SLAC–PUB–6271
SLAC/SSRL–0038
October 1993
(SSRL–M)

Time-Resolved Detection of X-rays Using Large-Area Avalanche Photodiodes *

A.Q.R. Baron and S.L. Ruby
Stanford Linear Accelerator Center
Stanford Synchrotron Radiation Laboratory
Stanford, CA 94309, USA

Abstract

We have measured the efficiencies and time responses of several large-area (200 and 80 mm^2) beveled-edge avalanche photodiodes to x-rays of energy 6–20 keV in a photon counting mode. The diodes were used with avalanche gains of 10^2 to 10^3. Detector efficiency varied from 50% at 8 keV to 5% at 20 keV, while the time resolution was 1/2 ns with a tail out to longer times (5 ns). The duration of the tail may be reduced (<2 ns) by selecting only larger pulses from the diode (increasing the discriminator threshold). These results are well explained by a simple physical picture of the device: the fast response is due to x-rays absorbed in the depletion layer, and the tail results from x-rays absorbed in a low-field undepleted region in front of the diode. For one device, the total active thickness (depletion and low field region) was measured to be 49 μm and the low field region to be 9 μm.

Submitted to Nuclear Instruments and Methods in Physics Research

---

* Work Supported by Department of Energy contract DE-AC03-76SF00515.

Time Resolved Detection of X-rays Using Large Area Avalanche Photodiodes *

A.Q.R. Baron and S.L. Ruby
Stanford Linear Accelerator Center
Stanford Synchrotron Radiation Laboratory
Stanford, CA 94309, USA

Abstract

We have measured the efficiencies and time responses of several large-area (200 and 80 mm^2) beveled-edge avalanche photodiodes to x-rays of energy 6–20 keV in a photon counting mode. The diodes were used with avalanche gains of 10^2 to 10^3. Detector efficiency varied from 50% at 8 keV to 5% at 20 keV, while the time resolution was 1/2 ns with a tail out to longer times (5 ns). The duration of the tail may be reduced (<2 ns) by selecting only larger pulses from the diode (increasing the discriminator threshold). These results are well explained by a simple physical picture of the device: the fast response is due to x-rays absorbed in the depletion layer, and the tail results from x-rays absorbed in a low-field undepleted region in front of the diode. For one device, the total active thickness (depletion and low field region) was measured to be 49 μm and the low field region to be 9 μm.

* Work Supported by Department of Energy contract DE-AC03-76SF00515.

I. Introduction

Avalanche photodiodes (APDs) are possible replacements for the traditional scintillator-photomultiplier (PMT) combination used in many photon-counting experiments for x-rays of energies <20 keV. They are intrinsically fast devices having good time resolution and large dynamic range [1]. When used in a spectroscopic setup, their energy resolution is better than an NaI(Tl)-PMT combination, and similar to a proportional counter (10% at 5.9 keV) [2][3][4].

An avalanche photodiode is essentially a reverse biased p-n junction operated near [5] its breakdown voltage (see [6] and [7] for detailed reviews). The electric fields are sufficiently high (10^5 V/cm in Si) that current reaching the high field region will be amplified through impact ionization, resulting in gains that may vary from 10^1 to 10^3. However, the gain is a sensitive function of the electric fields in the device, and it is only relatively recently that processing techniques have allowed sufficient control [8] to make large area (> 1 square cm) devices of good uniformity.

Presently APDs are being used in an increasing number of applications. APD/scintillator combinations [4][9], are used to detect high energy (> 100 keV) x-rays. Here the advantage of the APDs over PMTs is their immunity to magnetic fields (as high as 2 T [10]) and their small size and ruggedness. APDs may be used for direct detection of low-energy (1.5 keV [2][3]) and very low energy x-rays (200–700 eV, in a Geiger mode [11]). The time resolution of specially designed APDs has been pushed as low as 20 ps (cooled and operated in Geiger mode [12]). APDs have also been used as the amplifying element in a photomultiplier tube

[13][14]: the photocathode is followed by a high field region that accelerates photoelectrons to about 10 keV before they hit the surface of an APD, for a total gain of 10^6. Finally, multipixel arrays of APDs for x-ray detection have been fabricated [15].

II. Devices Measured

We have investigated the time responses of 4 large-area silicon APDs to x-rays of energy 6 to 20 keV. All diodes were of the "beveled edge" configuration [16] manufactured by Advanced Photonix, Inc. (API) [17]. Three of the devices were 16 mm in diameter, and one was 10 mm in diameter. Table 1 lists the voltages, gains, and dark currents for each diode just before breakdown, as measured by API before delivery. Note that there is some diode-to-diode variation.

The basic structure of these devices is shown in Fig. 1, while Fig. 2 shows approximate doping, field, and avalanche coefficient profiles. When high voltage (2200–2500 V) is applied across the APD, the depletion region spreads out from the p-n junction to within 5–10 μm of the surface and deep into the n region, resulting in the field profile shown qualitatively in Fig. 2b. The front portion of the p region has a residual field estimated to be 50 V/cm [18], resulting from the high doping gradient [19]. The peak field near the junction is estimated to be 1.8×10^5 V/cm [20], sufficient to cause impact ionization by electrons. Figure 2c shows the approximate distribution of the ionization coefficient (the inverse of the average length required for an electron in the high-field region to generate an electron-hole pair). This last is effectively the gain that an electron traveling through the APD will experience at each position. Note that in silicon the ionization coefficient for holes is much lower than that for electrons so, to a first approximation, hole amplification may be ignored [21].

III. Model for X-ray Response

X-rays incident from the p+ side of the device will be absorbed in the silicon with a characteristic (1/e) attenuation length $L_{abs}(E\gamma)$, where $E\gamma$ is the photon energy. This is plotted in Fig. 3 [22]. The majority of absorption events result in the production of a fast electron having nearly the energy of the incident photon. This fast electron slows, generating one electron-hole pair for every 3.6 eV of energy [23], over some characteristic distance: the range for a 10 keV electron in Si is approximately 1.4 μm [24]. On the scale of field variations within the diode, each x-ray absorption event may be considered a point-like deposition [25] of E_γ/3.6 eV electron-hole pairs.

The single photon, x-ray response of these APDs can be understood by considering the locations at which the x-rays are absorbed. There are three regions of interest (see Figs. 1 and 2).

(i) X-rays absorbed in the undepleted p region (the first 5–10 μm of the device), where there is a small electric field, will generate electrons that will be slowly (1.3 ns/μm for a field of 50 V/cm in pure Si [26]) transported to the edge of the high-field region, where they will then be swiftly carried to the high-gain region and amplified. However, there will be a delay of as much as several ns due to the time required to traverse the low-field region. In addition, traps [27] in this region may hold electrons for long periods of time (tens of ns to several μs

[28]) causing a reduction in the peak amplitude of the current pulse getting to the high-field area as well as lengthening its trailing edge as the electrons are released. These effects combine to give output pulses from the APD having lower peak height, a slower fall time, and, possibly, a slower rise time than pulses from absorption of x-rays in the high-field region ((ii), below). The pulses from this region may also have a lower integrated charge if some of the traps lead to recombination, as well as delay.

(ii) X-rays absorbed in the front portion of the depleted p region will generate pulses with a fast rise time and the complete gain of the device. Electrons will quickly (<1 ns) be transported to the high gain region of the device and amplified with approximately the average device gain M [29].

(iii) X-rays absorbed within the gain region of the device, either at the back of the depleted p region, or in the front of the depleted n region, will generate electrons that will only be partially amplified, resulting in lower amplitude pulses that will appear (up to a scale factor) much like pulses from region (ii) above. The amplitudes of these pulses should vary continuously to zero.

This simple model accounts nicely for the results described below.

IV Experimental Setup

X-Ray Source

All of these measurements were done at Stanford Synchrotron Radiation Laboratory (SSRL) on bending magnet beamline 2–3. A two-crystal silicon (220) monochromator was used to define the photon energy to < 0.1%. A slight angular shift between the two crystals was introduced to reduce the contribution of higher-order harmonics. Second-order harmonic contamination was less than 0.5% from 8 to 10 keV, and substantially less than that (<0.1%) above 10 keV [30]. The photon beam was collimated to a spot size of about 50 x 150 μm, and passed through an ion chamber before falling on the photodiode.

The time structure of the x-ray pulses is the same as that for the electron bunches in the storage ring (SPEAR). The radiofrequency (rf) of the accelerating field, 358 MHz, confines electrons to bunches of nearly Gaussian shape and small width (full width at half maximum = 0.13 ns [31]) separated by integer multiples of 2.8 ns = 1/358 MHz. Typically, only every fourth or fifth rf bucket was filled with electrons, around 2/3 of the ring circumference, with 1/3 of the circumference left empty. The electron distribution was then a series of bunches separated by 11.2 or 14 ns followed by a dead time of 250 ns. This was repeated at 1.28 MHz, or every 780 ns, corresponding to the ring circumference. Upon close examination, smaller bunches (by a factor of 10^{-3}) were discovered between the bunches described above. These "minibunches" were an irritant in the measurements described below, but not a problem, since they could be identified by their 2.8 ns periodicity (Fig. 9).

Electronics

Figure 4 is a schematic of the electronics used in these measurements. A more detailed description noting facts relevant to these measurements follows.

High Voltage Control

The APD was kept at positive high bias using a Bertan current-limited voltage supply. The limit was set so that if the APD drew more than 8 μA current, the high voltage would shut off. This happened perhaps five times during a 10 day experimentation period. Only once did the supply shut off at APD gains of less than M 500 (most shut offs were at gains of 1000 or higher and, therefore, at higher voltages as well). Shut offs seemed uncorrelated with high count rates.

Signal Amplification and Output Pulse Shape

The signal from the APD was immediately fed into a Phillips 6954 pulse pre-amplifier (x100). The output of the 6954 amplifier was sent through a 32 ns cable to another amplifier (EG&G FTA 420) which provided additional x20 gain. A scope photograph of the signal out of the FTA 420 amplifier is shown in Fig. 5 for APD#1 operated at a gain of M 300. Each trace corresponds to the pulse generated by the absorption of a single 14.4 keV photon. Note the appearance of two distinct types of pulses, one having large amplitude and the other with a smaller amplitude, slower decay and possibly a slower rise time. These may be identified with photons absorbed in regions (ii) and (i), respectively, as will be shown below [32]. In addition, pulses from region (iii) create a general blurring of the area under the larger amplitude peak. The signals from diodes 2, 3, and 4 were similar to that for diode 1, with the exception that, for diode 4, the rise and fall times were slightly shorter, and the intensity of the lower amplitude component was reduced.

Timing Electronics

The output of the 420 was fanned out (Phillips 744) and sent to a constant fraction discriminator (Ortec 934 CFD). The discriminators have adjustable thresholds from –30 to –1000 mV. The minimum (magnitude) threshold used was about 55 mV in order to discriminate against noise from the 6954 amplifier. For this threshold, noise from the amplifier triggered the discriminator at rates much less than 0.1 Hz.

The time response of the diodes was determined by measuring the interval between the CFD output and a signal synchronized to the 1.28 MHz ring frequency of the synchrotron (and therefore locked to the time at which photons hit the APD surface). The CFD output was used to start a time to amplitude converter (TAC), and the 1.28 MHz signal was used as a stop [33]. The TAC output was fed into a multichannel analyzer with conversion gain set so that one channel was 1/12 ns. A check of the ring timing signal against a delayed copy of itself showed the jitter in the electronics to be about 2 channels, 0.2 ns.

V. Efficiencies

The efficiency of each APD was calculated by dividing the number of pulses out of the device causing the CF discriminator to fire by the flux incident normally on the surface of the silicon (both integrated over

10 seconds). The latter number was calculated from the integrated current out of the argon-filled ion chamber immediately upstream of the APD and introduces an error of < 5% over the energy range of 8–20 keV [34]. Note that in all cases the efficiency quoted here is for photons incident on the surface of the silicon of the APD: there has been a correction included for the absorption in the thin Al window used as a visible light shield in front of the diode [35]. Count rates were typically 10^4/sec.

Figure 6 shows the efficiency as a function of energy for all four diodes at comparable average gains M 200 and a 55 mV discriminator threshold (about 10% of the peak signal height at 14.4 keV). The average gain M for each diode was calculated by finding the voltage across the APD and using tables provided by API for each diode. Note that these tables of gain as a function of voltage were generated using visible light that illuminated the entire surface of the APD. For point-like illumination, as was used here, these values are approximate: the gain varied (for a fixed voltage) by about 10% for M 200 when a diode was scanned through the 14.4 keV x-ray beam. Gain variations have been seen before (e.g. [6]) and tend to worsen at higher average gains.

The measured efficiency depends on the setting of the discriminator threshold. In general, it is desirable to set this threshold as low as possible to get the highest efficiency, but not so low that noise triggers the discriminator. The lower limit, set by noise from the 6954 amplifier, was 40 mV. Table 2 shows how this threshold affects the efficiency measured for APD#1 at two different gains at 14.4 keV. At the lower gain, one does not count all the photons absorbed within the active region of device, while at the higher gain, the efficiency seems to have saturated. In fact, for

APD#1, the gain was raised as high as M 2000 with an increase in efficiency at 14.4 keV to only 0.119. This suggests that for gains of M> 600 with APD #1, and energies ≥14.4 keV, the measured efficiency is a good estimate of the actual fraction of photons that are creating pulses within the active thickness of the APD.

A series of measurements of the efficiency of APD #3 with M 700 and a 55 mV discriminator threshold allow estimation of the active thickness of silicon in the diode. This is possible because the absorption length of x-rays in silicon varies from 65 μm at 8 keV to 983 μm at 20 keV (see Fig. 3). Assuming the 55 mV threshold is sufficient to collect all of the pulses that undergo amplification, one would expect the device to have 100% efficiency over some active thickness of silicon. The efficiency **E** should then be fit by the expression

$$\mathbf{E}(E_\gamma) = c(E_\gamma) \quad [\, 1 - \exp\{-L_{active} / L_{abs}(E_\gamma)\}\,], \qquad (1)$$

where $L_{abs}(E_\gamma)$ is the absorption length in silicon for x-rays of energy E_γ, and L_{active} is the active thickness of the photodiode. The quantity in square brackets is simply the fraction photons absorbed in the first L_{active} thickness of material (the 0.2 μm p+ layer and the 0.5 μm passivation layers at the device surface have been neglected). The term $c(E_\gamma)$ is a correction factor to compensate for the fact that not all absorption events generate electrons. It is the ratio of the photoelectric cross section to the total cross section (photoelectric + Compton + Raleigh) and varies from 0.99 at 8 keV to 0.93 at 20 keV [22][36].

The data and a fit to the data using Eq. (1) are plotted in Fig. 7. The agreement is quite good, giving L_{active} = 49 μm. Note that 49 μm is about

10 μm larger than x_j, the distance from the front of the device to the junction. This confirms that there is significant gain in the front portion of the depleted n region.

The efficiency of the diodes may be improved if the incident x-ray beam has a small area. The diode may be tipped, relative to the incident beam direction, so that the x-ray path length in the active region of silicon is increased. This has been done, with a factor three increase in efficiency (from 0.1 to 0.3) observed at 14.4 keV, for a beam incident at about 70 degrees from normal.

VI. Time Response

Response at 14.4 keV

The time response of APD #1 to 14.4 keV radiation (M 200, discriminator threshold 55 mv) is shown in Fig. 8. The plot shows the distribution of the number of events (each a single x-ray absorption) triggering the constant fraction discriminator as a function of the time after the signal synchronized to the electron storage ring. The plot may be interpreted as being the probability distribution for the time required for an x-ray entering the device to generate an output pulse larger than the discriminator setting. The responses of all diodes are similar, showing a sharp peak of FWHM 1/2 ns and a tail out to later times, lasting 5 ns.

The basic response pattern, sharp peak followed by a long tail, is easily explained using the model above. The counts occurring in the peak are from photons absorbed in the depletion region of the diode where photoelectrons are quickly transported to the avalanche region and amplified. The tail is from photons absorbed in the undepleted p region at the front of the device, where the field is low, so that electrons are only slowly transported to the depletion region.

For timing measurements, it is desirable to reduce the tail on the response curve. This is possible by raising the discriminator threshold, thereby selecting for larger pulses. Figure 9 shows the effects of higher discriminator thresholds on the response of APD #3: the tail is reduced in length, while the FWHM is unaffected. Note that in Fig. 9 the amplitude of the peak also decreases, suggesting that the higher discriminator threshold reduces the number of "fast" events. This is due to removal of the events in which x-rays were absorbed in the interior of the avalanche region where the photoelectrons were not fully amplified.

Similar results were obtained for all four diodes, as shown in Fig. 10, where the FWHM and full width at 100th maximum is plotted as a function of the efficiency obtained by raising the discriminator threshold (higher efficiencies corresponding to lower thresholds). Note that the full width at half maximum remains constant over this range.

The correlation of pulse height with delay time may be explained in terms of the model above by assuming that the front low field region has mechanisms that will trap electrons for short periods of time (at least a few ns) or even lead to recombination. If the number of electrons trapped increases with distance traveled in this region, then one would expect exactly the correlation noted: when absorption occurs in the undepleted region, the output pulse will be delayed due to the transit time in the low field and attenuated due to trapping. Both delay and trapping will increase as the point of absorption moves toward the surface [37].

Variation in Time Response With X-ray Energy

The model above suggests that the counts in the tail of the time response are due to absorption of x-rays in the front of the APD in the undepleted p region. A simple confirmation of this is the change in response with x-ray energy. In particular, one would expect that as the attenuation length for the x-rays in Si decreases at lower energy, a larger fraction of the x-rays would be absorbed in the front of the device, leading to larger tails. This is confirmed in Fig. 11 which shows the increase in tail amplitude as the x-ray energy varies from 14.4 to 5.5 keV for APD#1 at M 700.

A quantitative measure of the increase in tail size may be obtained by measuring the number of counts in the tail of the time response as a function of photon energy. Figure 12 shows a plot of the "slow efficiency" in the tail, calculated by multiplying the measured efficiency by the fraction of counts occurring more than 1/2 ns after the peak of the time response. The solid line is a fit to an expression of the form of Eq. (1). There is good agreement for a "slow" active length of L_{slow} = 9 μm, suggesting the low field region at the front of the device is 9 μm thick. In addition, the length of the tail, 4 ns, gives the average velocity over this region to be 2.3 μm/ns, corresponding to an average effective field of about 160 V/cm (mobility 1450 $cm^2/s/V$ [26]), somewhat different than the 50 V/cm value suggested by API [18]. The reason for this difference is not yet understood.

VII. Summary

This work presents a clear picture of the time response of one type of large-area avalanche photodiode. This response is easily and consistently understood relative to the device structure and the point like creation of carriers from x-ray absorption. The active thickness of silicon, 50 μm, may be conceptually and practically divided into three regions. The front 10 μm is not depleted and contributes a tail to the time response resulting from the low drift velocity in this region. The next 20 μm of silicon is depleted, having a high drift velocity. The field in this region, however, is not sufficient to cause significant gain, so photons absorbed in this region will generate electrons that are quickly transported to the high gain section of the device where they are fully amplified. The last 20 μm of the device is the gain region, and photons absorbed here will be quickly, but only partially amplified.

APDs used for direct detection of x-rays may be good replacements for the traditional scintillator/photomultiplier tube combination. APDs have a faster response, larger dynamic range, and better energy resolution. The efficiency, as measured here, may not be as good at high energies, due to the limited stopping power of the 50 μm of silicon [38]. However, this may be improved by a factor of 3 or more for small-area x-ray beams (e.g. synchrotron beams are 10 x 1 mm) simply by tipping the device and increasing the x-ray path length in the active thickness of silicon.

Acknowledgments

The authors gratefully acknowledge the assistance of Dr. Ernesto Gramsch of Advanced Photonix, Inc., who kindly answered our many questions about the APDs used in these experiments. We would like to thank George S. Brown for many useful and intellectually stimulating conversations and John Arthur for his help during the experiment. We would also like to acknowledge that it was the work of Shunji Kishimoto [1] which first introduced us to the use of APDs for x-ray detection. Support for this research was provided by the U.S. Department of Energy under contract DE–AC03–76F00515.

Notes

[1] S. Kishimoto. Rev. Sci. Instrumen. 63 (1992) 824–827.

[2] P. P. Webb and R. J. McIntyre. IEEE Trans. Nucl. Sci. NS-23 (1976) 138–1441.

[3] R. Farrell, K. Vanderpuye, G. Entine and M. R. Squillante. IEEE Trans. Nucl. Sci. 38 (1991) 144–147.

[4] M. J. Szawlowski, S. Zhang, A. DeCecco, M. Madden, M. Lindberg and E. Gramsch. IEEE Symposium on Nuclear Science, Orlando, Fl. (1992).

[5] These diodes were always operated below the avalanche breakdown voltage where there is a well defined average gain. Diodes may also be operated above breakdown voltage in a "Geiger" mode where a single carrier can trigger run-away gain until the circuit is externally quenched. See, for example: A. Lightstone, A. MacGregor, D. MacSween, R. McIntyre, C. Trottier and P. Webb. Electronic Engineering. 61 (1989) 37–45.

[6] P. P. Webb, R. J. McIntyre and J. Conradi. RCA Review. 35 (1974) 234–278.

[7] S. M. Sze. Physics of Semiconductor Devices (2nd ed). (John Wiley and Sons Inc., 1981).

[8] V. L. Gelezunas, W. Siebt and G. Huth. Appl. Phys. Lett. 30 (1977) 118–120.

[9] C. Carrier and R. Lecomte. IEEE Trans. Nucl. Sci. 37 (1990) 209–214.

[10] R.R. Raylman, Private Communication (1993).

[11] J. R. Palmer and G. R. Morrison. Rev. Sci. Instrumen. 63 (1992) 828–831.

[12] This was done with small-area devices and visible light, but presumably could be done with x-rays as well, though with poor efficiency due to the small active thickness in these devices. S. Cova, A. Lacaita, M Ghioni and G. Ripamonti. Rev. Sci. Instrumen. 60 (1989) 1104–1110.

[13] M. Szawlowski, E. Gramsch, S. Zhang, and M. Madden. Performance of the Vacuum Avalanche Photodiode (1993). Available from API.

[14] P. Cushman and R. Rusack. To be published in Nucl. Instr. and Meth. (1993).

[15] E. Gramsch, S. Zhang, M. Madden, M. Lindberg and M. Szawlowski. Presented at 1993 SPIE Conference in San Diego, California. To be published in the conference proceedings. See also M. R. Squillante, R. Farrell, J. C. Lund, F. Sinclair, G. Entine, and K. R. Keller, IEEE Trans. Nucl. Sci. 33 (1986) 336–339.

[16] All APD designs must somehow reduce the fields at the silicon surface, otherwise the large surface defect concentration will lead to large noise currents and breakdown. The beveled edge design is one of several possibilities. See [6] and [7].

[17] Advanced Photonix, Inc. (API) 1240 Avenida Acaso, Camarillo, California 93012. Telephone (805) 987-0146. Note that another company making similar devices (e.g. reference [3]) is Radiation Monitoring Devices, Inc. (RMD), 44 Hunt Street, Watertown, Massachusetts 02172. Telephone (617) 929–1167.

[18] "Response Time Characteristics of Advanced Photonix Avalanche Photodiodes". Application note from API.

[19] A. G. Jordan and A. G. Milnes. IEEE/IRE Transactions on Electron Devices. 7 (1960) 242–251.

[20] "Noise Characteristics of Advanced Photonix Avalanche Photodiodes". Application note from API.

[21] It is well known that the impact ionization coefficient for holes in silicon is much less than that for electrons, see for example: R. VanOverstraeten and H. DeMan. Solid-State Electronics. 13 (1970) 583–608. For detailed calculations of APD response, the average ratio of the electron coefficient to the hole coefficient is an important parameter, see [6]. However, for understanding the x-ray response at the level of this paper, it is sufficient to know this ratio is small, 0.002 [20], so that holes will not suffer much amplification.

[22] W. H. McMaster, N. K. Del Grande, J. H. Mallet, and J. H. Hubbell, Lawrence Radiation Laboratory, Compilation of X-ray Cross Sections (1969).

[23] R. D. Ryan. IEEE Trans. Nucl. Sci. NS 20 (1973) 473–480.

[24] This range was calculated in the continuous slowing down approximation (CSDA) from stopping power tables in J. C. Ashley, C. J. Tung, R. H. Ritchie and V. E. Anderson. IEEE Trans. Nucl. Sci. NS 23 (1976) 1833–1837. See also L. Katz and A. S. Penfold. Rev. Mod. Phys. 24 (1952) 28–44

where an empirical range-energy relation gives 0.7 μm for the range of 10 keV electrons in aluminum of density 2.3 g/cc.

[25] This point-like deposition is important in comparing the response of a diode used in x-ray photon counting to that in applications where pulses containing many photons are detected (e.g. visible, or near visible, light applications). A pulse containing many photons will generate electron-hole pairs that are initially distributed over a large (many micron) range of depths within the silicon, as each photon is absorbed at a slightly different depth. The response to a pulse of many photons is then an average over a fairly large portion of the device. As a result, there are qualitative differences between the output signals for visible light pulses which always create some carriers in the low-field undepleted p region and single x-rays, that may create carriers only within the depletion region.

[26] C. Jacoboni, C. Canali, G. Ottaviani and A. A. Quaranta. Solid-State Electronics. 20 (1977) 77–89.

[27] G. L. Miller, D. V. Lang and L. C. Kimerling. Annual Review of Material Science. 7 (1977) 377.

[28] S. Cova, A. Lacaita and G. Ripamonti. IEEE Electron Device Letters. 12 (1991) 685–687.

[29] There are several sources of gain variation that should be noted. In addition to the ±10% variation in the average gain seen at different positions in the device (see Section V), there will be statistical variation from event to event even at a fixed value of the average gain. This arises from the avalanche multiplication process. See [2] for a summary of the

relevant factors involved. A detailed treatment may be found in R. J. McIntyre. IEEE Transactions on Electron Devices. ED–13 (1966) 164. See also R. J. McIntyre. IEEE Transactions on Electron Devices. ED–19 (1972) 703–713.

[30] This contamination was measured using the APD, which has an efficiency that dies off at larger energies. Therefore the incident beam probably had a somewhat larger harmonic content.

[31] Private Communication from Heinz-Dieter Nuhn (1993).

[32] The presence of the separate band of pulses associated with the low-field region suggests that there is some mechanism that causes a reduction in the peak height of all pulses from x-rays absorbed in this region. One speculates that this may be trapping of electrons (or recombination) just after absorption. Also, the density of the E_γ /3.6 eV electron-hole pairs is comparable to the majority carrier density, so space charge may have an effect as well.

[33] This arrangement of start and stop pulses necessitates "reversing" the x-axis of the data, but prevents the TAC start from being triggered at 1.28 MHz, which would saturate the device.

[34] At energies below 8 keV, the high absorption of the Ar in the ion chamber and resultant harmonic contamination make this method incident flux measurement difficult.

[35] For other experiments, where high efficiency was important, the diodes were used with an aluminized mylar window, having nearly negligible absorption at these x-ray energies. However, these windows

were slightly transparent to visible light, resulting in a larger background current. This was not desirable for the measurements described here, which were occasionally carried out at very high gains. Therefore a 25 μm Al window was used in this work.

[36] Equation (1) ignores multiple scattering: the possibility that a photon might suffer non-destructive scattering at one location and then be photo-electrically absorbed someplace else. This is reasonable because the correction is only significant at higher energies at which the active device thickness is much less than the absorption length for the x-rays.

[37] Another partial explanation might be a non-uniform electric field profile (where the field decreased toward the front of the device), resulting in a larger transit time spread (due to the spatial extent of the initial charge cloud) and therefore lower peak height for absorption events occurring at the front of the layer. However, the flatness of the time response in the tail suggests that this is not the dominant cause.

[38] An alternative to the beveled edge design, the reach through design (see references [2], [6] and [7]), permits increasing the active thickness to > 100 μm of silicon, with a decrease in the average electric field strength. We have not tested any of these devices, but would expect improved x-ray efficiency at the cost of poorer time resolution. Larger-area reach through devices (5 x 5 mm) are available from EG&G Optoelectronics Canada, 22001 Dumberry Road, Vaudreuil, Quebec J&V 8P7 Canada. Telephone (514) 424–3300.

Tables

Table 1: APD parameters just before breakdown. From the data sheets supplied with each diode from API.

APD Number	Diameter (mm)	Voltage (V)	Gain (M)	Dark Current (uA)
1	16	2510	2500	1.7
2	16	2213	740	3.5
3	16	2223	500	3.6
4	10	2188	260	0.2

Table 2: Efficiency as a function of discriminator threshold for APD#1 operated at two different gains. Note saturation at high gain and low threshold.

Discriminator Voltage (mV)	Efficiency at M 200	Efficiency at M 600
42	0.102	0.115
50	0.097	0.112
60	0.077	0.110

Figures

Figure 1. Schematic cross section of a beveled edge APD. X-rays enter through a thin (0.2 μm) p+ window at the device surface (x=0) and are absorbed in the silicon, generating electron-hole pairs that are then amplified. The depletion region extends from x_p in the p-region to x_n in the n-region. The junction is located at x_j 38 μm. There is a 0.5 μm thick passivation layer on the entrance surface which is not shown.

Figure 2. Approximate (a) doping, (b) field and (c) ionization coefficient (local gain) for the APD's used here. The vertical dashed lines are approximate divisions for the regions discussed in Section III of this paper. Modified from a figure in [18], courtesy of API.

Figure 3. X-ray absorption (1/e) length in silicon as a function of x-ray energy. From fits in [22].

Figure 4. Schematic of electronics used in this experiment. Component values are R_p=8.2 M for the protection resistor and C=100 pF for the coupling capacitor. After pre-amplification (x100) the signal was sent to a NIM bin where it was further amplified (x20) and the sent to a constant fraction discriminator (CFD). The discriminator output was used to drive a time to amplitude converted (TAC).

Figure 5. Scope trace for APD#1 as seen on a Tektronics 475A oscilloscope looking at the output of the 420A (x20) amplifier. 14.4 keV

photons were incident at a rate 10 kcts/sec and the diode was run at 2435 V, M 300. One division is 200 mV and 2.5 ns.

Figure 6. Efficiency measured for each APD as a function x-ray energy at gain M 200 and 55 mV discriminator threshold.

Figure 7. Efficiency of APD #1 as a function of energy at M 700, discriminator threshold 55 mV. Solid line is a fit using equation (1) of the text with L_{active} = 49 μm.

Figure 8. Time response of APD #1 to single 14.4 keV x-rays for M 200 and a 55 mV discriminator threshold. This may be interpreted as the probability distribution for the length of time between when a 14.4 keV photon enters the APD and when an output pulse above the discriminator threshold is generated. The zero time has been arbitrarily set at the peak position.

Figure 9. APD time response with increasing discriminator threshold for APD#3 operated at M 220. All plots have been scaled to the same incident flux rate and counting time. Thresholds and efficiencies are : (a) 56 mV, 0.099 (b) 147 mV, 0.080 (c) 243 mV, 0.065 (d) 309 mV, 0.051 (e) 353 mV, 0.045 (f) 409 mV, 0.043 (g) 442 mV, 0.032 (h) 504 mV, 0.011. Note the appearance of the minibunch at 2.8 ns as the resolution improves.

Figure 10. Full width at 100th maximum (upper points) and FWHM (lower points) for all 4 APDs with 14.4 keV x-rays incident. Plotted as a function of increasing efficiency (or decreasing discriminator threshold). The FWHM remains nearly constant at 0.5 ns while

the full width at 100th maximum increases as the efficiency of the device is improved by lowering the discriminator threshold.

Figure 11. Time response of APD#1 (M 700, 55 mV discriminator threshold) at 5.5, 8.0, and 14.4 keV (absorption lengths 22, 65, and 370 μm, respectively). All plots have been normalized to the same area. Note the rising tail at late times as the energy decreases. The response at 20 keV (983 μm absorption length) would lie directly on top of 14.4 keV response.

Figure 12. Plot of the slow efficiency as a function of energy for APD#1 at M 700. The solid line is a fit using equation (1) with L_{slow} = 9 μm. See text.