

LETTERS

Reinventing germanium avalanche photodetector for nanophotonic on-chip optical interconnects

Solomon Assefa¹, Fengnian Xia¹ & Yurii A. Vlasov¹

Integration of optical communication circuits directly into high-performance microprocessor chips can enable extremely powerful computer systems¹. A germanium photodetector that can be monolithically integrated with silicon transistor technology^{2–8} is viewed as a key element in connecting chip components with infrared optical signals. Such a device should have the capability to detect very-low-power optical signals at very high speed. Although germanium avalanche photodetectors^{9,10} (APD) using charge amplification close to avalanche breakdown can achieve high gain and thus detect low-power optical signals, they are universally considered to suffer from an intolerably high amplification noise characteristic of germanium¹¹. High gain with low excess noise has been demonstrated using a germanium layer only for detection of light signals, with amplification taking place in a separate silicon layer¹². However, the relatively thick semiconductor layers that are required in such structures limit APD speeds to about 10 GHz, and require excessively high bias voltages of around 25 V (ref. 12). Here we show how nanophotonic and nanoelectronic engineering aimed at shaping optical and electrical fields on the nanometre scale within a germanium amplification layer can overcome the otherwise intrinsically poor noise characteristics, achieving a dramatic reduction of amplification noise by over 70 per cent. By generating strongly non-uniform electric fields, the region of impact ionization in germanium is reduced to just 30 nm, allowing the device to benefit from the noise reduction effects^{13–15} that arise at these small distances. Furthermore, the smallness of the APDs means that a bias voltage of only 1.5 V is required to achieve an avalanche gain of over 10 dB with operational speeds exceeding 30 GHz. Monolithic integration of such a device into computer chips might enable applications beyond computer optical interconnects¹—in telecommunications¹⁶, secure quantum key distribution¹⁷, and subthreshold ultralow-power transistors¹⁸.

For on-chip interconnects, the germanium (Ge)-based APD photodetector should be integrated into a silicon waveguide that can route near-infrared light on a silicon chip^{8,19–23}. Furthermore, ideally the APD should have a compact micrometre-scale footprint, operate at a voltage close to 1 V that is compatible with complementary metal oxide semiconductor (CMOS) technology, possess high avalanche gain, and detect very fast optical signals of up to 40 gigabits per second (Gb s^{-1}). These requirements are in strong contradiction with each other and are almost impossible to meet without significant innovations.

Figure 1 shows a schematic of a waveguide-integrated Ge APD, with modelling and scanning electron microscope (SEM) cross-sections. The thicknesses and widths of both the Ge and Si layers were optimized to ensure the highest responsivity within the smallest possible footprint. The resulting thicknesses of 140 nm and 100 nm, and widths of 750 nm and 550 nm for the Ge and Si layers, respectively, provide propagation of at most only two optical modes in the combined layer stack for the transverse electric field polarization at both the 1.3 μm

and 1.5 μm wavelengths^{8,20–23}. This optical design allows efficient coupling of light from the routing silicon waveguide up into the fundamental mode in the Ge layer, which is further facilitated by a short Ge taper, as shown in Fig. 1a. As can be seen from the calculated profile of the fundamental mode in Fig. 1b, the resulting optical power resides almost completely (confinement factor of over 77%) in the top Ge layer. Besides ensuring a very short absorption length of the order of 10 μm even for a wavelength of 1.5 μm , this design allows us to minimize the APD capacitance to as low as 10 fF (ref. 8).

However, growth of such a thin Ge layer directly on top of Si using epitaxial techniques is problematic because a large concentration of misfit dislocations are generated at the interface owing to a lattice mismatch of over 4%. Instead, we used a rapid melting growth technique²⁴ that results in the formation of a thin single-crystalline defect-free Ge layer on top of an insulating SiON film, isolating it from the underlying Si waveguide⁸ (see also Methods). The SEM images of lateral and longitudinal cross-sections of the Ge APD are shown in Fig. 1d and e, respectively.

The use of this very thin Ge layer in our waveguide APD ensures fast operation up to 40 Gb s^{-1} and beyond. Standard CMOS metallization includes the fabrication of a series of tungsten (W) plugs followed by the formation of the first level of copper (Cu) interconnect, which contacts the W plugs in an interdigitated manner as seen in the schematic of Fig. 1a. Because the W plugs are in direct contact with the Ge layer (see Fig. 1d and e), a series of metal–semiconductor–metal Schottky diodes are formed all along the 20 μm length of the photodiode. The small thickness of Ge means that very strong electric fields exceeding 30 kV cm^{-1} are generated throughout the Ge layer, with a bias of just 2.8 V, as seen from the calculated field profile in Fig. 1c. The high electric field strength ensures fast acceleration of both electrons and holes to their saturation velocities²⁵, resulting in photodetector operation in a drift-limited regime. Complete electrical isolation of the Ge layer from the Si waveguide also blocks the unwanted slow diffusion of photo-generated carriers away from the high-field region, thus facilitating fast response.

Figure 2a shows a series of impulse responses of an APD with 200 nm contact spacing (see also Methods). The transition from a diffusion-limited regime (slow exponential tail at 0.1 V bias) to a drift-limited regime (the fast component starts to dominate at 1 V bias) is evident. The total area under the impulse response, which is proportional to the total number of carriers collected at the electrodes, is plotted in Fig. 2b. The plateau between 0.5 V and 1.5 V biases indicates that all photo-generated carriers are being collected. The responsivity of 0.40 A W^{-1} and 0.14 A W^{-1} has been measured for this plateau for the 1.3 μm and 1.5 μm wavelengths, respectively⁸. The further exponential increase of the total pulse area at higher biases indicates a significant current gain, reaching a multiplication factor of $M = 10$ at 3.5 V. At voltages over 1 V, the fast component makes up almost 70% of the pulse area, indicating that the

¹IBM T. J. Watson Research Center, Yorktown Heights, New York 10598, USA.

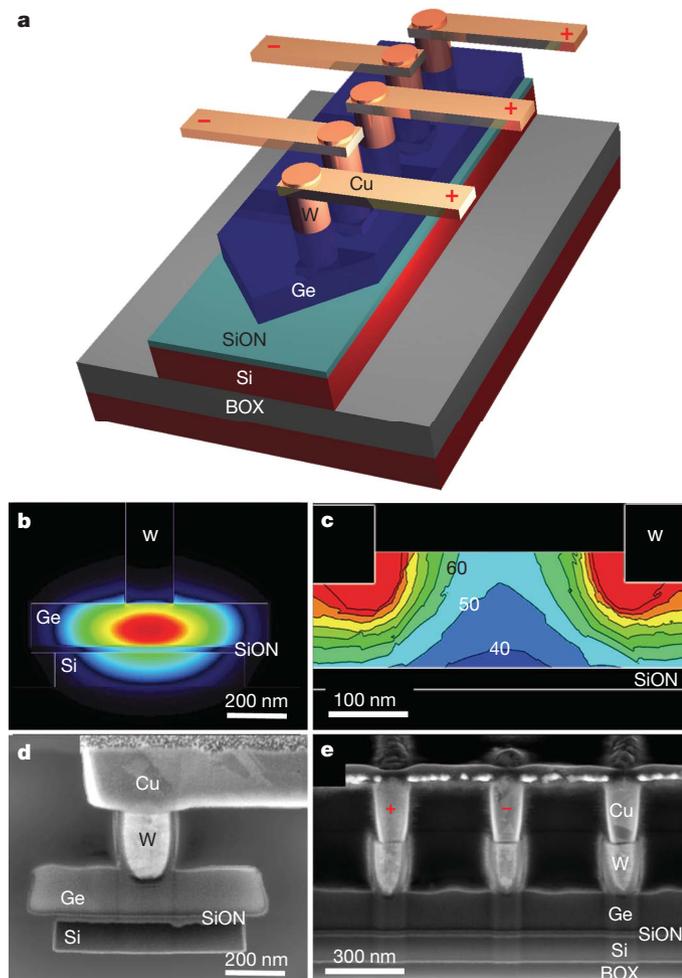


Figure 1 | The nanophotonics Ge waveguide-integrated APD. **a**, Schematic of the waveguide-integrated Ge APD. The Ge layer is deposited on top of a SiON insulating layer which overlays an Si waveguide. The detector is biased through metallic interdigitated contacts consisting of W plugs and Cu wires. **b**, Intensity profile of the fundamental optical mode in the Ge APD, calculated using the eigenmode expansion method with realistic metal parameters. The mode confinement factor reaches over 77% in the top Ge layer. **c**, The electric field distribution, calculated using the static approximation for an APD with 300 nm contact spacing and a 2.8 V bias between the contacts. Black contour lines (numbers on the contours are in kV per cm) correspond to a 10 kV cm^{-1} increase. Even at the bottom of the Ge waveguide the electric field strength exceeds the 30 kV cm^{-1} necessary to saturate the drift velocity for both electrons and holes. Areas coloured red in the immediate vicinity of the W plugs correspond to electric field strength in excess of 120 kV cm^{-1} high enough to impose avalanche amplification. **d**, An SEM image of a lateral cross-section of the Ge APD showing the Si and Ge waveguides, separated by a thin SiON dielectric layer, and the W plug and Cu wire. **e**, An SEM image of the longitudinal cross-section of the Ge APD taken along the waveguide length. Bias voltage is applied to Cu wires, as shown by red plus and minus signs. BOX denotes the buried oxide layer.

gain is fast. Furthermore, the gain is broadband, as is clear from the ‘open-eye’ diagram in the inset of Fig. 2b, which was measured at $M = 10$ for 40 Gb s^{-1} random non-return-to-zero (NRZ) bit stream (see Methods). Therefore, the gain in our photodetector can hardly be assigned to a slow photovoltaic gain observed in metal–semiconductor–metal Schottky photodiodes²⁶ caused by ionization of traps or surface states and subsequent carrier injection to maintain the charge neutrality. Rather, the exponential current increase results from an avalanche gain due to an impact ionization process. Indeed, as opposed to the uniform electric field distribution in conventional p–i–n diodes, the use of metal–semiconductor–metal contacts produces highly non-uniform fields with an extremely large field gradient towards the W plugs, as seen in Fig. 1c. The electric field strength in

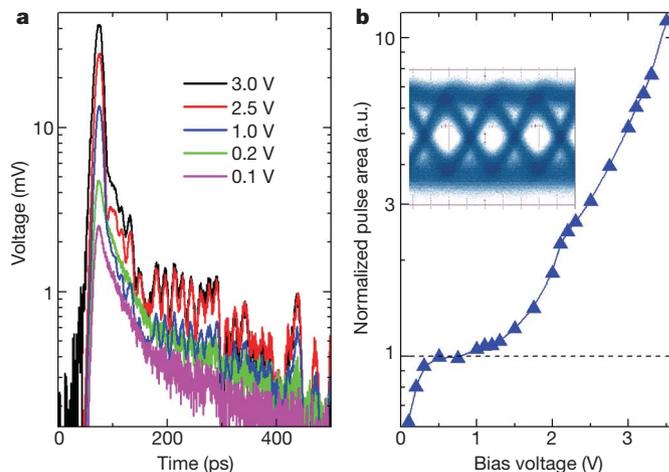


Figure 2 | Time-resolved photocurrent measurements. **a**, Semi-logarithmic plot of photocurrent transients measured on an APD with 200 nm contact spacing at different biases. **b**, Semi-logarithmic plot of the normalized photocurrent pulse area for different bias voltages extracted from Fig. 1a. Normalization is done with respect to the plateau between 0.5 V and 0.8 V that corresponds to the collection of all photo-generated carriers. The inset shows a 40 Gb s^{-1} ‘eye’ diagram measured at 3.4 V bias that corresponds to a gain of $M = 10$. The eye diagram was measured for a $2^{31} - 1$ PRBS NRZ bit stream.

the immediate proximity of the corners of the W plugs exceeds 120 kV cm^{-1} , suggesting a high probability of impact ionization.

Additional evidence for the avalanche origin of the observed gain can be found in a series of small-signal radio-frequency (r.f.) measurements of the S_{21} parameter (see Methods), shown in Fig. 3a for the APD with 200 nm contact spacing (a different sample from the one used in Fig. 2a). The r.f. traces are characterized by an almost flat frequency response from 10 MHz up to 1 GHz with the average amplitude significantly increasing with bias voltage. This r.f. amplitude is plotted in Fig. 3d (red curve), exhibiting a characteristic plateau between 0.4 V and 0.8 V that corresponds to the collection of all photo-generated carriers. The further exponential increase in r.f. amplitude is analogous to the gain observed in Fig. 2b, with the multiplication factor rising to $M = 9.3$ at 1.5 V. Similar measurements performed for the APD with 400 nm contact spacing (blue trace in Fig. 3d) also demonstrate similar high multiplication factors but at higher voltages of around 3.7 V, illustrating the dependence of field strength close to the W plugs on the contact spacing.

The 3 dB bandwidth of the APD with 200 nm contact spacing is extracted from the r.f. traces of Fig. 3a and plotted in Fig. 3c (red curve). The bandwidth exhibits a fast increase from 5 GHz up to 34 GHz at biases from 0.1 V to 1.1 V, while the gain just starts to ramp up. At higher biases, however, the bandwidth stays almost constant, owing to carriers reaching their saturation velocity as well as to intrinsic time delays associated with impact ionization. Despite the apparent saturation of the APD bandwidth, the gain \times bandwidth product continues to grow because of the rise in avalanche gain. The resulting gain \times bandwidth product (red curve in Fig. 3b) reaches 300 GHz, which is among the highest ever reported for APD photodetectors¹². The same trend is seen for the APDs with 400 nm contact spacing (shown in Fig. 3d and c) with the highest bandwidth of 39.5 GHz achieved at 2.5 V bias close to the end of the unity gain plateau. Saturation of the bandwidth before considerable gain is reached indicates that carrier transport and avalanche amplification are taking place in spatially separated areas within the APD. The avalanche multiplication occurs mostly in a small volume around the W plugs where the field is high enough for impact ionization, whereas photo-carriers are mostly generated in between the plugs and then drift to contacts at saturation velocity.

A large (10 dB) avalanche gain in the APD does not necessarily guarantee a corresponding increase in the detector sensitivity (that is,

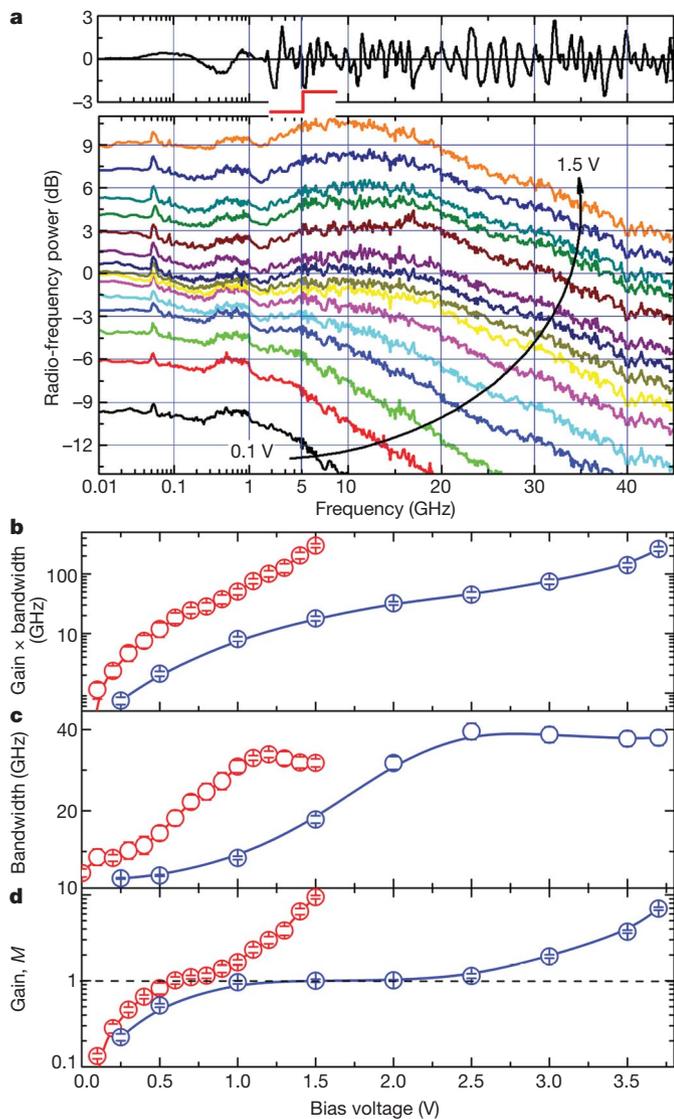


Figure 3 | Small signal r.f. measurements of the S_{21} parameter. **a**, Small-signal r.f. measurements of the S_{21} parameter of an APD with 200 nm contact spacing (the sample was different from that measured in Fig. 2a). Traces are measured for different bias voltages between 0.1 V and 1.5 V in 0.1 V steps. The power of the 1.3 μm laser light at the APD input was kept constant at 1 mW. After normalization on a calibration curve (see Methods), all experimental traces show an identical periodically oscillating pattern most probably coming from reflections at the pads/probe interface, which are not impedance matched. For ease of comparison this pattern, shown at the top of Fig. 3a, was subtracted from each of the experimental traces. The horizontal axis is split at 5 GHz with a logarithmic scale for lower frequencies down to 10 MHz and a linear scale for higher frequencies up to 45 GHz. **b**, Gain \times bandwidth product extracted from S_{21} r.f. traces as a function of bias voltage. Red circles correspond to the APD with 200 nm contact spacing measured in Fig. 1a. Blue circles correspond to the APD with 400 nm contact spacing. Lines are drawn to guide the eye. **c**, 3-dB bandwidth for different bias voltages measured from small signal r.f. traces of the S_{21} parameter. Colour coding as in **b**. **d**, Multiplication factor M for different bias voltages measured as an average amplitude in r.f. response at low frequency range between 10 MHz and 1 GHz. Colour coding as in **b**. Error bars in **b–d** represent one s.d.

an ability to detect low-power optical signals with a small bit error rate), because the latter can easily degrade as a result of the higher excess noise level^{27,28}. Initial evidence that the excess noise in our APD is relatively low even at a gain of $M = 10$ is supported by the 40 Gb s⁻¹ ‘eye’ diagram that remains open as shown in the inset of Fig. 2b. To confirm, a series of more detailed bit-error-rate measurements were performed at 10 Gb s⁻¹ on the APD with 300 nm contact spacing for different bias voltages. The results plotted in Fig. 4a show that the

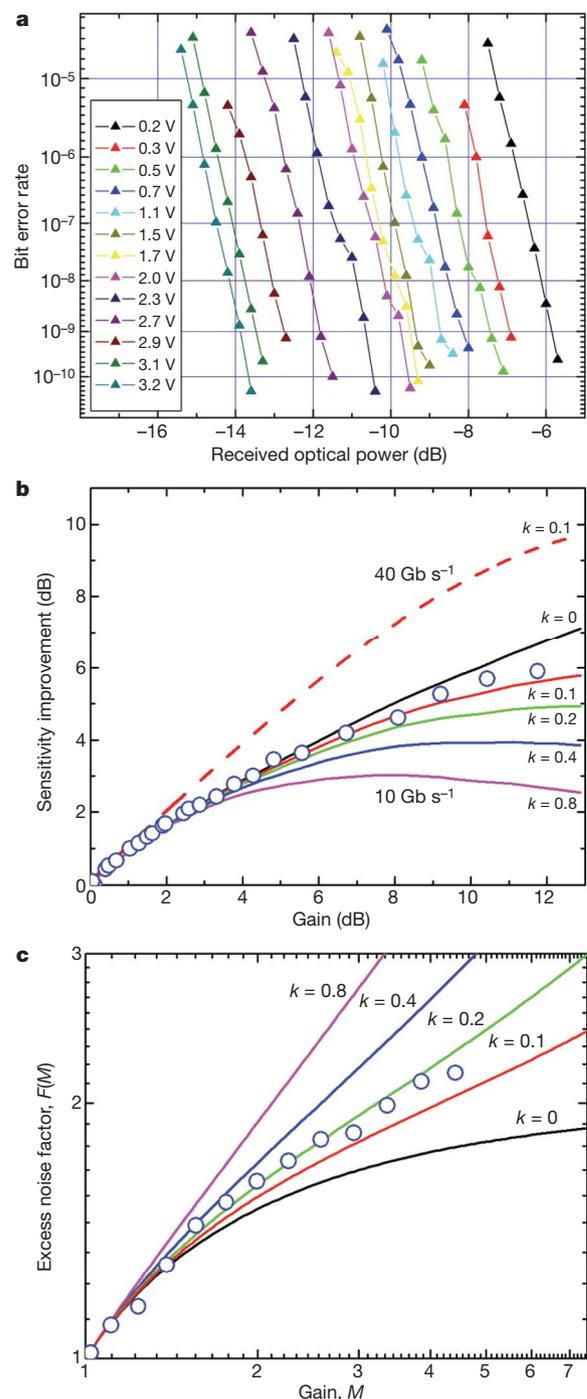


Figure 4 | Sensitivity and excess noise measurements. **a**, Bit-error-rate measurements of the APD with 300 nm contact spacing at 10 Gb s⁻¹. Sensitivity curves are measured for different bias voltages with a 2³¹ - 1 PRBS NRZ bit stream. **b**, Sensitivity improvement (blue circles) measured from the experimental curves of **a** at a bit error rate of 10⁻⁹ versus avalanche gain. The sensitivity value is normalized with respect to the unity gain plateau, at around 0.7 V. The avalanche gain is extracted from separately measured S_{21} r.f. traces similar to Fig. 3a. The solid curves are the result of theoretical fitting using the model described in the Supplementary Information, which takes into account experimentally measured dark current, gain and the noise figure of the receiver amplifier. The only adjustable parameter in the model is k_{eff} from equation (1). The dashed red line is a theoretical curve calculated from the same model using a dark current ten times smaller than that observed experimentally. **c**, The excess noise factor measured for different gain values in the APD with 300 nm contact spacing (the sample was different from that used in **a**). Avalanche gain is extracted from separately measured S_{21} r.f. traces similar to Fig. 3a. Experimental data (blue circles) are fitted with the McIntyre model equation (1) for different k_{eff} values.

sensitivity continues to improve even after the unity gain plateau is reached, at around 0.7 V. The open blue circles in Fig. 4b represent an improvement of sensitivity measured (see Methods) at a bit error rate of 10^{-9} as a function of an avalanche gain. Because bit-error-rate measurements were performed at a wavelength of 1.5 μm , where responsivity is only 0.15 A W^{-1} , and a 45 GHz bandwidth amplifier was used in the set-up (see Methods), the absolute sensitivity of only -8 dB m was measured at unity gain; however, a significant improvement of 5.9 dB at a bias of 3.2 V was achieved when the gain reached 11.8 dB.

High dark current as well as excess multiplication noise are usually the main factors resulting in saturation of sensitivity improvement with further increase of gain. Because it is a metal–semiconductor–metal photodetector, our APD is characterized by relatively large dark current reaching 50 μA at a unity gain (Supplementary Fig. 1). To understand the dominant contributions to the excess noise level, theoretical fitting (solid curves in Fig. 4b) was performed using non-local McIntyre's theory²⁷ (see Supplementary Information). The stochastic nature of the impact ionization process means that the excess noise factor $F(M)$ could be expressed for the case of a uniform electric field when electrons initiate multiplication as:

$$F(M) = k_{\text{eff}} [M + (2 - 1/M)](1 - k_{\text{eff}}) \quad (1)$$

where k_{eff} is an effective ratio of ionization coefficients for electrons and holes. Such coefficients are almost equal in bulk Ge giving a k_{eff} of about 0.9, which results in a very large excess noise, making conventional Ge APD uncompetitive for building digital optical links^{9–11}.

In contrast, the fitting of experimental data in Fig. 4b reveals $k_{\text{eff}} \approx 0.1$, which is much less than anticipated for bulk Ge. Moreover, as shown in Fig. 4b (dashed curve), an improvement in sensitivity of over 10 dB at even higher bit rates of 40 Gb s^{-1} can be expected in our APD if the dark current could be suppressed tenfold by, for example, increasing the Schottky barrier height (that is, unpinning the Fermi level at the W plug interface²⁹).

The direct experimental measurements of excess multiplication noise in the APD with 300 nm contact spacing are shown in Fig. 4c (see also Methods). Fitting the data with equation (1) further confirms very low excess noise with k_{eff} less than 0.2. The total reduction of noise spectral power density can be estimated as more than 70% with respect to the noise expected for a bulk Ge.

Several factors can account for the dramatic reduction of excess multiplication noise in our nanophotonic APD. First, as seen from the field profile in Fig. 1c, the avalanche multiplication is happening only in very close proximity to the W plug, within just 30 nm of its edge. Thinning the multiplication region has been shown to favour a more deterministic statistics of ionization process and a narrower ionization-path-length probability distribution function^{13–16}. Thus, a large portion of the excess noise reduction can be accounted for, owing to this 'dead space' effect¹³. Second, another possible contribution to noise suppression could be the 'initial energy' effect¹⁴ in which carriers entering the multiplication region have already acquired high energy. In fact, spatial separation of absorption and multiplication regions in our APD might result in significant narrowing of the energy distribution of photo-carriers entering the multiplication region because most of them are drifting at saturation velocity. These might then narrow the probability distribution functions and suppress excess multiplication noise. Third, the large electric field gradients generated in our APD in the proximity of the W plugs could result in further narrowing of the probability distribution functions owing to the fast acceleration of secondary carriers towards the ionization threshold¹⁵.

METHODS SUMMARY

The waveguide-integrated APD was fabricated from single-crystal Ge waveguides formed by the rapid melting growth method^{8,20}. The use of a thin SiGe buffer for the deposition of Ge layer resulted in unwanted incorporation of an Si concentration of about 10% in the final crystallized Ge waveguide layer during

rapid melting growth that resulted in decreased effective absorption at wavelength 1.3 μm . The metallization included W plugs followed by formation of Cu interdigitated contacts. SiON fibre couplers were formed at the edges of the silicon chip for efficient coupling of light from a microlensed fibre⁸.

High gain values, low excess noise and fast operation at low bias were common observations for all the devices measured. We found it difficult, however, to perform all sets of experiments on a single device owing to accidental device failures at gains higher than 10 dB, to which all devices were eventually driven. The experiments were done on a set of almost identical devices that were fabricated on the same wafer and were located in neighbouring dies. Therefore we believe that the findings reported in this work are generic to the device design we used.

Bandwidth and gain were characterized by time-resolved photocurrent measurements that were performed using optical pulses coupled to the SiON coupler while applying bias voltage to the APD. Further confirmation was obtained through small-signal r.f. measurements of the S_{21} parameter performed with a scalar network analyser while coupling in modulated laser light at 1.3 μm .

To confirm that sensitivity is not degraded when operating at high bandwidth and high gain, eye-diagrams were measured by detecting a $40 \text{ Gb s}^{-1} 2^{31} - 1$ pseudo-random binary sequence (PRBS) NRZ optical bit stream generated by modulating the laser beam at a wavelength of 1.5 μm . Furthermore, a similar set-up was used to perform bit-error-rate measurements by attenuating the input to the APD. Finally, noise measurements were performed with the noise figure meter at 90 MHz.

Full Methods and any associated references are available in the online version of the paper at www.nature.com/nature.

Received 22 September 2009; accepted 7 January 2010.

1. DARPA IPTO. Exascale Computing Study Report (http://users.ece.gatech.edu/~mrichard/ExascaleComputingStudyReports/ECS_reports.htm) 133 (28 September 2008).
2. Masini, G., Sahni, S., Capellini, G., Witzens, J. & Gunn, C. Waveguide photodetectors integrated in a CMOS process. *Adv. Opt. Technol.* **2008**, 1–5 (2008).
3. Ahn, D. *et al.* High performance, waveguide integrated Ge photodetectors. *Opt. Express* **15**, 3916–3921 (2007).
4. Yin, T. *et al.* 31GHz Ge n-i-p waveguide photodetectors on silicon-on-insulator substrate. *Opt. Express* **15**, 13965–13971 (2007).
5. Vivien, L. *et al.* 42 GHz p.i.n germanium photodetector integrated in a silicon-on-insulator waveguide. *Opt. Express* **17**, 6252–6257 (2009).
6. Chen, L. & Lipson, M. Ultra-low capacitance and high speed germanium photodetectors on silicon. *Opt. Express* **17**, 7901–7906 (2009).
7. Tang, L. *et al.* Nanometre-scale germanium photodetector enhanced by a near-infrared dipole antenna. *Nature Photon.* **2**, 226–229 (2008).
8. Assefa, S. *et al.* CMOS-Integrated 40 GHz germanium waveguide photodetector for on-chip optical interconnects. *Optical Fiber Communication Conf. OFC/NFOEC paper OMR4* (2009); *Opt. Express* (in the press).
9. Pearsall, T. P., Temkyn, H., Bean, J. C. & Luryi, S. Avalanche gain in GeSi/Si infrared waveguide detectors. *Electron. Device Lett.* **7**, 330–332 (1986).
10. Melchior, H. & Lynch, W. T. Signal and noise response of high speed germanium avalanche photodiodes. *Trans. Electron Devices* **13**, 829–838 (1966).
11. Ando, H. *et al.* Characteristics of germanium avalanche photodiodes in the wavelength region of 1–1.6 μm . *J. Quantum Electron.* **14**, 804–809 (1978).
12. Kang, Y. *et al.* Monolithic germanium/silicon avalanche photodiodes with 340GHz gain-bandwidth product. *Nature Photon.* **3**, 59–63 (2009).
13. Hayat, M. M., Saleh, B. E. A. & Teich, M. C. Effect of dead space on gain and noise of double-carrier multiplication avalanche photodiodes. *Trans. Electron Devices* **39**, 546–552 (1992).
14. Hayat, M. M. *et al.* Boundary effects on multiplication noise in thin heterostructure avalanche photodiodes: theory and experiment. *Trans. Electron Devices* **49**, 2114–2123 (2002).
15. Liew, S. C. *et al.* Modeling of avalanche multiplication and excess noise factor in InAlAs avalanche photodiodes using a simple Monte Carlo model. *J. Appl. Phys.* **104**, 013114 (2008).
16. Campbell, J. C. Recent advances in telecommunications avalanche photodiodes. *J. Lightwave Technol.* **25**, 109–121 (2007).
17. Kardynat, B. E., Yuan, Z. L. & Shields, A. J. An avalanche-photodiode-based photon-number-resolving detector. *Nature Photon.* **2**, 425–428 (2008).
18. Gopalakrishnan, K., Griffin, P. B. & Plummer, J. D. Impact ionization MOS (I-MOS)-Part I: device and circuit simulations. *Trans. Electron Devices* **52**, 69–76 (2005).
19. Vlasov, Y., Green, W. M. J. & Xia, F. High-throughput silicon nanophotonic wavelength-insensitive switch for on-chip optical networks. *Nature Photon.* **2**, 242–246 (2008).
20. Dehlinger, G., McNab, S. J., Xia, F. & Vlasov, Y. A. Waveguide photodetector. US patent US7515793 (February 2006).
21. Assefa, S., Bedell, S. W., Xia, F. & Vlasov, Y. A. Optoelectronic device with germanium photodetector. Pending US patent 11/925170 (October 2007).
22. Assefa, S. *et al.* Suspended germanium photodetector for silicon waveguide. Pending US patent 12/191687 (August 2008).
23. Assefa, S., Xia, F. & Vlasov, Y. A. Avalanche impact ionization amplification devices. Pending US patent 12/533521 (August 2009).

24. Liu, Y., Deal, M. D. & Plummer, J. D. High-quality single-crystal Ge on insulator by liquid-phase epitaxy on Si substrate. *Appl. Phys. Lett.* **84**, 2563–2565 (2004).
25. Reggiani, L., Canali, C., Nava, F. & Ottaviani, G. Hole drift velocity in germanium. *Phys. Rev. B* **16**, 2781–2791 (1977).
26. Carrano, J. C. *et al.* Comprehensive characterization of metal–semiconductor–metal ultraviolet photodetectors fabricated on single-crystal GaN. *J. Appl. Phys.* **83**, 6148–6160 (1998).
27. McIntyre, R. J. Multiplication noise in uniform avalanche diodes. *Trans. Electron. Devices* **13**, 164–168 (1966).
28. Agrawal, G. P. *Fiber-Optic Communication Systems* 3rd edn, Ch. 4 (Wiley, 2001).
29. Kobayashi, M. *et al.* Fermi level depinning in metal/Ge Schottky junction for metal source/drain Ge metal-oxide-semiconductor field-effect-transistor application. *J. Appl. Phys.* **105**, 023702 (2009).

Supplementary Information is linked to the online version of the paper at www.nature.com/nature.

Acknowledgements We are grateful to C. Schow for numerous discussions and insights, S. Bedell for Ge growth, Y. Zhang for his etch expertise, T. Topuria and P. Rice for TEM characterization of the samples, and the staff of Materials Research Laboratory at IBM Yorktown Heights for contributions to device fabrication.

Author Contributions All authors contributed equally to this work.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Correspondence and requests for materials should be addressed to S.A. (sassefa@us.ibm.com).

METHODS

Calculations. Optical mode profiles and attenuation coefficients were calculated using the fully vectorial three-dimensional (3D) eigenmode expansion method (Fimmwave, Photon Design). The results were confirmed with a fully vectorial 3D finite difference time domain method (Fullwave, RSoft) with realistic material numbers for metallization layers. The 3D electric field distribution for different biases was calculated in a static approximation using the finite element method (Maxwell3D, Ansoft).

Fabrication. Single-crystal Ge waveguides were obtained by the rapid melting growth method^{8,20}. After formation of silicon waveguides a thin SiON layer was deposited on top and a small seeding window down to Si layer was etched. Using rapid thermal chemical vapour deposition (RTCVD), a thin buffer layer of 30% SiGe was grown first, followed by a thick Ge layer. After Ge waveguide patterning and encapsulation the Ge was melted using rapid thermal annealing at around 1,000 °C. During fast cooling the crystallization of melted Ge starts from a seeding window and propagates along the Ge waveguide, leaving behind a high-quality single-crystalline SiGe strip 20 μm long with a total Ge concentration of over 90%. Subsequently the photodetector is metallized with W plugs followed by formation of Cu interdigitated contacts. SiON fibre couplers were formed at the edges of silicon chip to obtain efficient coupling of light from a fibre.

Impulse response measurements. Time-resolved photocurrent measurements were performed using optical pulses centred at wavelength 1.3 μm and having a pulse width of 120 fs generated from a Coherent MIRA optical parametric oscillator and coupled to the integrated SiON coupler at the edge of a silicon chip via a microlensed fibre. Bias voltage was applied to APD pads with a 40 GHz Picoprobe probe connected to a 50 GHz bias-tee and the collected photocurrent was further amplified with a 45 GHz Picosecond 5882 amplifier and measured with an Agilent 86100C DCAj oscilloscope with 86116B 80GHz electrical plug-in. Owing to timing jitter in oscilloscope triggering, the time resolution was limited to 15 ps, which precluded accurate measurements of optical bandwidth.

Eye-diagram measurements at 40 Gb s⁻¹. For the eye diagram, we detected a 40 Gb s⁻¹ 2³¹ - 1 PRBS NRZ optical bit stream generated by modulating the

laser beam at a wavelength of 1.5 μm with an external 40 GHz LiNbO₃ modulator. Modulated light was amplified with an erbium-doped fibre amplifier and then filtered with a 2 nm notch filter to provide 1 mW input power at the photodetector. The modulator was driven by electrical pulses from a SHF BPG44E pattern generator amplified with a SHF810 amplifier.

Small-signal r.f. S₂₁ measurements. Small-signal r.f. measurements of the S₂₁ parameter were performed with an Anritsu 56100A 50 GHz scalar network analyser. Laser light at 1.3 μm was modulated with an external 40 GHz LiNbO₃ modulator driven with an SHF810 amplifier, providing 0.150 mW of input power at the photodetector. The reference normalization trace was measured using the same set-up with a 45 GHz New Focus 1014 photodetector. Additionally, all components, including the reference photodetector, were characterized with an Agilent PNA E8361C network analyser connected to a N4373C 67 GHz Lightwave Component Analyser.

Bit-error-rate measurements at 10 Gb s⁻¹. Bit-error-rate measurements were performed with the same set-up as used for eye-diagram measurements. Modulated, optically amplified and filtered light from a 1.5 μm laser was attenuated with the Anritsu MN9610B variable optical attenuator. The electrical signals from an APD amplified with the Picosecond 5882 amplifier were analysed with the SHF EA40G bit-error analyser. The gain was measured separately for the same bias points using the set-up for r.f. measurements with the Anritsu 56100A 50 GHz scalar network analyser.

Excess noise measurements. Noise measurements were performed with an HP 8970B noise figure meter at 90 MHz. The alternating-current component of the APD photocurrent was amplified with a Picosecond 5882 amplifier. Full system calibration was performed with the 15 dB calibration noise source. Calibration curves of noise power density versus photocurrent as well as photocurrent versus optical input power were measured separately at unity gain and confirm the high linearity of these responses over four orders of magnitude. With these calibration curves the excess noise factor $F(M)$ was then extracted as a function of multiplication factor M . The latter was measured separately for the same bias points using the set-up for r.f. measurements with the Anritsu 56100A 50 GHz scalar network analyser.