Infrared Detector Arrays for Astronomy

G.H. Rieke

Steward Observatory, University of Arizona, Tucson, Arizona 85750;
email: grieke@as.arizona.edu

Key Words

astronomical techniques, bolometers, photoconducting devices,
photodiodes, impurity band conduction (IBC) detectors

Abstract

Use of infrared detector arrays in astronomy began roughly 20 years
ago, and our detection capabilities in parts of this spectral range
have doubled about every seven months since then. A variety of
approaches are now used for detector arrays operating from 1 μm
to 1 mm and beyond. They include direct hybrid arrays of InSb and
HgCdTe photodiodes that operate from 0.6 μm to 5 μm, and of Si:As
impurity band conduction detectors from 5 μm to 28 μm; a number
of approaches to photoconductive detector arrays in the far-infrared;
and bolometer arrays read out by transistors or superconducting
devices in the far-infrared through millimeter-wave spectral range.
The underlying principles behind these approaches are discussed.
The application of these principles is illustrated through detailed
discussion of a number of state-of-the-art detector arrays.

1. INTRODUCTION

We have many ways to detect photons. In selecting one of these options, most users feel constrained by practical considerations such as ruggedness and convenience. However, astronomers are nearly unique in the demands they make to approach fundamental levels of performance. For them, how quickly a detection can be achieved on faint sources is virtually all that matters. This behavior depends on a combination of detector characteristics such as (*a*) noise, (*b*) quantum efficiency and (*c*) the number of pixels that can be put into an array. Together, they can be gauged by the astronomical capability metric (Bahcall et al. 1991), which is proportional to

$$(\# \textit{ of pixels}) \times (\textit{sensitivity per pixel})^2. \qquad (1)$$

This metric indicates the number of positions on the sky that can be measured in a given time to a given detection limit. A similar expression applies for spectroscopy (where for a simple spectrograph, the free spectral range and slit length substitute for positions on the sky).

The foundations of modern infrared astronomy were laid in 1965–1985, based on observations with single detectors (or small collections of such detectors operated in parallel) (Low, Rieke & Gehrz 2007). The first astronomical use of true infrared detector arrays occurred about 20 years ago (e.g., Arens et al. 1984, Forrest et al. 1985). These early devices were small (typically 32×32 pixels) with read noises of more than 1000 electrons. They were described as having performance similar to that of the best single-pixel devices, but with the multiple-pixel advantages of an array (see Wynn-Williams & Becklin 1987 for a number of articles on early infrared arrays for astronomy). Starting just before these first demonstrations, the astronomical capability for near-infrared detection as measured by Equation 1 has doubled roughly every seven months.

Astronomers can only take a little credit for the incredible expansion of capabilities in the near- and mid-infrared (1 μm to 5 μm and 5 μm to 30 μm, respectively). The initial investments that established the basic architecture and processes to produce high-performance arrays came from the military (e.g., Harwit 2000). Further investments from NASA and the National Science Foundation have allowed the development of the incredible devices to be described below. The array manufacturers have taken a strong interest and have done far more for us than we might have anticipated. Some of these groups have been disbanded, but we must express our appreciation to the ones that are still developing improved detector arrays. They are Raytheon Vision Systems (RVS), DRS Technologies, and Teledyne Imaging Sensors (formerly Rockwell Scientific Company; the names occasionally change as a result of corporate reorganizations and acquisitions). In the far-infrared (30 μm to 200 μm) and submillimeter (200 μm to 1000 μm), however, astronomy has often been the major customer for detector arrays and astronomers have played far more central roles in developing the fundamental technologies.

This review first discusses hybrid infrared detector arrays operating from 0.6 through 120 μm. It describes the building blocks for these devices: Section 2 examines the detectors themselves; Section 3 examines the readouts; and Section 4 examines

the process to join the detectors and readouts together. Further discussion of these matters can be found in Rieke (2003). Section 5 discusses a few examples of state-of-the-art arrays. Section 6 describes bolometer arrays being constructed for the range from ∼50 μm to the millimeter-wave. Richards & McCreight (2005) provide a shorter review of all these topics that also gives more background information. Section 7 describes some aspects of the detailed performance of these devices, and Section 8 summarizes and illustrates briefly the impact of arrays on infrared astronomy.

2. DETECTORS

Detector arrays for the near- and mid-infrared detect photons through some form of photoconductivity. The most fundamental photoconductive detector is a small block of semiconductor material with electrodes on opposite faces that are used to establish an electric field within the material volume. When photons free charge carriers within the block, the carriers migrate toward the electrodes and produce a current that can be sensed by a high-performance amplifier. When the photon energy is at least as large as the binding energy of the electrons in the semiconductor crystal, detection occurs through intrinsic photoconductivity, where the charge carriers are freed by breaking a semiconductor crystal bond. In the language of solid state physics, an electron is promoted from the valence energy band across the band gap to the conduction energy band. Detectors based on intrinsic absorption can operate up to a cutoff wavelength, λ_c, corresponding to the band gap energy of the material, or

$$\lambda_c = \frac{hc}{E_g} = \frac{1.24\,\mu\text{m}}{E_g(\text{eV})}, \tag{2}$$

where h is Planck's constant, c is the speed of light, and E_g is the band gap energy. For example, InSb has a band gap of ∼0.22 eV at 77 K and hence $\lambda_c \sim 5.6$ μm.

Detectors can also be based on extrinsic photoconductivity. Here, a photon frees a charge carrier by interacting with an impurity atom within the semiconductor crystal. In the language of solid state physics, an electron may get elevated from the valence band to an impurity level to create a hole (a missing negative charge that can be described as if it were a positively charged particle). Alternatively, an electron may be freed by elevating it from an impurity level to the conduction band. Either the hole or the electron can then move through the material in response to the electric field in the detector. The advantage of detectors based on extrinsic photoconductivity is that they can operate to much longer wavelengths, because it takes less energy to free a charge carrier from an impurity atom than from an atom of the semiconductor crystal material. The necessary photon energy corresponds to the energy difference between the impurity level and either the top of the valence band or bottom of the conduction one. That is, the gap energy, E_g, in Equation 2 can be replaced by the smaller impurity ionization energy E_i. Intrinsic photoconductivity is a very efficient process because the absorption of photons is strong and therefore high quantum efficiency (the proportion of photons converted to detectable charge carriers) can be achieved in thin layers of material. Extrinsic photoconductivity is far less efficient because of limits in the amount of impurity that

can be introduced into the semiconductor without altering the nature of the impurity states.

Implicit in the treatment of low background detectors is that the generation of free charge carriers is dominated by photon absorption, not by thermal excitation. Lower temperature is required for this condition as the long wavelength cutoff of the detector increases, going approximately as (see Hoffman, Love & Rosbeck 2004)

$$T_{\max} = \frac{200\,\mathrm{K}}{\lambda_c(\mu\mathrm{m})}. \tag{3}$$

Second-order effects, such as construction of the detector as a photodiode with a high impedance depletion zone, can relax the temperature requirements, but only modestly.

For a simple, or bulk, photoconductor, the identical block of semiconductor must meet both the electrical and optical requirements. Conflicts in these requirements can limit the detector performance. The most fundamental electrical parameter is resistance, R. The thermally driven Brownian motion of charge carriers generates a fluctuating electrical current that produces Johnson (1928) noise:

$$\langle I_J^2 \rangle^{1/2} = \sqrt{\frac{4kT\,df}{R}}, \tag{4}$$

where k is Boltzmann's constant, T is the absolute temperature, and df is the frequency bandwidth of the electrical apparatus. Assuming the detector is cooled only to the level required by Equation 3, the resistance required to suppress this noise to an acceptable level can be very high. If the impurity concentration in an extrinsic detector is too high, quantum mechanical processes such as tunneling and hopping allow charge conduction (dark current) in the detector, reducing its resistance in ways that are immune to mitigation by further reduction of the operating temperature.

The low levels of impurity concentration permitted by the electrical requirements necessitate long absorption paths for reasonable quantum efficiency. The resulting detectors have high susceptibility to ionizing particles because of their large volumes. In addition, they are built with an electrode on each side of a high resistance dielectric detector volume. The capacitance of this structure can interfere with the performance as a detector. For example, when a charge is collected at one electrode, the need to inject a charge at the opposite electrode to maintain electrical neutrality is only conveyed at the speed of the dielectric time constant. This parameter is basically the RC time constant of the detector, which can be many seconds in a high performance detector operating at low background (such as on a cooled telescope in space).

An unavoidable consequence, then, of the combined requirements for photon absorption and high electrical resistance is that bulk extrinsic photoconductors at low backgrounds nearly always are operating out of electrical equilibrium. Consequently, their outputs consist of a fast component associated with the immediate release of free charge carriers by absorbed photons, plus a slow component that reflects the adjustment of the detector to the new interior condition imposed by the motion of the charge carriers. This behavior is sometimes described as a memory effect for previously detected sources. For similar reasons, the response of these devices is

modified when a cosmic ray (or other ionizing particle) creates a flood of charge within them. These effects obviously are a challenge for calibration. Nonetheless, such detectors have been used in the *Infrared Astronomical Satellite* (IRAS), the *Infrared Space Observatory* (ISO), and for the far-infrared channels on the *Spitzer Space Telescope* (*Spitzer*). In the latter case, rather elaborate strategies both in the instrument and in the data pipeline have tamed the calibration problems to a large extent.

In addition, there are inherent contradictions in the optical and electrical parameters of some semiconductors with otherwise desirable characteristics for bulk intrinsic photoconductive detectors. For example, InSb has large electron mobility, making it impossible to achieve high resistance in a geometry that allows for a reasonable bias voltage and good quantum efficiency.

The solution is to separate the electrical and optical functions of the device. Virtually all detectors operating between 1 μm and 5 μm are based on this approach. High performance detectors are manufactured as photodiodes in InSb and other semiconductors by forming a junction, with one side doped with donor impurities (n-type) and the other doped with acceptors (p-type). An n-type impurity has an additional electron beyond those needed to complete the crystal bonds, whereas a p-type impurity has one less electron than needed for these bonds. Because the extra electrons are relatively easily freed from the donor atoms, close to the junction they flow to the acceptors and bond with them to complete the crystal structure. As a result, there is a net positive voltage on the donor side and negative voltage on the acceptor side of the junction. The net voltage across the junction is called the contact voltage. It creates a contact potential field that sweeps any free charge carriers across the junction and depletes this region of them, causing it to have high resistance. Photons are absorbed in the material in the vicinity of the junction and produce free charge carriers that diffuse toward it; when they fall into the junction field, they are driven across the junction and produce a current. Because the photons have sufficient energy for intrinsic absorption, the absorption is efficient. To preserve the potential for high quantum efficiency, the absorbing layer must be thin enough for the diffusion to occur much more quickly than the carrier lifetime. Assuming this condition is met, quantum efficiencies approaching 100% can be obtained in devices coated to minimize reflection losses.

InSb was the first material used widely in high-performance near-infrared photodiodes. Its band gap at 77 K provides response out to about 5.6 μm (with a slight dependence on operating temperature, e.g., a cutoff at 5.2 μm at 2 K), thus nicely matching the JHKLM atmospheric windows. The material can be grown with good crystallography, high purity, and excellent uniformity, all of which contribute to its high performance in infrared detectors (Hoffman et al. 2004).

In addition to InSb, $Hg_{(1-x)}Cd_xTe$ is used widely for photodiodes. By varying the elemental composition of this material, the band gap can be changed from 1.55 eV for x = 1, corresponding to a cutoff wavelength of 0.8 μm, to a very small band gap energy, or even a metal. For years, this flexibility in the band gap was the main advantage of HgCdTe for detectors compared with simple semiconductors such as InSb. It made it possible to relax the detector cooling requirements when response was required only at wavelengths less than 5.6 μm (Equation 3). Early methods for manufacturing HgCdTe material had difficulties achieving the level of uniformity of InSb devices

IRAS: *Infrared Astronomical Satellite*

ISO: *Infrared Space Observatory*

with their inherently simpler material. However, superb HgCdTe detectors are now produced by Teledyne using molecular beam epitaxy (MBE; where the detector is built up literally one molecular layer at a time from a molecular beam in vacuum) and by RVS using advanced liquid phase epitaxy (where the crystal is grown from a liquid in which the detector material has been introduced at high concentration). The control now achieved over the internal structure of the photodiodes (see **Figure 1**) allows adjustment of the molecular composition and, hence, the band gap within the detector to improve the performance (Garnett et al. 2004a). In practice, very high-performance photodiodes with cutoff wavelengths in the 5 μm range can be manufactured in either InSb or $Hg_{0.70}Cd_{0.30}Te$, with HgCdTe providing the potential of even higher performance with relaxed temperature requirements where shorter values of λ_c are acceptable—for example, $Hg_{0.55}Cd_{0.45}Te$ provides $\lambda_c \sim 2.5$ μm.

Because of its adjustable band gap, HgCdTe would seem to be the logical approach for devices operating at wavelengths longer than 5 μm. The biggest issue is that material with a small band gap only allows a small contact voltage to maintain the diode depletion region. Therefore, long-wavelength detectors tend to have high dark current. Promising results have been reported by Bacon et al. (2004) for HgCdTe detectors fabricated to have a cutoff wavelength of 9.6 μm and operating at 30 K. They found that half the pixels in the test array had dark current below 0.05 e/s; one-sixth of the pixels with the highest dark current still fell between 4 and 30 e/s. HgCdTe diodes can be made to work reasonably well out to ~15 μm, but with too much dark current for the performance levels demanded by astronomers. If high-performance detectors can be developed with cutoff wavelengths somewhat longer than at present (e.g., 13 μm would allow operation across the entire 10 μm atmospheric window), this technology would find wide application in astronomy.

Beyond 10 μm, excellent performance is achieved using a different approach to separate the electrical and optical functions of extrinsic silicon detectors. These devices are termed impurity band conduction (IBC) detectors (see **Figure 2**). They have been built using Si:Ga (cutoff wavelength ~18 μm; Hogue et al. 2003), Si:As (cutoff wavelength ~28 μm—widely used on groundbased telescopes and in all three *Spitzer* instruments), and with Si:Sb (cutoff ~40 μm) for the *Spitzer* infrared spectrograph (Huffman et al. 1992, van Cleve et al. 1995). Phosphorus is another attractive dopant

MBE: molecular beam epitaxy

Figure 1

(a) Cross-sectional drawing of a HgCdTe photodiode. The junction is at the interface between p- and n-type material. The cap layer is grown with a larger band gap than the infrared-active n-type layer. The CdZnTe substrate is used to support the array of diodes while they are being grown, but can be removed afterward. The lower panel (b) shows schematically how a photon is detected. It is absorbed in the layer of n-type material to produce an electron-hole pair, each member of which then diffuses through the material. If the absorbing layer thickness, l, is small enough, the hole falls into the contact potential field, E, before it can recombine. It is then driven across the depletion region, producing the current that is used for the act of detection. The width of the depletion region is exaggerated for clarity. The upper panel (a) is based on the discussion in Garnett et al. (2004a). The lower panel (b) is based on a figure in Rieke (2003), reproduced by permission of Cambridge University Press.

a

p-type implant

HgCdTe cap layer

n-type HgCdTe

Transparent contact

CdZnTe substrate

Metallization

CdTe passivation

b

l

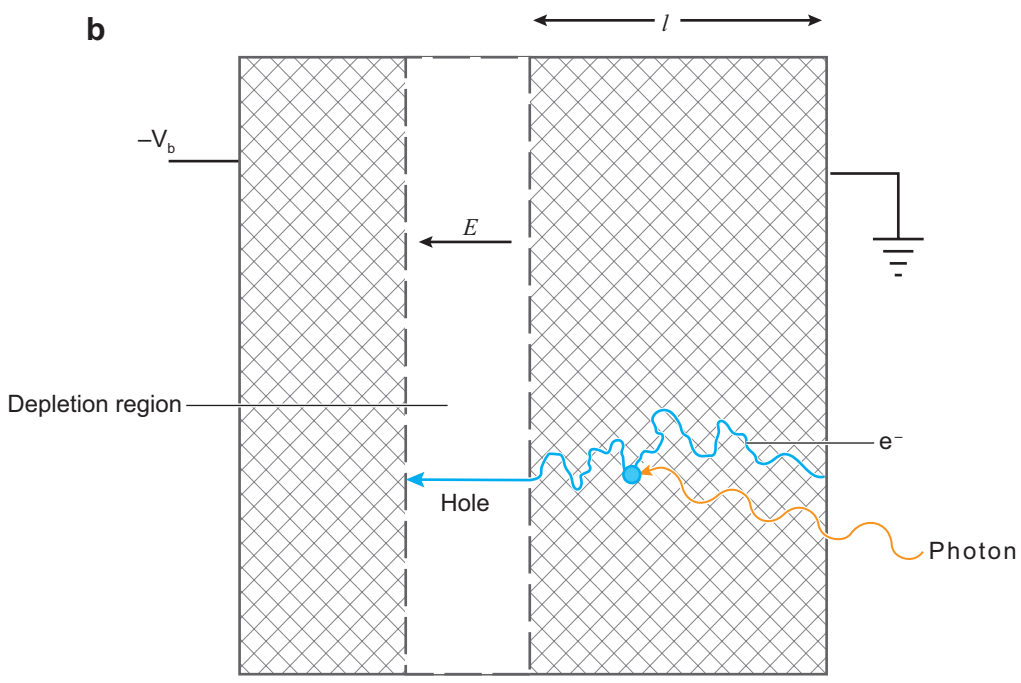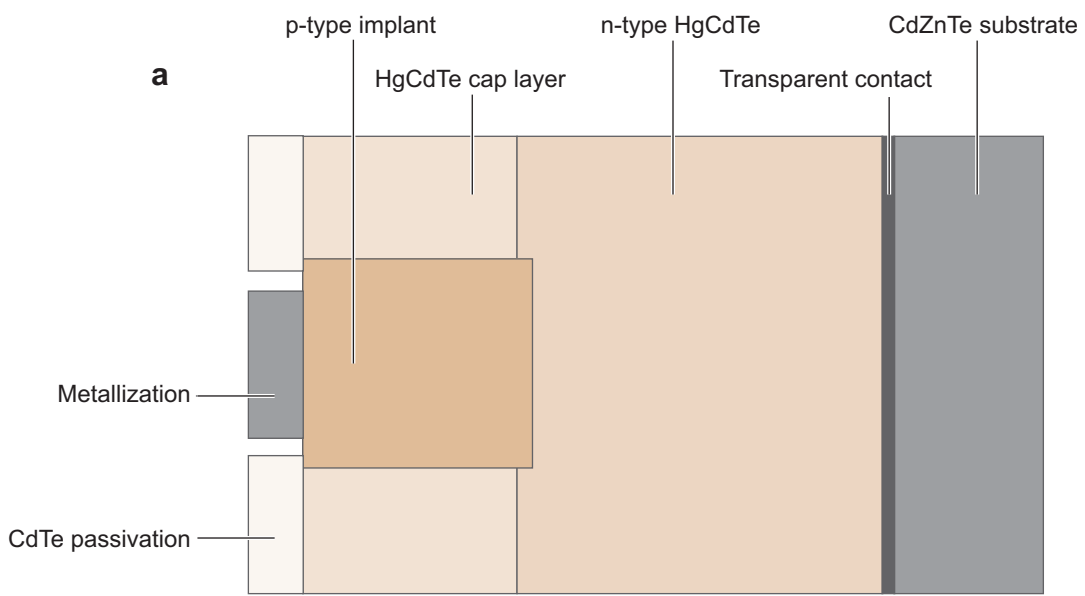$-V_b$

E

Depletion region

Hole

e^-

Photon

Figure 2

Cross-sectional drawing of an impurity band conduction (IBC) detector. The top panel (*a*) shows the physical construction; the V etch is to one side of an array (out of the photosensitive area) and allows a connection to the buried, transparent contact. The upper metallized region is the contact for a pixel. The center panel (*b*) shows schematically the application of the bias voltage to a pixel and the absorption of a photon and generation of a free electron in the infrared (IR)-active layer. The lower panel (*c*) shows the electric field and the extent of the depletion region in the infrared-active layer. For good performance, the bias would need to be increased to complete the depletion of this layer. Figure adapted from Rieke (2003), reproduced by permission of Cambridge University Press.
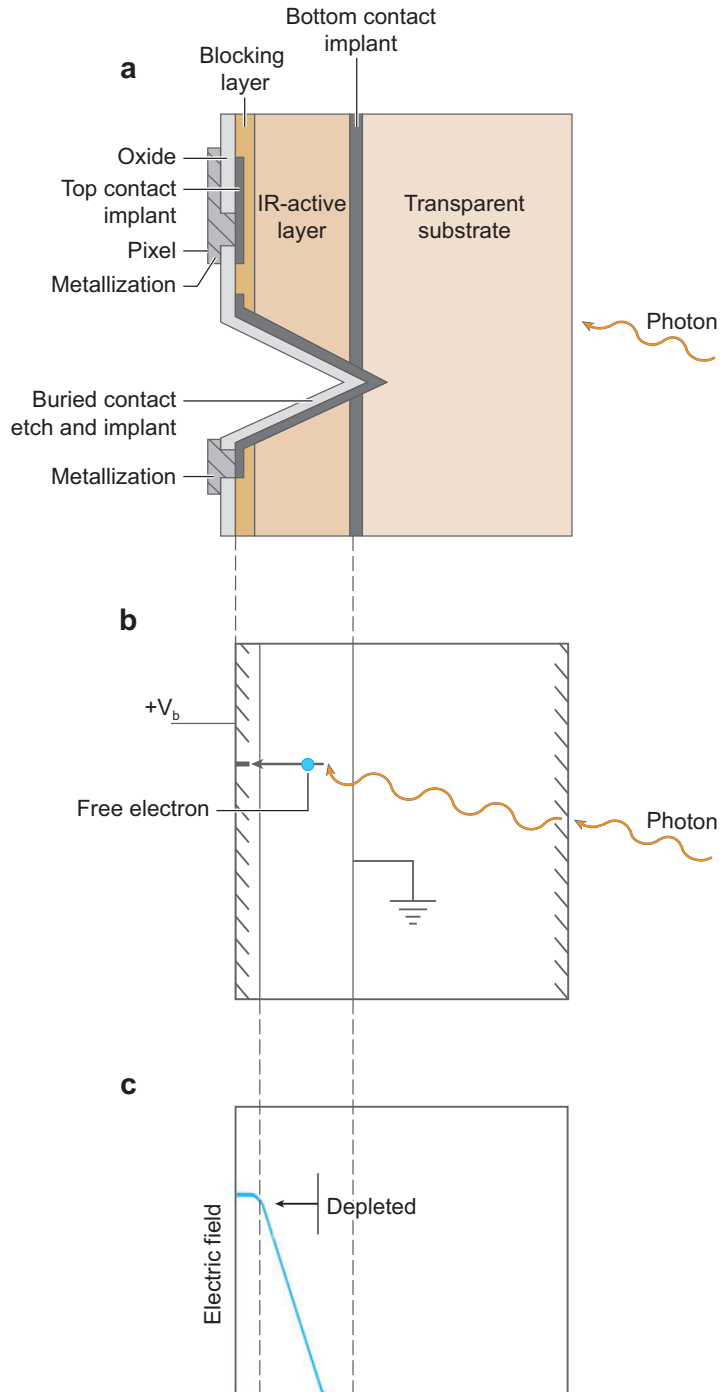
(cutoff ~34 μm; Hogue et al. 2003) because it is widely used in commercial integrated circuits and hence it would be relatively straightforward to fabricate detectors. IBC detectors have also been fabricated with Ge:Ga (Watson & Huffman 1988), Ge:B (Wu et al. 1991), Ge:Sb (Bandaru, Beeman & Haller 2002), and GaAs:Te (Reichertz et al. 2006), which together can extend this detector approach to beyond 300 μm. However, Si:As IBC detectors are the only form currently readily available in large format arrays.

The IBC detector in **Figure 2** consists of a transparent silicon substrate, with a buried transparent contact and then a relatively heavily doped infrared-active layer of thickness 25 μm to 35 μm, followed by an intrinsic layer about 3 μm to 4 μm thick, with the second contact on this layer. The bias voltage on the buried contact is established through a V-shaped etched trough, metal-coated to make it conductive, and placed to one side of the array. Impurity band conductivity through hopping in the active layer would produce unacceptable dark current if the infrared-active band had access to both detector contacts. However, the impurity band terminates at the blocking layer, so leakage can only occur through thermal excitation up into the conduction band (for this discussion, we assume an n-type dopant such as arsenic). Because the intrinsic layer blocks the dark current, the doping in the infrared-active layer can be two orders of magnitude greater than in a bulk photoconductor, enough to provide nearly complete absorption near the peak of the response. The highly doped, thin active layer overcomes the issues with bulk detectors of slow adjustment toward electrical equilibrium and of large cross section for ionizing radiation.

The high conductivity of the infrared-active layer must be suppressed to make it possible to generate a field across it and collect the photogenerated charge carriers. When a positive bias is applied across the blocking layer to the buried contact, negative carriers are collected at the interface between the intrinsic and infrared-active layers and positive ones are driven away from this interface. Thus, a high-resistance region depleted of free charge carriers is produced in the infrared-active material near the interface. However, beyond the depletion region there is virtually no field because of the high conductivity of this material. To gather photoelectrons efficiently that are produced throughout this layer requires that it be entirely depleted. The minority, p-type impurities attach electrons and maintain a negative space charge in this region that tends to cancel the effect of the positive bias. Thus, complete depletion requires that the level of the minority impurities be kept very low, roughly below 2×10^{12} cm^{-3}. That is, the impurity atoms must be at or below 1 in 10^{10} silicon atoms. As we discuss in more detail below, high performance IBC detectors depend on material growth with excellent minority impurity control.

IBC: impurity band conduction

CCD: charge coupled device

3. READOUTS

Initially, infrared arrays used silicon charge coupled device (CCD) readouts. This approach is not compatible with low read noise because of the low operating temperature required by the infrared detectors. Below about 70 K, the buried channels used to avoid trapping noise in high performance CCDs freeze out—there is no longer sufficient mobile charge to maintain the channel. The devices then operate in surface channel mode, with accompanying high noise. In addition, the sponsors for the

initial development of infrared arrays were worried about damage to the devices under extremely high doses of ionizing radiation, which can degrade the charge transfer in CCDs. Therefore, a different architecture was developed in which a readout amplifier is dedicated to each pixel, and transistor switches bring the signals to an output amplifier. Ironically, the development of these devices to replace CCD readouts in the infrared has come full circle with CMOS (complementary metal oxide semiconductor) detector arrays. They operate on the same amplifier-per-pixel basis and are replacing CCDs in many applications in the optical (for a review of CMOS detector arrays, see Hoffman, Loose & Suntharalingam 2005).

All modern near- and mid-infrared arrays have individual-amplifier-per-detector readouts based on metal oxide semiconductor field effect transistors (MOSFETs). MOSFETs are grown on a thin surface layer of a silicon substrate by laying down a pattern of insulators and conductors. They consist of a channel doped to allow current flow between two electrodes, a source, and a drain. A third electrode, the gate, is placed on a thin insulator over the channel. Any current that deposits electric charge onto the gate creates a field that modulates the current flow in the channel, thus amplifying the gate current. MOSFETs are the building block for integrated circuits in general, and when run cold have nearly ideal properties for receiving the tiny signals from very high resistance infrared detectors. They require very little power, have virtually infinite input resistance, can be turned on and off without hysteresis, and can be constructed into very large and complex circuits.

In an infrared array, an individual detector amplifier collects the charge generated by the detector on the combined capacitance of the detector and the MOSFET gate. The MOSFET amplifies this signal to provide enough drive power to convey it over circuit traces for reasonable distances. The ensemble of readout MOSFETs directs

MOSFET: metal oxide semiconductor field effect transistor

$\longrightarrow$

Figure 3

This figure shows a readout circuit—four cells of a detector array (the detectors are shown schematically as diodes). Signal is collected from the photodiode in the form of a current that deposits charge on the gate of transistor T_1 until it is judged time to measure the integrated level of charge. To read it out, power is applied to the row driver R_1 and, at the same time, to C_1. The transistors T_2, T_3, and T_4 conduct current as a result, and apply power to T_1 as well as connecting it to the output bus, which connects the signal to the output amplifier of the array (*lower left* in the figure) where it can be measured with an external circuit. There is now a choice. If one wants to continue integrating the signal, power is removed from C_1, and T_2, T_3, and T_4 turn off, removing power from T_1, so the pixel can continue to accumulate charge on its gate. Possibly, power would be applied to C_2 to read out the next pixel in the row of the array. In any case, these steps provide a nondestructive read of the upper left pixel, because they allowed determination of the level of detected charge without disturbing it. In the second case, one resets the collected charge and initiates a new integration. To do so, the reset line (below the row bus in the figure) is pulsed while T_2, T_3, and T_4 are still on, which sets the integrating node (the input to T_1) to the voltage V_R. Because the integrated charge is lost in this operation, and assuming one reads out T_1 before the reset, this operation has caused a destructive read. It is therefore possible to address each pixel in the array individually, read out the signal it has accumulated, and either continue through the array or reset the signal for a new integration. Figure from Rieke (2003), reproduced by permission of Cambridge University Press.

the signals from each detector sequentially to a small number of output amplifiers, multiplexing the signals down to a small number of output wires. In fact, the readouts are often termed multiplexers or MUXs. They are implemented in readout integrated circuits (ROICs), fabricated in standard integrated circuit foundries. The simple array circuit in **Figure 3** illustrates how the readout functions are performed.

In principle, the circuit in **Figure 3** allows addressing any pixel in the array, or any sequence of pixels in any order. Pixels can also be reset individually. This high degree of flexibility is usually considered to be too much of a good thing, and the array is controlled by on-board circuits that advance the signals in response to a simple clock pulse. Also, the readout amplifier can be simplified by eliminating the capability to reset individual pixels and therefore allowing T_3 to be removed from the circuit.
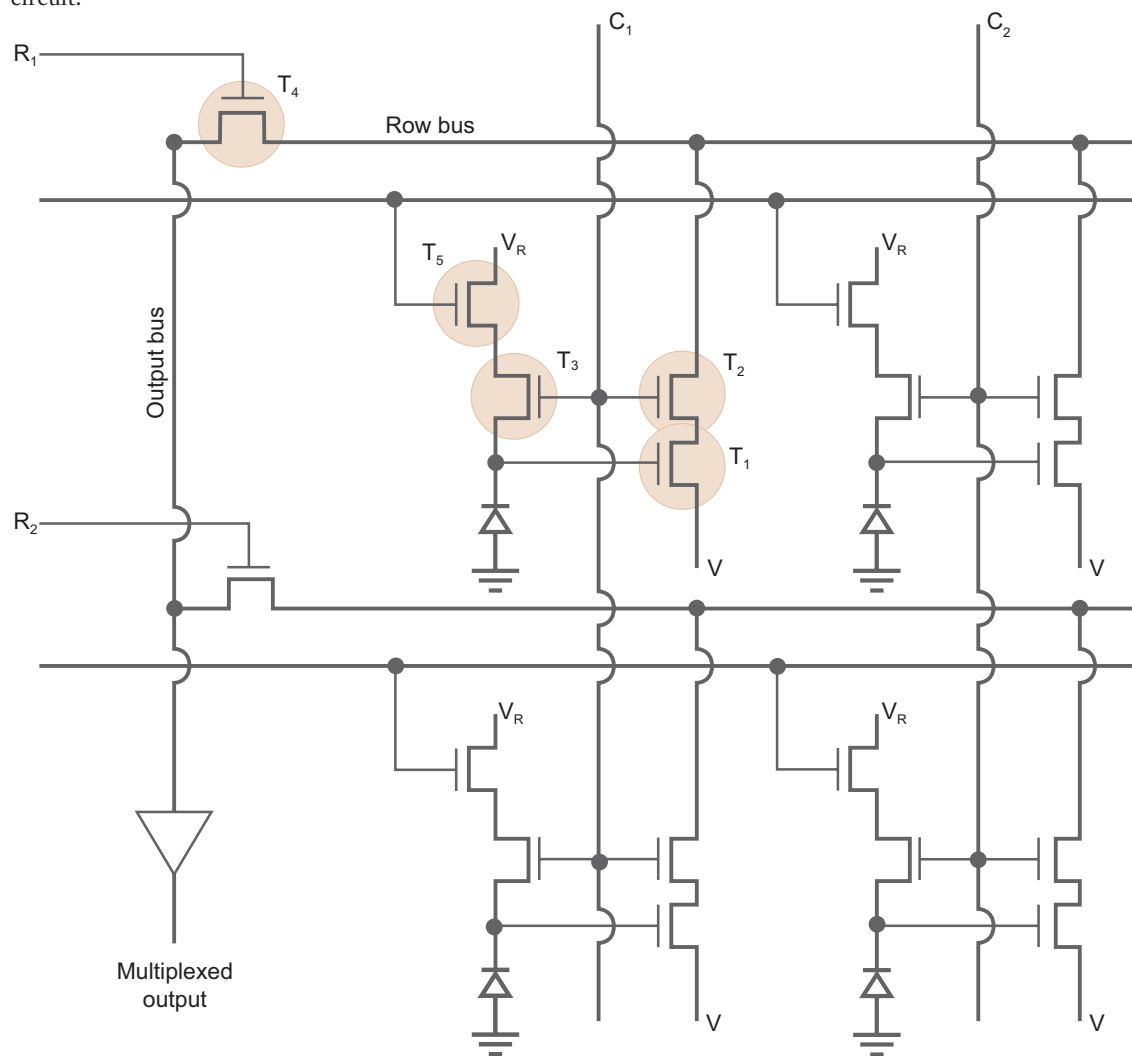
The fact that an integrating MOSFET can be turned on and then off without disturbing the charge collected on its gate, combined with the inherent flexibility of the circuit, allows for a number of ways to read out the array. One can address a pixel, read the accumulated signal, and then reset it, counting on the reset level always to be the same. If this degree of faith is lacking, one can read, reset, and then read the reset level and determine the signal as the difference. However, both of these approaches are subject to kTC, or reset, noise. It is caused by thermally driven currents that produce a randomly fluctuating charge level of average level $\sqrt{(kTC)}$ on the integrating capacitor, C. These unwanted signals fluctuate on a timescale $\tau = RC$, where R is the resistance attached to the integrating capacitor. When the reset switch is closed, R is small and the fluctuations are fast. However, when the switch is open, R is huge for a high-performance detector and the fluctuations are frozen in place. Thus, a better low-noise readout strategy is to reset the amplifier, read its output after the reset switch has been opened and the amplifier has settled, then turn it off and integrate, and finally turn it back on and measure its output at the end of the integration before it is reset. The signal is the difference of the two reads of the amplifier.

A disadvantage of this approach is that it requires enormous DC stability of the array and its readout electronics, so that any drifts are less than the tiny signals corresponding to the collection of just a few electrons. Fortunately, modern arrays and electronics are capable of such stability. Improved performance can be obtained, however, by including in the array a number of reference pixels that are electrically identical to the detector amplifiers but do not respond to light. The signals from these pixels are subject to many of the same sources of slow drifts that apply to the other amplifiers of the array. The output of the reference pixels is averaged (to increase the ratio of signal to noise) and can then be used in image processing to correct for drifts.

The ability to read out an amplifier without disturbing the signal can be used to enhance array performance. One approach is to read out each pixel of an array a number of times, N, in rapid succession at the beginning of an integration, then to turn off the array to minimize its power dissipation, and then to turn it back on and read it N times at the end of the integration ramp. Assuming that the read noise is dominated by relatively high frequency noise in the amplifier, which is usually the case, averaging the multiple reads can reduce the effective read noise as opposed to just taking single reads at the beginning and end. For moderate values such as N = 8 or 16, the effective read noise can often be reduced by nearly the square root of N. This approach is sometimes called Fowler sampling after Al Fowler, who was an early advocate. A standard terminology is Fowler-N sampling. A similar benefit can be obtained by turning the ROIC of the array on and reading the detector outputs periodically throughout the integration as signal is accumulated—an approach sometimes called multiaccum, or sampling up the ramp. Although for a given number of reads, this latter approach is theoretically less effective than Fowler sampling, it has an important additional benefit. If the integration is upset—say by an ionizing particle striking the detector, or because the readout amplifier experiences an output jump—it is possible to identify the approximate time of the event and reject it, while saving most of the integration before and after it. A composite scheme can also be used, in which N

repetitive reads are repeated in bursts throughout a sample-up-the-ramp integration. This latter approach can combine the benefits of both multiple read techniques.

4. HYBRID CONSTRUCTION

Silicon is the material of choice for the readouts, regardless of the material used for the detectors. In theory, it would be possible to build infrared arrays using the detector semiconductor type for the readout electronics. However, these materials generally do not have the properties needed for high-performance electronics, such as the easily formed rugged oxide of silicon that allows robust insulating layers. In addition, silicon readouts benefit from the huge investment in fabrication of integrated circuits on this material for commercial applications.

To complete an array requires a method to attach a grid of detectors to a matching grid of silicon-based amplifiers. The method generally used for direct hybrid arrays is illustrated in **Figure 4**. The amplifiers are grown on a thin wafer of silicon with a grid of input contacts, one for each amplifier. The detectors are grown on a wafer of suitable semiconductor, with a matching grid of output contacts. Indium bumps are deposited on each grid of contacts, the detector and amplifier grids are carefully
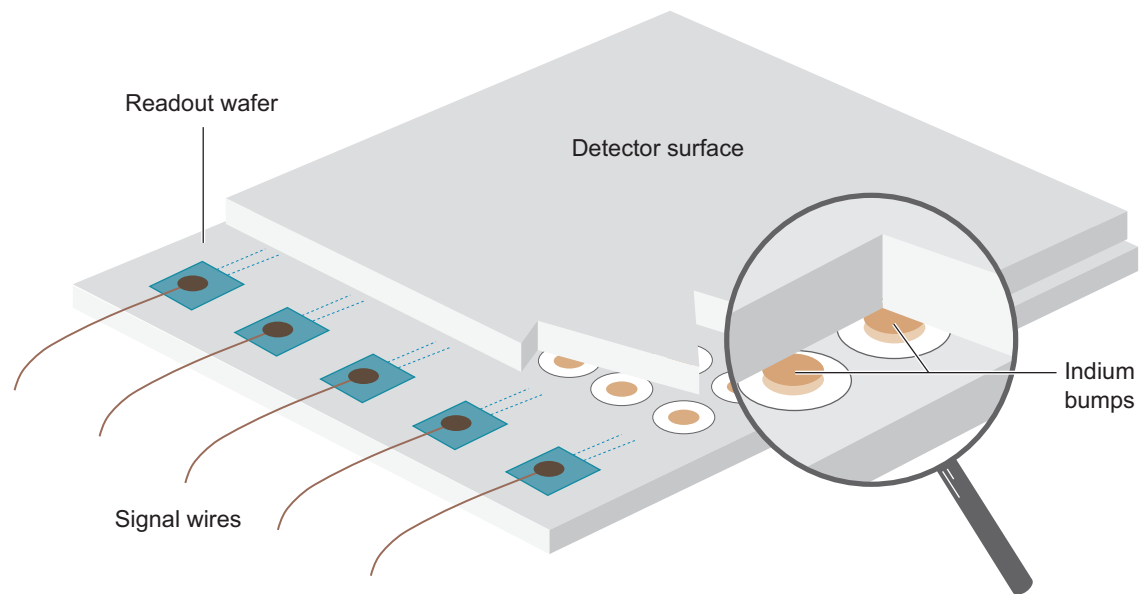

Figure 4

In direct hybrid construction, a wafer of detectors is attached to a silicon wafer (carrying the readout amplifiers and associated circuitry) through matching grids of indium bumps. When the wafers are aligned and pressed together, the bumps distort, their indium oxide skins crack, and the exposed indium metal welds the detector outputs to their individual amplifier inputs to complete the array. Figure from Rieke (2006), reprinted by permission of the University of Arizona Press.

aligned, and a large force is applied to push them together. Indium is soft, and when the indium bumps deform, their indium oxide skins fracture and the exposed metal cold welds, attaching each detector to its matching amplifier. This process to produce a many-pixel array is not simple because of the large forces that must be controlled precisely, without the occurrence of sideways motion, and while maintaining the planarity of the device. Another issue is that the thermal contraction of the detector material on cooling generally differs from that of the readout (unless, of course, the detectors are made of silicon). The resulting forces can rip apart the indium bump bonds, particularly at the corners of an array where the stresses are largest because the distance from the array center is largest. Something has to give. To allow the necessary stretching, either the detector wafer is thinned (chemically, or by diamond-tool machining after the array has been hybridized), or the readout is thinned. Although it can readily be shown that typical detector materials would rip apart from their silicon readouts at an array size of about 1 cm square for rigid parts, arrays are now made up to about 3.6 cm^2 (13 cm^2 in area) with appropriate thinning.

5. ARRAY EXAMPLES

Teledyne/Rockwell and RVS supply the majority of infrared arrays used in astronomy; DRS Technologies fills nearly all the remainder of the market (Simons et al. 2006). The characteristics of the most advanced infrared arrays for astronomy from these manufacturers are summarized in **Table 1**. The parameters have been taken from a variety of sources without establishment of a common test protocol, so small differences are not necessarily real. Operating conditions also influence the performance measures. For example, the DRS Technologies Si:As IBC array has been built to operate with solid hydrogen cryogen, so its dark current is a reflection of the relatively high operating temperature for its detector type.

Even under nominally similar operating conditions, array performance may be subject to variations depending on test details. Hall (2006) illustrates how the Teledyne HAWAII-2RG is sensitive to temperature changes at the 1 mK level. With care to maintain this level of stability, he finds that the read noise can be driven down to about 4 electrons (Fowler-32 sampling) and the dark current to 0.002 e/s. The problem is that such stability may be unachievable outside of the laboratory. Nonetheless, some of the benefits implied by these results can be achieved by operating arrays under the most uniform possible conditions. Specifically, the clocking cadences (the sequence and timing of control signals to the array) should be kept extremely uniform and other steps taken to maintain operating arrays in the most benign and constant environment possible.

Array performance is also strongly affected by background levels, which cover a huge range for astronomical infrared arrays. Moderately high-resolution ($\lambda/\Delta\lambda \sim 3000$) spectrometers operating on the ground in the near-infrared between the atmospheric airglow OH lines, or even more so in space, can deliver tiny background signals of 0.1 photons/(s m^2 arcsec^2), or less for higher spectral resolution. An

Table 1   Array Characteristics

Parameter	RVS VIRGO/VISTA	Teledyne H2RG JWST	RVS Orion	DRS Technologies WISE	RVS JWST
Detector type	HgCdTe	HgCdTe	InSb	Si:As IBC	Si:As IBC
Wavelength range (μm)	0.85–2.5	0.6–5.3	0.6–5.5	5–28	5–28
Format	2048×2048	2048×2048	2048×2048	1024×1024	1024×1024
Pixel pitch (μm)	20	18	25	18	25
Operating temperature (K)	78	37	32	7.8	6.7
Read noise (e rms)	6 (slow readout)	6 (slow readout) 30–40 (fast readout)	6	42 (Fowler-1; lower noise expected with more reads)	10
Dark current (e/s)	<0.1	<0.01	0.01	<5	0.1
Well capacity (e)	$>1.4 \times 10^5$	8×10^4	1.5×10^5	$>10^5$	2×10^5
Quantum efficiency	>70%	>80%	>80%	>70%	>70%
Outputs	4, 16	1, 4, 32	64	4	4
Frames/sec	0.4, 1.5	0.1 to 30	10	1	0.3
References	Love et al. (2004) Bezawada & Ives (2006)	Rauscher et al. (2004) Garnett et al. (2004)	McMurtry et al. (2003) Fowler et al. (2004)	Mainzer et al. (2005a)	Love et al. (2005)

Note: IBC, impurity band conduction; JWST, *James Webb Space Telescope*; RVS, Raytheon Vision Systems; VISTA, Visible and Infrared Survey Telescope for Astronomy; WISE, *Wide-Field Infrared Survey Explorer*.

imager in the OH-dominated H-band atmospheric window (1.6 μm) and at $\lambda/\Delta\lambda \sim 4$ will encounter of order 10^4 photons/(s m^2 arcsec^2). An imager operating in the 20-μm atmospheric window at spectral resolution $\lambda/\Delta\lambda \sim 4$ may encounter a background approaching 10^{10} photons/(s m^2 arcsec^2), whereas a similar imager on a cold telescope in space will see of order 10^4 photons/(s m^2 arcsec^2).

The examples in **Table 1** illustrate the features of available infrared detector arrays. Although all infrared arrays operate on similar principles, these examples have been optimized for low backgrounds (e.g., up to a few thousand photons per pixel per second). In such devices, the read noise in electrons is minimized by minimizing the detector plus gate integrating capacitance, C, so by $V = Q/C$ the voltage amplitude, V, for a given charge, Q, is maximized. Unfortunately, this optimization results in saturating the readout amplifiers if the detectors are exposed to high background levels. Solutions include simply increasing the integrating capacitance, or using alternative amplifier architectures that can handle relatively large signals. There is an inevitable degradation in the read noise, but at high background levels this loss is usually inconsequential.

The funding for astronomical arrays has generally been for low-background applications, such as cooled telescopes in space, so arrays with low read noise and dark current are more fully developed than are devices for high backgrounds. Pixels in conventional groundbased imagers may encounter thermal backgrounds up to

10^9 photons/second. There are no modern arrays optimized for such backgrounds in the 3 μm to 5 μm region, and the largest available high-background Si:As IBC arrays are 256×256 or 240×320 pixels. However, there are plans to use the investment in low-background 1024×1024 devices at DRS Technologies to leverage high-background arrays of the same format (Mainzer et al. 2005b; see also Love et al. 2006).

JWST: *James Webb Space Telescope*

5.1. Teledyne HgCdTe Arrays

Teledyne produces a variety of very high performance HgCdTe arrays. They include: (*a*) a line of devices initially used in the Near Infrared Camera and Multi-Object Spectrometer (NICMOS) on the *Hubble Space Telescope* (e.g., Rieke et al. 1993) and then in groundbased astronomy (including arrays far more advanced than the NICMOS ones), mostly with a cutoff at 2.5 μm; (*b*) the devices derived from the ground-based HAWAII-1 series (Hodapp et al. 1996) for the Wide Field Camera 3 on the *Hubble Space Telescope*, with a cutoff wavelength of 1.72 μm (Robberto et al. 2004); and (*c*) arrays developed for the *James Webb Space Telescope* (JWST), with cutoff wavelengths of 2.5 μm and 5.3 μm (as in **Table 1**). The detectors for the current line of devices are manufactured by MBE and have architecture similar to the diode shown in **Figure 1**. We describe in detail the 2048×2048 pixel HAWAII-2RG, but the other members of the family are generally similar.

Referring to **Figure 1** and Garnett et al. (2004a), growth of the detectors is carried out on a CdZnTe substrate. This substrate material has been selected because it has an excellent match to the HgCdTe crystal structure, which is important to minimize crystal defects in the grown material. Defects in the detector crystal structure can act as traps for free electrons and thus degrade the performance. A previous generation of detectors from Rockwell used sapphire as the substrate, which provides a poorer match.

First, the bulk HgCdTe is laid down, with the Hg and Cd fractions selected for the appropriate band gap and cutoff wavelength. These first layers of material will eventually become the back side of the detector through which the photons enter. This material is doped with indium to make it n-type. The final stages of detector growth cap the device with a HgCdTe layer with higher band gap energy. This cap avoids having open bonds and crystal structure flaws at the detector front side. Because of the higher band gap, it provides a potential barrier that reflects minority (photogenerated) carriers from the detector front side into the bulk material so they can diffuse into the junction (a compositional gradient in the bulk material can also help drive minority carriers away from the front side). This control of minority carriers reduces cross talk and latent images and improves the quantum efficiency.

Once the layers have been grown, the pixels are defined by using arsenic to create p-type implants. The diode junction is formed at the interface between the implanted region and the bulk n-type material. The detectors are completed by depositing a protective layer of CdTe to passivate the device, and by depositing a metal contact on each p-implant. Indium bumps are evaporated onto the metal contacts, in preparation for hybridization. The CdZnTe can be removed after the detector array is complete,

because it blocks visible light photons and can also glow when struck by cosmic rays.

For these arrays, the readout is fabricated in a high-volume state-of-the-art commercial semiconductor foundry operated by United Microelectronics Corporation, using the standard foundry processes with 0.25-μm design rules. The outer four pixels around the periphery of the ROIC are used as reference pixels, so the light-sensitive area is 2040 × 2040 pixels. The readout amplifiers are similar to those in **Figure 3**, but with additional MOSFETs to implement a sophisticated subarray readout mode. The ROIC contains a substantial amount of logic circuitry (MacDougal et al. 2006). This logic allows selection of 1-, 4-, or 32-output amplifiers, trading readout speed against power dissipation as the number of amplifiers increases. In addition, it is possible to select a subarray section and read it out independently of the rest of the device, taking the signals to a dedicated output amplifier. One use of this capability is to allow centroiding on a stellar image to determine guiding corrections while taking long exposures (e.g., Riopel, Teeple & Ward 2006). One can program different reset options, and select output gains between 1 and 16 times. The logic also allows for on-board identification and correction of single-event upsets from ionizing radiation.

Each short wavelength optical train for the Near Infrared Camera on JWST will use four of these arrays mounted in a mosaic as illustrated in **Figure 5**.
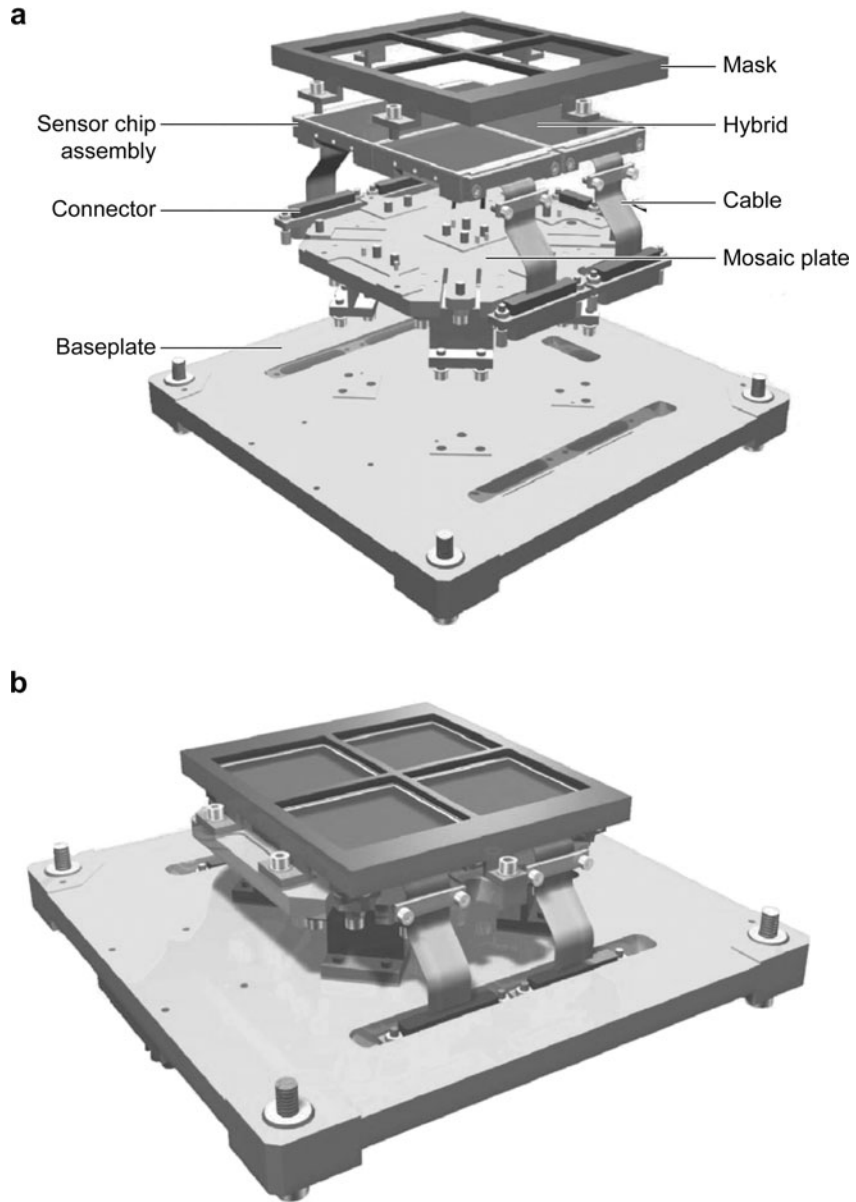
5.2. Raytheon Vision Systems Si:As IBC Arrays

We now describe the 1024 × 1024 pixel Si:As IBC array being provided for the Mid-Infrared Instrument on JWST by RVS and shown in **Figure 6**. This device illustrates the characteristics of other arrays using this detector type, which utilize similar detector processing and similar designs for their readouts.

These detectors are designed for the best broad-band response, i.e., highest possible infrared absorption, because they are to be used down to 5 μm, where the Si:As intrinsic absorption coefficient is relatively low. The detectors must also be designed to deplete the infrared-active layer fully. Otherwise, the quantum efficiency would be reduced; furthermore, with simple readout amplifiers as in **Figure 3**, the output would be highly nonlinear as charge accumulates and changes the bias on the detector. The resulting design constraints are discussed by Love et al. (2004a) and can be understood by refering to **Figure 2** and the surrounding discussion. The goal is to produce the thickest possible infrared-active layer doped as heavily as possible, subject to the limitation that this layer should not become so conducting that the detectors have a large dark current. This limit is reached at an arsenic concentration of about 10^{18} cm^{-3}. The layer thickness is set by the (unwanted) minority impurity concentration and by the requirement that the bias voltage should not be so high that it causes incipient avalanching of charge carriers and hence increased noise. The minority impurity concentration upper limits are 1.44×10^{12} cm^{-3} for a 45-μm thick layer, and 1.85×10^{12} cm^{-3} for a 35-μm thick layer (Love et al. 2004a). The detectors were therefore designed for an arsenic doping level of 7×10^{17} cm^{-3} and a thickness of 35 μm. The processing at Advanced Technology Materials, Inc., yielded devices very close to these targets. Measurements of the material characteristics, plus evidence

of full depletion at about 2 V (Love et al. 2006), show that the minority impurity
level is probably slightly above 10^{12} cm^{-3}, that is, within the range that yields high
performance detectors.

The readouts for these devices use a circuit similar to that in **Figure 3**.
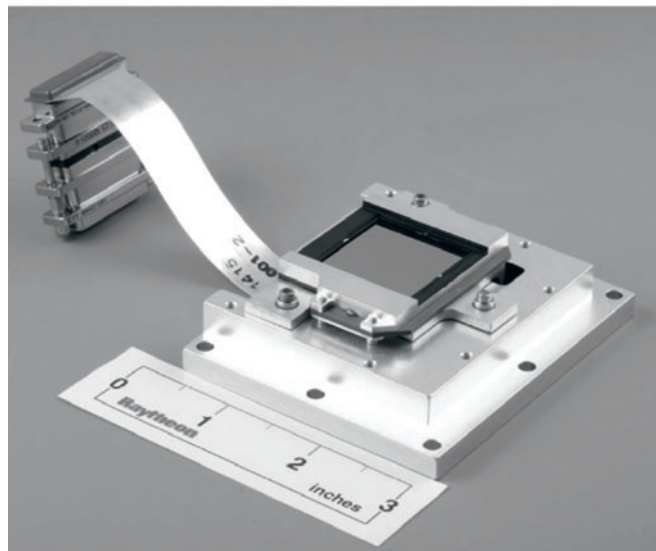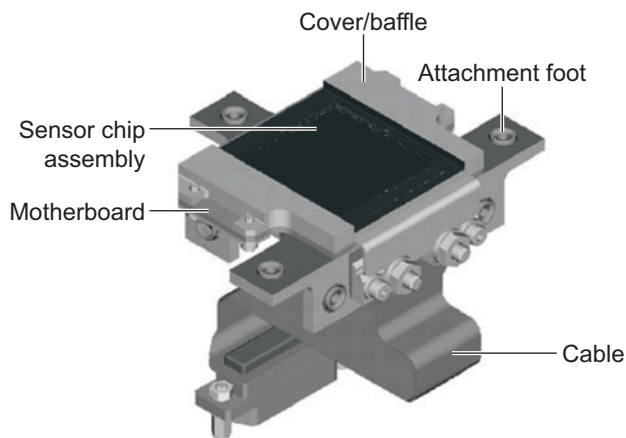Two columns of reference pixels are provided at either side of the 1024×1024

light-sensitive area. An additional reference output is also provided that does not go through the multiplexing for the array pixels and hence can be used to quantify high-frequency noise.

Silicon-based MOSFETs show a number of operational difficulties at the very low temperatures required for the readout circuits for these detectors (Glidden et al. 1992, Ando et al. 2003). They are all related to freeze-out of thermally generated charge carriers in implants and elsewhere, making the circuits unstable, increasing noise, and causing signal hysteresis. For example, any bias applied to the MOSFET gate sets up a potential across the substrate carrying the MOSFET that must be neutralized by thermally generated charge carriers in the substrate. If the substrate is frozen out, such neutralization can take a very long time, leading to hysteresis and poor DC

stability. Many of these issues can be mitigated by growing the circuits on wafers that are heavily—degenerately—doped except for a thin (3 μm to 6 μm) surface layer of high-purity silicon to carry the circuits and isolate them from the bulk material (Glidden et al. 1992). The heavy doping increases the concentration of free charge carriers in the substrate so it is conducting even at very low temperature. As a result, it can adjust quickly to changes in the MOSFET operating conditions. Such substrates can also conduct away the charge that would otherwise build up as a result of impact ionization at the MOSFET drain electrode. In addition to this wafer design, a number of modifications in the detailed circuit layout can also improve performance. Glidden et al. (1992) discuss some possible approaches in detail. The resulting devices can be DC-stable and have low noise down to 1.5 K or below (Young et al. 1995).

RVS has established a process optimized for deep cryogenic operation at the Supertex foundry, where the readouts have been fabricated. The ROICs are built with 2-μm design rules. Readouts with fine design rules (such as the 0.25 μm used for the Teledyne arrays) can also function at these low temperatures without the deep cryogenic substrate design, in part because of the high current densities in their small features. However, their noise tends to be higher than has been achieved for the deep cryogenic readouts (see, e.g., Lum et al. 1993 for a comparison of the noise with early versions of these approaches).

5.3. Other Near- and Mid-Infrared Arrays

The first true infrared detector arrays to fly on a space astronomy mission were the two channels in the ISOCAM instrument on the ISO. The longer-wavelength device worked well and was responsible for a substantial part of the science return of the mission. It was based on bulk Si:Ga photoconductors in a 32×32-pixel array readout by a MOSFET amplifier array similar to those discussed above (Lucas et al. 1988; Mottier, Agnèse & Lagage 1991).

There are a number of other arrays of interest to astronomers. RVS produces its own 2024×2024-pixel 2.5 μm cutoff HgCdTe VIRGO arrays (Love et al. 2004b; see **Table 1**). Like the Teledyne arrays, the detectors are grown on CdZnTe substrates, but using liquid phase epitaxy. If the substrate is removed and they are antireflection coated, they have very high quantum efficiency from the red to 2.5 μm. They also have excellent uniformity and can operate well up to temperatures of 70 K to 80 K. The readout design is modular in units of 512×1024 pixels, so it can support other array formats, and it can be used with HgCdTe detectors with other cutoff wavelengths. Sixteen of these arrays have been delivered in an open mosaic for the Wide-Field Visible and Infrared Survey Telescope for Astronomy (VISTA) imager; the devices can also be butted closely on three sides to make filled mosaics. In addition, RVS supplies a high-performance InSb array, also listed in **Table 1**.

The Si:As IBC array from DRS Technologies, described in **Table 1**, benefits from the close interaction of the detector group with their supplier of Si:As detector wafers, Lawrence Semiconductor Research Laboratory, Inc. This team has developed the means to control the minority impurity concentrations to levels significantly below 10^{12} cm^{-3}. Therefore, small bias voltages can fully deplete the infrared-active layer

VISTA: Visible and Infrared Survey Telescope for Astronomy

with minimum dark currents, allowing operation at relatively high temperatures. This attribute is important for use in the *Wide-Field Infrared Survey Explorer* (WISE) where the focal plane will be cooled in a solid hydrogen cryostat to 7.0 ± 0.5 K (Mainzer et al. 2005a).

In space, dislocation damage from cosmic ray hits within the blocking layer can degrade the dark current in Si:As IBC detectors (e.g., Heras et al. 2000). The small required bias for the ultrahigh purity DRS detectors improves the tolerance to such damage. The readout for this array is manufactured in a standard commercial foundry using fine design rules, rather than the optimized deep cryogenic architecture for the RVS array. The readout works well, and there is the potential that further optimization of its operation can bring the noise down into the range of the deep cryogenic devices (e.g., as for the JWST Si:As IBC array in **Table 1**).

5.4. Far-Infrared Photoconductive Arrays

At wavelengths beyond 40 μm, there is no suitable silicon-based photoconductor technology for high-performance arrays. Instead, IRAS, ISO, and *Spitzer* have all used bulk germanium photoconductors. The first two missions used individual sensors operated in parallel. However, for *Spitzer* a true 32×32-pixel array was developed for the 70-μm band (Young et al. 1998). To allow for the long absorption path in this material, the detectors are illuminated edge-wise with transverse contacts, and the readouts are hidden behind them (see **Figure 7**). This construction style is sometimes called Z-plane to indicate that the array has substantial size in the third dimension. The readouts use an ancestor of the RVS cryogenic process (discussed at the end of Section 5.2) to provide the necessary stability to maintain the small (45 mV) but sensitive bias settings required by these detectors. In addition, the amplifier design is modified from the simple integrator in **Figure 3** to include electrical feedback to stabilize the detector bias during integrations.

The poor absorption of the Ge:Ga detector material requires that the detectors in this array be huge—2 mm long. As a result, they have large cross sections for hits by ionizing particles, and in fact are hit on orbit every 10 s to 15 s. When a pixel is hit, electrical feedback on its readout amplifier quickly returns its bias to the nominal value (within a fraction of a second). As a result, cross talk of the hit to neighboring pixels is minimized. In addition, it is possible to recover much of the data in an integration both before and after the hit; there is a transient on the integration ramp, but it quickly returns to normal.

Far-infrared photoconductor arrays suffer from the standard bulk photoconductor photometric issues. In *Spitzer*, these problems have been addressed by flashing a calibrator to measure the response of all the pixels approximately every two minutes. The data pipeline interpolates between these calibrations to provide an instantaneous response map of the array. When measuring point sources, the pipeline removes a column-by-column offset determined by median averaging along the columns. It also allows digital filtering in the time domain to suppress the slow response components. The final result is photometry that repeats to about 5% (Gordon et al. 2005, 2007).
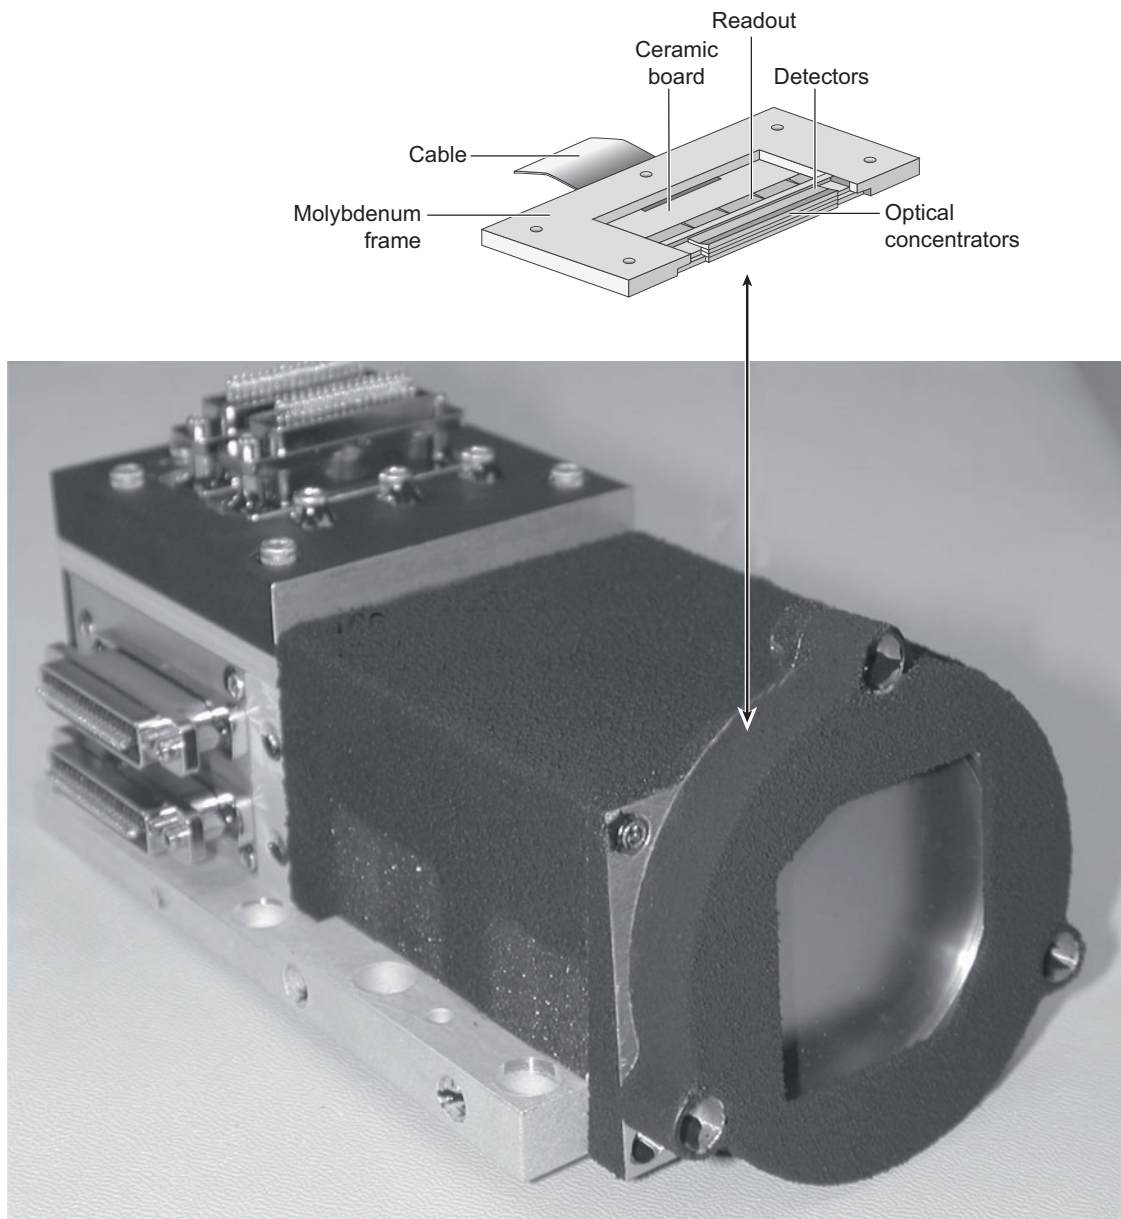
Figure 7

The *Spitzer Space Telescope* 70-μm array. The line drawing shows the design of a 4 × 32 module of the edge-illuminated detectors. Eight such modules are stacked to create the full 32 × 32 array. The sensitive area of the array is 24 mm on a side.

Akari is currently observing with four small arrays of bulk germanium photo-conductors (30 to 60 pixels each), covering together the 50 μm to 180 μm spectral range (Fujiwara et al. 2003, Shirahata et al. 2004). Two 16×25-pixel arrays of bulk germanium photoconductors will also be flown in the Photodetector Array Camera and Spectrometer (PACS) on Herschel (Poglitsch et al. 2003, Birkmann et al. 2004). Despite the photometric issues, these detectors have the advantages of rugged construction and of operation at relatively high temperature (about 1.5 K). However, a different detector type, bolometers, is capable of better performance so long as the detectors can be cooled to sub-Kelvin temperatures.

6. BOLOMETER ARRAYS

6.1. Bolometer Principles of Operation

Bolometers operate on a different principle from the detectors discussed above. Rather than individual photons creating free charge carriers, photons are absorbed and thermalized, and the resulting energy is sensed. This approach yields very high performance detectors for the submillimeter and millimeter spectral regions.

Bolometers are based on an absorber that is isolated by a thermal link from a heat sink, as shown schematically in **Figure 8**. Photons incident on the absorber raise its temperature, causing a sensitive thermometer attached to it to change resistance, producing a signal that can be amplified to achieve a detection. The strength of the thermal link is G (in W/K). Obviously, for a given power input, the temperature excursion that produces the signal varies in proportion to $1/G$.

Bolometer performance is usually described in terms of Noise Equivalent Power (*NEP*). The *NEP* is the signal power that yields a rms signal to noise of unity into a frequency bandwidth of 1 Hz (the smaller the *NEP*, the better the performance of the detector at low light levels). Bolometers must be operated at low temperatures to suppress thermal noise that arises as a result of thermodynamic fluctuations in the
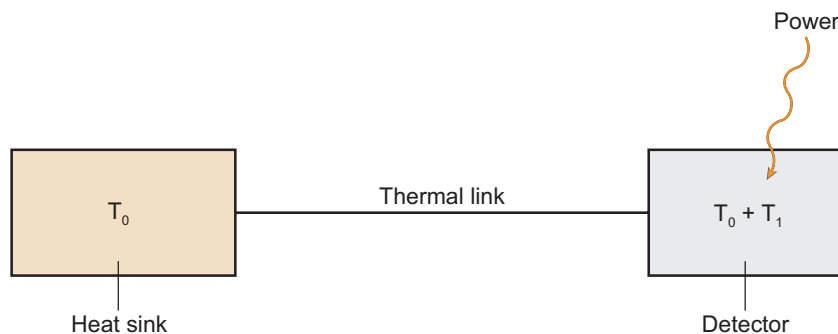
Figure 8

Thermal model of a bolometer. From Rieke (2003), reproduced by permission of Cambridge University Press.

flow of energy across the thermal link:

$$NEP_T = \frac{(4kT^2G)^{1/2}}{\eta}, \tag{5}$$

JFET: junction field effect transistor

where k is Boltzmann's constant, T is the temperature, and η is the quantum efficiency. Bolometers operating at sub-Kelvin temperatures can be designed so the thermal noise is the ultimate limit, but bolometer noise may also have a significant contribution from Johnson noise. The corresponding component of the *NEP* is roughly proportional to T^2G, even more steeply dependent on temperature than NEP_T. A logical reaction to Equation 5 (and to our earlier statement that the signal-producing temperature excursions grow inversely with G) would be that performance could be improved without limit by reducing G. However, the thermal time constant of the bolometer is

$$\tau_T = \frac{\Psi}{G}, \tag{6}$$

where Ψ is the heat capacity. Even though the realized speed of the detector can be increased through feedback from the bias circuit (discussed below), high performance with adequate time response demands low heat capacity. The specific heats of bolometer materials—dielectrics, semiconductors, and metals—all decrease with decreasing temperature. Bolometers are operated at very low temperature both to reduce thermal noise and to reduce the response time.

A bolometer operates best when the combination of heat dissipated in the thermometer plus the heat from the infrared background raises its temperature to about 1.5 times that of its heat sink. For linear response, the power dissipated in the thermometer must exceed the infrared power. These constraints set optimum values to G and the operating temperature even when the time response is not an issue. Empirically, it is found that the achievable *NEP* scales approximately as $T^{2-2.5}$ (Rieke 2003). Achieving photon-noise-limited performance requires temperatures of ~ 0.3 K on the ground and ~ 0.1 K when using cold optics in space.

6.2. Transistor Readout Bolometer Arrays

Very high performance bolometers have been built into small arrays for some time, but until recently these devices were based on parallel operation of single pixels (for a specific example, see Rownd et al. 2003, Nguyen et al. 2006). The obstacle to true array-type construction was that the very small signals required use of junction field effect transistor (JFET) amplifiers that needed to operate above about 50 K, far above the operating temperature of 0.3 K or below for the bolometers themselves. It is difficult to implement the simple integration of detector and amplifier that is the heart of array construction with this temperature difference.

With the development of adequately low-noise readouts that can operate near the bolometer temperature, the first true high-performance bolometer arrays for the far-infrared and submillimeter spectral ranges are just becoming available. For example, the Herschel/PACS instrument uses a 2048-pixel array of bolometers (Billot et al. 2006). The architecture of this array is vaguely similar to the direct hybrid arrays for

the near- and mid-infrared. One silicon wafer is patterned with bolometers, each in
the form of a silicon mesh, as shown in **Figure 9**. The delicate construction of the
detector depends on the ability to etch exquisitely complex miniature structures in
silicon. The development of silicon micromachining has enabled substantial advances
in bolometer construction generally and is central to making large-scale arrays. In this
instance, the silicon mechanical structure around the mesh region provides the heat
sink; the mesh is isolated from it with thin and long silicon rods. Numerical model-
ing allows optimization of the structure both for thermal properties and for efficient
photon absorption. The rods and mesh both need to be designed carefully to achieve
appropriate response and time constant characteristics. The mesh is blackened with
a thin layer of titanium nitride with sheet resistance matched to the impedance of
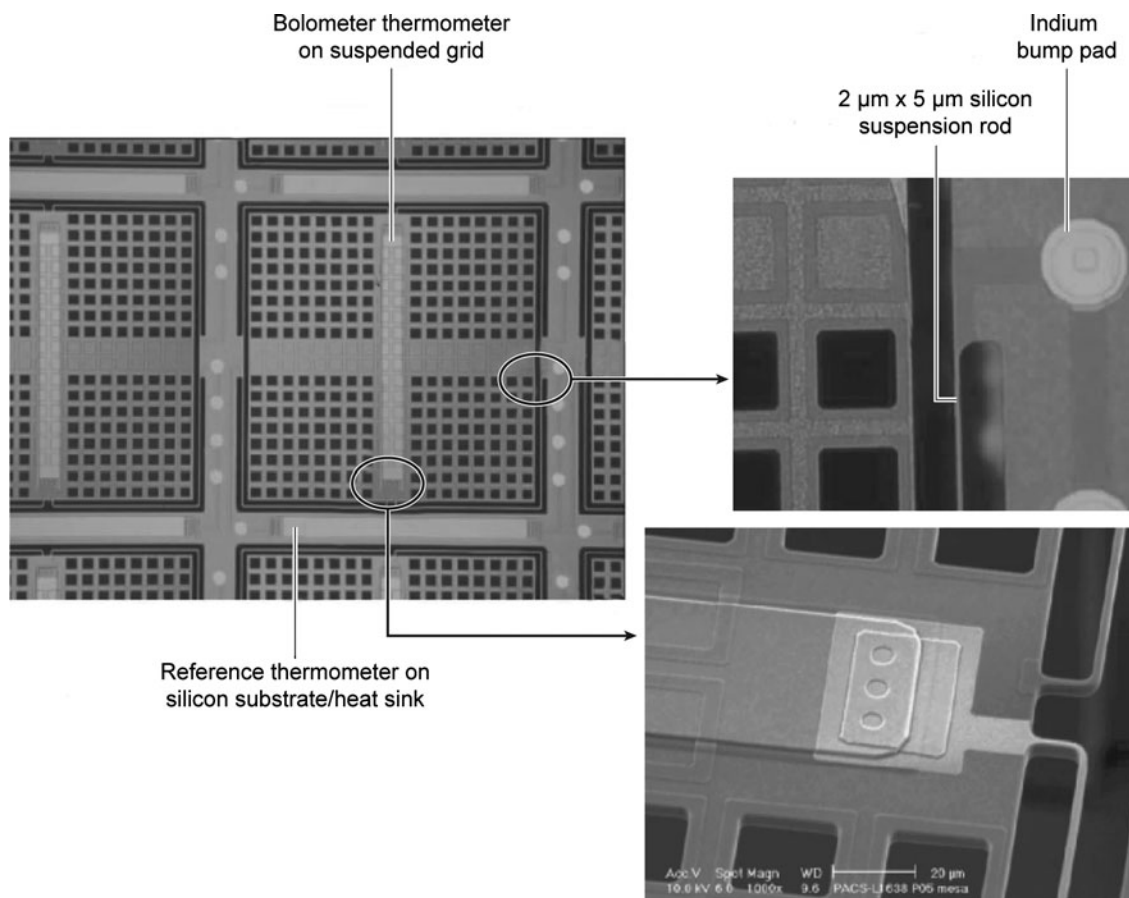free space (377 Ω/square section of film). This matching provides an efficiency of

Figure 9

A single pixel in the Herschel/PACS bolometer array, pixel size about 750 μm, from P. Agnèse
(private communication).

50% over a broad band in absorbing submillimeter- or millimeter-wave photons. Quarter-wave resonant structures can tune the absorption to higher values over limited spectral bands. For each bolometer a silicon-based thermometer doped by ion implantation to have appropriate temperature-sensitive resistance lies at the center of the mesh. Large resistance values are used so the fundamental noise is large enough to utilize MOSFET readout amplifiers. A second silicon wafer is used to fabricate the MOSFET-based readouts, and the two are joined by indium bump bonding. When far-infrared photons impinge on the array, they are absorbed by the grids and raise the temperatures of the thermometers. The resulting resistance changes are sensed by the readouts, amplified, and conveyed to the external electronics. To minimize thermal noise and optimize the material properties, the bolometer array is operated at 0.3 K. Further details are in Billot et al. (2006). Application of similar arrays on ground-based telescopes is discussed by Talvard et al. (2006).

6.3. Transition Edge Sensor Bolometers with Superconducting Readouts

Another approach is taken in transition edge sensor (TES) arrays such as the ones to be used in the submillimeter camera SCUBA (Submillimetre Common-User Bolometer Array)-2. A short discussion of these arrays can be found in Walton et al. (2004), while Irwin & Hilton (2005) and Clarke et al. (2006) provide thorough reviews of TES-based bolometers and readout electronics in general. The name of these devices is derived from their thermometers, which are based on thin superconducting films held within their transition region, where they change from the superconducting to the normal state over a temperature range of a few milliKelvin. In this region, the films have a stable but very steep dependence of resistance on temperature. The temperature of this transition can be set to a convenient point by using a bilayer film consisting of a layer of normal material and a layer of superconductor. The Cooper pairs from the superconductor can diffuse into the normal metal and make it weakly superconducting, a process called the proximity effect. As a result, the transition temperature is lowered relative to that for the pure superconducting film. Adjusting the film thicknesses changes the extent of the effect and adjusts the transition temperature to a convenient value (e.g., one appropriate for a specific type of low-temperature refrigerator).

However, the resistance of a TES is low, so it can deliver significant power only to low-input impedance amplifiers, which rules out JFETs and MOSFETs. Instead, the signals are fed into superconducting quantum interference devices (SQUIDs; see **Figure 10**). SQUIDs are the basis for a growing family of electronic devices that operate by superconductivity (e.g., Clarke & Braginski 2006).

Because of the steep temperature dependence of their resistance, TESs are most stable when biased with a constant voltage. The SCUBA-2 devices, as with all TES bolometers, are operated in this mode (see **Figure 10** for a typical bias circuit). In this state, when their temperature rises due to power from absorbed photons, their resistance rises, the bias current drops, and the electrical power dissipation in them decreases, partially canceling the effects of the absorbed power and limiting the net thermal excursion. This behavior is called electrothermal feedback. The steep

temperature dependence of the resistance of a TES makes the effect very strong. This feedback expedites operating arrays with TESs because minor variations in the transition temperature can be overcome by the tendency of the feedback to force each device to a suitable operating point. Electrothermal feedback can also make the bolometers operate tens or even hundreds of times faster than implied by Equation 6. In fact, if the TES is too fast, the bolometer/SQUID circuit can be unstable and measures must be taken to slow the response.

TES bolometer arrays use SQUIDs for the same readout functions that we have discussed for photodiode and IBC detector arrays. The operation of a simple SQUID time-domain multiplexer is illustrated in **Figure 10** (Benford et al. 2000). The biases across the SQUIDs are controlled by the address lines. Each SQUID can be switched from an operational state to a superconducting one if it is biased to carry about 100 μA. The address lines are set so all the SQUIDs in series are superconducting except one, and then only that one contributes to the output voltage. By a suitable series of bias settings, each SQUID amplifier can be read out in turn.

From our perspective, a major advance in these devices is that the superconducting readouts operate with very low power dissipation and at the ultralow temperature required for the bolometers. Therefore, integration of detectors and readouts is simplified and the architecture can potentially be scaled to very large arrays. Each SCUBA-2 array is made of four subarrays, each with 1280 transition-edge sensors. The design is illustrated in **Figure 11**. The detector elements are separated from their heat sinks by a deep etched trench that is bridged by only a thin silicon nitride membrane. The absorbing surface is blackened by implanting it with phosphorus to match the impedance of free space. The dimensions of the array pixels are adjusted through numerical modeling to form a resonant cavity at the wavelength of operation, to enhance the absorption efficiency. The superconducting electronics that read out the bolometers are fabricated on separate wafers. The two components are assembled into an array using indium bump bonding. Further details are in Griffin (2000), Walton et al. (2004), and Woodcraft et al. (2006).

In addition to the arrays for SCUBA-2, various forms of TES-based bolometers with SQUID readouts are under active development in a number of laboratories (e.g., Benford et al. 2006, May et al. 2006). For example, in the millimeter-wave range the pixel-based array geometry in **Figures 10** and **11** can be replaced with tiny antennae defined by photolithography (Myers et al. 2002). The operating spectral range of antenna-coupled bolometers lies between the infrared and radio and they make use of a mixture of technologies from both spectral regimes. The antenna feeds respond to a single polarization, an advantage if the detectors are planned for a polarimeter. Antennae can be arranged in a single focal plane to measure several polarization angles simultaneously. Microstrip transmission lines can bring the antenna signals outside the sensitive area of the array (a microstrip consists of a miniature circuit trace on an insulator and over a ground plane that can be designed to have some of the characteristics of a waveguide). There, the signals can be sent to a bank of microstrip filters that separate them into multiple spectral bands. Microstrip transmission lines carrying the signals are then terminated with normally conducting metal resistors and TESs sense the temperatures of the resistors as a measure of the power received
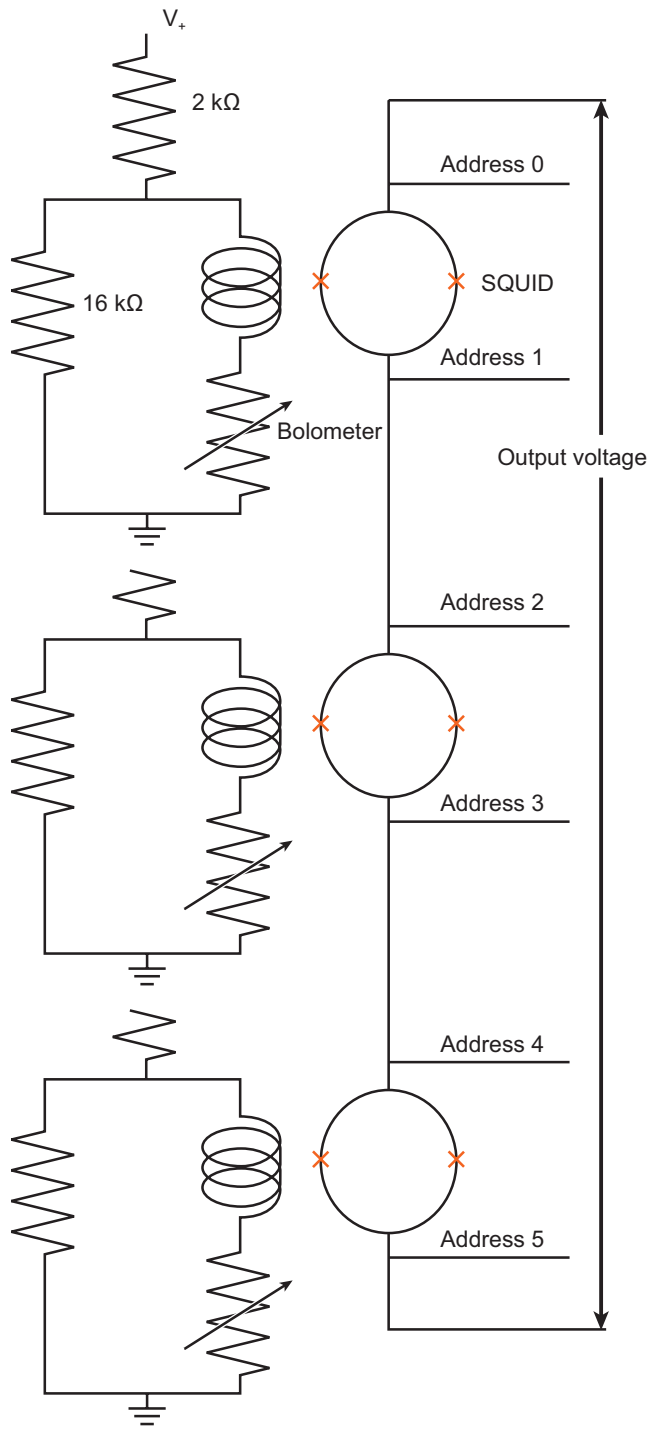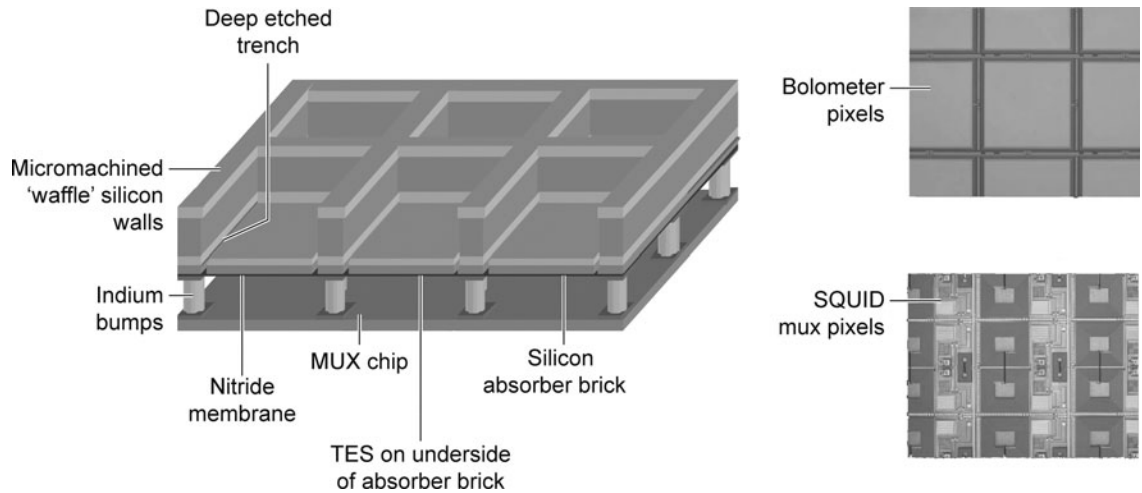
Figure 11

Design features of the SCUBA (Submillimetre Common-User Bolometer Array)-2 bolometer
array, pixel size about 1.1 mm. From Walton et al. (2004 and private communication) and
K. Irwin (private communication).

by the antennae in each band (Goldin et al. 2003). More information about this class
of array can be found in Myers et al. (2005) and Kuo et al. (2006).

There are two basic approaches to multiplexing TES signals. We have described
the time-domain approach, but multiplexing in the frequency domain is also possible.
In this case, each TES is biased with a sinusoidally varying voltage and the signals from
a number of TESs are encoded in amplitude-modulated carrier signals by summing
them. They are read out by a single SQUID and then brought to room-temperature
electronics that recover each of the signals by synchronous detection (e.g., Clarke
et al. 2006 and references therein; also Lanting et al. 2005).

←——————————————————————————————

Figure 10

Bias circuit for TES bolometer and superconducting quantum interference device (SQUID)
readout (*top circuit*). A SQUID consists of an input coil that is inductively coupled to a
superconducting current loop. Two Josephson junctions—junctions of superconductors with
an intervening insulator—interrupt the loop. The Josephson junctions are indicated with "X".
The Cooper pair current across a Josephson junction is a sinusoidal function of the
superconducting phase difference between the two sides of the junction. The superconducting
phase around the current loop is also a function of the magnetic flux through the loop, and
thus of the electrical current through the input coil. In a phenomenon analogous to a two-slit
optical interferometer, interference of the superconducting wavefunction around the loop
results in a voltage response on the output of the SQUID that is a very sensitive function of
the current applied to the input coil. Thus, changes in the bolometer current produce a large
modulation of the SQUID current—i.e., when its output is made linear by using feedback, the
device works as an amplifier. The circuit is repeated three times with appropriate address lines
to operate as a simple SQUID multiplexer. After Rieke (2003), reproduced by permission of
Cambridge University Press.

7. APPLICATION OF INFRARED ARRAYS

7.1. Basic Performance

"Sufficiently advanced technology is indistinguishable from magic" (Clarke 1962). To anyone who has worked in the single-pixel era of infrared astronomy, modern detector arrays carry an aura of magic. As with CCDs in the optical, we are approaching perfection in the near-infrared. That is, we can use very large format arrays of detectors, each of which has nearly 100% quantum efficiency and sufficiently low read noise and dark current that they do not degrade the photon-limited performance in most circumstances. These arrays have sufficient flexibility for most astronomical applications: their wells can accommodate the backgrounds, and they can operate at high speed by reading out subarrays. We are also approaching perfection with regard to arrays between 3 μm and 28 μm for space astronomy missions. The high backgrounds encountered on the ground require changes in array architecture that have not been implemented as broadly, but there are no fundamental technical obstacles to reach near-perfection for these applications also. Beyond 28 μm, array technology is advancing rapidly and, as a result, the possibilities for new astronomical breakthroughs are growing.

As optical astronomers have demonstrated, nearly perfect detector arrays are not the end of the road. Instead, it becomes sensible to invest in ambitious instrumentation knowing that a new array technology is unlikely to make the investment obsolete in a few years. For example, infrared arrays of the current generation have been designed to make two- or even three-sided butting possible. Hence, they can fill large imaging fields of view efficiently or can provide large formats for spectrometers, developments that are being pursued at a number of observatories (e.g., **Figure 5**).

7.2. Data Processing

In general, it becomes more difficult to achieve cosmetically uniform images from an array as the wavelength of operation is increased. As the absorbed photon energy decreases, the relative importance of small nonuniformities in the array detectors grows. Therefore, to appreciate the quality of the data provided with infrared arrays requires a series of careful processing steps. In comparison, a CCD can produce a cosmetically acceptable image without detailed processing, although scientific use of that image requires processing similar to that required with infrared arrays.

In the section below, we describe artifacts that may appear in infrared images even after a standard set of processing steps. Here we assume a reasonably ideal array subject only to nonuniformities that appear as noise but remain reasonably independent of conditions and that can be removed from the data in a standard series of processing steps. There are the following sources of such nonuniformities:

1. **Zero** variations are caused by electronic or temperature drifts that generally depend on the elapsed time and the temperature stability of the array and its readout electronics.

2. **Offset** or **bias** variations have amplitudes that are fixed, independent of signal level, integration time, and elapsed time. They can result from the layout of the ROICs, or from electronic signals that are coupled into the ROIC in a nonvariable way.

3. **Dark currents** depend on integration time but not on signal level. Signals due to photon leaks can sometimes be treated in the data analysis identically to true dark currents, if they also depend only on integration time.

4. Pixel **response variations** depend on both signal level and integration time, normally as the product.

Four types of calibration data are required to remove these sources of noise. Zero variations can be mitigated by appropriate use of reference pixel outputs. A set of very short exposures with the detector in the dark can be averaged to provide a high signal-to-noise image of the offset variations, which can be subtracted from the science image. Dark currents can in principle be removed by taking long exposures in the dark, scaling the signals to the actual exposure time, and subtracting. However, this approach requires ideal behavior of the entire system, so it is safer to take the dark frames with the identical exposure as the science image. In general, the dark frames will embed the offset variations, so removing the latter specifically may be unnecessary if dark frames are used (however, in some cases the dark current is negligible, and then it may be more efficient to ignore it and just concentrate on removing the offsets). Response variations are removed by dividing by an image formed by averaging many exposures on a uniform background.

All of these processing steps work best if the operating conditions of the array are kept as constant as possible throughout acquisition of the science data and calibration images. In this regard, we again emphasize the desirability of maintaining a constant readout cadence on the arrays. Where the infrared backgrounds are substantial compared with the signal (nearly always the case from the ground and often in space), high quality flat field frames must generally be obtained on the sky under identical conditions as those for the data frames. This goal can be accomplished for images by dithering the telescope pointing to obtain many exposures and then using image processing to remove the source signals to produce the flat field frames. Strategies for spectra are more complex and depend on the nature of the spectrometer. In any case, the flat field frames must have zeros, offsets, and dark currents removed as in the data frames, and then they can be divided into the data frames to yield reduced images.

7.3. Artifacts

As good as modern infrared arrays are, they still have a few quirks that are not necessarily removed by the processing steps described above. We try to give a reasonably complete catalog of them to help users identify possible problems in array data. We concentrate this discussion on near- and mid-infrared arrays rather than bolometers because of the far greater observing experience with the former devices (bolometer arrays are only in the first stages of use for astronomical observations).

7.3.1. Electronic. In general, science-grade detector arrays have more than 99% operative pixels. The exceptional pixels are either dead, producing no signals, or hot, with high dark currents. The pixels also have a distribution of read noises, approximately Gaussian in shape but with a tail extending to high values.

Most infrared arrays exhibit a small degree of nonlinearity. Even with proper bias settings, there is a slow reduction in response as signals increase. A typical readout amplifier reduces the detector bias as charge is collected, which directly reduces the response. Reduced bias can also change the detector characteristics to produce similar effects. For example, as the bias on a photodiode is reduced, the width of its depletion region decreases and its capacitance increases, thus reducing the gain of the detector/amplifier combination. Under typical operating conditions, these effects are generally no more than a few percent. If the bias is set too low, the effects on linearity can be much larger, as we have already discussed for IBC detectors.

Bright sources can produce another set of effects, such as duplicate images, pull down of the response along rows of detectors in the array, and bleeding of the signal into additional output channels. These issues can be mitigated by careful design of the array control and bias-supply electronic circuits.

Offsets of the signal are another category of electronic problems. A fixed pattern of offsets can arise from the layout of the ROIC. When arrays are reset, their outputs may end up significantly displaced from zero signal and with large-scale patterns over the array reflecting a gradient in the offset, a class of effect sometimes called a reset anomaly. In the case of the IBC arrays used in the *Spitzer* Infrared Spectrograph and Multiband Imaging Photometer instruments, there is droop, where the output includes a signal proportional to the total signal incident on the array as well as the usual component proportional to the signal on the pixel. There can also be fixed levels of electrical cross talk in the signal. Most of these effects are relatively benign, in that they do not cause noise in proportion to the signals they produce. Therefore, appropriate data processing strategies can remove them almost completely, although they do reduce the dynamic range of the device. However, such effects can vary if the operating cadence of the array is changed, so a constant cadence is helpful in removing them thoroughly from the data.

At the very demanding conditions for long integrations at low background, other forms of subtle electronic misbehavior can manifest themselves. One such phenomenon is popcorn noise, a name derived because similar noise in an audio system can produce a sound reminiscent of popcorn popping (other names are burst noise and telegraph noise). This behavior has been seen in prototype detectors for the near-infrared JWST instruments (Rauscher et al. 2004) and other devices using the same foundry process for the readouts. It takes the form of a rapid—apparently nearly instantaneous—transition between output levels, and can come in a broad variety of jump sizes and patterns (Bacon et al. 2005). It is hypothesized that the underlying cause is charge trapping in areas of the multiplexer that are particularly sensitive, but the details of the noise mechanism are not well understood (Bacon et al. 2005 give additional references regarding the underlying mechanism). Popcorn noise can result in a substantial degradation of the read noise. Its effects can be reduced by handling the array data in a way that allows rejection of the contaminated integrations (e.g., by

obtaining multiple samples along the integration ramp and rejecting segments with anomalous slopes).

Virtually all infrared arrays show latent images. For high-performance HgCdTe detectors, they can be as small as 0.03% of the original signal and decay in a few minutes (e.g., Garnett et al. 2004b). For InSb, Si:As IBCs, and other array types, they can start at up to 1% of the original signal in the first readout after a source has been removed from a pixel, and also decay in a few minutes. In some cases, when exposed to very strong inputs, there is a second class of latent image that requires many hours to disappear unless mitigation steps are taken. Such steps can include raising the temperature of the array sufficiently to thermalize it (termed annealing) or increasing its bias voltage to run a large current through it (bias boost). Latent images of the first kind can be readily identified in a series of images but those of the second kind are more difficult to remove because they tend to propagate into calibration frames used for flat fielding and other operations during the processing of the images.

The free charge carriers produced in an infrared detector upon exposure to ionizing radiation are, of course, phenomenally well detected as spikes of conductivity. Ionizing radiation can produce a number of other effects. In bulk photoconductors, these effects can dominate the calibration uncertainties, but some of them can also exist at very low levels in other detector types. For example, some of the free charge carriers can be captured by the minority impurities in the semiconductor, changing the electric properties of the detector. As a result, at very low background operation, there can be a persistent boost in the detector responsivity but also in its noise. At high backgrounds, the recovery time is relatively short and these effects may not be apparent. However, for good performance at low backgrounds, it may be necessary to remove them periodically by warming the detector (annealing) or other strategies. Fortunately, in detector types where the electrical and optical functions are separated and where the infrared-active regions can therefore have relatively high conductivity—e.g., photodiodes or IBC devices—the susceptibility to these long-term ionizing radiation effects is greatly reduced.

7.3.2. Optical. Because infrared arrays are constructed as thin parallel layers of semiconductor with reflective contacts, they can have strong Fabry-Perot-type fringing at wavelengths where the photon absorption is not complete. This behavior can make accurate calibration of spectra challenging. Also, their pixels behave as isolated detectors and there can be reduced response in the gaps between them. The pixels may also have a low level of nonuniform response over their faces. In optical systems where the image is not well-sampled by the pixels, the resulting fixed grid of low response regions and the dependence of the output on the exact placement of the source on the pixels can make accurate calibration difficult.

Early infrared array readout multiplexers emitted light when their amplifiers dissipated significant levels of power. The array detectors could intercept the resulting mux glow to produce a signal. Although these signals were fixed in strength and distribution over the array for a given readout pattern, and hence disappeared well in flat field image processing, they did contribute noise. Such glow is an inherent property of silicon-based integrated circuits, but it is well controlled in modern arrays. It can

be further reduced by applying power to the array ROIC only when the signals from it are being read out. Glow is at sufficiently low levels in the current generation of arrays that it is not a serious limitation so long as appropriate operational measures are taken to minimize it.

Spurious light signals can also result from luminescence when energetic, ionizing particles excite the materials in an array. Such glow can have time variability that makes it difficult to remove from the images. In addition, light can scatter from exposed bonding wires onto the array from bright sources imaged off it, producing a complex pattern of signal that is virtually impossible to remove completely. It is also sometimes possible for light to enter the detectors at the edge of an array from the side. Both of these issues can be removed by adding a light shield around the edge of the sensitive area. At wavelengths where the array absorption is low, light can also pass through the detectors into the space between them and the readout. There it becomes trapped and undergoes multiple scatterings before it is absorbed. If the final absorption occurs in pixels of the array, the result is a very extended halo of a low level of response around an image. Although these halos are handled well in the usual methods for obtaining photometry of point sources, they make calibration of extended source images complex.

8. SUMMARY

Very high performance near- and mid-infrared arrays have become available for astronomy over the past 20 years through a combination of (*a*) military investments to prove the hybrid array concept and to establish the infrastructure for manufacturing such devices, (*b*) sophisticated semiconductor detector architectures and the means to fabricate them with good control over material properties, and (*c*) advanced integrated circuits on silicon. From 1 μm to 5 μm, the arrays use photodiodes in InSb or HgCdTe, whereas detectors operating from 5 μm to 40 μm are based on extrinsic photoconductivity in silicon. It has been possible to apply some of these advances to the first far-infrared arrays using doped germanium. Similar dramatic progress is occurring in the submillimeter- and millimeter-wave with superconducting bolometers enabled by (*a*) silicon micromachining to manufacture complex thermally isolated detectors, (*b*) superconducting electronics for readout electronics that can operate at the same ultralow temperatures as the detectors, (*c*) numerical modeling to optimize absorption efficiencies, and (*d*) refrigerators that maintain the devices at the 0.1 K to 0.3 K temperatures required to suppress thermal noise and to provide the necessary material properties.

At one level, the advances associated with these infrared detector arrays are summarized in **Figure 12**. The original discovery map of the Galactic Center (**Figure 12*a***) was painstakingly put together from a few scans using a single detector on a 5-m telescope, and required a number of nights. It is now easy to surpass the information content of this map in a few seconds with a 1.3-m telescope (**Figure 12*b***, 2MASS survey; the field shown is approximately the 2MASS array field of view, so the nominal integration time is 8 s). **Figure 12*c*** was taken with a telescope of similar size (the 6.5-m Magellan Telescope) as was used for the discovery map and in less time, but
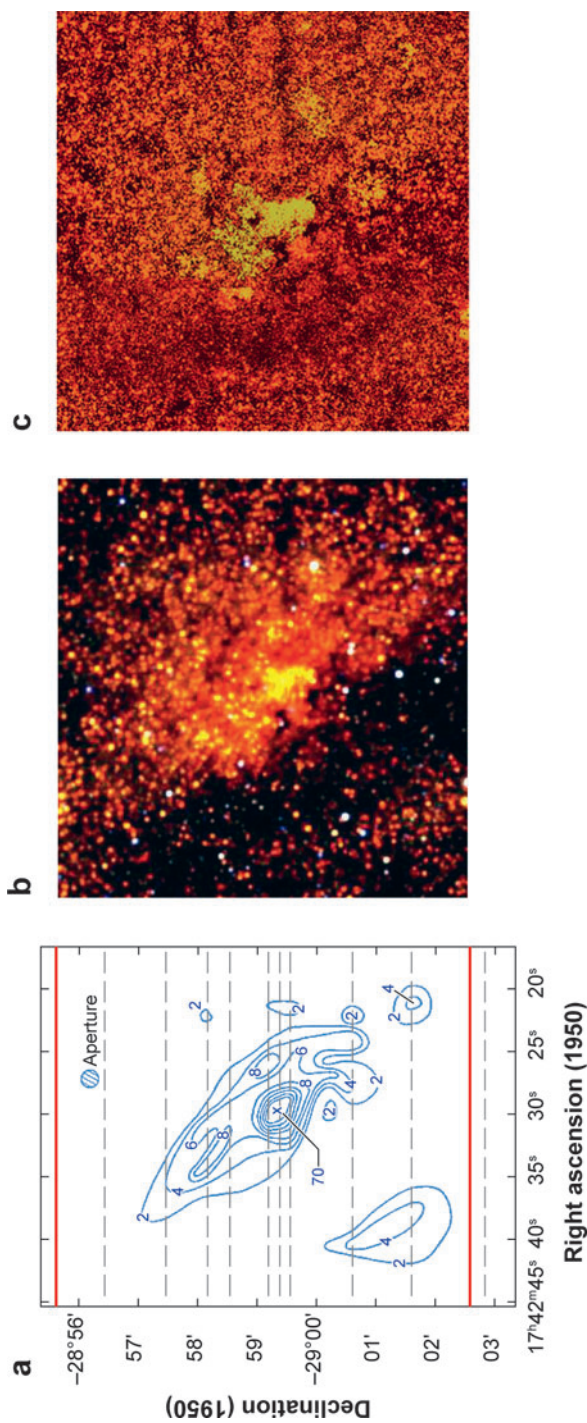
Figure 12

Images of the Galactic Center region at 2 μm. (*a*) The original map by Becklin & Neugebauer (1968), using a single detector with an aperture of 15 arcsec (reproduced by permission of the American Astronomical Society). The dashed lines show the actual scans across the region; the data were taken over three nights with the Palomar 200-in (5-m) telescope. The red lines delineate the roughly 7 × 7 arcmin field shown in panels *b* and *c*. (*b*) The 2MASS image of the same region, using a 256 × 256 detector array on a 1.3-m telescope and with integrations of 8 s per point. The image includes data at 1.25 μm, which identifies foreground stars as blue in the false color image (courtesy of 2MASS/UMass/IPAC-Caltech/NASA/NSF). (*c*) A mosaic obtained by Laycock et al. (private communication) on the 6.5-m Magellan Telescope with a 1 K × 1 K array. The typical resolution is 0.5 arcsec and the total time for the area shown was about an hour. At the depth of this image, Olber's Paradox comes to mind—nearly every line of sight seems to intersect a star.

with a 1024×1024 array. The gain in resolution is about a factor of 30 and in depth by more than three orders of magnitude (for a point source).

Figure 12 illustrates the spectacular advances using detector arrays from the ground, particularly in the near-infrared. At longer wavelengths, arrays have also brought great advances, but their ultimate performance is severely limited by the blinding thermal radiation of the telescope and atmosphere; also, much of the far-infrared is blocked by telluric atmospheric absorption. Modern mid- and far-infrared arrays can reach their full potential only when used on cold telescopes in space. The astronomical capability quantified in Equation 1 can be extended to include the telescope by allowing for its effects on the system sensitivity. The gain in capability for JWST over *Spitzer* is about a factor of 10^4 throughout the mid-infrared. Similarly, an array similar to that being built for SCUBA-2 on a space-borne 10-m telescope [e.g., the Single Aperture Far Infrared Telescope (SAFIR)] would be capable of an advance in capability by a factor of 10^4 to 10^5. These gains are of the same order as the advances from one version of the Galactic Center image to the next in **Figure 12**.

ACKNOWLEDGMENTS

LITERATURE CITED

Amico P, ed. 2004. *Scientific Detectors for Astronomy: The Beginning of a New Era.* Dordrecht: Kluwer Acad.

Ando KJ, Hoffman AW, Love PJ, Toth A, Anderson C, et al. 2003. *Proc. SPIE* 5074:648–57

Arens JF, Lamb GM, Peck MC, Moseley H, Hoffmann WF, et al. 1984. *Astrophys. J.* 279:685–93

Bacon CM, McMurtry CW, Pipher JL, Forrest WJ, Garnett JD, et al. 2004. *Proc. SPIE* 5167:313–19

Bacon CM, McMurtry CW, Pipher JL, Forrest WJ, Garnett JD. 2005. *Proc. SPIE* 5902:118–29

Bahcall JN, Beichman CA, Canizares C, Cronin J, Heeschen D, et al. 1991. *The Decade of Discovery in Astronomy and Astrophysics.* Washington, DC: Natl. Acad. Press

Bandaru J, Beeman JW, Haller EE. 2002. *Proc. SPIE* 4486:193–99

Becklin EE, Neugebauer G. 1968. *Astrophys. J.* 151:145–61

Beletic JE, Beletic JW, Amico P, eds. 2006. *Scientific Detectors for Astronomy 2005*. Netherlands: Springer-Verlag

Benford DJ, Allen CA, Chervenak JA, Grossman EN, Irwin KD, et al. 2000. *ASP Conf. Ser.* 217:134–39

Benford DJ, Steguhn JG, Ames TJ, Allen CA, Chervenak JA, et al. 2006. *Proc. SPIE* 6275:62751C, 12 pp.

Bezawada N, Ives D. 2006. See Beletic et al. 2006, pp. 499–506

Billot N, Agnèse P, Auguères JL, Béguin A, Bouère A, et al. 2006. *Proc. SPIE* 6265:62650D, 12 pp.

Birkmann SM, Eberle K, Grozinger U, Lemke D, Schreiber J, et al. 2004. *Proc. SPIE* 5487:437–47

Clarke AC. 1962. *Profiles of the Future*. New York: Harper & Row

Clarke J, Braginski AI, eds. 2006. *The SQUID Handbook*. Vol. II: Applications. Weinheim, Ger.: Wiley-VCH

Clarke J, Lee AT, Mück M, Richards PL. 2006. See Clarke & Braginski 2006, pp. 1–80

Forrest WJ, Moneti A, Woodward CE, Pipher JL, Hoffman AW. 1985. *PASP* 97:183–98

Fowler AM, Merrill KM, Ball W, Henden A, Vrba F, McCreight C. 2004. See Amico 2004, pp. 51–58

Fujiwara M, Hirao T, Kawada M, Shibai H, Matsuura S, et al. 2003. *Appl. Opt.* 42:2166–73

Garnett JD, Farris MC, Wong SS, Zandian M, Hall DNB, et al. 2004a. *Proc. SPIE* 5499:35–46

Garnett JD, Zandian M, Dewames RE, Carmody M, Pasko JG, et al. 2004b. See Amico 2004, pp. 59–79

Glidden RM, Lizotte SC, Cable JS, Mason LW, Cao C. 1992. *Proc. SPIE* 1684:2–39

Goldin A, Bock JJ, Hunt CL, Lange AE, LeDuc HG, et al. 2003. *Proc. SPIE* 4855:163–71

Gordon KD, Engelbracht CW, Fadda D, Stansberry JA, Wachter S, et al. 2007. *PASP* In press

Gordon KD, Rieke GH, Engelbracht CW, Muzerolle J, Stansberry JA, et al. 2005. *PASP* 117:503–25

Griffin MJ. 2000. *Nucl. Inst. Methods Phys. Res. A* 444:397–403

Hall DNB. 2006. See Beletic et al. 2006, pp. 491–98

Harwit M. 2000. ASP *Conf. Ser.* 195:3–11

Heras AM, Wieprecht E, Feuchtgruber H, Lahuis F, Leech K, et al. 2000. *Exp. Astron.* 10:177–97

Hodapp KW, Hora JL, Hall DNB, Cowie LL, Metzger M, et al. 1996. *New Astron.* 1:177–96

Hogue HH, Guptill ML, Reynolds D, Atkins EW, Stapelbroek MG. 2003. *Proc. SPIE* 4850:880–89

Hoffman AW, Corrales E, Love PJ, Rosbeck JP, Merrill KM, et al. 2004. *Proc. SPIE* 5499:59–67

Hoffman AW, Loose M, Suntharalingam V. 2005. *Exp. Astron.* 19:111–34

Hoffman AW, Love PJ, Rosbeck JP. 2004. *Proc. SPIE* 5167:194–203

Huffman JE, Crouse AG, Halleck BL, Downes TV, Herter TL. 1992. *J. Appl. Phys.* 72:273–75

Irwin KD, Hilton GC. 2005. In *Cryogenic Particle Detection*, ed. C. Enss, *Top. Appl. Phys.* 99:63–149. Berlin: Springer-Verlag

Johnson JB. 1928. *Phys. Rev.* 32:97–109

Kuo CL, Bock JJ, Chattopadthyay G, Goldin A, Golwalla S, et al. 2006. *Proc. SPIE* 6275:62751M, 12 pp.

Lanting TM, Cho HM, Clarke J, Holzapfel WL, Lee AT, et al. 2005. *Appl. Phys. Lett.* 86:2511–13

Love PJ, Beuville EJ, Corrales E, Drab JJ, Hoffman AW, et al. 2006. *Proc. SPIE* 6276:62761Y, 8 pp.

Love PJ, Hoffman AW, Gulbransen DJ, Murray MP, Ando KJ, et al. 2004a. *Proc. SPIE* 5167:134–42

Love PJ, Hoffman AW, Lum NA, Ando KJ, Ritchie WD, et al. 2004b. *Proc. SPIE* 5499:68–77

Love PJ, Hoffman AW, Lum NA, Ando KJ, Rosbeck JP, et al. 2005. *Proc. SPIE* 5902:590209, 58–66

Low FJ, Rieke GH, Gehrz RD. 2007. *Annu. Rev. Astron. Astrophys.* 45:43–75

Lucas C, Maillart P, Mother P, Pantigny P, Vilain M. 1988. *Proc. SPIE* 865:110–16

Lum NA, Ashbrock JF, White R, Kelchner RE, Lum L, et al. 1993. *Proc. SPIE* 1946:100–9

MacDougal MH, Bai Y, Loose M, Baletic JW. 2006. See Beletic et al. 2006, pp. 403–10

Mainzer AK, Eisenhardt P, Wright EL, Lin FC, Irace W, et al. 2005a. *Proc. SPIE* 5899:262–73

Mainzer AK, Hong J, Stapelbroek MG, Hogue H, Molyneux D, et al. 2005b. *Proc. SPIE* 5881:253–60

May T, Anders S, Fritzsch L, Boucher R, Zakosarenko V, et al. 2006. *Proc. SPIE* 6275:62751Q, 8 pp.

McMurtry CW, Forrest WJ, Moore A, Pipher JL. 2003. *Proc. SPIE* 4850:847–57

Mottier PL, Agnèse P, Lagage PO. 1991. *Proc. SPIE* 1512:60–67

Myers MJ, Holzapfel W, Lee AT, O'Brient R, Richards PL, et al. 2005. *Appl. Phys. Lett.* 86:4103–5

Myers MJ, Lee AT, Richards PL, Schwan D, Skidmore JT, et al. 2002. In *Low Temperature Detectors*, *AIP Conf. Proc.*, ed. FS Porter, D McCammon, M Galeazzi, CK Stahle, pp. 247–50. New York: Am. Inst. Phys. Press

Nguyen HT, Ringold P, Ade PAR, Battle J, Beeman JW, et al. 2006. *Proc. SPIE* 6275:627518, 5 pp.

Poglitsch A, Katterloher RO, Hoenle R, Beeman JW, Haller EE, et al. 2003. *Proc SPIE* 4855:115–28

Rauscher BJ, Figer DF, Regan MW, Boeker T, Garnett JD, et al. 2004. *Proc. SPIE* 5487:710–26

Reichertz LA, Beeman JW, Cardozo BL, Jakob G, Katterloher R, et al. 2006. *Proc. SPIE* 6275:62751S, 8 pp.

Richards PL, McCreight CR. 2005. *Phys. Today* 58(2):41–47

Rieke GH. 2003. *Detection of Light from the Ultraviolet to the Submillimeter*. Cambridge, UK: Cambridge Univ. Press

Rieke GH. 2006. *The Last of the Great Observatories*. Tucson: Univ. Ariz. Press

Rieke MJ, Winters GJ, Cadien J, Rasche RW. 1993. *Proc. SPIE* 1946:214–24

Riopel M, Teeple D, Ward J. 2006. See Beletic et al. 2006, pp. 159–64

Robberto M, Baggett SM, Hilbert B, MacKenty JW, Kimble RA, et al. 2004. *Proc. SPIE* 5499:15–22

Rownd B, Bock JJ, Chattopadhyay G, Glenn J, Griffin MJ. 2003. *Proc. SPIE* 4855:510–19

Shirahata M, Matsuure S, Makiuti S, Patrashin MA, Kaneda H, et al. 2004. *Proc. SPIE* 5487:369–80

Simons DA, Amico P, Baade D, Barden S, Randall C, et al. 2006. See Beletic et al. 2006, pp. 13–44

Talvard M, André P, Rodriquez L, Minier V, Benoit A, et al. 2006. *Proc. SPIE* 6275:627503, 12 pp.

van Cleve JE, Herter TL, Butturini R, Gull GE, Houck JR, et al. 1995. *Proc. SPIE* 2553:502–13

Walton AJ, Parkes W, Terry JG, Dunare C, Stevenson JTM, et al. 2004. *IEE Proc. Sci. Meas. Technol.* 151:110–20

Watson DM, Huffman JE. 1988. *App. Phys. Let.* 52:1602–4

Woodcraft AL, Hollister MI, Bintley D, Ellis MA, Gao X, et al. 2006. *Proc. SPIE* 6275:62751F, 11 pp.

Wu IC, Beeman JW, Luke PN, Hansen WL, Haller EE. 1991. *Appl. Phys. Lett.* 58:1431–33

Wynn-Williams CG, Becklin EE, eds. 1987. *Infrared Astronomy with Arrays*. Honolulu: Univ. Hawaii. 551 pp.

Young ET, Davis JT, Thompson CL, Rieke GH, Rivlis G, et al. 1998. *Proc. SPIE* 3354:57–65

Young ET, Rieke GH, Dang H, Barg I, Thompson CL. 1995. *Proc. SPIE* 2475:435–40

Contents