backpropagation techniques that allow for training the weight matrices of the neural network. In conventional backpropagation techniques, an error from a loss function is propagated backwards through individual weight matrix components using the chain rule of calculus. Backpropagation techniques compute the gradients of the elements in the weight matrix, which are then used to determine an update to the weight matrix using an optimization algorithm, such as stochastic gradient descent (SGD), AdaGrad, RMSProp, Adam, or any other gradient-based optimization algorithm. Successive application of this procedure is used to determine the final weight matrix that minimizes the loss function.

The inventors have recognized and appreciated that an optical processor of the type described herein enables the performance of a gradient computation by recasting the weight matrix into an alternative parameter space, referred to herein as a “phase space” or “angular representation.” Specifically, in some embodiments, a weight matrix is reparameterized as a composition of unitary transfer matrices, such as Givens rotation matrices. In such a reparameterization, training the neural network includes adjusting the angular parameters of the unitary transfer matrices. In this reparameterization, the gradient of a single rotation angle is decoupled from the other rotations, allowing parallel computation of gradients. This parallelization results in a computational speedup relative to conventional serial gradient determination techniques in terms of the number of computation steps needed.

An example photonic processing system that may be used to implement the backpropagation techniques described herein is provided above. The phase space parameters of the reparameterized weight matrix may be encoded into phase shifters or variable beam splitters of the photonic processing system to implement the weight matrix. Encoding the weight matrix into the phase shifters or variable beam splitters may be used for both the training and evaluation stages of the neural network. While the backpropagation procedure is described in connection with the particular system described below, it should be understood

that embodiments are not limited to the particular details of the photonic processing system described in the present disclosure.

As described above, in some embodiments, photonics processing system 100 may be used to implement aspects of a neural network or other matrix-based differentiable program that may be trained using a backpropagation technique.

An example backpropagation technique 2-100 for updating a matrix of values in a Euclidean vector space (e.g., a weight matrix for a layer of a neural network) for a differentiable program (e.g., a neural network or latent variable graphical model) is shown in FIG. 2-1.

At act 2-101, a matrix of values in a Euclidean vector space (e.g., a weight matrix for a layer of a neural network) may be represented as an angular representation by, for example, configuring components of photonics processing system 100 to represent the matrix of values. After the matrix is represented in the angular representation, the process 2-100 proceeds to act 2-102, where training data (e.g., a set of input training vectors and associated labeled outputs) is processed to compute an error vector by assessing a performance measure of the model. Process 2-100 then proceeds to act 2-103, where at least some gradients of parameters of the angular representation needed for backpropagation are determined in parallel. For example, as discussed in more detail below, the techniques described herein enable gradients for an entire column of parameters to be determined simultaneously, significantly speeding up the amount of time needed to perform backpropagation as compared to evaluating the gradient with respect to each angular rotation individually. Process 2-100 then proceeds to act 2-104, where the matrix of values in the Euclidean vector space (e.g., the weight matrix values for a layer of a neural network) is updated by updating the angular representation using the determined gradients. A further description of each of the acts illustrated in process 2-100 of FIG. 2-1 is provided below.

FIG. 2-2 illustrates a flowchart of how act 2-101 shown in FIG. 2-1 may be performed in accordance with some embodiments. At act 2-201, a controller (e.g., controller 107) may receive a weight matrix for a layer of a neural network. At act 2-202, the weight matrix may be decomposed into a first unitary matrix, V , a second unitary matrix, U , and a diagonal matrix of signed singular values, Σ , such that the weight matrix W , is defined as:

$$W = V^T \Sigma U,$$

where U is an $m \times m$ unitary matrix, V is an $n \times n$ unitary matrix, Σ is an $n \times m$ diagonal matrix with signed singular values, and the superscript “ T ” indicates the transpose of

a matrix. In some embodiments, the weight matrix W is first partitioned into tiles, each of which is decomposed into the triple product of such matrices. The weight matrix, W , may be a conventional weight matrix as is known in the field of neural networks.

In some embodiments, the weight matrix decomposed into phase space is a pre-specified weight matrix, such as those provided by a random initialization procedure or by employing a partially trained weight matrix. If there is no partially-specified weight matrix to use for initialization of the backpropagation routine, then the decomposition in act 2-202 may be skipped and instead, the parameters of the angular representation (e.g., singular values and parameters of the unitary or orthogonal decomposition) may be initialized by, for example, randomly sampling the phase space parameters from a particular distribution. In other embodiments, a predetermined initial set of singular values and angular parameters may be used.

At act 2-203, the two unitary matrices U and V may be represented as compositions of a first and second set of unitary transfer matrices, respectively. For example, when matrices U and V are orthogonal matrices, they may be transformed in act 2-203 into a series of real-valued Givens rotation matrices or Householder reflectors, an example of which is described above in Section V.

At act 2-204, a photonics-based processor may be configured based on the decomposed weight matrix to implement the unitary transfer matrices. For example, as described above, a first set of components of the photonics-based processor may be configured based on the first set of unitary transfer matrices, a second set of components of the photonics-based processor may be configured based on the diagonal matrix of signed singular values, and a third set of components of the photonics-based processor may be configured based on the second set of unitary transfer matrices. Although the processes described herein are with respect to implementing the backpropagation technique using a photonics-based processor, it should be appreciated that the backpropagation technique may be implemented using other computing architectures that provide parallel processing capabilities, and embodiments are not limited in this respect.

Returning to the process 2-100 in FIG. 2-1, act 2-102 is directed to processing training data to compute an error vector after the weight matrix has been represented as an angular representation, for example, using a photonics-based processor as described above. FIG. 2-3 shows a flowchart of implementation details for performing act 2-102 in accordance with some embodiments. Prior to processing training data using the techniques described herein, the training data may be divided into batches. Training data may take any form. In some

embodiments, the data may be divided into batches in the same way as is done for some conventional mini-batch stochastic gradient descent (SGD) techniques. The gradients computed with this procedure can be used for any optimization algorithm, including but not limited to SGD, AdaGrad, Adam, RMSProp, or any other gradient-based optimization algorithm.

In the process shown in FIG. 2-3, each vector in a particular batch of training data may be passed through the photonics processor and a value of a loss function may be computed for that vector. At act 2-301, an input training vector from a particular batch of training data is received. At act 2-302, the input training vector is converted into photonic signals, for example, by encoding the vector using optical pulses that have amplitudes and phases corresponding to the input training vector values, as described above. At act 2-303, the photonic signals corresponding to the input training vector are provided as input to a photonics processor (e.g., photonics processor 103), which has been configured to implement (e.g., using an array of configurable phase shifters and beam splitters) a weight matrix, as described above, to produce an output vector of pulses. The optical intensity of the pulses output from the photonics processor may be detected using, for example, homodyne detection, as described above in connection with FIGS. 9 and 10 to produce a decoded output vector. At act 2-304, the value of a loss function (also known as a cost function or error metric) is computed for the input training vector. The process of acts 2-301 to 2-304 then repeats until all input training vectors in the particular batch have been processed and a corresponding value of the loss function has been determined. At act 2-305, the total loss is computed by aggregating these losses from each input training vector, e.g., this aggregation may take the form of an average.

Returning to the process 2-100 in FIG. 2-1, act 2-103 is directed to computing in parallel gradients for the parameters of the angular representation (e.g., the values of the weights that have been implemented using the components of the photonics processor as Givens rotations). The gradients may be calculated based on the computed error vectors, input data vectors (e.g., from a batch of training data), and the weight matrix implemented using the photonics processor. FIG. 2-4 show a flowchart of a process for performing act 2-103 in accordance with some embodiments. At act 2-401, for the k -th set of Givens rotations, $G^{(k)}$ in the decomposition (e.g., a column of the decomposed matrix – referred to herein as “derivative column k ”), a block diagonal derivative matrix containing the derivatives with respect to each angle in the k -th set is computed. At act 2-402, a product of the error vector

determined in act 2-102 with all of the unitary transfer matrices between the derivative column k and the output is computed (hereby referred to as a “partial backward pass”). At act 2-403, a product of the input data vector with all unitary transfer matrices starting from the input up to and including the derivative column k is computed (hereby referred to as a “partial forward pass”) At act 2-404, the inner products between successive pairs of elements output from acts 2-402 and 2-403 are computed to determine the gradients for the derivative column k . The inner products between successive pairs of elements may be computed as $\frac{\partial W^{(l)}}{\partial \theta_{ij}^{(k)}} = x_i \delta_i + x_j \delta_j$, where superscript (k) represents the k-th column of photonic elements and i, j represent the i -th and j -th photonic mode that are being coupled by the unitary transfer matrix with parameter $\theta_{ij}^{(k)}$, x is the output of the partial forward pass and δ is the output of the partial backward pass. In some embodiments, an offset is applied before the successive pairing of the outputs (e.g., output pairs could be (1, 2), (3, 4), etc. rather than (0, 1), (2, 3)). The determined gradients may then be used as appropriate for a particular chosen optimization algorithm (e.g., SGD) that is being used for training.

Example pseudocode for implementing the backpropagation technique on a photonic processor having the left-to-right topology shown in FIGS. 1-4 in accordance with some embodiments is provided below:

- Initialize two lists for intermediate propagation results, x' , δ' .
- For each column of MZI, starting with the last column and going to the first column:
 - Rotate the angles in the column to correspond to a derivative matrix
 - Propagate the input data vector through the photonics processor
 - Store the result in x'
 - Make the current column transparent
- Column-by-column, progressively build up the transpose matrix. For each new column:
 - Propagate the error vector through the photonics processor
 - Store the result in δ'
- For each $x'[i]$, $\delta'[i]$
 - Compute the inner products between successive pairs, with the result being the gradients for the angles in the i th column of MZI.

According to some embodiments, rather than adjusting the weights of a weight matrix via gradient descent as is done in some conventional backpropagation techniques, the parameters of angular representation (e.g., the singular values of the matrix Σ and the Givens rotation angles of the decomposed orthogonal matrices U and V) are adjusted. To further demonstrate how backpropagation in the reparameterized space works according to some embodiments, what follows is a comparison of backpropagation within a single layer of a neural network using conventional techniques versus the method according to some embodiments of the present disclosure.

A loss function, E , measures the performance of the model on a particular task. In some conventional stochastic gradient descent algorithms, the weight matrix is adjusted iteratively such that the weight matrix at time $t + 1$ is defined as a function of the weight matrix at time t and a derivative of the loss function with respect to the weights of the weight matrix is as follows:

$$w_{ab}(t + 1) = w_{ab}(t) - \eta \frac{\partial E}{\partial w_{ab}(t)},$$

where η is the learning rate and (a,b) represent the a-th row and b-th column entry of the weight matrix W , respectively. When this iterative process is recast using the decomposed weight matrix, the weights w_{ab} are functions of the singular values σ_i of the matrix Σ and the rotation angles θ_{ij} of the orthogonal matrices U and V . Thus the iterative adjustments of the backpropagation algorithm become:

$$\theta_{ij}^{(k)}(t + 1) = \theta_{ij}^{(k)}(t) - \eta \frac{\partial E}{\partial \theta_{ij}^{(k)}(t)}$$

and

$$\sigma_i(t + 1) = \sigma_i(t) - \eta \frac{\partial E}{\partial \sigma_i(t)}$$

To perform iterative adjustments to the singular values and rotation angles, the derivatives of the loss function must be obtained. Before describing how this can be achieved in a system such as the photonic processing system 100, a description is first provided for

backpropagation based on iteratively adjusting the weights of the weight matrix. In this situation, the output result measured by the system for a single layer of the neural network is expressed as an output vector $y_i = f((Wx)_i + b_i)$, where W is the weight matrix, x is the data vector input into the layer, b is a vector of biases, and f is a nonlinear function. The chain rule of calculus is applied to compute the gradient of the loss function with respect to any of the parameters within the weight matrix (where for convenience of representation, the definition $z_i = (Wx)_i + b_i$ is used):

$$\frac{\partial E}{\partial w_{ab}} = \sum_{ij} \frac{\partial E}{\partial y_i} \frac{\partial y_i}{\partial z_j} \frac{\partial z_j}{\partial w_{ab}}$$

Computing the derivative of z with respect to w_{ab} results in:

$$z_i = (Wx)_i + b_i = b_i + \sum_j W_{ij}x_j$$

$$\frac{\partial z_j}{\partial w_{ab}} = \delta_{ja}x_b$$

Using this fact, the sum representing the gradient of the loss function can then be written as:

$$\frac{\partial E}{\partial w_{ab}} = \sum_i \frac{\partial E}{\partial y_i} \frac{\partial y_i}{\partial z_a} x_b$$

The first sum is defined as the error vector e and x is the input vector, resulting in the following expression:

$$\frac{\partial E}{\partial w_{ab}} = e_a x_b$$

Using the above equations from conventional backpropagation techniques, the equations can be extended to the case of weight matrices decomposed into a singular value matrix and unitary transfer matrices. Using the fact that the weight matrix is a function of rotation angles, the chain rule can be used to write:

$$\begin{aligned}
\frac{\partial E}{\partial \theta_{ij}^{(k)}} &= \sum_{ab} \frac{\partial E}{\partial w_{ab}} \frac{\partial w_{ab}}{\partial \theta_{ij}^{(k)}} \\
&= \sum_{ab} e_a x_b \frac{\partial w_{ab}}{\partial \theta_{ij}^{(k)}} \\
&= \vec{e}^T \frac{\partial W}{\partial \theta_{ij}^{(k)}} \vec{x}
\end{aligned}$$

Thus, the backpropagation in phase space involves the same components as in conventional backpropagation (the error vector and the input data), with the addition of a term that is the derivative of the weight matrix with respect to the rotation angles of the unitary transfer matrices.

To determine the derivative of the weight matrix with respect to the rotation angles of the unitary transfer matrices, it is noted that the derivative of a single Givens rotation matrix has the following form:

$$\begin{aligned}
\frac{\partial G_{ij}^{(k)}}{\partial \theta_{ij}^{(k)}} &= \frac{\partial}{\partial \theta_{ij}^{(k)}} \begin{pmatrix} 1 & 0 & \dots & & \\ 0 & \ddots & & & \\ \vdots & & \cos \theta_{ij}^{(k)} & -\sin \theta_{ij}^{(k)} & \\ & & \sin \theta_{ij}^{(k)} & \cos \theta_{ij}^{(k)} & \vdots \\ & & & & \ddots & 0 \\ & & & & & 0 & 1 \end{pmatrix} \\
&= \begin{pmatrix} 0 & 0 & \dots & & \\ 0 & \ddots & & & \\ \vdots & & -\sin \theta_{ij}^{(k)} & -\cos \theta_{ij}^{(k)} & \\ & & \cos \theta_{ij}^{(k)} & -\sin \theta_{ij}^{(k)} & \vdots \\ & & & & \ddots & 0 \\ & & & & & 0 & 0 \end{pmatrix}
\end{aligned}$$

As can be seen, the derivative for any entry of the Givens rotation matrix that is not in the i -th row and j -th column is zero. Thus, all derivatives for angles inside $G^{(k)}$ may be grouped into a single matrix. To compute the derivative with respect to all $\text{floor}(\frac{n}{2})$ angles inside a column of unitary transfer matrices may, in some

embodiments, be accomplished using a two-step process, as described above. First, the error vector is propagated through the decomposed matrix from the right (output) up to the current set of rotations being differentiated (partial backward pass). Second, the input vector is propagated from the left (input) up to the current set of rotations (partial forward pass), and then the derivative matrix is applied.

In some embodiments, the derivation for the singular values is achieved using a similar process. The derivative with respect to a singular value σ_i results in the element Σ'_{ii} to 1 and all other Σ'_{jj} to 0. Therefore, all of the derivatives for the singular values may be calculated together. In some embodiments, this may be done by propagating the error vector from the left (partial forward pass) and propagating the input vector from the right (partial backward pass), then computing the Hadamard product from the outputs of the forward pass and the backward pass.

In the implementation of act 2-103 described in FIG. 2-4, all of the angles in a column k are rotated by $\pi/2$ in order to compute the gradient term for that column. In some embodiments, this rotation is not performed. Consider the matrix for a single MZI:

$$\begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$$

Taking the derivative with respect to θ yields

$$\begin{pmatrix} -\sin \theta & -\cos \theta \\ \cos \theta & -\sin \theta \end{pmatrix}$$

While this matrix corresponds to adding $\pi/2$ to θ , it also corresponds to swapping the columns of the original matrix and negating one of them. In mathematical notation, this means

$$\begin{pmatrix} -\sin \theta & -\cos \theta \\ \cos \theta & -\sin \theta \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$$

Rather than rotating the angle of each MZI in a column by $\pi/2$ and then computing the inner products between successive pairs of elements output from acts 2-402 and 2-403 as described above (e.g., $x_1\delta_1 + x_2\delta_2$), to determine the gradients for a column of the decomposed unitary matrix, in some embodiments, the angles are not rotated by $\pi/2$ and instead the relation $x_1\delta_2 - x_2\delta_1$ is calculated to obtain the same gradients. In some embodiments, where the size of the matrix W ($n \times m$) matches the size of the photonics processor with matrix U of size $n \times n$ and matrix V of size $m \times m$, acts 2-401 – 2-404 allow the controller to obtain $\text{floor}(\frac{n}{2})$ gradients for a unitary/orthogonal matrix of size $n \times n$.

Consequently, on hardware such the photonic processing system 100 described above, where each matrix multiplication can be computed in $O(1)$ operations, the overall backpropagation procedure may be completed in $O(n + m)$ operations when the photonic processor is of sufficient size to represent the full matrix. When the photonics processor is not of sufficient size to represent the full matrix, the matrix may be partitioned into tiles, as described above. Consider a photonic processor of size N . If the task is to multiply a matrix of size $I \times J$ by a vector of size J , a single matrix-vector product will have complexity $O(IJ/N^2)$ (assuming that both I and J are divisible by N), because each dimension of the matrix must be partitioned into matrices of size N , loaded into the processor, and used to compute a partial result. For a batch of K vectors (e.g., a second matrix of size $J \times K$), the complexity is $O(IJK/N^2)$ for the matrix-vector product.

An embodiment of a photonic processor, as described above, with n optical modes naturally computes a matrix-vector product between a matrix of size $[n \times n]$ and an n -element vector. This is equivalently expressed as a matrix-matrix product between matrices of sizes $[n \times n]$ and $[n \times 1]$. Furthermore, a sequence of K matrix-vector product operations with K different input vectors and a single, repeated input matrix can be expressed as the computation of a matrix-matrix product between matrices of size $[n \times n]$ and $[n \times K]$. But the applications and algorithms described herein often involve the computation of general matrix-matrix multiplication (GEMM) between matrices of arbitrary size; i.e., the computation

$$c_{ik} = \sum_{j=1}^J a_{ij} \cdot b_{jk} \quad \forall i \in [1, I], k \in [1, K]$$

Where a_{ij} is the element in the i^{th} row and j^{th} column of an $[I \times J]$ matrix A , b_{jk} is the j^{th} row and k^{th} column of a $[J \times K]$ matrix B and c_{ik} is the element in the i^{th} row and k^{th} column of the $[I \times K]$ matrix $C = AB$. Due to the recursive nature of this computation, this can be equivalently expressed as:

$$C_i = \sum_{j=1}^{\lfloor J/n \rfloor} A_{ij} B_j \quad \forall i \in [1, I], k \in [1, K]$$

Where A has been first padded to a $\left[n \cdot \left\lceil \frac{I}{n} \right\rceil \times \left(n \cdot \left\lceil \frac{J}{n} \right\rceil \right) \right]$ matrix and then partitioned into a $\left[\left\lceil \frac{I}{n} \right\rceil \times \left\lceil \frac{J}{n} \right\rceil \right]$ grid of submatrices (each of size $[n \times n]$) and A_{ij} is the $[n \times n]$ submatrix in the i^{th} row and j^{th} column of this grid, B has been first padded to a $[(n \cdot \left\lceil \frac{J}{n} \right\rceil) \times K]$ matrix and then partitioned into a $\left[\left\lceil \frac{J}{n} \right\rceil \times 1 \right]$ grid of submatrices (each of

size $[n \times K]$) and B_j is the $[n \times K]$ submatrix in the j^{th} row of this grid, and C has been first padded to a $[(n \cdot \lceil J/n \rceil) \times K]$ matrix and then partitioned into a $[\lceil I/n \rceil \times 1]$ grid of submatrices (each of size $[n \times K]$) and C_i is the $[n \times K]$ submatrix in the i^{th} row of this grid.

Using this process, a photonic processor can compute any GEMM by loading $(\lceil I/n \rceil \cdot \lceil J/n \rceil)$ different matrices into the photonic array and, for each loaded matrix, propagating k different vectors through the photonic array. This yields $\lceil I/n \rceil \cdot \lceil J/n \rceil \cdot k$ output vectors (each comprised of n elements), a subset of which may be added together to yield the desired $[I \times K]$ output matrix, as defined by the equation above.

In the implementation of act 2-103 described in FIG. 2-4, a left-to-right topology of the photonics processor for implementing the matrices was assumed, where the vector was input on the left of the array of optical components and the output vector was provided on the right of the array of optical components. This topology requires the transpose of the angular representation matrix to be calculated when propagating the error vector through the photonics processor. In some embodiments, the photonic processor is implemented using a folded-over topology that arranges both the inputs and outputs on one side of the array (e.g., the left side) of optical components. Such an architecture allows the use of a switch to decide which direction the light should propagate - either from input to output or output to input. With the choice of direction configurable on the fly, the propagation of the error vector can be accomplished by first switching the direction to use output as input and then propagating the error vector through the array, which eliminates the need to negate phases (e.g., by rotating the angle of each photonic element in a column k by $\pi/2$) and transpose the columns when the gradients for column k are being determined, as described above.

Returning to the process 2-100 in FIG. 2-1, act 2-104 is directed to updating the weight matrix by updating the parameters of the angular representation based on the determined gradients. The process shown in FIG. 2-4, just described, is for computing the gradients of the angular parameters for a single set (e.g., column k) of gradients in the decomposed unitary matrix based on a single input data exemplar. In order to update the weight matrix, the gradients for each of the sets (e.g., columns) need to be computed for each input data vector in a batch. FIG. 2-5 shows a flowchart of a process for computing all of the gradients needed to update the weight matrix in accordance with some embodiments. At act 2-501, the gradients for one set of multiple sets of unitary transfer matrices (e.g., Givens rotations) is determined (e.g., using the process shown in FIG. 2-4). At act 2-502, it is then

determined whether there are additional sets of unitary transfer matrices (e.g., columns) for which gradients need to be calculated for the current input data vector. If it is determined that there are additional gradients to be calculated, the process returns to act 2-501, where a new set (e.g., column) is selected and gradients are calculated for the newly selected set. As noted below, in some computing architectures, all columns of the array may be read out simultaneously, such that the determination at act 2-502 is not needed. The process continues until it is determined at act 2-502 that gradients have been determined for all of the sets of unitary transfer matrices (e.g., columns) in the decomposed unitary matrix. The process then proceeds to act 2-503, where it is determined whether there are more input data vectors to process in the batch of training data being processed. If it is determined that there are additional input data vectors, the process returns to act 2-501, where a new input data vector from the batch is selected and gradients are calculated based on the newly selected input data vector. The process repeats until it is determined at act 2-503 that all input data vectors in the batch of training data have been processed. At act 2-504, the gradients determined in acts 2-501-2-503 are averaged and the parameters of the angular representation (e.g., the angles of the Givens rotation matrices of the decomposition of the weight matrix) are updated based on the an update rule that makes use of the averaged gradients. As a non-limiting example, an update rule may include scaling the averaged gradients by a learning rate or include “momentum” or other corrections for the history of parameter gradients.

As discussed briefly above, although the above example was applied to a real weight matrix in a single layer neural network, the results may be generalized to networks with multiple layers and complex weight matrices. In some embodiments, the neural network consists of multiple layers (e.g., ≥ 50 layers in a deep neural network). To compute the gradient for a matrix of layer L , the input vector to that layer would be the output of preceding layer $L-1$ and the error vector to that layer would be the error backpropagated from the following layer $L+1$. The value of the backpropagated error vector can be computed using the chain rule of multivariable calculus as before. Moreover, in some embodiments, complex U and V matrices (e.g., unitary matrices) may be used by adding an additional complex phase term to the Givens rotation matrix.

While the description above applies generally independent of the hardware architecture, certain hardware architectures provide more significant computation acceleration than others. In particular, implementation of the backpropagation

techniques described herein on a graphical processing unit, a systolic matrix multiplier, a photonic processor (e.g., photonic processing system 100), or other hardware architectures capable of parallel computations of the gradients are preferred for the greatest gains compared to conventional approaches.

As described above, the photonic processing system 100 is configured to implement any unitary transformation. A sequence of Givens rotations is an example of such a unitary transformation, and thus the photonic processing system 100 can be programmed to compute the transformations in the decomposition above in $O(1)$ time. As described above, the matrix may be implemented by controlling a regular array of variable beam-splitters (VBSs). The unitary matrices U and V^T may be decomposed into a tiled array of VBS, each of which performs a 2-by-2 orthogonal matrix operation (e.g., a Givens rotation). The diagonal matrix Σ , along with the diagonal phase screen D_U and D_V (in the form of the diagonal matrix $D_U \Sigma D_V$), can be implemented in the photonic processing system 100 by controlling the intensity and phase of the light pulses, as described above.

Each entry in the diagonal matrix Σ corresponds to the amplification or attenuation of each photonic mode. An entry with magnitude ≥ 1 corresponds to amplification and an entry with magnitude ≤ 1 corresponds to attenuation, and a combination VBS and gain medium would allow for either attenuation or amplification. For an n -by- n square matrix M , the number of optical modes needed to apply the diagonal matrix Σ is n . However, if the matrix M is not square, the number of optical modes needed is equal to the smaller dimension.

As noted above, in some embodiments, the size of the photonic processor is the same as the size of the matrix M and input vector being multiplied. However, in practice, the size of the matrix M and the size of the photonic processor often differs. Consider a photonic processor of size N . If the task is to multiply a matrix of size $I \times J$ by a vector of size J , a single matrix-vector product will have complexity $O(IJ/N^2)$ (assuming that both I and J are divisible by N), because each dimension of the matrix must be partitioned into matrices of size N , loaded into the processor, and used to compute a partial result. For a batch of K vectors (e.g., a second matrix of size $J \times K$), the complexity is $O(IJK/N^2)$ for the matrix-vector product.

The ability to work on small N -by- N matrix partitions can be advantageous if the matrix is non-square, especially if either $I \gg J$ or $J \gg I$. Assuming a non-square

matrix A , direct SVD of the matrix produces one $I \times I$ unitary matrix, one $J \times J$ unitary matrix, and one $I \times J$ diagonal matrix. If either $I \gg J$ or $J \gg I$, the number of parameters needed to represent this decomposed matrices are much larger than the original matrix A .

However, if the matrix A is partitioned into multiple $N \times N$ square matrices having smaller dimensions, SVD on these $N \times N$ matrices produces two $N \times N$ unitary matrices and one $N \times N$ diagonal matrix. In this case, the number of parameters needed to represent the decomposed matrices is still N^2 —equal to the size of the original matrix A , and the total non-square matrix can be decomposed with $\approx IJ$ total parameters. The approximation becomes equality when IJ is divisible by N^2 .

For a photonic processor having $2N+1$ columns, the partial results of backpropagating the error vector for each column may be computed. Therefore, for a batch of K vectors, the complexity of backpropagation using a photonic processor of size N is $O(IJK/N)$. By comparison, the computation of backpropagated errors using a matrix multiplication algorithm on a non-parallel processor (e.g., a CPU) would be $O(IJK)$.

The description so far has focused on the use of a matrix within a neural network layer with an input vector data and a backpropagated error vector. The inventors have recognized and appreciated that the data in deep neural network computations are not necessarily vectors, but they are in general multidimensional tensors. Similarly, the weight values that describe the connection between the neurons are in general also multidimensional tensors. In some embodiments, the method described above can be directly applied if the weight tensor is sliced into matrix slices with each matrix slice being independent of one another. Therefore, singular value decomposition and the Givens-like rotation decomposition can be performed to obtain a valid representation in terms of phases for a particular matrix slice. The same method of computing the gradient of the phases can then be applied with the proper arrangement of the input tensor data and the backpropagated error data as well. The gradients for a specific matrix slice should be computed with parts of the input and error data that would have contributed to that particular matrix slice.

For concreteness, consider a general n -dimensional weight tensor $w_{a_1 a_2 a_3 \dots a_i \dots a_n}$. Choose two indices out of $\{a_i\}_{i=1, \dots, n}$ that would constitute the matrix slice—say that choice is labeled by indices a_b and a_c —and perform the decomposition to obtain the phases $\theta_{ij}^{(k)}$ by computing $\partial E / \partial \theta_{ij}^{(k)}$. Importantly, b and c can be any value between 1 and n . Consider now a general k -dimensional input tensor $x_{a_1 \dots a_i \dots a_k}$. For a valid tensor operation, it must be the case that the output of the operation of the weight tensor on this input tensor produces an $(n - k)$ -dimensional output tensor. It can therefore be concluded that the backpropagated error tensor the weight tensor of this layer is an $(n - k)$ -dimensional tensor: $e_{a_{k+1} \dots a_i \dots a_n}$. Therefore, the gradient to be computed is

$$\partial E / \partial \theta_{ij}^{(k)} = \sum_{a_1 \dots a_n} e_{a_1 \dots a_b \dots a_n} w_{a_1 \dots a_b \dots a_c \dots a_n} x_{a_1 \dots a_c \dots a_k},$$

where, for simplicity (but not a necessary condition), the indices of the weight tensors have been ordered such that the first k indices operate on x , and the last $(n - k)$ indices operate on the error e .

In other embodiments, it may be more convenient to perform higher-order generalization of singular value decomposition such as the Tucker decomposition, where an arbitrary n -dimensional tensor can be decomposed as such: $w_{a_1 \dots a_n} = \sum_{b_1 \dots b_n} g_{b_1 \dots b_n} U_{a_1 b_1}^{(1)} U_{a_2 b_2}^{(2)} \dots U_{a_n b_n}^{(n)}$, where each $U_{a_i b_i}^{(i)}$ is an orthogonal matrix that can be decomposed into its Givens rotation phases and $g_{b_1 \dots b_n}$ is an n -dimensional core tensor. In some cases, the core tensor can be chosen to be superdiagonal using a special case of the Tucker decomposition called CANDECOMP/PARAFAC (CP) decomposition. The Tucker decomposition can be made similar to the 2-dimensional SVD form by multiplying the inverses (transposes or conjugate transposes) of some of the unitary matrices. For example, the decomposition can be rewritten as $w_{a_1 \dots a_n} = \sum_{b_1 \dots b_n} (U^T)_{b_1 a_1}^{(1)} \dots (U^T)_{b_m a_m}^{(m)} g_{b_1 \dots b_n} U_{a_{m+1} b_{m+1}}^{(m+1)} \dots U_{a_n b_n}^{(n)}$, where the first m unitary matrices are pushed to the left of the core tensor. The collection of unitary matrices on either side of the core tensor can be decomposed into their rotation

angles and the gradient of each rotation angle is obtained by the chain rule of calculus and the contraction of the gradients with the input tensor and the error tensor.

The inventors have recognized and appreciated that the gradients of the phases (e.g., for decomposed matrices U and V) and the gradients of the signed singular values (e.g., for matrix Σ) may have different upper bounds. Consider a task to compute the gradients of the scalar loss function L with respect to neural network parameters. In Euclidean space, the value of the gradients is given by $\frac{\partial L}{\partial W}$, where W is a matrix. In phase space, for a particular scalar phase θ_k , the chain rule provides:

$$\frac{\partial L}{\partial W} = \sum_{ij} \frac{\partial L}{\partial W_{ij}} \frac{\partial W_{ij}}{\partial \theta_k}$$

From the definition of the trace, this is equal to:

$$\frac{\partial L}{\partial \theta_k} = \text{Tr} \left(\frac{\partial L}{\partial W} \left(\frac{\partial W}{\partial \theta_k} \right)^T \right)$$

where $\frac{\partial L}{\partial \theta_k}$ and $\frac{\partial W}{\partial \theta}$ are both matrices. It is known that the trace is bounded by the

Frobenius norm product $\text{Tr}(AB) \leq \|A\|_F \|AB\|_F$ and that $\|A\|_F = \|A^T\|_F$. Therefore,

$$\frac{\partial L}{\partial \theta_k} \leq \left\| \frac{\partial W}{\partial \theta_k} \right\|_F \left\| \frac{\partial L}{\partial W} \right\|_F$$

Because differentiating with respect to θ does not change the singular values of W and thus does not change the Frobenius norm, the following is true:

$$\frac{\partial L}{\partial \theta_k} \leq \|W\|_F \left\| \frac{\partial L}{\partial W} \right\|_F$$

Differentiating with respect to a particular singular value σ_k , all of the singular values go to zero except for the one being differentiated, which goes to 1, which means that

$$\left\| \frac{\partial W}{\partial \theta_k} \right\|_F = 1$$

Therefore,

$$\frac{\partial L}{\partial \theta_k} \leq \left\| \frac{\partial L}{\partial W} \right\|_F$$

In some embodiments, the gradients of the phases and the singular values are scaled separately during updating the parameters of the angular representation to, for example, account for the differences in upper bounds. By scaling the gradients separately, either the gradients of the phases or the gradients of the singular values may be rescaled to have the same upper bound. In some embodiments, the gradients of

the phases are scaled by the Frobenius norm of the matrix. According to some update rules, scaling the gradients of the phases and the singular values independently equates to having different learning rates for the gradients of the phases and the gradients of the singular values. Accordingly, in some embodiments, a first learning rate for updating the sets of components for the U and V matrices is different than a second learning rate for updating the set of components for the Σ matrix.

The inventors have recognized and appreciated that once a weight matrix is decomposed into phase space, both the phases and the singular values may not need to be updated in every iteration to obtain a good solution. Accordingly if only the singular values (but not the phases) are updated some fraction of the overall training time, during those epochs only $O(n)$ parameters would need to be updated rather than $O(n^2)$, leading to improvements in overall runtime. Updating only the singular values or the phases during some iterations may be referred to as “parameter clamping.” In some embodiments, parameter clamping may be performed according to one or more of the following clamping techniques:

- **Fixed clamping** : train all parameters for a certain number of iterations, then only update the singular values subsequently
- **Cyclic clamping** : Train all parameters for a number of epochs M , then freeze the phases (i.e., only update the singular values) for a number of epochs N . Resume training all parameters for another M epochs, then freeze the phases for N epochs again. Repeat until the total number of desired epochs has been reached.
- **Warmup clamping** : Train all parameters for some number of epochs K , then begin cyclic clamping for the remaining number of epochs.
- **Threshold clamping**: Continue updating phases or singular values until their updates are smaller than a threshold value ϵ

The inventors have recognized and appreciated that the architecture of the photonics processor may influence the complexity of the calculations. For example, in the architecture shown in FIG. 1, the detectors are only at one end of the photonic processor array, resulting in the column-by-column approach (e.g., using partial forward and backward passes) for calculating the gradients, as described above. In an

alternate architecture, a detector may be arranged at every photonic element (e.g., at each MZI) in the photonics array. For such an architecture, the column-by-column approach may be replaced with a single forward pass and a single backward pass, where the output at every column is read out simultaneously, therefore providing additional computational acceleration. Intermediate solutions, where columns of detectors are placed intermittently throughout the array are also contemplated. Any addition of detector columns commensurately reduces the number of partial forward and backward passes required for the gradient computation.

The techniques described above illustrate techniques for performing an update of weight matrix parameters while keeping all of the computation in phase space (e.g., using the angular representation of the matrix). In some embodiments, at least some of the calculations may be performed in a Euclidean vector space, whereas other calculations are performed in phase space. For example, the quantities needed to perform the update may be computed in phase space, as described above, but the actual updating of the parameters may occur in a Euclidean vector space. The updated matrix calculated in the Euclidean vector space may then be re-decomposed into weight space for a next iteration. In Euclidean vector space, for a given layer, the update rule may be:

$$\frac{\partial L}{\partial W_{ij}} = x_i \delta_j$$

The δ in this computation can be calculated with a backward pass through the entire photonics processor in phase space. Then, the outer product above between x and δ can be computed separately (e.g., off-chip). Once the updates are applied, the updated matrix can be re-decomposed and the decomposed values can be used to set the phases for the photonic processor as described above.

III. Convolutional Layers

Convolution and cross-correlation are common signal processing operations with many applications such as audio/video encoding, probability theory, image processing, and machine learning. The terms convolution and cross-correlation generally refer to mathematical operations that accept, as input, two signals and produce, as output, a third signal which represents the similarity that exists between the inputs. The inventors have

recognized and appreciated that computing convolutions and cross-correlations may be computationally resource-intensive. In particular, the inventors have developed techniques for improving the computational speed and efficiency of convolutions and cross-correlations. Embodiments of these techniques include computing convolutions and cross-correlations by transforming convolution operations into a matrix-vector product and/or a product of multi-dimensional arrays. Embodiments of these techniques further include computing convolutions according to a discrete transform.

The inventors have further recognized and appreciated that computing convolutions and cross-correlations may be performed in a variety of ways depending on the intended application. Input and output signals may be discrete or continuous. The data values that the signals are composed of may be defined over a variety of numerical domains such as the real numbers, the complex plane, or a finite integer ring. The signals may have any number of dimensions. The signals may also have multiple channels, which is a technique commonly used in convolutional neural networks (CNNs). The embodiments described herein may be implemented to accommodate these variations in any combination.

Furthermore, embodiments of these techniques may be implemented in any suitable computational system configured to perform matrix operations. Examples of such computational systems which may benefit from the techniques described herein include central processing units (CPUs), graphic processing units (GPUs), field programmable gate arrays (FPGAs), application-specific integrated circuits (ASICs), and photonic processors. While embodiments described herein may be described in connection to photonic processors, it is to be appreciated that these techniques may be applicable to other computational systems such as, but not limited to, those described above.

Following below are more detailed descriptions of various concepts related to, and embodiments of, techniques for computing convolutions and cross-correlations. It should be appreciated that various aspects described herein may be implemented in any of numerous ways. Examples of specific implementations are provided herein for illustrative purposes only. In addition, the various aspects described in the embodiments below may be used alone or in any combination, and are not limited to the combinations explicitly described herein.

The inventors further recognized and appreciated that using light signals, instead of electrical signals, overcomes many of the aforementioned problems with electrical computing. Light signals travel at the speed of light in the medium in which the light is

traveling; thus the latency of photonic signals is far less of a limitation than electrical propagation delay. Additionally, no power is dissipated by increasing the distance traveled by the light signals, opening up new topologies and processor layouts that would not be feasible using electrical signals. Thus, light-based processors, such as a photonic processor, may have better speed and efficiency performance than some conventional electrical-based processors.

To implement a photonics-based processor, the inventors have recognized and appreciated that the multiplication of an input vector by a matrix can be accomplished by propagating coherent light signals, e.g., laser pulses, through a first array of interconnected variable beam splitters (VBSs), a second array of interconnected variable beam splitters, and multiple controllable electro-optic elements between the two arrays that connect a single output of the first array to a single input of the second array.

Some embodiments of a photonic processor, as described herein, with n optical modes naturally compute a matrix-vector product between a matrix of size $[n \times n]$ and an n -element vector. This may be equivalently expressed as a matrix-matrix product between matrices of sizes $[n \times n]$ and $[n \times 1]$. A sequence of K matrix-vector product operations with K different input vectors and a single, repeated input matrix may be expressed as the computation of a matrix-matrix product between matrices of size $[n \times n]$ and $[n \times K]$. But the applications and algorithms described herein often involve the computation of general matrix-matrix multiplication (GEMM) between matrices of arbitrary size; i.e., the computation:

$$c_{ik} = \sum_{j=1}^J a_{ij} \cdot b_{jk} \quad \forall i \in [1, I], k \in [1, K],$$

where a_{ij} is the element in the i^{th} row and j^{th} column of an $[I \times J]$ matrix A , b_{jk} is the j^{th} row and k^{th} column of a $[J \times K]$ matrix B and c_{ik} is the element in the i^{th} row and k^{th} column of the $[I \times K]$ matrix $C = AB$. Due to the recursive nature of this computation, this can be equivalently expressed as:

$$C_i = \sum_{j=1}^{\lfloor \frac{J}{n} \rfloor} A_{ij} B_j \quad \forall i \in [1, \lfloor I/n \rfloor],$$

where A has been first padded to a $[(n \cdot \lfloor I/n \rfloor) \times (n \cdot \lfloor J/n \rfloor)]$ matrix and then partitioned into a $[\lfloor I/n \rfloor \times \lfloor J/n \rfloor]$ grid of submatrices (each of size $[n \times n]$) and A_{ij} is the $[n \times n]$ submatrix in the i^{th} row and j^{th} column of this grid, B has been first padded to a $[(n \cdot \lfloor J/n \rfloor) \times K]$ matrix and then partitioned into a $[\lfloor J/n \rfloor \times 1]$ grid of submatrices (each of size $[n \times K]$) and B_j is the $[n \times K]$ submatrix in the j^{th} row of this grid, and C has been first

padded to a $[(n \cdot \lceil J/n \rceil) \times K]$ matrix and then partitioned into a $\lceil I/n \rceil \times 1$ grid of submatrices (each of size $[n \times K]$) and C_i is the $[n \times K]$ submatrix in the i^{th} row of this grid.

According to some embodiments, using this process, a photonic processor can compute any GEMM by loading $(\lceil I/n \rceil \cdot \lceil J/n \rceil)$ different matrices into the photonic array and, for each loaded matrix, propagating K different vectors through the photonic array. This yields $\lceil I/n \rceil \cdot \lceil J/n \rceil \cdot K$ output vectors (each comprised of n elements), a subset of which may be added together to yield the desired $[I \times K]$ output matrix, as defined by the equation above.

The inventors have recognized and appreciated that a photonic processor may accelerate the process of computing convolutions and cross-correlations, but that embodiments for computing convolutions and cross-correlations described herein may be implemented on any suitable computational system. Embodiments described herein are discussed in terms of 2-dimensional convolutions, but may be generalizable to any number of dimensions. For an $[I_h \times I_w]$ input (herein called the “image,” though it is to be understood that the input could represent any suitable data), G , and a $[K_h \times K_w]$ filter, F , the mathematical formula for a two-dimensional convolution is:

$$(G * F)[x, y] = \sum_{i=0}^{K_h} \sum_{j=0}^{K_w} F[i, j] \cdot \hat{G}[x - i, y - j]$$

The two-dimensional cross-correlation is given by:

$$(G \star F)[x, y] = \sum_{i=0}^{K_h} \sum_{j=0}^{K_w} \overline{F[i, j]} \cdot \hat{G}[x + i, y + j]$$

where $\hat{G}$ is a function of G determined by the boundary conditions, $\bar{F}$ denotes complex-conjugation, and $\cdot$ denotes scalar multiplication.

In some implementations, convolution and cross-correlation operations may be interchangeable, as the cross-correlation of complex-valued, two-dimensional signals G and F can be converted to a convolution via

$$(G * F)[x, y] = \overline{(G[x, y] \star F[-x, -y])}[x, y].$$

The embodiments described herein will focus on the convolution case, but it is to be understood that embodiments described herein may be used to compute convolutions and cross-correlations.

In both convolution and cross-correlation, different variants exist depending on how the boundary conditions are handled. Two boundary conditions described in some embodiments herein include circular:

$$\hat{G}[x, y] = G[x \% I_h, y \% I_w]$$

and padded:

$$\hat{G}[x, y] = G[x, y] \text{ if } (0 \leq x \leq I_h \text{ and } 0 \leq y \leq I_w); 0 \text{ otherwise}$$

where $a \% n$ indicates $a \bmod n$.

Additional boundary condition variants may be used, according to some embodiments. These boundary conditions include symmetric (also known as mirror or reflective) boundary conditions in which the image is reflected across the boundary. The padded boundary condition may variously be called linear or fill in some embodiments. The circular boundary condition is also known as wrapped.

Additionally, different output modes may be employed to determine which elements interact with the boundary condition. These output modes include valid, same (or half-padded), and full output modes. Valid output mode requires that the output consists of only the elements that do not depend on the boundary condition. Same output mode requires the output to be the same size as the input. Full output mode requires that the output consists of all elements that do not exclusively depend on the boundary condition.

Different output modes control the number of points $[x, y]$ on which the output is defined. Each embodiment described herein may therefore be modified to operate in any given output mode. While the embodiments described herein focus on the same-mode convolution case, it is to be understood that these implementations may be extended to compute cross-correlation and/or alternate output modes instead or in addition to the embodiments described herein.

In some implementations, such as in CNNs, these operations may be generalized such that they can be applied to and/or produce multi-channel data. As an example, an RGB image has three color channels. For an input, G , with two spatial dimensions and C channels, the multi-channel operation is defined as:

$$(G \circledast F)[m, x, y] = \sum_{c=0}^C F[m, c] \circledast G[c] ; \forall m \in M$$

where $\circledast$ represents either convolution or cross-correlation, M is the number of output channels, G is a three dimensional $[C \times I_h \times I_w]$ tensor, F is a four-dimensional $[M \times C \times I_h \times I_w]$ tensor, and $(G \circledast F)$ is a three-dimensional $[M \times I_h \times I_w]$ tensor. For the above, slice indexing notation is used, with spatial dimensions suppressed, such that $F[m, c]$ accesses a two-dimensional $[K_h \times K_w]$ spatial slice of F and $G[c]$ accesses a two-dimensional $[I_h \times I_w]$ spatial slice of G .

In general, techniques for expressing convolutions as matrix operations may follow the process of FIG. 3-1. In act 3-102, pre-processing of the image and/or filter matrices may occur prior to the matrix operation in order to make sure the matrices are, for example, of the correct dimensionality, obey boundary conditions, and/or output modes. In act 3-104, the core matrix or matrix-vector operation may be applied, creating an output of the convolution. In act 3-106, post-processing of the output of the convolution may occur in order to, for example, reshape the output, as will be discussed in more detail herein.

Some embodiments may use a photonic processor to compute convolutions as a matrix-vector product. The inventors have recognized and appreciated that an array of variable beam splitters (VBSs), such as those that may be included in some embodiments of a photonic processor as described previously herein, may be used to represent any unitary matrix. As an example, using those techniques to represent an expanded image G_{mat} , the matrix may be decomposed with singular value decomposition as

$$G_{mat} = V^T \Sigma U .$$

In some embodiments of a photonic processor, the two unitary matrices U and V may then be decomposed with the algorithm described previously. The phases that result from that computation are programmed into the photonic array, along with the singular values. In some embodiments, the processor decomposes the filter rather than the image so that the filter can stay loaded for an entire batch of images.

An example of a process for computing a convolution in a photonic processor is shown in FIG. 3-2, according to some embodiments. In act 3-202, the process constructs the matrix G_{mat} from the input image matrix G . The matrix G_{mat} may be constructed in any

suitable way in accordance with the chosen boundary conditions, including but not limited to constructing G_{mat} as a doubly-block circulant matrix or a doubly-block Toeplitz matrix.

In act 3-204, the decomposed matrix G_{mat} may then be loaded into the photonic array. For each filter F in the input batch, a loop is repeated, wherein the filter F is flattened into a column vector in act 3-206, passed through the photonic array in act 3-208 to perform the matrix-multiplication operation, and then reshaped into an output with an appropriate dimensionality in act 3-210. In act 3-212, it is determined whether any further filters F remain. If further filters F are to be passed through the convolutional layer, the process returns to act 3-206. Otherwise, the process ends. Because of the commutative nature of convolutions, process 3-200 may be performed with the filter F expanded into F_{mat} and the images G being flattened into column vectors and passed through the photonic array in act 3-208.

A photonic processor may be used to implement any suitable matrix-multiplication-based algorithm. Matrix-multiplication-based algorithms re-order and/or expand the input signals such that the computation can be expressed as a general matrix-matrix multiply (GEMM) with some pre- and/or post-processing. Some example matrix-multiplication-based algorithms which may be implemented on a photonic processor include im2col, kn2row, and memory-efficient convolution (MEC).

According to some embodiments, the im2col algorithm may be implemented on a photonic processor. In the im2col algorithm, during pre-processing, the image G may be expanded from an $[I_h \times I_w]$ matrix to a $[(K_h \cdot K_w) \times (I_h \cdot I_w)]$ matrix. The filter F may be flattened from a $[K_h \times K_w]$ matrix to a $[1 \times (K_h \cdot K_w)]$ row vector. The output may then be generated by a matrix-vector product of the image and the filter because this pre-processing step generates an expanded data matrix in which each column contains a copy of all $(K_h \times K_w)$ elements that may be scaled and accumulated for each location in the output. The im2col algorithm may therefore require $O(K_h \cdot K_w \cdot I_h \cdot I_w)$ data copies and $O(K_h \cdot K_w \cdot I_h \cdot I_w)$ temporary storage.

According to some embodiments, the kn2row algorithm may be implemented on a photonic processor. The kn2row algorithm computes an outer product of the unmodified image and filter signals, generating a temporary matrix of size $[(K_h \cdot K_w) \times (I_h \cdot I_w)]$. The kn2row algorithm then adds particular elements from each row of the outer product together

to produce a $[1 \times (I_h \cdot I_w)]$ output vector. The kn2row algorithm may therefore also require $O(K_h \cdot K_w \cdot I_h \cdot I_w)$ data copies and $O(K_h \cdot K_w \cdot I_h \cdot I_w)$ temporary storage.

According to some embodiments, the MEC algorithm may be implemented on a photonic processor. The MEC algorithm may expand the input image by a factor of only K_h or K_w , rather than a factor of $(K_h \cdot K_w)$ as in the im2col algorithm. If the smaller filter dimension is chosen for expansion, then the algorithm requires only $O(\min(K_h, K_w) \cdot I_h \cdot I_w)$ temporary storage and data copies. Unlike im2col or kn2row, which compute a single matrix-vector product, the MEC algorithm computes a series of smaller matrix-vector products and concatenates the results.

In the embodiments discussed above, the filter matrix may be expanded during pre-processing rather than the image because of the commutative nature of convolutions. The choice of whether the image or the filter is to be tiled and reshaped into a matrix may be determined by which operations are faster and/or require less computational energy.

VII. Multi-Dimensional Convolution via Two-Dimensional Matrix-Matrix Multiplication

The inventors have recognized and appreciated that the matrix-multiplication-based algorithms for computing convolutions discussed above may not be suitable for some computing architectures or applications. The inventors have further recognized and appreciated that an approach that could combine the computational efficiency of im2col or kn2row with the memory-efficient features of the MEC algorithm would be beneficial for the computation of convolutions and cross-correlations. In particular, the inventors have recognized that these benefits may be achieved by splitting the re-ordering and reshaping of input and output matrices between pre- and post-processing steps, and that such a method may be generalized to N -dimensional convolutions, where $N \geq 2$.

According to some embodiments, the Multi-Dimensional Convolution via Two-Dimensional Matrix-Matrix Multiplication algorithm (herein the “cng2” algorithm), includes three steps. At a high level, for a non-limiting example of a two-dimensional, circular convolution, a preprocessing step builds a $[K_w \times (I_h \cdot I_w)]$ matrix by replicating and rotating the rows of the $[I_h \times I_w]$ input matrix, wherein in some implementations “rotation” refers to a cyclic permutation of the elements of a vector, e.g. rotate $([1,2,3,4], -1) \Rightarrow [2,3,4,1]$. In the GEMM step, the product of the $[K_h \times K_w]$ filter matrix and the $[K_w \times (I_h \cdot I_w)]$ matrix from the pre-processing step is computed. In post-processing, the rows of the $[K_h \times (I_h \cdot I_w)]$ matrix created by the GEMM are rotated and added to build the output.

According to some embodiments, the cng2 algorithm may be modified to implement other boundary conditions. As an example, for the case of padded convolution during pre- and post-processing, the vector rows are shifted rather than rotated. That is, the elements that would otherwise wrap around the row vectors during the rotation step are set to zero. Other boundary conditions which may be implemented in the cng2 algorithm include, but are not limited to, symmetric or mirror boundary conditions.

Additionally, it may be noted that the preprocessing step of the cng2 algorithm is not limited to being applied only to the left-hand-side input (the image, herein), but could rather be applied to the right-hand-side input (the filter, herein) according to some embodiments. For full- or valid-mode convolution, the operation is commutative, and the pre-processing phase could be applied to either input. For same-mode convolution, the operation is non-commutative when $I_h \neq K_h$ or $I_w \neq K_w$, but the pre-processing phase can still be applied to the right-hand-side, though the filter must first be zero-padded and/or cropped in each dimension to match the output size.

In some implementations, the cng2 algorithm may include additional steps, as described in FIG. 3-3A. Prior to the pre-processing stage as described previously, the input filter matrix or matrices may need to be reshaped into filter matrix F with appropriate dimensions, as shown in act 3-302. This reshaping may be done by concatenating the input filter matrix or matrices in any suitable manner. In act 3-304, the pre-processing step of building circulant matrix H is performed, and will be described in more detail with reference to FIG. 3-3B. In act 3-306, the GEMM step, is performed and an intermediate matrix $X = F \times H$ is created. Next, post-processing steps may be performed. In act 3-308, vector rows of matrix X are rotated and/or shifted to form matrix X' . In act 3-310, vector row addition is performed on the rows of matrix X' to form matrix Z . Depending on the memory layout of the particular processing system, matrix Z may be reshaped into an at least one output matrix in act 3-312.

The method of building matrix H may depend on the desired boundary conditions, as shown in an expansion of act 3-304 in FIG. 3-3B, according to some embodiments. In act 3-314, it may be determined whether the boundary condition is circular. If it is determined that the boundary condition is circular, the processing system may proceed to act 3-316, wherein matrix H is created by replicating and rotating rows of the at least one input matrix. If, instead, in act 3-314, it is determined that the boundary condition is padded rather than

circular, the processing system may proceed to act 3-318. In act 3-318, matrix H is created by replicating and shifting the rows of the at least one input matrix, as discussed previously. It is to be understood that other boundary conditions than a padded boundary condition may be employed in act 3-318 in some embodiments.

Alternately, according to some embodiments, when computing the cross-correlation the problem may not need to be explicitly converted into a convolution as in process 3-300. Instead, the element-reversal step 3-302 may be omitted and the pre- and post-processing steps of the cng2 algorithm can be modified accordingly. That is, the element-reversal step may be combined with the pre- and post-processing steps of the cng2 algorithm. How this is done depends on whether the pre-processing expansion is applied to the left-hand-side or right-hand-side input. If the left-hand-side input is expanded, shifts or rotations in both the pre- and post-processing steps may be carried out in the opposite direction. If the right-hand-side input is expanded, each of the circulant matrices generated during the preprocessing phase may be transposed and concatenated in the reverse order and the i^{th} row of the GEMM output matrix may be shifted or rotated by $(i - n + 1) \cdot n$ elements rather than $i \cdot n$ elements in the post-processing phase. For complex-valued data the cross-correlation still requires complex conjugation of one input.

In some implementations, such as in CNNs, it may be desirable to generalize the above-described operations so that they can be applied to and/or produce multi-channel data. For a problem with C input channels and M output channels, the filter matrix takes the form $[(M \cdot K_h) \times (K_w \cdot C)]$, the input matrix takes the form $[(K_w \cdot C) \times (I_h \cdot I_w)]$, and the output matrix takes the form $[(M \cdot K_h) \times (I_h \cdot I_w)]$.

Referring to FIGs. 3-4A through 3-4F, an example of process 3-300 for multi-channel inputs and outputs is depicted, according to some embodiments. For $[2 \times 2]$ filters, f , comprising 4 output channels and $[3 \times 3]$ images, G , comprising 3 input channels, act 3-402 is visualized in FIG. 3-4A. However, any size filter matrices, image matrices, and/or number of input and/or output channels may be implemented. In the example of act 3-402, reshaping of the filters f into the $[6 \times 8]$ filter matrix F is performed. In this example, reshaping of filters f is done by concatenating filters f without otherwise altering the ordering of the matrix elements, though other methods such as rotating, shifting, or otherwise altering rows of filters f may be employed. The reshaping of filters f ensures that filter matrix F is of the appropriate dimensionality for the later GEMM operation. However, in some

implementations, wherein the memory is laid out appropriately, act 3-402 may not be required prior to act 3-404 as described below.

According to some embodiments, after act 3-402, pre-processing of image G may be performed in act 3-404, as depicted in FIG. 3-4B. In this example, image G is formed to perform same-mode convolution, though any output mode may be used. Circulant matrix H in this example is formed based on circular boundary conditions, though any boundary conditions such as padded boundary conditions, for example, may be used. After act 3-404, the GEMM operation $F \times H = X$ of act 3-406 may be performed, as depicted in FIG. 3-4C. In this example, intermediate matrix X has dimensions of $[9 \times 8]$.

After the GEMM operation, post-processing steps may occur, as depicted in FIG. 3-4D through 3-4E. In FIG. 3-4D, act 3-408 is depicted, wherein rows of the intermediate matrix X are rotated to form matrix X' in accordance with the circular boundary conditions of this example. Other boundary conditions, such as padded boundary conditions, as a non-limiting example, may be implemented in pre- and post-processing, as long as the boundary conditions of the pre- and post-processing steps are identical to one another.

Referring to FIG. 3-4E, in which act 3-410 is depicted, the next step of post-processing adds rows of the matrix X' to form output matrix Z . That is, in this example, $x_{00} + x_{16} = z_{00}$, $x_{01} + x_{17} = z_{01}$, etc. In some implementations, depending on how the memory of the processing system is laid out, reshaping of output matrix Z may have to be performed in act 3-412 after act 3-410. In the example visualization of FIG. 3-4F, output matrix Z is reshaped into four output matrices A of dimensions $[3 \times 3]$.

In addition to being generalizable to multiple input channels, the cng2 algorithm may be generalized to higher-dimensional signals (i.e. greater than two), according to some embodiments. For an n -dimensional convolution between a filter tensor of size $[K_n \times K_{n-1} \times \dots \times K_1]$ and an image of size $[I_n \times I_{n-1} \times \dots \times I_1]$, it is possible to compute the desired output using two-dimensional matrix multiplication with similar steps to those taken for two-dimensional signals. During pre-processing, the input tensor may be expanded by a factor of $(K_a \cdot K_{a-1} \cdot \dots \cdot K_1)$, where a may be thought of as the number of dimensions handled during the pre-processing phase and any value in the range $1 \leq a \leq n - 1$ may be chosen. In the GEMM step, a product of the filter tensor partitioned as a $[(K_n \cdot K_{n-1} \cdot \dots \cdot K_{a+1}) \times (K_a \cdot K_{a-1} \cdot \dots \cdot K_1)]$ matrix and the expanded matrix from the pre-processing step

may be performed. During the post-processing step, the subvectors of the matrix produced during the GEMM may be rotated and accumulated.

The expanded matrix produced by the pre-processing phase may consist of $(I_i \cdot I_{n-1} \cdot \dots \cdot I_{a+1})$ horizontally-concatenated submatrices where each submatrix is a nested Toeplitz matrix of degree a and the innermost Toeplitz matrices are defined as they are in a two-dimensional cng2 implementation. The post-processing phase may perform $(n - a)$ rounds of rotations and additions where the i^{th} round partitions the matrix produced by the previous round (or, initially, by the GEMM operation) into submatrices of size $[K_{a+i} \times (I_{a+i} \cdot I_{a+i-1} \cdot \dots \cdot I_1)]$. For each submatrix, the following operations are then performed. First, the j^{th} row may be rotated or shifted by $(j \cdot (I_{a+i-1} \cdot I_{a+i-2} \cdot \dots \cdot I_1))$ elements. Then, all rows may be added together.

While the above description handles the dimensions in order, that is the pre-processing phase expands the data along the first a dimensions and the post-processing phase reduces the data along the final $n - a$ dimensions, according to some embodiments, this does not need to be the case. The pre-processing phase could expand the data along any a dimensions by re-ordering the input and output data in the same manner as was described for the two-dimensional case.

The cng2 algorithm offers a flexible framework for computing convolutions, with several alternate embodiments described herein. In some implementations, the overlapping regions of the input signals for a given point in the output may be shifted by a constant offset. Such an offset may be applied regardless of output mode but is most often paired with same-mode output. For convolution (cross-correlation) operating in same-mode and the definitions given above, the boundary condition may be applied to $(K_h - 1) \cdot I_w$ elements along the top (bottom) edge and $(K_w - 1) \cdot I_h$ elements along the left (right) edge of the input image G . This behavior may be altered by redefining the operation with a constant offset between the filter and output locations. When computing the convolution (cross-correlation), this modification can be applied to cng2 by subtracting (adding) the offset to the shift or rotation amounts in the pre-processing phase and by subtracting (adding) $\text{offset} \cdot I_w$ to the shift or rotation amounts in the post-processing phase.

Additionally, methods that have been proposed for reducing both the time and storage requirements of the kn2row post-processing step may similarly be applied to the cng2 algorithm, according to some embodiments. For the kn2row algorithm, the GEMM operation

may be broken into a series of $K_w \cdot K_h$ smaller GEMM operations, wherein the results of those smaller GEMM operations are continually accumulated together. This enables post-processing additions to be performed both in an optimized fashion and in-place with respect to the final output's storage. In the case of kn2row, this only works if the boundary conditions can be ignored or if an additional (and generally inefficient) so-called hole-punching process is introduced. But, in the case of the cng2 algorithm, this process can be applied directly without sacrificing accuracy or additional processing, effectively eliminating the computational cost of the post processing step and reducing the required temporary storage for the cng2 algorithm to $O(K_w \cdot I_h \cdot I_w)$.

In some embodiments, the spatial dimensions could be handled in the opposite order as described in process 3-300. The cng2 algorithm could be augmented with transpose operations applied to both input signals at the start of process 3-300 as well as a transpose operation on the final output. This still produces the desired result but changes the behavior when the shape of the filter is strictly rectangular (i.e. $K_w \neq K_h$). In this case, the input image is expanded by a factor of K_h rather than K_w and the post-processing step consists of $O(K_w \cdot I_h \cdot I_w)$ additions rather than $O(K_h \cdot I_h \cdot I_w)$. An implementation that combines this variant with the low-memory integrated-post-processing variant above can further reduce the required temporary storage for the cng2 algorithm to $O(\min(K_h, K_w) \cdot I_h \cdot I_w)$.

As an alternative implementation, the rows and/or columns in the matrices that are passed to the GEMM operation may be re-ordered. If the GEMM operation is defined as $C = AB$, the rows and/or columns of either input matrix A or B , may be re-ordered so long as the appropriate permutation is applied to the other input matrix (in the case of re-ordering the columns of A or rows of B) or the output matrix (in the case of re-ordering the rows of A or the columns of B). In particular, re-ordering the rows of A in the case of multiple output channels may reorganize the data-level parallelism available in the post-processing phase in a manner that is well suited for vector processors or single-instruction-multiple-data (SIMD) architectures.

The convolution computation may also be computed with a stride, according to some embodiments. For stride S_x in the first dimension and stride S_y in the second dimension, the convolution operation is defined as follows:

$$(G * F)[x, y] = \sum_{i=0}^{K_h} \sum_{j=0}^{K_w} F[i, j] \cdot \hat{G}[x \cdot S_x - i, y \cdot S_y - j]$$

This definition reduces the size of the output signal by a factor of $S_x \cdot S_y$, and is equivalent to increasing the step size by which the filter is slid across the image for each output point. This may be implemented by computing the un-strided convolution and then down-sampling the result by the appropriate amount in each dimension, but this requires $O(S_x \cdot S_y)$ more computation steps than necessary. At a minimum, this computational penalty can be reduced in cng2 to $O(S_y)$ by modifying the pre-processing phase 3-304 to generate only every S_x^{th} column in each individual circulant matrix and modifying the post-processing phase to shift or rotate each row by $i \cdot (I_w/S_x)$ rather than $i \cdot I_w$. In some implementations, the computational penalty can be completely eliminated with additional modifications to each phase. First, the preprocessing step 3-304 may be modified to produce S_y expanded matrices rather than a single matrix, where the i^{th} circulant matrix is assigned to the j^{th} expanded matrix if $j = i \bmod S$. The core processing phase must then perform S_y GEMM operations—one GEMM operation per expanded input matrix—each of which uses only K_w/S_y rows of the filter matrix. The post-processing steps 3-308, 3-310 may then interleave the rows of the resulting matrices, add each group of S_y rows directly (i.e., without shifting or rotating the rows), and run the K_w/S_y rows through the standard post-processing logic with shift or rotation amounts of $i \cdot (I_w/S_x)$.

Alternately, the convolution may be dilated, according to some embodiments. For dilation D_x in the first dimension and dilation D_y in the second dimension, the convolution operation is defined as:

$$(G * F)[x, y] = \sum_{i=0}^{K_h} \sum_{j=0}^{K_w} F[i, j] \cdot \hat{G}[x - i \cdot D_x, y - j \cdot D_y]$$

Dilation increases the receptive field of the filter across a larger patch of the image for each output point, and may be viewed as inserting spaces between elements of the filter. The cng2 algorithm may be modified to implement dilated convolution by increasing the rotation or shift amounts in both the pre- and post-processing phases by D_x and D_y , respectively. A dilated convolution may be further restricted to being computed with a causal output mode.

The inventors have further recognized and appreciated that convolutions and cross-correlations may be computed by using a transform-based algorithm. Transform-based algorithms change the nature of the computational problem by first computing the equivalent representation of the input signals in an alternative numerical domain (e.g. frequency

domain), performing an alternative linear operation (e.g. element-wise multiplication), and then computing the inverse transform of the result to return to the signal's original numerical domain (e.g. time domain). Examples of such transforms include discrete Fourier transforms (DFT), discrete sine transforms, discrete cosine transforms, discrete Hartley transforms, undecimated discrete wavelet transforms, Walsh-Hadamard transforms, Hankel transforms, and finite impulse response (FIR) filters such as Winograd's minimal filtering algorithm. An example of a transform-based algorithm based on a DFT will be described herein, but any suitable transform may be implemented in a transform-based algorithm and on a photonic processor.

For unitary normalization, the discrete Fourier transform (DFT) of a one-dimensional signal is computed as

$$X(k) = \mathcal{T}_{1D}(x(n)) = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} x(n) \cdot e^{-\frac{2\pi i}{N} kn}.$$

The inverse of this transform $\mathcal{T}_{1D}^{-1}$ may be computed by taking the complex conjugate. Similarly, in two dimensions, the unitary normalized DFT may be computed as

$$X(k, l) = \mathcal{T}_{2D}(x(m, n)) = \frac{1}{\sqrt{N}} \frac{1}{\sqrt{M}} \sum_{m=0}^M \sum_{n=0}^N x(m, n) \cdot e^{-2\pi i (\frac{km}{M} + \frac{ln}{N})}.$$

Performing the one-dimensional DFT defined above on a vector of size N can be accomplished by computing a matrix-vector product $\mathcal{T}_{1D}(x) = X = Wx$. The matrix W is referred to as the transform matrix, given by

$$W = \frac{1}{\sqrt{N}} \begin{pmatrix} 1 & 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega & \omega^2 & \omega^3 & \cdots & \omega^{N-1} \\ 1 & \omega^2 & \omega^4 & \omega^6 & \cdots & \omega^{2(N-1)} \\ 1 & \omega^3 & \omega^6 & \omega^9 & \cdots & \omega^{3(N-1)} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{N-1} & \omega^{2(N-1)} & \omega^{3(N-1)} & \cdots & \omega^{(N-1)(N-1)} \end{pmatrix}$$

where $\omega \equiv e^{-\frac{2\pi i}{N}}$. The inverse transform may be computed by a similar matrix-vector product where the elements of W^{-1} are the complex conjugates. The DFT is a separable transform, so it may be regarded as computing two one-dimensional transforms along orthogonal axes. Thus, the two-dimensional DFT of an $[M \times N]$ (i.e. rectangular) input x may be computed via the following matrix triple product:

$$\mathcal{T}_{2D}(x) = X = WxY^T,$$

where W is an $[M \times M]$ transform matrix associated with the columns and Y is an $[N \times N]$ transform matrix associated with the rows, and the superscript T indicates the matrix

transpose. In the case of a square input x of size $[N \times N]$, the transform matrix for the columns W is the same as the transform matrix for the rows Y .

Equivalently, this may be computed by first flattening x row-wise into a column vector x_{col} of size $M \cdot N$ and computing the following matrix-vector product:

$$X_{col} = (W \otimes Y) x_{col},$$

where $\otimes$ is the Kronecker product. According to some embodiments, the result vector X_{col} may then be reshaped into an $[M \times N]$ two-dimensional array X :

$$\mathcal{T}_{2D}(x) = X = \text{reshape}(X_{col}).$$

A similar process may be performed for other discrete transforms where forward transform matrix W and the matrix W^{-1} associated with the inverse transform are defined in any suitable way in accordance with said other transforms.

In the case of one-dimensional DFT, the matrix W is a unitary matrix, and may therefore be programmed directly into the photonic array in accordance with previously described embodiments. For other discrete transforms, the matrix W may not be unitary and thus require decomposition before being programmed into the photonic array in accordance with previously described methods. A process 3-500, according to some embodiments, for performing a one-dimensional transform on a vector is shown in FIG. 3-5. In act 3-502, the matrix W is programmed into the photonic array, and in act 3-504, the vector x is propagated through the array.

In some implementations, a two-dimensional transform may be computed as described in process 3-600 of FIG. 3-6. For a two-dimensional input x that is of size N -by- N , process 3-500 may be modified to produce a two-dimensional transform as defined above. In act 3-602, a transform matrix W_{small} is created corresponding to a size N one-dimensional input. Next, in act 3-604, a block-diagonal matrix B of size N^2 -by- N^2 is created by tiling W along the diagonal. A column vector x_{col} of size N^2 is then created by flattening the input x row-wise in act 3-606. In act 3-608, by propagating the vector x through the photonic array, the multiplication $X_{partial} = Bx$ may be performed. In act 3-610, the matrix $X_{partial}$ may then be reshaped into an N -by- N matrix which can be further flattened column-wise into a size N^2 column vector in act 3-612. In act 3-614, the multiplication $X = BX_{partial}$ may be performed by propagating $X_{partial}$ through the photonic array. In act 3-616, the resulting vector X may be reshaped into an N -by- N matrix for output.

In some embodiments, the two-dimensional transform of an $[N \times N]$ input x may then be computed by first programming the matrix W into the photonic array. Second, computing $X_{\text{partial}} = Wx$ by propagating the columns of x through the photonic array. Third, transposing the partial result X_{partial} . Fourth, propagating the columns of X_{partial}^T through the array a second time to compute WX_{partial}^T . Finally transposing the result to produce $X = WxW^T$.

Some systems, such as one embodiment of the photonic-based systems described herein, are limited to implementing real unitary matrices (that is, orthogonal matrices). In such implementations, the transform can still be computed, but additional steps are needed. The system must keep track of the real and imaginary parts of the transform matrix and input vector or image separately. The embodiments defined above for computing the products can be adapted for orthogonal matrices, except that for every pass through the photonic array as described above, the algorithm must perform four passes. Denoting the real part of a variable as $Re(x)$ and imaginary part as $Im(x)$, the real part of the product is $Re(Wx) = Re(W)Re(x) - Im(W)Im(x)$ and similarly the imaginary part of the product is $Im(Wx) = Re(W)Im(x) + Im(W)Re(x)$. According to some embodiments, in the photocore of a photonic processor representing only real matrices, the process 3-700 of FIG. 3-7 may be carried out. According to the dimensionality of the input, either process 3-500 or 3-600 may be used in process 3-700. In act 3-702, $Re(W)$ is loaded into the photonic array. In act 3-704, either process 3-500 or 3-600 may be performed, depending on the dimensionality of the input, for $Re(x)$ and $Im(x)$. This produces $Re(W)Re(x)$ and $Re(W)Im(x)$. In act 3-706, $Im(W)$ is loaded into the photonic array. In act 3-708, either process 3-500 or 3-600 may be performed, depending again on the dimensionality of the input, for $Re(x)$ and $Im(x)$. This produces $Im(W)Re(x)$ and $Im(W)Im(x)$. In act 3-710, $Re(W)Re(x)$ and $Im(W)Im(x)$ are subtracted to produce $Re(Wx)$. And, in act 3-712, $Re(W)Im(x)$ and $Im(W)Re(x)$ are added to produce $Im(Wx)$.

With the above-described processes 3-500, 3-600, and 3-700, an input matrix may be converted into its transform. Once the convolutional filter F and image G are converted to their transform counterparts, the convolution theorem may be applied, according to some embodiments. The convolution theorem states that a convolution of two signals corresponds to the transform of an element-wise product of the two signals' transforms. Mathematically, this is represented by:

$$\mathcal{T}(G * F) = \mathcal{T}(G) \odot \mathcal{T}(F),$$

or, equivalently,

$$G * F = \mathcal{T}^{-1}(\mathcal{T}(G) \odot \mathcal{T}(F)),$$

where $\odot$ represents element-wise multiplication, and $\mathcal{T}^{-1}$ represents the inverse transform. In some embodiments, the dimensions of the image and of the filter may differ; in such a case it is to be appreciated that the appropriate dimension transform matrices may be used to compute each of the forward transforms and the inverse transform. The matrix-multiplication equation representing the one-dimensional convolution with a general transform and general dimensionality of the filter and image is thus:

$$G * F = W_A^T((W_B F) \odot (W_D^T G)),$$

where W_B is the matrix associated with the transform of the filter F , W_D^T is the matrix associated with the transform of the image G , and W_A^T is the matrix associated with the inverse transform of the combined signal.

Similarly, the matrix-multiplication equation representing the two-dimensional convolution with a general transform on rectangular filters and images is thus:

$$G * F = W_A^T((W_B F W_C^T) \odot (W_D^T G W_E)) W_F,$$

where W_B and W_C^T are the matrices associated with the transform of the filter F , W_D^T and W_E are the matrices associated with the transform of the image G , and W_A^T and W_F are the matrices associated with the inverse transform of the combined signal.

Referring to FIG. 3-8, some embodiments of performing convolutions in the photocore of a photonic processor using a transform-based algorithm may include the following acts. In act 3-802, a transform of image G may be performed using any one of processes 3-500, 3-600, and/or 3-700.

In some embodiments, the filter F may then be padded with zeros in act 3-804 to match the size of the image G , after which a transform may be performed on filter F using any one of processes 3-500, 3-600, and/or 3-700 in act 3-806. In act 3-808, the transformed filter F may then be loaded into the element-wise multiplier of the photonic array, and in act 3-810, an image G may be propagated through the photonic array. In act 3-812, an inverse transform may be performed on the result of the previous computation using any one of processes 3-500, 3-600, and/or 3-700. The result of act 3-812 may then be reshaped in act 3-814 into the size of G and cropped to produce the final convolved image, $G * F$.

In some embodiments, the convolution $G * F$ may be computed in a divide-and-conquer fashion where one input is partitioned into a set of tiles and each tile is convolved with the second input separately. The results of each individual convolution can then be recombined into the desired output, but the algorithms (e.g., overlap-add, overlap-save) for implementing this divide-and-conquer approach are non-trivial. When one input is much smaller than the other and a transform-based algorithm is used for the convolution operation, this approach can be much more efficient than computing the entire convolution in one operation as described above with the filter being padded to match the size of the image. It may be appreciated that by performing the transformations of the tiles on a photonic array, such a divide-and-conquer algorithm for transform-based convolutions may be implemented on a photonic processor.

In some embodiments, the filter F and the image G may have multiple channels. As defined above, this means each channel of the image is convolved with the corresponding channel of the filter tensor, and the results are added together element-wise. When a multi-channel convolution is computed with a transform-based method, the summation across channels may be performed in either the transform or the output domain. In practice, it is often chosen to perform the summation in the transform domain because this decreases the amount of data on which the output transform must be applied. In this case, the element-wise multiplication followed by channel-wise summation can be expressed as a sequence of matrix-matrix multiplications (GEMMs). Mathematically, this can be expressed as follows:

Let G be an input signal comprising C data channels of $N \times N$ images. Let F be an input signal comprising $M \cdot C$ data channels of $N \times N$ filters. Let C and M be the number of input and output data channels, respectively. Let $Q_{m,c}$ be transformed data of the m^{th} output channel and the c^{th} input channel of the filter tensor (i.e., $Q_{m,c} = W_B F_{m,c} W_C^T$). Let R be the transformed three-dimensional $[C \times N \times N]$ input tensor and R_c be the c^{th} channel of the transformed input tensor (i.e., $R_c = W_D^T G_c W_E$). Then, the convolution of F and G producing multiple output channels is:

$$(G * F)_m = W_A^T \left(\sum_{c=1}^C Q_{m,c} \odot R_c \right) W_F \quad \forall \quad m \in [1, M].$$

If S^{ij} denotes a column vector comprised of the C elements in the $(i, j)^{\text{th}}$ position of each channel in a three-dimensional $[C \times N \times N]$ tensor S , this can be equivalently expressed as:

$$(G * F)_m^{ij} = W_A^T (Q_m^{ij} R^{ij}) W_F \quad \forall \quad i \in [1, N], \quad j \in [1, N], \quad m \in [1, M]$$

Each of the $Q_m^{ij} R^{ij}$ matrix-matrix multiplications may be computed on a photonic processor as described above. This may further be combined with the divide-and-conquer approaches described above.

Aspects of the present application provide methods, procedures and algorithms which may be performed on a processing device, such as a CPU, GPU, ASIC, FPGA or any other suitable processor. For example, the processing device may perform the procedures described above to generate settings for the variable beam splitters and modulators of the photocore of the photonic processor described herein. The processing device may also perform the procedures described above to generate the input data to be input into the photonic processor described herein.

One example implementation of a computing device may include at least one processor and a non-transitory computer-readable storage medium. The computing device may be, for example, a desktop or laptop personal computer, a personal digital assistant (PDA), a smart mobile phone, a tablet computer, a server, or any other suitable computing device. The computer-readable media may be adapted to store data to be processed and/or instructions to be executed by processor. The processor enables processing of data and execution of instructions. The data and instructions may be stored on the computer-readable storage media and may, for example, enable communication between components of the computing device. The data and instructions stored on computer-readable storage media may comprise computer-executable instructions implementing techniques which operate according to the principles described herein.

A computing device may additionally have one or more components and peripherals, including input and output devices. These devices can be used, among other things, to present a user interface. Examples of output devices that can be used to provide a user interface include printers or display screens for visual presentation of output and speakers or other sound generating devices for audible presentation of output. Examples of input devices that can be used for a user interface include keyboards, and pointing devices, such as mice, touch pads, and digitizing tablets. As another example, a computing device may receive input information through speech recognition or in other audible format. As another example, a computing device may receive input from a camera, lidar, or other device that produces visual data.

Embodiments of a computing device may also include a photonic processor, such as the one described herein. The processor of the computing device may send and receive information to the photonic processor via one or more interfaces. The information that is sent and received may include settings of the variable beam splitters and modulators of the photonic processor and/or measurement results from the detectors of the photonic processor.

IV. Photonic Encoder

The inventors have recognized and appreciated that phase-intensity relationships, whereby phase and intensity modulation of optical signals are interdependent of one another, pose a challenge for precisely encoding vectors in light fields for optical processing.

Some optical modulators encode numeric vectors in the optical domain, where numbers are encoded onto the phase or the intensity of an optical signal. In an ideal world, an optical intensity modulator is capable of modulating the intensity of an optical signal while maintaining the phase of the optical signal unchanged, and an optical phase modulator is capable of modulating the phase of an optical signal while maintaining the intensity of the optical signal unchanged. In the real world, however, the phase and the intensity of an optical signal are mutually interdependent. As such, intensity modulation gives rise to a corresponding phase modulation and phase modulation gives rise to a corresponding intensity modulation.

Consider for example integrated photonics platforms, in which intensity and phase are related to one another by the Kramers-Kronig equations (bidirectional mathematical relations connecting the real and imaginary parts of a complex analytic function). In these cases, the optical field can be represented by a complex function where the real part and the imaginary part of the function (and as a result, the intensity and the phase) are related to one another. In addition, realistic implementations of intensity and phase modulators generally suffer from dynamic loss whereby the amount of modulation they impart (phase or intensity) depends on their current settings. These modulators experience a certain power loss when no phase modulation occurs, and experience a different power loss when phase modulation occurs. For example, the power loss experienced at no phase modulation may be L_1 , the power loss experienced at a $\pi/2$ -phase modulation may be L_2 , and the power loss experienced at a π -phase modulation may be L_3 , with L_1 , L_2 and L_3 being different from each other. This behavior is undesirable because, in addition to phase modulation, the signal further experiences intensity modulation.

An example of an optical modulator (4-10) is depicted in FIG. 4-1, where ψ_{in} represents an input optical field, ψ_{out} represents the optical field output by modulator 4-10, α represents the amplitude modulation factor and θ represents the phase modulation factor. Ideally, modulator 4-10 would be able to set the entries of the optical field ψ_{in} to any phase in combination with any attenuation as $\psi_{\text{out}} = \alpha e^{i\theta}$, with $\alpha \leq 1$ and $\theta \in [0, 2\pi]$. Because modulators typically have an associated Kramers-Kronig relation, and/or are subject to dynamic loss, simply setting the values of α and θ cannot be easily accomplished. Conventionally, two or more modulators with two independently controllable modulating signals are used to encode both the amplitude and phase of the optical field in a way that is representative of a real, signed number. However, having multiple controllable modulating signals calls for complicated multi-variable encoding schemes, often involving feedback loops for simultaneously controlling the phase and intensity of the optical signal.

Optical modulators that are capable of modulating intensity without affecting the phase exist, including modulators based on electro-optic materials. Unfortunately, electro-optic modulators are challenging to work with in large-scale commercial contexts due the fact that they involve use of materials impractical to fabricate.

Conventional computers do not natively perform complex arithmetic. Rather, they use signed arithmetic where numbers can be positive or negative and complex arithmetic can be built up by performing several calculations and combining the results. The inventors have recognized and appreciated that this fact may be exploited to dramatically simplify the implementation of real, signed linear transformations, even in the presence of dynamic loss and/or non-ideal phase or intensity modulators.

Some embodiments relate to techniques for encoding vectors of signed, real numbers using non-ideal optical modulators (modulators in which intensity modulation gives rise to phase modulation and phase modulation gives rise to intensity modulation). Contrary to other implementations, some such embodiments involve use of a single modulating signal to control the phase as well as the intensity of an optical signal. The inventors have appreciated that, when encoding signed, real numbers onto optical field using a single modulating signal, the exact location of the optical signal in phase space (e.g., the real and imaginary part, or the amplitude and phase) is not important for purposes of decoding. What matters for accurate decoding, according to some embodiments, is the projection of the optical signal on the real axis (or other arbitrary axes pre-selected as the measurement axis). To perform the projection on the pre-selected axis, in some embodiments, coherent detection schemes are used.

Optical domain encoding techniques of the types described herein may be used in a variety of contexts, including but not limited to high-speed telecommunications for short, mid and long-haul applications, on chip-phase sensitive measurements for sensing, communications and computing, and optical machine learning using photonic processors.

More generally, encoding techniques of the types described herein may be used in any context in which optical signals are processed according to real transformations (as opposed to complex transformations). Conventional computers do not directly perform complex arithmetic. Rather, conventional computers use signed, real arithmetic where numbers can be positive or negative. In some embodiments, complex arithmetic can be built up in the optical domain by performing several calculations involving real number arithmetic. Consider for example an optical linear system configured to transform optical signals according to the following expression:

$$\begin{aligned} y &= Mx = [Re(M) + iIm(M)][Re(x) + iIm(x)] \\ &= [Re(M)Re(x) - Im(M)Im(x)] + i[Re(M)Im(x) + Im(M)Re(x)] \end{aligned}$$

where x represents an input vector, M represents a transformation matrix, y represents the output vector and i represents the imaginary number defined such that $i^2 = -1$. Considering real transformations (such that $Im(M) = 0$), y can be rewritten as follows:

$$y = e^{i\gamma} M [Re(x) + iIm(x)] = e^{i\gamma} M Re(x) + e^{i(\gamma + \frac{\pi}{2})} M Im(x)$$

It should be noted that, based on this equation, the real and imaginary parts of the input field x contribute only to linear independent components of the resulting field y . The inventors have recognized and appreciated that y can be decoded using a coherent receiver by projecting y on any arbitrary axis in the complex plane.

In some embodiments, encoding techniques of the types described herein may be used as part of a photonic processing system. For example, in some embodiments, optical encoder 1-101 of FIG. 1-1 and/or optical encoders 1-1103 of FIG. 1-11 and/or optical encoders 1-1211 of FIG. 1-12 B may implement encoding techniques of the types described herein.

FIG. 4-2A is a block diagram of a photonic system implementing optical encoding techniques according to some embodiments. Photonic system 4-100 includes an optical source 4-102, an encoder 4-104, an optical modulator 4-106, and optical transformation unit 4-108, a coherent receiver 4-110, a local oscillator 4-112 and a decoder 4-114. In some embodiments, photonic system 4-100 may include additional or alternative components not

illustrated in FIG. 4-2A. In some embodiments, some or all of the components of FIG. 4-2A may be disposed on the same semiconductor substrate (e.g., a silicon substrate).

Optical source 4-102 may be implemented in a number of ways, including, for example, using optical coherent sources. In one example, optical source 4-102 includes a laser configured to emit light at wavelength λ_0 . The wavelength of emission may be in the visible, infrared (including near infrared, mid infrared and far infrared) or ultraviolet portion of the electromagnetic spectrum. In some embodiments, λ_0 may be in the O-band, C-band or L-band. Light emitted by optical source 4-102 is modulated using optical modulator 4-106. Optical modulator 4-106 is a non-ideal modulator, such that a phase modulation gives rise to an intensity modulation and an intensity modulation gives rise to a phase modulation. In some embodiments, the phase modulation may be related to the intensity modulation according to the Kramers-Kronig equations. Alternatively, or additionally, modulator 4-106 may suffer from dynamic loss where a phase shift leads to attenuation.

In some embodiments, modulator 4-106 is driven by a single electrical modulating signal 4-105. Thus, a single modulating signal modulates both the phase and the amplitude of the optical field. This is in contrast to modulators conventionally used in optical communications to encode symbols in the complex amplitude of an optical field, where each symbol represents more than one bit. In such types of modulators, in fact, multiple modulating signals modulate the optical field. Consider for example optical modulators configured to provide quadrature phase-shift keying (QPSK) modulation schemes. In these types of modulators, one modulating signal modulates the real part of the optical field and one modulating field modulates the imaginary part of the optical field. Thus, in QPSK modulators, the phase and the amplitude of the optical field are modulated, collectively, using two modulating signals.

Examples of modulators that may be used for modulator 4-106 include Mach Zehnder modulators, electro-optical modulators, ring or disk modulators or other types of resonant modulators, electro-absorption modulators, Frank-Keldysh modulators, acousto-optical modulators, Stark-effect modulators, magneto-optical modulators, thermo-optical modulators, liquid crystal modulators, quantum-confinement optical modulators, and photonic crystal modulators, among other possible types of modulators.

Encoder 4-104 generates modulating signal 4-105 based on the real number to be encoded. For example, in some embodiments, encoder 4-104 may include a table mapping real numbers to amplitudes for the modulating signal.

Optical transformation unit 4-108 may be configured to transform the intensity and/or the phase of the received optical field. For example, optical transformation unit 4-108 may include an optical fiber, an optical waveguide, an optical attenuator, an optical amplifier (such as an erbium-doped fiber amplifier), a beam splitter, a beam combiner, a modulator (such as an electro-optic modulator, a Franz-Keldysh modulator, a resonant modulator, or a Mach Zehnder modulator, among others), a phase shifter (such as a thermal phase shifter or a phase shifter based on the plasma dispersion effect), an optical resonator, a laser, or any suitable combination thereof. In the context of optical communications, optical transformation unit 4-108 may include a fiber optic communication channel. The optical communication channel may include for example optical fibers, and optionally, optical repeaters (e.g., erbium-doped fiber amplifiers). In the context of optical processing, optical transformation unit 4-108 may include a photonic processing unit, an example of which is discussed in detail further below. In some embodiments, optical transformation unit 4-108 implements a real transformation, such that the imaginary part of the transformation is substantially equal to zero.

The optical field output by optical transformation unit 4-108 is provided, directly or indirectly (*e.g.*, after passing through one or more other photonic components), to coherent receiver 4-110. Coherent receiver 4-110 may include a homodyne optical receiver or a heterodyne optical receiver. The reference signal with which the received signal is beaten may be provided by local oscillator 4-112, as shown in FIG. 4-2A, or may be provided together with the modulated optical field after passing through optical transformation unit 4-108. Decoder 4-114 may be arranged to extract a real number (or a vector of real numbers) from the signal output by coherent receiver 4-112.

In at least some of the embodiments in which photonic processing system 1-100 of FIG. 1-1 implements encoding techniques of the types described herein, optical modulator 4-106 may be part of optical encoder 1-101, optical transformation unit 4-108 may be part of photonic processor 1-103 and coherent receiver 4-110 may be part of optical receiver 1-105.

FIG. 4-2B is a flowchart illustrating a method for processing real numbers in the optical domain, in accordance with some embodiments. Method 4-150 may be performed using the system of FIG. 4-2A, or using any other suitable system. Method 4-150 begins at act 4-152, in which a value representative of a real number is provided. The real number may be signed (*i.e.*, may be positive or negative) in some embodiments. The real number may represent a certain environmental variable or parameter, such as a physical condition (*e.g.*, temperature, pressure, etc.), information associated with an object (*e.g.*, position, motion,

velocity, rate of rotation, acceleration, etc.), information associated with a multimedia file (e.g., acoustic intensity of audio files, pixel color and/or intensity of image or video files), information associated with a certain chemical/organic element or compound (e.g., concentration), information associated with financial assets (e.g., price of a certain security), or any other suitable type of information including information derived from the examples described above. The information represented by the signed, real number may be useful for a variety of reasons, including for example to train a machine learning algorithm, to perform forecasting, data analytics, troubleshooting, or simply to collect data for future use.

At act 4-154, the value representative of the real number may be encoded onto an optical field. In some embodiments, encoding the value onto an optical field involves modulating the phase and the intensity of an optical field based on the value. As a result, the phase and the amplitude of the optical field reflect the encoded value. Act 4-154 may be performed in some embodiments using encoder 4-104 and modulator 4-106 (see FIG. 4-2A). In some such embodiments, modulating the phase and amplitude based on the value involves driving a single modulator with a single electrical modulating signal. Thus, a single modulating signal modulates both the phase and the amplitude of an optical field. Referring back to FIG. 4-2A by way of example and not limitation, encoder 4-104 may drive optical modulator 4-106 using a single modulating signal 4-105.

It should be noted that the fact that a single modulating signal is used to drive the modulator does not preclude the use of other control signals for controlling the environment in which the modulator operates. For example, one or more control signals may be used to control the temperature of the modulator or the temperature of a certain portion of the modulator. One or more control signals may be used to power the operations of the modulator. One or more control signals may be used to bias the modulator in a certain regime of operation, such as to bias the modulator in its linear region or to set the wavelength of the modulator to match the wavelength of the optical source (e.g., λ_0 of FIG. 4-2A).

A specific type of modulator is illustrated in FIG. 4-3A, in accordance with some embodiments. Modulator 4-206 is a ring resonant modulator. It should be noted that ring modulators are described herein merely by way of example, as any other suitable type of modulator, including those listed above, may alternatively be used. Modulator 4-206 includes a waveguide 4-208, a ring 4-210, and a phase shifter 4-212. Ring 4-210 exhibits a resonant wavelength, the value of which depends, among other parameters, the length of the ring's perimeter. When an optical field ψ_{in} is launched into waveguide 4-208, the optical field may or may not couple to ring 4-210 depending on the wavelength of ψ_{in} relative to the resonant

wavelength of ring 4-210. For example, if the wavelength of ψ_{in} matches the resonant wavelength, at least part of the energy of ψ_{in} is transferred, via evanescent coupling, to ring 4-210. The power transferred to the ring will oscillate indefinitely inside the ring until it is completely scattered or otherwise dissipated. By contrast, if the wavelength of ψ_{in} does not match the resonant wavelength, ψ_{in} may proceed straight through waveguide 4-208 without any significant attenuation.

Thus, the optical field may vary depending on the wavelength of ψ_{in} relative to the resonant wavelength. More particularly, the phase and the intensity of the optical field may vary depending on the wavelength of ψ_{in} relative to the resonant wavelength. FIG. 4-3B is a plot illustrating how the intensity (top chart) and the phase (bottom chart) may vary depending on the wavelength of ψ_{in} relative to the resonant wavelength. The top chart illustrates the intensity spectral response α as a function of all possible wavelengths (λ) that ψ_{in} may have. The bottom chart illustrates the phase spectral response θ as a function of the same wavelength. At the resonant frequency, the intensity response exhibits a dip and the phase response exhibits an inflection point. Wavelengths substantially lower than the resonant frequency are subject to a low intensity attenuation ($\alpha \sim 1$) and a low phase change ($\theta \sim 0$). Wavelengths substantially greater than the resonant frequency are subject to a low intensity attenuation ($\alpha \sim 1$) and a change in sign ($\theta \sim -\pi$). At the resonant wavelength, the intensity of ψ_{in} is attenuated by the value corresponding to the dip and is subject to a phase shift equal to the value of the inflection point (i.e., $-\pi/2$). In the example of FIG. 4-3B, the wavelength (λ_0) of ψ_{in} is slightly offset relative to the resonant frequency. In this case, the intensity of ψ_{in} is attenuated by α_{v0} (where $0 < \alpha_{v0} < 1$) and the phase is shifted by θ_{v0} (where $-\pi < \theta_{v0} < -\pi/2$).

Referring back to FIG. 4-3A, modulation of the amplitude and phase of ψ_{in} may be performed using voltage V, which in this case embodies modulating signal 4-105 of FIG. 4-2A. When a voltage V is applied to phase shifter 4-212, the refractive index of the phase shifter changes and so does, as a result, the effective length of the ring's perimeter. This, in turn, leads to a variation in the resonant frequency of the ring. The extent to which the effective perimeter length (and accordingly, the resonant frequency) varies depends upon the amplitude of V.

In the example of FIG. 4-3B, voltage V is set to zero. In the example of FIG. 4-3C, the voltage is set to V_1 . As shown in FIG. 4-3C, application of V_1 to phase shifter 4-212 leads to a shift in the intensity and phase spectral responses of the modulator along the wavelength axis. In this case, the responses exhibit a redshift (i.e., a shift towards greater wavelengths).

The result is that the intensity of ψ_{in} is now attenuated by α_{V1} (α_{V1} being different than α_{V0}) and the phase is shifted by θ_{V1} (θ_{V1} being different than θ_{V0}).

Thus, varying voltage V leads to a change in intensity and phase of ψ_{in} . In other words, V may be viewed as a modulating signal. In some embodiments, voltage V may be the only modulating signal driving modulator 4-210. Voltage V may assume the following expression in some embodiments $V=V_{DC}+V(t)$, where V_{DC} is a constant and $V(t)$ is varied over time depending on the real numbers to encode. V_{DC} may pre-bias the ring so that resonant frequency is in proximity to the wavelength (λ_0) of ψ_{in} .

FIG. 4-3D is an encoder table (e.g., a look-up table) providing an example of how real numbers may be encoded, using modulator 4-206, into the phase and intensity of an optical field, in accordance with some embodiments. In this case it will be assumed, by way of example, that real numbers between -10 and 10 with increments equal to 1 are to be encoded in the optical domain (see column labeled “real number to encode”). Of course, any suitable set of real numbers may be encoded using the techniques described herein. The column labeled “ V (applied voltage)” indicates the voltage being applied to phase shifter 4-212. The column labeled “ α (amplitude modulation)” indicates the value of the intensity spectral response at the wavelength (λ_0) of ψ_{in} (see FIGs. 4-3B and 4-3C, top charts). The column labeled “ θ (phase modulation)” indicates the value of the phase spectral response at the wavelength (λ_0) of ψ_{in} (see FIGs. 4-3B and 4-3C, bottom charts). In this case, real number “-10” is mapped to a certain voltage V_1 , which leads to an amplitude modulation α_1 and a phase modulation θ_1 ; real number “-9” is mapped to a certain voltage V_2 , which leads to an amplitude modulation α_2 and a phase modulation θ_2 ; and so on. Real number “10” is mapped to a certain voltage V_N , which leads to an amplitude modulation α_N and a phase modulation θ_N . Thus, different real numbers are encoded with different intensity/phase pairs.

FIG. 4-3E provides a visual representation, in the complex plane, of the encoding table of FIG. 4-3D, in accordance with some embodiments. Different points along the line 4-300 represent different intensity/phase pairs, in accordance with the spectral responses of FIGs. 4-3B and 4-3C, as the voltage is varied. The symbol S_{in} , for example, represents a specific symbol characterized by intensity modulation α and a phase modulation θ . S_{in} represents a certain row from the table of FIG. 4-3D. Each row of the table is mapped on a different point of line 4-300.

Referring back to FIG. 4-2B, at act 4-156, a transformation is applied to the modulated optical signal output by the modulator. The transformation may involve an intensity transformation and/or a phase transformation, depending on the way one desires to

process the encoded real numbers. In some embodiments, optical transformation unit 4-108 may be used for the transformation of the signal output by modulator 4-106. Optical transformation unit 4-108 may be configured to transform a symbol S_{in} , which is characterized by an intensity α and phase θ , into a symbol S_{out} , which is characterized by an intensity $\alpha\beta$ and phase $\theta+\Delta\theta$. In other words, optical transformation unit 4-108 introduces an intensity modulation β and a phase shift $\Delta\theta$. The values of β and $\Delta\theta$ depend upon the specific optical transformation unit being used.

Considering by way of example the input symbol S_{in} of FIG. 4-3E, FIGs. 4-4A through 4-4C provide examples of how the optical transformation unit may perform the transformation. The transformation of FIG. 4-4A is such that $\beta<1$ and $\Delta\theta=0$. That is, the optical transformation unit introduces an attenuation but does not alter the phase of the input optical field. This may be the case, for example, of an optical fiber with no optical repeaters, the optical fiber having a length selected to maintain the input phase. The result is that line 4-300 is transformed to line 4-401, where line 4-401 is a compressed version of line 4-300. The output sample, S_{out} , has the same phase as the S_{in} , but the intensity is equal to $\alpha\beta$ (i.e., is attenuated by β).

The transformation of FIG. 4-4B introduces an attenuation as well as a phase shift. This may be the case, for example, of an optical fiber with no optical repeaters, the optical fiber having an arbitrary length. The result is that line 4-300 is transformed to line 4-402, where line 4-402 is a compressed and rotated version of line 4-300. The output sample, S_{out} , has phase $\theta+\Delta\theta$ and intensity $\alpha\beta$.

The transformation of FIG. 4-4C involves a multi-mode transformation. This may occur when the input optical field is combined with one or more other optical fields. With multi-mode transformations, line 4-403 can be re-shaped in any suitable manner. Examples of multi-mode transformations include optical processing units, some of which are described in detail further below.

Referring back to FIG. 4-2B, at act 4-158, a coherent receiver may mix the transformed, modulated optical field with a reference optical signal to obtain an electric output signal. In some embodiments, the mixing may be performed using coherent receiver 4-110 of FIG. 4-2A. The reference optical signal may be a signal generated by a local oscillator (e.g., local oscillator 4-112 of FIG. 4-2A), or may be a signal transmitted together with the modulated optical field through the optical transformation unit. The effect of the mixing is illustrated in FIG. 4-5 by way of example and not limitation. Line 4-403 (the same line shown in FIG. 4-4C) represents all possible intensity/phase pairs coming out of a certain multi-mode

transformation. FIG. 4-5 may depict, for example, the effect of propagating an optical signal through photonic processor 1-103 of FIG. 1-1.

Suppose that S_{out} is the symbol coming out of the transformation at a certain time. As described above, symbol S_{out} is characterized by an intensity $\alpha\beta$ and a phase $\theta+\Delta\theta$. When the transformed field is mixed with the reference signal, points along line 4-403 are projected on a reference axis 4-500. Thus, for example, symbol S_{out} is projected onto point A on the axis 4-500. Point A represents the electric output signal arising out of the mixing, and is characterized by an amplitude equal to the distance between 0 and A ($\overline{OA}$). It should be noted that the angle of the axis 4-500 relative to the real axis (i.e., angle φ), depends upon the phase of the reference signal relative to the transformed, modulated optical field. In this example, $0 < \varphi < \pi/4$.

Referring back to FIG. 4-2B, at act 4-160, a decoder may obtain a value representative of a decoded real number based on the electric output signal obtained at act 4-158. In some embodiments, when decoding symbols obtained through an optical transformation, the exact location of the symbol in the complex plane may be unimportant. What matters for accurate decoding, according to some embodiments, is the projection of the optical signal on a known reference axis. In other words, symbols along the line 4-403 can be mapped to points along the reference axis, much like symbol S_{out} is mapped to point A. Thus, some embodiments implement optical decoding schemes in which points along a reference axis correspond to specific symbols in the complex plane.

The manner in which points along the reference axis are mapped to symbols in the complex plane may be determined using a calibration procedure. During the calibration procedure, a set of input symbols of known intensity and phase (the symbols representative of a set of real numbers), are passed through a certain optical transformation unit, and the resulting symbols are coherently detected using a reference signal having a known phase. The amplitude of the resulting electric output signal (i.e., the amplitude of the projection along the reference axis) is recorded and stored in a table (e.g., a look-up table). The table may subsequently be used during operation to decode real numbers based on the amplitude of the projection along the reference axis.

An example of such a table is illustrated in FIG. 4-6, in accordance with some embodiments. This table includes a column for the phase of the reference signal (“reference phase φ ”), a column for the projection along the reference axis (“projection $\overline{OA}$,” i.e., the amplitude of the electrical output signal), and a column for the decoded real number

(“decoded real number”). The table may be populated during a calibration procedure involving 1) passing a symbol of known intensity and phase through a known optical transformation unit, 2) setting the phase of a reference signal to a known value, 3) mixing the transformed symbol with the reference signal thereby projecting the transformed symbol onto the known reference axis, 4) determining the amplitude of the projection along the reference axis, and 5) recording the expected output real number. During operation, a user may infer the decoded real number using the table based on the value of the reference phase and the measured projection.

In the example of FIG. 4-6, only two reference phases are considered by way of example and not limitation—0 and $\pi/6$. At $\phi=0$, projections -1, -0.8 and 1 are mapped to real numbers -9.6, 0.2 and 8.7, respectively. At $\phi=\pi/6$, projections -0.4, 0.6 and 0.9 are mapped to real numbers 3.1, -5 and 10, respectively. Thus, for example, if a user obtains a projection of -0.4 when using a reference phase of $\pi/6$, the user can infer that the decoded real number is 3.1.

FIG. 4-7 illustrates a specific example of a photonic system 4-700 in which the optical transformation unit includes an optical fiber 4-708, in accordance with some embodiments. This example illustrates how encoding techniques of the types described herein may be used in the context of optical communications. In this implementation, optical fiber 4-708 separates transmitter 4-701 from receiver 4-702.

Transmitter 4-701 includes optical source 4-102, encoder 4-104 and optical modulator 4-106 (described above in connection with FIG. 4-2A) and receiver 4-702 includes coherent receiver 4-110 and decoder 4-114 (also described above in connection with FIG. 4-2A). As described above, in some embodiments, optical modulator 4-106 may be driven using a single modulating signal 4-105. In some such embodiments, optical modulator 4-106 may implement an on-off keying (OOK) modulation scheme.

It should be noted that, in this implementation, the reference signal used for the coherent detection is provided directly from transmitter 4-701, rather than being generated locally at receiver 4-702 (though other implementations may involve local oscillators at receiver 4-702). To illustrate this concept, consider for example the plot of FIG. 4-8. FIG. 4-8 illustrates an example of a power spectral density of an optical field at the output of transmitter 4-701. As shown, the power spectral density includes a signal 4-800 as well as a carrier 4-801. Signal 4-800 represents the information encoded using modulator 4-106. Carrier 4-801 represents a tone at the wavelength of optical source 4-102 (λ_0). Upon

transmission of this optical field through optical fiber 4-708, coherent receiver 4-110 may perform the mixing using carrier 4-801 itself as the reference signal.

Some embodiments relate to methods for fabricating photonic systems of the types described herein. One such method is depicted in FIG. 4-9, in accordance with some embodiments. Method 4-900 begins at act 4-902, in which a modulator is fabricated. The modulator may be fabricated to be driven with a single electrical modulating signal. An example of a modulator is modulator 4-106 (FIG. 4-2A). At act 4-904, a coherent receiver is fabricated. In some embodiments, the coherent receiver is fabricated on the same semiconductor substrate on which the modulator is fabricated. At act 4-906, an optical transformation unit is fabricated. The optical transformation unit may be fabricated on the same substrate on which the coherent receiver is fabricated and/or on the same semiconductor substrate on which the modulator is fabricated. The optical transformation unit may be fabricated to be coupled between the modulator and the coherent receiver. An example of an optical transformation unit is optical transformation unit 4-108 (FIG. 4-2A).

V. Differential Receiver

The inventors have recognized and appreciated that some conventional optical receivers are particularly susceptible to noise generated from voltage supplies, to noise arising from the fact that photodetectors inevitably produce dark currents, and to other forms of noise. The presence of noise reduces the signal-to-noise ratio, and therefore, the ability of these photodetectors to accurately sense incoming optical signals. This can negatively affect the performance of the system in which these photodetectors are deployed. For example, this can negatively affect the system's bit error rate and power budget.

The inventors have developed optical receivers with reduced susceptibility to noise. Some embodiments of the present application are directed to optical receivers in which both the optical-to-electric conversion and the amplification are performed in a differential fashion. In the optical receivers described herein, two separate signal subtractions take place. First, the photocurrents are subtracted from one another to produce a pair of differential currents. Then, the resulting differential currents are further subtracted from each other to produce an amplified differential output. The inventors have recognized and appreciated that having an optical receiver involving multiple levels of signal subtraction results in multiple levels of noise cancelation, thus substantially reducing noise from the system. This can have

several advantages over conventional optical receivers, including wider dynamic range, greater signal-to-noise ratio, larger output swing, and increased supply-noise immunity.

Optical receivers of the types described herein can be used in a variety of settings, including for example in telecom and datacom (including local area networks, metropolitan area networks, wide area networks, data center networks, satellite networks, etc.), analog applications such as radio-over-fiber, all-optical switching, Lidar, phased arrays, coherent imaging, machine learning and other types of artificial intelligence applications, as well as other applications. In some embodiments, optical receivers of the types described herein may be used as part of a photonic processing system. For example, in some embodiments, optical receivers of the types described herein may be used to implement optical receiver 1-105 of FIG. 1-1 (e.g., one or more homodyne receivers 1-901 of FIG. 1-9).

FIG. 5-1 illustrates a non-limiting example of an optical receiver 5-100, in accordance with some non-limiting embodiments of the present application. As shown, optical receiver 5-100 includes photodetectors 5-102, 5-104, 5-106 and 5-108, though other implementations include more than four photodetectors. Photodetector 5-102 may be connected to photodetector 5-104, and photodetector 5-106 may be connected to photodetector 5-108. In some embodiments, the anode of photodetector 5-102 is connected to the cathode of photodetector 5-104 (at node 5-103), and the cathode of photodetector 5-106 is connected to the anode of photodetector 5-108 (at node 5-105). In the example of FIG. 5-1, the cathodes of photodetectors 5-102 and 5-108 are connected to voltage supply V_{DD} and the anodes of photodetectors 5-104 and 5-106 are connected to the reference potential (e.g., to ground). The opposite arrangement is also possible in some embodiments. The reference potential may be at a potential equal to zero or having any suitable value, such as $-V_{DD}$. V_{DD} may have any suitable value.

Photodetectors 5-102 through 5-108 may be implemented in any of numerous ways, including for example with pn-junction photodiodes, pin-junction photodiodes, avalanche photodiodes, phototransistors, photoresistors, etc. The photodetectors may include a material capable of absorbing light at the wavelength of interest. For example, at wavelengths in the O-band, C-band or L-band, the photodetectors may have an absorption region made at least in part of germanium, by way of a non-limiting example. For visible light, the photodetectors may have an absorption region made at least in part of silicon, by way of another non-limiting example.

Photodetectors 5-102 through 5-108 may be integrated components formed monolithically as part of the same substrate. The substrate may be a silicon substrate in some

embodiments, such as a bulk silicon substrate or a silicon-on-insulator. Other types of substrates can also be used, including for example indium phosphide or any suitable semiconductor material. To reduce variability in the characteristics of the photodetectors due to fabrication tolerances, in some embodiments, the photodetectors may be positioned in close proximity to one another. For example, the photodetectors may be positioned on a substrate within an area of 1mm^2 or less, 0.1mm^2 or less or 0.01mm^2 or less.

As further illustrated in FIG. 5-1, photodetectors 5-102 through 5-108 are connected to a differential operational amplifier 5-110. For example, photodetectors 5-102 and 5-104 may be connected to the non-inverting input (“+”) of DOA 5-110 and photodetectors 5-106 and 5-108 may be connected to the inverting input (“-”) of DOA 5-110. DOA 5-110 has a pair of outputs. One output is inverting and one output is non-inverting.

In some embodiments, as will be described in detail in connection with FIG. 5-2, photodetectors 5-102 and 5-106 may be arranged to receive the same optical signal “t” and photodetectors 5-104 and 5-108 may be arranged to receive the same optical signal “b.” In some embodiments, photodetectors 5-102 through 5-108 may be designed to be substantially equal to each other. For example, photodetectors 5-102 through 5-108 may be formed using the same process steps and using the same photomask patterns. In these embodiments, photodetectors 5-102 through 5-108 may exhibit substantially the same characteristics, such as substantially the same responsivity (the ratio between the photocurrent and the received optical power) and/or substantially the same dark current (the current generated when no optical power is received). In these embodiments, the photocurrents generated by photodetectors 5-102 and 5-106 responsive to reception of signal t may be substantially equal to each other. Such photocurrents are identified as “ i_t ” in FIG. 5-1. It should be noted that, due to the orientations of photodetectors 5-102 and 5-106, the photocurrents generated by photodetectors 5-102 and 5-106 are oriented in opposite directions. That is, the photocurrent of photodetector 5-102 is directed towards node 5-103 and the photocurrent of photodetector 5-106 is oriented away from node 5-105. Furthermore, the photocurrents generated by photodetectors 5-104 and 5-108 responsive to reception of signal b may be substantially equal to each other. Such photocurrents are identified as “ i_b .” Due to the orientations of photodetectors 5-104 and 5-108 relative to each other, the photocurrents generated by photodetectors 5-104 and 5-108 are oriented in opposite directions. That is, the photocurrent of photodetector 5-108 is directed towards node 5-105 and the photocurrent of photodetector 5-104 is oriented away from node 5-103.

In view of the orientations of the photodetectors, a current with amplitude $i_t - i_b$ emerges from node 5-103 and a current with amplitude $i_b - i_t$ emerges from node 5-105. Thus, the currents have substantially the same amplitudes, but with opposite signs. Photodetectors 5-102 through 5-108 may produce dark currents. Dark currents are typically due to leakage and arise from a photodetector regardless of whether the photodetector is exposed to light or not. Because dark currents arise even in the absence of incoming optical signals, dark currents effectively contribute to noise in the optical receiver. The inventors have appreciated that the negative effects of these dark currents can be significantly attenuated thanks to the current subtraction described above. Thus, in the example of FIG. 5-1, the dark current of photodetector 5-102 and the dark current of photodiode 5-104 substantially cancel out one another (or at least are mutually substantially reduced), and so do the dark currents of photodetector 5-106 and 5-108. Consequently, noise due to the presence of the dark currents is greatly attenuated.

FIG. 5-2 illustrates a photonic circuit 5-200 arranged for providing two optical signals to photodetectors 5-102 through 5-108, in accordance with some non-limiting embodiments. Photonic circuit 5-200 may comprises optical waveguides for routing the optical signals to the photodetectors. The optical waveguides may be made of a material that is transparent or at least partially transparent to light at the wavelength of interest. For example, the optical waveguides be made of silicon, silicon oxide, silicon nitride, indium phosphide, gallium arsenide, or any other suitable material. In the example of FIG. 5-2, photonic circuit 5-200 includes input optical waveguides 5-202 and 204 and couplers 5-212, 5-214 and 5-216. As further illustrated, the output optical waveguides of photonic circuit 5-200 are coupled to photodetectors 5-102 through 5-108.

In the example of FIG. 5-2, couplers 5-212, 5-214 and 5-216 comprise directional couplers, where evanescent coupling enables transfer of optical power between adjacent waveguides. However, other types of couplers may be used such as Y-junctions, X-junctions, optical crossovers, counter-direction couplers, etc. In other embodiments, photonic circuit 5-200 may be implemented with a multi-mode interferometer (MMI). Couplers 5-212, 5-214 and 5-216 may be 3dB couplers (with a 50%-50% coupling ratio) in some embodiments, though other ratios are also possible, such as 51%-49%, 55%-45% or 60%-40%. It should be appreciated that, due to fabrication tolerances, the actual coupling ratio may deviate slightly from the intended coupling ratio.

Signal s_1 may be provided at input optical waveguide 5-202 and signal s_2 may be provided at input optical waveguide 204. Signals s_1 and s_2 may be provided to the respective

input optical waveguides using for example optical fibers. In some embodiments, s_1 represents a reference local oscillator signal, such as the signal generated by a reference laser, and s_2 represents the signal to be detected. As such, the optical receiver may be viewed as a homodyne optical receiver. In some such embodiments, s_1 may be a continuous wave (CW) optical signal while s_2 may be modulated. In other embodiments, both signals are modulated or both signals are CW optical signals, as the application is not limited to any particular type of signal.

In the example of FIG. 5-2, signal s_1 has amplitude A_{LO} and phase ϑ , and signal s_2 has amplitude A_S and phase φ . Coupler 5-212 combines signals s_1 and s_2 such that signals t and b emerge at respective outputs of coupler 5-212. In the embodiments in which coupler 5-212 is a 3dB coupler, t and b may be given by the following expression:

$$\begin{pmatrix} t \\ b \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & i & 1 \end{pmatrix} \begin{pmatrix} A_{LO} e^{i\vartheta} \\ A_S e^{i\varphi} \end{pmatrix}$$

and the powers T and B (of t and b , respectively) may be given by the following expressions:

$$\begin{aligned} T &= [A_{LO}^2 + A_S^2 + 2A_{LO}A_S \sin(\vartheta - \varphi)] \\ B &= [A_{LO}^2 + A_S^2 - 2A_{LO}A_S \sin(\vartheta - \varphi)] \end{aligned}$$

Thus, in the embodiments in which couplers 5-214 and 5-216 are 3dB couplers, photodetectors 5-102 and 5-106 may each receive a power given by $T/2$ and photodetectors 5-104 and 5-108 may each receive a power given by $B/2$.

Referring back to FIG. 5-1, and assuming that the responsivities of photodetectors 5-102 through 5-108 are all equal to each other (though not all embodiments are limited in this respect), the currents emerging from node 5-103 and 5-105, respectively, may be given by the following expressions:

$$\begin{aligned} i_t - i_b &= 2A_{LO}A_S \sin(\vartheta - \varphi) \\ i_b - i_t &= -2A_{LO}A_S \sin(\vartheta - \varphi) \end{aligned}$$

DOA 5-110 is arranged to amplify the differential signal received at the “+” and “-” inputs, and to produce an amplified differential output, represented in FIG. 5-1 by voltages $V_{out,n}$ and $V_{out,p}$. In some embodiments, DOA 5-110, in combination with impedances z , may be viewed as a differential transimpedance amplifier, in that it produces a differential pair of voltage ($V_{out,n}$, $V_{out,p}$) based on a differential pair of current ($i_b - i_t$, $i_t - i_b$). In some embodiments,

each of $V_{out,n}$, $V_{out,p}$ may be proportional to the difference between current $i_t - i_b$ and current $i_b - i_t$, thus giving rise to the following expressions:

$$\begin{aligned} V_{out,p} &= 2z(i_t - i_b) \\ V_{out,n} &= 2z(i_b - i_t) \end{aligned}$$

This differential pair of voltages may be provided as input to any suitable electronic circuit, including but not limited to an analog-to-digital converter (not shown in FIG. 5-1). It should be noted that optical receiver 5-100 provides two levels of noise rejection. The first level of noise rejection occurs thanks to the subtraction of the photocurrents, the second level of noise rejection occurs thanks to the subtraction taking place in the differential amplification stage. This results in a significant increase in noise rejection.

In the example of FIG. 5-1, impedances z are shown as being equal to each other, however different impedances may be used in other embodiments. These impedances may include passive electric components, such as resistors, capacitors and inductors, and/or active electronic components, such as diode and transistors. The components constituting these impedances may be chosen to provide a desired gain and bandwidth, among other possible characteristics.

As discussed above, optical receiver 5-100 may be integrated monolithically on a substrate. One such substrate is illustrated in FIG. 5-3A, in accordance with some non-limiting embodiments. In this example, photodetectors 5-102 through 5-108, photonic circuit 5-200 and DOA 5-110 are monolithically integrated as part of substrate 5-301. In other embodiments, photodetectors 5-102 through 5-108 and photonic circuit 5-200 may be integrated on substrate 5-301 and DOA 5-110 may be integrated on a separate substrate 5-302. In the example of FIG. 5-3B, substrates 5-301 and 5-302 are flip-chip bonded to one another. In the example of FIG. 5-3C, substrates 5-301 and 5-302 are wire bonded to one another. In yet another example (not illustrated), photodetectors 5-102 through 5-108 and photonic circuit 5-200 may be fabricated on separate substrates.

Some embodiments of the present application are directed to methods for fabricating optical receivers. One such method is depicted in FIG. 5-4, in accordance with some non-limiting embodiments. Method 5-400 begins at act 5-402, in which a plurality of photodetectors are fabricated on a first substrate.

Once fabricated, the photodetectors may be connected together, for example in the arrangement shown in FIG. 5-1. In some embodiments, the photodetectors may be positioned on the first substrate within an area of 1mm^2 or less, 0.1mm^2 less or 0.01mm^2 or less. At act 5-404, a photonic circuit is fabricated on the first substrate. The photonic circuit may be

arranged to provide a pair of optical signals to the photodetectors, for example in the manner shown in FIG. 5-2. At act 5-406, a differential operational amplifier may be fabricated on a second substrate. An example of a differential operational amplifier is DOA 5-110 of FIG. 5-1. At act 5-408, the first substrate may be bonded to the second substrate, for example via flip-chip bonding (as shown in FIG. 5-3A), wire bonding (as shown in FIG. 5-3B), or using any other suitable bonding technique. Once the substrates are bonded, the photodetectors of the first substrate may be electrically connected to the differential operational amplifier of the second substrate, for example in the manner shown in FIG. 5-1.

Examples of fabrication processes are depicted schematically at FIGs. 5-4A through 5-4F, in accordance with some embodiments. FIG. 5-4A depicts a substrate 5-301 having a lower cladding 5-412 (e.g., an oxide layer such as a buried oxide layer or other types of dielectric materials) and a semiconductor layer 5-413 (e.g., a silicon layer or a silicon nitride layer, or other types of material layers). At FIG. 5-4B, semiconductor layer 5-413 is patterned, for example using a photolithographic exposure, to form regions 5-414. Regions 5-414 may be arranged to form optical waveguides in some embodiments. In some embodiments, the resulting pattern resembles photonic circuit 5-200 (FIG. 5-2), where waveguides 5-202 and 204, and couplers 5-212, 5-214 and 5-216 are embedded into one or more regions 5-414. At FIG. 5-4C, photodetectors 5-102, 5-104, 5-106 and 5-108 (and optionally, other photodetectors) are formed. In this example, an optical absorbing material 5-416 is deposited to be adjacent a region 5-414. The optical absorbing material 5-416 may be patterned to form the photodetectors. The material used for the optical absorbing material may depend on the wavelength to be detected. For example, germanium may be used for wavelengths in the L-Band, C-Band or O-Band. Silicon may be used for visible wavelengths. Of course, other materials are also possible. The optical absorbing material 5-416 may be positioned to be optically coupled to regions 5-414 in any suitable way, including but not limited to butt coupling, taper coupling and evanescent coupling.

At FIG. 5-4D, DOA 5-110 is formed. In some embodiments, DOA 5-110 includes several transistors formed via ion implantation. FIG. 5-4D depicts implanted regions 5-418, which may form part of one or more transistors of DOA 5-110. While only one ion implantation is illustrated in FIG. 5-4D, in some embodiments, formation of DOA 5-110 may involve more than one ion implantations. Additionally, DOA 5-110 may be electrically connected to the photodetectors, for example via one or more conductive traces formed on substrate 5-301.

The arrangement of FIG. 5-4D is such that photonic circuit 5-200, photodetectors 5-102 through 5-108 and DOA 5-110 are formed on a common substrate (as shown in FIG. 5-3A). Arrangements in which DOA 5-110 is formed on a separate substrate (as shown in FIG. 5-3B or FIG. 5-3C) are also possible. In one such example, DOA 5-110 is formed on a separate substrate 5-302, as shown in FIG. 5-4E, where implanted regions 5-428 are formed via one or more ion implantations.

Subsequently, substrate 5-301 is bonded to substrate 5-302, and photodetectors 5-102 through 5-108 are connected to DOA 5-110. At FIG. 5-4F, conductive pads 5-431 are formed and placed in electrical communication with optical absorbing material 5-416, and conductive pads 5-432 are formed and placed in electrical communication with implanted regions 5-428. The conductive pads are electrically connected via wire bonding (as shown in FIG. 5-4F) or via flip-chip bonding.

Some embodiments are directed to methods for receiving input optical signals. Some such embodiments may involve homodyne detection, though the application is not limited in this respect. Other embodiments may involve heterodyne detection. Yet other embodiments may involve direct detection. In some embodiments, reception of optical signals may involve optical receiver 5-100 (FIG. 5-1), though other types of receivers may be used.

An example of a method for receiving an input optical signal is depicted in FIG. 5-5, in accordance with some embodiments. Method 5-500 begins at act 5-502, in which the input signal is combined with a reference signal to obtain first and second optical signals. The input signal may be encoded with data, for example in the form of amplitude modulation, pulse width modulation, phase or frequency modulation, among other types of modulation. In some of the embodiments involving homodyne detection, the reference signal may be a signal generated by a local oscillator (e.g., a laser). In other embodiments, the reference signal may also be encoded with data. In some embodiments, the input signal and the reference signal are combined using a photonic circuit 5-200 (FIG. 5-2), though other types of optical combiners may be used, including but not limited to MMIs, Y-junctions, X-junctions, optical crossovers, and counter-direction couplers. In the embodiments in which photonic circuit 5-200 is used, t and b may represent the signals obtained from the combination of the input signal with the reference signal.

At act 5-504, the first optical signal is detected with a first photodetector and with a second photodetector and the second optical signal is detected with a third photodetector and with a fourth photodetector to produce a pair of differential currents. In some embodiments, act 5-504 may be performed using optical receiver 5-100 (FIG. 5-1). In some such

embodiments, the first optical signal is detected with photodetectors 5-102 and 5-106, and the second optical signal is detected with photodetectors 5-104 and 5-108. The produced pair of differential currents is represented, collectively, by currents i_b-i_t and i_t-i_b . Being differential, in some embodiments, the currents of the pair may have substantially equal amplitudes, but with substantially opposite phases (e.g., with a π -phase difference).

At act 5-506, a differential operational amplifier (e.g., DOA 5-110 of FIG. 5-1) produces a pair of amplified differential voltages using the pair of differential currents produced at act 5-504. In the embodiments that use DOA 5-110, the produced pair of differential voltages is represented by voltages $V_{out,n}$ and $V_{out,p}$. Being differential, in some embodiments, the voltages of the pair may have substantially equal amplitudes, but with substantially opposite phases (e.g., with a π -phase difference).

Method 5-500 may have one or more advantages over conventional methods for receiving optical signals, including for example wider dynamic range, greater signal-to-noise ratio, larger output swing, and increased supply-noise immunity.

VI. Phase Modulator

The inventors have recognized and appreciated that certain optical phase modulators suffer from high dynamic loss and low modulation speed, which significantly limit the range of applications in which these phase modulators can be deployed. More specifically, some phase modulators involve significant trade-offs between modulation speed and dynamic loss, such that an increase in modulation speed results in an increase in dynamic loss. As used herein, the phrase “dynamic loss” refers to optical power loss experienced by an optical signal that depends on the degree to which its phase is modulated. Ideal phase modulators are such that power loss is independent of the phase modulation. Real-world phase modulators, however, experience a certain power loss when no modulation occurs, and experience a different power loss when modulation occurs. For example, the power loss experienced at no phase modulation may be L_1 , the power loss experienced at a $\pi/2$ -phase modulation may be L_2 , and the power loss experienced at a π -phase modulation may be L_3 , with L_1 , L_2 and L_3 being different from each other. This behavior is undesirable because, in addition to phase modulation, the signal further experiences amplitude modulation.

Some such phase modulators, in addition, require several hundreds of microns in length to provide sufficiently large phase shifts. Unfortunately, being so long, such phase modulators are not suitable for use in applications requiring integration of several phase

shifters on a single chip. The phase modulators alone may take up most of the space available on the chip, thus limiting the number of devices that can be co-integrated on the same chip.

Recognizing the aforementioned limitations of certain phase modulators, the inventors have developed small footprint-optical phase modulators capable of providing high modulation speeds (e.g., in excess of 6-100MHz or 1GHz) while limiting dynamic loss. In some embodiments, a phase modulator may occupy an area as small as $300\mu\text{m}^2$. Thus, as an example, a reticle having an area of 1cm^2 can accommodate as many as 15,000 phase modulators while saving an additional 50mm^2 for other devices.

Some embodiments relate to Nano-Opto-Electromechanical Systems (NOEMS) phase modulators having multiple suspended optical waveguides positioned adjacent to one another and forming a plurality of slots therebetween. The dimensions of the slots are sufficiently small to form slot waveguides, whereby a substantial portion (e.g., a majority) of the mode energy is confined in the slots themselves. These modes are referred to herein as slot modes. Having a substantial portion of the mode energy in the slots enables modulation of the effective index of the mode, and a result, of the phase of an optical signal embodying the mode, by causing variations in the dimensions of the slots. In some embodiments, phase modulation can be achieved by applying mechanical forces that cause variations in the dimensions of the slots.

The inventors have recognized and appreciated that the modulation speed achievable with the NOEMS phase modulators described herein can be increased, without significant increases in dynamic loss, by decoupling the mechanical drivers from the region where optical modulation takes place. In phase modulators in which the mechanical drivers are decoupled from the optical modulation region, electric driving signals are applied on the mechanical drivers, rather than being applied on the optical modulation region itself. This arrangement removes the need to make the optical modulation region electrically conductive, thus enabling a reduction in the doping of this region. The low doping results in a reduction of free carriers which may otherwise lead to optical absorption, thus lowering dynamic loss.

Furthermore, decoupling the mechanical drivers from the optical modulation region enables a greater modulation per unit length, and as a result a shorter modulation region. Shorter modulation regions, in turn, enable, greater modulation speed.

The inventors have further recognized and appreciated that including multiple slots in the modulation region can enable a further reduction in the length of the phase modulator (thereby decreasing its size). Having more than one slot, in fact, enables a substantial reduction in the length of the transition region through which light is coupled to the

modulation region. The result is a substantially more compact form factor. Thus, NOEMS phase modulators of the types described herein can have shorter modulation regions and/or shorter transition regions. Phase modulators of the types described herein can have lengths as low as 20 μ m or 30 μ m, in some embodiments.

As will be described in detail further below, some embodiments relate to phase modulators in which a trench is formed in the chip, and is arranged so that the modulating waveguides are suspended in air and are free move in space.

The inventors have recognized a potential drawback associated with the use of trenches that results from the formation of cladding/air interfaces. When a propagating optical signal enters (or exits) a trench, it encounters a cladding/air interface (or an air/cladding interface). Unfortunately, the presence of the interface can give rise to optical reflections, which in turn can increase insertion losses. The inventors have appreciated that the negative effect of such interfaces can be mitigated by reducing the physical extension of the optical mode in the region where it passes through the interface. This can be achieved in various ways. For example, in some embodiments, the extension of the optical mode may be reduced by tightly confining the mode within a rib waveguide. A rib waveguide may be sized so that only a small fraction of the mode energy (e.g., less than 20%, less than 10%, or less than 5%) is outside the edges of the waveguide.

NOEMS phase modulators of the types described herein may be used in a variety of applications, including for example in telecom and datacom (including local area networks, metropolitan area networks, wide area networks, data center networks, satellite networks, etc.), analog applications such as radio-over-fiber, all-optical switching, coherent Lidar, phased arrays, coherent imaging, machine learning and other types of artificial intelligence applications. Additionally, the NOEMS modulators may be used as part of amplitude modulators, for example if combined with a Mach Zehnder modulator. For example, a Mach Zehnder modulator may be provided in which a NOEMS phase modulator is positioned in one or more of the arms of the Mach Zehnder modulator. Several modulation schemes may be enabled using NOEMS phase modulators, including for example amplitude shift keying (ASK), quadrature amplitude modulation (QAM), phase shift keying (BPSK), quadrature phase shift keying (QPSK) and higher order QPSK, offset quadrature phase-shift keying (OQPSK), Dual-polarization quadrature phase shift keying (DPQPSK), amplitude phase shift keying (APSK), etc. Additionally, NOEMS phase modulators may be used as phase correctors in applications in which the phase of an optical signal tends to drift unpredictably. In some embodiments, NOEMS phase modulators of the types described herein may be used

as part of a photonic processing system. For example, in some embodiments, NOEMS phase modulators of the types described herein may be used to implement phase modulators 1-207 of FIG. 1-1, and/or to implement part of the variable beam splitters 1-401 of FIG. 104, and/or to implement the phase shifters 1-505, 1-507 and 1-509 of FIG. 1-5, and/or to implement phase modulator 1-601 of FIG. 1-6, and/or to implement part of the amplitude modulators 1-603 of FIG. 1-6, and/or to implement part of the amplitude modulators 1-205 of FIG. 1-2.

FIG. 6-1A is a top view illustrating schematically a Nano-Opto-Electromechanical Systems (NOEMS) phase modulator, in accordance with some non-limiting embodiments. NOEMS phase modulator 6-100 includes input waveguide 6-102, output waveguide 6-104, input transition region 6-140, output transition region 6-150, suspended multi-slot optical structure 6-120, mechanical structures 6-130 and 6-132, and mechanical drivers 6-160 and 6-162. NOEMS phase modulator 6-100 may be fabricated using silicon photonic techniques. For example, NOEMS phase modulator 6-100 may be fabricated on a silicon substrate, such as a bulk silicon substrate or a silicon-on-insulator (SOI) substrate. In some embodiments, NOEMS phase modulator 6-100 may further include electronic circuitry configured to control the operations of mechanical drivers 6-160 and 6-162. The electronic circuitry may be fabricated on the same substrate hosting the components of FIG. 6-1A, or on a separate substrate. When disposed on a separate substrate, the substrates may be bonded to one another in a any suitable way, including 3D-bonding, flip-chip bonding, wire bonding etc.

At least part of NOEMS phase modulator 6-100 is formed in a trench 6-106. As will be described in detail further below, trenches of the types described herein may be formed by etching a portion of the cladding. In the example of FIG. 6-1A, trench 6-106 has a rectangular shape, though trenches of any other suitable shape may be used. In this example, trench 6-106 has four sidewalls. Sidewalls 6-112 and 6-114 are spaced from one another along the z-axis (referred to herein as the propagation axis), and the other two sidewalls (not labeled in FIG. 6-1A) are spaced from one another along the x-axis.

In some embodiments, the separation along the z-axis between sidewalls 6-112 and 6-114 may be less than or equal to 50 μ m, less than or equal to 30 μ m, or less than or equal to 20 μ m. Thus, the modulation region of this NOEMS phase modulator is significantly shorter than other types of phase modulators, which require several hundreds of microns for modulating the phase of an optical signal. The relatively short length is enable by one or more of the following factors. First, having multiple slots improves coupling to the optical modulation region, which in turn enables a reduction in the length of the transition region. The improved coupling may be the result of enhanced mode symmetry in the multi-slot

structure. Second, decoupling the mechanical drivers from the optical modulation region enables a greater modulation per unit length, and as a result a shorter modulation region.

During operation, an optical signal may be provided to input waveguide 6-102. In one example, the optical signal may be a continuous wave (CW) signal. Phase modulation may take place in suspended multi-slot optical structure 6-120. A phase modulated optical signal may exit NOEMS phase modulator 6-100 from output waveguide 6-104. Transition region 6-140 may ensure loss-free or nearly loss-free optical coupling between input waveguide 6-102 and suspended multi-slot optical structure 6-120. Similarly, transition region 6-150 may ensure loss-free or nearly loss-free optical coupling between suspended multi-slot optical structure 6-120 and output waveguide 6-104. Transitions regions 6-140 and 6-150 may include tapered waveguides in some embodiments, as described in detail further below. As discussed above, the length of the transitions regions may be shorter relative to other implementations.

The input optical signal may have any suitable wavelength, including but not limited to a wavelength in the O-band, E-band, S-band, C-band or L-band. Alternatively, the wavelength may be in the 850nm-band or in the visible band. It should be appreciated that NOEMS phase modulator 6-100 may be made of any suitable material, so long as the material is transparent or at least partially transparent at the wavelength of interest, and the refractive index of the core region is greater than the refractive index of the surrounding cladding. In some embodiments, NOEMS phase modulator 6-100 may be made of silicon. For example, input waveguide 6-102, output waveguide 6-104, input transition region 6-140, output transition region 6-150, suspended multi-slot optical structure 6-120, and mechanical structures 6-130 and 6-132 may be made of silicon. Given silicon's relatively low optical bandgap (approximately 1.12eV), silicon may be particularly suitable for use in connection with near infrared wavelengths. In another example, NOEMS phase modulator 6-100 may be made of silicon nitride or diamond. Given silicon nitride's and diamond's relatively high optical bandgaps (approximately 5eV and approximately 5.47eV, respectively), these materials may be particularly suitable for use in connection with visible wavelengths. However, other materials are also possible, including indium phosphide, gallium arsenide, and or any suitable III-V or II-VI alloy.

In some embodiments, input waveguide 6-102 and output waveguide 6-104 may be sized to support a single mode at the wavelength of operation (though multi-mode waveguides can also be used). For example, if a NOEMS phase modulator is designed to operate at 1550nm (though of course, not all embodiments are limited in this respect), input

and output waveguides 6-102 and 6-104 may support a single mode at 1550nm. In this way, the mode confinement within the waveguide may be enhanced, thus reducing optical losses due to scattering and reflections. Waveguides 6-102 and 6-104 may be rib waveguides (e.g., with rectangular cross sections) or may have any other suitable shape.

As described above, part of NOEMS phase modulator 6-100 may be formed within a trench 6-106, so that the waveguides in the modulation region are surrounded by air and are free to move in space. The drawback of including a trench is the formation of a cladding/air interface and an air/cladding interface along the propagation path. Thus, the input optical signal passes a cladding/air interface (in correspondence with sidewall 6-112) before reaching the region where modulation occurs and passes an air/cladding interface (in correspondence with sidewall 6-114) after the modulation region. These interfaces may introduce reflection losses. In some embodiments, reflection losses may be reduced by positioning transition region 6-140 inside, rather than outside, trench 6-106 (as shown in FIG. 6-1A). In this way, the mode expansion associated with the transition region takes place where the optical signal has already passed the cladding/air interface. In other words, the mode is tightly confined as it passes the cladding/air interface, but is expanded in the trench, using the transition region, for purposes of coupling to the suspended multi-slot structure 6-120. Similarly, transition region 6-150 may be formed inside trench 6-106, thereby spatially re-confining the mode before it reaches sidewall 6-114.

FIG. 6-1B illustrates suspended multi-slot optical structure 6-120 in additional detail, in accordance with some non-limiting embodiments. In the example of FIG. 6-1B, multi-slot optical structure 6-120 includes three waveguides (6-121, 6-122 and 6-123). Slot 6-124 separates waveguide 6-121 from waveguide 6-122 and slot 6-125 separates waveguide 6-122 from waveguide 6-123. The width of the slots (d_1 and d_2) may be less than the critical width (at the wavelength of operation) for forming slot modes, whereby a substantial portion of the mode energy (e.g., more than 40%, more than 50%, more than 60%, or more than 75%) is within the slots. For example, each of d_1 and d_2 may be equal to or less than 200nm, equal to or less than 150nm, or equal to or less than 100nm. The minimum width may be set by the photolithographic resolution.

FIG. 6-1C is a plot illustrating an example of an optical mode supported by the waveguides 6-121, 6-122 and 6-123, in accordance with some non-limiting embodiments. More specifically, the plot illustrates the amplitude of a mode (e.g., the electric field E_x , E_y or E_z , or magnetic field H_x , H_y or H_z). As illustrated, a substantial portion of the overall energy is confined within the slots, where the mode exhibits peaks of amplitude. In some

embodiments, there is more optical energy in one of the slots than in any one of the individual waveguides. In some embodiments, there is more optical energy in one of the slots than in all the waveguides considered together. Outside the outer walls of the exterior waveguides, the mode energy decays (for example exponentially).

Widths d_1 and d_2 may be equal to, or different than, one another. The widths of the slots and the waveguides may be constant along the z -axis (as in FIG. 6-1B) or may vary. In some embodiments, the widths of waveguides 6-121, 6-122 and 6-123 may be less than the width of input waveguide 6-102. In some embodiments, when the wavelength of operation is in the C-band, the widths of waveguides 6-121, 6-122 and 6-123 may be between 200nm and 400nm, between 250nm and 350nm, or within any other suitable range, whether within or outside such ranges.

While the example of FIG. 6-1B illustrates suspended a multi-slot optical structure 6-120 with three waveguides and two slots, any other suitable number of waveguides and slots may be used. In other examples, a suspended multi-slot optical structure 6-120 may include five waveguides and four slots, seven waveguides and six slots, nine waveguides and eight slots, etc. In some embodiments, the structure includes an odd number of waveguides (and consequently, an even number of slots) so that only symmetric modes are excited, while antisymmetric modes remain unexcited. The inventors have appreciated that enhancing the symmetry of the mode enhances coupling into the slotted structure, thus enabling a substantial reduction in the length of the transition region. However, implementations with even number of waveguides are also possible.

As will be described in detail further below, phase modulation occurs by causing the exterior waveguides (6-121 and 6-123 in FIG. 6-1B) to move relative to the center waveguide (6-122 in FIG. 6-1B) along the x -axis. When waveguide 6-121 moves in the x -axis relative to waveguide 6-122, the width of slot 6-124 varies, and the shape of the mode supported by the structure varies accordingly. The result is a change in the effective index of the mode supported by the structure, and consequently, a phase modulation takes place. Motion of the exterior waveguides may be induced using mechanical structures 6-130 and 6-132.

An example of a mechanical structure 6-130 is illustrated in FIG. 6-1D, in accordance with some non-limiting embodiments. Mechanical structure 6-132 (see FIG. 6-1A) may have a similar arrangement. In the example of FIG. 6-1D, mechanical structure 6-130 includes beams 6-133, 6-134, 6-135 and 6-136. Beam 6-133 connects mechanical driver 6-160 to beam 6-134. Beams 6-135 and 6-136 connect beam 6-134 to the exterior waveguide. To limit optical losses, beams 6-135 and 6-136 may be attached to the exterior waveguide in the

transition regions 6-140 and 6-150, respectively, rather than in the modulation region (as shown in FIG. 6-1E, which is discussed below). However, attaching beams 6-135 and 6-136 to the exterior waveguide to the modulation region is also possible. Beams with different shapes, sizes and orientations may be used in alternative or in addition to those illustrated in FIG. 6-1D.

Mechanical structure 6-130 may transfer mechanical forces generated at mechanical driver 6-160 to waveguide 6-121, thereby causing waveguide 6-121 to move relative to waveguide 6-122. Mechanical drivers 6-160 and 6-162 may be implemented in any suitable way. In one example, the mechanical drivers may include piezoelectric devices. In one example, the mechanical drivers may include conductive fingers. When a voltage is applied between adjacent fingers, the fingers may experience acceleration, thus imparting a mechanical force to the mechanical structures. In some embodiments, the mechanical drivers may be driven with an electrical signal having a pattern encoded thereon. In this way, modulation results in the pattern being imparted onto the phase of an input optical signal.

It should be appreciated that, because the waveguides of suspended multi-slot optical structure 6-120 are driven using external mechanical drivers, rather than being directly supplied with electrical signals as in certain conventional phase modulators, the conductivity of the waveguides can be relaxed, thus reducing free carrier absorption loss, and consequently, dynamic loss. This is different than some conventional phase modulators, where the waveguide itself is doped to act as a heater or a carrier accumulation region. In some embodiments, waveguides 6-121, 6-122 and 6-123 may be made of an undoped, or low-doped, semiconductor material (e.g., undoped silicon or silicon with a doping concentration less than 10^{14}cm^{-3}). In some embodiments, the resistivity of the material forming the waveguides may be greater than $1300\ \Omega\text{cm}$.

FIG. 6-1E illustrates an example of a transition region 6-140, in accordance with some non-limiting embodiments. In this implementation, waveguide 6-122 is contiguous to (e.g., is the continuation of) input waveguide 6-102. As shown, waveguide 6-122 is tapered in the transition region such that its width reduces as it approaches the suspended multi-slot optical structure 6-120. By contrast, waveguides 6-121 and 6-123 are tapered in the transition region such that their widths increase as they depart from suspended multi-slot optical structure 6-120. The tapered waveguides may allow adiabatic coupling between the mode of input waveguide 6-102 and the mode of suspended multi-slot optical structure 6-120, thereby limiting coupling losses. A similar arrangement may be used for transition region 6-150. Due to the enhanced symmetry of the mode supported by the multi-slot structure, transition

regions 6-140 and 6-150 are significantly shorter than other implementations. In some embodiments, the transition regions may be as short as 10 μ m or less, or 5 μ m or less, though other values are also possible.

FIG. 6-2 is a cross sectional view of a NOEMS phase modulator 6-100 taken in a yz-plane passing through waveguide 6-122 (see plane 6-190 in FIG. 6-1B), in accordance with some non-limiting embodiments. Input waveguide 6-102 and output waveguide 6-104 are surrounded by a cladding made of a material (e.g., silicon oxide) with a refractive index lower than the refractive index of the core material. Lower cladding 6-202 is between the waveguide and the underlying substrate 6-201. Upper cladding 6-206 is formed over the waveguide.

To enable free motion of the waveguides of the suspended multi-slot optical structure 6-120, a trench 6-106 is formed through part of upper cladding 6-206. In some embodiments, a portion of the lower cladding 6-202 is removed under the suspended multi-slot optical structure 6-120, thus forming an undercut 6-204. As a result, waveguides 6-121, 6-122 and 6-123 are suspended in air and are free to move responsive to mechanical forces. A cladding/air interface exists at trench sidewall 6-112 and an air/cladding interface exists at trench sidewall 6-114. The sidewalls may be substantially vertical, for example if the trench is formed by reaction ion etching (RIE), or may alternatively be angled. Undercut 6-204 may have curved sidewalls, as illustrated in FIG. 6-2, if an isotropic etch is used, or may alternatively be substantially vertical. In some embodiments, trench 6-106 and undercut 6-204 may be formed as part of the same etch, while in other embodiments, they be formed using separate etches.

FIG. 6-3 is a cross sectional view of a NOEMS phase modulator 6-100 taken in a xy-plane passing through waveguides 6-121, 6-122 and 6-123 (see plane 6-191 in FIG. 6-1B), in accordance with some non-limiting embodiments. FIG. 6-3 shows that waveguides 6-121, 6-122 and 6-123 and beams 6-134, are co-planar (at least in this example), and are suspended in air above substrate 6-201. As further illustrated in this figure, waveguides 6-121, 6-122 and 6-123 do not contact lower cladding 6-202 at this cross section. When mechanical drivers 6-160 and 6-162 are actuated, beams 6-134 and waveguides 6-121 and 6-123 oscillate along the x-axis, thus varying the widths of the slots 6-124 and 6-125. An example of an oscillatory motion of waveguides 6-121 and 6-123 is illustrated, collectively, in FIGs. 6-4A through 6-4C, in accordance with some non-limiting embodiments. FIG. 6-4A illustrates a case in which no mechanical force is applied. As a result, the widths of the slots are unperturbed. In FIG. 6-4B, a pair of forces is applied such that both waveguides 6-121 and 6-123 move towards waveguide 6-122, as illustrated by the arrows. As a result, the widths of the slots are

reduced. In FIG. 6-4C, a pair of forces is applied such that both waveguides 6-121 and 6-123 move away from waveguide 6-122, also illustrated by the arrows. As a result, the widths of the slots are increased. In some embodiments, the forces may be applied in a periodic fashion, and/or following the pattern of the driving electrical signals. In some embodiments, the forces may be applied to waveguides 6-121 and 6-123 differentially, such that the same intensity is applied to both waveguides but with opposite signs.

FIG. 6-5 is a plot illustrating how the effective refractive index (N_{eff}) of the mode propagating in the suspended multi-slot optical structure 6-120 varies as a function of width d_1 (the width of the slot between waveguides 6-121 and 6-122), in accordance with some non-limiting embodiments. A similar response may be plotted as a function of d_2 . The effective index variation is caused by the fact that, as the separation between the waveguides varies under the effect of an applied mechanical force, the shape of the mode deviates relative to the one illustrated in FIG. 6-1C. As the width varies over time, so does the mode effective index, and consequently, the phase of the mode.

FIG. 6-6 is a flowchart illustrating an example of a method for fabricating a NOEMS phase modulator, in accordance with some non-limiting embodiments. It should be appreciated that the steps of the method described below may be performed in any suitable order, as fabrication processes are not limited to the specific order illustrated in FIG. 6-6.

Fabrication method 6-600 begins at step 6-602, in which a chip is obtained. In some embodiments, the chip may be a silicon-on-insulator chip, or a bulk silicon chip. The chip may have a substrate and any of the following layers: a lower cladding layer, a semiconductor layer and an upper cladding layer. The lower cladding layer may comprise silicon oxide in some embodiments. The semiconductor layer may comprises silicon, silicon nitride and/or doped silicon oxide in some embodiments. The upper cladding layer may comprise the same material forming the lower cladding layer, or a different material. FIG. 6-3. Illustrates an examples of a substrate (substrate 6-201) having a lower cladding layer (cladding 6-202), a semiconductor layer (the layer of waveguides 6-121, 6-122 and 6-123) and an upper cladding layer (cladding 6-206). It should be appreciated that any of the layers identified above may already be present on the chip when the chip arrives at the fabrication facility (where the NOEMS phase modulator is fabricated), or may be formed at the facility as part of the fabrication process.

At step 6-604, the semiconductor layer is patterned form a multi-slot optical structure having first and second slots (or any other number of slots greater than two). In the example of FIG. 6-3, waveguides 6-121, 6-122 and 6-123 may be formed at step 6-604. Patterning the

semiconductor layer may involve deposition of a photoresist layer, a photolithographic exposure and etching through the semiconductor layer. In some embodiments, any one of mechanical structures 6-130 and 6-132, mechanical drivers 6-160 and 6-162, waveguides 6-102 and 6-104 and transition regions 6-140 and 6-142 (see FIG. 6-1A) are fabricated as part of the same photolithographic exposure, though not all embodiments are limited in this respect as one or more separate photolithographic exposures may be used. In some embodiments, at step 6-604, mechanical drivers 6-160 may be doped, for example using ion implantation. In some embodiments, the multi-slot optical structure may remain undoped.

At step 6-606, a trench may be formed through the upper cladding layer. An example of a trench (trench 6-106) is illustrated at FIG. 6-3. The trench may be formed, for example, using a dry etch such as a reactive ion etch. However, wet etches may alternatively or additionally be used. Formation of the trench may involve removal of a portion of the upper cladding layer in a region above the multi-slot optical structure formed at step 6-604. As a result, the multi-slot optical structure may be exposed, partially or entirely, to air.

At step 6-608, an undercut may be formed in the lower cladding layer. An example of an undercut (undercut 6-204) is illustrated at FIG. 6-3. The undercut may be formed, for example, using a wet etch, though dry etches may alternatively or additionally be used. Formation of the undercut may involve removal of a portion of the lower cladding layer in a region under the multi-slot optical structure. As a result, at least part of the multi-slot optical structure may be suspended over air.

Having thus described several aspects and embodiments of the technology of this application, it is to be appreciated that various alterations, modifications, and improvements will readily occur to those of ordinary skill in the art. Such alterations, modifications, and improvements are intended to be within the spirit and scope of the technology described in the application. It is, therefore, to be understood that the foregoing embodiments are presented by way of example only and that, within the scope of the appended claims and equivalents thereto, inventive embodiments may be practiced otherwise than as specifically described. In addition, any combination of two or more features, systems, articles, materials, and/or methods described herein, if such features, systems, articles, materials, and/or methods are not mutually inconsistent, is included within the scope of the present disclosure.

Also, as described, some aspects may be embodied as one or more methods. The acts performed as part of the method may be ordered in any suitable way. Accordingly, embodiments may be constructed in which acts are performed in an order different than

illustrated, which may include performing some acts simultaneously, even though shown as sequential acts in illustrative embodiments.

All definitions, as defined and used herein, should be understood to control over dictionary definitions, definitions in documents incorporated by reference, and/or ordinary meanings of the defined terms.

The indefinite articles “a” and “an,” as used herein in the specification and in the claims, unless clearly indicated to the contrary, should be understood to mean “at least one.”

The phrase “and/or,” as used herein in the specification and in the claims, should be understood to mean “either or both” of the elements so conjoined, *i.e.*, elements that are conjunctively present in some cases and disjunctively present in other cases.

As used herein in the specification and in the claims, the phrase “at least one,” in reference to a list of one or more elements, should be understood to mean at least one element selected from any one or more of the elements in the list of elements, but not necessarily including at least one of each and every element specifically listed within the list of elements and not excluding any combinations of elements in the list of elements. This definition also allows that elements may optionally be present other than the elements specifically identified within the list of elements to which the phrase “at least one” refers, whether related or unrelated to those elements specifically identified.

The terms “approximately” and “about” may be used to mean within $\pm 20\%$ of a target value in some embodiments, within $\pm 10\%$ of a target value in some embodiments, within $\pm 5\%$ of a target value in some embodiments, and yet within $\pm 2\%$ of a target value in some embodiments. The terms “approximately” and “about” may include the target value.

CLAIMS

What is claimed is:

1. A photonic processor comprising:
 - a first array of interconnected variable beam splitters (VBSs) comprising a first plurality of optical inputs and a first plurality of optical outputs;
 - a second array of interconnected VBSs comprising a second plurality of optical inputs and a second plurality of optical outputs; and
 - a plurality of controllable optical elements, each of the plurality of these controllable optical elements coupling a single one of the first plurality of optical outputs of the first array to a respective single one of the second plurality of optical inputs of the second array.
2. The photonic processor of claim 1, wherein each of the VBSs of the first and second array comprises a Mach-Zehnder interferometer comprising:
 - a first beam splitter;
 - a second beam splitter; and
 - at least one phase modulator configured to modulate a phase of light of an optical mode that couples the first beam splitter and the second beam splitter.
3. The photonic processor of claim 1, wherein each of the controllable optical elements comprises a variable attenuator and at least one phase modulator.
4. The photonic processor of claim 3, wherein each of the variable attenuators of the controllable optical elements comprises a Mach-Zehnder interferometer.
5. The photonic processor of claim 1, further comprising at least one fold between the first array of interconnected VBSs and the second array of interconnected VBSs, wherein a propagation direction of an optical mode of each of the first plurality of optical outputs is opposite a propagation direction of an optical mode of a respective optical input of the second plurality of inputs.
6. The photonic processor of claim 4, further comprising a plurality of columns of VBSs, wherein at least a subset of the plurality of columns comprises at least one of the

VBSs of the first array of interconnected VBSs and at least one of the VBSs of the second array of interconnected VBSs such that at least a portion of the first array of interconnected VBSs is interleaved with at least a portion of the second array of interconnected VBSs.

7. A photonic processing system comprising:
 - an optical encoder configured to encode an input vector into a first plurality of optical signals;
 - a photonic processor configured to:
 - receive the first plurality of optical signals, each of the first plurality of signals received by a respective input spatial mode of a plurality of input spatial modes of the photonic processor;
 - perform a plurality of operations on the first plurality of optical signals, the plurality of operations implementing a matrix multiplication of the input vector by a matrix; and
 - output a second plurality of optical signals representing an output vector, each of the second plurality of signals transmitted by a respective output spatial mode of a plurality of output spatial modes of the photonic processor; and
 - an optical receiver configured to detect the second plurality of optical signals and output an electrical digital representation of the output vector.
8. The photonic processing device of claim 7, wherein the optical receiver is configured to detect the second plurality of optical signals using phase sensitive detectors.
9. The photonic processing device of claim 8, further comprising a light source configured to:
 - provide first light to the optical encoder for use in encoding the first plurality of optical signals; and
 - provide second light to the optical receiver for use as a local oscillator by the phase sensitive detectors, wherein:
 - the local oscillator is phase coherent with each of the first plurality of optical signals;and
 - a first path length of the first plurality of optical signals from the light source to the optical receiver is substantially equal to a second path length of the local oscillator from the light source to the optical receiver.

10. The photonic processing device of claim 7, wherein the optical encoder is configured to:

encode an absolute value of a vector component of the input vector into an amplitude of a respective optical signal of the first plurality of optical signals; and

encode a phase of the vector component of the input vector into a phase of the respective optical signal of the first plurality of optical signals.

11. The photonic processing device of claim 7, wherein the matrix is an arbitrary unitary matrix.

12. The photonic processing device of claim 7, further comprising:

a plurality of frontends, wherein each of the plurality of frontends is associated with one input spatial mode of the plurality of input spatial modes of the photonic processor, wherein each of the plurality of frontends comprises:

a plurality of optical encoders, each of the optical encoders configured to encode a respective component of an input vector into an optical signal, wherein each optical encoder is configured to output an optical signals of a wavelength different from wavelengths output by the other optical encoders; and

an input wavelength division multiplexer (WDM) configured to receive each of the optical signals from each of the plurality of optical encoders in a separate spatial mode and output each of the optical signals in a single spatial mode connected to a respective input spatial mode of the plurality of input spatial modes of the photonic processor; and

a plurality of backends, wherein each of the plurality of backends is associated with one output spatial mode of a plurality of output spatial modes of the photonic processor, wherein each of the plurality of backends comprises:

an output wavelength division multiplexer (WDM) configured to receive optical signals of different wavelengths from a respective one of the plurality of output spatial modes of the photonic processor and output each of the optical signals of different wavelengths in a respective spatial mode of a plurality of output spatial modes of the WDM;

a plurality of optical receivers, each of the optical receivers configured to determine a respective component of an output vector by detecting a respective optical signal associated with a respective output spatial mode of the WDM.

13. The photonic processing device of claim 7, wherein the photonic processor comprises:
a first array of interconnected variable beam splitters (VBSs) comprising a first plurality of optical inputs corresponding to the first plurality of input spatial modes and a first plurality of optical outputs;

a second array of interconnected VBSs comprising a second plurality of optical inputs and a second plurality of optical outputs corresponding to the plurality of output spatial modes; and

a plurality of controllable optical elements, each of the plurality of these controllable optical elements coupling a single one of the first plurality of optical outputs of the first array to a respective single one of the second plurality of optical inputs of the second array.

14. The photonic processing device of claim 13, further comprising a controller configured to:

perform a singular value decomposition (SVD) of the matrix to determine a first, second, and third SVD matrix;

control the first plurality of interconnected VBSs to implement the first SVD matrix;

control the second plurality of interconnected VBSs to implement the second SVD matrix

control the plurality of controllable optical elements to implement the third SVD matrix, wherein the third SVD matrix is a diagonal matrix.

15. The photonic processing device of claim 14, wherein the controller further comprises at least one digital-to-analog converter (DAC) to adjust one or more parameters of the first plurality of interconnected VBSs and the second plurality of interconnected VBSs.

16. The photonic processing device of claim 15, wherein:

each of the VBSs of the first plurality of interconnected VBSs and each of the VBSs of the second plurality of interconnected VBSs is associated with a respective address; and

the at least one DAC includes a single DAC that controls a plurality of the VBSs of the first and/or second plurality of interconnected VBSs using the addresses.

17. The photonic processing device of claim 7, wherein the matrix is a first matrix and the photonic processing device further comprises a controller configured to control the photonic processing device to perform multiplication of a second matrix by the first matrix by:

- (a) determining a plurality of input vectors from each column of the second matrix;
- (b) selecting an input vector from the plurality of input vectors;
- (c) encoding the selected input vector into the first plurality of optical signals using the optical encoder;
- (d) performing the plurality of operations on the first plurality of optical signals associated with the first input vector;
- (e) detecting the second plurality of optical signals associated with the selected input vector;
- (f) storing digital detection results based on the detected second plurality of optical signals;
- (g) repeating acts (b) – (f) for the other input vectors of the plurality of input vectors;
- (h) digitally combine the digital detection results to determine a resulting matrix resulting from the multiplication of the second matrix by the first matrix.

18. The photonic processing device of claim 7, wherein the optical receiver comprises a low-pass filter configured to perform an analog summation of multiple subsequent signals associated with each output spatial mode of the plurality of output spatial modes of the photonic processor.

19. A method of optically performing matrix-vector multiplication, the method comprising:

- receiving a digital representation of an input vector;
- encoding, using an optical encoder, the input vector into a first plurality of optical signals;
- performing, using a processor, a singular value decomposition (SVD) of a matrix to determine a first, second, and third SVD matrix;
- controlling photonic processor comprising a plurality of variable beam splitters (VBS) to optically implement the first, second, and third SVD matrix;
- propagating the first plurality of optical signals through the photonic processor;

- detecting a second plurality of optical signals received from the photonic processor;
and
determining an output vector based on the detected second plurality of optical signals,
wherein the output vector represents a result of the matrix-vector multiplication.
20. The method of claim 19, wherein the detecting of the second plurality of optical signals is performed using phase-sensitive detectors.
21. The method of claim 20, further comprising:
providing first light, from a light source, to the optical encoder for encoding the first plurality of optical signals;
providing second light, from the light source, to the phase-sensitive detectors, wherein the second light is used as a local oscillator by the phase-sensitive detectors.
22. The method of claim 21, wherein:
the local oscillator is phase coherent with each of the first plurality of optical signals;
and
a first path length of the first plurality of optical signals from the light source to the phase-sensitive detectors is substantially equal to a second path length of the local oscillator from the light source to the phase-sensitive detectors.
23. The method of claim 19, wherein encoding the input vector comprises:
encoding an absolute value of a vector component of the input vector into an amplitude of a respective optical signal of the first plurality of optical signals; and
encoding a phase of the vector component of the input vector into a phase of the respective optical signal of the first plurality of optical signals.
24. The method of claim 19, wherein the matrix is an arbitrary unitary matrix.
25. The method of claim 19, further comprising performing multiple matrix-vector multiplications simultaneously using wavelength division multiplexing.
26. The method of claim 25, wherein:
the input vector is one of a plurality of input vectors:

the method further comprises encoding each of the plurality of input vectors into a respective one of a first plurality of optical signal of a particular wavelength, wherein each wavelength associated with each one of the first plurality of optical signals is different from the other wavelengths of the other ones of the first plurality of optical signals.

27. The method of claim 19, wherein the photonic processor comprises:

a first array of interconnected variable beam splitters (VBSs) comprising a first plurality of optical inputs corresponding to the first plurality of input spatial modes and a first plurality of optical outputs;

a second array of interconnected VBSs comprising a second plurality of optical inputs and a second plurality of optical outputs corresponding to the plurality of output spatial modes; and

a plurality of controllable optical elements, each of the plurality of these controllable optical elements coupling a single one of the first plurality of optical outputs of the first array to a respective single one of the second plurality of optical inputs of the second array.

28. The method of claim 27, wherein:

the first plurality of interconnected VBSs implements the first SVD matrix;

the second plurality of interconnected VBSs implements the second SVD matrix

the plurality of controllable optical elements implements the third SVD matrix,

wherein the third SVD matrix is a diagonal matrix.

29. The method of claim 27, wherein:

each of the VBSs of the first plurality of interconnected VBSs and each of the VBSs of the second plurality of interconnected VBSs is associated with a respective address; and

the VBSs of the first and/or second plurality are controlled by at least one digital to analog converter (DAC) that controls a plurality of the VBSs using the addresses.

30. The method of claim 7, wherein the matrix-vector multiplication is one of a plurality of matrix-vector multiplications performed to perform matrix-matrix multiplication, wherein the matrix is a first matrix and the matrix-matrix multiplication comprises multiplication of a second matrix by the first matrix by, the method further comprising:

(a) determining a plurality of input vectors from each column of the second matrix;

(b) selecting an input vector from the plurality of input vectors;

- (c) encoding the selected input vector into the first plurality of optical signals using the optical encoder;
- (d) performing the plurality of operations on the first plurality of optical signals associated with the first input vector;
- (e) detecting the second plurality of optical signals associated with the selected input vector;
- (f) storing digital detection results based on the detected second plurality of optical signals;
- (g) repeating acts (b) – (f) for the other input vectors of the plurality of input vectors;
- (h) digitally combine the digital detection results to determine a resulting matrix resulting from the multiplication of the second matrix by the first matrix.