

64 Gb/s low-voltage waveguide SiGe avalanche photodiodes with distributed Bragg reflectors

BINHAO WANG,*  ZHIHONG HUANG, YUAN YUAN,  DI LIANG, XIAOGE ZENG,  MARCO FIORENTINO,  AND RAYMOND G. BEAUSOLEIL

Hewlett Packard Laboratories, Hewlett Packard Enterprise, Palo Alto, California 94304, USA

*Corresponding author: binhao.wang@hpe.com

Received 13 February 2020; revised 11 April 2020; accepted 4 May 2020; posted 4 May 2020 (Doc. ID 390339); published 5 June 2020

We demonstrate low-voltage waveguide silicon-germanium avalanche photodiodes (APDs) integrated with distributed Bragg reflectors (DBRs). The internal quantum efficiency is improved from 60% to 90% at 1550 nm assisted with DBRs while still achieving a 25 GHz bandwidth. A low breakdown voltage of 10 V and a gain bandwidth product of near 500 GHz are obtained. APDs with DBRs at a data rate of 64 Gb/s pulse amplitude modulation with four levels (PAM₄) show a 30%–40% increase in optical modulation amplitude (OMA) compared to APDs with no DBR. A sensitivity of around -13 dBm at a data rate of 64 Gb/s PAM₄ and a bit error rate of 2.4×10^{-4} is realized for APDs with DBRs, which improves the sensitivity by ~ 2 dB compared to APDs with no DBR. © 2020 Chinese Laser Press

<https://doi.org/10.1364/PRJ.390339>

1. INTRODUCTION

Data communications in data centers and high-performance computing (HPC) have grown tremendously due to emerging applications in social media, video streaming, artificial intelligence (AI), and the internet of things. Hyperscale data centers and Exascale HPC require high-bandwidth and energy-efficient optical interconnects [1,2]. In this context, silicon photonic interconnects are attractive, thanks to the high integration and low cost [3–5]. In particular, the monolithic integration of silicon photonic devices with complementary metal-oxide-semiconductor (CMOS) circuitry has shown that this technology is promising [6,7]. At the same time, to double the data rate, the advanced four-level pulse amplitude modulation (PAM₄) format has become the industry standard for 400G Ethernet and other next-generation optical interconnects [8]. When implementing PAM₄ in an optical link, one has to consider that, due to the transmission of two bits per clock period, the optical modulation amplitude (OMA) of PAM₄ signaling is only one third of on-off keying (OOK) modulation. Therefore, PAM₄ requires either a higher-power laser or a higher-sensitivity receiver to obtain the same bit error rate (BER) in an optical link. Using a receiver with better sensitivity can yield a lower link total power consumption compared to using a high-power laser and consequently improve the energy efficiency [9]. Particularly, high-sensitivity detectors relax the link budget requirements for on-chip lasers with limited output power [10].

An avalanche photodiode (APD) with internal gain is the ideal candidate to increase the receiver sensitivity [11].

Compared to most III–V compound devices, silicon-germanium (SiGe) APDs have lower noise and higher bandwidth due to the low impact ionization coefficient ratio in silicon ($k = 0.02$) [12–15]. APDs with a low breakdown voltage have attracted our interests due to the requirement of less than 12 V power rails in current computer architectures [16]. Although Ge APDs with a breakdown voltage of less than 10 V (and even down to 5 V) have been demonstrated [17,18], the large impact ionization coefficient ratio of Ge ($k = \sim 1$) results in large excess noise in Ge APDs. To take advantage of the large absorption in Ge in the near-infrared (NIR) region and the low multiplication noise in Si, APDs with a separate absorption and charge multiplication (SACM) structure were proposed [12–15]. A 10 V breakdown voltage SiGe APD with a gain bandwidth product (GBP) of 276 GHz has been demonstrated [14,19]. Recently, a three-terminal SiGe APD with independently controlled electric fields in the absorption and multiplication regions has been demonstrated with an ultralow breakdown voltage of 6 V [20]. Due to the small footprint and the decoupling of quantum efficiency and carrier transit time, waveguide APDs can usually achieve higher quantum efficiency and bandwidth compared to normal incident APDs [14,15,19,20]. Moreover, waveguide APDs can be integrated in complex photonic integrated circuits (PICs) such as wavelength division multiplexers (WDMs) for many applications. However, in the design of a waveguide APD, a trade-off is required among quantum efficiency, bandwidth, and a low breakdown voltage. A thin SACM structure is required to reduce the breakdown voltage, but this results in a decrease in quantum efficiency that cannot be compensated by

using a larger waveguide area because it would increase the RC time constant and reduce the bandwidth.

This paper proposes a low-voltage waveguide SiGe APD integrated with a distributed Bragg reflector (DBR) to enhance the quantum efficiency while retaining high bandwidth without additional process steps. Two types of DBRs are designed using first-order and second-order gratings, respectively. The APD quantum efficiency is improved from 60% to 90% at 1550 nm by integrating a 6-period first-order DBR, achieving a 50% increase in quantum efficiency. Moreover, APDs with DBRs can still achieve a 25 GHz bandwidth, which is comparable to APDs with no DBR. We successfully demonstrate experimental open eye diagrams at 64 Gb/s PAM4, which show the increased OMA for APDs with DBRs. We also show that APDs with DBRs improve the sensitivity by ~2 dB in BER measurements. A sensitivity of around -13 dBm at a BER of 2.4×10^{-4} is obtained for APDs with a 6-pair, first-order DBR at a reverse bias of 10 V and a data rate of 64 Gb/s PAM4. Error-free transmission (BER = 10^{-12}) at 64 Gb/s PAM4 is successfully demonstrated for APDs integrated with DBRs, achieving a sensitivity of -4 dBm at a reverse bias of 10 V.

2. DEVICE DESIGN

Figure 1(a) shows the schematic of a low-voltage waveguide SiGe SACM APD integrated with a DBR. The SACM structure consists of a p-doped Ge absorption layer, a p-doped silicon charge layer, and a silicon multiplication layer, which are epitaxially grown on a silicon-on-insulator (SOI) substrate. As shown in Fig. 1(b), the scanning electron microscope (SEM) image of the device cross section illustrates the dimension of each layer, where the Ge is 400 nm thick doped with higher than $1 \times 10^{19} \text{ cm}^{-3}$ boron for P-contact and the labeled silicon region consists of 50 nm thick $2 \times 10^{18} \text{ cm}^{-3}$ p-doped silicon, 100 nm thick intrinsic silicon, and 220 nm thick silicon doped with $1 \times 10^{20} \text{ cm}^{-3}$ phosphorus for N contact. The doping

profile creates a low electric field in Ge to avoid undesired breakdown and confines the high electric field in the silicon multiplication region. Moreover, the majority holes generated in the p-doped Ge reduce the transport distance of holes, resulting in a shorter carrier transit time and therefore a higher speed [14].

As shown in Fig. 1(a), the DBR sits behind the SACM structure to reflect back the residual unabsorbed light. It was fabricated in the waveguide etching step with no additional fabrication process. Since the waveguide APD was designed with a TM mode grating coupler at 1550 nm, the DBR structure was also optimized for reflecting TM modes. Two DBR designs are integrated with conventional waveguide SiGe APDs, where DBR1 adopts a second-order grating design and DBR2 design is based on a first-order grating. A higher scattering loss and lower reflectivity are expected for the second-order grating in DBR1. However, this design is able to relax the fabrication tolerance and the photolithography requirement due to the larger dimension. The dimensions of the two DBR designs are listed in Table 1. The calculated effective indices for the TM₀ mode are ~1.4528 in groove and ~2.5118 in teeth, respectively, using the dimension of the fabricated devices.

The simulation for the APD absorption profile is conducted using a Lumerical FDTD. Figure 1(c) shows the absorption top views of 4 μm wide by 10 μm long devices with no DBR, DBR1, and DBR2, respectively. The light in a Si waveguide can be evanescently coupled into Ge for absorption along the propagation direction thanks to the larger index of Ge. However, due to the multimode waveguide design in the light absorption area, the optical modes oscillate between Si and Ge, resulting in nonuniform absorption profiles, as shown in Fig. 1(c). The light can be fully absorbed within a 4 μm × 50 μm APD with no DBR based on simulation results. In contrast, only 60% of the light at 1550 nm is absorbed when traveling in a 4 μm × 10 μm device, which is consistent with measurement. As shown in Fig. 1(c), the integrated DBR indeed reflects the unabsorbed light back and causes interference fringes in absorption profiles.

Figure 1(d) shows calculated reflectivities versus the number of grating periods for two types of DBRs using the T-matrix method, where the inset is the reflectivity spectra for the two DBR designs with 2, 4, and 6 periods, respectively. The reflectivity for both designs increases sharply with the number of periods and starts to saturate with a 6-period grating. A 6-period grating achieves reflectivities of 85% and 95% for DBR1 and DBR2, respectively. DBR2 has 10% higher reflectivity than DBR1 at 1550 nm due to the smaller scattering loss of the first-order grating in DBR2 and a slight shift of the peak wavelength, as shown in the inset of Fig. 1(d). The peak wavelength offset is caused by the Si thickness tolerance in device fabrication. The 1 dB reflectivity bandwidth for a 6-period grating is over 200 nm and 500 nm for DBR1 and DBR2,

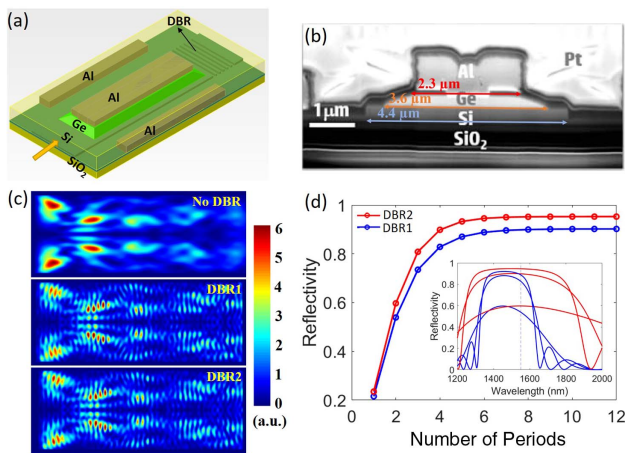


Fig. 1. (a) Schematic of a waveguide SiGe APD integrated with a DBR. (b) SEM cross-section view of an APD. (c) Top views of FDTD simulated photo-carrier generation profiles for 4 μm wide by 10 μm long APDs with no DBR, DBR1, and DBR2. (d) Calculated reflectivity versus number of period for two DBR designs. The inset is the reflectivity spectra for the two types of DBRs with 2, 4, and 6 periods, respectively.

Table 1. Design Parameters for DBR1 and DBR2

Design	Period (nm)	Groove (nm)	Teeth (nm)
DBR1	695	267	428
DBR2	384	211	173

respectively, making the design suitable for a wide range of applications (e.g., WDM). Each type of DBR was designed and fabricated with 2-, 4- and 6-period gratings on the chip for comparison. The measured quantum efficiency agrees with the calculation that APDs with 6-period DBRs realize the highest quantum efficiency. Therefore, only APDs with no DBR and 6-period DBRs will be studied here.

3. DEVICE CHARACTERIZATION

Figure 2(a) shows four characteristics for a typical low-voltage waveguide SiGe APD measured with no optical input and -10 , -5 , and 0 dBm optical power (before the grating coupler), respectively. The punch-through voltage is 2 – 3 V and the breakdown voltage is around 10 V. As shown in the curves of gain versus bias voltage, a multiplication gain of more than 10 is achieved with reverse bias of as low as 10 V. Due to the thin Si multiplication layer and the optimum doping profile, 10 V reverse bias is sufficient to generate an electric field that exceeds the impact ionization threshold for Si multiplication. The low operating voltage reduces the power consumption on APDs and makes them compatible with computer power rails [16]. The dark current is around 1 nA at 1 V reverse bias and as

low as 1 μ A even at 10 V reverse bias, which contributes negligible noise to an APD receiver for applications in direct detection. Figure 2(b) shows the measured photocurrent versus received optical power for three types of APDs, all with a waveguide dimension of 4 μ m \times 10 μ m. APDs with 6-period DBR1 and DBR2 show a consistently 30% and 40% higher photocurrent than an APD with no DBR. Figure 2(c) shows the measured responsivities at unity gain and extracted quantum efficiencies for three types of APDs with the same waveguide width of 4 μ m, but various waveguide lengths of 10 μ m, 25 μ m, and 50 μ m. A similar trend is observed for APDs integrated with DBRs. They achieve higher quantum efficiencies than APDs with no DBR, and APDs with DBR2 obtain the highest quantum efficiencies. It is worth noting that the quantum efficiency measured from 4 μ m \times 50 μ m APDs is smaller than 4 μ m \times 25 μ m APDs due to the unoptimized mode coupling and confinement design for 4 μ m \times 50 μ m APDs. As a result, the 4 μ m \times 25 μ m APDs consistently show the highest responsivity among three APD dimensions, where the one with 6-period DBR2 achieves near 100% quantum efficiency.

The APD bandwidth is characterized by impulse response measurement. A Calmar femtosecond impulse laser was injected into an APD as an optical input, and the APD response was recorded by a Keysight DCA sampling scope. Figure 2(d) shows impulse responses of an APD with a 6-period DBR2 at several gains. One can see that the full width at half-maximum (FWHM) is around 14.5 ps, except for operating at unity gain, which is similar to the FWHM of APDs with no DBR, as reported previously [14]. The oscillation at the falling edge of impulse responses is due to the space charge effect, which enhances the device bandwidth [21–23]. The model in Ref. [23] indicates that the long tail of the falling edge at unity gain results from the long carrier transit time rather than electrical parasitics. When a low-bias voltage is applied to the device the Ge is not fully depleted, which results in a slow velocity for Ge electrons. By doing a fast Fourier transform (FFT) of the impulse responses, the device responses in the frequency domain can be obtained to estimate the 3 dB electrical bandwidth. Figures 2(e) and 2(f) show the bandwidth versus gain for APDs with DBR1 and DBR2, respectively. A bandwidth of 25 GHz is achieved at gains of larger than 6 for both DBR designs. The gain bandwidth product (GBP) of near 500 GHz is obtained for the APD with DBR2, as shown in Fig. 2(f). The bandwidth enhancement near the device breakdown is achieved due to increased carrier drift velocities in the absorption region [24] and/or the inductive peaking caused by the space charge effect [23,25]. The ~ 10 GHz bandwidth is observed at low gain, as expected, due to the slow falling edge response. The device bandwidth can be enhanced at low gain with a thinner or less doped Ge layer. This approach, however, may cause other trade-offs. Additionally, it is not the gain region of interest for high-speed optical interconnects [9].

4. EYE DIAGRAM AND BER MEASUREMENT

As shown in Fig. 3, the radiofrequency (RF) signal for APD high-speed testing was generated by a Keysight 92 GS/s arbitrary waveform generator (AWG). A PAM4 $2^9 - 1$ pseudo random bit sequence (PRBS9) signal from the AWG was amplified

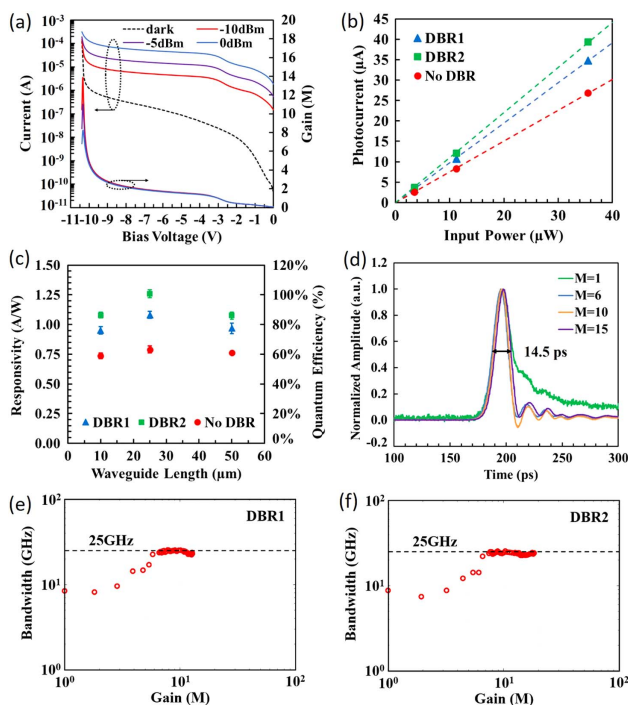


Fig. 2. (a) Dark and photo current versus bias voltage with input optical power of -10 , -5 , and 0 dBm for a 4 μ m \times 10 μ m APD with a 6-period DBR2. (b) Photocurrent versus input optical power for 4 μ m \times 10 μ m APDs with no DBR, DBR1, and DBR2, respectively. (c) Responsivity at unity gain and quantum efficiency for three types of APDs, each with the same waveguide width of 4 μ m but various waveguide lengths of 10 μ m, 25 μ m, and 50 μ m. (d) Measured impulse responses of a 4 μ m \times 10 μ m APD with DBR2 at various multiplication gains. The shortest pulse has an FWHM of 14.5 ps. (e) and (f) Device bandwidth versus multiplication gain for APDs with DBR1 and DBR2, respectively. A 25 GHz bandwidth was achieved for both designs.




Fig. 3. Experimental setup for eye diagram and BER measurement: MZ, Mach–Zehnder; EDFA, erbium-doped fiber amplifier; BPF, bandpass filter; DUT, device under test; Scope, sampling oscilloscope; BERT, bit error rate tester.

via an electrical amplifier to drive a 25 GHz commercial LiNbO₃ Mach–Zehnder modulator (MZM). A continuous wave (CW) 1550 nm optical carrier was provided by a Santec tunable laser. To compensate insertion losses from the MZM (~6 dB) and the TM-mode grating coupler (~12 dB), the MZM output was amplified using an erbium-doped fiber amplifier (EDFA) followed by a tunable optical bandpass filter (BPF) before being coupled into the waveguide SiGe APD. The APD reverse bias was applied by a Keithley source meter. The APD photocurrent signal was collected by either a Keysight DCA sampling scope for recording eye diagrams or a BER tester (BERT) after another electrical amplifier for BER analysis. The AWG calibration function was applied to the measurement system to equalize the bandwidth and signal distortion from the MZM, the electrical amplifier that drives the MZM, and RF cables.

Figure 4 shows 64 Gb/s PAM4 eye diagrams for 4 μm × 10 μm APDs with and without DBRs at a reverse bias of 8 V and 10 V, respectively. One can clearly see that eye openings are achieved for three types of APD designs, and the eye

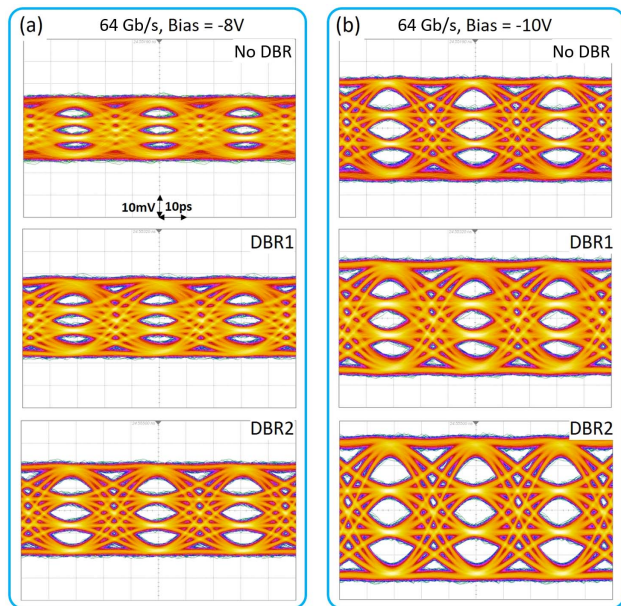


Fig. 4. Measured 64 Gb/s PAM4 eye diagrams for 4 μm × 10 μm waveguide APDs with and without DBRs at reverse bias of (a) 8 V and (b) 10 V, respectively. The DCA sampling scope was set with 10 ps/div in the *x* axis and 10 mV/div in the *y* axis.

openings for APDs with DBRs are 30%–40% larger than that of devices without DBRs at both bias conditions, which are consistent with the results from the device responsivity measurement. Compared to APDs with DBR1, APDs with DBR2 achieve larger eye openings thanks to the higher reflectivity of DBR2. Due to the higher multiplication gain at a higher reverse bias, the eye diagrams at a reverse bias of 10 V obtain around a 100% larger eye opening than at a reverse bias of 8 V.

To explore the sensitivity improvement for APD receivers assisted with DBRs, the BER measurement with non-return-to-zero (NRZ) and PAM4 signaling was performed. Figure 5 shows the BER versus the received optical average power for APDs with and without DBRs at data rates of 32 Gb/s NRZ and 64 Gb/s PAM4. Two reverse bias conditions at 8 V and 10 V are illustrated in Figs. 5(a) and 5(b), respectively. The received OMA has an extinction ratio (ER) of ~5 dB for NRZ or the outer OMA of PAM4. Compared to NRZ modulation, PAM4 signaling has 4–5 dB worse sensitivity at a high BER since each eye opening for PAM4 signaling has only one third of the eye opening with NRZ signaling. As the BER decreases, the power penalty for the device using PAM4 signaling is larger than NRZ signaling due to the harsher requirements of PAM4 for the receiver noise and the signal distortion from bandwidth limitation and time jitters [9]. As a result, BER curves with PAM4 are not as steep as with NRZ modulation.

Compared to the sensitivities of various APD designs, the overall sensitivities of APDs with DBRs are 1–2 dB better than APDs with no DBR under different modulation formats and bias conditions, as shown in Fig. 5. Figure 5(a) shows that, at reverse bias of 8 V and a data rate of 32 Gb/s NRZ, the APD with DBR2 shows ~0.5 dB better sensitivity than the APD with DBR1, achieving sensitivities of -13.3 dBm and -11.5 dBm at BERs of 2.4×10^{-4} and 10^{-12} , respectively. The BER of 2.4×10^{-4} is the threshold for error-free transmission with KP4 forward error correction (FEC), which may result in tens of nanoseconds latency for 400 gigabit Ethernet optical interconnects [8]. While doubling the data rate to 64 Gb/s with PAM4 modulation, a sensitivity of -8.7 dBm at a BER of 2.4×10^{-4} is achieved for APDs with DBRs at reverse bias of 8 V, resulting in a ~2 dB better sensitivity compared to the APD with no DBR.

When the reverse bias is increased from 8 V to 10 V, a 2–4 dB sensitivity improvement, depending on various modulation formats and APD designs, is realized, as shown in Figs. 5(a) and 5(b). For APDs with DBRs at a reverse bias of 10 V, the sensitivity achieves as low as -16 dBm at a data rate of 32 Gb/s NRZ and a BER of 2.4×10^{-4} for error-free transmission under KP4 FEC, as shown in Fig. 5(b). A sensitivity of -13.5 dBm at the same data rate is achieved for low-latency, error-free (BER = 10^{-12}) transmission without FEC. When operating at reverse bias of 10 V and 64 Gb/s PAM4, a sensitivity of around -13 dBm is achieved for the APD with DBR2 at a BER of 2.4×10^{-4} . A 4 dB better receiver sensitivity is achieved when compared to operating at reverse bias of 8 V. As shown in Fig. 5, the only situation for low-latency, error-free transmission at a data rate of 64 Gb/s PAM4 is demonstrated by the APD with DBR2 at a reverse bias of 10 V, which

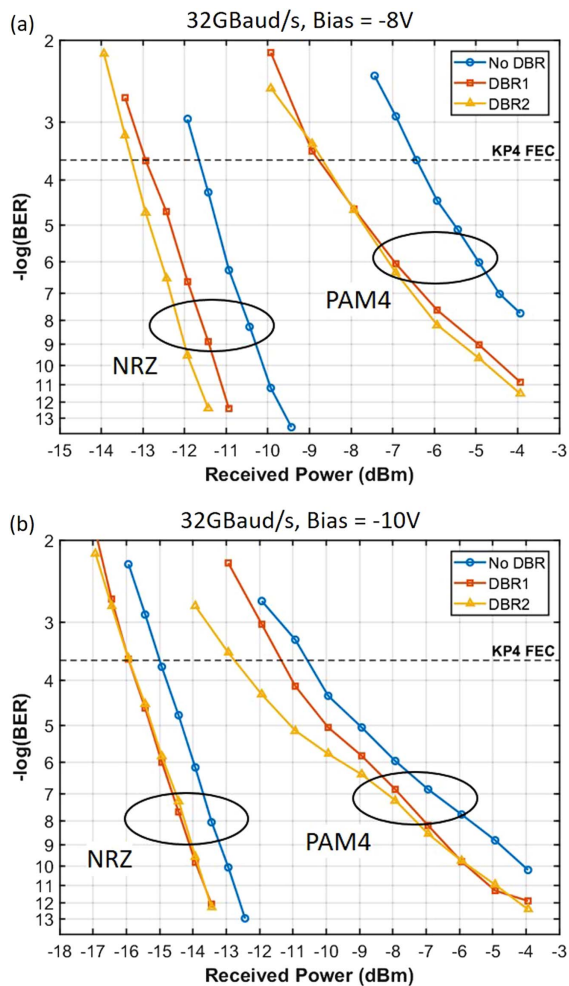


Fig. 5. 32 Gb/s NRZ and 64 Gb/s PAM4 BERs versus received average optical power with a 5 dB ER for $4\ \mu\text{m} \times 10\ \mu\text{m}$ waveguide APDs with and without DBRs at a reverse bias of (a) 8 V and (b) 10 V, respectively. These BER curves were measured without a transimpedance amplifier (TIA).

achieves a sensitivity of $-4\ \text{dBm}$. Here, the results were demonstrated without a transimpedance amplifier (TIA) and the sensitivity of APD receivers would be improved by integrating a TIA after the APD.

5. DISCUSSION

Four important metrics for APD design are the quantum efficiency, the multiplication gain, the bandwidth, and the excess noise factor. They typically have trade-offs that must be optimized carefully. As we discussed, using waveguide APDs decouples quantum efficiency and bandwidth compared to normal incidence APDs. However, we found that the trade-off still exists to further improve both the quantum efficiency and bandwidth when comparing waveguide SiGe APDs with various sizes [14]. In addition to varying the area of waveguides, the change of the Ge thickness yields the same result [14]. In contrast, APDs integrated with DBRs can increase the quantum efficiency while preserving the device bandwidth. The APD performance would be further improved using TE mode

instead of TM mode due to less scattering loss from the waveguide sidewall.

Achieving a low effective impact ionization coefficient is crucial to design APDs with high GBP and low noise. For SiGe APDs, the layer thicknesses and doping profiles should be optimized to make the avalanche multiplication mostly occur in Si rather than Ge. Since the total APD receiver noise mainly consists of the laser noise, the shot noise from APDs, and TIA input referred noise [9], the sensitivity for an APD receiver is determined by the received OMA and the overall noise. As discussed in Ref. [9], the optimum sensitivity of an SiGe APD receiver is achieved at a gain of around 10 for NRZ signaling and less than 10 for PAM4 signaling with reasonable noise consumption for lasers and TIAs. It indicates that APD design with an extremely high multiplication gain is not necessary for applications in high-speed optical interconnects.

In the gain region of interest, the bandwidth of APDs is mainly determined by electrical parasitics and the carrier transit time [23]. A thinner SACM structure can reduce the carrier transit time, but the device bandwidth may be limited by the electrical parasitics due to the larger capacitance. Another challenge is to achieve a low breakdown voltage. A thin Si multiplication layer is useful to achieve a low breakdown voltage. However, it may result in a lower gain due to the shorter multiplication path. To avoid that, the doping profile of the SACM structure, particularly for the charge layer, becomes critical to achieve low breakdown voltage as well as high multiplication gain.

6. CONCLUSION

High-bandwidth and energy-efficient optical interconnects are essential to make the bandwidth and cost scalable for data communications in mega data centers and high-performance computing. A high-sensitivity receiver is important to reduce the total link power consumption and relieve the link budget requirement. We have demonstrated low-voltage waveguide SiGe APDs integrated with DBRs to break the trade-off between quantum efficiency and bandwidth. The internal quantum efficiency is improved from 60% to 90% at 1550 nm while remaining a bandwidth of 25 GHz compared to APDs with no DBRs. The APDs with DBRs at a data rate of 64 Gb/s PAM4 show a 30%–40% increase in OMA. The sensitivities are improved by $\sim 2\ \text{dB}$ at a data rate of 64 Gb/s PAM4, achieving $-13\ \text{dBm}$ at a BER of 2.4×10^{-4} and $-4\ \text{dBm}$ at a BER of 10^{-12} . This design realizes overall better performance without additional fabrication steps.

Disclosures. The authors declare no conflicts of interest.

REFERENCES

1. Q. Cheng, M. Bahadori, M. Glick, S. Rumley, and K. Bergman, "Recent advances in optical technologies for data centers: a review," *Optica* **5**, 1354–1370 (2018).
2. B. Wang, Q. Huang, K. Chen, J. Zhang, G. Kurczveil, D. Liang, S. Palermo, M. R. Tan, R. G. Beausoleil, and S. He, "Modulation on silicon for datacom: past, present, and future," *Prog. Electromagn. Res.* **166**, 119–145 (2019).
3. A. V. Krishnamoorthy, R. Ho, X. Zheng, H. Schwetman, J. Lexau, P. Koka, G. Li, I. Shubin, and J. E. Cunningham, "Computer systems

- based on silicon photonic interconnects," *Proc. IEEE* **97**, 1337–1361 (2009).
4. J. E. Bowers, T. Komljenovic, M. Davenport, J. Hulme, A. Y. Liu, C. T. Santis, A. Spott, S. Srinivasan, E. J. Stanton, and C. Zhang, "Recent advances in silicon photonic integrated circuits," *Proc. SPIE* **9774**, 977402 (2016).
 5. T. Komljenovic, M. Davenport, J. Hulme, A. Y. Liu, C. T. Santis, A. Spott, S. Srinivasan, E. J. Stanton, C. Zhang, and J. E. Bowers, "Heterogeneous silicon photonic integrated circuits," *J. Lightwave Technol.* **34**, 20–35 (2016).
 6. C. Sun, M. T. Wade, Y. Lee, J. S. Orcutt, L. Alloatti, M. S. Georgas, A. S. Waterman, J. M. Shainline, R. R. Avizienis, S. Lin, B. R. Moss, R. Kumar, F. Pavanello, A. H. Atabaki, H. M. Cook, A. J. Ou, J. C. Leu, Y. H. Chen, K. Asanović, R. J. Ram, M. A. Popović, and V. M. Stojanović, "Single-chip microprocessor that communicates directly using light," *Nature* **528**, 534–538 (2015).
 7. V. Stojanović, R. J. Ram, M. Popović, S. Lin, S. Moazeni, M. Wade, C. Sun, L. Alloatti, A. Atabaki, F. Pavanello, N. Mehta, and P. Bhargava, "Monolithic silicon-photonic platforms in state-of-the-art CMOS SOI processes [Invited]," *Opt. Express* **26**, 13106–13121 (2018).
 8. "IEEE P802.3bs 400GbE," <http://www.ieee802.org/3/bs>.
 9. B. Wang, Z. Huang, W. V. Sorin, X. Zeng, D. Liang, M. Fiorentino, and R. G. Beausoleil, "A low-voltage Si-Ge avalanche photodiode for high-speed and energy efficient silicon photonic links," *J. Lightwave Technol.* (2019).
 10. G. Kurczveil, D. Liang, M. Fiorentino, and R. G. Beausoleil, "Robust hybrid quantum dot laser for integrated silicon photonics," *Opt. Express* **24**, 16167–16174 (2016).
 11. J. C. Campbell, "Recent advances in avalanche photodiodes," *J. Lightwave Technol.* **34**, 278–285 (2016).
 12. Y. Kang, H. D. Liu, M. Morse, M. J. Paniccia, M. Zadka, S. Litski, G. Sarid, A. Pauchard, Y. H. Kuo, H. W. Chen, W. S. Zaoui, J. E. Bowers, A. Beling, D. C. McIntosh, X. Zheng, and J. C. Campbell, "Monolithic germanium/silicon avalanche photodiodes with 340 GHz gain-bandwidth product," *Nat. Photonics* **3**, 59–63 (2009).
 13. N. Duan, T.-Y. Liow, A. E.-J. Lim, L. Ding, and G. Q. Lo, "310 GHz gain-bandwidth product Ge/Si avalanche photodetector for 1550 nm light detection," *Opt. Express* **20**, 11031–11036 (2012).
 14. Z. Huang, C. Li, D. Liang, K. Yu, C. Santori, M. Fiorentino, W. Sorin, S. Palermo, and R. G. Beausoleil, "25 Gbps low-voltage waveguide Si-Ge avalanche photodiode," *Optica* **3**, 793–798 (2016).
 15. M. Huang, P. Cai, S. Li, G. Hou, N. Zhang, T. I. Su, C. Y. Hong, and D. Pan, "56 GHz waveguide Ge/Si avalanche photodiode," in *Optical Fiber Communications Conference and Exposition (OFC)* (2018), paper W4D.6.
 16. M. Ware, K. Rajamani, M. Floyd, B. Brock, J. C. Rubio, F. Rawson, and J. B. Carter, "Architecting for power management: the IBM® POWER7™ approach," in *International Symposium on High-Performance Computer Architecture* (2010), pp. 1–11.
 17. L. Virost, P. Crozat, J. M. Fédéli, J. M. Hartmann, D. Marris-Morini, E. Cassan, F. Boeuf, and L. Vivien, "Germanium avalanche receiver for low power interconnects," *Nat. Commun.* **5**, 4957 (2014).
 18. H. T. Chen, J. Verbist, P. Verheyen, P. De Heyn, G. Lepage, J. De Coster, P. Absil, B. Moeneclaey, X. Yin, J. Bauwelinck, J. Van Campenhout, and G. Roelkens, "25-Gb/s 1310-nm optical receiver based on a sub-5-V waveguide-coupled germanium avalanche photodiode," *IEEE Photon. J.* **7**, 7902909 (2015).
 19. B. Wang, Z. Huang, X. Zeng, D. Liang, W. V. Sorin, M. Fiorentino, and R. G. Beausoleil, "60 Gb/s PAM4 low-voltage waveguide Si-Ge avalanche photodiode," in *European Conference on Optical Communication (ECOC)* (2019), paper Tu.1.E.4.
 20. X. Zeng, Z. Huang, B. Wang, D. Liang, M. Fiorentino, and R. G. Beausoleil, "Silicon-germanium avalanche photodiodes with direct control of electric field in charge multiplication region," *Optica* **6**, 772–777 (2019).
 21. J. W. Shi, F. M. Kuo, F. C. Hong, and Y. S. Wu, "Dynamic analysis of a Si/SiGe-based impact ionization avalanche transit time photodiode with an ultrahigh gain-bandwidth product," *IEEE Electron Device Lett.* **30**, 1164–1166 (2009).
 22. D. Dai, H.-W. Chen, J. E. Bowers, Y. Kang, M. Morse, and M. J. Paniccia, "Equivalent circuit model of a Ge/Si avalanche photodiode," in *6th International Conference on Group IV Photonics* (2009), pp. 1–3.
 23. B. Wang, Z. Huang, X. Zeng, W. V. Sorin, D. Liang, M. Fiorentino, and R. G. Beausoleil, "A compact model for Si-Ge avalanche photodiodes over a wide range of multiplication gain," *J. Lightwave Technol.* **37**, 3229–3235 (2019).
 24. W. S. Zaoui, H.-W. Chen, J. E. Bowers, Y. Kang, M. Morse, M. J. Paniccia, A. Pauchard, and J. C. Campbell, "Frequency response and bandwidth enhancement in Ge/Si avalanche photodiodes with over 840 GHz gain-bandwidth-product," *Opt. Express* **17**, 12641–12649 (2009).
 25. D. Dai, M. J. Rodwell, J. E. Bowers, Y. Kang, and M. Morse, "Derivation of the small signal response and equivalent circuit model for a separate absorption and multiplication layer avalanche photodiode," *IEEE J. Sel. Top. Quantum Electron.* **16**, 1328–1336 (2010).