

RAW Processing Using Adobe Camera RAW (ACR)

By Robert Lay

December 2, 2005

Revision 2, December 3, 2005

Revision 3, January 17, 2006

Revision 4, January 22, 2006

Introduction

This paper presents an explanation of how ACR works. It includes an analysis of the interactions between the primary ACR processing controls. It also provides a set of detailed initial conditions or settings recommended for use as Camera Defaults, which are intended to provide you with an unbiased starting point for examining your RAW image.

An understanding of the interactions between the controls combined with a well defined set of initial conditions and a set of procedural rules allows the user to better utilize ACR – especially in those more extreme situations of over-exposure or a high range of scene brightness.

Due to their importance in understanding the functions of the ACR controls, it is also necessary that we include topics on the subject of the camera sensor, the concept of “clipping” and some basic understanding of how each control works, both individually and in conjunction with other controls.

The purpose of this document is to provide the reader with a better understanding of how the ACR plug-in works internally and is not intended as a comprehensive procedural guide.

In Search of a Model

There is a big difference between what the user sees and what is really going on behind the scenes in Adobe Camera RAW. As Frank Morgan, Wizard, says in the “Wizard of Oz”, “Pay no attention to that man behind the curtain!”

In order to promote a better understanding of what really goes on behind the scenes, we not only have to have a “model” that includes those aspects of RAW capture that we know about, we also have to include certain characteristics or attributes of the system, such as soft and hard clipping and recoverable highlight data that we know are there only through observing its behavior as we use the program.

Our model of the internal workings of ACR is very idealistic – it accounts for all aspects of behavior that we can observe in normal use, without any regard for exactly how those characteristics or attributes are implemented. We know, for example, that the internal

computations are carried out in 16 bit precision using a range of values from 0 through 32768 for a total of 32769 distinct values. We also know that the RAW pixel values are quantized to 4096 discrete values. On the other hand the user only sees the Preview display and the Histogram and the individual Red, Green and Blue channel values as 8 bit values (0 to 255).

The most intriguing aspect of the model is that it must account for three separate categories of pixels for its internal representation of what is being displayed to the user. The user is presented with display information which unambiguously distinguishes between clipped pixels and unclipped pixels.

Definition: Clipped pixels have individual channel values that are either 0 (black) or 255 (white).

In addition to displaying information that indicates clipping, the model must also support a behavior such that some pixels (“soft clipped”) have the ability of passing from an unclipped state into a clipped state and then returning to an unclipped state. This is described as recoverable detail, whereas pixels that are identified as clipped at the time of the exposure in the camera are identified as “hard clipped” and cannot ever change to an unclipped state.

Definition: Hard clipping of pixels is setting them permanently to the maximum or minimum possible value.

Definition: Soft clipping of pixels is causing one or more channels of a pixel to be set, temporarily, to the maximum or minimum possible value.

Unfortunately, ACR does not provide a display that distinguishes between hard and soft clipping. We are only able to make that distinction through observation of the behavior of ACR as we process the RAW files using the controls provided.

An Ultra Simple Model of the Camera Sensor and Analog/Digital (A/D) Converter

In order that the reader be able to relate the clipping mechanisms discussed under section “*In Search of a Model*”, it is necessary to have at least a cursory knowledge of the camera’s photographic sensor and A/D converter.

The array of sensors is essentially an array of photo-sensitive devices that accumulate an electrical charge when they receive light energy. The voltage developed in a given sensor site is proportional to the light intensity and the length of time over which the light is present (the integration of the light energy over time). If the light intensity is very low (as from a lens opening of f/22) and the exposure time is very long (say 1 second), then the voltage due to the electrical charge at that site might be “x” volts. The same voltage could be developed with a larger lens opening (say f/5.6) and a correspondingly shorter exposure time (1/15 seconds). The important thing is that the voltage developed at the

given site will be linearly related to the integrated light; which means that if we either double the exposure time, or open the lens by one f-stop, the voltage will then be “2x” volts rather than “x” volts. Whatever the voltage is, it is converted to a binary number in the camera, so that the virtual image produced in the camera becomes an array of 12 bit numbers, each of which represents a brightness level (0 – 4095) at one of the sensor sites for that image being photographed.

Such sensors and their ancillary A/D converters have limits on their sensitivity (a threshold value for the lowest value of light detectable, below which they cannot record a value) and a maximum voltage output, which represents a hard limit on the maximum level of integrated light energy that they can record. These two limits effectively determine the maximum range of brightness that can be recorded by the camera. Whatever that range of brightness is, it is scaled and expressed as a value between 0 (black) and 4095 (white), assuming a 12 bit encoder. In this document we will refer to that range as 0 through 255, instead – reason being that in ACR, the brightness levels are scaled and displayed as 8 bit per channel values in order to be consistent with the JPG format. This does not mean that the converted RAW image cannot be stored in 16 bit per channel formats. On the contrary, both 8 bit and 16 bit per channel formats are available for storing the converted image.

The only reason that the reader needs to understand anything about the sensor array and its A/D converter is that whenever the integrated light energy falling on a give sensor site during an exposure exceeds the maximum value that can be recorded, the value stored in the virtual image array will be the maximum value possible (“4095” in a 12 bit encoder) and not a number that is linearly related to the integrated light. In other words, the numerical value stored in the array is “clipped” (to use the photographic term) or “clamped” to use the electronics term.

In the user interface to ACR, that maximum possible level is 255 decimal (FF hex, or 11111111 binary, or 377 octal). It is the maximum unsigned value that can be stored in an 8 bit byte. If you choose a 16 bit per channel format for storage of the converted image, the maximum values stored in that image file are scaled to a modified 16 bit data width, which means they will be stored with values from 0 to 32768 (32768 values, total).

ACR Initial Settings and Camera Default Settings

With any RAW image opened in ACR, in **Settings**, select Custom.

Set White Balance=As Shot, Exposure=0, Brightness=0, Contrast=0, Shadows=0, Saturation=0, Sharpening=0, Luminance Smoothing=0, and Color Noise Reduction=0.

[Note that for the Powershot G5 camera the camera defaults are as follows: White Balance=As Shot, Exposure=0, Brightness=50, Contrast=+25, Shadows=5, Saturation=0, Sharpening=25, Luminance Smoothing=0, and Color Noise Reduction=25.]

Having established a Custom setting, use the menu arrow (->) at the right of Settings and select “**Set Camera Default**”. Note that you can return to the original Camera Default at any time by using **Reset Camera Default**.

Having replaced the Camera Default settings with the above custom settings, any time a RAW file is opened for the first time in ACR, these settings will be applied initially. Note that if a RAW file has already been opened in ACR, it will display with whatever processing settings that were previously applied (settings automatically saved in Adobe database and Sidecar files and are applied whenever a RAW file is re-opened).

The reasons for suggesting that the user set the Camera Defaults as described above are two-fold. First, it provides a common reference configuration for comparing our results with your results in the tutorial topics. Second, the above Camera Default Settings are closer to what is commonly referred to as “Linear Capture”, -i.e, the image as it would appear after the demosaicing is complete but before the gamma encoding is applied.

Basic Procedures

Setting the White-point with the Exposure Control and Finding the Onset of Soft Clipping

The RAW file will often include some hard clipped pixels. ACR shows both hard clipped and soft clipped pixels in the Highlight or Shadow clipping displays. The Highlight and Shadow clipping displays are invoked by holding down the Alt key as the left mouse button is pressed on the Exposure or Shadow slider, respectively.

When using the Highlight clipping display, the clipped channels display their own colors when clipped and display as white when all three channels are clipped – all on a black background.

When using the Shadow clipping display, the clipped channels display their own colors when clipped and display as black when all three channels are clipped – all on a white background.

Note that the clipping displays show the combined hard and soft clipped pixels – there is no display that distinguishes between hard and soft clipped pixels, directly.

To find the onset of soft clipping, reduce the Exposure control to the extreme left and check the Highlight clipping display. Whatever pixels are showing white at this point are sure to be hard clipped. As the control is moved to the right, somewhere short of a setting of zero, the clipping display will show a sudden increase in clipping. This is where soft clipping begins. Up until that point, all “hard clipped” pixels are shown in white and are essentially constant as the control is moved to the right. That is the critical point of adjustment of the Exposure control. Fully blown out pixels will never leave the clipping display, regardless of how far to the left the Exposure control is set. It is important to adjust the Exposure control throughout its range to determine the critical point of adjustment.

Setting the black point with the Shadows Control

Setting the black point follows a procedure essentially the same as setting the white point, except the Shadows control is used. All pixels with value 0,0,0 in the Shadow clipping display are shown in black on a white background. It is important to adjust the Shadows control throughout its range to determine the critical point of adjustment. The critical point is that point where you first encounter a sudden avalanche of pixels becoming clipped to black. Those pixels which were clipped with the Shadows control at full left position should be ignored, as they cannot be changed by any controls.

Except for those few pixels that show as clipped when the Shadows control is at zero, all pixels which become clipped by moving the Shadows control to the right are soft clipped and can be recovered.

Interactions between the Controls

Open a RAW file and start with the controls at the initial settings (the customized Camera Default settings). Those settings will display an image that is dark in comparison to what you normally expect. That is because these initial settings do not give the image its normal brightness and gamma settings, and they were deliberately chosen to give a standardized configuration for observing the interactions between the controls. The Brightness control is at 0 instead of 50 and the Contrast control is at 0 instead of +25. (The Brightness control has a range of 0 – 150 while the Contrast control has a range of (-25) to (+100). Those two initial settings create an image with less contrast and brightness than that to which we are accustomed.

We will not discuss the following controls – Tint, Saturation, Sharpening, Luminance Smoothing and Color Noise Reduction. Those controls have essentially no significant effect on the final image brightness, contrast or range of tones.

The interactions between the controls will be discussed in the following topics using the standardized initial settings of the controls as the point of departure.

Interactions of the Shadows Control

The Shadows control is primarily for setting the black point. That is, set the Shadows control so that the lowest value pixels are as close to the left boundary of the histogram as possible without becoming clipped. You will find that as you move the Shadows control to the right it forces pixels to bunch up and become clipped at the left boundary of histogram but does not modify the white point. There is essentially no interaction between the Exposure control and the other controls.

Interactions of the Exposure Control

The Exposure control is primarily for setting the white point. The objective is to set the Exposure control so that the highest value pixels are as close to the right boundary of the histogram as possible without becoming clipped. You will find that the Exposure control spreads the pixels out over a greater or lesser range of values but does not affect the black point. When moving either the Exposure control or the Shadows control, it is as if the

other control is nailed down or locked. There is essentially no interaction between the Exposure control and the Shadows control. However, as will be shown, there can be considerable interaction between the Exposure, Contrast and the Brightness controls.

Interactions of the Brightness Control

The primary difference between the Brightness control and the Exposure control is that the latter operates right through the right boundary of the histogram as if the boundary were not there. As the Exposure control is set to higher values, the pixels are simply spread further to the right – right through the right boundary, where they are stacked up as clipped pixels. Compare this with the Brightness control which elevates (brightens) the lower and middle tones but forces the higher values to block up or become compressed but not clipped (this is a significant difference between the Exposure and Brightness controls – one which is easily overlooked).

The Brightness control modifies all of the pixels by applying a multiplicative factor, thus increasing the value of each pixel by a fixed percentage. That effect is almost identical to the behavior of the Exposure control except that the Brightness control is constrained to not push any pixels through the right boundary. In other words, as the Brightness control is increased to the right, a higher multiplicative factor is applied to all values, increasing each by a fixed percentage, but the product is constrained from reaching the clipping point. The net effect of this is to cause a brightening of the image to the extent that the higher values become more and more blocked (compressed), just as would happen when reaching the shoulder of a film characteristic curve. Compare this with the Exposure control which just pushes pixels right through the histogram right boundary (temporarily clipping them to 255,255,255). The difference is that the Exposure control actually represents the abrupt and hard limit reached by the sensor when it reaches its saturation point, whereas the Brightness control simulates the soft shoulder of overdeveloped film.

Interactions of the Contrast Control

The interactions of the Contrast control are the most complex of all the controls. This is because the Contrast control operates so as to either

- Spread pixels outwards from the center of the distribution towards the ends
- Or, force pixels back towards the center from the ends

Consider the first of those two possibilities, which is the one that we would describe as increasing the overall contrast. The pixels below the midpoint are reduced in value by an amount proportional to their distance from the midpoint, and the pixels above the midpoint are increase in value by an amount proportional to their distance from the midpoint. In both directions the spreading is tempered by the proximity to the boundary of the histogram. In other words, increasing the contrast will not normally cause clipping either in the highlights or the shadows, but it will definitely cause blocking up of the highlights and blocking up of the shadows. Note that we can also describe this as an improvement in the contrast for the mid range tones at the expense of contrast in the highlights and shadows. In other words, one is traded for the other – there is no free lunch!

The latter possibility is the opposite of the first. Now the contrast of the mid range will be decreased and the contrast in the highlights and shadows will be improved. We are again trading one for the other.

The Contrast control can also be described as turning an otherwise linear relationship between the input (the RAW data) and the output (the final image data) into an “S” curve or a backwards “S” curve. The “S” curve with the higher contrast for the mid-tones is considered the normal processing objective in that it usually gives us the image that we find most normal and most pleasing. We should keep in mind at this point that our choice of initial conditions for studying the interactions between the controls is a configuration of neutral contrast – i.e., neither mid-tones nor highlights nor shadows are given any emphasis.

Clearly, the Contrast control can and will conflict with both the Exposure control and the Brightness control.

Interaction between Contrast and Brightness Controls

If the Contrast control and the Brightness control are both allowed to depart from our arbitrary initial conditions, it is clear that the two controls have the potential to nullify each other in the highlight region (When Brightness is increased and Contrast is moved to negative values). Conversely, if both Brightness and Contrast are increased, there is little effect on shadows but the effect on mid-tones is significant.

Interaction between Contrast and Exposure Controls

If the Contrast control and the Exposure control are both allowed to depart from the initial conditions, one would expect that they could nullify each other or strengthen each other in the highlight region, just as is the case when the Contrast and Brightness controls are both changed. Absent “soft clipping”, that is exactly what would happen – they would nullify or strengthen one another in the highlights region. However, in practice, “soft clipping” prevails, and the highlights are darkened by a decrease in exposure, regardless of the Contrast setting.

When the Contrast control is operating with the Exposure control there is little effect on the shadow region if Contrast and Exposure are both reduced or both increased. However, if Contrast and Exposure controls are working in opposition, that will cause the Shadow region to either darken or lighten. The following table helps to summarize the interactions between the Contrast and Exposure controls. (Note that all incremental changes are performed from a base configuration of initial settings as defined in “ACR Initial Settings and Camera Default Settings”.)

	Shadows	Mid-tones	Highlights
increase brightness & decrease contrast	Lightens	Lightens	0
increase brightness & increase contrast	0	Lightens	Lightens
reduce contrast & reduce exposure	0	0	Darkens
increase contrast & increase exposure	0	Lightens	Lightens
increase contrast & decrease exposure	Darkens	Darkens	Darkens

decrease contrast & increase exposure	Lightens	Lightens	0
---------------------------------------	----------	----------	---

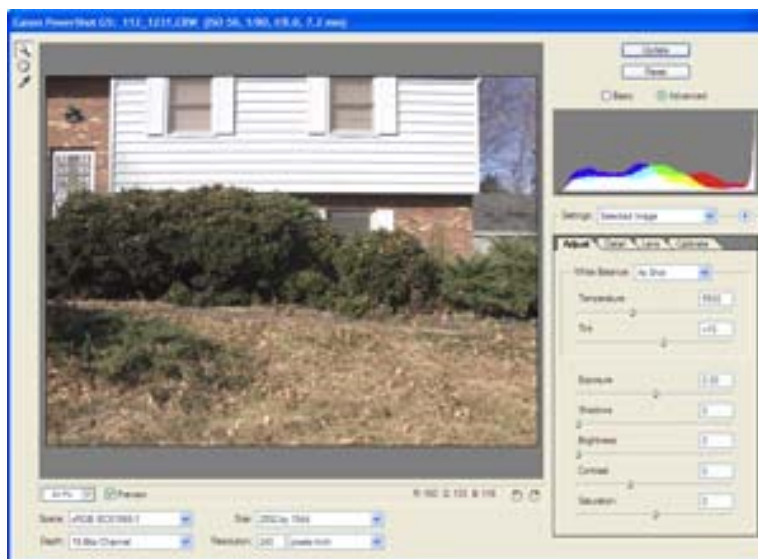
Interactions of the Color Temperature and Tint Controls

The interactions of the Color Temperature and Tint controls with other controls should be negligible if there is no clipping. However, no satisfactory tests have been designed to verify this.

Dealing with a Problem Image

This screen shot shows a problem image opened in ACR. Remember, we are using the initial settings as described above under “ACR Initial Settings and Camera Default Settings”. Under those conditions the image does not have much appeal, because its brightness and contrast have been reduced to minimum. In general, properly exposed images will look flat and dark in comparison with the image presented using the ACR Camera Defaults. The image that we will use as an example is a slightly over-exposed image with a large range of brightness, which just exceeds the camera’s ability to capture all of the tones present in the scene.

A low contrast, underexposed image will be treated in another section.



Before making any adjustments, we should analyze the image to determine exactly what problems it has. The first thing that stands out is the severe clipping of highlights indicated by the tall white line at the right boundary of the histogram.


The badly clipped highlights correspond to the white shutters on the side of the house facing the early morning sun.

We can examine this in more detail using the Highlight clipping display. The problem is immediately recognized as being confined to the white shutters facing the early morning

sun. Aside from that, the rest of the image is OK, except for not having much contrast in the mid-tones. Since this shot was intentionally overexposed by one stop, our analysis at this point simply confirms what we already expected. We can determine the extent or severity of the problem with the Exposure control.

Sadly, we find that this image has such severely blown out highlights that even reducing the exposure to the -4.00 value (the limit of our Exposure control's range of adjustment) is of no help, and that means that it cannot be salvaged. How could that happen with only a one f-stop overexposure? The answer is simple. We metered on a standard 18% Gray card and arrived at a correct exposure for an average scene. We then purposely overexposed by one f-stop. However, this scene has an abnormally high dynamic range, such that the white shutters are already blown out completely even with only 1 stop overexposure. You may have thought that there would be more exposure latitude than that.

The next picture shows the Shadow clipping display.



This is why Ansel Adams and others developed the Zone System - specifically to provide for proper placement of scene values to be followed by less than normal development of the negative.

In our example, we should have been more aware of the fact that even with a correct exposure, the shutters were clipping.

However, there is another problem. Why is it that even with an application of negative 4 f-stops of Exposure control correction in ACR, we still cannot affect the severe clipping seen in the white shutters? The answer is that the exposure level on the shutters siding was sufficient to cause "hard clipping", which is explained fully, elsewhere in this document. The fact is that once the areas that were "hard clipped" have been identified, there is no point in trying to recover beyond that point, because the information is not there and cannot be recovered.

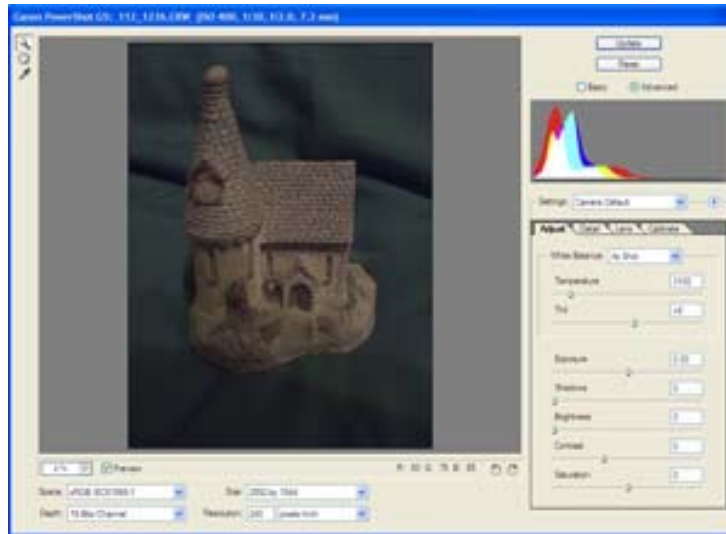
The following will provide the procedure for properly recognizing or identifying the area that is “hard clipped” as opposed to the areas that are soft clipped and can be recovered. It is also important to realize that while the “hard clipped” areas cannot be recovered at all, the “soft clipped” areas can be completely recovered without any degradation. The procedure to follow is essentially the same as was discussed under “**Basic Procedures**”.

Using the technique for observing the clipped areas, as described above, move the Exposure control from a high setting all the way to the -4.0 limit. As you move the control, note the exact point where appearance of the clipping screen stops changing. Keep that picture in mind, because that is the picture of the “hard clipped” pixels. As you make whatever adjustments of the various controls that you deem necessary, or to your liking, just remember to keep coming back to the Exposure control, periodically, and check to insure that it is still at the critical point where the clipping is at the verge of changing. If you do that properly, then when you are all done, you can be assured that you have recovered all that can be recovered.

A reasonable question at this point would be, “why not use the settings that were provided for your camera instead of the “initial settings” that I have recommended in section “**ACR Initial Settings and Camera Default Settings**”? In order to answer that question properly, we have to look at the example image in ACR as it would have looked when using the Camera Default Settings that are initially provided in ACR for this camera. The image would look significantly overexposed, as one would expect (remember, it was purposely overexposed by one f-stop!). What would we do at that point? We would probably adjust the Exposure control downwards in order to make the picture darker, and that would work – up to a point. The problem that would arise is that we would not have a good understanding of why the adjustment of the Exposure control has no effect on recovering the detail in the blown out shutters. We also would not have an appreciation of exactly where to set the Exposure control in order to recover exactly that detail which is recoverable. This is not to say that those techniques could not be learned while using the Camera Defaults for your camera. Once you have understood the nature of clipping, and once you have understood how each control works, and once you have understood how the controls interact with one another, then, of course you can use any initial settings that you like. Our purpose in configuring ACR with the recommended settings is to give you a better appreciation of how the unprocessed RAW data appears while it is still relatively unmodified by the algorithms that are used to provide the gamma encoding, which is what the original Camera Defaults were intended to provide. If at the end of your adjustments, you feel that you have independently arrived at the same thing as was provided by the Camera Defaults, then by all means stick with the Camera Defaults, because the purpose of this document is to provide an educational tool, not a procedural guide.

Dealing with a Low Contrast or Underexposed Image

The screen shot below illustrates a typical image having a lower range of brightness, which may have been the result of under-exposure, or the result of an overcast or low-contrast scene.




The histogram shows clearly that the image is low-key, which may have been due to accidental under-exposure. In any case, the overall contrast is low, which is another way of saying that the range of brightness is very conservative or low (in comparison to the camera's capability).

The objective is to brighten the image and extend the range of tones by use of the basic controls, Exposure, Brightness, and Contrast.

First, we set the white point using the Exposure control. This requires about +1.5 f-stops of correction. Next we increase the Brightness control to 50. Next, increase the contrast to +25, which gives us an overall appearance of a properly exposed, full tone image. It is not a coincidence that the corrections applied are virtually the same as the originally supplied ACR Camera Defaults for this camera. The only unique adjustment was the Exposure correction. Therefore, there is no reason to not use the originally supplied Camera Defaults, once you have gained experience in RAW processing.

Here, then, is the final result.



End Notes

Additional Notes on the Camera Sensor – Hard Clipping

Having an understanding of how the Camera Sensor and its A/D converter work is one thing. Finding out what its operational performance parameters are is something else. It would be helpful to know how many f-stops range of scene brightness can be captured with a given camera's sensor and A/D converter.

In order to discuss the theoretical performance characteristics and make an informed guess about the operational parameters, one can theorize that the operational range of the sensor is comprised of two ranges – one that is bounded by the full range of voltage that can be generated by the integrated light energy, V_{max} , and the other, somewhat lesser range of voltage that is converted to the mathematical range of values from 0 to 255, which I will call V_o - the operational or output range.

V_{max} would be that voltage that corresponds to “hard clipping”. When the integrated light from a given pixel reaches V_{max} , that pixel is in hard clipping, and it is impossible to recover any detail or meaningful information from that pixel other than the fact that it is hard clipped. This condition can arise only in the camera as the image is being captured. By the time the virtual image file is transferred from the camera to a computer, that pixel is permanently categorized as “hard clipped” and unrecoverable.

On the other hand, V_o represents a voltage level corresponding to an arbitrary level of integrated light that the designer has designated as the operational range of integrated light values, or exposure. From the design standpoint, V_o will be a lower level of voltage than V_{max} but not so low as to be inefficient or wasteful of a critical resource – namely, the dynamic range of the system. In other words, the operational or useful range is purposely held at a slightly lower level than the total sensor capability, simply for the purpose of providing a small unused portion of the transfer characteristic for other

purposes. Therefore, we should assume that the operational or useful range of values from the sensor will use almost but not all of the dynamic range capability of the sensor.

In order to put this small unused portion of the full dynamic range into perspective, we will assume that it will be used to provide a region of over-exposure that supports recovery of lost data. Before we go any further in that direction, we need to re-visit the concept of “soft clipping”.

Additional Notes on the Camera Sensor – Soft Clipping

The range of output voltage from 0 to V_o (scaled to 0 – 255) represents the useful or operational range of the combined sensor and A/D converter. Everything that the user can ultimately do with the virtual image during RAW processing is constrained to pixel values between 0 and 255 inclusive. That’s the range of numeric values that the JPG and other 8 bit per channel formats use for displaying images on a monitor screen or printing them on a printer. (For purposes of this discussion, we are simply ignoring as not germane to the issues, the fact that there are also 16 bit per channel formats available for image processing and storage, although not all output devices are capable of working with 16 bit/channel image formats.)

In order for the RAW processor to perform the conversion from the RAW format to end-use formats, it is necessary to support a concept of “soft clipping” and buffering of “soft clipped” pixels. While using the Exposure, Brightness, and Contrast controls in ACR, it becomes apparent that pixel counts in the histogram can be freely translated up and down the brightness scale. Not only that, but pixel counts can be promoted into a region beyond the right boundary of the histogram. Any pixel count that is translated into that region is considered to be soft-clipped.

It is fundamental to our model of the RAW processor that “soft clipped” pixels be freely translated back into the histogram when needed. Obviously, a pixel that is returned to the histogram must come back with its previous value – otherwise the entire RAW processing concept would be non-linear and irreversible – attributes or characteristics that we know are not true of the RAW processor.

Equally fundamental and critical to the model of the RAW processor is that “hard clipped” pixels are not capable of being translated back and forth through that right boundary. There is no information available about hard clipped pixels other than the fact that they are hard clipped. They are recorded in the camera when the RAW file is written to memory and forever carry the value 255,255,255, (pure white).

Additional Notes on the Camera Sensor – Recoverable Detail and Dynamic Range Issues

In the two previous sub-sections it was mentioned that there is a small region of the sensor dynamic range that is held in reserve – that is, a region of higher brightness levels that are supported by the camera but not available as values to be associated with pixels in the output virtual image file from the RAW processor.

From another perspective, it is those brightness levels that could be supported by the hardware up through the point of delivering the data as a file written to the camera's memory card. That small range of brightness levels is purposely kept small because it is part of a critical resource – the system dynamic range. Its purpose is very simply to provide the Camera RAW processing algorithms with a small amount of head room that can be used for recovery of detail from over-exposure. We can only guess at how much range of brightness is being reserved for this purpose in a given camera, since it is not a published attribute. From our own measurements and experience with the Canon Powershot G5, we believe that to be approximately 1.00 f-stops.

My informed guess for the Powershot G5 camera, based on my own experiments and experience, would be that it can capture a range of no more than approximately 8 - 9 f-stops over base noise level. An interesting aspect of that is that there is reason to believe that a certain amount of "Automatic Gain Control" is applied electronically at the time of exposure, if the ISO speed setting is set to **Auto**. When in **Auto** the camera chooses the lowest ISO setting that it can use without causing underexposure but raises it to the maximum available when necessary. That would suggest that there is virtually no provision for automatically reducing the dynamic range of a high contrast scene but that it will expand the dynamic range of a low contrast scene if set to Auto.