Sky on a Chip: The Fabulous CCD

James Janesick, Caltech, and Morley Blouke, Tektronix, Inc.

TELESCOPES today are no bigger than those in use two generations ago. Yet, curiously, astronomers can now study objects that are hundreds or even thousands of times fainter than those accessible in the 1930's.

The explanation is simple: the last few decades have brought an unprecedented improvement in the sensitivity of light detectors. The changes in our own era are as significant as those caused by the development of the telescope in the 17th century or the photographic plate in the 19th century.

Of all the new astronomical tools, one in particular stands out because it is an almost perfect radiation detector. Known as a charge-coupled device, or CCD, it has the ability to register almost every photon (particle of light) that strikes it over wavelengths ranging from the near infrared to X-rays. In contrast, the most sensitive photographic emulsion used by astronomers records two or three percent of the light striking it, and over a much narrower range of wavelengths.

The CCD revolution came about indirectly. It began in the late 1960's when two researchers at Bell Telephone Laboratories, Willard S. Boyle and George E. Smith, sought to invent a new type of computer memory circuit. In particular, they wanted to construct an electronic analog to magnetic bubbles.

While Boyle and Smith conceived the CCD as a memory element, it rapidly became apparent that the tiny chip of semiconducting silicon they first demonstrated in 1970 had many other applications, including signal processing and imaging (the latter because silicon responds to visible light).

In recent years the CCD's initial promise as a memory element has vanished. It has, however, established itself as the premier imaging device for scientific applications, especially astronomical ones. Indeed, CCD's are at the forefront of an ongoing revolution in all types of electronic photography. They are already common in lightweight video systems and auto-focus cameras.

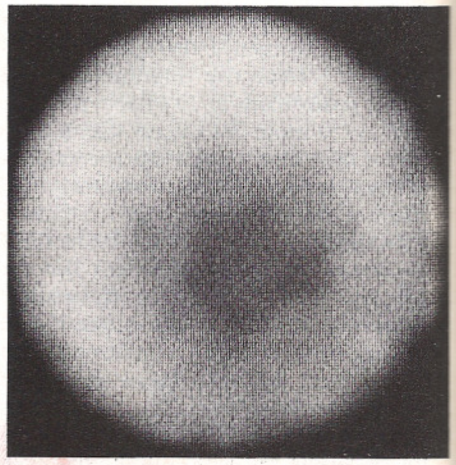

This picture of Uranus is thought to be the first astronomical image made with a charge-coupled device, or CCD. It was obtained in 1975 by scientists from the Jet Propulsion Laboratory and the University of Arizona, using the 61-inch telescope in the Santa Catalina Mountains near Tucson. Recorded at a wavelength of 8900 angstroms in the near infrared, it shows a region of enhanced methane absorption (dark area) near Uranus' south pole.

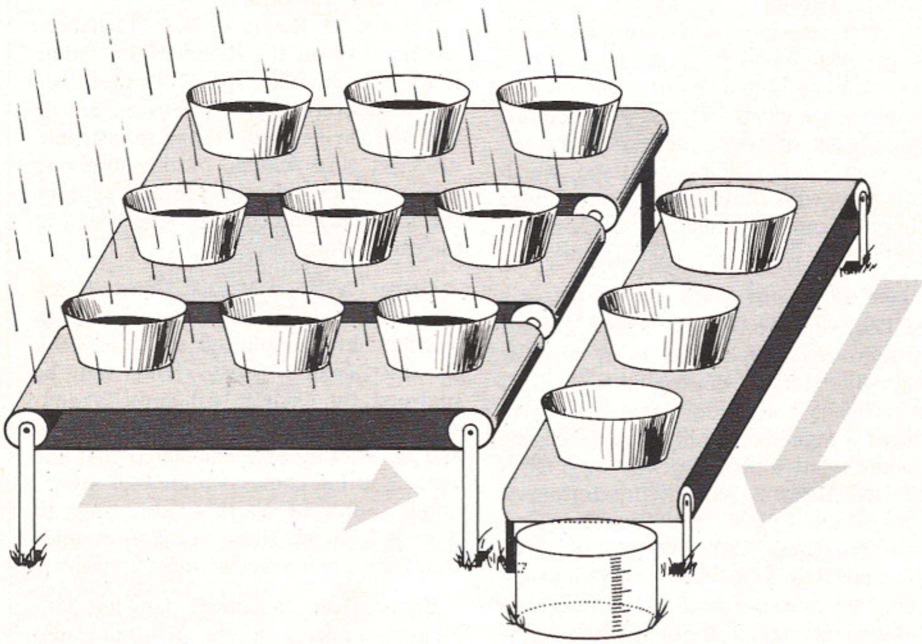

Determining the brightness distribution in a celestial object with a charge-coupled device can be likened to measuring the rainfall at different points in a field with an array of buckets. Once the rain has ceased, the buckets in each row are moved horizontally across the field on conveyor belts. As each one reaches the end of the conveyor, it is emptied into another bucket on a belt that carries it to the metering station where its contents are measured. Artwork by Steven Simpson.

PIONEERING USE

Astronomers were among the first to recognize the extraordinary imaging capabilities of the CCD. In 1972 researchers at the Jet Propulsion Laboratory (JPL) established a program to develop CCD's for space astronomy. Three years later the JPL team, working with scientists from the University of Arizona, took what is probably the first astronomical image obtained with a CCD (see the photograph above).

Today every major telescope in the world has a CCD camera as part of its observing arsenal. Such a device can be used either for direct imaging or as the detector in a spectrometer or other instrument. In addition, the Vega, Giotto, and Suisei probes to Halley's comet, and many future space projects — including the Hubble Space Telescope, Galileo, and the Advanced X-ray Astrophysics Facility — will have CCD's at the heart of their camera systems.

The operation of a CCD is quite simple in principle. An elegant analogy by Jerome Kristian of the Carnegie Insti-

Sky on a Chip: The Fabulous CCD

James Janesick, Caltech, and Morley Blouke, Tektronix, Inc.

TELESCOPES today are no bigger than those in use two generations ago. Yet, curiously, astronomers can now study objects that are hundreds or even thousands of times fainter than those accessible in the 1930's.

The explanation is simple: the last few decades have brought an unprecedented improvement in the sensitivity of light detectors. The changes in our own era are as significant as those caused by the development of the telescope in the 17th century or the photographic plate in the 19th century.

Of all the new astronomical tools, one in particular stands out because it is an almost perfect radiation detector. Known as a charge-coupled device, or CCD, it has the ability to register almost every photon (particle of light) that strikes it over wavelengths ranging from the near infrared to X-rays. In contrast, the most sensitive photographic emulsion used by astronomers records two or three percent of the light striking it, and over a much narrower range of wavelengths.

The CCD revolution came about indirectly. It began in the late 1960's when two researchers at Bell Telephone Laboratories, Willard S. Boyle and George E. Smith, sought to invent a new type of computer memory circuit. In particular, they wanted to construct an electronic analog to magnetic bubbles.

While Boyle and Smith conceived the CCD as a memory element, it rapidly became apparent that the tiny chip of semiconducting silicon they first demonstrated in 1970 had many other applications, including signal processing and imaging (the latter because silicon responds to visible light).

In recent years the CCD's initial promise as a memory element has vanished. It has, however, established itself as the premier imaging device for scientific applications, especially astronomical ones. Indeed, CCD's are at the forefront of an ongoing revolution in all types of electronic photography. They are already common in lightweight video systems and auto-focus cameras.

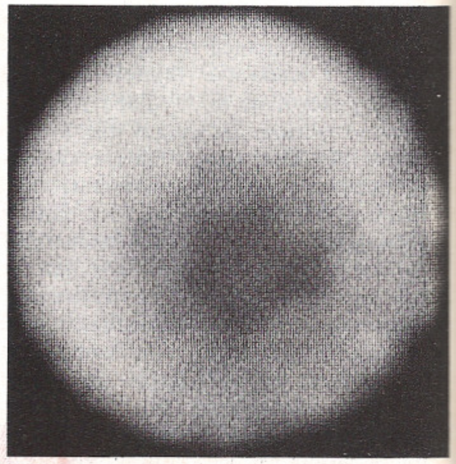

This picture of Uranus is thought to be the first astronomical image made with a charge-coupled device, or CCD. It was obtained in 1975 by scientists from the Jet Propulsion Laboratory and the University of Arizona, using the 61-inch telescope in the Santa Catalina Mountains near Tucson. Recorded at a wavelength of 8900 angstroms in the near infrared, it shows a region of enhanced methane absorption (dark area) near Uranus' south pole.

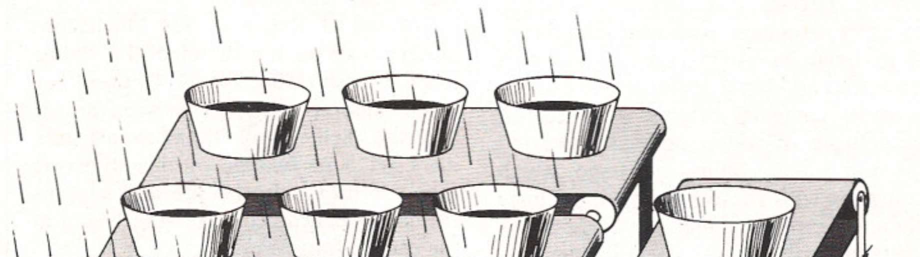

Determining the brightness distribution in a celestial object with a charge-coupled device can be likened to measuring the rainfall at different points in a field with an array of buckets. Once the rain has ceased, the buckets in each row are moved horizontally across the field on conveyor belts. As each one reaches the end of the conveyor, it is emptied into another bucket on a belt that carries it to the metering station where its contents are measured. Artwork by Steven Simpson.

PIONEERING USE

Astronomers were among the first to recognize the extraordinary imaging capabilities of the CCD. In 1972 researchers at the Jet Propulsion Laboratory (JPL) established a program to develop CCD's for space astronomy. Three years later the JPL team, working with scientists from the University of Arizona, took what is probably the first astronomical image obtained with a CCD (see the photograph above).

Today every major telescope in the world has a CCD camera as part of its observing arsenal. Such a device can be used either for direct imaging or as the detector in a spectrometer or other instrument. In addition, the Vega, Giotto, and Suisei probes to Halley's comet, and many future space projects — including the Hubble Space Telescope, Galileo, and the Advanced X-ray Astrophysics Facility — will have CCD's at the heart of their camera systems.

The operation of a CCD is quite simple in principle. An elegant analogy by Jerome Kristian of the Carnegie Institution of Washington is often used to describe how it works (see the bottom diagram on the facing page). Imagine an array of buckets covering a field. After a rainstorm, the buckets are sent by conveyor belts to a metering station where the amount of water in each is measured. Then a computer could display a picture of how much rain fell on each part of the field. In a CCD system, the "raindrops" are photons. (See the box on page 240 for a more technical description.)

ADVANTAGES OF THE CCD

To understand why a CCD is so useful and powerful, we need to understand what a good astronomical detector must do. Its most important requirements are:

- High resolution — it must be able to see fine detail.
- High quantum efficiency — as much of the radiation falling on the sensor as possible must be actually detected.
- Broad spectral response — the detector must be sensitive to radiation over a wide range of wavelengths.
- Low noise — any spurious signals it generates must be far weaker than the celestial signals it is trying to measure.
- Large dynamic range — the difference between the brightest and faintest objects it can view simultaneously must be as great as possible.
- Photometric accuracy — its output should be in such a form that the brightnesses of the objects viewed can easily be measured with high precision in absolute units, such as magnitudes.
- Linearity — the output should be uniform; if one object is twice as bright as another, its image should be twice as strong.

Unlike many other detectors currently in use, the CCD excels in every one of these areas.

A CCD's resolution is determined primarily by the number of picture elements, or pixels, making up the imaging area. More pixels mean better resolution, as evident in the photograph at right.

The first large image-forming CCD's were introduced by Fairchild Semiconductor in 1973. They contained 10,000 pixels arranged in 100 rows and 100 columns. A short time later, RCA Corp. made versions achieving the standard television resolution of 512 rows by 320 columns (163,840 pixels). Detectors like this are typical of the CCD's now in use at observatories around the world.

The CCD's made especially for Galileo and the Hubble Space Telescope by Texas Instruments are now becoming available to astronomers in limited numbers. These have 640,000 pixels arranged in an 800-by-800 array. Texas Instruments has also made chips with over one million picture elements for the cameras of the proposed High-Resolution Solar Observatory.

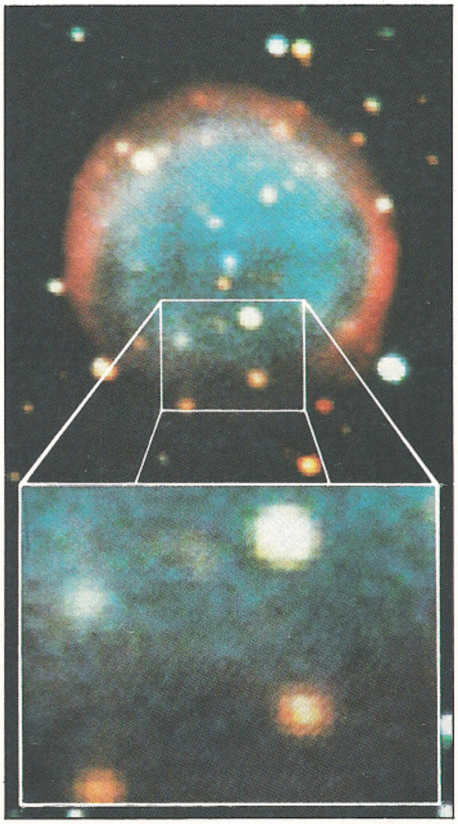

This image of the planetary nebula NGC 6781 in Aquila was made by combining three CCD images, each made with a different color filter. Like all CCD images, it is built up from an array of tiny squares (*inset*), each representing the data collected by one pixel. The more pixels that make up the image, the finer the detail visible. This image was obtained with a CCD having 512 rows and 320 columns of picture elements. Courtesy Rudolph Schild.

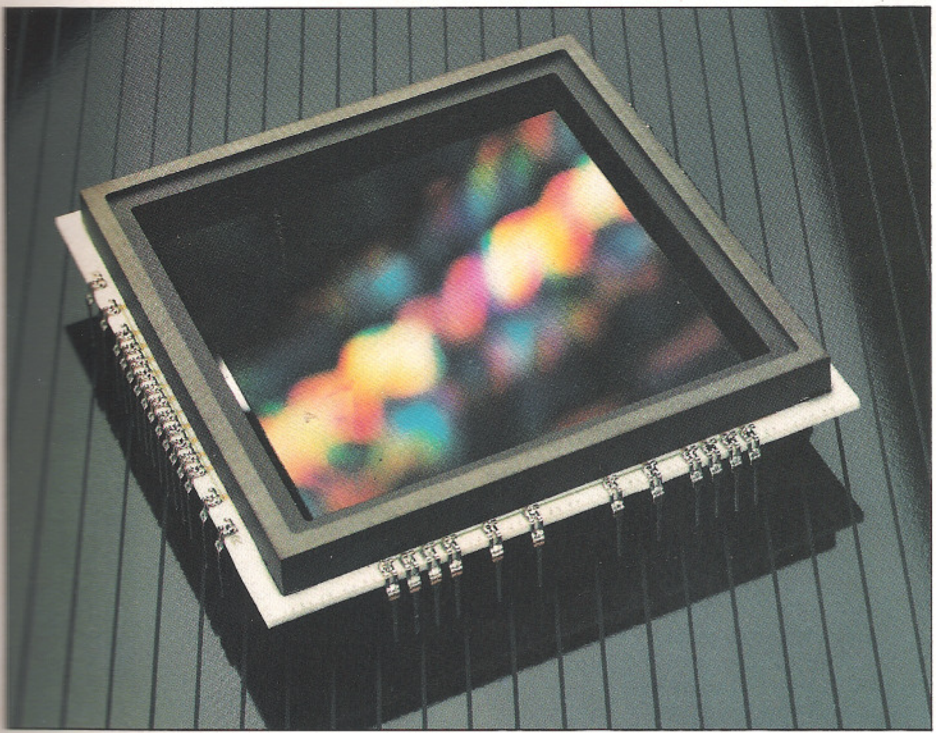

The largest charge-coupled device currently available is made by Tektronix, Inc. Comprising an array of 2,048 rows and 2,048 columns of light-sensitive picture elements, the chip measures over 2.5 inches square. All illustrations supplied by the authors unless stated otherwise.

Today the largest CCD is made by Tektronix, Inc. (see left and next page) and incorporates over 4,000,000 pixels (2,048 by 2,048). Its active image area measures over 2.5 inches across, making it the world's largest integrated circuit.

Impressive as these numbers sound, they are small compared to the size and resolution of photographic emulsions. Even the coarsest-grain 35-millimeter film might have a resolution equivalent to 25,000,000 pixels. Fine-grained astronomical plates can do a lot better.

However, CCD images have other advantages that render them far superior to photographs. For example, the bottom left diagram on page 241 shows how the quantum efficiency of a typical CCD exceeds that of several other detectors used by astronomers.

The eye, the first astronomical "detector," has a quantum efficiency of about one percent. In other words, of every 100 photons entering the pupil, only one is detected. In contrast, for every 100 photons hitting a CCD, 50 to 70 or even more are registered. In addition, the range of wavelengths to which the eye can respond is much smaller than the CCD's. The same limitation is true of photocathodes, film, and vidicon-type television detectors.

It has recently been shown that the next generation of CCD's will respond usefully from 1 angstrom in the "soft" X-ray region of the spectrum to greater than 10,000 angstroms in the near-infrared (see the box on page 242). This huge wavelength span is far greater than that of any other detector known; it is one of the properties that makes CCD's useful for such diverse missions as the Advanced X-ray Astrophysics Facility and Hubble Space Telescope.

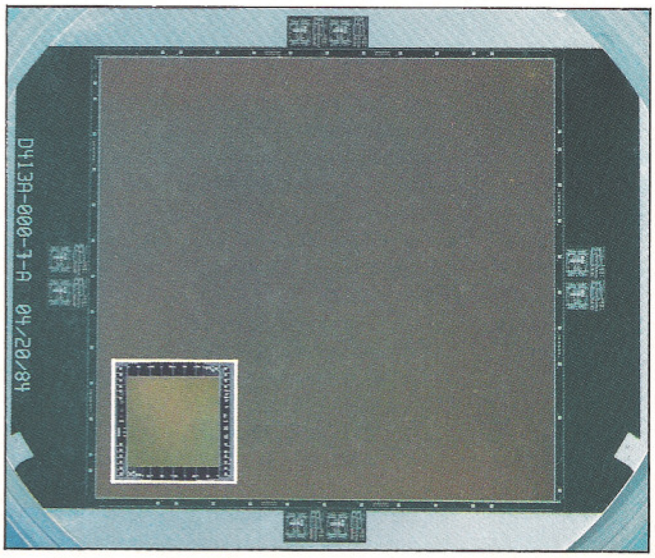

A Tektronix 2,048-by-2,048-pixel charge-coupled device shown life size. This monster chip, the world's largest, is currently under development for astronomical applications. Also shown here full size is the 800-by-800-pixel CCD (*inset*) made especially by Texas Instruments for the Hubble Space Telescope and the Galileo Jupiter probe. While perfect examples are worth dozens of times their weight in gold, cosmetically defective models have already been made available to the astronomical community.

How a Charge-Coupled Device Works

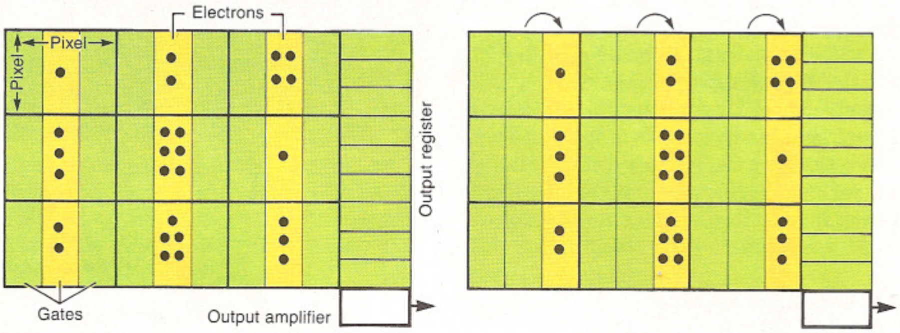

The operation of a CCD is illustrated schematically by this simplified chip consisting of nine pixels, an output register, and an amplifier. Every pixel is subdivided into three regions, or gates; each is an electrode whose voltage can be varied. *Left:* During an exposure the central gate of each pixel is "on" (yellow areas) and its neighbors are "off" (green areas). This creates "electron buckets" under the middle electrodes and barriers between adjacent pixels. *Right:* At the end of the exposure the gate voltages are changed and electrons shifted one gate to the right as new potential wells are created and old ones are destroyed.

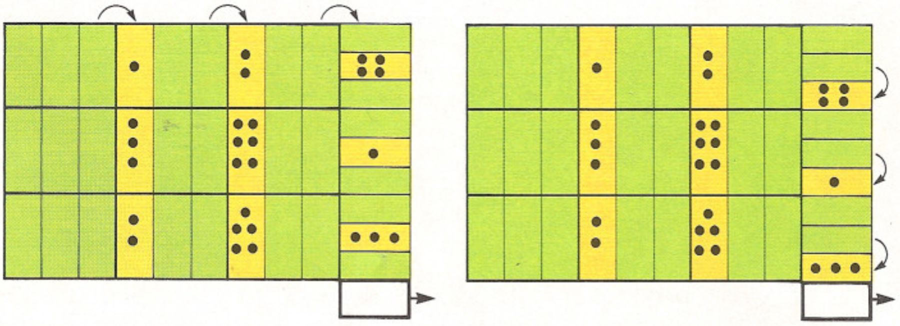

***Left:* As the voltages are cycled again, electrons flow from the right-most gate of one pixel to the left-most gate of its neighbor. The electrons in the right-hand pixels are transferred to the output register. *Right:* Before the pixel array can be shifted again, charge must be transferred, one pixel at a time, through the output register and amplifier. When this register has been completely emptied, another cycle of pixel-array transfers is executed. These steps continue in a systematic way until every charge packet has moved horizontally along its row, vertically down the output register, and into the amplifier, where it is measured. In a real astronomical CCD, this process can take as long as 10 seconds.**

A charge-coupled device must perform four tasks to generate an image. These operations are:

- Photoelectron generation (rain).
- Electron collection (buckets).
- Charge transfer (conveyor belts).
- Readout (metering station).

The first operation relies on a physical process known as the photoelectric effect — when light hits certain materials, electrons are generated. When a pattern of light from, say, a distant galaxy falls on a CCD, the photons that make up the incoming image are absorbed within the chip's silicon lattice, where they give rise to free electrons.

In the second step these so-called photoelectrons are collected in the nearest discrete collecting sites, or pixels (picture elements). These collection sites are defined by an array of electrodes, called gates, formed on the surface of the CCD. In our rainfall analogy, the buckets represent the pixels.

The third stage, charge transfer, is accomplished by changing the voltages on each gate in a systematic way so that the electrons move horizontally from one pixel to the next, conveyor-belt fashion.

At the end of each row is the so-called output register, which is in effect a conveyor running perpendicular to all the others. It transports the charge packets one by one to the output amplifier where, in the final operating steps, the electrons are counted and converted to a form suitable for storage in a computer, display on a television monitor, and analysis by astronomers.

But a high quantum efficiency and broad spectral response are useless if the detector introduces false signals that swamp those from a faint astronomical source. The electrons produced in each pixel of a CCD are counted by a single transistor built into the chip's output circuitry. The accuracy of the counting is determined by the noise introduced by the transistor. Today's CCD's have amazingly low noise and can reliably detect charge packets containing as few as 10 electrons (generated by the capture of 15 to 20 visible-light photons).

Small as these numbers seem, astronomy demands even greater sensitivity. Counting errors are typically comparable to the number of photons collected every second when a 24th-magnitude object is viewed with a 60-inch telescope. Substantial effort is now being expended to enable CCD's to count charge packets containing fewer than four electrons.

One of the most significant advantages a CCD has over other detectors is its large dynamic range. For CCD's this important property is defined as the ratio of the largest charge packet that can be measured to the smallest.

The larger the size of a pixel, the more electrons it can hold before overflowing. For example, the 15-micron-square pixels on Texas Instruments' 800-by-800 CCD can each hold about 75,000 electrons. Picture elements twice as wide have a capacity of nearly 1,000,000 electrons. These numbers reveal that current CCD's can have dynamic ranges up to 100,000 (1,000,000 electrons in a full pixel divided by 10 electrons of noise).

Large dynamic range is particularly important for an astronomical detector because of the large brightness differences commonly found between, say, the faint outer parts of galaxies and their brilliant nuclei. The greater the dynamic range, the greater the CCD's ability to image *both* bright and dim objects in the same field. In stark contrast, the best photographic plates can reproduce a brightness range of only 100 to 1.

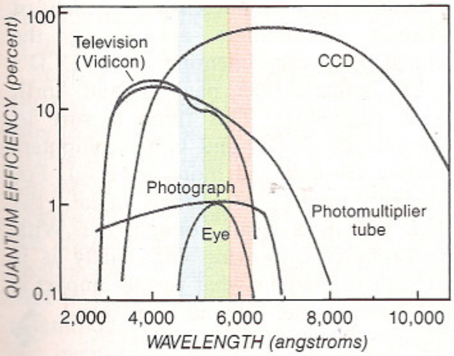

The superiority of the charge-coupled device (CCD) over almost all other image-forming detectors is illustrated here. In the red and near-infrared portions of the spectrum, CCD's are more efficient at detecting radiation than any other type of device. In the ultraviolet, however, the performance of older CCD's illustrated here is much worse than television-type detectors and even photographic film.

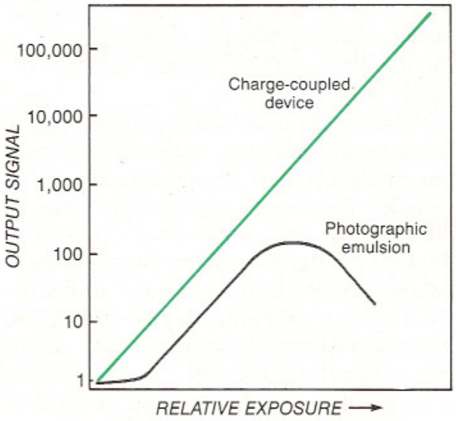

CCD's are exceptionally linear; doubling the exposure yields a signal exactly twice as strong. This holds over an enormous exposure range, enabling CCD's to detect simultaneously very bright and very faint objects in a single image. This behavior contrasts strikingly with that of photographic emulsions, which are only linear over a narrow range of exposure. Emulsions must also overcome a threshold before they start to form an image at all. Worse, beyond a certain point emulsions saturate; their ability to detect light actually decreases if the exposure is too long.

MORE THAN PRETTY PICTURES

Even the most sensitive detector is essentially worthless to astronomers if it can't be used to determine reliably the brightness of what is being observed in some sort of absolute units such as magnitudes. This requires that the detector be stable over long periods of time so that it always produces the same output signal for a given light input.

Because they are solid-state devices, CCD's are inherently stable. Once a chip has been calibrated by observing standard stars, photometric measurements made during the course of a night can be held to an absolute accuracy of 0.5 percent (equivalent to an uncertainty of about 0.005 magnitude).

The CCD also exhibits remarkable linearity. In other words, the number of electrons collected in a pixel is precisely proportional to the number of incident photons (see the diagram above). This is in sharp contrast to photographic film and vidicon-type television detectors. The former has a very complicated response to

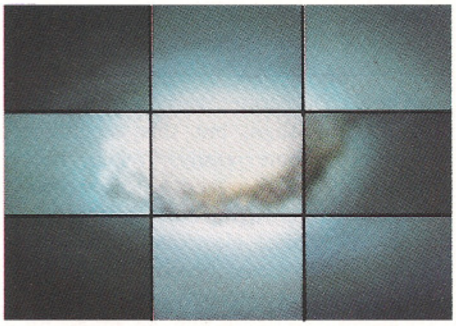

This simulated "raw" image from a charge-coupled device does not give a realistic representation of the well-known Black-eye galaxy, M64 in Coma Berenices, because individual pixels have varying sensitivities to light. This effect is shown here schematically by dividing the image into nine regions of grossly different sensitivities. Images such as this are obviously not very useful to astronomers or anybody else.

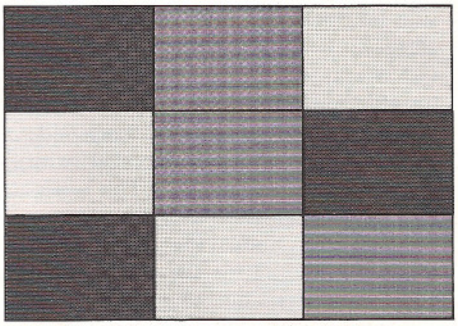

To overcome this nonuniformity problem, observers use a technique called flat-fielding. During an observing run an exposure is made of a uniform source such as the dawn sky or inside of the telescope dome to map the variations among the CCD's pixels. This image represents the flat-field frame corresponding to the simulated raw image of the Black-eye galaxy shown above.

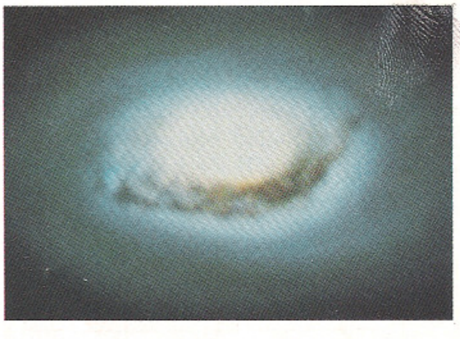

Once the response of each pixel is known, observers can scale the data by an amount determined from the flat-field exposure and thereby correct for each pixel's individual response. The result is the same as that which would be registered by a CCD with pixels of identical sensitivity; it can now be analyzed by astronomers. Image courtesy Rudolph Schild.

The All-Seeing CCD

The CCD operates over an enormous wavelength range, but this versatility does not come easily. When illuminated from the front, it is rather insensitive to ultraviolet and X-ray radiation. The pixel electrodes ("gates") are opaque to these photons.

A 20-percent improvement in short-wavelength response can be made by coating the front surface with a phosphor that converts incident ultraviolet radiation to longer wavelengths where the device is more efficient. A better method involves thinning the CCD and illuminating it from the back. Photons can then enter the active region unimpeded by the gates.

Further treatment of the thinned CCD is required because photoelectrons can get trapped by a small "back side potential well" at the chip's rear surface. Researchers have developed two ways to "push" these electrons toward the front-surface electrodes.

In one approach, called back side charging, the CCD is flooded with intense ultraviolet radiation. This generates excess photoelectrons that destroy the well and charge the rear surface negatively as long as the chip is kept cold. The CCD's on the Hubble Space Telescope will employ this technique.

The "flash gate" is a more permanent solution. An electrical potential generated between the CCD's silicon and a thin layer of platinum or gold deposited on its back side eliminates the unwanted well. Second-generation instruments for the Space Telescope will probably use this technology.

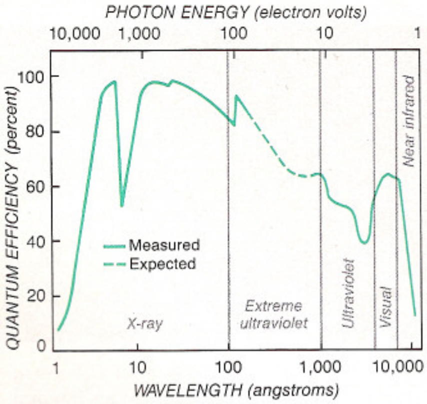

Recent technological advances enable charge-coupled devices to respond to radiation across the spectrum from near-infrared to soft-X-ray wavelengths. No other detector so efficiently responds over such a broad spectral range.

light and at best can achieve a photometric accuracy of only 5 percent over a small portion of its already narrow dynamic range.

Vidicons are also highly nonlinear. The relationship between their input and output signals must be approximated by a complex mathematical expression that is different for each point on the detector. This makes the extraction of photometric information difficult and time consuming. Even when the required procedures are followed, vidicon photometry is only accurate to about two percent.

In contrast, CCD's are linear to 0.1 percent over most of their dynamic range, far superior to any other device. This makes it relatively simple to remove the effects of pixel-to-pixel variations in sensitivity. In practice this is achieved by means of a technique known as "flat-fielding" (illustrated on page 241).

Vidicons and related detectors have other serious flaws that limit their astronomical usefulness. They all form images on a light-sensitive medium that is scanned by a beam of electrons just as in a television camera. The first problem encountered is to know exactly where the beam is at any instant, for its position can't be measured directly. Thus the relative positions of points in the image may not correspond exactly to those in the object viewed.

Compounding these positional uncertainties are the distortions introduced by the magnetic fields used to focus the readout beam. Taken together, these effects compromise the fidelity with which vidicons reproduce a given scene. These problems do not apply to a CCD. The position of each pixel is fixed rigidly when the chip is manufactured.

BEYOND THE VISIBLE

In the past decade more and more astronomers have turned their attention to those parts of the electromagnetic spectrum invisible to our eyes. Increasingly, observations are being made at ultraviolet (1200-3500 angstroms), extreme ultraviolet (120-1200 angstroms), and soft-X-ray (1.2-120 angstroms) wavelengths.

Photons in these short-wavelength portions of the spectrum are more energetic than their visible counterparts. When absorbed in the CCD's silicon lattice these high-energy photons generate many electrons, the exact number depending on the wavelength (energy) of the radiation. For example, a 2.1-angstrom X-ray photon will spawn, on average, 1,620 electrons.

For wavelengths less than 100 angstroms, therefore, a *single* photon can be detected and its energy (or wavelength) determined by directly measuring the amount of charge generated. This raises an interesting possibility: in a single image (actually, many very short exposures superimposed), a CCD can provide an X-ray astronomer with simultaneous high-resolution spatial *and* spectroscopic data. No other detector can do this.

CCD's FOR EVERYBODY?

Even though CCD's are almost perfect detectors, they are not suitable for all astronomers. They are very expensive; a 2,048-by-2,048-pixel chip presently lists for over $80,000, or 19 cents per pixel! Devices with standard TV resolution range from $2,000 to $10,000 (1.2 to 6 cents per pixel), depending on their quality.

CCD's generally remain beyond the reach of most amateurs, though the increasing popularity and decreasing price of lightweight video cameras is changing this somewhat. Several Japanese manufacturers sell CCD's for about $200 each. (Cosmetic rejects go for about half this price.) But even chips this cheap are far more expensive than photographic film.

Perhaps the most important disadvantage of CCD's is the amount of ancillary electronics needed to make them work. Prime among these accessories is a powerful computer to analyze the tremendous amount of data these tiny silicon chips create. A single image from a 2,048-by-2,048 CCD will contain 10 megabytes of information. The amount of data gathered during a single observing run could rapidly grow to one billion bytes or more. This must be stored, calibrated, and analyzed — a very expensive enterprise.

Lastly, a CCD, like a racing car, requires a full-time engineering support team if it is to work at peak performance.

Nevertheless, an increasing number of amateur astronomers are taking the plunge and experimenting with CCD's (*S&T:* January, 1985, page 71, and April, 1986, page 407). This trend is sure to continue as CCD's and home computers become cheaper and more powerful. For amateurs and professionals alike, the CCD is truly a dream coming true. With performance improving continually it is difficult to imagine what further impact it will have on the astronomical community in the next decade.

James Janesick is head of the advanced imaging sensors group at the Jet Propulsion Laboratory. He has developed CCD's for NASA space missions and ground-based astronomy for the last 14 years. Morley Blouke is principal scientist in the CCD engineering group at Tektronix, Inc. He is currently involved in designing, fabricating, and testing large CCD's for scientific applications.