From Light to Numbers

Radu Corlan

2006

Contents

1 The Sensor

Image formation begins with the camera lens projecting an image of the recorded scene on the surface of the sensor. The photons of light that are absorbed on the surface of the sensor are converted into electric charge (through the photoelectric effect). This charge is accumulated on a capacitor in each pixel and thus converted into a voltage.

Not all photons that fall on a pixel are converted into charge; some are absorbed in protection layers on top of the sensor, or opaque structures in the pixels. The percentage of photons that generate useful charge is the *quantum efficiency* or *QE* of the sensor. Quantum efficiency depends markedly on the wavelength (color) of light, and can reach peak values of around 40% in the sensors used in Phantom cameras.

Before an image frame is taken, the voltage on the pixel capacitor is brought to a preset value, which sets the *dark level* of the sensor. Then for a preset amount of time, the *exposure* or *integrating time*, the pixel is left to accumulate any charge generated by illumination. Finally, at the end of the exposure, the resulting voltage is routed to the sensor output, the pixels are reset and the cycle repeats.

There is a maximum value the voltage on the pixels can reach before being limited by the sensor's circuitry. This level is the *saturation level* or *full well voltage* of the sensor. The saturation level is usually expressed in electrons, typical values being between 20000 and 50000 electrons.

Below saturation, the difference between the voltage at the end of integration and the dark level is proportional to the illumination of the respective pixel. The relation between the two is linear to a high degree.

2 Digitization

For images to be stored in a digital camera, the voltages at the outputs of the sensor have to be converted into integer numbers. Since the amount of information (bits) that can be stored for each pixel is limited, the voltages cannot be represented exactly. Instead, the values are rounded to the nearest level that can be stored.

The number of possible levels depends on the *bit depth* of the camera system. As the number of levels increase, the distance between levels, or *quantisation step* is decreased, and the sensor output can be represented more precisely, as shown in the table below:

Bit Depth	Levels	Step size
8	256	0.0039
10	1024	0.00976
12	4096	0.000244
14	16384	0.000061
16	65536	0.000015

The maximum error that can occur when rounding exact values to a limited bit depth is one half of the quantisation step size. It can be seen that the error decreases rapidly with the number of bits (only 0.003% of full scale for 14 bits).

As an example, let's consider that the sensor output varies from 0V (dark level) to 1.0V (saturation level), and that a particular output level is exactly 0.734696V. At various bit depths, this level will generate the following codes:

Bit Depth	Code	Rounded value	Error
8	187	0.7333	0.0013
10	751	0.73412	0.00058
12	3008	0.73491	0.00022
14	12036	0.734664	0.000032
16	48148	0.734691	0.000004

3 Memory

The memory system is the most complex and expensive part of a high-speed digital camera. As such, it is desirable that memory is used as efficiently as possible. In a camera that stores uncompressed images, the main choice is in selecting the precision (or bit depth) of stored data. Phantom cameras use memory systems supporting various bit depths (from 8 to 16 bits/sample). In all cases, changing the memory bit depth does not modify the range of input values that are stored (which always span the full dynamic range of the sensor); it's the precision of values that is changed. Or, in other words, when saving less than the full bit depth, the bits that are stored are the most significant ones.

4 How Many Bits?

With so many bit depth options, a natural question arises: how many bits are needed? The answer depends on the nature of the application. In many cases in the industry 8-bit systems provide good images that yield excellent data when analysed. In addition to providing longer recording times, using the 8-bit depth minimizes file sizes and speeds up all image processing steps.

In more demanding applications, such as low light or high-dynamic range situations, substantial image quality improvments can be achieved by increasing the bit depth. For cameras that spec a dynamic range between 60 and 66 dB, going from 8 bits to 10 will show a significant increase in image quality. The next step, from 10 to 12 is less obvious but still measurable. Going to 14 bits makes all quantisation effects negligible; this greatly simplifies error modelling in complex data reduction chains.

The above considerations assume the digitisation is "perfect", so that all quantisation steps are equal, and transitions noiseless. In reality, common data converter industry practice is to allow some deviation from the perfect A/D response. This can be mitigated by the fact that using just the most significat bits of a converter creates the equivalent of an substantially perfect part, albeit of lower resolution. This is the reason why some cameras use 10 or 12-bit converters even though only 8 bits are stored in memory.

5 Display

A camera with a dynamic range of 60dB can detect light variations over a range of 1000:1 and more. Common display devices are much more limited however. Computer monitors show a maximum of 256 levels even in the best viewing conditions and printed images are even more limited.

The phantom software maps the wide range of input values present in a high bit depth image to the 256 levels of display devices using the dark and white point sliders. The dark point slider specifies what input level maps to zero (full black) on the display scale. All input levels below the value of the dark point slider will also map to black, so any image content below the black level will be invisible on the display. The White point slider work similarly. A pixel having an input level equal or higher than the white point slider will display as 100% white. Levels between the sliders are converted linearly to display values. When adjusting the sliders, the resulting transfer function is displayed on top of the image histogram.

6 Files

The primary way of saving camera data is to uninterpolated (raw) .cin files. These files always contain the original data from the camera, as well as the positions of all image adjustment controls. This way, when the file is later open, images look the way they did when captured. However, the adjustments are reversible. if for instance the black/white point sliders' values caused the image highlights to be clipped, moving the high slider back to higher values will recover them.

When images are saved to RGB file formats or exported to e.g. AVI or JPEG files, all image adjustments are applied before saving, so image content lost to clipping as well as the reduction in the number of level due to using a lower bit depth are permanent.