

APPLICATION NOTE

CCD Image Sensor Noise Sources

January 10, 2005
Revision 2.1

Introduction:

In an ideal digital camera, the CCD limits the noise performance, not the system electronics. The application note provides a brief description of the noise sources that determine this performance limit.

We can classify noise sources into two types: temporal and spatial. Temporal noise can be reduced by frame averaging, while spatial noise cannot. However, some spatial noise can be removed by frame subtraction or gain/offset correction techniques. Examples of temporal noise that are discussed in this document include shot noise, reset noise, output amplifier noise, and dark current shot noise. Spatial noise sources include photo response non-uniformity and dark current non-uniformity.

Shot Noise:

Shot Noise is the noise associated with the random arrival of photons at any detector. It is nature's fundamental limit on noise performance in light detection systems. Since the time between photon arrivals is governed by Poisson statistics, the uncertainty in the number of photons collected during a given period of time is simply:

$$\sigma_{shot} = \sqrt{S}$$

where, σ_{shot} is the shot noise and S is the signal, both expressed in electrons. So a 10,000-electron exposure will have a shot noise of 100 electrons. This implies that the best signal-to-noise ratio possible for a 10,000-electron signal is 10,000/100 = 100.

Reset Noise:

At the device output, the signal from an image sensor is typically converted from the charge domain to the voltage domain by means of a sense capacitor and source-follower amplifier. Before measuring each pixel's charge packet, the CCD's sense capacitor is reset to some reference level, V_{RD} . There is an uncertainty in this voltage related to thermal noise generated by the channel resistance of the reset FET. In volts, this noise is:

$$\sigma_{reset} = \sqrt{4kTBR}$$

where, k is Boltzman's constant (J/K), T is the temperature (K), B is the noise power bandwidth (Hz) and R is the effective channel resistance (Ω). In terms of electrons this becomes

$$\sigma_{reset} = \frac{\sqrt{kTC}}{q}$$

where, C represents the sense node capacitance (F) and q is the fundamental charge (C). This equation is the origin of the term 'kTC noise', which is a commonly used synonym for reset noise.

Since this noise is around 50 electrons, most manufacturers of CCD cameras include circuitry that completely eliminates it. Most CCD cameras use a correlated double sampler (CDS), a circuit that measures the difference between the reset voltage and the signal voltage for each pixel, eliminating the need to reset to the same level each time.

Output Amplifier Noise:

The two primary sources of noise in the output amplifier, white noise and flicker noise, are discussed below. Together, these make up the CCD's 'read noise.'

White Noise:

Like the reset FET, the output amplifier has a resistance that causes thermal noise. The effective resistance in this case is the output impedance of the source follower, R_{out} . This type of thermal noise is sometimes called 'Johnson noise,' after its pioneer, or simply 'white noise,' since its magnitude is independent of frequency. In volts, the noise is:

$$\sigma_{white} = \sqrt{4kTBR_{out}}$$

and, in electrons:

$$\sigma_{white} = \frac{\sqrt{4kTBR_{out}}}{\left(\frac{\Delta V}{\Delta N}\right)A_v}$$

where, $\left(\frac{\Delta V}{\Delta N}\right)$ is the sensitivity (V/e-) and A_v is the output amplifier gain.

Flicker Noise:

Flicker Noise, also called 1/f noise, is noise that has an approximately inverse dependence on frequency. The higher the frequency or pixel rate, the lower the noise.

More specifically, the noise power $\left(\frac{V^2}{Hz}\right)$ decreases by a factor of 10 for each decade increase in frequency. Many natural systems exhibit 1/f behavior. They have in common a collection of states that turn on and off individually with randomly distributed time constants. In the case of the MOSFET, the states are traps at the silicon-oxide interface, and the time constants are emission time constants associated with those traps. When electrons are in the traps, they act like pebbles in a stream, affecting the flow of current in the channel. The superposition of all of these traps, each with its own time constant, generates the familiar 1/f noise spectrum. The frequency at which the noise levels off, indicating that the amplifier is operating white noise limited, is called the 1/f corner frequency. For cameras in which pixels are read out slowly ($\sim 1\text{MHz}$), 1/f noise usually determines the noise floor.

In general, white noise increases with amplifier area. Assuming a constant drain current, flicker noise decreases with amplifier area. The goal of amplifier design is to find the lowest-noise compromise between competing geometries for the desired operating frequency.




Figure 1.
Output amplifier noise of a typical Kodak linear CCD

Clocking Noise:

A number of clocks are required to transfer the signal through a CCD and process its output. Many of these clocks are high frequency, specifically the HCCD and reset clocks. Generation of these clocks and the power associated with driving the CCD loads can generate feed through signals to the output waveform. Varying loads due to signal level and variability in clock edge placement due to jitter result in a variation in the level of this feed through we term clock noise. Fast data rates imply higher bandwidth driver circuits and wider bandwidth amplifiers; hence clock noise is shown to be a function of the clocking frequency, typically following a square root relationship.

Dark current noise:

Dark current is the result of imperfections or impurities in the depleted bulk silicon or at the silicon-silicon dioxide interface. These sites introduce electronic states in the forbidden gap which act as steps between the valence and conduction bands, providing a path for valence electrons to sneak into the conduction band, adding to the signal measured in the pixel. The efficiency of a generation center depends on its energy level, with states near mid-band generating most of the dark current. The generation of dark current is a thermal process wherein electrons use thermal energy to hop to an intermediate state, from which they are emitted into the conduction band. For this reason, the most effective way to reduce dark current is to cool the CCD, robbing electrons of the thermal energy required to reach an intermediate state.

Surface Dark Current:

There are far more generation centers at the sensor's surface than in the depleted bulk, typically by a factor of 100. These centers are surface states formed at the silicon-silicon dioxide interface. The number of these states is reduced by proper thermal treatment during and after oxide growth, when hydrogen is allowed to diffuse into the interface and eliminate dangling bonds caused by the mismatch of the Si and SiO₂ lattices. But even with proper processing, surface states remain the major source of dark current.

Although surface states cannot be eliminated, the dark current generated by them can be greatly reduced by inverted operation, also known as accumulation mode clocking. In this clocking scheme, the low vertical voltage is set negative enough to create an inversion layer of holes between the surface and the collecting well. Electrons emitted from surface states recombine with these holes rather than being collected by the well, so that surface dark current is eliminated under barrier phases.

Bulk Dark Current:

Most of the dark current generated in the bulk silicon and collected into pixels is generated at or near the depleted region of the pixel. The mean level of dark current in the bulk is attributed primarily to defects in the silicon. Localized dark current spikes are attributed to trace amounts of metallic impurities. The number of dark current-generating defects depends on the quality of the starting material and the processing of the material during fabrication of the CCD.

Gettering techniques during the processing are employed to move impurities from the active sensor regions but once the device is completed, the dark charge generated in the bulk can be reduced only by cooling the sensor.

The relationship between bulk dark current and temperature follows the empirical formula:

$$D = 2.5 \times 10^{15} \cdot A \cdot I_d \cdot T^{1.5} \cdot e^{-\frac{E_g}{2kT}}$$

where D is the dark current (electrons/pixel/s), A is the pixel area (cm²), I_d is the dark current measured at 300K (nA/cm²), E_g is the bandgap at temperature T, and T is the temperature in Kelvin. The bandgap of silicon varies with temperature according to

$$E_g = 1.1557 - \frac{7.021 \times 10^{-4} \cdot T^2}{1108 + T}$$

Noise Associated with Dark Current:

Dark current generated two types of noise: dark current non-uniformity and dark current shot noise.

Dark current non-uniformity is a noise that results from the fact that each pixel generates a slightly different amount of dark current. This noise can be eliminated by subtracting a dark reference frame from each image. The dark reference frame should be taken at the same temperature and with the same integration time as the image.

Although the dark signal can be subtracted out, the shot noise associated with this signal cannot. As in the case of photon shot noise, the amount of dark current shot noise is equal to the square root of the dark signal.

$$\sigma_{dark} = \sqrt{D}$$

The dark noise in an image resulting from the subtraction of a raw image and a dark frame is more than this by a factor of $\sqrt{2}$.

There exist sources of dark current that do not follow the general dark current equation and cannot be reliably subtracted out. Examples include dark current spikes, generated by proton-induced cluster damage or by various metallic contaminants, contained in the bulk silicon.

Photo Response Non-Uniformity (PRNU):

Due to process variations, not all pixels demonstrate the same sensitivity to light. The result at the pixel-to-pixel level is a faint checkerboard pattern in a flat-field image. Usually this variation is on the order of a percent or two of the average signal, and is linear with average signal.

The noise associated with this variation in sensitivity can be removed by ‘flat-fielding,’ a process by which a previously captured flat-field image is used to calibrate out the differences between pixels. Although this process removes the photo response non-uniformity, the subtraction of images introduces a $\sqrt{2}$ increase in shot noise.

Noise Measurement (Photon Transfer Curve):

Generation of a Photon Transfer Curve is a useful CCD camera characterization technique. In addition to a camera’s noise floor, photon transfer also provides the sensor’s full well, and the camera’s conversion constant (electrons/DN). The full well and noise floor can then be used to derive the system’s dynamic range.

Each point in a photon transfer curve represents a portion of a flat field image (e.g. 100 x 100 pixels) taken with a different exposure time. Generally the integration time is fixed, but the amount of time the light source is on varies for each image. This way the dark current noise is constant throughout the curve. For each 100 x 100 pixel square, the noise, i.e. the standard deviation in the pixel values, is plotted against the average signal. The plot is traditionally done on a log-log scale.

The resulting curve has three sections. At the lowest signal levels, the curve is flat. This portion of the curve can be extended to the noise axis to give the camera’s read noise floor. The middle portion of the curve has a slope of $\frac{1}{2}$ and represents the part of the camera’s dynamic range over which its operation is shot-noise limited. The last third of the curve has a slope of 1 and corresponds to the range in which the camera’s operation is pattern-noise limited. This portion of the curve can be eliminated by using data from differences of frames rather than individual

frames. A comparison of the photon transfer curves using individual frames and differenced frames can be used to measure PRNU.

An example of a photon transfer curve based on individual frames is shown below. The noise floor comes in near 25 electrons rms and the sensor’s full well is around 40,000 electrons. This implies a dynamic range of

$$DR = 20 \cdot \log\left(\frac{40,000}{25}\right) = 64dB$$




Figure 2
Photon transfer curve for a full frame Kodak CCD and evaluation board at 28 MHz.

Revision Changes:

Revision Number	Description of Changes
1	Initial formal version.
2	Updated to current device types covered by this application note.
2.1	Updated format only.