40 Gb/s Data Transmission Over a 1-m-Long Multimode Polymer Spiral Waveguide for Board-Level Optical Interconnects

(article starts on next page)

40 Gb/s Data Transmission Over a 1-m-Long Multimode Polymer Spiral Waveguide for Board-Level Optical Interconnects

Nikolaos Bamiedakis, Jian Chen, Petter Westbergh, Johan S. Gustavsson, Anders Larsson, *Fellow, IEEE*, Richard V. Penty, *Senior Member, IEEE*, and Ian H. White, *Fellow, IEEE*

Abstract—**Optical interconnects have attracted considerable attention for use in short-reach communication links within high-performance electronic systems, such as data centers, supercomputers, and data storage systems. Multimode polymer waveguides, in particular, constitute an attractive technology for use in board-level interconnects as they can be cost-effectively integrated onto standard PCBs and allow system assembly with relaxed alignment tolerances. However, their highly multimoded nature raises important concerns about their bandwidth limitations and their potential to support very high on-board data rates. In this paper, we report record error-free (BER $< 10^{-12}$) 40 Gb/s data transmission over a 1-m-long multimode polymer spiral waveguide and present thorough studies on the waveguide bandwidth performance. The frequency response of the waveguide is investigated under a wide range of launch conditions and in the presence of input spatial offsets, which are expected to be highly-likely in real-world systems. A robust bandwidth performance is observed with a bandwidth-length product of at least 35 GHz×m for all launch conditions studied. The reported results clearly demonstrate the potential of this technology for use in board-level interconnects, and indicate that data rates of at least 40 Gb/s are feasible over waveguide lengths of 1 m.**

Index Terms—**Board-level optical interconnects, multimode waveguides, polymer waveguides, waveguide bandwidth.**

I. INTRODUCTION

OPTICAL technologies have attracted considerable interest in recent years for use in short-reach communication links. Conventional copper-based interconnection technologies struggle to keep up with the ever increasing demand for interconnection bandwidth in high-performance electronic systems, such as data centres, supercomputers and data storage systems [1]–[3]. The inherent disadvantages of copper-based interconnects when operating at high data rates, such as electromagnetic interference (EMI), limited bandwidth and large power consumption,

have led to the consideration of the use of optical technologies in high-speed board-level communication links [4]–[7]. Optics can offer large bandwidths, immunity to EMI, reduced power consumption and relaxed thermal management requirements. Active optical cables have already successfully penetrated in rack-to-rack interconnection links (lengths 1–10 m) replacing copper cables and are also increasingly being used inside the rack [8], [9]. For shorter communication links (<1 m) various optical technologies have been proposed and are currently being developed, including fibre [10]–[12] and planar waveguide [13]–[17] technologies, free-space optical interconnects [18], [19], plasmonics [20]–[22] and Si photonics [23], [24]. However, important challenges for their adoption in real-world systems still remain including the cost-efficient integration of optics and electronics onto standard printed circuit boards (PCBs) and within electronic modules, and the development of appropriate system architectures.

Multimode polymer waveguides are a promising candidate for use in board-level interconnections as they can be directly integrated onto PCBs and offer relaxed alignment tolerances in the system assembly owing to their relatively large dimensions [25]–[27]. Various waveguide-based system demonstrators have been reported in recent years featuring large number of parallel polymer waveguides and achieving high aggregate data capacities of ∼1 Tb/s [3], [28], [29]. The typical dimensions of the waveguides employed in the demonstrated optical interconnection systems are in the range of 30 to 70 μm, while the refractive index step Δn used is ∼0.02–0.03. These values enable 1 dB alignment tolerances of ∼ $\pm$10 μm and good compatibility with standard multimode fibre (MMF) patchcords and vertical-cavity surface-emitting lasers (VCSELs). VCSELs are the main sources of interest for such applications since they are sufficiently low-cost, can be formed in large arrays, exhibit large bandwidth and low threshold current and don't require any temperature control for reliable operation [30]. In recent years, there has been a continuous improvement in their high-speed performance with most recent reports demonstrating up to 64 Gb/s direct modulation operation [30]–[33]. These improvements however, in conjunction with the highly-multimoded nature of the polymer waveguides, raise important questions about the bandwidth limits of this technology for high-speed board-level optical interconnects and its capability to support high data rates over distances beyond few tens of centimetres. Moreover, the technological requirement for cost-efficient assembly and connectorisation of opto-electronic (OE) PCBs further generates concerns over the robustness of the waveguide bandwidth performance as large spatial offsets are expected to

be highly-likely in real-world optically-interconnected systems. In order to address these concerns, the use of single mode waveguides for such applications has been proposed [34]–[36]. Single-mode waveguide technology can offer important performance advantages over multimode waveguides, namely large bandwidth and accurate mode control, but does not match well the application requirements (relaxed fabrication and alignment tolerances, cost-effective assembly and packaging, compatibility with low-cost sources such as MM VCSELs).

Some studies on the bandwidth performance of multimode polymer waveguides have been published [37]–[40], but these have been focussed on restricted launch conditions [37], [38], [40] or on graded-index waveguides [39]. Moreover, limited number of data transmission experiments on relatively long (≥ 1 m) multimode polymer waveguides have been reported [27], [39], [41]–[43], with demonstrated data rates being up to 12.5 Gb/s. We have recently presented thorough bandwidth studies on a 1.4 m long polymer multimode spiral waveguide and reported a bandwidth-length product of at least 25 GHz×m, even under overfilled launch conditions [44]. Moreover, we have demonstrated error-free (BER $< 10^{-12}$) 25 Gb/s data transmission over this relatively long waveguide [45]. The high-speed performance of this particular waveguide link was however power-limited rather than bandwidth-limited due to the non-optimised loss performance of the spiral waveguide structure. In this paper, we utilise a new 1 m long polymer multimode spiral waveguide with improved loss characteristics and demonstrate record error-free (BER $< 10^{-12}$) 40 Gb/s data transmission over such a waveguide length. The bandwidth performance of this long polymer multimode waveguide is studied under various launch conditions and in the presence of spatial input offsets. The obtained results indicate a bandwidth-length product of at least 35 GHz×m for all launch conditions studied and a robust waveguide performance in the presence of input offsets.

The results demonstrate the potential of this technology for use in board-level optical interconnects at data rates of 40 Gb/s or higher and address concerns about its suitability for use in next-generation high-performance electronic systems. The improvement achieved over the previously published result of 25 Gb/s in [44] is significant (1.6 times better) indicating that large aggregate capacity interconnections (>1 Tb/s) could be achieved with a significantly reduced number of links, OE components and required board area. Moreover, the reported bandwidth measurements demonstrate that there is enough bandwidth to support 40 Gb/s or even higher data rates over such multimode waveguides and suggest that power budget, rather than bandwidth, becomes the important parameter for the implementation of such high-speed links in real-world systems. In the following sections, the polymer multimode spiral waveguide employed in this work is presented in Section II, the bandwidth studies carried out on the waveguide are described in Section III and the data transmission experiments are reported in Section IV. Section V briefly discusses the theoretical bandwidth limit of such multimode waveguides, while Section VI provides the conclusions.

II. 1 M LONG MULTIMODE POLYMER SPIRAL WAVEGUIDE

The spiral waveguide employed in this work is fabricated from silicone polymer materials developed by Dow Corning [Dow


Fig. 1. Images of (a) the 1 m long spiral waveguide illuminated with red light and (b) the waveguide output facet illuminated with 850 nm light.

Corning WG-1020 Optical Waveguide Core and Dow Corning OE-4141 Cured Optical Elastomer (cladding)]. These polymer materials have been appropriately engineered to fit the application requirements and allow direct integration of the waveguides onto standard PCBs [13]. The materials can withstand the temperatures in excess of 250 °C required for solder reflow processes and exhibit low loss of ∼0.04 dB/cm at the 850 nm data communications wavelength [13], [46], [47]. The refractive index of the core and cladding materials is ∼1.52 and 1.50 respectively at 850 nm. The spiral waveguide structure has a length of 105.5 cm and is fabricated on an 8-in silicon wafer with conventional photolithography [48]. The waveguide core has a cross section of approximately $32 \times 50 \ \mu m^2$ while the waveguide facets are exposed with a dicing saw. No polishing steps are carried out to improve the quality of the waveguide facets. Fig. 1 shows photographs of the spiral waveguide illuminated with red light and the waveguide output facet.

III. BANDWIDTH STUDIES

The bandwidth performance of the 1 m long spiral waveguide is assessed under varying launch conditions and in the presence of spatial input offsets. In order to obtain the waveguide frequency response, the frequency response of the optical link with and without (i.e., back-to-back) the waveguide is recorded for each launch condition and input position. By subtracting the two obtained responses, the frequency response of the waveguide can be extracted.

A. Input Launch Conditions and Experimental Setup

The launch conditions studied cover a wide range of input conditions likely to be encountered in real-world systems and range from restricted to practically overfilled launches: (i) 4/125 μm SMF, (ii) "typical" 50/125 μm MMF [numerical aperture (NA) of 0.2], (iii) a "quasi-overfilled" 50/125 μm MMF, (iv) a 100/140 μm MMF (NA of 0.29). In general, overfilled launches are expected to couple larger percentage of optical power to higher order modes at the waveguide input and therefore, result in increased multimode dispersion and reduced waveguide bandwidth. Restricted launch conditions typically exhibit larger bandwidths but can lead to significant performance variations for different input positions. Spatial input offsets result in the excitation of different mode groups at the waveguide input and therefore, varying levels of multimode dispersion can be obtained

Fig. 2. Experimental setup used for the S_{21} measurement of the optical link (a) with and (b) without the spiral waveguide. For launch conditions (iii) and (iv), the mode mixer is inserted in the link.

Fig. 3. (a) Far-field intensity profile and (b–e) near-field image of the input fibre for the different launch conditions studied: (b) 4 μm SMF, (c) "typical" 50 μm MMF, (d) "quasi-overfilled" 50 μm MMF and (e) 100 μm MMF.

at the waveguide output. As a result, for each launch condition, the waveguide frequency response is also studied for different positions of the input fibre in order to assess the robustness of the waveguide bandwidth performance. Although the results presented in the following sections are obtained for horizontal input offsets, similar plots and behaviour have also been obtained for vertical offsets.

More specifically, the 4/125 μm SMF input is employed at waveguide input to generate a restricted launch condition into the waveguide. The small size of the fibre core ensures real single-mode fibre operation at 850 nm and results in the excitation of small number of modes at the waveguide input. The "typical" 50/125 μm MMF launch consists of coupling light from the VCSEL source into the waveguide via a 50/125 μm MMF patch-cord without the use of any launch conditioning elements. Such a launch is expected to be encountered in real-world systems and in short MMF-waveguide interfaces and connectors. The "quasi-overfilled" 50/125 μm MMF launch is achieved using a mode mixer that generates a more uniform power distribution inside the input fibre. This launch can provide a useful insight in the waveguide behaviour when a larger volume of modes is excited at its input. Finally, the use of the 100/140 μm MMF generates a practically overfilled launch condition at the waveguide input and can therefore be considered to be a "worst-case" launch condition into the waveguide as far as multimode dispersion is concerned.

The basic link setup used in the S_{21} measurements is shown in Fig. 2. An 850 nm VCSEL with a bandwidth of ∼25 GHz [32] is employed as the light source while a fibre-coupled 30 GHz photodiode (PD) (VI D30-850M) is used as the receiver. A cleaved fibre [9/125 μm SMF for launch condition (i), 50/125 μm MMF for launch conditions (ii) and (iii), or 100/140 μm MMF for launch condition (iv)] is used to collect the VCSEL emitted light. Another cleaved fibre of the appropriate type is employed to couple the light into the spiral waveguide. For the overfilled launch conditions (iii) and (iv), a mode mixer (Newport FM-1) is used before the waveguide. The far-field intensity profile of the input fibre is measured and the near-field profile at the fibre end is recorded for each launch condition prior to the S_{21} measurements (Fig. 3). The cleaved end of the input fibre is positioned on a precision translational stage and a displacement sensor is used to control the induced input spatial offsets. At the waveguide output, a cleaved 50/125 μm MMF (NA of 0.2) is used to collect the transmitted light and deliver it to the PD. Index-matching gel is employed at both input and output waveguide facets to minimise Fresnel losses and scattering losses due

to facet surface roughness. A vector network analyzer (Agilent 8722ET) is used to measure the S_{21} frequency response of the waveguide and back-to-back link. The RF signal is fed to the VCSEL via a high-bandwidth bias tee and a 40 GHz RF probe, while the electrical signal received at the PD end is amplified with a 40 GHz RF amplifier. A multimode variable optical attenuator (MM VOA, Agilent NA7766A) is inserted in the back-to-back link in order to adjust the received optical power to similar levels as the ones obtained for the respective waveguide link. The total insertion loss of the 1 m long spiral waveguide is measured to be 8.2, 9.3, 10 and 12.5 dB for the 4/125 μm SMF, "typical" 50/125 μm MMF, "quasi-overfilled" 50/125 μm MMF and 100/140 μm MMF input respectively. This includes the input and output coupling loss and the propagation loss of the spiral waveguide. The loss performance of this spiral waveguide is significantly improved over that of the 1.4 m long spiral waveguide employed in [44] (∼6 dB lower) offering therefore a larger power budget in the link for high-speed data transmission.

It should be noted that the use of the 50/125 μm MMF (NA of 0.2) at the waveguide output is not ideal for the bandwidth measurements as the fibre has a slightly smaller NA than that of the waveguide (NA ∼0.25). Its use in the experiments however cannot be avoided as the high-bandwidth PD has a similar (50/125 μm OM3 MMF) fibre-coupled input. This can result in the suppression of the power received at the waveguide output from higher-order modes and therefore, in reduced observed dispersion. In order to assess this effect on the measurements, the optical power received with the 50/125 μm MMF (NA: 0.2) at the waveguide output is compared with that obtained when an output fibre with larger dimensions and NA is used (100/140 μm MMF with an NA of 0.29). The input type is chosen to be the "quasi-overfilled" 50/125 μm MMF and the 100/140 μm MMF inputs, as these launch conditions are relatively overfilled and therefore, ensure that considerable optical power is coupled to higher-order modes at the waveguide input. The power received at the waveguide output is recorded for varying positions of the input fibre for both types of output fibre and the obtained normalised plots are compared (Fig. 4). Any significant difference between the two should indicate high mode selective loss at the waveguide output due to the use of the 50 μm MMF with the smaller NA.

Fig. 4 indicates a very small difference between the power variation plots obtained for the two fibres employed at the waveguide output and for both launch conditions studied. The results suggest that the mode selective loss due to the use of the 50/125 μm MMF at the waveguide output is very low and therefore, it should not have an important effect on the dispersion measurements. It is believed that the spiral structure of the
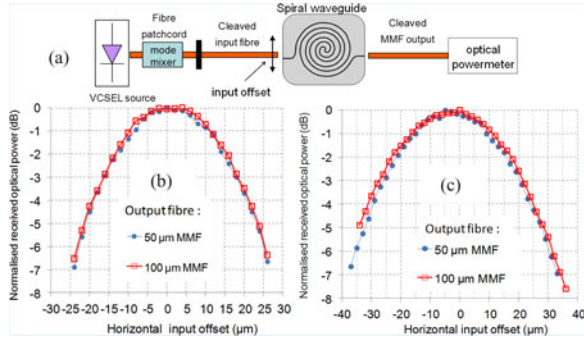
Fig. 4. (a) Measurement setup and normalized received optical power as a function of the input position for a 50 μm MMF and 100 μm MMF output fibre for a (a) "quasi-overfilled" 50/125 μm and (b) a 100/140 μm MMF input.

waveguide suppresses the very high order modes that would not couple efficiently to the 50/125 μm MMF output fibre and therefore the effect is minimal. Moreover, the refractive index profile of the waveguide exhibits some graded variation near the waveguide edges. As a result, the fabricated waveguide is believed to have an NA smaller than expected and support fewer waveguide modes than the respective step-index waveguide with identical dimensions. Similar observations have been made on other waveguide samples fabricated from the same polymer materials and it has been found that this refractive index variation depends on the fabrication conditions [49]. A more thorough analysis is underway to characterise the obtained refractive index profiles and assess their effect on the light guiding properties of the fabricated waveguides. Accurate control of the waveguide refractive index profile by adjusting the fabrication parameters can enable dispersion engineering according to the application requirements.

B. Waveguide Frequency Response

The obtained waveguide frequency responses for the different launch conditions and input positions are shown in Fig. 5. The 4/125 μm SMF launch yields a flat frequency response for a well-aligned input up to the 35 GHz instrumentation limit. Spatial offsets result in a slight degradation of the bandwidth performance due to the excitation of higher order modes at the waveguide input, but no significant transmission impairments are observed up to the 35 GHz instrumentation limit. More overfilled launch conditions result in a reduced (less flat) frequency response as expected, due to the coupling of larger optical power to higher order modes at the waveguide input. In all cases however, the 3 dB frequency is beyond the instrumentation limit of 35 GHz, yielding a bandwidth-length product of at least 35 GHz×m for all launch conditions studied and in the presence of input offsets. The observed waveguide bandwidth performance indicates that the transmission of data rates of 40 Gb/s or higher over such long waveguide structures is possible without any significant transmission impairments.

When comparing the frequency response plots in Fig. 5 with the ones obtained from the slightly longer, but more lossy, spiral waveguide reported in [44], it can be noticed that the curves shown here are less flat. This indicates that the spiral waveguide studied here exhibits slightly larger multimode dispersion than the older non-optimised spiral waveguide sample. The observed difference can potentially be justified by the different
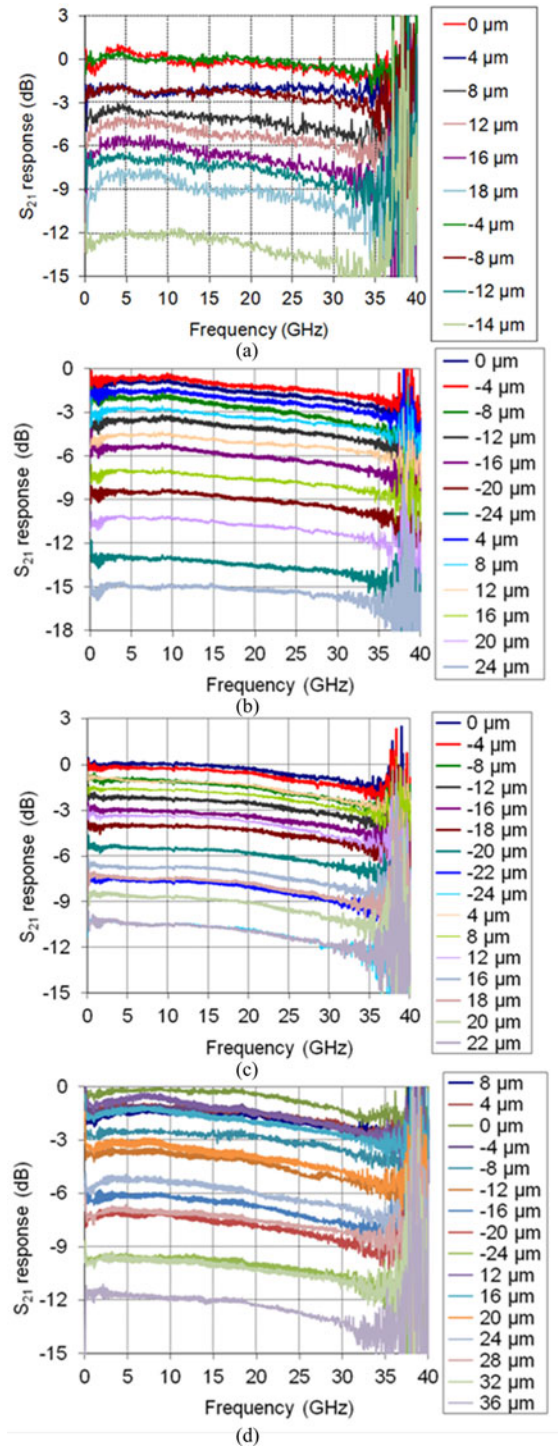

Fig. 5. Waveguide frequency response for different input positions for (a) the 4/125 μm SMF, (b) the "typical" 50/125 μm MMF, (c) the "quasi-overfilled" 50/125 μm MMF and (d) the 100/140 μm MMF input.

spiral structure and bending loss performance of the two waveguide samples. The older spiral sample exhibits higher bending losses than the newer optimised spiral employed here, and as a result, higher-order modes may get more suppressed in the former than in the latter as they propagate along the structure. Consequently, slightly increased multimode dispersion would be expected in the newer spiral waveguide, justifying therefore the observed difference in the shape of the obtained frequency

Fig. 6.   Near field images of the waveguide output for different input positions and for the different launch conditions studied.


Fig. 7.   Normalised far-field intensity profile at waveguide output for different types of input fibres and input offsets: (a) $\Delta x = 0$ μm and (b) $\Delta x = +8$ μm.

Fig. 8.   Normalised far-field profiles at waveguide output for different input positions and fibre type: (a) "typical" 50 μm MMF and (b) 100 μm MMF input.

Fig. 9.   Experimental setup of (a) the waveguide and (b) the back-to-back link.

responses. This observation illustrates an important trade-off in waveguide layout design as the use of relatively lossy waveguide bends might offer improved bandwidth performance at the cost of some additional optical loss.

C. Near Field Images and Far-Field Profiles

Near field images and far-field profiles of the waveguide output are recorded for the different launch conditions and for varying input offsets. These provide some insight on the mode propagation inside the waveguide. Samples of the obtained near-field images are shown in Fig. 6. It can be observed that for the 4/125 μm SMF input, despite being a highly-restricted launch, a significant number of mode groups are noticed at the waveguide output. This can be attributed to the spiral structure and long length of the waveguide and mode mixing inside the waveguide due to surface roughness and local imperfections. As expected, for the other types of MMF input with increasing overfilled launch conditions, a more uniform output power distribution is observed at the waveguide output and a reduced speckle contrast. Although spatial input offsets result in some power redistribution at the waveguide output, the changes observed in the images are not significant, justifying therefore the similarities in the shape of the frequency responses obtained across input offsets.

Far-field profiles of the waveguide output are also recorded and a few plots obtained for the different input fibres and launch positions are shown in Figs. 7 and 8. The obtained results agree with the observations on mode power distribution inside the waveguide made above. The SMF input results in the narrower far-field beam profile at the waveguide output, indicating therefore the excitation of a small group of modes at the waveguide input, while the 100 μm MMF input yields the widest profile (Fig. 7). Moreover, no significant changes in the far-field profile are observed at the waveguide output for the different positions

of the input fibre when MMF launches are employed at the waveguide input (Fig. 8). The 5% intensity value (-13 dB) is noted in the plots for clarity.

IV. HIGH-SPEED DATA TRANSMISSION EXPERIMENTS

The setup used for the data transmission experiments is shown in Fig. 9. The VNA is replaced with a bit-error-rate (BER) test set (Anritsu MP18000A) and a digital communication analyser (Agilent Infiniium 86100A). The VCSEL is directly modulated by a short $2^7 - 1$ pseudorandom bit sequence in order to emulate the short run codes typically employed in data communication links (e.g., 8B10B). A pair of $16\times$ microscope objectives (NA of 0.32) is used to couple the light into the waveguide to minimise coupling losses, while the 50/125 μm MMF is used at the waveguide output. The total insertion loss of the 1 m long spiral waveguide [power difference between points A and B in [Fig. 9(a)] is measured to be 9 dB. The respective back-to-back link without the spiral waveguide is also setup and tested [Fig. 9(b)]. The MM VOA is used to adjust the average optical power level at the receiver end and introduce optical loss in the back-to-back link which is comparable to the insertion loss of the waveguide.

The data transmission experiments are carried out at 25, 36 and 40 Gb/s. The VCSEL operating conditions (current bias, modulation amplitude) are slightly adjusted for each data rate in order to obtain optimum link performance. The received eye diagrams for the waveguide and back-to-back link are shown in Fig. 10. Open eye diagrams are obtained for all data rates, while minimum additional dispersion and noise can be observed in the link due to the insertion of the waveguide. The relative eye closure observed in the recorded eye diagrams at 36 and 40 Gb/s is due to the bandwidth limitation of the active devices in the link and can be noticed in both the waveguide and back-to-back link.

BER measurements are also carried out on both optical links. For all data rates, error-free (BER $< 10^{-12}$) data transmission is achieved over the waveguide link. The obtained BER curves as a function of the average received optical power are shown in
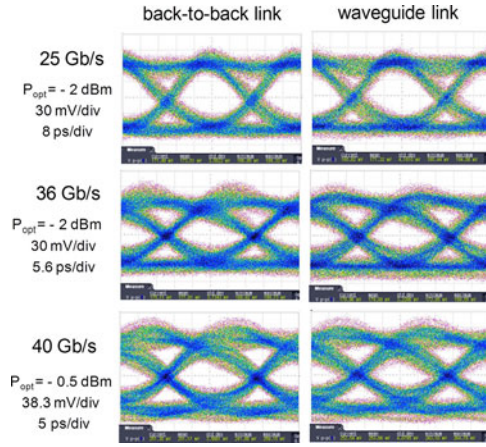
Fig. 10. Received eye diagrams for the back-to-back and waveguide links at all data rates studied. Average received optical power level (P_{opt}) noted for each data rate as well as the voltage and time scale of the recorded waveforms.
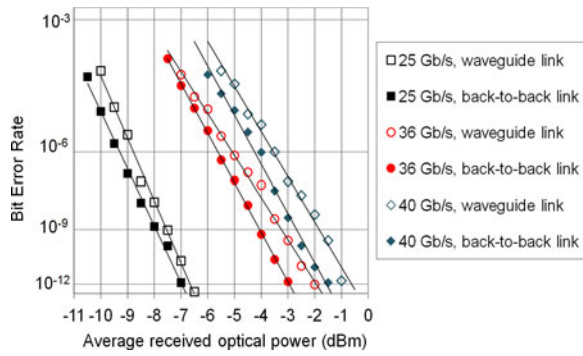

Fig. 11. BER curves for the back-to-back and waveguide link obtained at 25, 36 and 40 Gb/s.

Fig. 11 for both the waveguide and back-to-back link. The power penalty for a BER of 10^{-9} due to the insertion of the waveguide in link is found to be ~0.5 dB for 25 Gb/s data transmission and ~0.85 dB for 36 and 40 Gb/s. The results clearly demonstrate the potential to use this technology for use in board-level optical interconnects at data rates of 40 Gb/s, despite the multimode nature of the waveguide and its relatively long length.

V. DISCUSSION

The results presented here demonstrate the potential of multimode polymer waveguides to support data rates ≥40 Gb/s over distances of 1 m, but raise the interesting question of the ultimate bandwidth limit of these multimode waveguides. Basic modelling studies have been carried out to study the bandwidth of such multimode polymer waveguides. The values obtained however through simulation can vary significantly depending on the employed waveguide parameters (dimensions, refractive index profile), assumed launch conditions and magnitude of important transmission phenomena taken into consideration in the model (mode loss due to surface roughness and waveguide structure, mode mixing etc.). Assuming the simplest "ideal" waveguide (uniform mode loss profile, no mode mixing) with a step index profile and typical waveguide parameters (50×50 μm^2, $\Delta n = 0.02$) yields bandwidth-length product values that range from ~10 GHz×m for an overfilled

launch to ~150 GHz×m for a highly-restricted launch. The large bandwidth-length product values obtained for restricted launches have been experimentally verified as most of the studies reported have employed such launches (e.g., SMF inputs) [37], [38], [40]. However, limited experimental work on waveguide bandwidth has been reported with MMF or more overfilled launches, and these experiments show much larger bandwidth-length products than simulation would predict [39], [44]. As mentioned above, variations in the shape of the refractive index profile and waveguide structures such as long large-radius bends, such as in the spiral waveguide employed here, can have a beneficial effect on the obtained waveguide bandwidth. Mode mixing can also significantly affect the observed bandwidth performance [37]. A basic waveguide model is therefore inadequate to properly estimate the waveguide bandwidth. A detailed study is underway to correlate the experimentally-observed bandwidth results with an accurate waveguide model.

VI. CONCLUSION

Multimode polymer waveguides have attracted considerable attention for use in board-level optical interconnections as they allow direct integration onto low-cost PCBs and enable cost-effective system assembly. The relatively waveguide large dimensions (width: 30 to 70 μm) and numerical aperture (NA: 0.2–0.3) typically employed in such OE PCBs, in conjunction with the continuous improvement in bandwidth performance of VCSELs, raise however important questions on the bandwidth limitations of this technology and its capability to support very high data rates. Bandwidth studies on a 1 m long multimode spiral waveguide demonstrate a bandwidth-length product of at least 35 GHz×m for even overfilled launch conditions and a robust bandwidth performance in the presence of input offsets. Record error-free (BER $< 10^{-12}$) 40 Gb/s data transmission is demonstrated over this 1 m long polymer waveguide with a small power penalty of 0.85 dB for a 10^{-9} BER. The reported results address the concerns about the potential of this technology for use in very high speed board-level interconnects and clearly demonstrate their capability to support data rates of at least 40 Gb/s over distances of 1 m.

ACKNOWLEDGMENT

REFERENCES

[1] A. F. Benner, M. Ignatowski, J. A. Kash, D. M. Kuchta, and M. B. Ritter, "Exploitation of optical interconnects in future server architectures," *IBM J. Res. Dev.*, vol. 49, pp. 755–775, 2005.

[2] M. A. Taubenblatt, "Optical interconnects for high-performance computing," *J. Lightw. Technol.*, vol. 30, no. 4, pp. 448–457, Feb. 2012.

[3] K. Schmidtke, F. Flens, A. Worrall, R. Pitwon, F. Betschon, T. Lamprecht, and R. Krahenbuhl, "960 Gb/s optical backplane ecosystem using embedded polymer waveguides and demonstration in a 12G SAS storage array," *J. Lightw. Technol.*, vol. 31, no. 24, pp. 3970–3975, Dec. 2013.

[4] D. A. B. Miller, "Rationale and challenges for optical interconnects to electronic chips," *Proc. IEEE*, vol. 88, no. 6, pp. 728–749, Jun. 2000.

[5] D. W. Huang, T. Sze, A. Landin, R. Lytel, and H. L. Davidson, "Optical interconnects: Out of the box forever?," *IEEE J. Sel. Topics Quantum Electron.*, vol. 9, no. 2, pp. 614–623, Mar./Apr. 2003.

[6] S. Borkar, "Role of interconnects in the future of computing," *J. Lightw. Technol.*, vol. 31, no. 24, pp. 3927–3933, Dec. 2013.

[7] C. Hoyeol, P. Kapur, and K. C. Saraswat, "Power comparison between high-speed electrical and optical interconnects for interchip communication," *J. Lightw. Technol.*, vol. 22, no. 9, pp. 2021–2033, Sep. 2004.

[8] A. F. Benner *et al.*, "Optics for high-performance servers and supercomputers," in *Proc. Opt. Fiber Commun. Conf.*, 2010, pp. 1–3.

[9] J. A. Kash *et al.*, "Optical interconnects in future servers," in *Proc. Opt. Fiber Commun. Conf.*, 2011, pp. 1–3.

[10] I.-K. Cho, J.-H. Ryu, and M.-Y. Jeong, "Interchip link system using an optical wiring method," *Opt. Lett.*, vol. 33, pp. 1881–1883, 2008.

[11] R. Pitwon *et al.*, "Demonstration of fully enabled data center subsystem with embedded optical interconnect," *Proc. SPIE, Opt. Interconnects XIV*, vol. 8991, pp. 1–11, 2014.

[12] J. Matsui *et al.*, "High bandwidth optical interconnection for densely integrated server," in *Proc. Opt. Fiber Commun. Conf.*, 2013, pp. 1–3.

[13] R. Dangel, F. Horst, D. Jubin, N. Meier, J. Weiss, B. J. Offrein, B. W. Swatowski, C. M. Amb, D. J. DeShazer, and W. K. Weidner, "Development of versatile polymer waveguide flex technology for use in optical interconnects," *J. Lightw. Technol.*, vol. 31, no. 24, pp. 3915–3926, Nov. 2013.

[14] M. Immonen, M. Karppinen, and J. K. Kivilahti, "Fabrication and characterization of polymer optical waveguides with integrated micromirrors for three-dimensional board-level optical interconnects," *IEEE Trans. Electron. Packag. Manuf.*, vol. 28, no. 4, pp. 304–311, Oct. 2005.

[15] R. C. A. Pitwon, W. Kai, J. Graham-Jones, I. Papakonstantinou, H. Baghsiahi, B. J. Offrein, R. Dangel, D. Milward, and D. R. Selviah, "FirstLight: Pluggable optical interconnect technologies for polymeric electro-optical printed circuit boards in data centers," *J. Lightw. Technol.*, vol. 30, no. 21, pp. 3316–3329, Nov. 2012.

[16] N. Bamiedakis, A. Hashim, R. V. Penty, and I. H. White, "A 40 Gb/s optical bus for optical backplane interconnections," *J. Lightw. Technol.*, vol. 32, no. 8, pp. 1526–1537, Apr. 2014.

[17] R. Kinoshita, K. Moriya, K. Choki, and T. Ishigure, "Polymer optical waveguides with GI and W-shaped cores for high-bandwidth-density on-board interconnects," *J. Lightw. Technol.*, vol. 31, no. 24, pp. 4004–4015, Nov. 2013.

[18] K. Wang, A. Nirmalathas, C. Lim, E. Skafidas, and K. Alameh, "High-speed free-space based reconfigurable card-to-card optical interconnects with broadcast capability," *Opt. Exp.*, vol. 21, pp. 15395–15400, 2013.

[19] M. Jarczynski, T. Seiler, and J. Jahns, "Integrated three-dimensional optical multilayer using free-space optics," *Appl. Opt.*, vol. 45, pp. 6335–6341, 2006.

[20] S. Papaioannou, K. Vyrsokinos, O. Tsilipakos, A. Pitilakis, K. Hassan, J. Weeber, L. Markey, A. Dereux, S. I. Bozhevolnyi, A. Miliou, E. E. Kriezis, and N. Pleros, "A 320 Gb/s-throughput capable 2×2 silicon-plasmonic router architecture for optical interconnects," *J. Lightw. Technol.*, vol. 29, no. 21, pp. 3185–3195, Nov. 2011.

[21] M. L. Brongersma, R. Zia, and J. A. Schuller, "Plasmonics—The missing link between nanoelectronics and microphotonics," *Appl. Phys. A, Mater. Sci. Process.*, vol. 89, pp. 221–223, 2007.

[22] J. A. Conway, S. Sahni, and T. Szkopek, "Plasmonic interconnects versus conventional interconnects: A comparison of latency, crosstalk and energy costs," *Opt. Exp.*, vol. 15, pp. 4474–4484, 2007.

[23] Y. Vlasov, W. M. J. Green, and F. Xia, "High-throughput silicon nanophotonic wavelength-insensitive switch for on-chip optical networks," *Nature Photon.*, vol. 2, pp. 242–246, 2008.

[24] M. Lipson, "Guiding, modulating, and emitting light on silicon-challenges and opportunities," *J. Lightw. Technol.*, vol. 23, no. 12, pp. 4222–4238, Dec. 2005.

[25] N. Bamiedakis, A. Hashim, J. Beals, R. V. Penty, and I. H. White, "Low-cost PCB-integrated 10-Gb/s optical transceiver built with a novel integration method," *IEEE Trans. Compon, Packag. Manuf. Technol.*, vol. 3, no. 4, pp. 592–600, Apr. 2013.

[26] I. Papakonstantinou, D. R. Selviah, R. Pitwon, and D. Milward, "Low-cost, precision, self-alignment technique for coupling laser and photodiode arrays to polymer waveguide arrays on multilayer PCBs," *IEEE Trans. Adv. Packag.*, vol. 31, no. 3, pp. 502–511, Aug. 2008.

[27] R. Dangel, C. Berger, R. Beyeler, L. Dellmann, M. Gmur, R. Hamelin, F. Horst, T. Lamprecht, T. Morf, S. Oggioni, M. Spreafico, and B. J. Offrein, "Polymer-waveguide-based board-level optical interconnect technology for datacom applications," *IEEE Trans. Adv. Packag.*, vol. 31, no. 4, pp. 759–767, Nov. 2008.

[28] F. E. Doany, C. L. Schow, B. G. Lee, R. A. Budd, C. W. Baks, C. K. Tsang, J. U. Knickerbocker, R. Dangel, B. Chan, L. How, C. Carver, H. Jianzhuang, J. Berry, D. Bajkowski, F. Libsch, and J. A. Kash, "Terabit/s-class optical PCB links incorporating 360-Gb/s bidirectional 850 nm parallel optical transceivers," *J. Lightw. Technol.*, vol. 30, no. 4, pp. 560–571, Feb. 2012.

[29] J. Beals *et al.*, "A terabit capacity passive polymer optical backplane based on a novel meshed waveguide architecture," *Appl. Phys. A, Mater. Sci. Process.*, vol. 95, pp. 983–988, 2009.

[30] A. Larsson, "Advances in VCSELs for communication and sensing," *IEEE J. Sel. Topics Quantum Electron.*, vol. 17, no. 6, pp. 1552–1567, Nov./Dec. 2011.

[31] D. Kuchta *et al.*, "64 Gb/s Transmission over 57 m MMF using an NRZ modulated 850 nm VCSEL," in *Proc. Opt. Fiber Commun. Conf.*, 2014, pp. 1–3.

[32] P. Westbergh *et al.*, "High-speed 850 nm VCSELs with 28 GHz modulation bandwidth operating error-free up to 44 Gbit/s," *Electron. Lett.*, vol. 48, pp. 1145–1147, 2012.

[33] D. M. Kuchta *et al.*, "A 56.1 Gb/s NRZ modulated 850 nm VCSEL-based optical link," in *Proc. Opt. Fiber Commun. Conf.*, 2013, pp. 1–3.

[34] E. Zgraggen, I. M. Soganci, F. Horst, A. La Porta, R. Dangel, B. J. Offrein, S. A. Snow, J. K. Young, B. W. Swatowski, C. M. Amb, O. Scholder, R. Broennimann, U. Sennhauser, and G. L. Bona, "Laser direct writing of single-mode polysiloxane optical waveguides and devices," *J. Lightw. Technol.*, vol. 32, no. 17, pp. 3036–3042, Sep. 2014.

[35] A. Boersma *et al.*, "Polymer-based optical interconnects using nanoimprint lithography," *Proc. SPIE, Optoelectronic Interconnects XIII*, vol. 8630, pp. 1–9, 2013.

[36] A. Sugama, K. Kawaguchi, M. Nishizawa, H. Muranaka, and Y. Arakawa, "Development of high-density single-mode polymer waveguides with low crosstalk for chip-to-chip optical interconnection," *Opt. Exp.*, vol. 21, pp. 24231–24239, 2013.

[37] F. E. Doany, P. K. Pepeljugoski, A. C. Lehman, J. A. Kash, and R. Dangel, "Measurement of optical dispersion in multimode polymer waveguides," in *Proc. LEOS Summer Top. Meet.*, 2004, pp. 31–32.

[38] P. Pepeljugoski *et al.*, "Comparison of bandwidth limits for on-card electrical and optical interconnects for 100 Gb/s and beyond," *Proc. SPIE, Optoelectronic Integr. Circuits X*, vol. 6897, pp. 1–7, 2008.

[39] T. Kosugi and T. Ishigure, "Polymer parallel optical waveguide with graded-index rectangular cores and its dispersion analysis," *Opt. Exp.*, vol. 17, pp. 15959–15968, 2009.

[40] W. Xiaolong, W. Li, J. Wei, and R. T. Chen, "Hard-molded 51 cm long waveguide array with a 150 GHz bandwidth for board-level optical interconnects," *Opt. Lett.*, vol. 32, pp. 677–679, 2007.

[41] F. Mederer *et al.*, "10 Gb/s data transmission with TO-packaged multimode GaAsVCSELs over 1 m long polymer waveguides for optical backplane applications," *Opt. Commun.*, vol. 206, pp. 309–312, 2002.

[42] Y. Takeyoshi and T. Ishigure, "High-density 2×4 channel polymer optical waveguide with graded-index circular cores," *J. Lightw. Technol.*, vol. 27, no. 14, pp. 2852–2861, Jul. 2009.

[43] S. Takenobu and T. Okazoe, "Heat resistant and low-loss fluorinated polymer optical waveguides at 1310/1550 nm for optical interconnects," in *Proc. Eur. Conf. Opt. Commun.*, 2011, pp. 1–3.

[44] N. Bamiedakis, J. Chen, R. Penty, and I. White, "Bandwidth studies on multimode polymer waveguides for ≥ 25 Gb/s optical interconnects," *IEEE Photon. Technol. Lett.*, vol. 26, no. 20, pp. 2004–2007, Sep. 2014.

[45] N. Bamiedakis, R. V. Penty, I. H. White, P. Westbergh, and A. Larsson, "25 Gb/s data transmission over a 1.4 m long multimode polymer spiral waveguide," presented at theProc. *Conf. Lasers Electro-Optics*, San Jose, CA, USA, 2014, Paper STu1G.7.

[46] J. V. DeGroot Jr., "Cost-effective optical waveguide components for printed circuit applications," *Proc. SPIE, Passive Compon. Fiber-Based Devices IV*, vol. 6781, pp. 1–12, 2007.

[47] B. W. Swatowski *et al.*, "Flexible, stable, and easily processable optical silicones for low loss polymer waveguides," *Proc. SPIE, Organic Photon. Mater. Devices XV*, vol. 8622, pp. 1–11, 2013.

[48] N. Bamiedakis, J. Beals, R. V. Penty, I. H. White, J. V. DeGroot, and T. V. Clapp, "Cost-effective multimode polymer waveguides for high-speed on-board optical interconnects," *IEEE J. Quantum Electron.*, vol. 45, no. 4, pp. 415–424, Mar. 2009.

[49] B. W. Swatowski, C. M. Amb, M. G. Hyer, R. S. John, and W. K. Weidner, "Graded index silicone waveguides for high performance computing," in *Proc. IEEE Opt. Interconnects Conf.*, 2014, pp. 1–3.

Authors' biographies not available at the time of publication.