# Chip-to-chip Optical Data Communications using Polarization Division Multiplexing

Darko Ivanovich
Chenfeng Zhao
Xuan Zhang
Roger D. Chamberlain
*McKelvey School of Engineering*
*Washington University in St. Louis*
{di3,chenfeng.zhao,xuan.zhang,roger}@wustl.edu

Amit Deliwala
Viktor Gruev
*Dept. of Electrical and Computer Engineering*
*University of Illinois Urbana-Champaign*
{amitd2,vgruev}@illinois.edu

***Abstract*—Short distance optical communication is challenging in significant part because the expense of constructing effective systems is high. We describe an optical data communication system that is designed to operate over very short distances (neighboring chips on a board) and is compatible with traditional CMOS fabrication, substantially decreasing the cost to build relative to previous approaches. Polarization division multiplexing is exploited to increase the achievable data rates.**

***Index Terms*—photonics, optical links, polarization**

## I. Introduction

Optical communication has been used extensively for long haul applications (city-to-city, or even around the world). However, its use for shorter distances has been much more problematic, in significant part due to the manufacturing costs involved [1]. Yet, the need for short distance high-bandwidth communications continues to grow in a seemingly unbounded way. The memory wall, a significant limitation for high-performance computing performance, is due in large part to limited chip-to-chip bandwidth capacity [2]. The use of optics to ameliorate this limitation has a long research history [3], [4], but hasn't yet proved commercially viable.

Here, we assess the viability of using polarization division multiplexing techniques for an inexpensive optical chip-to-chip data path. The benefit we can achieve through the use of polarization is increased data rates. The ability to build an $n$-channel system enables an $n$-fold increase in data rate. A set of VCSELs (vertical emitting lasers) integrated into the transmitter launch a set of light signals through a set of polarizing filters (polarimeters), each of which is oriented at a distinct polarization angle relative to the others. The light is reflected off a pair of $45^\circ$ mirrors, travels through a corresponding set of polarimeters on the receiver, and impinges on a set of CMOS photodiodes. The basic elements are illustrated in Figure 1.

The reason this optical system is cost effective is that the lasers, polarimeters and detectors can all be fabricated on-chip. These on-chip components (particularly the polarimeters) are not nearly as high quality as those available for optical bench work. The question we attempt to address is, "are they good enough for the task at hand?"

This work supported by NSF grants CNS-1739643 and CNS-1763503.

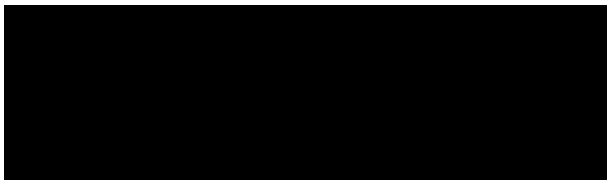

Fig. 1. Chip-to-chip optical communications path.

To answer this question, we model the optical data path and quantify the relationship between optical power, number of channels, and noise margin at the receiver. Exploiting empirical data for calibration purposes (both reported here and in the literature), we investigate the viability of both two- and three-channel designs. To assess one potential performance implication of optical chip-to-chip interconnects, we simulate the use of a polarization division multiplexed optical path as the link between a processor's last level cache and main memory.

## II. Background and Related Work

In this work, we will extensively use Stokes vectors to represent the polarization signature of a light signal and Mueller matrices to represent the transformation of that polarization signature by one or more optical elements [5], [6]. As such, an input light signal

$$S_{IN} = \begin{bmatrix} S_0 \\ S_1 \\ S_2 \\ S_3 \end{bmatrix}, \tag{1}$$

where $S_i$ represents the $i$th Stokes parameter of $S_{IN}$, that is sent through an optical element modeled via Mueller matrix $M \in \mathbb{R}^{4\times4}$ results in the output light signal

$$S_{OUT} = MS_{IN}. \tag{2}$$

Note that we will not always normalize the Stokes vector so that $S_0 = 1$, but will allow the magnitude of the light intensity to be included as $S_0$.

Recent work using Division of Focal Plane (DoFP) polarimeters integrated with CMOS technology has enabled compact, real-time polarization imaging [7]–[9]. The aluminum nanowire polarimeters were empirically measured by Powell and Gruev [10]. Their fabrication is compatible with standard CMOS processing steps. Figure 2 is an SEM image of one of these polarimeters oriented at $45°$.

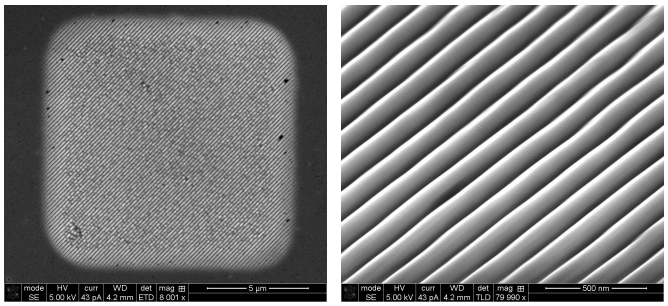

Fig. 2. $45°$ oriented aluminum nanowire polarimeter (including both wide view and zoomed in).

The benefit of this construction is that the cost-to-build in volume is dramatically lower than previously constructed systems using an optical bench [11].

Receiver subsystems for 2-, 3-, and 4-channel optical links that exploit these polarimeters are described by Ivanovich et al. [12], concluding that both 2- and 3-channel systems (each channel being associated with a unique polarization angle in the transmitted light) can be effective using a free-space optical path.

PAM-4 receiver electronics have been demonstrated by Elhadidy et al. [13] that operate at 32 Gb/s rates, and high-speed CMOS laser drivers have been reported by Hyun [14], Illing and Kernel [15], and Kern [16]. With the above elements (CMOS compatible polarimeters as well as high-speed receiver circuits and laser drivers), we have the constituent components for an optical communications system that exploits polarization division multiplexing and that is highly integrated with CMOS fabrication.

Polarization division multiplexing communications has been demonstrated several times in the context of an optical bench. Gigabit per second data rates have been reported by Wang et al. [11], Hsu et al. [17], Morant et al. [18], Yao et al. [19], and Kwon et al. [20], all using the two orthogonal channels that are separated by $90°$. In addition, Herard and Lacourt [21] showed the feasibility of three channels with different polarization angles, and Chen et al. [22] demonstrated an approach for four independent channels.

Finally, the use of optics for processor-to-memory interconnect has been evaluated by Fritts and Chamberlain [3] and processor-to-processor interconnect (exploiting wavelength division multiplexing rather than polarization division multiplexing) by Kirman et al. [23].

## III. Path Model Description

### A. Optical Path Model

As stated in Section II, we will model optical signals via their Stokes vectors and optical path elements via their Mueller matrices. The symbols we use are summarized in Table I

TABLE I
Symbol Definitions.

| Symbol | Definition |
|---|---|
| $C$ | Set of channels |
| $b_i$ | binary bit transmitted on channel $i$ |
| $S_T$ | Transmitted light on each channel |
| $S_{F.i}$ | Light out of transmit polarimeter $i$ |
| $S_F$ | Complete light signal from polarimeters |
| $S_{M1}$ | Reflected light from first mirror |
| $S_{M2}$ | Reflected light from second mirror |
| $S_{R.i}$ | Light impinging on photodiode $i$ |
| $I_{R.i}$ | Intensity of light delivered to photodiode $i$ |
| $M_{T.i}$ | Mueller matrix for transmit polarimeter $i$ |
| $M_M$ | Mueller matrix for each mirror |
| $M_{R.i}$ | Mueller matrix for receiver polarimeter $i$ |

If (when sending a binary 1) the transmitting laser provides a narrow beam of intensity $I_T$, the transmitted light on each channel is $S_T = [I_T \quad 0 \quad 0 \quad 0]^T$. After being filtered by the transmit polarimeters, the set of light signals that mix in the path are as follows

$$S_{F.i} = M_{T.i} S_T, \qquad i \in C \tag{3}$$

and the resulting transmitted optical signal is

$$S_F = \sum_{i \in C} b_i S_{F.i} \tag{4}$$

where $b_i$ is the bit transmitted on channel $i$ (constrained to a value of either 0 or 1).

This optical signal is reflected twice

$$S_{M1} = M_M S_F \tag{5}$$

$$S_{M2} = M_M S_{M1} = M_M M_M S_F \tag{6}$$

and filtered at the receiver's polarimeters,

$$S_{R.i} = M_{R.i} S_{M2}, \qquad i \in C. \tag{7}$$

The light intensity delivered to the photodiode, $I_{R.i}$, is simply the first Stokes parameter of $S_{R.i}$,

$$I_{R.i} = [1 \quad 0 \quad 0 \quad 0] S_{R.i}, \qquad i \in C. \tag{8}$$

Due to a combination of manufacturing variation and (more significantly) interference from the other channels (all of which is included in the model above [12]), a high-valued input might be as low as $I_{R.i-H(MIN)}$ and a low-valued input might be as high as $I_{R.i-L(MAX)}$.

*1) Two-channel System:* For a two-channel system, we will exploit the two orthogonal polarization angles of $0°$ and $90°$ (i.e., $C = \{0°, 90°\}$). The ideal transmitted signals for each channel are

$$S_{F.0°} = \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} \tag{9}$$

and

$$S_{F.90°} = \begin{bmatrix} 1 \\ -1 \\ 0 \\ 0 \end{bmatrix} \tag{10}$$

where the intensity has been normalized to unity. In Section V, we will compute these values from (3) using empirical data for the transmit polarimeters, $M_{T.i}$.

*2) Three-channel System:* For a three-channel system, we follow the theoretical guidance of Tyo [24] and evenly space the three channels at $0°$, $60°$, and $120°$ (i.e., $C = \{0°, 60°, 120°\}$). In this case, the ideal transmitted signals for each channel are

$$S_{F.0°} = \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \tag{11}$$

$$S_{F.60°} = \begin{bmatrix} 1 \\ -\frac{1}{2} \\ \frac{\sqrt{3}}{2} \\ 0 \end{bmatrix}, \tag{12}$$

and

$$S_{F.120°} = \begin{bmatrix} 1 \\ -\frac{1}{2} \\ -\frac{\sqrt{3}}{2} \\ 0 \end{bmatrix}. \tag{13}$$

### *B. Electrical Path Model*

For the electrical path, we assume a straightforward receiver as illustrated in Figure 3.

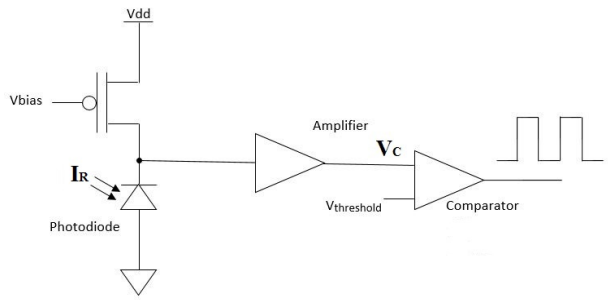


Fig. 3. Receiver block diagram.

The voltage at the input to the comparator, $V_C$, can be represented as

$$V_C = G_{IE} G_{OI} I_R \pm \epsilon \tag{14}$$

where $G_{OI}$ specifies the gain of the photodiode that converts the input light signal $I_R$ into current, $G_{IE}$ specifies the gain of the transimpedance amplification that converts photodiode current into voltage, and $\epsilon$ represents the output-reflected electrical noise due to, e.g., the noise in the readout electronics, the shot noise of the photodiode, and the manufacturing variation from one channel to another.

This yields worst-case electrical inputs to the comparator as follows

$$V_{C.H(MIN)} = G_{IE} G_{OI} I_{R-H(MIN)} - \epsilon \tag{15}$$

and

$$V_{C.L(MAX)} = G_{IE} G_{OI} I_{R-L(MAX)} + \epsilon. \tag{16}$$

## IV. Model Calibration

To provide empirical data to calibrate the model, we performed measurements using the experimental setup illustrated in Figure 4.

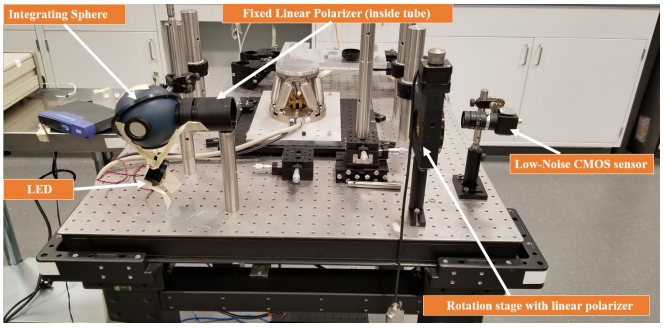


Fig. 4. Optical bench experiment.

As illustrated in the figure, a 530 nm uniform green LED light source illuminates light that travels to an integrating sphere where the light signal is diffused and then an aperture is used to create a narrow light beam. After this process a collimating lens is used to create a straight, uniform beam.

This beam is polarized with a fixed linear polarizer (manufactured by Newport) located inside the tube that is adjacent to the integrating sphere. The polarized light passes through a rotation stage with linear polarizer also made by Newport that is manually rotated from $0°$ to $180°$ in steps of 2 degrees. The light is then passed to a linear low-noise monochrome CMOS sensor [25] and the data is recorded. The image received by the CMOS sensor is limited to a 100x100 pixel square array to ensure uniformity of illumination.

The measured data is a 4D array structure with 91 angles from $0^o$ to $180^o$, every two degrees, 100 pixels in the x-axis direction, 100 pixels in the y-axis direction, and 64 frames of data for each angle. Internal to the CMOS sensor, the light signal is digitized using a 12-bit analog-to-digital converter (ADC) with a 3.5 V reference voltage.

The mean value for each angle, along with its standard deviation, was computed and is shown in Figure 5. The data is, as expected, showing Malus's law distribution with noise approximately equivalent to that predicted by shot noise.

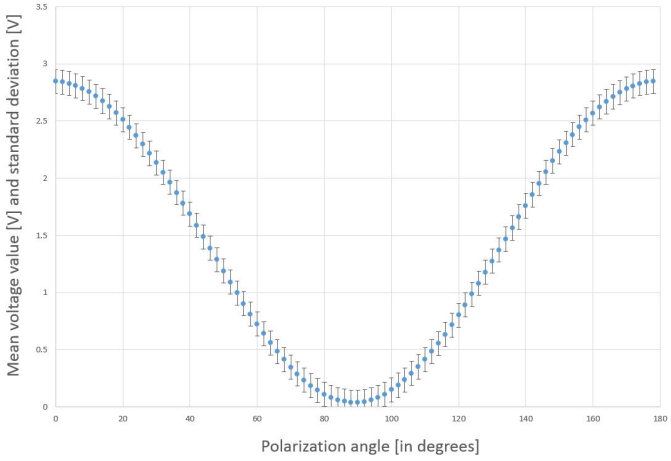


Fig. 5. Receiver signal level and std. dev.

Figure 6 shows the relationship between the standard deviation and received signal for the linear CMOS sensor. A least mean squared error curve fit to this data yields the expression

$$\sigma = y = 0.0107x + 0.0036 \tag{17}$$

where $y$ is the std. dev. of the sensor reading and $x$ is the mean reading for a given angle.

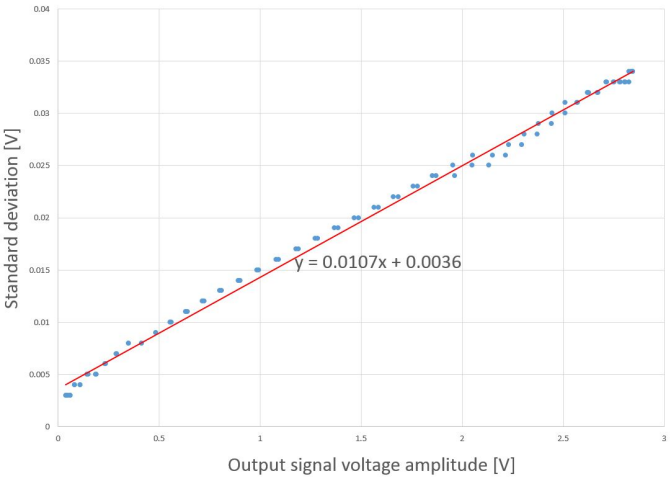


Fig. 6. Std. dev. as a function of signal level.

What is clear from the figure is that the output reflected noise is dominated by pixel-to-pixel variations, not the shot noise of the individual detectors (which have Poisson statistics rather than the linear relationship that is measured).

In what follows, we will use this measured relationship as a calibration (i.e., a proxy) for the electrical noise in linear receiver circuits. Although for practical applications it will be necessary to use the noise properties of the actual receiver electronics in use, we anticipate that receiver-to-receiver variation will still dominate the shot noise of the detector.

## V. Data Link Evaluation

In this section we will perform a quantitative assessment of the relationship between input optical power, number of channels, and resultant noise margins for the system of Figure 1. Where we have (and use) quantitative data for component calibration, we will point out its source.

Since the basic system design is linear, we start with a normalized set of laser sources, each providing unpolarized light, i.e.,

$$S_T = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}. \tag{18}$$

As such, the resulting noise margins will need to be scaled, appropriately.

The assessment will consider two-channel and three-channel designs in succession.

### A. Two-channel System

In the two-channel system, the two channels are separated by a polarization angle of 90°. We will identify each channel by its polarization angle. From [10], the Mueller matrices, $M_{T.i}$, for the aluminum nanowire polarimeters are

$$M_{T.0^\circ} = \begin{bmatrix} 0.25 & 0.2475 & 0 & 0 \\ 0.2475 & 0.25 & 0 & 0 \\ 0 & 0 & 0.071 & 0 \\ 0 & 0 & 0 & 0.035 \end{bmatrix} \tag{19}$$

and

$$M_{T.90^\circ} = \begin{bmatrix} 0.25 & -0.2475 & 0 & 0 \\ -0.2475 & 0.25 & 0 & 0 \\ 0 & 0 & 0.071 & 0 \\ 0 & 0 & 0 & 0.035 \end{bmatrix} \tag{20}$$

which yields the following two light signals being transmitted

$$S_{F.0^\circ} = M_{T.0^\circ} S_T = \begin{bmatrix} 0.25 \\ 0.2475 \\ 0 \\ 0 \end{bmatrix} \tag{21}$$

$$S_{F.90^\circ} = M_{T.90^\circ} S_T = \begin{bmatrix} 0.25 \\ -0.2475 \\ 0 \\ 0 \end{bmatrix} \tag{22}$$

and summed to

$$S_F = b_0 S_{F.0^\circ} + b_{90} S_{F.90^\circ} \tag{23}$$

which is the light incident on the first mirror. Table II gives the four possible values for $S_F$.

TABLE II
$S_F$ FOR ALL POSSIBLE BINARY INPUTS.

| $b_{90}$ | 0 | 0 | 1 | 1 |
|---|---|---|---|---|
| $b_0$ | 0 | 1 | 0 | 1 |
| $S_F$ | $\begin{bmatrix}0\\0\\0\\0\end{bmatrix}$ | $\begin{bmatrix}0.25\\0.2475\\0\\0\end{bmatrix}$ | $\begin{bmatrix}0.25\\-0.2475\\0\\0\end{bmatrix}$ | $\begin{bmatrix}0.5\\0\\0\\0\end{bmatrix}$ |

Each mirror reflects the light signal via the Mueller matrix $M_M$, which comes from [26], [27].

$$M_M = \begin{bmatrix} 1 & -0.734 & 0 & 0 \\ -0.734 & 1 & 0 & 0 \\ 0 & 0 & 0.646 & -0.21 \\ 0 & 0 & 0.21 & 0.646 \end{bmatrix} \tag{24}$$

Therefore,

$$S_{M1} = M_M S_F \tag{25}$$

and

$$S_{M2} = M_M S_{M1}. \tag{26}$$

The four values for $S_{M2}$ are shown in Table III.

TABLE III
$S_{M2}$ FOR ALL POSSIBLE BINARY INPUTS.

| $b_{90}$ | 0 | 0 | 1 | 1 |
|---|---|---|---|---|
| $b_0$ | 0 | 1 | 0 | 1 |
| $S_{M2}$ | $\begin{bmatrix}0\\0\\0\\0\end{bmatrix}$ | $\begin{bmatrix}0.0214\\0.0138\\0\\0\end{bmatrix}$ | $\begin{bmatrix}0.748\\-0.748\\0\\0\end{bmatrix}$ | $\begin{bmatrix}0.769\\-0.734\\0\\0\end{bmatrix}$ |

The incident light at each photodiode is then

$$I_{R.0^\circ} = [1 \quad 0 \quad 0 \quad 0] M_{R.0^\circ} S_{M2} \tag{27}$$

$$I_{R.90^\circ} = [1 \quad 0 \quad 0 \quad 0] M_{R.90^\circ} S_{M2} \tag{28}$$

and the resultant values are shown in Table IV.

TABLE IV
INCIDENT LIGHT ON PHOTODIODES.

| $b_{90}$ | 0 | 0 | 1 | 1 |
|---|---|---|---|---|
| $b_0$ | 0 | 1 | 0 | 1 |
| $I_{R.90^\circ}$ | 0 | 0.00191 | 0.372 | 0.374 |
| $I_{R.0^\circ}$ | 0 | 0.00877 | 0.00191 | 0.0107 |

If we (arbitrarily) assign $G_{IE}G_{OI}$ to balance the electrical signals in each channel and yield a 3.5 V top of scale, and assume that the interference from the other channel dominates the noise in the optical system,

$$V_{C.H(MIN)} = 2.87 \text{ V} - \epsilon \tag{29}$$

$$V_{C.L(MAX)} = 0.62 \text{ V} + \epsilon. \tag{30}$$

Using $6\sigma$ effective bounds for the electrical noise (i.e., $\epsilon = 6\sigma$), from equation (17),

$$V_{C.H(MIN)} = 2.87 - 0.21 = 2.66 \text{ V} \tag{31}$$

$$V_{C.L(MAX)} = 0.62 + 0.06 = 0.68 \text{ V}. \tag{32}$$

Setting the threshold voltage at the midpoint, 1.67 V, the effective noise margins are

$$N_H = V_{C.H(MIN)} - V_T = 0.99 \text{ V} \tag{33}$$

$$N_L = V_T - V_{C.L(MAX)} = 0.99 \text{ V}. \tag{34}$$

While there are a number of caveats to include in the above analysis (e.g., the output-reflected noise will vary with the gain of the electrical subsystem, while we have assumed above that it does not), the end result is that the noise margins are sufficiently high (together over half of the 3.5 V total range) that a system of this type is certainly achievable. While we don't perform the quantitative analysis here, it should be clear that alternative modulation schemes (e.g., PAM-4) are also quite viable on the two-channel system.

### *B. Three-channel System*

In the three-channel system, the channels are each separated by $60^\circ$ of polarization angle. Here, due to significantly more crosstalk between the channels, we expect the resulting noise margins to be substantially lower than is the case for the two-channel design.

From [10], the Mueller matrices, $M_{T.i}$, for the two new transmit polarimeters are

$$M_{T.60^\circ} = \begin{bmatrix} 0.25 & -0.124 & 0.214 & 0 \\ -0.124 & 0.089 & -0.093 & 0 \\ 0.214 & -0.093 & 0.196 & 0 \\ 0 & 0 & 0 & 0.035 \end{bmatrix} \tag{35}$$

and

$$M_{T.120^\circ} = \begin{bmatrix} 0.25 & -0.124 & -0.214 & 0 \\ -0.124 & 0.089 & 0.093 & 0 \\ -0.214 & 0.093 & 0.196 & 0 \\ 0 & 0 & 0 & 0.035 \end{bmatrix} \tag{36}$$

and $M_{T.0^\circ}$ is given by equation (19). This yields the three light signals being transmitted ($S_{F.0^\circ}$ being given by equation (21)),

$$S_{F.60^\circ} = M_{T.60^\circ} S_T = \begin{bmatrix} 0.25 \\ -0.124 \\ 0.214 \\ 0 \end{bmatrix} \tag{37}$$

$$S_{F.120^\circ} = M_{T.120^\circ} S_T = \begin{bmatrix} 0.25 \\ -0.124 \\ -0.214 \\ 0 \end{bmatrix}, \tag{38}$$

and summed to

$$S_F = b_0 S_{F.0^\circ} + b_{60} S_{F.60^\circ} + b_{120} S_{F.120^\circ} \tag{39}$$

which is the light incident on the first mirror.

Following the same procedure as for the two-channel system, we can determine the nominal light signal at each photodiode. These values are given in Table V.

TABLE V
INCIDENT LIGHT ON PHOTODIODES.

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| $b_{120}$ | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 |
| $b_{60}$ | 0 | 0 | 1 | 1 | 0 | 0 | 1 | 1 |
| $b_0$ | 0 | 1 | 0 | 1 | 0 | 1 | 0 | 1 |
| $I_{R.120°}$ | 0 | 0.00363 | 0.194 | 0.197 | 0.228 | 0.231 | 0.421 | 0.425 |
| $I_{R.60°}$ | 0 | 0.00363 | 0.228 | 0.231 | 0.194 | 0.197 | 0.421 | 0.425 |
| $I_{R.0°}$ | 0 | 0.00877 | 0.0036 | 0.0124 | 0.0036 | 0.0124 | 0.0073 | 0.016 |

Again assigning $G_{IE}G_{OI}$ to balance the electrical signals in each channel and yield a 3.5 V top of scale, the voltage signals at the comparators are as follows,

$$V_{C.H(MIN)} = 1.88 \text{ V} - \epsilon \tag{40}$$

$$V_{C.L(MAX)} = 1.62 \text{ V} + \epsilon. \tag{41}$$

Again using $\epsilon = 6\sigma$ for the electrical noise, from equation (17),

$$V_{C.H(MIN)} = 1.88 - 0.14 = 1.74 \text{ V} \tag{42}$$

$$V_{C.L(MAX)} = 1.62 + 0.13 = 1.75 \text{ V}. \tag{43}$$

This, unfortunately gives a negative noise margin, which leads us to conclude that the three-channel system would, at the very least, require channel coding to be viable.

The above negative conclusion is in direct contrast to our earlier work [12], in which an evaluation using a free-space optical path without mirrors was deemed reasonable. It is apparent that the addition of a pair of mirrors in the path can swing the decision from one conclusion to its opposite.

## VI. USE CASE EVALUATION

The benefits of increased capacity in an optical data link are ubiquitous, with usage ranging from rack-to-rack communications in a machine room, to processor-to-processor interconnections (see, e.g., [23]), to a processor-to-memory data path. Here, we will investigate the performance implications of the latter, the use of an optical data path between the last-level cache of a processor and the main memory. The performance evaluation that follows is not intended to be comprehensive; rather, it is intended to be illustrative of one potential use case.

Using techniques similar to those of [3], we develop an architectural simulation model (see Figure 7) of a baseline system whose properties are given in Table VI using the gem5 toolsuite [28]. Hybrid Memory Cube (HMC) technology is used for the memory, which consists of a stack of memory chips coupled with a single logic chip (for control and interface purposes) [29]. Conveniently, HMC technology uses serial links for data communications, making it easier to compare the baseline electrical data paths with an optical data path.

We compare the execution time on this reference system to a comparable system in which the memory link has been replaced with an optical path that has the same control latency and 4× the data rate, as proposed by Pawlowski [31]. The physical structure of this new system is illustrated in Figure 8, in which the logic chip is presumed to be at the top of the memory stack (rather than at the bottom). The set of applications are drawn from the DIBS benchmark suite [32].

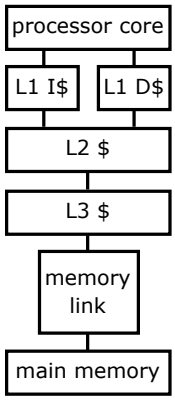


Fig. 7. Architectural simulation model.

TABLE VI
COMPUTE SYSTEM PROPERTIES.

| **Parameter** | **Value** |
|---|---|
| Processor core | Arm OoO at 4 GHz [30] |
| L1 I-cache and D-cache | 32 kB each |
| L2 cache | 2 MB |
| L3 cache | 8 MB |
| Memory link | 40 GB/s per SerDes link [29] |
| Main memory | 4 GB HMC [29] |

Figure 9 shows the performance improvement (relative to the baseline system) of each of the benchmark applications. The geometric mean of all 12 applications shows a speedup of 1.13×, with individual performance improvements varying from only 1% (`edgelist->csr`) to as much as 27% (`ebcdic->txt`). Additionally, with only a single core being modeled, the overall utilization of the memory link is low (under 5%), even in the baseline architecture. Of course, in a practical system upwards of 16 cores (or more) would be present. The optical data path keeps the memory link from saturating (which would have substantial negative performance impact) at even larger core counts.

## VII. CONCLUSIONS

We have presented a model that describes the salient characteristics of a chip-to-chip communications pathway that utilizes polarization division multiplexing to provide multiple distinct channels. The entire system is compatible with modern CMOS fabrication techniques, thereby enabling an ability to build the system in a cost-effective manner. Laser light that is launched from the transmitter is filtered via aluminum nanowire polarimeters, reflected twice off of aluminum mirrors (including a thin aluminum oxide layer), and filtered again via polarimeters on the receiver impinges on photodiodes for conversion back into an electrical signal.

The model is calibrated using experimental measurements that are reported in the literature, some by us and some by others. This calibration specifically includes the characterization of the Mueller matrices for the transmitter and receiver polarimeters as well as the Mueller matrices for the mirrors.

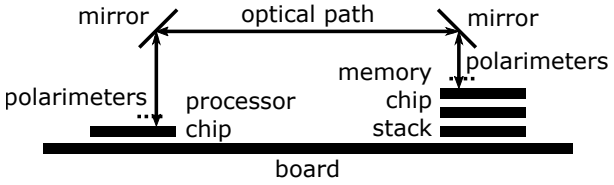


Fig. 8. Processor-to-memory optical communications path.

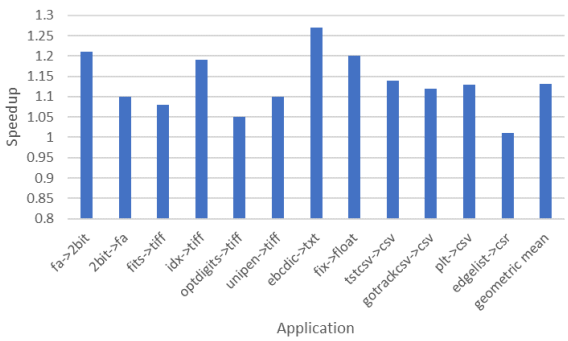


Fig. 9. Application speedup with increased memory data rate.

For the two-channel system, in which the two data channels are separated by a polarization angle of 90°, there is excellent channel separation and a relatively limited amount of optical signal interference between the channels (even given the fact that the aluminum nanowire polarimeters and the mirrors are far from ideal). As such, the noise margins seen at the input to the comparators that decode the received signals are plenty wide.

The three-channel system, however, in which the three data channels are separated by polarization angles of 60° between channels, does not yield a viable system. This is due to the combination of imperfections in the components that make up the optical path as well as the fact that there will be inherently a larger amount of optical signal interference simply due to the fact that the 3-channel system is fundamentally trading noise margin for additional channels. The noise margins at the comparators are less than 0, implying that if we are to practically make a system like this work, either some design choices need to be reconsidered (e.g., better polarimeters or better mirrors) or channel coding will be required.

We also explored the performance implications of an optical data path inserted between the last level cache and the main memory as an example use case. Architectural simulation results showed an average speedup of 13% (over 12 applications) under circumstances where the data path is not saturated (i.e., the performance gains are not due to queueing delays waiting for a contended resource, they are directly due to the improved data rate in the cache-to-memory data path).

There are a number of additional investigations to pursue in this work. First, we have assumed that the source lasers can launch a sufficiently narrow beam that it is effectively detectable at the receiver. This might be accomplished via microlenses prior to the polarimeters, however, the construction of this is unproven. Second, the mechanical stability of the two mirrors in the optical path must be ascertained and assessed. Third, the electrical noise properties we used in our analysis were based on measurements from an image sensor. This will need to be replaced by a high-speed receiver (e.g., as presented by Elhadidy et al. [13]).

The above investigations will help us be more confident of the conclusions we have made. The directions articulated next enable the system to be used in a wider range of applications as well as potentially yield even higher data rates. While the optical path in the present work is assumed to be free space, it is clearly worth investigating the viability of a fiber-based optical path. A Mueller matrix for a 10 km fiber channel can be derived from measurements by Dong et al. [33], however, a more interesting use case might be shorter distances (e.g., within a machine room). Hollow-core fiber provides one more potential path [34]. Another intriguing option would be the use of polymer waveguides as part of the optical path.

Finally, this paper assumes the use of pulse code modulation, or the transmission of light to signal a binary 1 and the absence of light to signal a binary 0. Given the noise margins available in the 2-channel system, it is clearly plausible to expand this system to using PAM-4 modulation, giving 4 symbols per unit time on each of the two polarization channels, for a total of 16 symbols per unit time or 4 bits per unit time. In yet another approach, the work of Garcia et al. [9] could enable simultaneous polarization division multiplexing and wavelength division multiplexing on chip-level scales.

## References

[1] A. F. J. Levi, "Silicon photonics stumbles at the last meter," *IEEE Spectrum*, pp. 38–43, Sep. 2018.

[2] W. A. Wulf and S. A. McKee, "Hitting the memory wall: Implications of the obvious," *ACM SIGARCH Computer Architecture News*, vol. 23, no. 1, pp. 20–24, 1995.

[3] J. E. Fritts and R. D. Chamberlain, "Breaking the memory bottleneck with an optical data path," in *Proc. of 35th Annual Simulation Symposium*, Apr. 2002, pp. 352–362.

[4] D. Qian and D. Zhu, "Challenges and possible approaches: towards the petaflops computers," *Frontiers of Computer Science in China*, vol. 3, no. 3, pp. 273–289, Sep. 2009.

[5] W. H. McMaster, "Matrix representation of polarization," *Reviews of Modern Physics*, vol. 33, no. 1, pp. 8–28, 1961.

[6] ——, "Polarization and the Stokes parameters," *American Journal of Physics*, vol. 22, no. 6, pp. 351–362, 1954.

[7] V. Gruev, R. Perkins, and T. York, "CCD polarization imaging sensor with aluminum nanowire optical filters," *Optics Express*, vol. 18, no. 18, pp. 19 087–19 094, 2010.

[8] R. Perkins and V. Gruev, "Signal-to-noise analysis of Stokes parameters in division of focal plane polarimeters," *Optics Express*, vol. 18, no. 25, pp. 25 815–25 824, 2010.

[9] M. Garcia, C. Edmiston, R. Marinov, A. Vail, and V. Gruev, "Bio-inspired color-polarization imager for real-time in situ imaging," *Optica*, vol. 4, no. 10, pp. 1263–1271, 2017.

[10] S. B. Powell and V. Gruev, "Calibration methods for divison-of-focal-plane polarimeters," *Optics Express*, vol. 21, no. 18, pp. 21 039–21 055, 2013.

[11] Y. Wang, C. Yang, Y. Wang, and N. Chi, "Gigabit polarization division multiplexing in visible light communication," *Optics Letters*, vol. 39, no. 7, pp. 1823–1826, 2014.

[12] D. Ivanovich, S. B. Powell, V. Gruev, and R. D. Chamberlain, "Polarization division multiplexing for optical data communications," in *Proc. of SPIE 10538, Optical Interconnects XVIII*, 2018.

[13] O. Elhadidy, A. Roshan-Zamir, H. Yang, and S. Palermo, "A 32 Gb/s 0.55 mW/Gbps PAM4 1-FIR 2-IIR tap DFE receiver in 65-nm CMOS," in *Proc. of Symposium on VLSI Circuits*, Jun. 2015, pp. C224–C225.

[14] S. H. Hyun, "Design of High-Speed CMOS Laser Driver Using a Standard CMOS Technology for Optical Data Transmission," Ph.D. dissertation, Georgia Institute of Technology, Atlanta, GA, USA, 2004.

[15] L. Illing and M. B. Kernel, "Shaping current waveforms for direct modulation of semiconductor lasers," *IEEE Journal of Quantum Electronics*, vol. 40, no. 5, pp. 445–452, 2004.

[16] A. M. Kern, "CMOS Circuits for VCSEL-Based Optical IO," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, USA, 2004.

[17] C.-W. Hsu, C.-H. Yeh, and C.-W. Chow, "Using adaptive equalization and polarization-multiplexing technology for gigabit-per-second phosphor-led wireless visible light communication," *Optics & Laser Technology*, vol. 104, pp. 206–209, 2018.

[18] M. Morant, J. Pérez, and R. Llorente, "Polarization division multiplexing of OFDM radio-over-fiber signals in passive optical networks," *Advances in Optical Technologies*, vol. 2014, p. Article ID 269524, 2014.

[19] S. X. Yao, L. Yan, B. Zhang, A. Willner, and J. Jiang, "All-optic scheme for automatic polarization division demultiplexing," *Optics Express*, vol. 15, no. 12, pp. 7407–7414, 2007.

[20] D.-H. Kwon, S.-J. Kim, S.-H. Yang, and S.-K. Han, "Optimized pre-equalization for gigabit polarization division multiplexed visible light communication," *Optical Engineering*, vol. 54, no. 7, p. 076101, 2015.

[21] C. Herard and A. Lacourt, "New multiplexing technique using polarization of light," *Applied Optics*, vol. 30, no. 2, pp. 222–231, 1991.

[22] Z.-Y. Chen, L.-S. Yan, Y. Pan, L. Jiang, A.-L. Yi, W. Pan, and B. Luo, "Use of polarization freedom beyond polarization-division multiplexing to support high-speed and spectral-efficient data transmission," *Light: Science & Applications*, vol. 6, no. 2, p. e16207, 2017.

[23] N. Kirman, M. Kirman, R. K. Dokania, J. F. Martinez, A. B. Apsel, M. A. Watkins, and D. H. Albonesi, "Leveraging optical technology in future bus-based chip multiprocessors," in *Proc. of 39th IEEE/ACM International Symposium on Microarchitecture*, 2006, pp. 492–503.

[24] J. S. Tyo, "Optimum linear combination strategy for an N-channel polarization-sensitive imaging or vision system," *Journal of the Optical Society of America A*, vol. 15, no. 2, pp. 359–366, 1998.

[25] *6.5μm, 4MP Scientific BSI CMOS Image Sensor*, GpixelINC, Aug. 2017, rev. 0.2.

[26] G. V. Harten, F. Snik, and C. Keller, "Polarization properties of real aluminum mirrors, I. influence of the aluminum oxide layer," *Publications of the Astronomical Society of the Pacific*, vol. 121, no. 878, pp. 377–383, 2009.

[27] J. Krijger, R. Snel, G. van Harten, J. Rietjens, and I. Aben, "Mirror contamination in space I: mirror modelling," *Publications of the Atmospheric Measurement Techniques*, vol. 7, p. 3387–3398, 2014.

[28] N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, T. Krishna, S. Sardashti *et al.*, "The gem5 simulator," *ACM SIGARCH Computer Architecture News*, vol. 39, no. 2, pp. 1–7, 2011.

[29] *Hybrid Memory Cube Specification 2.1*, Hybrid Memory Cube Consortium, Nov. 2015.

[30] M. Lin, T. Huang, C. Tsai, K. Tam, C. Hsieh, T. Chen, W. Huang, J. Hu, Y. Chen, S. K. Goel, C. Fu, S. Rusu, C. Li, S. Yang, M. Wong, S. Yang, and F. Lee, "A 7nm 4GHz Arm®-core-based CoWoS® chiplet design for high performance computing," in *Proc. of Symposium on VLSI Circuits*, 2019, pp. C28–C29.

[31] J. T. Pawlowski, "Hybrid memory cube (HMC)," in *Proc. of IEEE Hot Chips 23 Symposium (HCS)*, 2011.

[32] A. M. Cabrera, C. J. Faber, K. Cepeda, R. Derber, C. Epstein, J. Zheng, R. K. Cytron, and R. D. Chamberlain, "DIBS: A data integration benchmark suite," in *Proc. of ACM/SPIE Int'l Conf. on Performance Engineering Companion*, Apr. 2018, pp. 25–28.

[33] H. Dong, P. Shum, M. Yan, J. Zhou, G. Ning, Y. Gong, and C. Wu, "Measurement of Mueller matrix for an optical fiber system with birefringence and polarization-dependent loss or gain," *Optics Communications*, vol. 274, no. 1, pp. 116–123, 2007.

[34] A. Taranta, E. N. Fokoua, S. A. Mousavi, J. R. Hayes, T. D. Bradley, G. T. Jasion, and F. Poletti, "Exceptional polarization purity in antiresonant hollow-core optical fibres," *Nature Photonics*, May 2020.